EUROPEAN CENTRAL BANK

RISK MANAGEMENT FOR CENTRAL BANK FOREIGN RESERVES

EDITORS:
CARLOS BERNADELL,
PIERRE CARDON,
JOACHIM COCHE,
FRANCIS X. DIEBOLD AND
SIMONE MANGANELLI

RISK MANAGEMENT FOR CENTRAL BANK FOREIGN RESERVES

EUROPEAN CENTRAL BANK

EUROPEAN CENTRAL BANK

RISK MANAGEMENT FOR CENTRAL BANK FOREIGN RESERVES

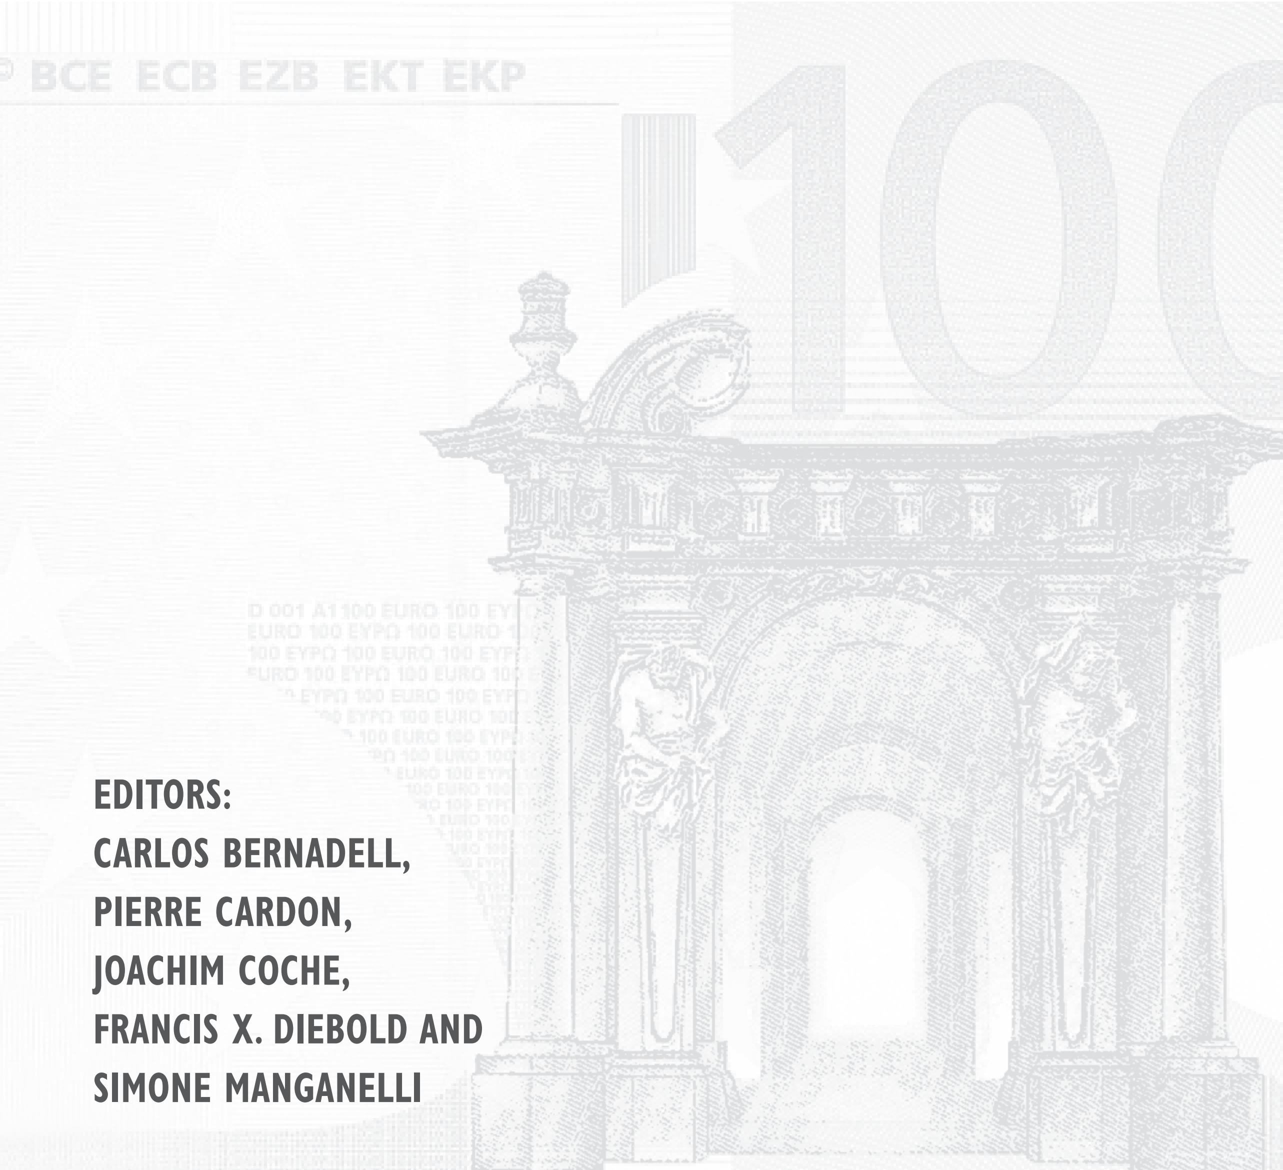

EDITORS:
CARLOS BERNADELL,
PIERRE CARDON,
JOACHIM COCHE,
FRANCIS X. DIEBOLD AND
SIMONE MANGANELLI

Published by:

Address	Kaiserstrasse 29 60311 Frankfurt am Main Germany
Postal address	Postfach 16 03 19 60066 Frankfurt am Main Germany
Telephone	+49 69 1344 0
Internet	http://www.ecb.int
Fax	+49 69 1344 6000
Telex	411 144 ecb d

This publication is also available as an e-book to be downloaded from the ECB’s website.

Editors:
Carlos Bernadell (ECB), Pierre Cardon (BIS), Joachim Coche (ECB),
Francis X. Diebold (University of Pennsylvania) and Simone Manganelli (ECB)

Typeset and printed by:
Kern & Birner GmbH + Co.

ISBN 92-9181-497-0 (print)
ISBN 92-9181-498-9 (online)

Table of Contents

Foreword

Risk management is a key element of sound corporate governance in any financial institution, including central banks. In particular, central banks, in performing their policy tasks, are exposed to a variety of financial and non-financial risks, which they may want to manage. One such key risk concerns foreign reserves, because central banks' main activity, namely ensuring price stability, needs to be backed by an adequate financial position.

Efficient management of foreign exchange reserves is vital if a central bank's credibility is to be maintained. For many central banks, a significant part of the financial risks inherent in their balance sheet arises from foreign reserve assets. Successful foreign reserves management ensures that the capacity to intervene in the foreign exchange markets exists when needed, while simultaneously minimising the costs of holding reserves. Risk management of foreign reserves contributes to these objectives by strategically managing and controlling the exposure to financial and operational risks.

Undoubtedly, foreign reserves risk management can benefit from methodologies and tools applied in the private asset management industry, as well as from developments of leading firms in competitive markets. However, the motivation for a volume addressing risk management from a central bank's point of view is that not all private sector concepts are directly applicable to foreign reserves management. Central banks are idiosyncratic investors, because policy objectives induce specific portfolio management objectives and constraints and prescribe a generally prudent attitude towards market, credit and liquidity risk. Foreign reserves management deviates in terms of the investment universe, available risk budgets, investment horizons, management of liquidity risk, and the role and scope of active portfolio management.

This volume gathers valuable contributions by academics and practitioners that reflect the specific nature of central bank reserves management. The contributions highlight the important role risk management plays in the continuous validation and improvement of central banks' investment processes.

Traditionally, reserves were mainly invested in liquid sovereign bonds. A changing investment universe makes it possible or even requires holdings to be more diversified. While observing liquidity and other policy requirements, highly-rated non-government instruments are added to the investment universe. These developments change the role of risk management: beyond a pure risk control perspective, proactive risk management must on a strategic level be involved when transforming policy requirements into strategic investment decisions.

Despite a broadened investment universe, holding foreign reserves implies opportunity costs, as investments must necessarily deviate from a broadly diversified market portfolio. In recent years, many central banks have started using active management to further minimise these costs. Strategies and methods applied in the private asset management industry have therefore found their way into reserves management. These developments should go hand-in-hand with a further strengthening of risk management functions.

This volume contributes to the development of methodologies and best practices in a changing environment for reserves management. In so doing, it strengthens the belief that risk management functions in central banks need comprehensive mandates to assure an efficient allocation of resources, development of sound governance structures, improved accountability, and a culture of risk awareness across all operational activities.

Gertrude Tumpel-Gugerell
Member of the Executive Board of the European Central Bank

Introduction

Carlos Bernadell (ECB), Pierre Cardon (BIS), Joachim Coche (ECB),
Francis X. Diebold (University of Pennsylvania), and Simone Manganelli (ECB)

The management of foreign exchange reserves is an important task undertaken by central banks. Depending on the design of exchange rate arrangements and the requirements of monetary policy, foreign reserve assets may serve a variety of purposes, ranging from exchange rate management to external debt management. Hence central banks' efficient management of foreign reserves is vital if they are to fulfil their mandates comprehensively. In particular, efficient allocation and management of foreign reserves will promote the liquidity needed to fulfil policy mandates while at the same time minimising the costs of holding reserves. Central bank foreign reserves *risk management* can contribute to these objectives by managing and controlling the exposure to financial and operational risks.

In recent years, many central banks have expanded their risk control units into comprehensive risk management functions, beneficially independent to some extent from the bank's risk-taking activities, and supporting decisions at all stages of the foreign reserves investment process. In addition to supporting traditional control functions such as compliance monitoring, foreign reserves risk management can contribute to the translation of policy goals into specific and efficient strategic asset allocations that focus not only on risk, but also on return.

Indeed, it is precisely the risk-return interface, and the tension that arises for central banks navigating that interface, that motivate this volume. On the one hand, it is probably socially wasteful for a central bank to hold only sovereign bonds, accepting their relatively low risk-free return, which suggests the desirability of more aggressive central bank investment strategies. On the other hand, central banks are unique institutions with very particular mandates, which suggests that naively importing private sector asset management strategies may be misguided. So, then, what should a central bank do? In this volume, we attempt to progress toward an answer.

Our approach contains three components, corresponding to the volume's three parts: (I) General Framework and Strategies, (II) Specifics of Risk Measurement and Management, and (III) Case Studies. Part I sets the stage in broad terms, suggesting and evaluating various alternatives, and making it clear that an appropriate framework must respect the unique aspects of central banking environments, characterised by a high degree of risk aversion and institutional constraints. Part II contains a variety of rather more technical contributions focusing on risk measurement and optimisation of the risk-return trade-off as appropriate for central banks. Finally, Part III contains descriptions of current practice at a variety of central banks worldwide, which are designed to provide context and perspective.

Part I, General Framework and Strategies, begins with **Cardon and Coche,** who stress the importance of good corporate governance and a sound organisational design. Their paper views strategic asset allocation as a three-step process. First, an appropriate organisational design should be developed to ensure a smooth implementation of daily reserves risk management. The paper argues for a three-tier governance structure, with clearly distinguished and segregated strategic asset allocation, tactical asset allocation, and actual portfolio management responsibilities. Second, the general policy and institutional requirements should be translated into specific, precise and quantifiable investment guidelines. Finally, these investment guidelines should be transformed into an optimal long-term risk-return profile.

Putnam dwells on the second step of the process described by Cardon and Coche. He argues that for central bank foreign risk management, it is crucial to understand the interplay between investment objectives and investment guidelines. A thorough examination of the commonly-employed investment guidelines may uncover the existence of strategies that actually work against the complex long-term investment objectives of central banks. In concrete terms, he suggests addressing the trade-off between short-term and long-term needs by dividing the foreign reserves portfolio into two sections: "liquid" and "liquidity-challenged". This would permit the central bank to withstand sudden shocks to the market environment, while at the same time earning liquidity, complexity and volatility premia which are typically only available to long-term investors.

The remaining five contributions in Part I provide different examples of how central banks' investment guidelines can be embedded in a well-structured mathematical framework. **Claessens and Kreuser** suggest a numerical approach in order to solve a dynamic stochastic optimisation model that incorporates both macro aspects of policy objectives (e.g. monetary policy needs and foreign exchange management) and micro elements (e.g. the definition of portfolio benchmarks and the evaluation of investment managers).

Fisher and Lie criticise risk management strategies based on exogenous ad hoc restrictions of the investment universe. This typical asset allocation process generally leads to overconstrained portfolios and to significant efficiency losses. They suggest an asset allocation framework that maximises portfolio returns, given a risk target and subject to constraints on liquidity, credit quality and currency allocations.

Remolona and Schrijvers examine three alternative strategies. The first focuses on duration, the second on default risk and corporate bonds, and the third on higher-yielding currencies. They find that the trade-off between risk and return as measured by the Sharpe ratio points to a recommended duration of not longer than two years. In the case of corporate bonds, the key issue is how to achieve a proper diversification, given the significant asymmetries that characterise the distribution of these portfolio returns. For higher-yielding currencies, empirical evidence suggests that yield differentials are generally not offset, but rather reinforced, by currency movements.

Ferket and Zwanenburg quantify the risk and return characteristics of some of the most popular asset classes in the private asset management industry (long-term and global government bonds, investment-grade credits, high-yield bonds and equities), which they then compare to those of a cash benchmark – the lowest risk portfolio. Their empirical results suggest several diversification strategies that may have attractive risk-return trade-offs for central banks.

Scherer and Gintschel look at currency allocation. The literature focuses on two problems – wealth preservation and liquidity preservation – that are typically solved separately. Rather than following each approach in isolation, the authors model the currency allocation decision as a multi-objective optimisation problem, making explicit the trade-off between the two objectives, and incorporating political constraints into the decision-making process.

Part II of the volume, Specifics of Risk Measurement and Management, contains contributions that deal with more technical aspects of the risk management process. **Nugee and Dwyer** introduce the concept of "whole enterprise" risk management. They first describe risk management from a narrow financial risk control perspective. Then, they examine the typical financial risks faced by a central bank, and critically review the traditional risk methodologies in use. In the second part of the paper, they argue in favour of a wider framework of risk management and corporate governance for the entire central bank, incorporating aspects of legal, operational and reputational risks in addition to the common financial risks.

The papers by **Grava** and by **Ramaswamy** discuss issues related to diversification towards corporate bonds and measurement of credit risk. Given the current environment – characterised by low- yield, highly-rated government bonds – managers of official foreign exchange reserves have started to consider higher-yielding alternative instruments. The overall message is that the potential inclusion of higher-yielding securities in a central bank reserves portfolio should not be discarded a priori, provided that the related risks are properly measured and managed.

Grava studies the effects of adding corporate securities to reserves portfolios. He considers only highly-rated investment-grade bonds, on the grounds that investment in lower-rated securities might require specialised skills and resources not typically available at a central bank. The main finding is that adding spread risk leads to better risk-return profiles than increasing portfolio duration. Moreover, a long-term passive allocation to credit sectors, coupled with the ability to tolerate short-term underperformance, generates significantly higher returns in the long run.

Ramaswamy provides a framework to implement an internal credit risk model for reserves management in a central bank. The model uses as input only publicly available information, thereby providing a good compromise between accuracy and simplicity. The paper also provides indicative values for the credit risk model parameters required for quantifying credit risk.

The next three papers deal with market risk. **Dynkin and Hyman** describe the Lehman Brothers market risk model. This is a multi-factor model, with the factor loadings rather than the factors viewed as observables. The paper illustrates the advantages of such a methodology and provides a good overview of its usefulness for risk management.

Bortje and van der Hoorn present a balance sheet approach to managing market risk. The paper distinguishes two dimensions along which the financial strength of a central bank can be measured: its profit-generating capacity, and its ability to absorb losses. A central bank's profitability can be gauged from the profit and loss account. Under simplifying assumptions about exchange rates and yield curves, the paper argues that profitability is largely driven by the size of the monetary base, the interest rate level, and operating costs. On the other hand, the ability to absorb losses is found by comparing the potential loss (as measured by the Value-at-Risk of the portfolio) with the total amount of reserves. The composition of the balance sheet is subsequently optimised within a constrained maximisation framework.

Puschkarski develops a general procedure for decomposing time variation in portfolio Value-at-Risk from one reporting period to the next. This decomposition occurs across three main dimensions: time, market developments, and changes in portfolio allocation. A fourth element, taking into account the interaction between these three dimensions, is also described. Such analysis will help managers to set and monitor risk limits, and to understand how and why they are occasionally breached.

Finally, we conclude Part II with a very general contribution on operational risk as relevant to central banks. **Embrechts, Kaufmann and Samorodnitsky** note that operational risk arises from inadequate internal processes and/or unanticipated external events, both of which are highly relevant for central banks. Hence proper quantification of operational risk may affect central bank reserve management, because the bank will generally want to react when confronted with unanticipated catastrophic events. The paper first discusses issues related to the availability and characteristics of operational risk data, and then, exploiting analogies between the nature of operational risk data and insurance losses, argues that statistical tools from extreme value theory can be successfully applied to the modelling of operational risk.

Part III of the volume, Case Studies, contains contributions by risk managers from a selected sample of central banks around the world. **Rogers** gives a broad non-technical overview of the implementation of risk management at the ECB. The paper describes the ECB's financial position by examining a stylised balance sheet, illustrating the monitoring and management of the main risks related to currency and interest rate movements.

Schmidt, Bauer, Koblas and Mochan describe how Česká národní banka (CNB) manages its foreign exchange reserves. CNB has the explicit objective of maximising returns on its foreign reserves, subject to liquidity, market risk and credit risk constraints. The paper describes the strategies adopted to achieve this objective, with special attention given to the currency composition and the duration of the portfolio.

Ho illustrates the framework and the application of risk management to Hong Kong's foreign reserves portfolio. The paper starts with a brief historical overview of the Exchange Fund – the body in charge of safeguarding the value of the Hong Kong dollar. It then moves on to describe the risk management framework, the implementation of the strategic asset allocation, and measurement of performance attribution.

The issue of performance attribution is taken up by the two subsequent papers. **Zajac and Simicak**, from Národná banka Slovenska, and **de Almeida**, from the Central Bank of Brazil, discuss in detail the methods used in their respective central banks to identify the sources of differential returns in a portfolio with many assets and currencies. As these contributions clearly point out, performance attribution is a key element in the risk management process. It enhances the transparency of the investment process and ultimately leads to portfolios that more closely reflect the general investment guidelines.

The volume ends with two contributions from the Central Bank of Venezuela (CBV) and the Bank of Israel. **Delgado, Martínez, Osorio and Pabón** discuss the methodology in place at the CBV for the risk, return and liquidity management of CBV international reserves. Liquidity management is particularly challenging in Venezuela since the country's international reserves are mainly determined by oil exports, which represent a significant source of volatility. **Assouline** presents a method to determine the target duration of the Bank of Israel's dollar portfolio using a shortfall approach. The method requires the portfolio manager to set three preference parameters, which reflect the bank's risk aversion, and calculates the optimal portfolio duration implied by these parameters.

In closing, we would like to thank all of the authors who contributed to this volume. In compiling it, we have attempted to convey a sense of the excitement presently associated with risk management in central bank foreign reserves contexts, as cutting-edge techniques from private sector asset management are adapted to central bank environments. Indeed, the contributions make it clear that best practice central bank reserves management is already in a state of flux, owing to improvements in asset management techniques and decreasing supplies of government bonds. To minimise the costs of holding reserves while observing liquidity and other constraints, central banks are now adding non-government bonds to their investment universe, and are increasingly using active asset management strategies, employing modern performance attribution and risk decomposition methods to evaluate performance. We hope that the volume stimulates additional discussion and provides a blueprint for additional improvements.

1 GENERAL FRAMEWORK AND STRATEGIES

1

Strategic asset allocation for foreign exchange reserves[1]

Pierre Cardon and Joachim Coche

Abstract

This paper discusses a possible blueprint for the management of the foreign reserves' strategic asset allocation. At the outset we address the importance of a sound organisational set-up. A three-tier governance structure comprising an oversight committee, investment committee and actual portfolio management is one approach whereby asset allocation decisions can be efficiently implemented. In a second step, we focus on the design of investment philosophies, which translate general policy requirements into concrete objectives and constraints required when establishing the long-term risk return profile. Finally, a quantitative framework for deriving the actual asset allocation is developed.

1 Introduction

Central banks hold foreign exchange reserves for a variety of reasons, one of which is to maintain the capacity to intervene in exceptional circumstances in currency markets. Another is to provide liquidity to support currency boards and fixed exchange rate regimes. With the aim of reducing external vulnerability, foreign reserves holdings also take into consideration the country's external debt. Furthermore, reserves serve as a store of national wealth. Central banks have to choose an appropriate strategic asset allocation of the foreign reserves in agreement with these general policy objectives. An important consequence of the chosen asset allocation is its impact on overall performance and risk over time, as shown by many empirical studies.[2]

Strategic asset allocation can be defined as the long-term allocation of capital (wealth) to different asset classes such as bonds, equity and real estate. The aim is to optimise the risk/return trade-off given the specific preferences and goals of an individual or an organisation. For central banks' foreign reserves portfolios, the asset allocation process typically comprises decisions on the currency composition and, within each currency, on the allocation to various fixed income asset classes, mainly government bonds and other highly liquid, highly secure instrument types.

Although it is to be expected that a strategic asset allocation decision will be effective over the medium to long term, the allocation might be reviewed and revised in the light of changing investment opportunities. Despite these revisions, strategic asset allocation does not aim to generate superior returns compared to a market index by moving in and out of asset classes at the most beneficial time.[3] Rather, the strategic asset allocation process transforms

[1] The views expressed in this article are those of the authors (Pierre Cardon (BIS) and Joachim Coche (ECB), and do not necessarily reflect those of the Bank for International Settlements or the European Central Bank.

[2] For example, Ibbotson and Kaplan (2000) show that asset allocation decisions explain about 90% of the variability of returns over time.

[3] Such a market timing strategy may be implemented by using a tactical asset allocation over a short to medium-term horizon.

goals and risk return preferences into the long-term optimal proportions of individual asset classes.[4] In the context of reserves management, strategic asset allocation may be seen as a three-step process.

In the first step, we will show the importance of a sound organisational set-up for managing reserves efficiently. In terms of an active investment style, we will argue for a three-tier governance structure where the responsibilities for strategic, tactical asset allocation and actual portfolio management are clearly segregated. Once in place, this framework will facilitate a disciplined implementation of the asset allocation decision and should help in clarifying accountability, managing risks and promoting a risk awareness culture across the organisation.

In the second step, we will discuss three alternative investment philosophies for central banks whereby policy requirements can be translated into investment principles. We will first look at the individual currency approach, where the primary objective for reserves management is to ensure efficient risk-return combinations on the level of individual currency sub-portfolios. Alternatively, the base currency approach explores diversification effects on the level of aggregated reserves as measured in the central bank's domestic currency or another base currency such as Special Drawing Rights (SDRs) issued by the IMF. In contrast to these first two asset-only approaches, the asset and liability perspective seeks to derive objectives by taking into consideration central banks' ability to bear financial risks and or the country's external debt.

In the third step, the reserves' long-term risk-return profile is derived from the previously established investment principles. To this end, we discuss a model-based approach in order to establish a strategic asset allocation and risk budgets for active management. Such a quantitative process, in our example essentially a basic one-period mean-variance optimisation, would be the most objective, long-term estimate for fulfilling the investment principles. It also offers the advantage of disburdening decision-makers, who are responsible for the design of monetary policy, from having to make concrete investment decisions beyond specifying preferences and policy requirements. However, we will argue that such a quantitative investment process should not be followed mechanically. Instead, all results should be subject to an extensive validation process before a strategic asset allocation is finally decided.

The paper is organised as follows. Section 2 describes an organisational set-up which ensures effective governance and implementation of day-to-day reserves management. Section 3 discusses how policy requirements can be transformed into concrete objectives and how constraints for strategic asset allocation can be codified in the Statement of Investment Principles. Based on this, Section 4 outlines a quantitative process for strategic asset allocation. Finally, Section 5 concludes this paper.

---

[4] In recent years, discussions have focused on broadening the strategic asset allocation towards a more comprehensive risk budgeting approach. When establishing asset class weights, the risk budget approach simultaneously determines the optimal leeway for active management, and thereby explicitly accounts for diversification effects between benchmark risk and active management risk (Chow and Kritzman, 2001). Risk budgeting could also be seen as a technique for tracking the risk per unit of return.

2 Organisational set-up

Most central banks today are subject to stringent reporting requirements vis-à-vis the general public and, more specifically, parliamentarian or governmental bodies. In order to satisfy these commitments, it is of the utmost importance that the central bank's balance sheet and the implied financial risks are managed efficiently. In particular, given the significance of the foreign reserve assets in the financial statements of many central banks, a transparent and accountable reserves management framework should be in place to ensure effective governance and implementation of the agreed strategic asset allocation. The framework should rely on a sound organisational set-up and appropriate measures to manage and control financial risks.

A three-tier governance structure

A necessary requirement for transparency and accountability is a clearly specified investment process in combination with a sound governance structure. If the central bank decides on an active investment style, an increasingly popular practice is to have a three-tier governance structure comprising an Oversight Committee, an Investment Committee, and Portfolio Management units that are responsible for strategic and tactical asset allocation and actual portfolio management. Figure 1 illustrates this three-tier governance structure with an investment process that allows active reserves management. Starting from a passively managed strategic asset allocation, a tactical asset allocation is added, followed by actual portfolio mandates. The aim of this set-up is to improve the risk-return profile of the strategic benchmark by providing the necessary flexibility to take advantage of short to medium-term investment opportunities.

Figure 1: Organisational set-up

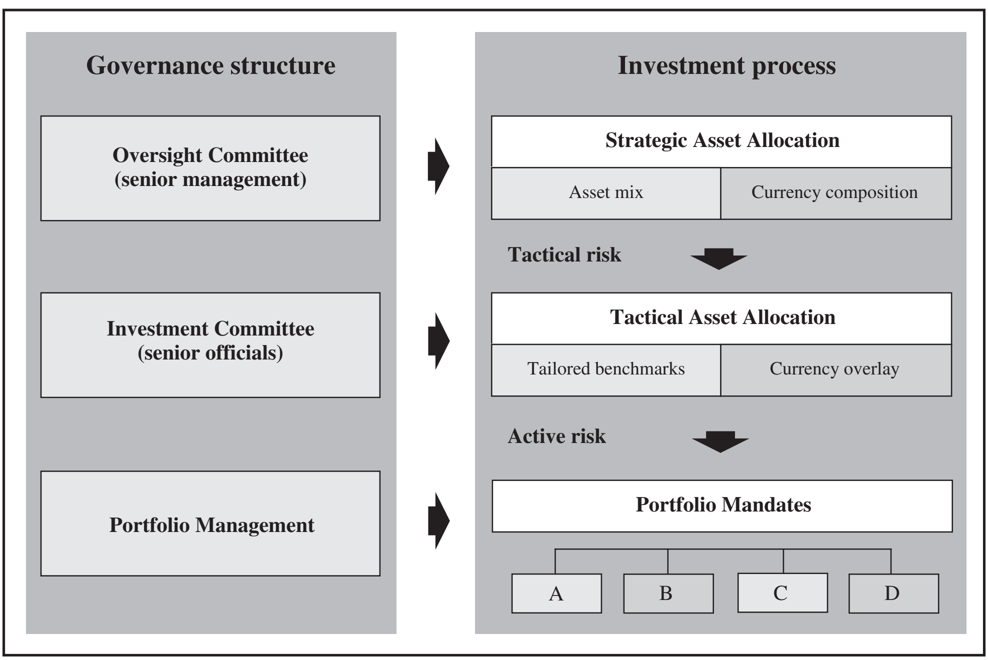

In this structure, the key responsibilities of the oversight committee are to lay down the Statement of Investment Principles, articulate the long-term risk-return preferences and objectives for the management of foreign reserves, and oversee the efficient implementation of the asset allocation. The oversight committee's main vehicle for conveying the institution's risk-return preferences to the remaining tiers in the governance structure is the strategic asset allocation. Usually, the strategic asset allocation will be set for the medium to long term. However, the oversight committee should be able to review it to reflect shifts in the central bank's risk preferences and structural changes in investment opportunities (see Foley, 2003). In addition, through a comprehensive risk budgeting approach, the oversight committee simultaneously establishes the asset mix, the currency composition as well as the leeway for active management. Both the asset mix and the currency composition are defined independently of each other as passive strategies. The leeway for active management is often defined in terms of forward-looking tracking error or relative Value-at-Risk (VaR)[5]. It is usually broken down into a tactical risk budget, monitoring the extent to which the tactical benchmark is allowed to deviate from the strategic one, and an active risk budget, monitoring the deviation of the actual portfolios mandates from the tactical benchmark (see Winkelmann, 2000).

To avoid conflicts of interest at the policy level, the oversight committee should not be involved in actual implementation issues. Typically, the committee comprises Executive Board members responsible for risk management, portfolio management and internal finance as well as senior officials from these areas. This structure allows the reserves' risk neutral position to be established with a clear allocation of responsibilities.

At the second tier, an *investment committee* is responsible for the implementation and monitoring of the strategic asset allocation. The committee should also be allowed to deviate from the strategy by using the tactical risk budget to exploit movements in risk premia between asset classes and currencies. To this end, it is responsible for establishing a tactical asset allocation that can take advantage of changing investment opportunities over the short to medium term. The tactical asset allocation will be reviewed more frequently than the strategic asset allocation (see Anson, 2004). The investment committee can either have a total return objective, or aim at outperforming the strategic benchmarks. In the former, the tactical asset allocation would be seen as a long-term overlay programme, deviating from the strategic asset allocation mainly to protect against downside risk. The latter represents a way to relax the constraint that the portfolio mandates have to be managed against the strategic benchmark. These tactical decisions could be implemented either through tailored benchmarks in line with the portfolio managers' skills, or through currency (or asset) overlay strategies, using mostly derivatives instruments to reduce the implementation costs as far as possible. To be successful, tactical asset allocation requires a strong governance structure that must also be comfortable with the risk of short-term losses. It should only be attempted if there is broad consensus within the central bank in favour of such an approach.

Finally, internal or external *portfolio management* units implement the tactical asset allocation decisions and take active risks versus their respective tailored benchmarks. This is the second layer of active management after the tactical asset allocation, which is governed by the Investment Committee. These two layers should be as far as possible independent from each other in order to diversify investment styles. The Investment Committee assigns

[5] Forward-looking (ex ante) tracking error is normally expressed as the standard deviation of the possible future difference in portfolio and benchmark return over the coming year. It measures both upside and downside risk, whereas relative VaR only measures the downside risk.

portfolio mandates to portfolio managers with a credible and replicable investment strategy supported by a solid risk model and a sound portfolio construction process. Portfolio mandates are allocated with a relative return objective against a tailored benchmark within guidelines broad enough to allow the active managers to apply their skills. Usually, portfolio managers will have a shorter time frame than in tactical asset allocation. The trend is to use the available active risk budget as efficiently as possible by investing in strategies where the probability of adding outperformance is maximised and where inefficiencies can create opportunities. On an aggregate basis, the objective is to put together a group of managers whose information ratio[6] should be higher than at the individual manager's level (see Winkelmann, 2001).

Implementation of the asset allocation decision

The strategic asset allocation process, as described so far in this paper, results in asset class weightings and risk budgets for active management. The implementation therefore first has to address how to derive a strategic benchmark from the asset class weights; second, how to define the rebalancing rules; and third, how to specify the leeway for active management, if any, around this benchmark.

Within the investment process, *strategic benchmarks serve three main functions* in addition to reflecting the bank's long-term risk-return preferences. First, the strategic benchmark establishes the risk-neutral position for active management. For example, tactical asset allocation may retreat to the strategic benchmark in the absence of specific views on outperformance opportunities or at times of exceptional uncertainty. Second, the strategic benchmark provides the yardstick for measuring and attributing the success of any active or passive management strategy. However, the strategic benchmark should be intended as a guide and not as an index that is tracked too closely. Third, establishing strategic benchmarks is a precondition for effective risk control, as they are the long-term position against which the reserves' risks and returns are measured.

Strategic benchmarks can be customised either based on notional portfolios of instruments (internal benchmarks) or using publicly-available market sub-indices. In the case of internal benchmarks, the notional portfolios are implemented following clear rules, which specify rebalancing frequencies and securities selection within the individual asset classes. For example, one of these rules could always specify the inclusion of the latest issue of 2, 5 and 10-year US Treasuries. Conversely, the asset allocation might be implemented on the basis of market indices. Indices for individual market sectors (issuer classes and maturity buckets) are widely available. Thus in this case, the strategic benchmark is composed of these sector indices, weighted according to the previously determined asset allocation. When comparing both options, internal benchmarks are more costly but also more precise. For example, internal benchmarks accurately reflect the investor's preferred investment universe, minimise pricing mismatches between benchmarks and actual portfolios, and can be designed to have stable characteristics over time. On the other hand, market indices are more diversified and are more transparent if communicated to the public.

In both cases, the strategic benchmark, whether using internal benchmarks or market indices, should not be allowed to drift too far, as it will otherwise lose its anchoring role. However, in view of the transaction costs, it is not optimal to rebalance the strategic

[6] The information ratio (the excess return over the tracking error) indicates if managers are achieving sufficient additional return from taking active risk.

benchmark constantly. There is therefore a trade-off between trading costs and acceptable deviation. The selection of *rebalancing rules* has been widely covered in the financial literature, with each rule producing different risk and return characteristics. For instance, if the benchmark allocation is rebalanced on a monthly or quarterly basis, this will be a form of a contrarian's strategy that forces the portfolio manager to buy in falling markets and to sell in rising markets. However, a benchmark can be rebalanced free of transaction costs and is therefore harder to replicate. In practice, the underlying reserves will gradually drift from the strategic benchmark. The Investment Committee should therefore be responsible for managing this drift and should have the discretion to decide if and when to rebalance the tactical benchmark within the deviation ranges. One rule could be to rebalance the tactical benchmark only if the tactical ranges are breached.

In reserves management, *risk budgets for active management* are traditionally defined in terms of maximum deviations in modified duration and/or maximum exposures to individual asset classes and maturity buckets. Modified duration is a powerful and well-established concept for the management of fixed income securities belonging to one asset class in the short run. However, there are important drawbacks for longer investment horizons and especially for the management of more than one asset class. For example, duration does not take changes in spreads between asset classes and changes in the shape of the yield curve into account. Limits based on tracking error, defined as the standard deviation of excess returns, are more effective as they comprehensively limit positions in terms of interest rate, credit and currency risk.

Generally risk budgets, no matter whether they are defined as tracking error or modified duration limits, can be implemented as *hard limits* or *soft limits.* Hard limits require the portfolio manager to be in line with such a limit at each point in time. Exceeding the limit would constitute a breach, which would be followed up by risk control units. Conversely, soft limits might be exceeded for a short period of time. Soft limits would therefore reduce unnecessary trading costs as tracking error could be impacted by short-term developments in volatility and correlation. They would also provide portfolio managers with an orientation of the sponsor's appetite for active risk and excess return objectives. In the private sector, the hard limit concept is often found on the trading floor, whereas the soft limit idea typically prevails in an asset management environment. Although the reserves management of central banks is in many aspects comparable to asset management, a rigorous implementation based on hard limit concepts currently appears to be more in line with the general prudent attitude of many central banks.

From risk control to risk management

Originally, the risk control functions in central banks were designed for reserves which had to remain highly liquid, using the Treasury function in *commercial banks as a model.* Accordingly, risk control functions were implemented focusing on the computation of daily profit and loss figures and risk measures. Furthermore, transactions were checked on an intraday or day-by-day basis in order to identify as quickly as possible any "rogue traders" who were not compliant with the investment framework. The principal objective of risk control was thus to capture short-term anomalies and ensure that the trading books were matched. As reserves have grown, the adequacy of commercial banks as the sole model for the design of risk control functions has become gradually less appropriate, especially for the investment tranche of the reserves. In this context, the investment horizon is medium to long term and portfolio managers are expected to keep open positions against a benchmark, unlike a typical trader in a commercial bank environment. In this respect, reserves management activities are more comparable to those of the private asset management industry.

Following the model from *the asset management industry*, risk control units expanded into risk management functions, supporting decisions on all levels of the investment process. The more comprehensive mandate comprises, in addition to the traditional measurement and compliance tasks, four main features. Firstly, the integration of risk management aspects right from the start when transforming policy requirements into a strategic benchmark decision. Secondly, on the level of active management, risk management supports the tactical asset allocation and the actual portfolio management processes by providing models to diversify the various types of risks and advising portfolio managers on the optimal allocation of skills to risk budgets. For example, risk management will have to find ways of supporting portfolio managers to take more credit risks if deemed appropriate. Thirdly, performance and risk measurement would be extended to an in-depth assessment of portfolio management skills through performance attribution analyses. Fourthly, risk management will play an important role in the continuous validation and improvement of the bank's investment process.

In the end, *central banks' reserves management* is situated somewhere between a commercial bank's Treasury operation and a private asset manager. Given the generally prudent attitude of central banks, requirements from both worlds have to be fulfilled. On the one hand, a rigorous limit framework following the commercial bank model has to be implemented, which ensures the compliance of portfolio management with the decision-maker's guidelines. On the other hand, risk management has to contribute to an efficient usage of the available risk budgets. Therefore, risk management has a role at all levels of the investment process, ranging from supporting the oversight committee when translating preferences and policy goals into the reserves strategic asset allocation, to the bottom of the process when monitoring the compliance and the success of portfolio management.

3 Policy objectives and investment principles

To derive the reserves' strategic asset allocation, general policy objectives, such as the provision of liquidity, the reduction of external vulnerability or the storage of national wealth, have to be transformed into more specific, preferably quantifiable, objectives and constraints. These objectives and constraints are to be laid down in the Statement of Investment Principles and form the basis for determining the optimal combination between reserves' expected return, liquidity (risk) and security (market and credit risk). In addition, the investment principles also define the organisational set-up and make specific provisions for the conduct of tactical asset allocation, portfolio implementation and risk management. In this section we discuss three alternative investment philosophies which may support, depending on the decision-maker's preferences, the formulation of concrete objectives and constraints. Subsequently in Section 4, we introduce a specific example of a strategic asset allocation process which picks up these objectives and constraints, quantifies them if necessary, and derives the reserves' strategic asset allocation.

As a first investment philosophy, *the individual currency approach* might be implemented, in which the strategic asset allocation for individual currency sub-portfolios is established independently of the reserves currency distribution. This therefore separates the decisions about the allocation of the overall reserves to individual currencies and the allocation to individual asset classes within these currencies. The subordinated treatment of exchange rate risks implied by this approach reflects the notion that these risks are of a special nature for central banks[7]. Therefore, finding asset allocation strategies that minimise exchange rate exposure might not be the primary objective. Regarding the specification of concrete objectives, this approach requires the objectives for each currency sub-portfolio to be defined together with an objective for determining the reserves' currency distribution. For example, the objectives for the individual sub-portfolio might be to maximise expected returns (for each sub-portfolio separately), given liquidity constraints and a maximum tolerance for credit and interest rate risk. With regard to the currency distribution, the objective might be to find a risk-minimising strategy, subject to policy constraints such as minimum allocations to individual currencies. One potential difficulty of such an approach is to establish the levels of maximum risk tolerance or the required minimum returns. Given the assumed asset-only perspective, these specifications do not necessarily concur with the institution's ability to bear financial risks. Furthermore, it is unclear how the specifications for the individual sub-portfolios and for the currency distribution relate to each other, as this approach disregards diversification effects on the level of country sub-portfolios. These drawbacks might be particularly relevant for central banks that have rapidly accumulated reserves, as observed in a number of Asian and eastern European countries in recent years. In such cases, the prospect of exchange valuation losses might require the reserve portfolio to be diversified beyond the level that is attainable by the individual currency approach.

A second investment philosophy, *the base currency approach*, explores diversification effects on the level of aggregated reserves as measured in the central bank's domestic currency or in another base currency (such as SDRs). This approach maintains the asset-only perspective of the first alternative, but establishes the currency distribution and the asset composition of the sub-portfolios simultaneously. It requires the definition of objectives only at the level of the overall reserves in base currency returns. For example, the objective could be to minimise the overall risk exposure in domestic currency subject to predefined return requirements, as well as to meet policy and operational needs (e.g. liquidity requirements and a minimum allocation to specific currencies). This approach utilises diversification effects between currency returns and local returns. Compared to the first alternative, this approach is likely to result in more extreme allocations within individual sub-portfolios (e.g. higher modified durations in some sub-portfolios). However, it could improve the reserves' risk-return profile measured in base currency. Furthermore, this approach mitigates the problem of establishing adequate levels of risk tolerance or required returns, as such specifications are now only required on the level of the overall reserves and no longer on the level of the respective sub-portfolios.

A possible third alternative is *the asset and liability approach,* which seeks to integrate the management of reserves portfolios with either the central bank's ability to bear financial risks or the country's external debts. This approach, which is growing in importance, might be considered as the overriding principle for determining the objectives and constraints of reserves management. Asset liability management refers either, in a narrow definition, to a joint consideration of asset and liability items on the central bank's books or, more broadly, incorporates the management of the country's sovereign or external debt. In a narrow definition, the size of foreign reserves[8], the currency distribution and the strategic

[7] Unlike interest rate and credit risk, currency risk is considered to be inescapably linked to the reserves functions. Moreover, fluctuations in the reserves' market value induced by exchange rate movements may rather concur with liquidity needs. The reserves' market value (in the bank's domestic currency) is high when there is an increased likelihood of action in support of the domestic currency and vice versa.

benchmarks would be determined simultaneously in a balance sheet context. Such an approach exploits, in addition to diversification effects within the foreign reserves, the risk reduction or income enhancement potential stemming from the reserves' correlation with other balance sheet items. In addition to increasing balance sheet efficiency, such an approach is also useful for establishing the central bank's risk tolerance on the grounds of capital adequacy considerations. In concrete terms, the following objectives are conceivable from an asset and liability perspective: maximising income from the reserves given the entity's ability to bear financial risks, or minimising the reserves' contribution to the overall balance sheet risk, alternatively minimising the shortfall of income from the reserves with respect to its funding.

In particular, when the institution's responsibilities encompass both reserves management and debt management, asset-liability management may be applied in a broader sense. In such a case, objectives and constraints for deriving the strategic asset allocation can be determined against the size and characteristics of foreign currency liabilities. For example, reserves management may strive to find the optimal risk return combination in a set-up where risk is measured relative to the composition of external public sector debt.

With any of the three above investment philosophies, specific policy and operational requirements enter the reserves management process in the form of investment constraints. Such constraints usually determine first the eligible investment universe, and second, may impose direct restrictions on the reserves' currency allocation and instrument composition.

Deviating from the market portfolio recommended by modern portfolio theory, central banks typically restrict the eligible investment universe to highly liquid government bonds and instruments issued by international financial institutions, government-sponsored institutions and supranationals. This narrow definition is primarily based on the overriding importance of liquidity considerations. In addition, the sponsor's preference for prudent reserves management may enter the investment process not only in the form of parameters in the utility function, but also in the definition of eligible instrument classes.

Despite this narrow definition, the specification of the investment universe is a dynamic process. The characteristics of asset classes change over time. For example, at the end of the 1990s, the long-term outlook for the liquidity of the US Treasury market was in question, whereas the liquidity of other market segments, such as the US agency market, had substantially improved. Furthermore, new instrument classes suitable for central banks' reserves management emerged, such as the Medium-Term Instruments (MTI) issued by the BIS. In practice, proposals for the inclusion of new instrument classes are often driven by portfolio managers. Portfolio managers are close to the market and assess the characteristics of various instrument classes from their daily trading activities. A general requirement might be that changes in the investment universe should be made in connection with a review of the strategic asset allocation. It is, however, debatable whether every change in the investment universe necessarily induces a change in the passively managed strategic benchmarks to retain a level playing field with active management.

A second class of constraints, also motivated by liquidity considerations, may restrict the reserves currency composition or, within the individual country sub-portfolios, define thresholds for the allocation to individual instrument classes. For instance, minimum currency weights can be imposed on the basis of the relative size and depth of markets alongside the needs of intervention operations.[9] On the sub-portfolio level, liquidity

[8] The adequacy of reserves is beyond the scope of this contribution.

considerations can be dealt with, given a previously-defined investment universe, by introducing explicit constraints on the minimum share of highly liquid instruments or by creating a liquidity tranche separated from the investment portfolio. In any case, constraints on the currency composition, minimum and maximum shares of individual instrument classes should be established on the basis of a predefined strategy, which defines the order in which to liquidate liquid and less liquid assets.[10]

4 Deriving a strategic asset allocation process

Having defined objectives and constraints, we will now discuss how a strategic asset allocation process is determined in practice by following a four-step process, comprising the quantification of risk tolerance and investment horizon, the formation of return and risk expectations, optimisation, and the actual selection of the strategic asset allocation. More specifically, we focus on a model-based approach that offers the advantage of being efficient on the one hand, while also disburdening decision-makers of the need to make concrete investment decisions beyond specifying preferences and policy requirements.

Step 1: Risk tolerance and investment horizon

Central banks' investment principles, while expressing a general preference for prudent management, are often vague as to the exact degree of risk aversion. Ideally, within a quantitative asset allocation process, the investment principles are translated into a utility function. Various alternative forms of utility functions have been proposed in the literature. A standard example of such a utility function is quadratic utility over wealth, which implies a linear trade-off between expected returns and risks. From the perspective of keeping the investor indifferent, the trade-off between expected returns and squared volatility is governed by the risk aversion parameter λ. This parameter has to be specified by the investor, whereby a highly risk-averse investor has a high λ and a less risk-averse investor has a low λ. In practice, however, λ is difficult to calibrate. Although the literature gives indications for λ of an average investor, it is not obvious what the parameter would be for a more risk-averse investor such as a central bank.

Therefore, instead of formal utility functions, preferences are often expressed in the form of simple rules such as maximising the benchmarks' expected returns subject to a maximum risk tolerance. For example, one such rule requires expected portfolio returns to be maximised subject to the condition that there are no negative returns at a given confidence level. Again, the sponsor has to specify the degree of risk aversion, thus the confidence level. However, in this case there is an intuitive interpretation. Assuming a confidence level of 95% and annualised returns, losses in the reserves' market value are only tolerated in one out of 20 years. The confidence level can either be specified by the decision-makers directly or can be derived, under the decision-makers' guidance, from the institution's ability to bear financial

[9] Generally, these policy and operational considerations are difficult to capture in an unconstrained investment framework. As an alternative to imposing hard limits, Ramaswamy (1999) proposes a quantitative framework which determines, based on fuzzy decision theory, the currency distribution against acceptable ranges for the proportion for respective currencies.

[10] For example, a strategy aimed primarily at minimising transaction costs may stipulate the liquidation of the most liquid instruments first. Taking tail losses and sufficiency of reserves into account, however, the most appropriate strategy may be to sell less liquid instruments earlier in order to keep a "cushion" for periods of heightened stress (Duffie and Ziegler, 2003).

risks. In the latter case, the ability to bear financial risks would be determined by the bank's capital and provisions.

Analogous to the risk tolerance, the investment horizon for the reserves' strategic asset allocation is closely linked to the objectives and constraints laid down in the investment principles. If reserves are held to provide ready liquidity for financing foreign exchange policy operations, the investment horizon is related to the probability, timing and volume of such operations. Alternative objectives, such as the management of external debt, may require different specifications. In addition to these policy considerations, investment considerations also have an influence on the horizon. In particular, there might be a trade-off between the stability of the benchmark's main risk parameters over time, which is often a desired property from a decision-maker's perspective, and necessary reactions to changes in the underlying economic fundamentals.

Step 2: Formation of return and risk expectations

While asset allocation decisions are the most important determinant of the investment success, the expected return distribution is the most important driver of the asset allocation decision. In particular, the expected mean return is crucially important. For example, Chopra and Ziema (1993) show that, depending on the investor's risk tolerance, errors in expected returns are about ten times as important as errors in variances and covariances.

Regarding the nature of expected returns, it should be remembered that the primary objective of strategic asset allocation is not to outperform the markets but to find a risk-return combination that maximises the sponsor's utility. Thus it is not necessary to derive forecasts that are superior to those of other market participants and which could serve as a basis for the generation of excess returns, but only to derive expectations that reflect consensus views. Against this background, historic average returns might represent a starting point for the generation of return and risk expectations. However, the usefulness of historic returns as input for strategic asset allocation is limited with regard to two aspects. First, the gradual decrease of yields since the early 1980s in most markets results in a systematic overestimation of returns. Furthermore, the available historic sample might not be sufficiently long for some asset classes, or structural breaks in the data prevent the use of earlier observations.

Second, strategic asset allocation may aim to generate risk and return expectations that are forward-looking. This would take into consideration the growing literature on the predictability of bond returns (e.g. Fama and Bliss, 1987, Ilmanen, 2003) and on the relation between bond returns and the business cycle (see Fama and French, 1989). However, the expected returns to be used within reserves management at the strategic as well as tactical level should be generated independently of those macroeconomic assessments that are the basis for the central bank's monetary policy decisions. Superior investment performance should not rely on non-public information, and monetary policy decisions should not be revealed unintentionally to the markets when counterparties are in a position to extract information from readjustments in reserves portfolios.

In addition to the risk-return expectations on individual asset classes, assumptions about the effectiveness of active management are required. Ideally such expectations would be in the form of a trade-off between risk budgets available for active management, e.g. duration or tracking error limits, and expected outperformance. In practice such a trade-off might be difficult to estimate, as there is little experience with the performance of active management for alternative risk budgets. Therefore expectations often take the form of an expected outperformance for a given, constant risk budget. These expectations should be made against the level of market inefficiencies in each individual asset class, the set of portfolio management skills as well as the governance structure to carry out active management. A reasonable approach might be to extrapolate past outperformance at an unchanged risk budget. In any case, expectations should be made in close cooperation between risk and portfolio management.

Step 3: Optimisation

Those central banks that primarily care about risk-return efficiency in local or foreign currency frequently establish their strategic asset allocation on the basis of Markowitz's mean-variance analysis (Markowitz, 1952). That is, the derived expectations about the assets' risks and returns are transformed into an efficient frontier, which is a set of portfolios that maximise expected returns at each level of portfolio risk. In the context of an asset and liability approach, the analysis can be conducted within the familiar mean-variance space, but with two important modifications: asset returns are calculated as excess return over the liabilities, and a shortfall risk constraint is added. However, in the rest of this paragraph, we will focus on an asset-only perspective.

As an example, Figure 2 presents a possible process for deriving the efficient frontier and determining the optimal allocation for an actively managed reserves portfolio comprising US Treasuries of different maturities. The top left panel shows the efficient frontier for an annual analysis horizon, assuming that the investments are passively managed and that there are no

Figure 2: Mean-standard deviation analysis

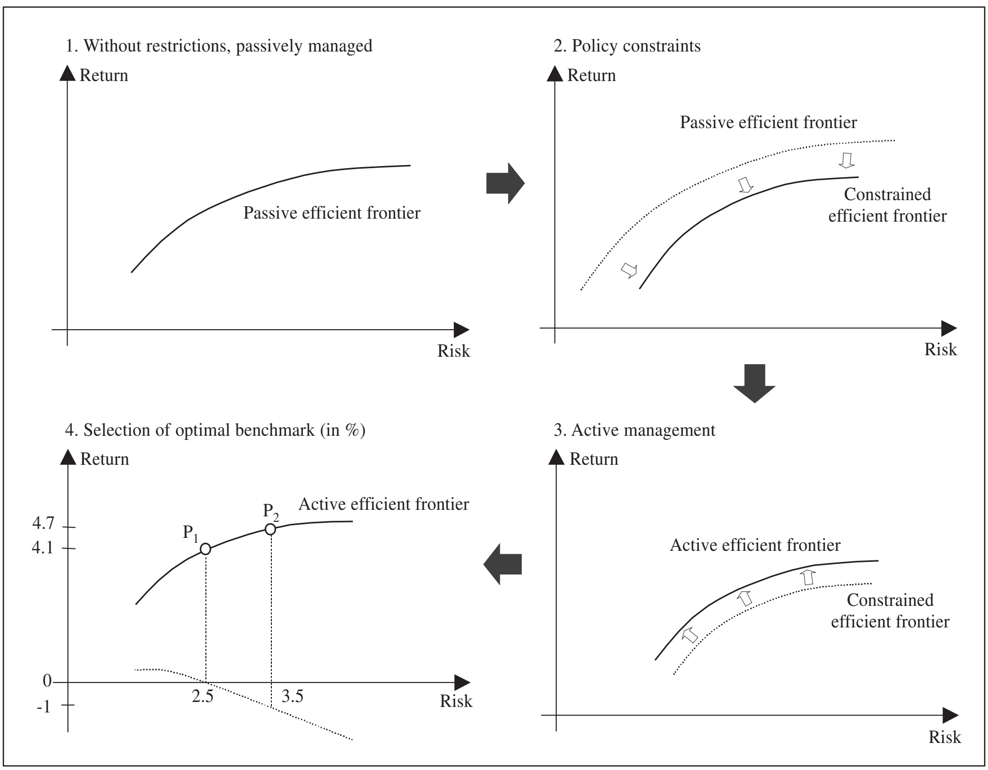

further constraints beyond short-sale restrictions. Subsequently, policy considerations can be taken into account. According to Panel 2, the introduction of liquidity constraints, such as minimum allocation to those maturity buckets of highest liquidity, shortens the frontier and shifts it downwards. In Panel 3 active reserves management is taken into consideration, which improves the feasible risk-return combinations. The objective is to assess to what extent active management improves the risk-return characteristics of the actual reserves portfolio by extrapolating, for example, realised outperformance and actual consumption of active risk in the past. From the active efficient frontier, the optimal asset mix can be identified based on the previously defined objective function.

To find the optimal portfolio we assume that the sponsor's preferences, in an asset-only perspective, can be represented by the rule of maximising expected returns subject to the constraint that there are no negative returns with a 95% confidence level. The dotted line in Panel 4 shows the implied minimum return at a 95% confidence level for each portfolio from the active efficient frontier. For example, portfolio P_2 has an annualised expected return of 4.7% at a volatility of 3.5%, which results in a minimum return of -1% at a 95% confidence level. Thus, this portfolio would be too risky assuming the above objective. On the other hand, portfolio P_1 offers the highest expected return while the minimum return still is positive. Thus the asset allocation underlying this portfolio could be the basis for the design of the strategic benchmark.

The approach outlined above does not consider risk-free assets. In the presence of such assets, e.g. Treasury bills with a maturity corresponding to the analysis horizon, the efficient frontier would be provided by a straight line starting in the location of the riskless asset, and be tangential to the curved line that comprises the efficient combinations of risky assets. In this case different levels of risk aversion will result in alternative combinations of the tangential portfolio and the riskless asset. Modigliani and Sutch (1966, pp. 183-84) point out that the identification of the riskless asset depends on the investment horizon. That is, an investor who cares about the return after *n* periods would consider an *n*-period government bond as a riskless asset. However, in the context of reserves management with uncertain and varying investment horizons, the identification of the riskless asset is not a clear-cut.

Step 4: Selection of a strategic asset allocation

The purpose of taking a quantitative approach to the investment process is not necessarily to follow the results of the optimisation exercise mechanically. Rather, the model-based framework provides a starting point for discussions among risk managers, portfolio managers and senior management on why and to what extent the final strategic asset allocation should deviate from the optimisation results. These discussions, for example, could introduce further policy requirements and try to balance additional constraints, which are difficult to capture within a purely quantitative framework. Part of this process is an assessment of implementational aspects of the envisaged asset allocation from a portfolio management point of view. In the end, against the background of the models' limitations, policy requirements and portfolio management considerations, the process should result in a strategic asset allocation with an overall risk-return profile senior management is most comfortable with.

This discussion can be supported by further quantitative analysis. Thus the risk-return characteristics of the model-based and possible alternative asset allocations are checked independently from the specific assumptions used in the optimisation exercise. For example, this validation process may comprise analysis of historic properties, stress testing, risk decomposition and an analysis of the potential impact of active management.[11] At the end of the validation process, a decision is made on a strategic asset allocation that best reflects the institution's investment philosophy, its risk tolerance and all operational considerations. Typically, the strategic asset allocation will be expressed in terms of a percentage weighting of asset classes as well as deviation ranges. The design of the strategic benchmark, and thus the implementation of the agreed asset weighting by means of market indices or notional portfolios, is outside the actual asset allocation process, and should be left to the implementation phase (as described in Section 2).

5 Conclusions

This paper has discussed a possible blueprint for the management of foreign reserves' strategic asset allocation. At the outset we addressed the importance of a sound organisational set-up. Actively managed reserves need a clearly segregated management of strategic, tactical asset allocation and actual portfolio management. The described three-tier governance structure comprising an oversight committee, investment committee and actual portfolio management is just one example of an organisational set-up that can meet these criteria. It facilitates a disciplined implementation of asset allocation decisions, improves accountability and promotes a risk-awareness culture across the organisation.

In a second step, the design of an investment philosophy is proposed, which translates general policy requirements into concrete objectives and constraints for day-to-day reserves management. Depending on the respective nature of policy requirements (e.g. supporting fixed exchange rates, maintaining the capacity to intervene under flexible exchange rates or managing external debt), a central bank may opt for alternative priorities. We have discussed how objectives and constraints are derived when the efficiency of individual currency sub-portfolios, of overall reserves, or of the bank's balance sheet is prioritised.

Finally, a quantitative framework for the actual asset allocation process was developed. It was argued that such a process might offer the advantage of being efficient on the one hand, but would also disburden decision-makers, who are responsible for the design of monetary policy, of the need to make concrete investment decisions beyond the specification of preferences and policy requirements. A comprehensive fulfilment of this mandate requires an effective IT infrastructure to generate timely reporting of performance figures and risk exposures. Furthermore, organisational provisions should be in place, which allow the escalation of any non-compliance to the decision-makers and the Oversight Committee. All this requires a certain degree of independence of the Bank's reporting and compliance functions from the bank's risk taking activities.

Despite having a clearly-defined investment philosophy, quantitative support for decision-making and a sound governance structure, people are key to success in strategic allocation as well as in the remaining stages of the investment process. The ability to recruit and retain highly qualified professionals will depend on the extent to which a challenging working environment is provided. To remain at the forefront of market practices, central banks have to continue promoting strong business ethics and encouraging staff to constantly update their skills set.

---

[11] Active management brings an additional dimension to reserve management by allowing diversification in investment strategies with an independent source of outperformance (Waring and Siegel, 2003).

References

Anson, M. (2004), "Strategic versus Tactical Asset Allocation", *Journal of Portfolio Management*, pp. 8-22, Winter 2004.

Chopra, V. K. and W. T. Ziemba, 1993, "The Effect of Errors in Means, Variances and Covariances on Optimal Portfolio Choice", *Journal of Portfolio Management*, pp. 6-11.

Chow, G. and M. Kritzman (2001) "Risk Budgets", *Journal of Portfolio Management*, Winter 2001, pp. 56-60.

Duffie, D. and A. Ziegler (2003): "Liquidation Risk", *Financial Analysts Journal*, May-June 2003, pp. 42-51.

Foley, T. (2003): "Strategic Benchmarks Are Almost Always Static and Should Be Almost Always Dynamic", State Street Global Advisors Research, July 2003.

Ibbotson, R. and P. D. Kaplan (2000), "Does Asset Allocation Policy Explain 40, 90, or 100 Per Cent of Performance?", *Financial Analysts Journal* 56 (1), pp. 26-33.

Ilmanen, A. (2003), "Expected Returns on Stocks and Bonds", *Journal of Portfolio Management*, Winter 2003, pp. 7-27.

IMF (2003), *Guidelines for Foreign Exchange Reserve Management: Accompanying Document*, International Monetary Fund, Washington.

Fama, E. F. and R. R. Bliss (1987), "The Information in Long-maturity Forward Rates", *American Economic Review* 77, pp. 680-92.

Fama, E. F. and K. R. French (1989), "Business Conditions and Expected Returns on Bonds and Stocks", *Journal of Finance* 54, pp. 1325-90.

Markowitz, H. (1952), Portfolio Selection, *Journal of Finance* 7, pp. 77-91.

Modigliani, F. and R. Sutch (1966), "Innovations in Interest Rate Policy", *American Economic Review* 56, pp. 178-97.

Ramaswamy, S. (1999), "Reserve Currency Allocation: An Alternative Methodology", Bank for International Settlements, Working Paper, No. 72, August 1999.

Waring B. and L. Siegel (2003), "The Dimensions of Active Management", *Journal of Portfolio Management*, Spring 2003.

Winkelmann, K., "Risk Budgeting: Managing Active Risk at the Total Fund Level", Goldman Sachs & Co. investment management research, February 2000.

Winkelmann, K. and R. Howard, "Developing an Optimal Active Risk Budget", Goldman Sachs & Co. investment management research, July 2001.

2

Thoughts on investment guidelines for institutions with special liquidity and capital preservation requirements

Bluford H. Putnam

Abstract

Management of foreign reserves must weigh the needs of liquidity and capital preservation during difficult market environments. Dividing foreign reserves into two sections, "liquid" and "liquidity-challenged", is recommended. Within the "liquid" section for capital preservation, the guidelines should encourage using assets that have return probability distributions which can be approximated with normal distributions. They should avoid embedded options, should not use credit quality constraints as a substitute for liquidity risk management, and should be willing to consider the use of special purpose investment vehicles that would fail "look-through" tests but offer special liquidity features. In the "liquidity-challenged" section of the central bank foreign reserves portfolio, used for long-term investment, more straightforward risk-return criteria can apply, which allows for long-term efficient allocation of scarce resources, including the ability to earn liquidity, complexity, or volatility premiums available to long-term investors.

1 Introduction

Certain types of institutions, especially central banks, have unique investment needs. Because of the nature of the role that central banks play in regulating the financial system and helping to insure low inflation and a stable currency, the investment objectives for the management of central bank foreign reserves portfolios must weigh the needs of liquidity during difficult market environments and the needs of capital preservation over both short and long-term time horizons as paramount concerns, rather than focus solely on return objectives. This means that for central banks, the objective of earning the maximum possible expected rate of return given a risk target is much too simple an approach to the development of investment guidelines.

Straightforward risk-return performance objectives have great appeal, and they are enshrined in the academic literature in the path-breaking portfolio optimisation work of Professor Harry Markowitz (1959) and the investment theory work of Professor William Sharpe (1964), among others. Unfortunately, straightforward return-to-risk measures fail to address the issues of market liquidity and capital preservation during short-term periods. The times of financial stress are when central banks may have critical national duties and special needs. So, central banks require a more complex approach to setting investment objectives. This also means that setting investment guidelines for the management of central bank foreign reserves is substantially more difficult compared to the portfolio management world of private sector financial institutions, pension funds, or other long-term pools of private sector capital.

Managing liquid capital for the long term, with multiple objectives and constraints, is a tough task in any market environment. What makes managing central bank foreign reserves over the long term that much harder is that one must prepare for a larger range of market environments, including some especially challenging ones, since the longer the time horizon, the more diverse the market environments one is likely to experience, for good and for bad.

The practical reality is that to manage long-term liquid capital successfully requires both more discipline and greater risk tolerance than many fiduciaries, trustees, boards of governors, regulatory overseers or supervisory boards often possess. To make the long-term capital management task more manageable and to protect the overseers or supervisory boards, the standard approach is to adopt a set of investment guidelines that give the appearance and hopefully the reality of providing a prudent framework for financial risk management in the wide variety of market environments that can and will occur over a long time horizon.

Unfortunately, many investment guidelines that appear "conservative" or "prudent" on the surface can work against the long-term risk and return objectives of the institution whose supervisory board has adopted and approved these guidelines. The reasons for this conflict between the appropriate desire to adopt a set of rules to prohibit investment decisions that might cause large losses, and the common outcome of a very poor performance record in terms of risks and rewards, not to mention the all too frequent financial disasters in certain sectors of the portfolio, is that there are a myriad of indirect effects that negatively impact portfolio management decisions which emanate from seemingly straightforward and "conservative" rules.

In considering central bank foreign reserves management, therefore, a critical element is to understand the interplay between investment objectives and investment guidelines. Commonly-employed investment guidelines need to be examined to identify those that may work unexpectedly and indirectly against the complex long-term investment objectives of these institutions.

The question of the short-term versus the long-term needs and requirements of central banks and official international agencies and institutions, in terms of their foreign reserves management, leads to the division of the foreign reserves portfolio into two sections: "liquid" and "liquidity-challenged", each with different investment objectives and different investment guidelines. Within the "liquid" section of the central bank foreign reserves portfolio, the investment guidelines should encourage the use of asset classes and securities that have return probability distributions which can be reasonably approximated with normal or log-normal distributions, should avoid securities with embedded options, should not use credit quality constraints as a substitute for liquidity risk management, and should be willing to consider the use of certain types of special purpose investment vehicles that would fail "look-through" tests, but do offer special liquidity and return features as a package. In the "liquidity-challenged" section of the central bank foreign reserves portfolio, which is most likely the smaller portion of the portfolio, more straightforward risk-return criteria can apply, and a variety of risk measures, stress tests and scenario tests should govern the portfolio rather than asset class or security-specific constraints. This allows a more efficient allocation of scarce resources, including the ability to earn liquidity, complexity, or volatility premiums which are only available to long-term investors.

2 The interplay of objectives and guidelines

The first step in any capital management task is to understand as clearly as possible the objectives and constraints associated with the needs and requirements of the institution. This initial process of identifying objectives and constraints cannot be underestimated for complex institutions such as central banks, which may have inherent conflicts in certain objectives and constraints that must be addressed at the outset to set objectives and then to construct investment guidelines in a proper manner for foreign reserves management.

Most central banks emphasise the complexity of their multiple duties to regulate and assist in the maintenance of a stable and secure national financial system, to maintain a low level of inflation, to influence as far as is reasonable the stability of the international purchasing power of their currency, and to promote the economic health of their nation. Regardless of the ranking in terms of priorities that a central bank sets for these critical national tasks, the combination of these tasks strongly argues that central bank reserves management must be handled in a manner that sets liquidity and preservation of capital as key constraints in determining the appropriate risks that the institution can tolerate while seeking long-term expected returns on its foreign reserves. Further complicating the process of setting risk and return objectives, central bank foreign reserves management processes may often be viewed as a model of prudent financial management. Hence, there is the need to be even more disciplined and to set forth investment guidelines with greater clarity than might be required of a private sector pension fund or financial institution which is not quite so obviously setting an example for the nation.

Specifically, central bank foreign reserves management must set investment guidelines that are consistent with their special tolerances for shorter-term risk of large losses (capital preservation) and the ability to access funds in a market crisis or period of national emergency (liquidity management). This means that central bank foreign reserves management objectives need to be stated explicitly in terms of the required liquidity needs of the institution and the maximum losses that are tolerable in any short-term period in the pursuit of long-term returns, before long-term return objectives can be set. Thus, before moving to the task of setting long-term return objectives and investment guidelines, we need to explore carefully the practical implications of putting a high priority on capital preservation and liquidity management needs, as these are the major elements of central bank foreign reserves management that are not generally considered by private sector portfolio managers with only simplified long-term risk-return objectives.

2.1 Liquidity management challenges

Liquidity management involves the ability to access capital for emergency purposes in large quantities during difficult and potentially turbulent market conditions. For a central bank, it serves no purpose to have plenty of liquid capital in good times, because these are exactly the times when it does not need liquidity. Central banks need their foreign reserves to be liquid exactly when national markets may be facing broad-based liquidity problems. Whether the cause of the lack of liquidity in national financial markets stems from international pressures on the currency, domestic loan losses in banking portfolios, or from some unexpected national disaster, the central bank must be able to serve as the lender of last resort, thereby supporting the security and stability of the nation's financial system. This may include a substantial call on foreign reserves to support the currency or the international borrowing capability of the nation.

This means that the percentage of investment exposures in the portfolio of the central bank's foreign reserves that are subject to liquidity difficulties in challenging market environments must be specified clearly. The percentage may, in fact, be relatively small. While there is no standardised approach for determining liquidity needs, central banks know that countries prone to fiscal budget excesses, ballooning trade deficits, large external debts, wide swings in political processes, or potential moves toward restrictive capital controls, can easily experience disruptive financial market conditions requiring the central bank to use its foreign reserves. In the end, the percentage of foreign reserves assets that must be in the "liquid" section of the portfolio and the percentage that may be allowed in the "liquidity-challenged" section of the portfolio will be set based on the considered judgement of the governing board of the central bank. In some countries, it may well be appropriate for the "liquid" section percentage to be 100 and the "liquidity-challenged" section to be 0, while other countries may be more comfortable with a 70-30 split (or even in a few countries a 50-50 split).

It is preferable to divide a portfolio into two sections with different investment objectives and guidelines to achieve mixed short-term and long-term objectives, rather than use a single portfolio with an overly complex set of investment guidelines that compromises the multiple purposes of central bank foreign reserves. One can more efficiently achieve the overall objective by having two "pure" plays and then spending one's time achieving the correct combination or percent weights of the two sections of the investment portfolio.

Some may think the preceding discussion mostly affects the central banks of emerging market countries, but this is not the case. The need for liquid foreign reserves is not at all limited to developing countries. One need only look at the currency intervention activities over the past 30 years of certain very large, politically stable and industrially mature countries to see that their central banks still require a high level of liquidity in their foreign reserves portfolio.

The "liquidity-challenged" section of the central bank foreign reserves portfolio, small as this section may be, should be allowed to focus on long-term, straightforward risk-return criteria with little or no constraints on asset classes or securities so long as reasonable overall portfolio diversification and risk criteria are met. The "liquid" section of the central bank portfolio, of course, is an entirely different story. Regardless of the percentage selected for the "liquid" section of the foreign reserves portfolio, the focus here is on setting appropriate investment guidelines for those parts of the portfolio that specifically deal with the challenges of liquidity and short-term capital preservation objectives. Critically, the investment guidelines associated with this "liquid" portion of the foreign reserves portfolio must take a realistic view of what types of investment exposures are prone to liquidity problems during difficult financial market environments. The answers to this may surprise some people.

(i) Credit ratings

In practice, high-quality credit ratings may not provide much assurance of liquidity in difficult market conditions. Liquidity is created by many market participants being willing to buy and sell a given security in a wide variety of market environments. Credit ratings do not address liquidity, because they are designed to answer the question of the probability of default, assuming the security is held to maturity.

A debt security, for example, can be reasonably and appropriately given a superior or "AAA" credit rating and still be exceedingly illiquid in a credit squeeze type of financial market environment. While this point has been demonstrated in virtually every credit crunch environment, one need only remember the huge price declines suffered by AAA and AA asset-backed and mortgage-backed securities in the midst of summer 1998 when Long-Term Capital Management (LTCM), the large hedge fund, collapsed.[1]

Many debt securities that are given high credit ratings are actually very complex, structured investment vehicles. The nature of the customised structure is often designed to provide (1) for a structured bond with a very high probability of the timely repayment of interest and

---

[1] See Bookstaber, 1999.

a full repayment principle at maturity, and (2) a residual equity tranche should any money be left over at bond maturity. Generally, there is a very illiquid portfolio of specific asset claims in these multi-part structures that have been packaged or pooled together. This illiquid portfolio is set as collateral against the interest and principal of the investment-grade bonds. Indeed, usually the portfolio must be over-collateralised to obtain the "AAA" or "AA" rating for the bonds. At maturity, once the interest and principal obligations of the investment-grade bonds have been paid off, what is left over belongs to the equity holders. The equity tranche in the structure is highly risky and may or may not have much value at the time of bond maturity, although in good times the residual equity value can be quite spectacular.

Investment-grade bonds in such complex special purpose vehicles will have more liquidity than the individual asset claims in the collateralised portfolio underlying the bonds. Nevertheless, bonds that achieve this high quality credit rating through an over-collateralised portfolio of illiquid asset claims are still likely to have only a limited number of market participants willing to do the extensive and expensive homework required to decide what is an accurate and fair price for this structured security, and then to monitor the security over time and different market environments.

The conclusion is that complexity in bond structure, regardless of a high credit rating, usually means a severe lack of liquidity in any kind of volatile or challenging market conditions. Thus, the portion of the portfolio of the central bank's foreign reserves that is set aside to meet liquidity needs should be constrained to straightforward securities of the "vanilla" variety. Structured debt securities, even if given an "AAA" credit rating, should be avoided in this part of the portfolio. By contrast, the long-term, "liquidity-challenged" section of the portfolio would allow for illiquid securities explicitly for the purpose of earning a long-term liquidity premium in return for taking the short-term risks of illiquidity.

(ii) Embedded options

Securities with embedded options are just as important as complex structured bonds, and should also be avoided in the "liquid" section of the central bank foreign reserves portfolio. Embedded options are quite common in many debt securities. Mortgage-backed securities originated in the United States and, with a credit rating of AAA or AA, often consist of mortgage pools that offer the borrower an option to prepay the loan at any time, usually without penalty. Many corporate bonds, including investment-grade bonds, include clauses giving the borrower a call option to prepay the loan at a specified price or under specified conditions. Corporate convertible bonds with equity option clauses attached give lenders the option to convert to equity at a given equity strike price.

Embedded options are a direct cause of liquidity problems in times of financial market uncertainty. The reason is very intuitive. Option pricing models require a reasonable estimate of the future price volatility of the underlying security to obtain a good estimate of the fair price of the option.[2] If future market volatility is unknown or hard to estimate, as it always is in difficult market environments, then any security with an embedded option will be just that much harder to price. Hard-to-price securities are never liquid in difficult market conditions. In challenging market conditions, complex securities trade only by appointment, and appointments may be very hard to arrange. Consequently, it should be forbidden to invest the portion of the portfolio of the central bank's foreign reserves that is held for liquidity reasons in securities with explicit or embedded options. Conversely, direct and embedded options are

[2] See Black and Scholes, 1973.

totally appropriate in the "liquidity-challenged" part of the portfolio in order to earn the volatility premium offered by options in that part of the portfolio that can handle fat-tailed risk on average and over time.

2.2 Value-at-risk and capital preservation

The special needs of central banks also means that they must manage the "liquid" section of their foreign reserves in a way that keeps the probability of large and unexpected losses very low even during short-term periods. This is a surprisingly difficult requirement for which to construct investment guidelines and still leave the portfolio with a reasonable opportunity to make money on average and over time. The best approach may be for the guidelines for the "liquid" section of the foreign reserves portfolio to exclude investment exposures known to be prone to low-probability large losses, or so-called "fat-tailed" risk.[3]

The term "fat-tailed" risk refers to a probability distribution of expected returns that differs in shape from a normal probability distribution because of the "fat" nature of the tails or extremes of the distribution. Many risk measures are in concept and practice based on the assumption of normal, or sometimes log-normal, probability distributions of expected returns. As the assumption of a normal distribution becomes less and less realistic, the ability of many traditional risk assessment techniques to estimate the probability of large losses is dramatically reduced, often to the point where reliance on such simplified risk measures is more likely to cause losses than prevent them. The problem is that when one uses a risk technique that is based on a normal distribution assumption, one can be lulled into a false sense of confidence as to both the likelihood and magnitude of large losses, and thus fail to take appropriate risk management steps to prevent the types of exposures and portfolio structures that could cause these large losses.

Any risk measurement approach that uses the standard deviation of historical returns as its key input into the risk assessment calculation should be treated as only being able to answer one type of risk question, namely: what is the likely range in which most return observations will occur? Typical standard deviation-based risk systems, such as the way Value-at-Risk (VaR) is commonly implemented and practised, provide a very good estimate of the frequency at which returns will fall between plus or minus one standard deviation of the expected mean of the distribution.

In the mid-1800s a famous Russian mathematician, Pafnuty Chebyshev (1821-1894) proved a theorem, known as "Chebyshev's Inequality", that explains why VaR is so good at answering the question of the frequency of returns that fall within a given range, and yet so bad at answering the question of just how large the losses can be if the return event falls outside the given range.

What Chebyshev proved was that one did not even need to know the shape of the probability distribution of expected returns to set forth upper and lower boundaries and to estimate the frequency that returns would fall in the range between the boundaries.[4] If one assumes a normal distribution of expected returns, then about two-thirds of the time series of return observations will be within plus or minus one standard deviation of the expected mean. If one drops the assumption of normality and admits total ignorance of the shape of the probability distribution, then using Chebyshev's Inequality the upper and lower bands are

[3] See Putnam, Wilford and Zecher, 2002.

[4] For the basic ideas of Chebyshev's Inequality, see Downing et al., 1997; for a more intense study of Chebyshev's body of work, see Kolmogorov (ed.), 1998.

widened relative to the normal distribution case to more like three-quarters or four-fifths for plus or minus one standard deviation. Even if one chooses not to make some adjustment to VaR to compensate for the possibility of unusually shaped probability distributions, one can still use VaR as a reasonable estimate for understanding the typical daily and monthly swings in portfolio value that will regularly occur about two-thirds of time.

Unfortunately, even if one knows with great confidence the upper and lower boundaries within which two-thirds or three-quarters of the return observations are going to fall, one can say precious little about just how large the losses can be in the other one-third or one-quarter of cases where the return observation fall outside the specified boundaries. In short, one needs to have a very good idea of the shape of the tails of the probability distribution of returns if one wants to estimate the likelihood of extreme losses and to guess at their magnitude.

This mathematical observation, well-known for 150 years, means that in the "liquid" section of the portfolio where one is most concerned about preservation of capital and liquidity, one may want to consider investment guidelines that exclude securities which come with difficult-to-estimate return probability distributions. This will focus the "liquid" section of the central bank foreign reserves portfolio on those types of securities that are much more likely to maintain a normal probability distribution of returns even in difficult market environments. One of the implications of adopting this conservative approach to investment guidelines in the "liquid" section of the foreign reserves portfolio is that VaR can represent a very effective measure of risk, given that the likelihood of fat tails has been minimised.

By contrast, the "liquidity-challenged" section of the central bank foreign reserves portfolio, where illiquid securities and options are allowed, would have much too much likelihood of fat tails to depend on VaR as traditionally practised as the sole measure of risk. Since VaR typically emphases the standard deviation as the key measure of risk, and embedded assumptions of normal distributions of expected returns, risk managers will need to supplement VaR measures to assess risk in the "liquidity-challenged" section of the portfolio. Thus, it is of paramount importance that central bank investment objectives and investment guidelines go well beyond VaR concepts in setting risk limits, which means employing a host of modern finance and statistics tools.

Monte Carlo simulations are extremely useful, so long as one does not defeat the purpose by incorporating a heavy burden of normality assumptions into the inputs of the simulation.[5] In addition, a wide variety of econometric applications, some based on behavioural finance, can be very useful for determining how markets may respond to different types of risk aversion assumptions.[6] Finally, dynamic Bayesian statistical processes have proven extremely useful in estimating the predictive standard deviation and correlations for a set of securities.[7] Measuring risk in complex portfolios with liquidity issues and direct and embedded options requires a much more extensive discussion than can be covered here, but the message is clear: do not depend solely on one measure of risk in this part of the portfolio.

(i) Equities

As in the liquidity management discussion, options present special challenges, and are a problem because they are structured with the intent to create a return probability distribution at maturity that is asymmetric. Furthermore, as noted earlier, since one needs an estimate of

---

5 See Handzy, 2003.
6 See Campbell, Lo and MacKinlay, 1997.
7 See Zellner, 1997, as well as Quintana and Putnam, 1996.

the underlying securities' future return volatility to determine the fair price of an option, in any market conditions in which volatility itself is hard to predict or estimate, then option prices are likely to swing widely and not follow anything remotely approaching a normal distribution.

Securities with direct options or embedded options are not, however, the only types of securities that are vulnerable to "fat-tailed" risk and likely to invalidate the assumption of normal return probability distributions. Equities are also a prime candidate.

The likelihood of asymmetric return distributions for equities is easily understood once one realises that equities can be analysed as a call option on the value of the assets of the company. One pays a price for corporate ownership in a legal structure that limits the losses and liability of the investor to the size of the initial investment, with unlimited upside if the company can put together a profitable, long-run growth path. This is the shape of the payout profile for the purchasers of a call option, where one of the key pricing factors is the ability to estimate the volatility of future earnings of the company.[8]

As the future earnings stream becomes more predictable, the call option (owning equity) becomes more and more "in-the-money" and the stock price follows a more normally distributed return pattern. For new companies or companies whose industries present special challenges, then the volatility of future earnings may be very hard to predict. This means the pattern of stock price returns will behave much more like a call option (asymmetric pattern) than a normal distribution.[9]

For this reason, equities are generally not an appropriate investment for the "liquid" section of the central bank's foreign reserves portfolio aimed at capital preservation. Please note that this conclusion is not based on the risk of equities, but rather on the inability to assume that a normal distribution applies to the expected returns, so that the probability of large losses and their associated magnitude is virtually impossible to estimate.

(ii) Risky securities are allowed

While this discussion has highlighted what types of securities and exposures should be excluded from that portion of the central bank's foreign reserves portfolio that is allocated for liquidity management and capital preservation, one also needs to explore what types of securities are appropriate. Based on the discussion here, the key criterion is the confidence that one can make the assumption of a normal probability distribution of expected returns.

This means that the actual level of risk for any given security or asset class is not the problem or the issue, even if the asset class or security is considered extremely risky. If one can be reasonably confident that the probability distribution of expected returns is not too far away from a normal distribution, and if one can reasonably estimate the standard deviation of those expected returns, then the actual level of risk can be managed by adjusting the ratio of the risky asset to risk-free cash. This is an absolutely critical point for setting portfolio management guidelines.

In the "liquid" section of the central bank foreign reserves management, one needs to focus first on the predictability of the shape of the probability distribution of returns. The level of risk is easy to adjust upwards or downwards for the whole portfolio, using more or less cash. Adjusting the risk level, however, can be done with confidence only so long as the risk in the extremes of the probability distribution of expected returns is relatively predictable. And

---

8 See Smithson and Smith with Wilford, 1998.
9 See Norland and Wilford, 2002.

since the special concerns of central banks place extra emphasis on avoiding liquidity problems, even in rare cases of unexpectedly large losses, it is even more important for central banks to focus on the predictability of the probability distribution of returns and not on the level of risk involved in a given security. The latter is especially important, as focusing on the level of risk can create enormous problems by giving a false sense of confidence that the magnitude of those rare unexpectedly large losses is somehow predictable even when the shape of the probability distribution is unknown.

(iii) Benchmark Bonds

Following this general observation, the portion of the central bank's foreign reserves set aside for liquidity and capital preservation therefore needs to include securities such as short-term government debt instruments, medium and long-term investment-grade bonds that do not include embedded options and do not exhibit unreasonable liquidity traits. These would include all non-callable bonds that serve as "benchmarks" for the government, international agency or private debt markets, as benchmark status conveys a certain confidence in liquidity and assures more active trading even in tough market conditions.

(iv) Currencies

Interestingly, the list of acceptable securities should also potentially include short-term currency exposures for the mature economies of the world, such as currency futures and forward contracts. Countries with mature and diversified economies tend to provide the appropriate underlying economic fundamentals that are necessary for one to make a confident assumption that the expected return pattern of short and long currency exposures will follow a normal distribution. The currencies of smaller countries with economies that are highly tilted to commodity production, such as mining, energy, or agriculture, would generally not qualify, nor would the currencies of countries in which democratic principles for periodically changing the leadership of the country are not enshrined.

One implication of this analysis that will be explored later is that the foreign reserves portfolio can be expanded to include significant exposures to currencies not generally considered reserve currencies. The question here is how much exposure can be allocated away from the euro or the US dollar as the primary international reserve currencies. So long as liquidity concerns are met, and they usually are for a small set of major currencies, then there may be a place for more active currency portfolio management than many central banks now allow in their foreign reserves portfolio. This issue will be covered in more detail when we move to the practical implications of setting guidelines for the "liquid" section of the portfolio.

(v) Commodities

Commodities and commodity futures contracts (but not options contracts) have also become good candidates for inclusion on the list of securities with more or less predictable probability distributions. Cash markets and futures contracts-based energy, metals, and agricultural products have displayed some extreme price movements in the past, but as suppliers have become more diverse and scattered across the globe, and as modern technology has improved the stability of production processes, the probability distribution of the expected returns on commodities can now be argued to resemble more closely a normal distribution than, say, the probability distribution for high-yield bonds or even some equities for new and developing companies. The point is that the very high risk of commodities can be managed with cash offsets, and the maturity and liquidity now provided in certain commodity futures contracts can make them useful diversification instruments, so long as one can become comfortable with the forward-looking view that the return probability distributions will more closely approximate normal distributions in the future, regardless of the level of risk. An interesting implication of this analysis is that gold, which is often owned by central banks, should probably be more actively traded along with other commodities, rather than held as a barren asset in the portfolio.

(v) Special purpose investment vehicles

On the surface of this discussion, it might seem that almost all structured investments would be excluded from the "liquidity" section of the central bank foreign reserves portfolios owing to liquidity concerns. This conclusion, however, should not be adopted without addressing the ability of special purpose investment vehicles to contain explicit liquidity or capital preservation provisions. Central banks should focus on the whole of the structure of the special purpose investment vehicle, regardless of the securities and exposures contained within the special purpose vehicle.

What some times occurs with highly restrictive investment guidelines is that they require a special purpose vehicle to be analyzed based on the individual components in the underlying portfolio. This is known as a "look through" provision in the guidelines. If the portfolio that is held inside a special purpose vehicle includes a prohibited security or type of exposure, then the whole special purpose vehicle is deemed an unacceptable investment. This "look through" approach should always be avoided, since it ignores the possibility that certain provisions of the structure may be able to compensate for the risks taken within the portfolio and make the whole investment very attractive and useful for foreign reserve management. Specifically, special purpose vehicles may come with provisions for accessing liquidity in emergency cases or they may have some form of protection from capital losses if the investment is held to maturity. Any special purpose vehicle with these provisions should be allowed in the "liquid" section of the reserve portfolio.

There are some caveats. In cases where liquidity provisions are attached to the investment structure, the liquidity provisions are only as good as the assets in the underlying portfolio or the borrowing capability of the institution standing behind the liquidity provisions. Nevertheless, when liquidity provisions are appropriately structured and backed by a structure or institution of investment-grade credit quality, then portfolios structured in this way may well allow central banks to be invested in a type of security that would otherwise fall outside the liquidity constraints of the institution.

The critical point here is that the guidelines for purchasing structured investments, such as specially managed portfolios or pools of securities, should allow the structure to meet the liquidity requirements rather than requiring the securities within the structure's underlying portfolio to meet the liquidity requirements. That is, "look-through" guidelines on the underlying portfolio of the special purpose vehicle need to be avoided, as they may prevent the central bank foreign reserves portfolio from taking exposures that offer appropriate return expectations relative to the risks, once the whole structure, with borrowing capabilities and capital preservation clauses, is considered.

Liquidity protection, such as the ability to redeem at short notice or to borrow against the net asset value of the structure, does not however answer the equally fundamental question of the predictability of the probability of returns of the underlying portfolio. Suppose the underlying portfolio contains securities that have liquidity issues and may involve direct or embedded options. One must recognise that a portfolio option will have a very different value and probability distribution of its returns than options on a portfolio.[10] If the underlying portfolio within the structure is well-diversified, it can conceivably go a long way towards achieving some degree of predictability of the return distribution even if the portfolio contains some direct options, embedded options, or illiquid securities.

In cases where reasonable comfort as to the predictability of the probability distribution of expected returns of the whole structured portfolio cannot be attained, then the underwriter of the structure should have the ability to add portfolio insurance clauses or to provide some measurable degree, if not necessarily 100% protection, against losses. For example, a structured portfolio that has, say a 90% guarantee of return of invested principle at maturity date, and also comes with a borrowing facility based on the initial investment amount, may now meet the criteria for inclusion in a portfolio of central bank foreign reserves, regardless of the nature of the securities that are included in the underlying portfolio.

This is a special case that needs further investigation. If the underlying portfolio contains securities that have liquidity issues or asymmetric return distributions, and the structure (backed by a high credit quality institution) comes with a borrowing facility to remove liquidity problems and portfolio loss protection (at maturity) to make the probability of a large loss predictable, then this structured investment is likely to have an asymmetric distribution of expected returns (i.e. the probability of losses has been eliminated). However, it may well have given up a significant part of the underlying expected return in the fees and costs associated with setting up the liquidity and portfolio guarantees for the structure. In this case, investment in these types of structures should not run afoul of investment guidelines aimed at limiting losses, although they might be a problem in terms of earning reasonable returns over time.

Risk guidelines are commonly established to prevent large losses and to set a reasonable expectation of the volatility of returns in the typical market environments. However, central banks need to be especially aware that structures that solve the issues associated with the risk guidelines may also eliminate opportunities for profits (that is, they can bleed to death, instead of exploding). It therefore falls on the investment committee or portfolio manager within the central bank to determine if the expected profits of a structured investment are still reasonable after the costs and fees of the structuring have been deducted from the gross expected returns of a portfolio that, without the structure, would not meet the risk guidelines. This is not an easy determination, but properly constructed risk guidelines should at least offer portfolio managers the choice of such investments.

3 Guidelines in practice

With the background and insights from this discussion on how to handle the special needs of liquidity management and capital preservation involving setting the risk guidelines for central bank foreign reserves management, we can now examine common guidelines to see whether they help or hurt risk management. In particular, we want to analyse credit quality requirements, leverage prohibitions, asset class restrictions and currency management constraints, among others. Rules involving each of these issues are typically included in the investment guidelines for central bank foreign reserves management.

[10] See Bookstaber, 1996.

3.1 Credit

Restricting the portfolio to investment-grade rated credit securities is a typical investment guideline, and solves one important issue. Low credit quality debt securities are more like short put option exposures than anything else. That is, the lender or owner of the high-yield bond (low credit quality bond) has given the borrower the right but not the obligation to put the underlying assets to the lender and go into bankruptcy in extreme cases.[11] Being short on a put option means that one has all the downside with only limited upside potential, and the distribution of expected returns is sharply skewed to the left (losses) with very fat tails.

Investors willingly buy these high-yield fixed income securities because the expected return, on average and over time, can be extremely high precisely to compensate for the low (and hard to estimate) probability of extreme losses. Insurance companies take risks like this all the time, and they compensate in part by constructing large pools of diversified risks so that the return distribution of the whole portfolio is much more predictable than that of any one (asymmetric) security. On the whole, though, low credit, high-yield bonds would fall well outside the liquidity and capital preservation needs of central bank foreign reserves management since their return distributions are potentially skewed toward low probability large losses.

Credit quality restrictions have a potentially negative impact on portfolio performance, however, because they are often viewed a substitute for liquidity concerns and for capital preservation concerns. As discussed earlier, a high credit rating implies absolutely nothing about market liquidity, and can give the investor an totally unwarranted comfort level about future expected returns if the security contains embedded options. Thus, the problem with constraining portfolios to contain only investment-grade credit quality is that it does not address liquidity or the predictability of returns associated with embedded options.

In the "liquid" section of the portfolio, central banks should consider limiting the portion of their portfolio that can contain any securities with embedded options (i.e. prepayment options, call options, conversion options, or other option-type mathematical formulas determining returns) regardless of credit quality. And, in that part of the portfolio that does allow securities with embedded options, the risk assessment techniques and monitoring processes must focus on stress tests and scenario testing, and must avoid relying on standard deviation-based risk measures, such as commonly-practised VaR approaches.

Central banks should also consider allowing structured portfolios that are wrapped with principal protection guarantees and borrowing facilities from high-quality credit institutions. That is, the investment guidelines for the securities in the foreign reserves portfolio should not be automatically extended to structured investments, if those special purpose investment structures can meet reasonable liquidity and capital preservation requirements, regardless of the securities in the underlying portfolio.

3.2 Leverage

Leverage is viewed by many conservative investors as dynamite, and they would accordingly prefer not to get anywhere near it. This view should however be reconsidered. Modern portfolio management techniques and portfolio risk monitoring systems have made leverage a meaningless concept because leverage can no longer be defined without significant controversy. The intent of leverage constraints was always to put a cap or upper boundary on

[11] See Kealhofer, 1999.

the risk level of the whole portfolio. This made a lot of sense when methods of risk measurement were in their infancy or when theoretical measures of risk were impossible to implement owing to technological or information constraints. However, this is no longer the case, and thus investment guidelines are more appropriately focused on the real objective, setting a maximum risk level rather than by constraining the use of leverage.[12]

The reason leverage is no longer measurable has to do with the use of futures, forwards, options, and various hedging techniques. Take a two security portfolio based on capital of 50 units, for example, with (a) a 100 units of long exposure in the (FX hedged) Australian stock market obtained through buying futures contracts that required a margin deposit of 10% of the notional value of the underlying security, and (b) a short exposure of 100 units in the German stock market also obtained by selling a futures contracts that required only 10% of the notional value of the underlying security as the margin requirement. Now let us calculate leverage.

One view of leverage is that of adding up the gross exposures, which in this case would be 200 units divided by the 50 units of capital, yielding a leverage ratio of 4.

Another view would be to take the average correlation of the stock market returns from the Australian stock index futures contract and the German stock futures contract over a historical period, and then use that to adjust the leverage calculation. In the extreme, if the correlation was positive and perfect, that is a correlation coefficient of positive one, then the positions would effectively collapse on each other for zero exposure, producing a leverage ratio of zero.

Alternatively, one can also observe that some calculations of leverage focus on the margin requirements of futures and options contracts. This would yield a leverage ratio of 2:5, indicating that the margin rules allow for much greater exposure than is being used in the portfolio at this time.

More realistically, a mean-variance approach using modern portfolio theory could be invoked as part of assessing leverage in a level of risk sense. Assuming a positive correlation of 0.5, and volatilities of 20% (annualised standard deviation) for both securities, then using a mean-variance approach the total risk of the combination of the short and long exposures in this case would be more or less in the range of 50 units, meaning that this relative value trade would have about half the risk of a directional exposure in either stock market. This means the leverage ratio would be about 1. Assuming a zero correlation would yield a risk assessment that the long-short combination had about a 7:5 leverage ratio. Assuming a perfect correlation of 1.0 would yield no exposure at all, the “non-existent” perfect hedge case.

Of course, if the portfolio had used a call option on one of the stock markets, then the leverage ratio would be a moving target depending on the relationship of the option to its strike price.

All of this may sound confusing, but that of course is the point. Calculating leverage requires a very high burden of assumptions about correlations and volatilities if it is to be consistently defined. Moreover, different assumptions may be more or less appropriate depending on the risk question being asked. All of this means that simple leverage definitions have absolutely no place in investment guidelines. The old and traditional definitions of leverage will rule out too many effective trading strategies and risk-reducing portfolio management techniques. Thus, leverage guidelines will generally lead to more risk and lower returns that could be obtained by using more sophisticated measures of risk, such as VaR, so long as the likelihood of fat tails has been minimised.

[12] See Norland and Wilford, 2002.

For example, not allowing for hedging techniques or some degree of offsetting exposures will greatly overstate leverage and actually force the portfolios to take more risks than they should, since hedges and offsetting exposures are penalised to the detriment of the portfolio. Leverage restrictions are outdated and need to be replaced by risk level limits with risk measurement systems appropriate to the securities in question.

This last caveat means that for securities with a reasonable expectation of a normal distribution, then VaR limits make sense, and if there is reason to suspect fat tails, then scenario and stress tests must replace the VaR calculations as the primary focus for setting risk limits.[13] Stress tests need to examine cases where correlations shift unexpected to values closer to one, which is common within a given asset class in a crisis period, as well as using a zero correlation assumption to guard against assuming that the hedging or offsetting techniques are much better than they really are. Hedges are never perfect, and in a crisis some hedging approaches break down, which is why the zero correlation stress test is important.

3.3 Asset class restrictions

Some asset classes are viewed as riskier than others. This is certainly correct. Since the overall risk level can easily be adjusted by adding cash to the portfolio to lower risk, there is no justification for restricting asset classes based on inherent risk levels. Again, the proper issue to focus on is the confidence with which one can estimate the shape of the probability distribution of returns.

Asset classes, or possibly sectors within asset classes, should be excluded from the "liquid" section of the central bank foreign reserves portfolios only when there is a very low predictability of the shape of the probability distribution of expected returns, or when there is a high confidence that the shape of the probability distribution is asymmetric in the extremes of the distribution. Asset classes such as equities probably do not meet the capital preservation needs of central banks, while low-quality corporate bonds have too much implied optionality in their return profiles to meet either the liquidity management or capital preservation requirements. By contrast, non-callable private sector notes and bonds, government-issued or government agency-issued notes and bonds without embedded options, and short-term foreign currency exposures from futures and forward contracts, as well as certain commodities, are all well within the scope of central bank foreign reserves management.

Furthermore, asset class restrictions should not apply to special purpose investment vehicles or structures in a "look-through" method, where the structure is excluded from the portfolio only on the basis of the securities held within the structure. The point here, discussed in detail earlier, is that if the whole structure contains clauses and covenants from an investment-grade institution that address the liquidity management and capital preservation needs of central banks, then it is appropriate for such structures to be considered for potential inclusion in the foreign reserves portfolio, regardless of the securities held inside the structure.

3.4 Currency management

Currency investment programmes are often excluded from central bank foreign reserves portfolios. In some cases, hedging is allowed only for the existing currency exposures contained in a global bond portfolio. This is much too restrictive and removes the ability to

---

[13] See Putnam, 2000.

earn reasonable risk-adjusted returns from currency management activities, which typically are highly liquid and not as prone to "fat-tailed" risks as certain other investments. Based on the liquidity management and capital preservation issues raised in this article, which focuses on the predictability of the shape of the probability distribution of expected returns, short-term currency futures and forward contracts associated with the currencies of mature economies are appropriate investments from a risk management sense.

This means that the euro and the US dollar, as the primary international reserve currencies, need not be the only currencies included in the foreign reserves portfolio. Significant diversification can be obtained by adding a few more currencies with excellent liquidity characteristics, such as the British pound, Swiss franc and Japanese yen, as well as a growing number of other currencies that are developing good liquidity. In addition, it is critical to note that so long as the overall risk of the "liquid" section of the portfolio is within the guidelines, there is no financial reason not to allow short positions in liquid currencies expected to depreciate relative to others. Of course, central banks may not wish to be put in the position of actively trading, even going short, in the currency of a close trading partner and political ally. The political and diplomatic dimension of portfolio management is totally at the discretion of the central bank's governing board, but from a financial perspective there is no reason not to include some active currency trading in the foreign reserves portfolio.

Another use which goes beyond that of the traditional international reserve currencies is how to hedge (or not) a multi-currency bond portfolio. In general, hedging restrictions on currency overlay programmes associated with global bond portfolios or global bond benchmarks are not a good idea, given the analysis presented here. It is common for a currency overlay programme to be designed to mitigate the currency risk associated with owning foreign currency-denominated bonds, and even to add expected profits to the foreign bond portfolio. A common investment restriction in currency overlay programmes allows only for the decision whether to hedge, partially hedge, or not hedge at all the currency exposure from a given foreign currency-denominated bond. This is too restrictive and violates many of the concepts of modern portfolio risk management.

A better approach is to set a risk target or upper bound estimated risk limit for the total risk of the currency portfolio based on some measure of VaR or a standard-deviation method of assessing risk. VaR is appropriate here, as are other standard deviation methods of assessing risk, since the set of currencies included in the investment programme are only from mature countries, and the exposures can be reasonably assumed to have an expected return distribution that does not differ too far from a normal distribution.

In practice, this means that if the British gilts, for example, were allowed in the bond portfolio, then the currency overlay portfolio could be long, short, hedged, un-hedged, over-hedged, etc., in that any exposure, long or short, to the British pound would be allowed so long as the overall risk of the currency portfolio is within the overall risk guideline. What this does is free the hands of the currency portfolio manager to design a long-short currency portfolio that is not pre-constrained and radically inhibited by the design of the global bond portfolio. In terms of modern portfolio theory, the currency manager can now employ cross-trades and correlated hedging strategies to reduce risk or enhance returns based on the overall structure of currency volatilities and correlations. This is a much more efficient process in terms of long-term return expectations relative to a risk target than the alternative of not allowing cross-rate hedges and various relative value strategies based on interest rate differentials and expected correlations and volatilities.

3.5 Duration Limits on Bonds

Duration calculations on bonds and bond portfolios are extremely helpful in understanding how bond prices will move upon a given change in interest rates. Duration calculations, however, are not a substitute for a VaR type of analysis. For example, government bonds from two different countries can have the same coupon, the same coupon payment dates, and the same maturity date, and still have different risks even if the duration calculation for each bond is exactly the same. The reason is that interest rates in one country can easily be more volatile than in another country, especially if the currency markets are volatile. Thus, duration limits are a crude way to constrain risk in a bond portfolio. They do not do a good job at risk control, while eliminating some less risky and potentially profitable trades. A better solution is to use duration analysis on the portfolio manager's desk to choose securities and to assist in scenario and stress test analysis. VaR measures can then be used to set an overall risk limit for the bond portfolio, or better still for the section of the whole portfolio that allows currency position risk and bond position risk to be evaluated as a package.

4 Central bank investment guidelines

Owing to the special needs of central banks for liquidity management and capital preservation, it is critical for central banks to focus their risk management guidelines on much more than just risk-return objectives. Focusing on liquidity issues highlights the special problems of complex securities, while focusing on capital preservation highlights the question whether the probability distribution of expected returns from a given security, asset class sector, or structured investment can be reasonably approximated by a normal or log-normal distribution, or whether it is reasonably predictable in the extreme tails. These concerns, when followed through in the analysis, make an argument for dividing the foreign reserves portfolio into two parts: a conservative, short-term portfolio known as the "liquid" section, and a long-term, more aggressive portfolio referred to here as the "liquidity-challenged" section. Each section will have its own return objectives, its own investment guidelines, and its own risk monitoring and risk management methods and processes.

For the short-term "liquidity" section of the central bank foreign reserves portfolio, special attention needs to be paid as to whether complex securities with known liquidity issues or securities with embedded options should be allowed at all. Complex securities almost always involve liquidity risks. It is well-known that options and securities with embedded options have asymmetric expected return probability distributions. VaR types of risk measures based on standard deviation measures of risk may be inappropriate for determining whether these investments are acceptable for use within the portfolio. Indeed, the preferred course is to focus the central bank foreign reserves portfolio only on those securities, sectors of asset classes, and special purpose investment structures that can meet with a reasonable degree of confidence the special liquidity and capital preservation needs of the central bank.

Interestingly, following this logic means that central banks can and should use currency management programmes and certain commodity investment programmes, in addition to traditional investment-grade fixed income programmes. As already noted, fixed income investment programmes need to prohibit or significantly constrain the use of fixed income securities with embedded options regardless of the credit rating of the security, owing to liquidity concerns that are simply not addressed by credit quality guidelines. Indeed, investment guidelines concerning credit quality must not be used as a substitute for avoiding liquidity problems, as credit ratings do not speak about liquidity issues at all. And finally, investment guidelines in the "liquid" section of the portfolio should avoid "look-through" provisions to the underlying securities in structured investment vehicles, and instead focus on whether the structure, viewed as a whole, has sufficient additional provisions to meet the liquidity management and capital preservation needs of the central bank.

The long-term, "liquidity-challenged" section of the central bank foreign reserves portfolio can and should take much more risk while seeking higher returns on average and over time. This means that certain types of securities that should probably be prohibited in the "liquid" section are actually quite appropriate for this section. Indeed, in many circumstances markets pay a liquidity premium, a complexity premium, or a volatility premium, and this is the part of the reserves portfolio where these returns can be earned.

When the long-term "liquidity-challenged" section of the portfolio includes substantial positions involving liquidity issues and embedded options, then the likelihood of fat-tailed risk is very high. This means it is critical to realise that straightforward VaR measures are not sufficient for risk monitoring or risk management. Multiple risk measures must instead be employed, including stress tests and Monte Carlo methods, among others. The requirement to use much more sophisticated risk measures and monitoring processes is simply part of the cost of investing in more complex portfolio strategies. Moreover, when evaluating the long-term returns on this complex part of the portfolio, it is critical to adjust those returns downwards for the higher management and monitoring costs with which the central bank is burdened. These are still worthwhile investments, but management and monitoring costs are non-trivial in complex portfolios and need to be part of the process of measuring returns.

All of this leaves the critical question of how to divide the central bank reserves portfolio between "liquid" and "liquidity-challenged" sections. In the end, this critical decision constitutes a judgement call for the central bank's governing board. By focusing the choice between these two distinct parts of the portfolio, however, one can ensure liquidity and capital preservation needs more directly, while more efficiently earning higher returns on average and over time in the more aggressive section of the portfolio.

References

Black, Fischer, and Scholes, Myron. "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* (May-June, 1973), pp. 637-54.

Bookstaber, Richard M. *Option Pricing and Investment Strategies,* Probus Publishing, 1996.

Bookstaber, Richard M. "A Framework for Understanding Market Crisis," *Risk Management: Principles and Practices,* AIMR Conference Proceedings, Association for Investment Management and Research, 1999, pp. 7-19.

Campbell, John Y., Lo, Andrew W., and MacKinlay, A. Craig. *The Econometrics of Financial Markets*, Princeton University Press, 1997.

Downing, Douglas, and Clark, Jeffrey. *Statistics: The Easy Way,* Barrons Educational Series, 1997.

Handzy, Damian. "Monte Carlo Techniques," *Improving the Investment Process through Risk Management,* Association for Investment Management and Research (AIMR), Charlottesville, Virginia, 2003.

Kealhofer, Stephen. "Credit Risk and Risk Management," *Risk Management: Principles and Practices,* AIMR Conference Proceedings, Association for Investment Management and Research, 1999, pp. 80-90.

Kolmogorov, A. N. *Mathematics of the 19th Century; Function Theory According to Chebyshev, Ordinary Differential Equations, Calculus of Variations, Theory of Finite Differences,* Springer Verlag, 1998.

Norland, Eric, and Wilford, D. Sykes. "Leverage, Liquidity, Volatility, Time Horizon, and the Risk of Ruin: A Barrier Option Approach," *Review of Financial Economics*, Volume 11, Number 3, 2002, pp. 225-39.

Markowitz, Harry. *Portfolio Selection: Efficient Diversification of Investments* (John Wiley & Sons), 1959.

Quintana, Jose M. and Putnam, Bluford H. "Debating Currency Market Efficiency Using Dynamic Multiple-Factor Models," *American Statistical Association 1996: Proceedings of the Section on Bayesian Statistical Science.* 1996, pp. 55-60, with Jose Mario Quintana.

Putnam, Bluford H. *Integrating Risk Management with Asset Allocation* Global Investor Euromoney plc., 2000.

Putnam, Bluford H., Wilford, D. Sykes, and Zecher, Philip D. "A Short Note on the Concept of Risk Management and VaR for Asset Management Firms," *Review of Financial Economics* (Volume 11, Number 3, 2002), pp. 205-12.

Sharpe, William F. "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk," *Journal of Finance* (September 1964), pp. 425-42.

Smithson, Charles, and Smith, Clifford, with Wilford, D. Sykes. *Managing Financial Risk: A Guide to Derivative Products, Financial Engineering, and Value Maximization* (McGraw-Hill), 1998.

Wilford, D. Sykes. "Risk Measurement versus Risk Management," *Improving the Investment Process through Risk Management,* Association for Investment Management and Research (AIMR), Charlottesville, Virginia, 2003.

Zellner, Arnold. *Bayesian Analysis in Econometrics and Statistics* , Edward Elgar Publishers, 1997.

3

A framework for strategic foreign reserves risk management

Stijn Claessens[1], Jerome Kreuser[2]

Abstract

We present a framework for active foreign exchange reserves management that integrates risk-return objectives with macroeconomic, macro-prudential and sovereign debt management concerns. Our framework allows for very general objective functions, does not restrict the class of eligible stochastic processes or limit the investment universe, and incorporates many types of macroeconomic concerns. Feedback between outcomes and decisions is easy using tools that reshape distributions and functions of the outcomes. Moreover, the model can be run on a PC-based platform. We apply the framework to several common reserves management problems and also show how it provides institutional guidance through developing benchmarks, portfolio evaluation criteria and management reporting.

1 Introduction

Central bank foreign exchange reserves risk management involves balancing many objectives and issues, from broad macroeconomic policy objectives, such as monetary policy and foreign exchange management, to more micro aspects, such as defining portfolio benchmarks and evaluating investment managers. Furthermore, constraints arising from legal and human resources, asset markets, institutional and other aspects affect the actual achievability and implementation of reserves management objectives. While the macroeconomic aspects of reserves management have been much analysed, they have continued to attract attention in recent years, especially the relationship between financial crises and reserves and debt management. Many crises in the last decade have underscored possible relationships between foreign exchange reserves and public debt management, and have shown how private debt management can exacerbate economic or financial shocks. The micro aspects of improving reserves management have also received much attention, thanks to the wider investment universe and the broader range of financial tools that have become available. These developments have led to more active reserves management strategies and methods, as applied by private asset managers, finding their way into foreign exchange reserves management.

While these macro and micro issues have been under consideration for some time and have been analysed in much detail in various strands of the literature, various aspects are nevertheless often not considered jointly, at least not within one common analytical and empirical framework. The analysis of tactical investment management, for example, often

---

[1] Professor of International Finance Policy, Department of Economics and Econometrics, The University of Amsterdam, stijn.claessens@uva.nl.

[2] Director, The RisKontrol Group GmbH, kreuser@riskontroller.com. This paper draws on experiences with developing foreign exchange reserves and other risk management tools for central banks and governments of several countries. The basis for the framework was developed as part of a research project undertaken at the World Bank, and is discussed in Claessens et al. (1998). The rapid development of models within this framework is discussed in Kreuser (2002). The technical details and equations left out of this paper are described and discussed in Claessens and Kreuser (2004). The RisKontrol Group GmbH in Bern, Switzerland is continuing to develop this framework, and additional papers and information may be found at http://RisKontroller.com. We would like to thank the referees for useful comments.

takes the achievement of certain macro objectives as given, and instead mainly focuses on assessing more narrowly-defined risk-return tradeoffs. Conversely, the macro literature has often focused on the link between foreign exchange reserves management and exchange rate determination, but has given little consideration to the investment choices made in terms of currency and maturity, for example.

Recent experiences suggest that these assumptions of separation cannot be taken for granted, as there are many relationships and feedbacks between various macro and microeconomic aspects and between foreign exchange reserves and liability management. The literature on financial crises, for example, has drawn attention to the need to consider the overall balance sheet of the government, and indeed of the country as a whole, when conducting foreign exchange reserves and debt management. The literature has also highlighted how private debt management and financial sector vulnerabilities can affect public debt and reserves management. The starting point to integrate the various, sometimes disjointed approaches has to be an analysis of what a general framework should look like. After this, one can then assess what is computationally and organisationally feasible. We will show in this paper that recent technological and analytical advances allow for very general approaches that at the same time can lead to outputs that are easily understandable and tools that can easily be implemented as well as adjusted.

Our approach is based on a general, dynamic stochastic optimisation model with a tree-based uncertainty structure. The numerical approach allows us to move beyond the classical methods of risk and return, efficient frontiers, and utility functions by considering much more general objective functions, behavioural relationships, equalities, and inequalities. The approach is also innovative in that it allows the user to view the full density functions of any outcomes that depend upon the optimal decisions taken. The user can then reshape these density functions as much as possible to any desired profiles and thereby obtain the resultant decisions to be taken today to achieve these profiles. As we will show, this approach allows senior management to define relatively broad objectives and constraints, and have these translated into an analytically rigorous approach, without giving up basic intuition and understanding.

The outline of the paper is as follows. We first review the typical approaches used for foreign exchange reserves risk management. We then present our general approach to asset/liability management (ALM), including reserves management, using a dynamic stochastic optimisation model. This is followed by several applications of our approach to the problem of foreign exchange reserves management by a central bank, along with some specific examples. We conclude with suggestions for institutional responsibilities for risk management that are consistent with our approach. Throughout the paper we use the acronym "ALM" in the most general sense, that is, to refer to the analysis and management of risks related to any combination of assets, liabilities, off-balance sheet items, and any cash flows that could affect the market value of the portfolio or the objective function, including contingent liabilities and catastrophic events. We do not restrict ALM to mean asset-liability matching; rather, our analytical framework allows a much richer universe of options, constraints and objectives in the analysis.

2 Typical approaches to foreign exchange reserves risk management

Approaches to reserves management vary along a wide spectrum. At opposite ends of the spectrum lie the traditional, pure macroeconomic-only oriented approach and the pure micro-based, risk-return-only oriented approach. We discuss the macroeconomic approach first.

Traditional reserves management objectives have mostly been formulated with respect to monetary policy and exchange rate management. These macroeconomic concerns are most relevant in the context of fixed or managed exchange rate regimes with some degree of capital mobility. In such cases, foreign exchange reserves need to act as a buffer against capital outflows in excess of the trade balance (and reserves will increase as capital flows exceed the trade deficit and vice versa). Reserves management is then secondary to the macroeconomic objectives, and mainly involves liquidity management, which assures the availability of sufficient free reserves at any given moment in time to intervene in the foreign exchange market.

More generally, under fixed as well as floating exchange rate regimes, the macroeconomic literature views holding reserves as a way of smoothing short-run shocks in external transactions, such as variations in imports due to terms of trade shocks, or variations in the capital account due to financial shocks. Rules have been developed for these circumstances, along the lines of optimal cash inventory management for corporations. The simplest practical application of this approach has taken the form of targeting a minimum ratio of foreign exchange reserves to imports, e.g. holding foreign exchange reserves that are at least equivalent to 12 months of imports. These and other, more sophisticated rules for a country's demand for international reserves have mostly considered only real variables, such as imports, exports, and the severity of possible terms of trade shocks, as well as some monetary policy considerations (e.g. Frenkel, 1980 and Frenkel and Jovanovic, 1981; see Bahmani-Oskooee and Brown, 2002, for a recent review). Most of this research, however, does not consider financial or balance sheet variables.

In the last decade, the relationship between (external) sovereign debt and reserves management has attracted considerable interest. Besides the total level of (external) sovereign debt, the importance of the maturity structure of debt relative to available reserves has been much highlighted. Governments typically find it relatively inexpensive to borrow on a short-term basis, since spreads usually contain a term premium, especially so for governments in emerging markets. However, a high proportion of short-term debt tends to increase the probability of self-fulfilling crises, as investors might suddenly decide not to roll over maturing debt or increase required yields on new debt. With foreign exchange reserves low relative to debt payments falling due, the risk of a currency or financial crisis can increase sharply. The structure of public debt, especially its maturity and currency composition, has indeed had key implications in many recent crises or near-crises. In those cases, a significant share of public debt was issued in foreign currency and/or linked to foreign currency, and contributed to rapid increases in debt ratios when exchange rates depreciated, inviting further pressure on reserves and the exchange rate because it generated doubts about the sustainability of public or private sector finances.

The literature that studies these relationships between government debt levels, maturity structure and reserves management and financial crises is rapidly growing and includes Calvo and Mendoza (1996), Sachs, Tornell and Velasco (1996), Rodrik and Velasco (1999), and Jeanne (2000). The literature has led to various rules on what are acceptable levels of debt payments falling due relative to the level of foreign exchange reserves (named after their principal or first advocates, the Calvo, Guidotti, and Greenspan rules). One rule of thumb has been that a country's governmental external debt repayments falling due in the next 12 months should not exceed its foreign exchange reserves. In some crisis countries, private asset and liability management has also contributed to pressures on reserves by not providing sufficient buffers against shocks. Consequently, some rules have included private as well as public debt payments falling due relative to foreign exchange reserves. In general, the structure of balance sheets, whether of the sovereign, the banking system or the corporate sector, have received more attention among reserves managers lately as the key role of reserves as a buffer in relation to both public and private debt has become clearer.

Less macroeconomic or balance sheet-oriented approaches and more microeconomic-oriented approaches to reserves management are feasible when monetary policy, exchange rate and debt management issues are of less concern, and when vulnerabilities in the financial and corporate sectors are small. This may be the case when the government pursues a flexible exchange rate policy, when it has a credible fiscal policy and institutional framework, and when domestic financial markets are well developed. A more active approach to reserves management may then put greater emphasis on profit objectives, although still within certain bounds. The implementation of this objective may involve the division of the reserves portfolio into active and passive parts. The passive portfolio would be used to deal with macroeconomic objectives, and would be mainly managed with liquidity objectives in mind. The active portfolio could be used for the purpose of profit, possibly taking into account liability management objectives. Its management would consider the breadth of tactical risk management tools available, including many types of borrowings and assets, forwards, swaps, plain vanilla and exotic options, etc., and would use concepts such as Value-at-Risk.

The division into two portfolios may, explicitly or implicitly, involve the separation of objectives, as in the case of separate management, where it may be assumed that there will be little or no spillover between the two parts, at least over the horizons set for the managers. It will typically also involve different institutional arrangements for the management of the two portfolios, as well as different investment tools or instruments to be used. In terms of ALM, the management of the passive portfolio could be similar to the general, more macro-oriented liquidity management raised above. The ALM of the active portfolio would appear to come very close to the ALM issues facing a commercial bank or a corporation. While correct in many ways, considerable differences still remain, however, such as the difficulty of choosing which measure to judge the performance of the portfolio (e.g. in which currency should the profits and risks of the central bank be measured: the local currency or one of the intervention currencies, such as the US dollar or euro?). Importantly, the (implicit) assumption of separation would need to be regularly reviewed.

In practice, these macro and micro-based objectives and accompanying tools for reserves management will often be mixed or will take a hierarchical form. Typically, approaches used for central bank reserves management involve combinations of many objectives, with four main concerns often mentioned, usually in this order: security, liquidity, profit and stability. Correspondingly, analytical ALM models may try to define one or more of these objectives. For example, the central bank may decide that its ALM objective should be to try to achieve those asset and currency allocations for which the probability of negative unrealised losses is very small. Alternatively, the objective may be to have sufficient liquidity available for intervention in 99% of cases. The objective could also be to ensure that the rates of return are between 3% and 4%, while at the same time the ratio of short-term debt to reserves should not exceed 1.05 for 99% of the time. Finally, the objectives might be stated as a combination of these or other goals.

In the next two sections we will describe our analytical framework and its application to specific reserves management problems, providing strategic allocations of assets, currencies and liabilities to meet these kinds of mixed macro and micro-based objectives. The analytical framework is essentially a dynamic stochastic optimisation model.[3] Whereas these kinds of

---

[3] Birge and Louveaux (1997) discuss dynamic stochastic optimisation, or stochastic programming, in general and include some financial examples.

models may have in the past taken a long time to build and deploy, this is no longer the case thanks to the confluence of several new developments in mathematics, computers, algorithms, modelling languages, and tools for insight and intuition. Many of these new developments have been recently combined into an applied framework. An application of this framework to a country case, the Republic of Colombia, is provided in Claessens et al. (2000), while the development of the conceptual framework, which took place at the World Bank, is provided in Claessens et al. (1998). These new developments are discussed further in Kreuser (2002).

In Section 5 we will discuss how this framework can be implemented into the institutional framework of a central bank, how these often-used objectives can be translated into a dynamic approach, and how the issues of consistency of objectives can be addressed.

3 A general analytical framework for ALM

The review of the approaches taken by central banks already shows that it is unlikely that a single reserves management model will be applicable in all situations. No single model can be expected to be optimal for every central bank's reserves management requirements. Every model developed today will need to change tomorrow. Flexibility must therefore be built into any approach to reserves management modelling. The purpose therefore should be to provide a framework that allows for substantial flexibility in model development in both theory and application. A framework has to be dynamic as well. Any medium to long-term analysis of reserves must include procedures for the dynamic rebalancing of the reserves portfolio for two reasons: the density functions of outcomes in the distant future depend upon decisions taken beforehand, that is in the near future; and changing regime conditions mean future decisions need to be made dependent on future outcomes. We will discuss these desirable features in turn.

(a) General ALM approaches at the country level

The starting point of our discussion is that the development of a strategic approach for ALM at the country level has lagged behind approaches in the corporate and financial sectors. Typical approaches to country ALM are copied from approaches for firms and financial institutions, but do not incorporate countryspecific factors, while strategic aspects are usually missing. They often exclude, for example, trade flows and fiscal dimensions. Modelling flexibility is very limited, with country adaptation often happening through a piecemeal approach using ad hoc analysis rather than optimisation and starting from first principles. Their perspective is often also the development of benchmarks, rather than strategic asset allocations or liability choices. By requiring a benchmark, which is constant over time, they fail to incorporate the dynamic realignment of portfolios. The treatment of uncertainty is typically also very limited and constraints are often not included in the optimisation process itself, but rather through iterating around the solution.

The need for strategic ALM for sovereigns is all the more necessary given the fact that sovereigns often have to consider risks on a much broader scale than corporations or financial institutions do. Risks for a sovereign concern not only the government's own direct financial exposures, such as those arising from debt and reserves management, but also those arising from contingent liabilities owing to risks in the banking system, the restructuring of state-owned enterprises, or restructuring and reform of the corporate sector. Approaches also need to relate to measures of the government's earning potential. This may mean risks need to be defined differently. Instead of measuring nominal variability in government debt service, for example, risk measures may need to take into account the sensitivity of fiscal revenues to global factors, such as interest rates. Without these factors, approaches to risk can ignore the existence of natural hedges in the external and fiscal sectors, limit the analysis to "on-balance" liabilities only, and ignore many important constraints. It is also essential that the sovereign adopts a truly dynamic approach. Countries often face, for example, many constraints in rapidly adjusting their assets and liabilities: transactions costs can be high, and market access may vary with general financial market conditions and investors' sentiment. This is especially true for developing countries, but also applies to developed countries. Typical ALM strategies pursued for financial institutions and corporations can thus clearly be less than optimal for sovereigns or central banks; moreover, measuring exposures only in terms of duration, asset composition and currency composition may even add to risk.

(b) The stochastic structure

In any ALM model, the variables that define the future states of the world belong to two classes. The first are the exogenous variables that include interest rates, exchange rates, current account deficits, and others that one wishes to stochastically estimate. The second are the "decision variables", namely those determined by the model. These could include the amount of reserves and debt, the maturity structure of foreign exchange reserves and debt, the percent of each reserve currency, etc. The initial step in any analytical framework with regard to the exogenous variables is to generate a collection of scenarios of possible realisations. These chosen exogenous variables could include interest rates, exchange rates, bond prices, GDP, the trade deficit, fiscal deficit, commodity prices, as well as the rare events triggering contingent liabilities (e.g. a banking system crisis), and other factors.[4]

One option to describe the future values of the exogenous variables is a scenario structure, where the evolution of the future is in the form of a horse tail (i.e. at time zero (today), a draw is made of a particular path of future events and all future events to follow are "chosen" at date zero). Scenarios then consist of multiple draws of such paths, all originating at date zero. For ALM purposes, however, instead of a scenario structure, a tree structure that considers the conditional probabilities of events at each point in time is more applicable. A simple example is depicted in Figure 1.

Each box in the tree of Figure 1, after the first, represents possible outcomes of the exogenous variables and the probability of the event (also called a node) or the values in that box occurring. Here the exogenous variables are taken to be an interest rate and an exchange rate. At each future point or at the end of each period, here taken to be six months and one year, we can measure the distribution of any of the exogenous variables as we have their probabilities. We call the movement from one event to the next a branch (this is also sometimes called an arc). We term a set of events from the beginning (today) to the model horizon (here one year) a path or one particular scenario of the tree.

[4] Not all of these variables will be fully exogenous to the reserves and debt management decisions taken. In principle one could make some of them endogenous, for example, by allowing reserves management to affect the behaviour of the local exchange rate. There will be choices involved as to how broad one wants the modelling and the objectives chosen.

Figure 1: A simple tree structure for interest and exchange rates

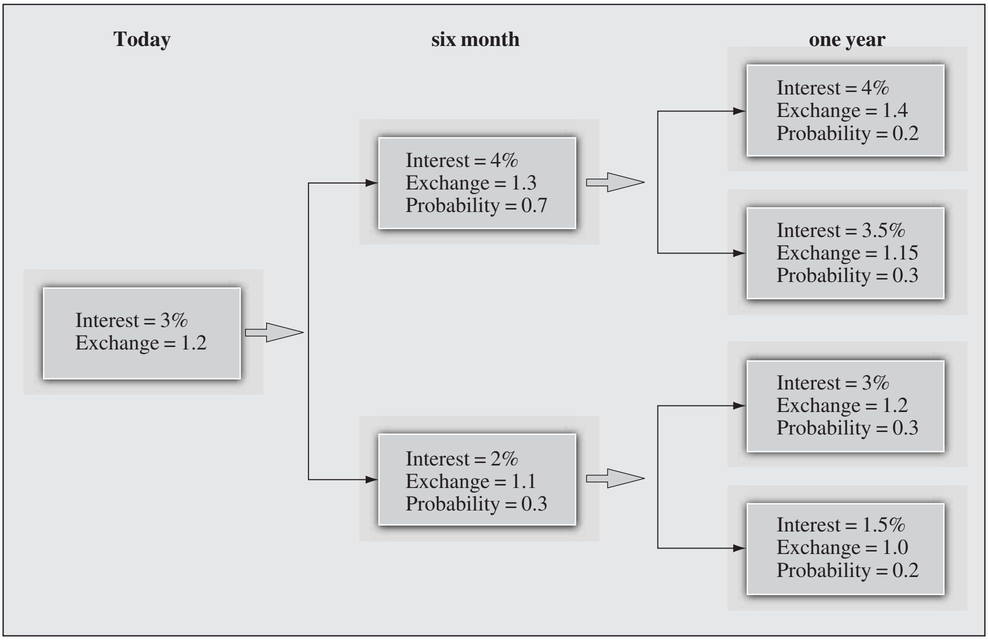

There are several reasons why a tree structure is better than a scenario structure. First, the tree branching structure provides a good discrete approximation of the stochastics of most processes, allows many interdependent factors to be generated, can be adjusted easily to satisfy the stochastic properties of processes, and is a natural representation of the way uncertainty unfolds. Second, the tree branching structure can be solved by moment matching using branches with variable probability. This means that low-probability but high-impact branches (in terms of affecting decisions made) can be added easily while keeping the tree small. Third, the exogenous variable values in the events can be constrained to satisfy many requirements. Theoretical equilibrium conditions, such as interest rate parity conditions, can be added to the moment-matching problem as constraints. Fourth, this structure is necessary for optimisation models for the reasonable calculation of future distributions and densities of endogenous variables. We discuss this in more detail in the next section.

In order to generate trees, we need a representation of the behaviour of the exogenous variables. One very useful representation is to define the movement of the exogenous variables as multi-factor stochastic partial differential equations. This representation underlies commonly-used multi-factor models, including those with mean-reversion, such as Hull-White and Heath, Jarrow and Morton. The parameters of these equations can be determined from historical data (see for example Göing, 1996). The stochastic partial differential equations can be formulated so as to transform from transient processes in the short term to stationary processes in the long term. The correlations can also be made to vary by time and level, allowing the generation of a very rich possibility of outcomes. The parameters can furthermore be adjusted so that the stochastic partial differential equations are consistent with current market information and implied values, such as those obtained from values of derivatives. They can also be adjusted to satisfy theoretical concepts, such as uncovered interest parity that determines how exchange rates move as a function of interest rates. Lastly, the parameters can be calculated so that the equations match expert opinion on expected values of some processes, volatility or correlations over time. The latter is particularly helpful for stress tests.

Then the tree is built up event-by-event by solving the moment-matching problem with variable probabilities using the stochastic partial differential equations at each event. The probabilities and the respective values of the exogenous variables are determined so as to match the first, second, and possibly higher moments. A tree structure with variable probabilities is very important because it means that scenarios of low probability and high impact can be added without requiring thousands of new scenarios to be generated. In addition, by building the tree event-by-event, one can trigger events to occur at certain times or at certain levels of variables, such as low levels of GDP growth triggering contingent liabilities in state-owned enterprises, or guarantees on deposit insurance falling due. This allows us to capture tail events and yet keep the tree relatively small. It allows for easy generation by stochastic processes that can also vary by time and level. Using this variable probability process, trees of 40 to 50 correlated factors can be generated over several periods.

After the tree has been generated, it will exactly satisfy all the stochastic properties (mean and covariance and possibly higher moments) and theories as specified for the processes. Furthermore, since the tree may remain relatively small, it is possible to build and solve the dynamic stochastic programming model on a PC using technology as discussed in Kreuser (2002). Figure 2 provides an example of an actual tree for one of the many factors defined in a model used as part of the study for the Republic of Colombia.

Figure 2: Sample tree[5] defined over 3 periods

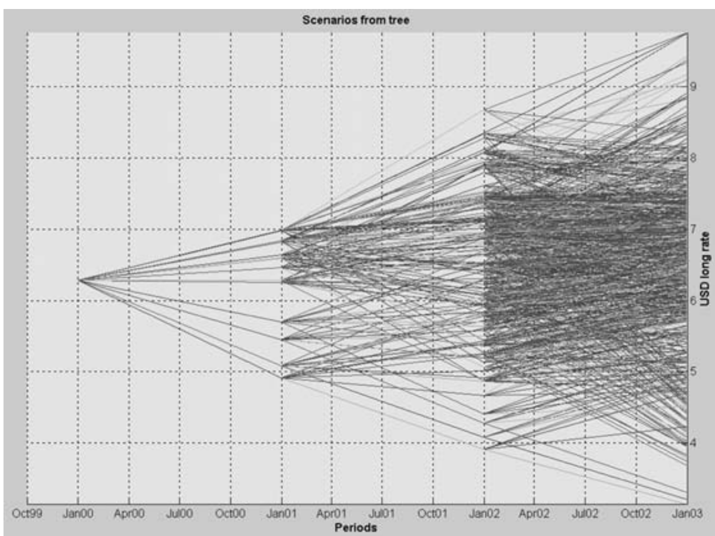

[5] The branches on this tree are not of equal probability. In order to visualise where the probability mass lies, a second graph is generated head-on into the branches with the size of the branch or dot scaled depending on its probability mass. Alternatively, the density function or distribution is generated at each time period. For more details, see Claessens, Kreuser and Wets, 2000.

(c) The structure of the dynamic stochastic programming models

The next step is to build a model to derive the decision variables. The model is defined independently of the stochastic processes, the tree, or the events on the tree. Put differently, the choice of the stochastic processes ought not to drive the choice of the model. Typically, however, in ALM models this separation will not be valid, as for example in a mean-variance model, which is only valid under certain assumptions about utility preferences. The more general approach to follow will be a dynamic stochastic optimisation model with objective, equalities and inequalities (examples of these kinds of models can be found in Ziemba and Mulvey, 1998). Allowing for more general model formulation combined with tree generation for the stochastic processes makes the system very powerful. It becomes trivial to make changes. One can adjust constraints or objectives without affecting the tree. Or one can modify the tree (make it larger or smaller or introduce specific stress tests) without changing the model. Decision variables are defined for each event with respect to the levels of assets, liabilities, currencies, alternative investments (derivatives) and other cash flows. Decision variables are all handled separately, allowing the specification of transaction costs, spreads and limits to be individually imposed, and thus increasing the stability and realism of the model. The selection of what decision variables need to be defined depends on the specific issues to be analysed and the objectives to be pursued. Since the model allows analysis at the strategic level, assets in the portfolio are usually defined in terms of classes, such as short, medium and long-duration assets, or buckets of different currency classes, rather than at the level of individual assets. We might for example define as an asset class those short duration assets of no more than six months and make the percentage of the portfolio holdings of those assets the decision variable.

As mentioned above, the tree structure provides a view of future distributions and densities that differs from a scenario structure. If we had a scenario structure, then the model "would know" at each event at the end of the first period what the interest rates etc. at the end of the second period would be as these would be defined by the particular scenario chosen at time zero. Similarly, at the end of the second, it would know events in the third period, and so on. Put differently, once a particular scenario is chosen in the first period, the future states are known with certainty and the model can choose under perfect foresight the "best" asset class at any point. This does not correspond to reality, however, and one should not allow the model to know the future in advance; rather, it must be non-anticipative. Models that use scenarios often get around this problem by requiring a rule such as always rebalancing to a benchmark. However, again this is unrealistic, as in some events a new strategy would, or at least should, probably be implemented. The tree structure allows complete flexibility to choose strategies depending on the time and level of any exogenous and endogenous variables.

The model constraints that might be imposed can be numerous and varied as well. Some possibilities include legal limits on asset classes, portfolio rollover constraints, transaction cost limits, cash-flow requirements, currency transfer constraints, market access constraints, liquidity constraints, and other policy constraints. These conditions constrain the set of feasible strategies. Then there is a whole collection of constraints to shape the distributions of factors that are deemed important. We can call these "objectives", although the distinction between constraints and objectives is not always clear. A simple example of an objective and the constraints that typically form the basis for many asset portfolio allocation problems is to maximise the expected return at the horizon with constraints governing the allowable purchase and sale of assets, including the formulation of transaction costs. One can add to this stochastic cash outflow constraints and limits on negative returns on investment portfolios.

An example for a central bank's asset and currency allocation problem would be to state as objectives the need to achieve a high return, a low probability of unrealised negative returns and/or a low probability of portfolio sales, while at the same time maintaining a liquid portfolio to best meet intervention requirements. An example of objectives for the debt management of a Ministry of Finance would be to minimise the expected ratio of external debt to GDP, to limit the upside volatility, to hedge against market liquidity, and/or to best meet new financing needs.[6]

Once the objectives and constraints have been defined, the model is solved simultaneously for all decision variables at once (for the computational approaches to solve this type of models, see Claessens et al., 1998). Solutions are not obtained sequentially event-by-event, because future decisions, transaction, costs and constraints will impact the decisions to be taken today. For example, constraints on liquidity for events that might occur in the future will affect the optimal portfolio composition today. What exactly to include as decision variables, constraints and objectives will depend on each central bank's or Ministry of Finance's policies, constraints, operating environments, and the issues it wishes to address. These can take the form of the definition of a particular objective function. Typically, however, experiences have shown that the precise definition of the objective is of less importance for the final solution. Rather, as we shall also show, in this framework the shaping of the distribution and density functions of various factors will be the most important analytical tool for the manager.

(d) Shaping of distributions and density functions of outcomes

We now discuss the shaping of distribution and density functions. Once the model has been solved (say under the objective of maximising the expected return at the end of the horizon), the density and distribution functions for all the factors under consideration can be estimated. Typically, one will have preferences regarding the shapes of the density and distribution functions. These preferences can be defined formally as objective functions, with or without constraints, as noted above. Another way to state the role of preferences is that they in some sense "shape" the preferred density and distribution functions of the outcomes. The advantages of more explicit shaping is that, unlike what is typically done in objective functions, one does not need to specify or estimate certain parameters, such as utility preferences. Rather, one allows the policy-makers to review the actual density and distribution functions obtained and then state their preferences as criteria related to concepts such as "a less fat left-tail" or "less probability mass in this or that region".

The shaping of a distribution function implies the shaping of its corresponding density function, and vice versa. As such, it is immaterial what is being shaped, and in our framework, we can do either. In some contexts, however, it may be more appropriate to do one rather than the other. Particularly, visualising density functions will often provide information more readily than distributions. Our techniques for shaping operate on densities or distributions of

[6] Several different indicators or measures exist that may be used and applied in our framework. Some references where these are discussed include Bank for International Settlements (BIS, 2002), Asian Development Bank (ADB, 2001), World Bank and International Monetary Fund (2001), and International Organization of Securities Commissions (IOSCO) (2002) for both debt and reserves management.

any shape, including non-normal.[7] The actual techniques applied to change the shape of densities to fit the objectives of the policy-makers can vary. We may shape the left or the right of the density, and in fact we can specify several points on the densities or distributions that we would like to shape. Furthermore, we may shape several densities or distributions at the same time and several referring to different periods of time.

Mathematically, we can characterise the preferences for the shape of the density as a function $P(s,I;X)$ with parameter s a slope, I an interval, and X the random variable of interest. The function P is like a utility function in that it has a positive first derivative with respect to X and a non-positive second derivative with respect to X. Therefore in maximising P, we will be maximising X (the function can also be defined for a minimisation problem, in which case we will be minimising X). The value of s can be interpreted as the intensity of shaping, and can be rather arbitrary.[8] I is the interval in which one would like to push as much of the probability mass of the random variable X as possible, subject to the trade-off arising from the overall maximisation. There are many possible mathematical forms of the function P (the details of some examples discussed in this paper can be found in Claessens and Kreuser, 2004). The important aspect is that the parameters s and I are defined explicitly and are not estimated in any way. The main purpose of the "optimising" function P is to determine the intensity by which one would like to push the probability mass into the interval I.

One example of an objective for a debt manager is the minimisation of a preference function defined over the expected end-of-horizon ratio of total external debt to GDP. The preference function may, for example, be stated so as to push the probability mass of debt/GDP as much as possible below 60%, while being rather indifferent to pushing it below 10%. The 60% value is perhaps determined by policy, and is thus relatively hard, whereas the 10% may be a softer, preference figure. Viewing the outcome of the density function, however, the policy-makers may decide to make the 10% target a harder target, as the cost of doing so may be relatively low. The point is that the policy-makers can observe the outcome and then decide in which way to shape the density function.

Figure 3 provides an example of the density functions of the reserves portfolio of a central bank at a point in the future. The density on the left hand is a measure of the current portfolio value, measured in tens of billion currency basket units, in this case made up primarily of US dollars, euro, and yen. It was obtained by maximising the expected value of the portfolio at the horizon measured in the currency basket. The density on the right hand was obtained after visualising the density on the left hand, and then deciding to squish its shape by applying a preference function to push as much of the probability mass of the portfolio wealth into a range between 14 to 15 billion basket currency units. To achieve the reshaping, the optimal currency and duration of the portfolio was redistributed. The process thus becomes very intuitive: define an interval where much of the probability mass is preferred to lie, visualise and evaluate the result, and finally obtain the outcome in terms of the new portfolio composition.

[7] Another step in the process is moving from discrete distributions to continuous densities. This is not a trivial one, but new techniques make the process efficient and accurate. The estimation of continuous densities using a modest amount of discrete data is an important part of our framework, and is discussed in more detail in Kreuser (2002).

[8] The parameter s is a slope factor. A small value of s pushes the probability mass into the interval I less than a large value. The actual value is not so important; usually s=100 is a good starting point. After visualizing the density, if one wants to try to force more (less) probability mass into the interval I, then the parameter s may be increased (decreased) to see if it has any further effect. This process is discussed and explained in more detail in Kreuser (2002).

Figure 3: Density functions before and after shaping

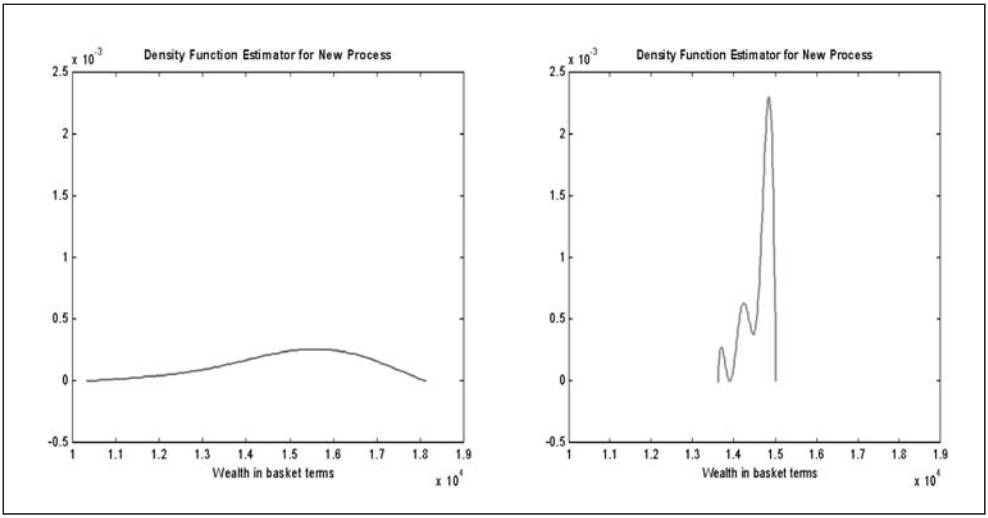

Another important measure for shaping distributions or densities is conditional Value-at-Risk (CVaR). CVaR is equivalent to the expected shortfall of a target at a specified level probability or confidence level $\alpha\%$. Mathematically, it is the concept of the average loss in the worst $\alpha\%$ cases.[9] CVaR is related to Value-at-Risk (VaR) and, for loss distributions, CVaR is always greater than or equal to VaR. Therefore, if we put a constraint on CVaR (an upper bound), the VaR value is also constrained to be below that value. CVaR may be thought of as a more robust measure than VaR, as it provides the average loss for those a% of cases rather than the minimum loss, and therefore takes extreme losses more fully into account.

The CVaR has some useful properties for modelling purposes. Specifically, CVaR constraints can be modelled as linear constraints, whereas VaR constraints are non-convex and thus more complex to model and solve. This makes CVaR constraints the method of choice for shaping the distributions or densities for dynamic stochastic optimisation models, especially when they are also very large. In addition, since CVaR always bounds the VaR measure, it is possible to obtain distributions that are VaR-constrained by sequentially relaxing the CVaR constraint. Further mathematical descriptions of CVaR constraints for the reserves management problems discussed in this paper can be found in Claessens and Kreuser (2004); a more detailed discussion of CVaR constraints in general can be found in Rockafellar and Uryasev (2001).

The CVaR concept allows us to translate intuitive descriptions of constraints on distributions into mathematical descriptions for use in our modelling framework. One example of a CVaR constraint can be that in the 1% of worst cases, average losses should not exceed 10% of the portfolio value. Another one could be that in the 5% of worst cases, the average shortfall in liquidity should not exceed 20%. And, as we can incorporate VaR-type constraints into this framework as mentioned by relaxing the CVaR constraints, a VaR-type constraint could be that with a 95% confidence level, the external debt to GDP ratio is less than 0.6. All of these may be incorporated into our analytical framework.

---

9 This is explained very precisely mathematically in Acerbi and Tasche (2002). This paper also discuss the relationship between CVaR, expected shortfall, worse conditional expectation, tail conditional expectation, and VaR.

Figure 4: Constraints on the distribution

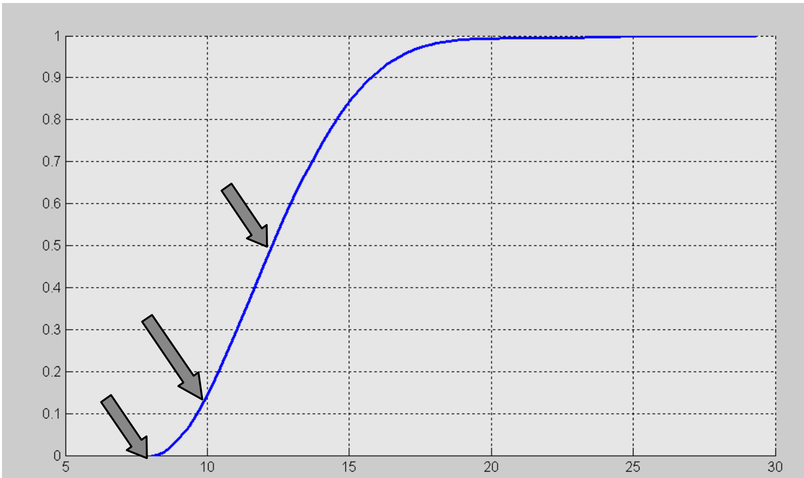

Reserve Portfolio US$ Billions at End of Period

We may apply several of these constraints at one time to one single distribution. We may, for example, constrain a distribution of returns to certain loss levels at the 50% confidence level, the 85% confidence level, and the 99.9% confidence level. Such a situation is illustrated in Figure 4, in which we depict these constraints as arrows on the distribution preventing it from moving too far to the left and thus effectively shaping it.

We can also constrain several distributions at the same time. In doing so, however, we may impose excessively tight constraints, which could mean there is no longer a feasible solution. This is why a two-stage process may be necessary that first of all attempts to find a feasible solution to all the constraints that have been imposed (or at least a solution that is as close as possible given the imposed constraints). If the solution is only close, then the close solution can be accepted or the constraints can be relaxed. Having found a feasible solution, the procedure can then continue to find an optimal solution.

Another way of shaping densities is through general and hard constraints. For example, one may wish to have no unrealised loss of principal over a specified period. This constraint can be applied directly and the result will be to cut off the density function of returns completely at zero. Its effect will be to reduce the riskiness of the investment portfolio. Another example is to constrain the expected value of returns and the downside returns to some specific levels. The disadvantage of these general and hard constraints is that they can be quite costly. To ensure, for example, no chance of loss whatsoever may imply a very conservative portfolio composition with a very low expected rate of return. In practice, there is likely to be more flexibility, with management willing to take an acceptable level of risk of a minor loss of, say, 1%. For this reason, the technology of shaping densities is closer to the problem that managers face.

Lastly, to shape the density functions, one can expand the set of investments and/or allow for the use of derivatives. One can use a derivative, for example, to flatten the density function, that is, to reduce “risk”. However, since derivatives may have an impact on the portfolio as a whole – a derivative may for example reduce the expected return or other expected values – one will need to evaluate the use of derivatives in the overall context. One can do so easily under this framework by adding the derivative to the model as another instrument or decision variable. Derivatives are particularly easy to add as their pricing and volatility are determined by those of the underlying process, which will already be included in the modelling.

4 Specific reserves management problems

Several of the techniques discussed can and will probably be applied simultaneously in order to obtain a solution that satisfies the various criteria of safety, liquidity, returns and stability as they are defined, explicitly or implicitly, by a particular central bank or government. In this section, some specific examples of techniques as applied to commonly-stated foreign exchange reserves management problems are discussed. We start with a macroeconomic approach that also considers balance sheets. We then consider improving the rates of return on foreign exchange reserves while satisfying some constraints. We next review a policy of portioning the foreign exchange reserves into active and passive portfolios.

(a) A macroeconomic approach

As an example of how our framework can be applied in the case of foreign exchange reserves risk management when macroeconomic issues are most important, we consider the objective of managing the ratio of debt payments falling due relative to foreign exchange reserves. In this case, we assume it is desirable to manage short-term external debt (STD) relative to reserves so that debt payments falling due within the next 12 months do not exceed foreign exchange reserves. We explore two specific possibilities of constraints. In the first, we assume that the central bank manages its reserves level and needs to obtain benchmarks for reserves currency and asset class composition, but has no influence over variables regarding public debt. In this case we assume STD is exogenous from the point of reserves management, although it may be dependent on other variables such as trade and GDP. Given this problem, we formulate the preference function P following the previous convention as

$$P\left(100, [1,1.05]; \frac{reserves}{STD}\right).$$

The objective of managing the ratio of reserves to shortterm debt can be defined as applying only at some point in the future (the horizon), or as a weighted sum of the ratio at all intermediate points in time. The latter case would allow for some smoothing of the ratio over time, if for example there were an unexpected temporary shortfall.

As *STD* is assumed to be exogenous, but still stochastic, maximising the function *P* has the effect of pushing much of the probability mass of the ratio of reserves to STD into the interval range between 1 to 1.05, while maximising the value of reserves. In other words, the function has the effect of trying to manage reserves to best meet the targeted ratio, but does allow for uncertainty. This could result in some volatility in returns, and one may want to add some constraints restricting this volatility. This could take the form of a CVaR constraint formulated as follows: in 1% of the worst cases, the average shortfall in returns below 3% should not exceed 2%.[10]

We can add other constraints as well, such as maintaining a minimum amount of foreign exchange reserves to meet expected liquidity requirements necessary for intervention. We may want to estimate these liquidity requirements from historical data and generate future stochastic requirements exogenously and/or dependent on factors such as GDP and terms of trade shocks. We would then set aside a portion of our portfolio for liquidity so that, for example, in 2% of the worst cases, the average shortfall in the level of the liquidity portfolio should not exceed 10%. The model will then determine the level of this liquidity portfolio, while also meeting the other constraints and objectives.

These examples of objectives and constraints are summarised in Figure 5. The full model consists of many other policy and legal constraints as required for each central bank, along with the necessary budget constraints defining cash flows at every event.

Figure 5: Summary of main model equations for a macroeconomic approach

Macro issue: Managing reserves relative to short-term debt			
Form	**Type**	**Criteria**	**Description**
Objective Maximise	Preference on $\frac{reserves}{STD}$	Stability of ratio and returns	Shape density of outcomes so that ratio most likely falls into the range [1,1.05] while maximising reserves.
Risk Constraint 1	Inequality on returns	Safety of principal	Limit the unrealised losses so that the portfolio return is always $\geq$ 0.
Risk Constraint 2	CVaR constraint on returns	Safety on extreme losses	In 1% of the worst cases, the average shortfall in returns below 3% should not exceed 2%.
Risk Constraint 3	CVaR constraint on liquidity	Liquidity	In 2% of the worst cases, the average shortfall in the liquidity portfolio should not exceed 10%.

[10] We formulate this CVaR constraint as a loss constraint. The result of imposing this constraint is that the VaR value at the 99% confidence level will be less than or equal to 2%. Or, in 99% or more of cases, returns will exceed 1%. This latter will hold but is not equivalent to the CVaR statement, which is actually stronger.

In the second case, we introduce decision variables for both debt and reserves management, and assume that the Ministry of Finance and the central bank coordinate their actions to some degree. We now disaggregate both reserves and debt into appropriate classes and currencies and allow both sets of variables as decision variables. We cannot apply the previous preference function in this case as this would lead the model to set STD to zero as debt is now also a decision variable. Instead, we apply an alternative minimising preference function,

$$P\left(s, [\alpha, \beta]; \frac{TB - TA}{Y}\right),$$

where *TB* is the total debt, *TA* is the total amount of foreign exchange reserves, and *Y* is a scaling variable, in this case assumed to be exports (the external debt-to-exports ratio is often used as an indicator of the repayment capacity of a country). The interval [α, β] would have to be defined appropriate to each individual country. Debt accumulation and reserves levels will be related through the budget constraints of the government and the central bank. Net new debt accumulation will now be determined by the model in such a way as to manage both the country's liquidity needs as well as the debt-to-export ratio. The model will also determine the currency and maturity structure of debt and the structure and currency of foreign exchange reserves.

In this case, minimising the preference function has the effect of seeking low borrowing costs and a high return on reserves. However, we did not introduce the volatility of returns or costs directly into the preference function. We can do so easily by introducing CVaR constraints for both returns and costs, that limit the downside risks for returns and the upside risks for costs. Alternatively, instead of introducing CVaR constraints, we can specify explicit constraints on costs and returns. We can also introduce an explicit constraint on STD, for example that STD is always less than foreign exchange reserves. Alternatively, this can be formulated as a CVaR constraint, such as "in the worst 5% of cases, the average shortfall in reserves below STD should not exceed 300 million". The advantage of the latter constraint is that in extreme circumstances, reserves may be allowed to fall below short-term debt. These constraints are summarised in Figure 6.

Figure 6: Summary of main model equations for a macroeconomic approach II

Macro issue: Managing debt and foreign exchange reserves			
Form	**Type**	**Criteria**	**Description**
Objective Maximise	Preference on $\frac{TB - TA}{Y}$	Stability of ratio costs and returns	Shape density of outcomes so that ratio most likely falls into the range $[\alpha, \beta]$ while reducing costs and increasing returns.
Risk Constraint 1	CVaR constraint on liquidity	Liquidity	In 2% of the worst cases, the average shortfall in our liquidity portfolio should not exceed 10%.
Risk Constraint 2	CVaR constraint on STD	Stability of STD	In 5% of the worst cases, the average shortfall in reserves below STD should not exceed 300 million.
Risk Constraint 3	CVaR constraint on returns	Returns	In 5% of the worst cases, the average shortfall in returns below 5% should not exceed 2.5%.

Additional constraints to satisfy policy objectives or limits can also be introduced. We can, for example, introduce constraints on the volatility of transaction costs or on the percentage of the portfolio rolled over, either/both on reserves and/or on debt. Legal constraints on the type of investments may also be introduced. The results will be reserves and debt allocations satisfying the required constraints, e.g. on safety, liquidity, returns and stability. In all cases, the reserves and debt variables will again be disaggregated into maturity and currency classes, thus allowing investment benchmarks to be determined as outcomes of the optimisation procedure.

The specific preference function chosen depends on each particular situation and the relevant indicator. Alternative indicators appropriate to specific situations (as discussed in IMF, 2000) include the ratios of reserves to short-term external debt, reserves to imports, reserves to broad money, external debt to exports, external debt to GDP, and measures such as the average interest rate on external debt, the average maturity of debt or reserves, and the share of foreign currency external debt in total debt. We do not advocate one particular indicator over another – indeed, many may be used at the same time – but rather emphasise that this framework can be used to incorporate all these indicators.

(b) Improving returns on foreign exchange reserves

In this section we assume that the central bank has substantial foreign exchange reserves so that its macroeconomic objectives are met. Its main concern rather is assumed to be to improve returns through active management and using various (new) financial instruments, including derivatives. It might choose as its objective function the maximisation of the expected value of the terminal level of reserves, where the terminal period may be any period, but will typically be after one or two years. One of the main policy issues is likely to be what currency the value of foreign exchange reserves and the realised rates of return should be measured in. An approach that is often used is a currency basket with weights equal to the proportion of each currency in the country's imports. When the weights are constant, a basket or reference currency results. In turn the basket is then assumed to represent a risk-neutral position vis-à-vis currency risk (i.e. investing in currency in the same proportion as in the basket represents a risk-neutral investment position). The returns in this model could be volatile, however, so an explicit constraint on the downside returns might be introduced, with the same reference basket currency used to measure the constraint. A limit on transaction costs may also be imposed, so that trading does not become excessive. This may take the form of a constraint that in the worst 10% of cases, the average total transaction costs in any period should not exceed 10 million. The constraints and objective are described in Figure 7 with "Risk constraint 2" representing the constraint on downside returns. The objective is described as a preference function, with two parameters, α and β. With α large, the preference function becomes equivalent to maximising returns. And changing the size of the interval $[\alpha, \beta]$ can be used to shape the density function as per Figure 3.

With more active reserves management, the universe of possible asset variables is typically expanded, including assets of longer duration and lower credit quality. Not only price fluctuation, but also credit risk considerations will then arise. The latter can be handled in the model by simply including asset limits on those assets with a lower credit rating. They can also be handled by estimating the stochastic processes driving the actual default risks, and incorporating the default risks explicitly into the model. As the framework is flexible, the difficulty in implementing credit risk arises mainly from the difficulty in determining the stochastic processes for default risks, and not from the technology of incorporating them into the model.

Figure 7: Equations for improving returns

Macro issue: Maximise returns with liquidity and safety			
Form	**Type**	**Criteria**	**Description**
Objective Maximise	Preference on returns	Returns	Shape density of returns measured in the currency basket so that they most likely fall into the range [α, β].
Risk Constraint 1	CVaR constraint on liquidity	Liquidity	In 2% of the worst cases, the average shortfall in our liquidity portfolio should not exceed 10%.
Risk Constraint 2	CVaR constraint on returns	Safety	In 5% of the worst cases, the average shortfall in returns below 6% should not exceed 2.5%.
Risk Constraint 3	CVaR constraint on transactions	Stability	In 10% of the worst cases, the average transaction costs should not exceed 10 million.

Derivatives may similarly be introduced. These sometimes represent a cheaper way to manage interest rate and currency risks than trading the underlying assets (as discussed in Rigaudy, 2000). These trade-offs can also identified through the model, for example when the model chooses a derivative over the underlying assets. Derivatives may also help satisfy specific constraints, such as managing the risk of downside returns. Again, in our model, derivatives will only be chosen if they are cheaper than the underlying assets themselves or more effective in reducing certain risks. For strategic exchange rate intervention purposes and other reasons, one can also allow the use of derivatives, especially if liquidity constraints are a binding factor (see Blejer and Schumacher, 2000). Finally, derivatives can be used to meet risk constraints that require the density function to be flattened in relation to returns.

(c) Partitioning into active and passive portfolios

A central bank may decide to partition its portfolio into active and passive parts and manage each part separately. We use the terms "active" and "passive" to mean that the active portfolio is managed more aggressively than the passive part. The two parts may have different sets of rules and instruments, and the number of transactions may be higher on the active portfolio than on the passive part. In this case the active portfolio may be managed as discussed in 4(b) and the passive part as in 4(a). Instead of a fixed allocation between the two portfolios, another approach could be to develop a model that determines the optimal proportion of foreign exchange reserves to be allocated to the active portfolio. To analyse this situation, we assume that the central bank takes a macroeconomic approach as in 4(a), whereby the model would be structured with the addition of a constraint limiting the volume of asset sales to a level that would be reasonable for the institution to handle. The limit to the level of trades would be chosen to be consistent with the central bank's objectives and its ability to manage the institutional environment demands (such as people, back office and reporting and control functions). A second constraint that limits total transaction costs as in the previous section could also be added. The solution to the model would then give us the current investment choices and the distributions of all future outcomes as well as the decision variables that are contingent on those choices.

With the solution to this model, we will also obtain the density function of the level of trading at any time in the future. Instead of predefining the active and passive portfolios, we can define our active portfolio as consisting of those assets in the overall portfolio that are actively traded or sold and purchased. This then gives us a natural split between an active and a passive portfolio, with the size of the active portfolio and choice of assets that make up the active portfolio chosen determined so as to satisfy the overall objectives and other constraints. The result is an optimal portfolio whereby the active part has been determined consistent with the central bank's objectives and its ability to manage trading within an uncertain future environment, and is moreover consistent in terms of constraints and objectives across both the active and the passive parts of the portfolio.

5 Application of the framework to foreign exchange reserves management

In this section we discuss how the analytical framework can be integrated into a central bank's operations from an institutional perspective, covering all topics from policy analysis and definition to operational modelling and full implementation, including back office and control functions.

(a) An analysis of the central bank's operations

The first step is a thorough review of the central bank's current operations, investment classes, policies, legal and operational constraints, and requirements for safety, liquidity, returns and stability. This can be followed by a review of how the central bank currently handles risks, with a subsequent discussion on how it might handle risks based upon a more formal risk control framework. For example, it may be determined that the important objectives include obtaining a reasonable level of returns, reducing the probability of negative unrealised losses, maintaining a satisfactory level of liquidity, reducing the necessity of active portfolio turnover, and seeking a reserves level that keeps, as best as possible, the ratio of reserves to short-term external debt above 1. Such a review and discussion should take into account the absorptive capacity of the central bank staff and resources available in terms of analytical modelling. As mentioned above, this process itself is very valuable as it forces many implicit policies and constraints to be reviewed. The review of objectives and requirements can then be translated into objectives and constraints and, on that basis, a prototype model along the lines discussed above can be built.

(b) The process of developing the baseline case

Assuming a model has been constructed, the next step is to analyse the portfolio choices. A baseline case would represent an initial solution to the model and a starting point from whence to understand and gain intuition into the analysis of the asset and liability risks. The baseline case is meant to be conservative and one in which no particular view is taken regarding the direction of the markets other than what the financial markets imply. This baseline can be used to compare other solutions to when one imposes a specific view on future asset prices or introduces new constraints or policies. The baseline case may, for example, make the following more neutral behavioural assumptions about interest and exchange rates behaviour:

- Expected future interest rates satisfy the interest rates implied in the current term structure. In other words, in the sense of expected value, the rates of return on long maturity assets are no different from those obtained by rolling over investments in short assets. This principle would apply to the current rates as well as to any future rates by calibrating the expected interest rates in the future to the forward rates implied today.

- Similarly, one can derive the expected volatility of and correlations among and between the interest and exchange rates and other prices from the current market prices of existing derivatives. If not available, then the historical volatility over the last 12 months or over other periods of these asset prices can be taken. Alternatively, the behaviour of volatility and correlations can be explicitly modelled and estimated.
- For some variables, arbitrage and equilibrium conditions can be incorporated. Arbitrage conditions can be used to determine prices or expected rates of return for some assets, in particular derivatives. Equilibrium conditions can be taken from financial markets or economic theory. In terms of exchange rates, for example, expected values can be computed such that the Uncovered Interest Parity hypothesis holds, i.e. exchange rates are expected to appreciate or depreciate by the differential in interest rates.
- One can allow the stochastic processes to change over time, for example, by introducing time-varying volatility to reflect the possibility of regimes in which there is greater or lesser uncertainty. One can also incorporate the mean-reversion found in many rates of return.

Based on these assumptions, the tree with all possible future exogenous variables can be computed and the model can then be solved (for examples, see Claessens et al., 1998). Expected balance sheets will be generated over time from the model solution along with density functions on factors of interest, including measures of safety, liquidity, returns and stability. Other appropriate measures can also be obtained: for example, risks can be classified into those due to movements in exchange rates, interest rates, etc. Additional measures can be extracted from the density functions, such as VaR-type measures for any particular risk category and over any part of the distribution, including on the upside.

(c) Presenting results and creating benchmarks

Having created the baseline case, it can be presented to and discussed with senior management. Discussions will cover many dimensions, but may focus on two aspects: the preferences senior management have regarding the shapes of the density functions of some of the factors; and the assumptions or views they would like to impose on the behaviour of the exogenous variables. The preferences can be in the form of refining the objective function or the desire to reshape the density functions. In terms of the views on the exogenous variables, management may prefer the assumption that markets will behave as they did last year, or rather that there will be a downturn in the markets, etc. These density-shaping methods and revisions to the assumptions are then translated into revisions to the model, and a new set of outcomes is obtained after solving the model.

This process may take a number of iterations. Once the model has been finalised and solved, it can provide the full density function of each of the factors under consideration, not only returns, but also measures such as unrealised losses, liquidity requirements, portfolio turnover and the ratio of reserves to short-term debt, for example. The results along with the analysis will then typically be presented to higher management, including possibly the board of the bank, using easily understandable concepts and measures, such as VaR measures. Management and the board may request some further revisions but, after a few iterations, a final strategy should be finalised. This process can be repeated every six months to a year or whenever there are substantial changes in market structures, characteristics that alter the investment environment, significant macroeconomic events, or institutional changes.

(d) Example of multiple density shaping

In order to understand how the density-shaping process might proceed, it is instructive to look at an actual example. Figure 9 illustrates the shaping of four density functions simultaneously by replicating a panel used by the software system RisKontroller to shape them. The four variables analysed are wealth measured in dollars, wealth measured in the currency basket, the portfolio average maturity, and the dollar rate of return on assets. This is a simplified example from an actual central bank study with many assumptions similar to those outlined above and using actual simulations, although conducted a few years ago.[11] The first column depicts the four density functions obtained by maximising the expected return of wealth measured in the currency basket with no constraints on portfolio rollovers or any other constraints. The second column represents outcomes of the four densities when maximising the same preference function on wealth again measured in the currency basket, but now conditional on a VaR constraint on the downside risk in terms of dollar wealth, a constraint on negative returns measured in dollars, and with some limits on the allowable maturity structure of the portfolio. As can be seen, the "reshaped" density functions have probability masses for wealth that are much narrower than the original, unconstrained solutions. The tools by which the densities are reshaped include preference functions, CVaR constraints, and general constraints. The tables at the bottom of each column summarise the expected values of the aggregated portfolios in each case. This example produced a final solution with little risk measured in terms of the currency basket. The actual portfolio solution was a 80%, 15%, 5% mix of US dollars, euro and yen respectively that was close to the actual currency basket chosen (which itself represented the proportion of imports in each currency). Interestingly, and with hindsight, the long-duration portfolio as determined here turned out to be very desirable for the country's overall ALM.

(e) Translating into operations and control

Any strategy needs to be translated into operational guidelines for reserves managers. Typically, as part of the normal management process, a benchmark is formulated on the basis of which managers are allowed to invest and are evaluated ex post. Rather than assuming a benchmark, the benchmark will follow directly from the model solution and will specify percentages of different types of assets to be invested. This benchmark can be static, but can be revisited from time to time with reconsideration by management and/or the board, or may be allowed to change over time, with or without the need for reconsideration by management and/or the board.

The next step is to actively manage the portfolio, either to move it towards the benchmark, to rebalance it to the benchmark when there was a deviation from it, or to accommodate changes when using a dynamic benchmark. Guidance on how to implement these changes, or trades, can be provided by a tactical model, which may indicate, for example, the cheapest way to achieve a certain asset mix. The tactical model itself can be based on analysis of short-term market movements and trends, and may use analysis from third parties specialising in, for example, exchange rate forecasting or investment banking analysis and reports. Alternatively, the tactical model can also use an approach similar to that of the strategic asset allocation model. The advantage of the latter is that it would be consistent with the strategic model in the assumptions used and the resulting risk trade-offs. It would require some

[11] In this example only the objective function is measured in the currency basket, with the other constraints measured in dollars (as the central bank reports in dollars).

Figure 9: Density shaping example

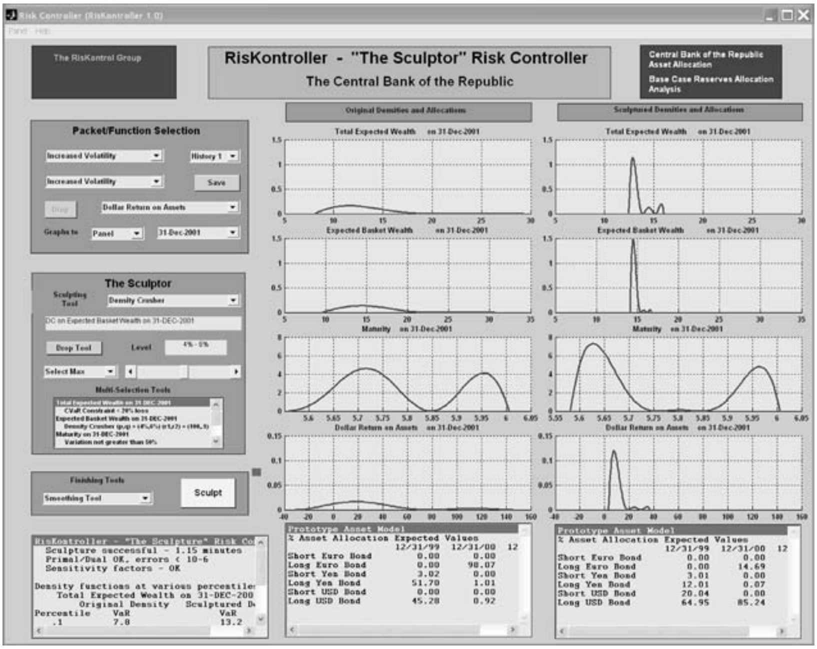

additional modelling, however, and there may be preferences among traders for other tools. The challenge will be to ensure that the use of other tools does not introduce new sources of risks.

The last step is to put in place operational control and reporting formats. Some of these operational control measures can be derived from the tactical and/or strategic models. A matrix of maximum allowable asset changes for each overall reserves size/current asset combinations can be derived, for example, from the strategic model, which will limit overall risks to within a certain desired level for a specified probability. Large trades or trades exceeding the maximums would require special approval. While these limits would only provide order-of-magnitude measures, as they do not take into account the correlation structures, using them will maintain the strategic benchmark, while still providing the traders in asset selection with some latitude. Reporting formats can similarly be derived from the model, and take the form of charts of marginal changes in density functions given a set of investment choices taken or to be taken. Similarly, the model can be used to develop a system to measure the performance of investment managers relative to the benchmark, as it can correct performance for risk taken. Scaling rewards by risk is somewhat complicated in this framework because of the various risk constraints; however, using the model can lead to a more consistent measure of return per unit of risk than by employing commonly-used measures, such as the Sharpe ratio (the ratio of return to standard deviation). This involves computing a return per unit of risk that takes into account perhaps several CVaR constraints on the same distribution, as well as several different distributions, with potentially a combination of both of these over time. However, the description of these risk measures goes beyond the scope of this paper.

(f) Stress tests

The framework can also provide for innovative stress tests, i.e. the risk of some major insolvency or financial shock due to unforeseen events. One can adjust any of the stochastic processes and change their expected values, volatility or correlations. It is even possible to allow these parameters to change over time or circumstances, or to allow volatilities to increase or decrease over time. Alternatively, one may run tests on correlation convergence or divergence, such as that experienced in the fall of 1998 (triggering the LTCM crisis). Events may also be conditioned upon the level and time of any of the exogenous variables. For example, a liquidity crisis triggered by a drop in GDP or a shock to the country's terms of trade can be modelled as an increase in volatility and increased yields on short-term debt and/ or constraints that limit the rollover of short-term debt. Other suggestions, such as those reported in IMF (2001b), can easily be accommodated. More generally, the kinds of quantitative stress tests that can be incorporated are unlimited as the approach is very flexible.

(g) Policy comparison and analysis

This framework provides a detailed analysis of the comparison of alternative stress tests and the quantification of their costs. It also enables the asset and liability portfolio to be reallocated where necessary to hedge against specific risk of financial stress. This in turn provides a quantitative measure of the sensitivity of costs and portfolio allocation in specific financial stress situations, whether financial, economic, or political.

This framework also enables different policies that affect decision variables to be analysed, and their costs to be quantified. A simple policy might contain limits on portfolio duration, for example. The model can then be run with limits, and again with no or more relaxed limits. The two policies can then be compared, not only in terms of total costs (or rates of return) or portfolio allocations, but also in terms of the density functions of outcomes under the two policies. Such a comparison motivated an earlier study (see Figure 10 below), where the density function of the difference between two policies was explicitly compared.

Figure 10: Policy comparison

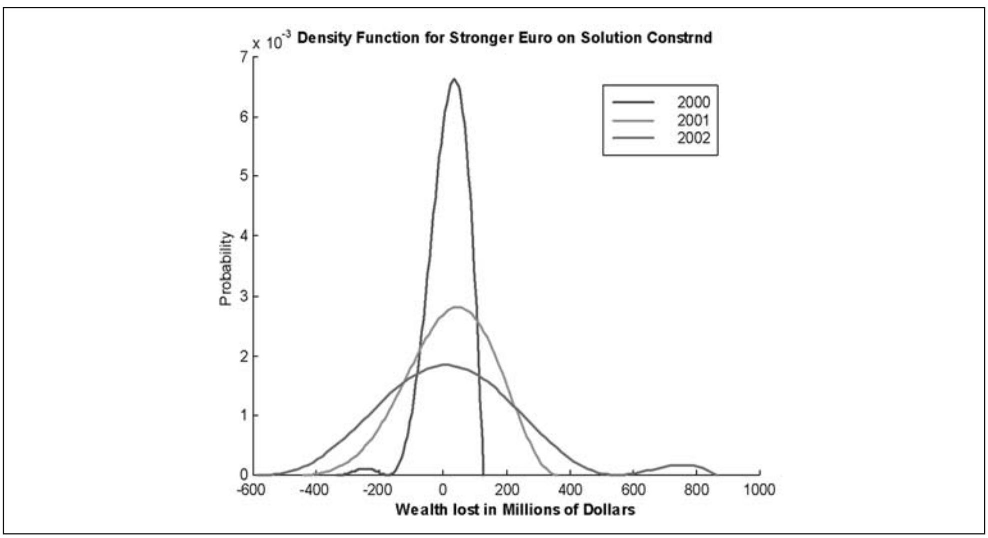

In this case, the density functions plotted the wealth from choosing one policy over another at different points in time. The densities are fairly symmetrical around zero, indicating that at this point the policy did not have a lot of impact on average, except for a slightly greater loss in the last year. However, there was a difference in distribution between the two policies, which may have meant that one was preferable to the other.

(h) Institutional organisational considerations

Finally, the key to any successful foreign exchange reserves risk management programme is a proper organisational structure that integrates various aspects and facilitates tasks. It is of course difficult to generalise as to what constitutes an appropriate organisational division of risk management responsibilities, as institutions differ in many dimensions; however, some of the following elements are likely to be needed.

A Middle Office (MO) would undertake much of the analytical work, with the other units liaising with it. The investment managers would execute within their limits most trades without direct intervention from the MO. When trades are large or limits need to be exceeded, then the traders would need further analysis from the MO and/or approval from executive managers. Regular strategy meetings would take place between traders, MO staff, and selected executive managers. Executive managers of the central bank would liaise with those of the Ministry of Finance and other relevant ministries in order to incorporate issues of mutual concern into the analysis.

An important aspect of risk management would be model evaluation and modification. MO staff would do much of the testing and reporting, as they would possess the correct analytical skills. Specifically, the MO would be responsible for estimating and adjusting the uncertainty structure of the exogenous variables described by the tree; for formulating the models capturing the necessary concerns, capturing the density functions of the various factors; and for properly executing identified stress tests. However, another unit, here named the Audit Department, would conduct the review and report on the results of the tests. Key final roles would remain with senior management, such as the responsibility to identify the appropriate specific issues for safety, liquidity, returns and stability. Senior management would also have the responsibility to identify stress tests and the appropriate assumptions under which the benchmark should be created. The Audit Department would have the responsibility to report back on how successful these stress tests were. The intention of the entire evaluation process is to be able to identify weak spots in the risk management process and to modify them on an ongoing basis. How these proposed responsibilities are allocated among departments and managed will vary greatly from one central bank to the next, depending on the degree of hierarchical structure among other factors.

6 Conclusions

We have presented a framework for foreign exchange reserves management that combines tactical asset allocation considerations with broad macroeconomic, macro-prudential risk and sovereign debt management considerations. Our framework for foreign exchange reserves management allows for very general objective definitions and does not restrict the class of eligible stochastic processes or limit the investment universe. It also allows for easy feedback between outcomes and decision variables by including various tools that can reshape densities. The model can also be operated using a PC-based platform. We show the possible application of our approach to several common reserves management problems, and demonstrate how it provides institutional guidance through developing benchmarks, portfolio evaluation criteria, stress tests and management reporting formats.

We see our approach as an important complement to the various ALM tools currently being used by central banks and commercial banks around the world. Our approach is more demanding than other approaches in terms of analytical modelling with regard to objectives, constraints and assumptions. We think this is worthwhile not only because of the improved quality of solutions, but also because it requires a very explicit process of model development that will help clarify the strategic aspects involved in risk management. Often, these strategic aspects, including linkages between macroeconomic, microeconomic and financial risks are left out or only implicitly treated in other approaches. At the same time, the use of numerical approaches to solving using a dynamic stochastic framework allows for analytical rigour while maintaining ample modelling flexibility.

Finally, our approach not only provides the complete density functions of outcomes of a selected portfolio (instead of just summary measures), but also defines criteria for the shape of density functions in terms that senior management can understand. Then, satisfying these criteria, our approach determines the strategic portfolio allocation that is optimal over time with respect to an objective – such as returns or the likelihood of an indicator, such as reserves to short-term debt, falling within a specified range. The intuitive approach allows our approach to be supported at each level in an institution, while its flexibility allows it to be adapted to the unique requirements of each individual central bank. Because of its extremely flexible structure, the framework is not only applicable to developed countries, but also to developing countries that may face greater economic and institutional constraints.

References

Acerbi, C. and D. Tasche, 2002. "On the Coherence of Expected Shortfall", RiskLab paper, ETH Zurich, 19 April 2002.

Asian Development Bank, 2001. *Government Bond Market Development in Asia*, ed. by K. Yun-Hwan, Manila.

Bahmani-Oskooee, M. and F. Brown, 2002. "Demand for International Reserves: A Review Article", *Applied Economics*, 34, pp. 1209-26.

Bank for International Settlements, 2002. "Central Bank Survey of Foreign Exchange and Derivatives Market Activity", Press Release, Basel.

Bank for International Settlements, 2002. "The Development of Bond Markets in Emerging Economies", BIS Papers No. 11, June, Basel.

Birge, J. R. and F. Louveaux, 1997. *Introduction to Stochastic Programming*, Springer Series in Operations Research, Springer, New York.

Blejer, M. I. and L. Schumacher, 1999. "Central Bank Vulnerability and the Credibility of its Commitments: A Value-at-Risk Approach", *Journal of Risk*, Volume 2/Number 1, Fall.

Blejer, M. I. and L. Schumacher, 2000. "Central Banks' Use of Derivatives and Other Contingent Liabilities: Analytical Issues and Policy Implications", *IMF Working Paper*, Asia and Pacific Department and Monetary Exchange Affairs Department, March, Washington.

Calvo, G. and E. Mendoza, 1996. "Mexico's Balance-of-Payments Crisis: A Chronicle of a Death Foretold", *Journal of International Economics*, 41, pp. 235-64.

Calvo, G., 1996. *Money, Exchange Rates, and Output*, MIT Press, Cambridge.

Cassard, M. and D. Folkerts-Landau, 2000. "Risk Management of Sovereign Assets and Liabilities", in *Sovereign Assets and Liabilities Management*, ed. by M. Cassard and D. Folkerts-Landau, Proceedings of a Conference held in Hong Kong SAR. IMF, Washington.

Claessens, S. and J. Kreuser, 2004. "Strategic Foreign Reserves Risk Management – A Framework and Applications", mimeo, University of Amsterdam. Submitted for publication, 2004.

Claessens, S., J. Kreuser and R. J.-B. Wets, 2000. "Strategic Risk Management for Developing Countries: The Colombia Case Study", Paper presented at the Second Inter-regional Debt Management Conference, UNCTAD, Geneva, 3-5 April 2000.

Claessens, S., J. Kreuser, L. Seigel and R. J.-B. Wets, 1998. "A Tool for Strategic Asset Liability Management"*, World Bank Working Paper,* Research Project Ref. No. 681-23, Washington, 12 March 1998.

Claessens, S., 1991. "Integrating Commodity and Exchange Rate Risk Management: Implications for External Debt Management", in *Commodity Risk Management and Finance,* ed. by T. Priovolos and R. C. Duncan, John Hopkins University Press/World Bank.

Claessens, S., 1994. *Risk Management in Developing Countries*, World Bank Technical Paper No. 235.

De Broeck, M., D. Guillaume and E. Van der Stichele, 1998. "Structural Reforms in Government Bond Markets", IMF Working Paper No. 108, August, Washington.

Dooley, M., 2000. "A Model of Crises in Emerging Markets", *Economic Journal*, 110, No. 460, January, pp. 256-72.

Dupačová, J., G. Consigli and S. W. Wallace, 2000. "Scenarios for Multi-stage Stochastic Programs", *Baltzer Journals*, 2 June 2000.

Frenkel, Jacob, 1980. "International Reserves under Pegged Exchange Rates and Managed Float: Corrections and Extensions", *Journal of Monetary Economics*, 6, pp. 295-302.

Frenkel, J. and B. Jovanovic, 1981. "Optimal International Reserves: A Stochastic Framework", *Economic Journal 91*, pp. 507-14.

Göing, A., 1996. "Estimation in Financial Models", RiskLab Report, ETH Zurich, Department of Mathematics, January.

Harwood, A., 2000. "Building Local Bond Markets: An Asian Perspective", International Finance Corporation, September, Washington.

Herring, R. J. and N. Chatusripitak, 2000. "The Case of the Missing Market: The Bond Market and Why It Matters for Financial Development", Asian Development Bank Institute Working Paper No. 11, July.

IMF, 2000. "Debt- and Reserve-Related Indicators of External Vulnerability", Policy Development and Review Department, 23 March 2000.

IMF, 2001a. "Financial Soundness Indicators: Policy Paper", prepared by the Monetary and Exchange Affairs Department and the Statistics Department, 4 June 2001, Washington.

IMF, 2001b. "Stress Testing of Financial Systems: An Overview of Issues, Methodologies, and FSAP Experiences", by W. Blaschke, M. T. Jones, G. Majnoni, and S. Martinez Peria, June, Washington.

IMF, 2001c. "Guidelines for Foreign Exchange Reserves management", 20 September 2001, updated 13 June 2001, International Monetary Fund, Washington.

International Monetary Fund/World Bank, 2001. "Guidelines for Public Debt Management", 21 March 2001, Washington.

International Organization of Securities Commissions, 2002. “The Development of Corporate Bond Markets in Emerging Market Countries”, Report of the Emerging Markets Committee of IOSCO, May.

Jeanne, O., 2000. “Debt Maturity and the Global Financial Architecture”, European Economic Review, 44:4-6, pp. 719-27.

Jorion, P., 2000. *Value at Risk, Second Edition*, McGraw-Hill.

Kreuser, J. L. 2002. “The Rapid Implementation of Asset/Liability Models for Sovereign Risk Management”, in *Programming Languages and Systems in Computational Economics and Finance*, Kluwer Academic Publishers.

Rigaudy, J.-F., 2000. “Trends in Central Bank Reserves Management”, in *Sovereign Assets and Liabilities Management*, ed. by M. Cassard and D. Folkerts-Landau, Proceedings of a Conference held in Hong Kong SAR, International Monetary Fund, Washington.

Rockafellar, R. T. and S. Uryasev, 2001. “Conditional Value-at-Risk for General Loss Distributions”, Risk Management and Financial Engineering Lab, University of Florida, Gainesville, 3 July 2001.

Rodrik, D. and A. Velasco, 1999. “Short-term Capital Flows”, in J. Stiglitz and B. Plescovik (eds), 11th Annual Bank Conference on Development Economics, World Bank, Washington.

Sachs J., A. Tornell and A. Velasco, 1996. “Financial Crises in Emerging Markets: The Lessons from 1995”, NBER Working Paper No. 5576, May.

Schneider, M. and A. Tornell, 2000. “Balance Sheet Effects, Bailout Guarantees and Financial Crises”, NBER Working Paper No. 8060, December.

Sundaresan, S. M., 2000. “Continuous-Time Methods in Finance: A Review and an Assessment”, *Journal of Finance*, Vol. 55, No. 4, August.

Turner, P., 2002. “Bond Markets in Emerging Economies: An Overview of Policy Issues”, in *The Development of Bond Markets in Emerging Economies*, BIS Papers No. 11, June.

Uryasev, S., 2000. “Conditional Value-at-Risk: Optimization Algorithms and Applications”, *Financial Engineering News*, February.

Wijnholds, J. O. de Beaufort and A. Kapteyn, 2001. “Reserve Adequacy in Emerging Market Economies”, IMF Working Paper, WP/01/143, September.

World Bank/International Monetary Fund, 2001. *Developing Government Bond Markets: A Handbook*, May, Washington.

Ziemba, W. T. and J. M. Mulvey, 1998. *Worldwide Asset and Liability Modeling*, Cambridge University Press, UK.

4

Asset allocation for central banks: optimally combining liquidity, duration, currency and non-government risk

Stephen J. Fisher and Min C. Lie

Abstract

We present a strategic asset allocation framework focused on central bank reserves management. The framework (a) explicitly examines total portfolio risk and return, (b) is inclusive of a wide variety of assets such as government bonds, non-government bonds (MBS, ABS, corporate bonds, etc.), equities and currency, and (c) ensures there is sufficient liquidity in the portfolio for trade and intervention requirements. Examples of potential optimal asset allocations are presented. The analysis shows that a typical central bank can significantly increase the efficiency of its portfolio. A reasonable calibration of the model suggests that relaxing liquidity, currency, duration and credit constraints could achieve up to an additional 140 basis points worth of return for current risk levels or reduce risk by around 150 basis points and still achieve the same expected return.

1 Introduction

A country's foreign reserves represent the national wealth of the nation. Foreign reserves can be employed in the following ways, depending on the objective of the nation:

(i) To purchase foreign goods and services

(ii) To service the nation's foreign debt

(iii) To manage the level of the exchange rate through market intervention

(iv) To invest reserves to generate future wealth

(v) To fund domestic fiscal spending programmes

(vi) To rebate reserves to citizens through lump-sum transfers and tax cuts.

A central bank (CB) is responsible for managing the nation's foreign reserves. In Asia, for example, the assets of many CBs have grown rapidly since the Asian crisis in 1997, reflecting a healthy trade position and the decision by the authorities to accumulate foreign assets. This decision was partly driven by the desire to appear financially strong to the external economy and avoid a speculative attack on the currency. Collectively, the external financial position of Asian CBs is ranked highest in the world, with total assets in excess of USD 1.7 trillion as of June 2003.

In general, CBs view their stewardship role in a highly conservative way and their investment philosophy reflects this. CBs are highly risk-averse investors in general, with the bulk of their assets invested in short-dated securities such as T-bills, time deposits and highly-rated government bonds. The desire to protect the level of reserves from short-term volatility has led these investors into relatively low-return investment programmes. This approach would seem to be acceptable for CBs with relatively low reserve levels and whose main

We would like to thank Michael Granito, Patrick Thomson and two anonymous referees for their helpful comments.

concern is maintaining sufficient liquidity to cover import costs, debt service obligations and intervention requirements. Capital preservation is a common investment objective for the majority of CBs.

Recently, however, a trend has been developing amongst the wealthier nations, who are seeking higher returns on their reserve assets. Nations with significant reserve positions have little difficulty financing imports or funding debt service obligations. The level of 'free reserves' (i.e. the amount of reserves in excess of what is needed for import, debt and intervention purposes) can be many times greater than the nation's liabilities. This means that the investment horizon for planning purposes is greater than several months or a year, as short-term negative wealth shocks can be absorbed without affecting a nation's capacity to finance its external transactions. For these nations, intergenerational wealth accumulation is the primary factor driving reserves management investment policy, and their investment horizon can extend to many years or even decades. Long-term investors can generate significantly higher returns by bearing short-term risk.

This paper focuses on how a CB can formulate and implement policies governing its reserves management that are portfolio-efficient. Our experience is that the majority of CBs approach the benchmark setting process in a similar manner. The typical approach imposes significant constraints on currency allocation, country allocation, portfolio maturity and non-government exposure. The results in this paper show that these constraints, while considered prudent, severely limit the opportunities for a CB to diversify risk. We demonstrate an alternative strategic asset allocation framework consistent with a CB's objectives that can significantly reduce risk and/or increase return.

The paper is organised as follows. Section 2 describes and provides a critique of what we see as the typical method currently employed by CBs to determine the strategic asset allocation benchmark for their reserves. The significant features of this method are as follows:

1. Investments are restricted to government and high-rated bonds.
2. A liquidity tranche is set aside for intervention purposes.
3. Currency and country allocation are treated identically so that the currency target equals the bond target.
4. Within each country's bond market, portfolio duration is restricted to control downside risk.

We contend that this typical asset allocation process is overconstrained, leading to portfolio inefficiency. In Section 3 we describe an alternative framework for the asset allocation decision, which offers the following key enhancements:

1. A broad investment universe is specified encompassing government bonds, non government bonds (MBS, ABS, corporate bonds, etc.), equities and currency.
2. The currency and country asset allocations can differ.
3. Portfolio liquidity is controlled through a transaction cost constraint. This enables the optimisation to trade off the costs and benefits of investing in less liquid securities.
4. Downside risk is controlled at the total portfolio level instead of the individual country bond market level.
5. Downgrade and default risk is controlled at the total portfolio level rather than by imposing minimum credit rating constraints.

The combined effect of these enhancements is to provide more flexibility in achieving a diversified risk-controlled portfolio.

Section 3 starts by constructing a base case scenario that is consistent with typical central bank practice. The section then goes on to relax the constraints and compares several portfolio outcomes which either reduce risk or add return. The main finding is that relaxing liquidity, currency, duration and credit constraints suggests that the portfolio could achieve up to an additional 140 b.p. worth of return relative to the base case or reduce risk by around 150 b.p. and still achieve the same expected return.

We must point out that, while we advocate the use of modern quantitative techniques in central bank strategic asset allocation, adopting this approach has its drawbacks. In particular, quantitative models make assumptions about investor risk tolerance and the structure of capital markets. They also require the user to estimate parameters that may exhibit instability during periods of market stress. We discuss these drawbacks in more detail in Section 3.

Section 4 concludes the paper.

2 Current central bank asset allocation practice

The majority of CBs around the world approach the asset allocation problem in a similar way. This section describes some of the common features used by CBs when determining the currency and country allocation for their portfolios. This process is explained in detail so that a clear comparison can be made with the optimal asset allocation approach described in Section 3.

It is striking how little has been published about the asset allocation process employed by CBs for reserves management, despite the substantial sum of wealth that these institutions control. The great majority of public research has been focused on determining the optimal *level* of reserves as opposed to the way, once accumulated, reserves are *managed*.[1] A fascinating paper by Nugee (2000) from the Bank of England's Centre for Central Bank Studies describes the factors that a CB must consider in the asset allocation process.[2] Nugee's paper does not, however, provide concrete recommendations or a quantitative framework to anchor the process.[3] One contribution of our paper is to draw on our experience to deliver a quantitative framework for CB asset allocation, and compare current practice with efficiency-improving enhancements.

We start by building a picture of a typical CB that reports its performance with a USD base.[4] A typical method for setting the CB's strategic benchmark asset allocation can be summarised as follows:

1. *The investment objective is to preserve capital and liquidity*. While a CB's investment policy objective is not clearly stated in general, we would classify the typical asset allocation as being conservative and focused on protecting the portfolio from any losses on an annual basis. The desire to preserve capital and invest in low-risk, highly liquid securities is very common amongst CBs. Historically, our impression is that achieving higher portfolio returns has always been a secondary objective for a CB.

---

[1] See, for example, Aizenman and Marion (2002).

[2] The Bank of England's Centre for Central Bank Studies can be likened to the Harvard Business School for Central Bankers.

[3] A typical and practical example is briefly described in the Reserve Bank of Australia's Annual Report 2003.

[4] By far the majority of central banks we have contacted report their reserve performance with a USD base. A small number report performance using a multi-currency base such as the IMF's Special Drawing Right or a custom currency mix. Occasionally, a CB is required to report performance in its local currency, an unusual requirement as foreign reserves are managed outside the local currency system.

The capital preservation objective implies that a CB has an extremely short investment horizon. As investors, central bankers carry many responsibilities. First, they are custodians of the financial system and must be in a position to intervene in the market during periods of crisis. This calls for access to short-term liquidity. Second, in general, the return on the reserves portfolio is reviewed annually by the government and is publicly disclosed. A negative return in any year can easily be portrayed as mismanagement in a public forum so that there is little reward for risk-taking.

For these reasons we choose to characterise the investment horizon for a typical CB as around one year in this paper, although there are several high profile exceptions.[5] For countries with large reserve holdings, there is recognition that a short-term investment horizon is costly in terms of long-run return. These countries have set aside longer-term investment portfolios with growth objectives, investing instead in bonds, property and equities.

Elsewhere, it has been argued that longer-term bonds, as opposed to short-term bonds, are more natural assets for investors such as CBs who maintain significant fixed income positions over many years.[6] This would offer some additional yield in an upward sloping curve environment. Many CBs would agree with this view, but are still concerned about reporting a negative mark-to-market return in any year. A common practice is to set aside a 'hold-to-maturity' portfolio of longer-dated bonds that is not marked-to-market and not available for liquidation.

2. *Restrict investments to fixed income securities only rated AA- or better.* Fixed income securities are relatively low volatility investments compared with, say, equity investments, whereas high-rated bonds are virtually default-free. While restricting investments to high-rated fixed income securities is common practice amongst CBs, it is not necessarily the most efficient way to construct a portfolio. The fact that securities are not perfectly correlated with each other means that their individual volatilities will be offset in the portfolio and hence total portfolio risk may be reduced.[7] The important guiding principle for portfolio construction is that *more securities are preferred to less irrespective of their structure,* and thus any rule that restricts the investment opportunity set is a priori potentially sub-optimal. The convention that limits the CB to investing in fixed income securities alone is potentially harmful to portfolio efficiency.
3. *Establish a liquidity tranche to cover day-to-day currency flows.* Our experience is that CBs often create a 'liquidity tranche' portfolio that is separate from the 'investment portfolio'. The liquidity tranche is a different investment portfolio designed to finance the day-to-day foreign exchange requirements, facilitating trade and financial flows. The liquidity tranche is generally managed entirely in T-bills and time deposits, which is a different asset allocation benchmark than the remainder of the fund. As these securities carry very low risk, the long-run rate of return on the liquidity tranche will be lower than for the rest of the reserves under management.

   Whether establishing a liquidity tranche is an optimal allocation of risk depends on (a) the practical requirements for liquidity in the CB's daily operating procedures and (b) exactly how much of this tranche is traded. The CB needs to trade off the marginal expected

[5] This is especially the case in Asia and the oil-rich nations.

[6] Cochrane (1999) reviews the literature on this subject. The premise is that long-term bonds create more return certainty for investors with long-term horizons. Fitting the bond maturity to the return horizon locks in a certain yield (assuming no default).

[7] This is the fundamental contribution made by Markowitz in 1956.

additional return from other, less liquid investments versus the marginal transaction costs incurred.

To illustrate this point, suppose that a liquidity tranche of, say, USD 1 billion in short-term, highly liquid investments is earning 2% per year. Let us suppose further that the opportunity cost of investing in, say, mortgage-backed securities (MBS) brings an additional 1% return per year. Suppose also that the transaction cost associated with trading MBS is 6 b.p. per trade, compared with zero for T-bills. Given the 1% additional return and the 6 b.p. additional trading costs for MBS, successfully balancing benefits and costs comes down to the question of how many times does the CB have to use the liquidity tranche to *intervene* in the currency market before it becomes optimal to invest in T-bills? The answer to this question is exactly 100 b.p. / 6 b.p. = 16.6 times. If the CB turns over the USD 1 billion tranche less than 16.6 times per year, then it is optimal to invest in MBS and incur the transaction costs from liquidating these securities when liquidity is required rather than hold T-bills.

The approach we adopt in Section 3 is to determine the liquidity allocation and the allocation to other investment assets simultaneously as part of the strategic asset allocation. An optimisation determines the optimal liquidity holdings with regard to the investment opportunities and transaction costs on all assets, and the overall objective.

4. *Determine the currency allocation based on four factors*. Our experience is that CBs determine their currency allocation based on four factors as follows:
   (a) Component weights of the JP Morgan Global Bond Index
   (b) Current payment currency proportions
   (c) Foreign debt service currency proportions
   (d) The observed currency allocation of other CB's in their peer group.

   We will refer to this as a 'four-factor model'. Each CB places different emphasis on each factor depending on its reserve position and objectives. These factors are motivated by several implicit themes. First, the use of the JP Morgan Global Bond Index implies a desire to invest in a diversified high-grade bond portfolio proportional to the relative supply conditions in the developed countries. Second, the current payment and foreign debt service factors imply an asset–liability matching objective. Third, the use of other CB allocations implies a desire to mimic the CB's peer group.[8]

   In contrast to the four-factor approach, our approach in Section 3 treats the currency allocation question within a risk/return framework. The four-factor model says nothing about expected return or risk in currency markets, even though these would seem to be critical factors.

5. *Fix the bond market allocations in the same proportion as the currency allocations.* The majority of CBs contacted in the context of this study determine their bond allocations from the currency decision, so that *the country weights equal the currency weights*. The reason this occurs is because the CB's portfolio is invested in physical securities in each currency target, and hence the country bond allocation must reflect the currency targets. Restricting investments to physical securities implies that the country interest rate risk decision and the currency risk decision cannot be separated.

   This feature of the benchmark process is perhaps the weakest in our view, as it can lead to a significant loss in diversification. For instance, if the currency mix is heavily skewed towards one country (for example, 90% to the US), then the portfolio will not be able to take advantage of the significant risk reduction associated with a globally diversified

---

[8] Data are available from the IMF, which publishes the asset allocation for each member CB at the end of the year.

portfolio.[9] Treating currency and country as separate benchmark decisions can lift the country bond allocation restriction. Currency targets can then be implemented using 'currency overlay'.[10]

6. *Within each country, set the bond portfolio duration so that no negative returns are experienced in any year with a high level of confidence.* This step in the process is designed to ensure that the CB's reserves do not experience a negative return in any given year. The interesting feature of this step is that the capital preservation objective is applied to *each country investment* rather than to the portfolio as a whole. When aggregated, this means that the portfolio will be over-protected on the downside as the country positions already have a tendency to diversify, thus limiting downside risk. The enhanced framework in Section 3 controls downside risk at the total portfolio level. The benefit is that slightly more risk can be taken within each country so that slightly more return can be earned for the portfolio as a whole.

In summary, the typical CB strategic asset allocation setting process is a set of sequential, independent steps that may be costly in terms of portfolio efficiency. The process can be improved in the following ways:

1. Provide a clear investment policy objective in terms of risk, return and investment horizon.
2. Broaden the investment opportunity set to include assets other than fixed income securities.
3. Determine the liquidity and longer-dated investment portfolios simultaneously as part of one optimisation.
4. Treat currency and country allocation separately so that the currency allocation decision does not drive the country allocation decision. A corollary to this is for the CB to develop expertise in currency overlay investing.
5. Apply the downside risk objective to the entire portfolio rather than on a country-by-country basis.

This is the focus of the next section.

3 A strategic asset allocation framework for central banks

The discussion up to this point has examined the typical CB benchmark setting process and highlighted some areas for enhancement. In sub-section 3.1 we first outline the asset allocation framework. Sub-section 3.2 sets up a base case that shows the set of efficient investment portfolios under the typical central bank approach to asset allocation. Sub-sections 3.3 to 3.5 then relax the currency, downside risk and credit constraints respectively.[11]

---

[9] For instance, the full maturity US component of the JP Morgan Global Bond Index (GBI) has a historical return standard deviation of about 6.5%, whereas the full maturity G5 components of the JP Morgan GBI (hedged into USD) have a historical return standard deviation of about 4%. This risk reduction comes primarily from country bond diversification.

[10] Currency overlay is a technique whereby the currency exposure target of the portfolio is achieved by trading currency forwards. Trading forward contracts means that there is no need to trade the underlying physical bonds to, say, reduce the exposure to the USD and increase exposure to the euro.

[11] We relax the constraints in this order for the following reason. The currency and bond correlations are low so that breaking the equality constraint primarily affects the ability to control risk. Next, allowing the portfolio to exploit the entire yield curve primarily affects expected return. Finally, adding in non-government securities is generally a last step for a central bank, so it is natural to judge the efficiency gain from relaxing this constraint after the gains from traditional government assets have been fully exploited.

We demonstrate that it is possible to achieve significant efficiency gains for reserves management.

3.1 Description of the asset allocation framework

The aim of the asset allocation framework is to derive a target portfolio that maximises return given a risk target, subject to constraints on liquidity, credit quality and currency allocations. The model we present is a version of the static mean-variance optimisation framework that is commonly employed in strategic asset allocation as set out in the Technical Appendix.[12]

Inputs to the model are expected returns, expected risks, covariances across asset classes and investment constraint parameters. These constitute the set of 'capital market assumptions' for the model, and are formulated with a long horizon in mind (for example five or more years). An example of the capital market assumptions employed by JP Morgan Fleming in this analysis is provided in the Technical Appendix.[13]

The choice of model parameters is critical in determining the portfolio allocations within the framework. These parameters must be estimated, and this exposes the main weakness in the analysis. The model we use is a static one, which means that we do not examine the time dependencies that exist in the real world. One example is the tendency for return correlations to be unstable in the short run; these moreover tend to rise during periods of extreme volatility, as in August 1998 or July 2003. In the event of the latter, portfolio diversification is lost during those periods when it is needed most.[14] In practice, when using quantitative models, we recommend stress testing of chosen allocations to changes in key assumptions, and the application of an overall 'reasonableness test' to a chosen portfolio.[15]

The output from the model is a set of portfolios that combine assets efficiently to maximise return consistent with various levels of risk. These portfolios provide a set of investment options to the CB's senior managers for review and debate. The final portfolio choice rests with the Investment Policy Committee, which is responsible for deciding on the strategic objective for the CB's reserves. This constitutes the benchmark.

The model is solved using constrained quadratic programming techniques. We use proprietary software based on the solution algorithm in Markowitz (1999). However, there are a number of other packages that are available to conduct the analysis.[16]

---

[12] We choose the simplest mean-variance model we can envisage to emphasise the power of the results from relaxing currency, duration and non-government bond constraints on the problem. The base case example can be viewed as a "straw man" that is easily deconstructed using standard analysis. We recognise, however, that there are more complex optimisation techniques that can enrich the framework. For instance, extensions to a multi-period analysis, asymmetric risk tolerance and adding in non-normal distributional assumptions are worthy enhancements.

[13] JP Morgan Fleming's capital market assumptions represent our view of the relative risks and returns across markets. They are not definitive, however, and other market participants may have different views of the long-run relative returns on offer to investors. The assumption set is intended to be a starting point from which a CB can debate the relative merits of the assumptions and develop its own view of market opportunities.

[14] We thank the referee for pointing this out.

[15] A 'reasonable' portfolio is one which, when compared to the current asset allocation, exploits a logical opportunity that flows from defensible assumptions. Put another way, central bankers need to be able to sleep at night.

[16] The framework we present differs from the Black-Litterman optimisation model in the following way. Simply put, the Black-Litterman model takes a given asset allocation and covariance matrix, and computes the expected returns that must be earned to support the asset allocation. In contrast, we take a set of expected returns and a covariance matrix, and then compute the asset allocation implied by these assumptions. The mean-variance analytical framework is common to each approach. Each approach has its advantages and disadvantages. Nevertheless, we prefer to focus our assumption set on the capital market opportunities in the investment markets and then solve for asset allocation, since this encourages debate as to the level of equilibrium returns.

Figure 3.1: Efficient frontier for baseline constraint set

(in %)

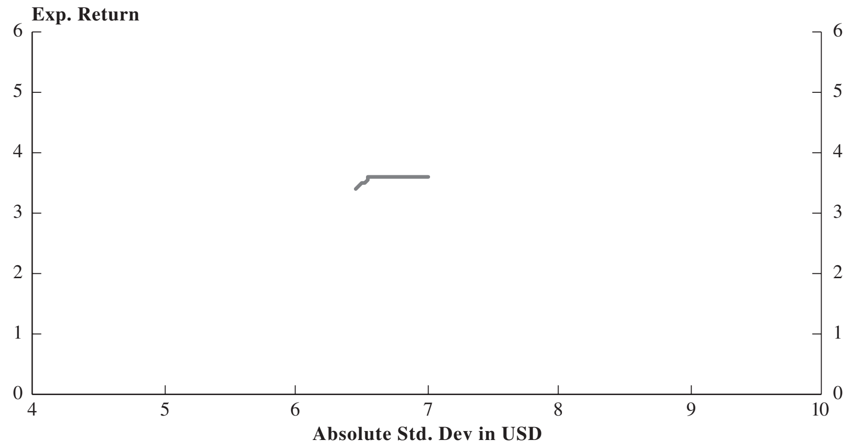

Currency weight set to bond allocation, within country duration constraints, zero allocation to non government bonds.

3.2 Base case: the typical central bank approach

The following example demonstrates the potential for efficiency gains in terms of increasing return and decreasing risk for the CB's portfolio. The example is designed to reflect our understanding of the CB's actual investment portfolio.

We assume the CB reports returns in USD. The base case example assumes that assets are invested with the following constraints:

1. The US allocation is 30%, the European allocation is 40% and the Japanese allocation is 30%.
2. Cash is constrained to be at least 5% of assets in each country.
3. Currency is constrained to be equal to the country bond allocations.
4. Country-level bond portfolio durations are constrained to a maximum of 1.7 years in the US and Europe and 0.5 years in Japan to avoid a negative return in each sub-portfolio.
5. There is a zero allocation to non-government bonds.

The following diagram shows the efficient frontier that is generated for the base case scenario, which is highly constrained.

Figure 3.1 shows the expected risk and return in USD terms for the set of optimal portfolios solving the quadratic optimisation. As the optimisation is highly constrained, the efficient frontier shows very little variation, with expected returns ranging between 3.4% and 3.6%, while expected risk for the portfolio set is tightly clustered between 6.5% and 7.0%. In our view, this is a realistic reflection of the portfolio choices that a typical CB currently confronts given its current benchmark setting process. Absolute volatilities for these portfolios in the order of 6.5-7.0% are consistent with many CBs' current portfolio positions.

Selecting a portfolio on the efficient frontier depends on the CB's risk aversion. CBs with higher risk aversion choose portfolios with lower risk. For expositional purposes, an example of the asset allocation for one portfolio on the efficient frontier reads as follows:

Table 3.1 An example of one portfolio on the efficient frontier for the baseline case: Risk = 6.5%.

Expected Return	3.48%
Votality	6.50%
Sharpe Ratio	0.54

Portfolio Summary	
Currency	
USD	31%
Euro	40%
JPY	29%
Country	
US	30%
Europe	41%
Japan	29%
Hedging (non-US)	
Unhedged	48%
Hedged to other currency (for consistency)	22%
Investment Type	
Treasury	100%
Non-Treasury Bonds	0%
Equity	0%
Liquidity	
Cash	59%
Govt 1-3	29%
US Duration	1.10
Europe Duration	1.70
Japan Duration	0.50
Transaction Cost – Normal	0.06%
Transaction Cost – Distressed	0.12%
Credit Quality	
Mortgages	0%
ABS	0%
AAA/AA	0%
A	0%
BBB	0%
Downgrade Probability	0.00%
Default Probability	0.00%

This portfolio has an expected return of 3.48% with a volatility of 6.5%. The country and currency allocations are the same. Note that given the expected return assumptions, the optimisation chooses to invest up to maximum duration (1.7 years) in Europe and Japan while constraining the US to 1.10 years. 59% of the portfolio is invested in cash instruments. The portfolio is also very cheap to liquidate, amounting to 6 b.p. under normal market conditions, indicating a high degree of liquidity.

We will refer back to this portfolio for comparative purposes in the next few sub-sections.

Figure 3.2 Efficient frontier relaxing the currency and country bond allocation constraints *(in %)*

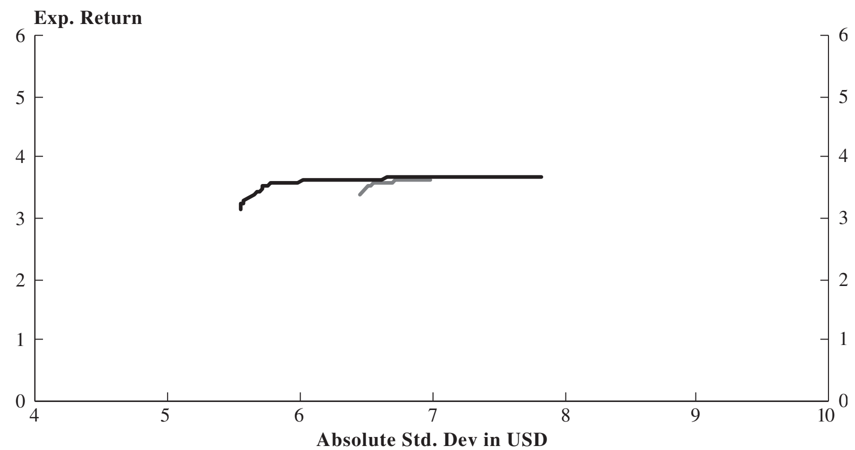

Allowing currency and bonds to reside in a 25-40% range and delinking bond allocation from currency.

3.3 Relaxing the currency and bond constraints

Relaxing the currency constraint shows how the set of optimal portfolios expands by both increasing expected return and reducing risk. The range of risk/return options for the CB expands significantly as there are more dimensions within which it can trade off risk and return in markets.

First, Figure 3.2 shows the effect of removing the constraint that the currency allocation should equal the country bond allocation in the CB's portfolio. This is a common constraint for CBs and, as the diagram shows, the cost to portfolio efficiency of this practice is quite significant. Relaxing the currency and bond constraints is achieved by allowing the allocation to each currency and bond allocation to fluctuate in the range 25-40%. This is a slight relaxation compared with the base case scenario. As the currency and bond returns have a relatively low correlation (e.g. the correlation of hedged and unhedged European 10-year bond returns in USD terms is 0.55), there is significant scope for diversification.

Figure 3.2 shows how the efficient frontier rises and expands once the currency constraint is relaxed. According to our assumptions, separating out the currency and bond allocation decisions for the portfolio provides a significant opportunity to reduce risk. Most of this benefit comes from simply having the ability to diversify country and currency risk in the portfolio rather than from investing in the extended asset markets such as non-government securities. The minimum risk portfolio in this scenario records a volatility of 5.55%, which is almost a full 1% lower than the minimum risk portfolio under the base case.

Table 3.2: The minimum risk portfolio after relaxing the currency constraint: Risk = 5.5%

Expected Return	3.17%
Volatility	5.55%
Sharpe Ratio	0.57
Portfolio Summary	
Currency	
USD	40%
Euro	35%
JPY	25%
Country	
US	35%
Europe	40%
Japan	25%
Hedging (non-US)	
Unhedged	46%
Hedged to other currency (for consistency)	19%
Investment Type	
Treasury	100%
Non-Treasury Bonds	0%
Equity	0%
Liquidity	
Cash	87%
Govt 1-3	0%
US Duration	0.00
Europe Duration	1.49
Japan Duration	0.00
Transaction Cost – Normal	0.04%
Transaction Cost – Distressed	0.08%
Credit Quality	
Mortgages	0%
ABS	0%
AAA/AA	0%
A	0%
BBB	0%
Downgrade Probability	0.00%
Default Probability	0.00%

The main difference in this table, compared with the base case, is that while the portfolio has a 40% USD currency target, it only invests 35% in USD bonds. A higher USD currency target has the effect of reducing risk *from the USD base currency perspective*. The cash allocation is also quite high, which is again a method for controlling total risk. In terms of portfolio efficiency, the Sharpe ratio for the minimum risk portfolio is 0.57, which is greater than that of the base case (0.54).

Figure 3.3: Efficient frontier relaxing within country duration constraints
(in %)

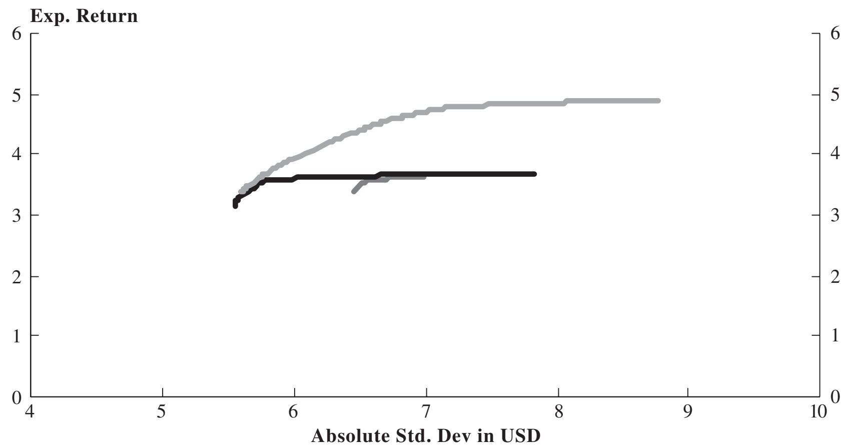

Risk controlled at the total portfolio level.

3.4 Relaxing the country duration constraints

We next relax the within-country duration constraints so that the portfolio can invest in the various maturity sectors across countries and trade off the relative risks contained in country yield curves. We do this by turning off the maximum duration constraints for each country. Figure 3.3 shows a significant additional extension of the efficient frontier.

The efficient frontier in Figure 3.3 controls portfolio risk by focusing on total volatility rather than from constraining individual duration constraints within each country sector. The effect of focusing on *total* portfolio risk significantly extends the efficient frontier relative to the base case, particularly for higher volatility portfolio allocations. This is because the optimisation can take advantage of the various shapes of yield curves as well as the different covariances amongst the maturity sectors. Controlling risk at the total portfolio level delivers opportunities to both increase return and reduce risk.

To illustrate the higher return opportunity, the following portfolio has the same volatility as the base case (6.5%), with the freedom to invest across the yield curve:

Table 3.3: An example of one portfolio on the efficient frontier after relaxing the duration constraints: Risk = 6.5%

Expected Return	4.40%
Volatility	6.50%
Sharpe Ratio	0.68
Portfolio Summary	
Currency	
USD	40%
Euro	35%
JPY	25%
Country	
US	35%
Europe	40%
Japan	25%
Hedging (non-US)	
Unhedged	20%
Hedged to other currency (for consistency)	45%
Investment Type	
Treasury	100%
Non-Treasury Bonds	0%
Equity	0%
Liquidity	
Cash	15%
Govt 1-3	2%
US Duration	5.15
Europe Duration	4.27
Japan Duration	3.70
Transaction Cost – Normal	0.20%
Transaction Cost – Distressed	0.40%
Credit Quality	
Mortgages	0%
ABS	0%
AAA/AA	0%
A	0%
BBB	0%
Downgrade Probability	0.00%
Default Probability	0.00%

The main difference in this example is the duration extension that is registered in each of the three bond markets. The optimal duration is 5.15 years in the US, 4.27 years in Europe and 3.70 years in Japan. Currency weights continue to deviate from the bond allocations. The ability to hold more USD in currency relative to the base case allows the portfolio to take on greater risk in the bond markets. The bottom line is a portfolio which is expected to earn 0.92% more than for the base case for an identical risk level. The Sharpe ratio rises to 0.68, compared with 0.54 for the base case.

3.5 Allowing non-government securities

We recognise that a central bank needs to be cautious when investing in non-government securities. Nevertheless, allowing credit rating agencies to dictate the structure of portfolios is unsatisfactory. Imposing a minimum credit rating, for instance, means that there may be forced selling of securities that have been downgraded. It also reduces the flexibility for managing credit risk *at the total portfolio level*.[17]

We approach the credit risk issue in a more innovative way by constraining the *probability of downgrade* and *probability of default* for the portfolio as a whole. Each asset class is assigned a downgrade and a default probability. Constraining default and downgrade probabilities at the portfolio level allows the investor to control the proportion of the portfolio that is subject to this risk, and trade off the potential benefits in terms of additional return. Technical details are explained in more detail in the Technical Appendix.

The advantage to controlling downgrade and default probabilities is that there may be circumstances when the BBB market offers significant opportunity relative to AA credits. In this case, barbelling the credit composition of the portfolio by holding proportionally more AAA securities to offset the higher downgrade and default probability of the BBB risks will lead to a more efficient portfolio.

In the following example, we control credit risk by the following constraints:

(a) Restricting the probability of downgrade to 1% for the portfolio as a whole.
(b) Restricting the probability of default to 0.10% for the portfolio as a whole.

Default and downgrade probabilities are calibrated from Standard and Poor's data for the ten years to December 2002. The assumptions employed are the following:

	Downgrade	To Junk/Default
Mortgage	0	0
Asset Backed	0	0
Corporate AAA	3.4%	0
Corporate AA	6.86%	0.06%
Corporate A	4.78%	0.33%
BBB	4.27%	4.27%

The table shows that MBS and ABS securities have zero default risk, while the highest downgrade risk applies to AA-rated securities. BBB-rated securities have the highest downgrade risk to junk or default.

[17] Many CBs impose an AA- or A minimum credit rating for constituent securities.

Figure 3.4 shows the effect of relaxing the credit constraint on the portfolio:

Figure 3.4: Efficient frontier allowing additional investment in investment-grade credit securities
(in %)

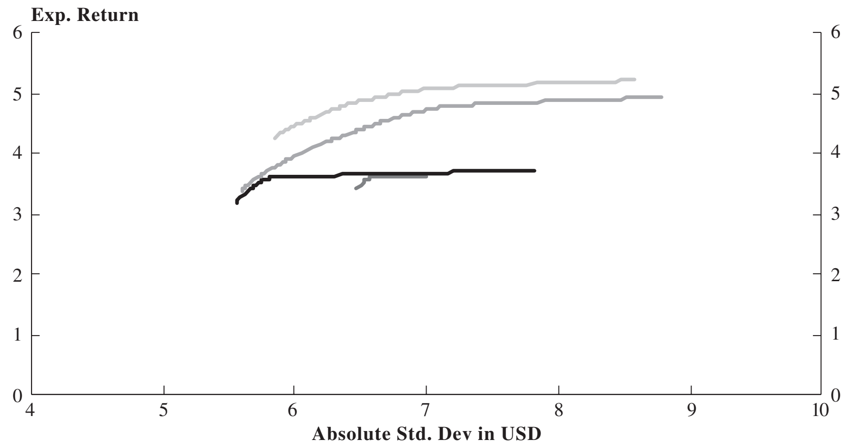

Maximum non-sovereign exposure 25% of portfolio, downgrade risk tolerance 1% and default risk tolerance 0.1%.

Allowing the portfolio to extend into investment-grade credits is achieved by increasing the limit on non-sovereign exposure in the portfolio to 25% and permitting a higher downgrade and default risk tolerance for the portfolio. Figure 3.4 shows that relaxing the credit constraint on the portfolio only has a marginal benefit for the risk and return on the portfolio. This is interesting as the benefit turns out to be less dramatic than the effect of relaxing the currency and duration constraints on the portfolio. This means that a significant increase in portfolio efficiency can be achieved by central banks without having to access non-government segments of the bond markets.

The following portfolio shows the incremental benefit from relaxing the non-government bond constraint compared with the base case. Again, we select a portfolio with similar volatility to the base case.

Table 3.4: An example of one portfolio on the efficient frontier from allowing non-government securities: Risk = 6.5%

Expected Return	4.86%
Volatility	6.50%
Sharpe Ratio	0.75

Portfolio Summary	
Currency	
USD	40%
Euro	35%
JPY	25%
Country	
US	40%
Europe	35%
Japan	25%
Hedging (non-US)	
Unhedged	16%
Hedged to other currency (for consistency)	44%
Investment Type	
Treasury	75%
Non-Treasury Bonds	25%
Equity	0%
Liquidity	
Cash	11%
Govt 1-3	0%
US Duration	5.05
Europe Duration	4.77
Japan Duration	4.44
Transaction Cost – Normal	0.25%
Transaction Cost – Distressed	0.66%
Credit Quality	
Mortgages	0%
ABS	4%
AAA/AA	0%
A	20%
BBB	1%
Downgrade Probability	1.00%
Default Probability	0.10%

The interesting feature of this portfolio is that total overall duration increases relative to the previous example; moreover, the maximum allocation to non-government bonds is optimal. The duration allocation to Japan and Europe increases significantly, whereas the US duration position falls marginally. This means that the proportion of both government and non-government risk increases, because adding non-government bonds to the opportunity set increases the ability to diversify risk. Contrary to popular belief in the central banking community, investing in non-government bonds actually reduces risk for the portfolio as a whole.

In terms of liquidity, the transactions costs associated with normal and distressed markets rise significantly relative to the base case. The distressed market cost of liquidating the portfolio would be 0.66% versus 0.12% for the base case. This reflects the higher cost of trading non-government bonds.

In terms of credit quality, given the assumptions the optimal portfolio invests 20% of the fund in A-rated names, with a 1% allocation to BBB-rated securities counterbalanced by a 4% allocation to AAA-rated asset backed securities. Notably, the optimisation makes full use of the 1% downgrade probability limit and the 0.1% default probability limit in achieving this outcome.

The expected return on the portfolio is 4.86% versus 3.48% for the base case, holding risk constant in each case. This represents a significant pick-up in portfolio efficiency.

3.6 Summary

In summary, the opportunity to increase return from relaxing the liquidity, currency, duration and credit constraints on the portfolio, for similar levels of risk as contained in the base case portfolio, is clear. Figure 3.4 shows that around 1.4% worth of additional return can be earned for similar levels of risk in the portfolio, or that risk can be reduced by around 1.5% while still maintaining the same level of return. This increase in portfolio efficiency is very powerful and can be achieved with minimal changes to the way CBs currently manages their portfolios.

4 Conclusion

This paper introduces a simple mean-variance approach for central bank strategic asset allocation. This framework allows us to compare the portfolio efficiency that is delivered by a "typical" central bank approach to asset allocation that is highly constrained. We progressively relax constraints on liquidity holdings, duration, currency allocation and non-government bond allocations. The paper demonstrates that portfolio efficiency is significantly enhanced without radically changing the portfolio structure.

Central bank portfolios tend to be invested in short-duration physical securities. This means that their portfolios are confined to the short end of the yield curve in each country and that their currency and bond allocations are equal. The source of efficiency gains comes primarily from exploiting the correlations between (a) currency and bond returns, (b) full maturity yield curves across countries, and (c) non-government securities.

A reasonable calibration of the model suggests that relaxing liquidity, currency, duration and credit constraints could achieve up to an additional 140 b.p. worth of return for current risk levels, or reduce risk by around 150 b.p. and still achieve the same expected return.

The results in the paper suggest that quantitative methods can be usefully applied to a central bank's strategic asset allocation decision. Adopting an optimisation framework for determining asset allocation, however, introduces model and estimation risk into the asset allocation process. Recommended portfolios therefore require stress testing and should be scrutinised for their reasonableness before being implemented.

Technical Appendix

Optimisation framework and assumptions

This note describes an asset allocation framework consistent with a CB's joint objective for (a) maximising returns on its reserve holdings for a given level of risk and (b) being prepared to intervene in FX markets.

App A.1 Optimisation framework

The formal model is set out as a single-period optimisation for the sake of simplicity. More detailed developed models can be developed.

$$\underset{\{w(i),w(L)\}}{Max} \; E\sum_{i=1}^{J}\left[R(i)w(i)+R(L)w(L)\right]-\lambda\mathbf{wVw}'$$

$$s.t.$$

$$R(L)w(L) \geq \Gamma$$

$$\sum_{i=1}^{I} w(i,s,t)+w(L,s,t)=1$$

The notation in the equation system is the following. R(i) is the return to asset i, R(L) the return to the liquidity asset, λ a risk aversion coefficient, **w**(s,t) a vector of portfolio weights to simplify notation, **V** a covariance matrix and Γ is the required liquidity.

Additional constraints added to the optimisation are the following:

1. Currency is controlled by assuming that hedged and unhedged returns are separate assets.
2. Country durations are controlled by assuming for each country, c, **w'd** ≤ D_c where **d** is a vector of durations and D_c is the duration limit on each country.
3. Downgrade and default probabilities are controlled by assigning a downgrade and default probability to each asset. We then impose the constraints **w'g** ≤ G and **w'f** ≤ F where the notations **g** and **f** are vectors of downgrade and default probabilities respectively while G and F are the respective upper limits.
4. Transaction cost constraints are controlled by assigning transaction costs to each asset under normal and distressed market conditions and imposing upper limits.

The optimisation is solved using proprietary software based on Markowitz (1999).

App A.2 Calibrating the model

In order to make the optimisation model operational, a number of parameters have to be calibrated. These include the expected return assumptions, the covariance structure, liquidity and other constraints.

The capital market assumptions we employ in our analysis are set out below:[18]

Capital market assumptions

	US		Europe		Japan	
	Expected Return	Percent Volatility	Expected Return	Percent Volatility	Expected Return	Percent Volatility
Govt Cash	2.60%	0.45%	3.70%	0.85%	0.10%	0.80%
Govt 1-3 yr	3.51	1.70	4.06	1.71	0.59	1.81
Govt 3-5 yr	4.30	3.41	4.43	2.75	1.06	3.45
Govt 5-10 yr	5.05	5.21	4.95	4.43	1.67	6.52
Govt 10+ yr	6.21	7.71	6.21	8.03	2.89	9.60
Mortgage	4.69	3.08	NA	NA	NA	NA
Asset Backed	4.83	2.76	NA	NA	NA	NA
AAA/AA	5.78	4.21	NA	NA	NA	NA
A	6.40	4.51	NA	NA	NA	NA
BBB	7.35	4.69	NA	NA	NA	NA
Equity	8.90	14.42	8.90	17.40	6.40	21.40

[18] We do not have assumptions for non-government bonds in either Europe or Japan. This is because the non-government sectors in these markets are nascent with little data to draw firm conclusions. Going forward, we will add these assets to the analysis.

This matrix shows the expected returns and risks associated with various maturity sectors and credit sectors across the major regional fixed income and equity markets – the US, Europe and Japan. A more finely-grained breakdown of risk and return is possible; however, for the purposes of the strategic asset allocation exercise, these assumptions serve as a suitable starting point. The assumptions have been derived by JP Morgan Fleming's Strategic Investment Advisory Group based on a combination of historic data analysis and logical or theoretical overlay. We do not claim to have a monopoly on the correct set of assumptions for the likely future risks and returns in capital markets, so the numbers should be interpreted as a starting point for discussion and analysis.

Transaction costs are relevant to the calculation of the optimal portfolio positions. Transaction cost assumptions for the various sectors are shown below.

Transaction costs
(in %)

	Normal market	Distressed market
Govt Cash	0	0
Govt 1-3 yr	0.025	0.05
Govt 3-5 yr	0.10	0.20
Govt 5-10 yr	0.20	0.40
Govt 10+ yr	0.25	0.50
Mortgage	0.10	0.40
Asset Backed	0.15	0.60
AAA/AA	0.25	1.00
A	0.35	1.40
BBB	0.5	2.0
Equity	0.5	2.0

Finally, default and downgrade probabilities are specified. We use the default and downgrade transition matrices calculated by Standard and Poor's, as follows:

	Downgrade	To Junk/Default
Mortgage	0	0
Asset Backed	0	0
Corporate AAA	3.4%	0
Corporate AA	6.86%	0.06%
Corporate A	4.78%	0.33%
BBB	4.27%	4.27%

References

Aizenman, J. and Marion, N. (2002) "The High Demand for International Reserves in the Far East: What Is Going on?", University of California, Santa Cruz, Working Paper

Carmel, J., Dynkin, L., Hyman, J. and Phelps, B. (2002) "Total Return of Reserves Management", Fixed Income Research Paper, Lehman Brothers

Cochrane, J. (1999) "Portfolio advice for a Multifactor World", Centre for Research into Security Prices, Working Paper No. 491, University of Chicago Business School

International Monetary Fund (2001) "Guidelines for Foreign Exchange Reserves Management", at www.imf.org

Markowitz, H. (1999) *Mean Variance Analysis in Portfolio Choice and Capital Markets*, 2nd edition

Nugee, J. (2000) "Foreign Exchange Reserves Management", *Handbooks in Central Banking No. 19,* Centre for Central Bank Studies, Bank of England

Reserve Bank of Australia Annual Report 2003

5

Reaching for yield: selected issues for reserves managers[1]

Eli M. Remolona and Martijn A. Schrijvers

Abstract

In an environment of historically low yields on highly-rated government securities, reserves managers have found themselves seeking instruments with higher yields in an effort to enhance returns. We focus on three cases for which issues related to higher yields are particularly interesting: a longer-duration portfolio, a corporate bond portfolio, and a portfolio of higher-yielding currencies. In the case of longer durations, we ask whether the present low-yield environment implies a new trade-off between duration and volatility. Concerning corporate bonds, we focus on the challenge of managing a portfolio in which risk is characterised by low probabilities of heavy losses. Finally, we revisit the issue of whether higher yields in certain currencies tend to be offset by movements in exchange rates.

1 Introduction

Managers of official foreign exchange reserves have been facing historically low yields on highly-rated government securities, the instruments to which they have traditionally devoted the bulk of their investment portfolios. In March 2004, the yield on the 2-year US Treasury note still stood at close to 1.5%, down from a peak of nearly 17% in 1981. It is true that much of the decline since 2001 had been the result of cuts in monetary policy rates, which had served to shift down yield curves. Nevertheless, even adjusting for the monetary policy cycle, yields in the major currencies have tended to be substantially lower in recent years compared with those in the previous decade. In these conditions, reserves managers have found themselves seeking instruments with higher yields in an effort to maintain the investment returns to which they had become accustomed.

In considering higher-yielding alternative instruments, reserves managers must ask two basic questions. First, do higher yields actually lead to higher returns?[2] Second, to the extent that higher expected returns are a compensation for taking on greater risk, what is the nature of the risk that is entailed? In this special feature, we focus on a few selected cases for which these questions seem particularly interesting. Three alternative portfolios that offer higher yields are selected, namely a longer-duration portfolio, a corporate bond portfolio and a portfolio of higher-yielding currencies. We discuss the issue of increased risk-taking with respect to durations and corporate bonds. In the case of durations, we ask the specific question of whether the present low-yield environment implies a new trade-off between duration and volatility. In the case of corporate bonds, we focus on the challenge of managing a portfolio in which risk is characterised by low probabilities of heavy losses. We finally examine the question of yield and return with respect to currencies. Specifically, do higher yields offered by instruments in certain currencies tend to be offset by movements in exchange rates?

---

[1] This article was originally published in the September 2003 issue of the *BIS Quarterly Review*. The views expressed in this article are those of the authors and do not necessarily reflect those of the BIS or De Nederlandsche Bank.

[2] Yields differ from returns, because the latter include capital gains or losses, which will depend on duration. For the relationship between yield and return, see footnote 5 below. In the case of foreign currencies, returns may also differ from yields because of exchange rate changes.

In the discussions below, we limit ourselves to issues of strategic investment over the medium to long term. Hence, we conduct our analysis in terms of averages of returns and measures of risk over extended periods of time. This focus allows us to avoid the tactical question of timing, i.e. the issue of when precisely reserves managers should undertake a change in positions. Timing depends on when yields or spreads may be expected to rise or fall, and on this issue we offer no guidance. Our focus on investment strategy also means saying nothing about issues of liquidity. While central banks often hold liquid reserves for intervention purposes, the reach for yield really pertains to the investment part of the portfolio.

2 Duration and volatility: have lower yields changed the trade-off?

The zero lower bound should lead to lower volatility

For default-free debt securities without the possibility of prepayment, risk is represented primarily by duration. A change in the level of interest rates would affect the market value of longer-duration securities more than that of shorter securities. One possible implication of a low-yield environment is a thinner yield cushion against capital losses. If interest rate volatility has remained the same, then a reserves manager who wishes to avoid negative returns would set a shorter duration target. But is it true that volatility is invariant to the level of yields? From a technical standpoint, the zero lower bound on nominal interest rates should naturally lead to lower volatility.[3] From an economic point of view, an environment of low interest rates may simply be an environment of low inflation. As lower levels of inflation tend to be associated with reduced variability of inflation, this may lead to lower interest rate volatility. Low interest rates may also reflect a more transparent monetary policy reaction function, which may also serve to dampen volatility.[4]

Volatilities are lower across the yield curve ...

Indeed, there is evidence that as yields have declined, so have the volatilities of returns.[5] In the left-hand panel of Graph 1, we compare for one-year investment horizons average volatilities between two periods, a high-yield period from January 1984 to December 1993, and a low-yield period from January 1994 to December 2002. As we would expect, the graph shows that in both periods longer duration is associated with higher volatility. More importantly, the graph shows consistently lower volatilities across the duration spectrum during the low-yield period. On average, volatility in recent years is about three-quarters of the average volatility in 1984-93. Assuming this volatility pattern continues to hold, a reserves manager with a given volatility target – or a given Value-at-Risk (VaR) standard – would now be able to extend duration without taking on more risk.

---

[3] This is one reason why models of interest rate movements incorporate the so-called "square-root process", in which volatility is specified to be proportional to the square root of the level of interest rates. In this case, an interest rate close to zero would imply a volatility close to zero. See, for example, Cox et al. (1985) and Gong and Remolona (1997).

[4] Indeed, Ait-Sahalia (1996) provides evidence that such volatility depends on both the monetary regime and, within a regime, on how far the interest rate is from its mean. There is also strong evidence for mean reversion in interest rates within a regime, suggesting that when interest rates are close to the trough in a period of monetary easing, the distribution of interest rate changes is likely to be skewed to reflect the likelihood of a reversal in the policy stance. Moreover, Borio and McCauley (1996) document that bond yield volatility depends asymmetrically on the direction of price changes, whereby rising yields lead to higher volatility.

[5] Note that the concept of volatility relevant to investors is the volatility of *returns*, not the volatility of percentage changes in yields. The relationship of return to yield is well approximated by $r_{t+1} = y_t + D_t\,(y_t - y_{t+1})$, where r_{t+1} is the return at the end of the holding period, y_t and y_{t+1} are the yields at the end and beginning of the holding period respectively and D_t is the duration. The relationship is exact for zero coupon bonds.

Graph 1: Volatility and Sharpe ratios for different durations[1)]

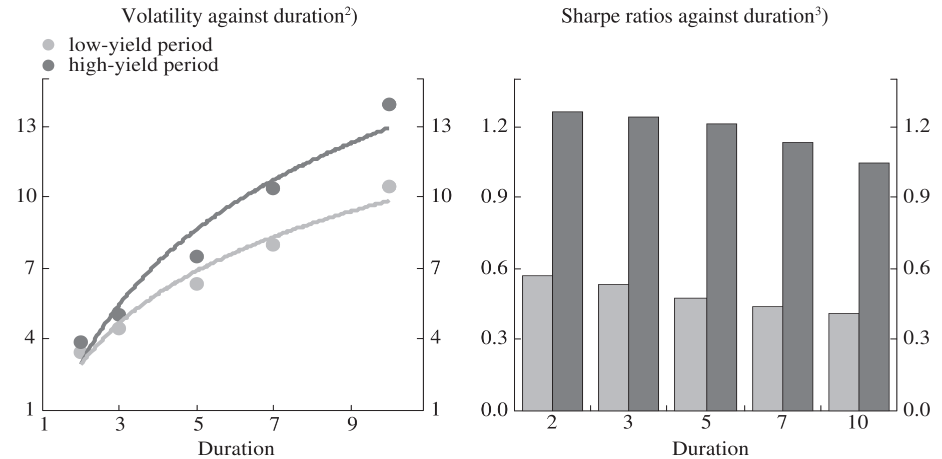

Sources: National data; BIS calculations.
1) The sample is split into two periods. The high-yield regime is defined as the period between January 1984 and December 1993, whereas the low yield period corresponds to the period between January 1994 to December 2002; for the calculation of returns or yields, zero coupon government yields were used.
2) Volatility on the vertical axis; Calculated as the standard deviation of the return, as a percentage.
3) Sharpe ratios on the vertical axis; shorter bars for the low yield period*.

... but Sharpe ratios are better for shorter durations

Another way to decide on duration is to consider the trade-off between risk and return in deviating from a benchmark portfolio. This trade-off may be measured by the Sharpe ratio, which consists of the excess return achieved by deviating from the benchmark, divided by the volatility of this excess return. To illustrate the problem, we consider a benchmark portfolio of 3-month US Treasury securities, and calculate Sharpe ratios for a shift into longer durations. We calculate excess returns by taking the average of realised monthly excess returns resulting from adding different durations to the benchmark portfolio.[6] We consider the addition of 2, 3, 5, 7 and 10-year durations. We undertake the calculations for both the high-yield period of January 1984 to December 1993 and the low-yield period of January 1994 to December 2002. Note that if similar calculations are done with other benchmarks, the Sharpe ratios may change. As shown in the right-hand panel of Graph 1, the calculated Sharpe ratios for the more recent period range from about 0.40 to 0.60, with the shorter durations providing the higher ratios. The computed Sharpe ratios were higher for the earlier period, in which declining yields produced unusually high returns. Nonetheless, it was also true that the trade-off between risk and return tended to be more favourable for portfolios with shorter durations.

---

[6] This is an ex post calculation of excess returns. In theory, the Sharpe ratio is about *expected* excess returns, and the calculation assumes that these returns can be measured by past experience. See, for example, Sharpe (1966).

Hence, while volatility seems to be lower in general, one gets less "bang for the buck" as duration increases. In the above analysis, the desirability of extending duration would depend on whether the reserves manager focuses on meeting a volatility target or on maximising a measure of the trade-off between risk and return. These two decision rules give different answers in the data set investigated here.

3 Credit risk and skewness: the challenge of diversification

Corporate spreads are wide relative to expected losses ...

Another way to increase expected returns is to take on credit risk. Corporate bond spreads tend to be much wider than would be implied by expected losses from default, so corporate bond portfolios do offer a high potential for enhanced returns. For example, as shown in the left-hand panel of Graph 2, the spread between yields on triple-B corporate bonds and US Treasury securities averaged about 203 basis points during 1998-2002. During the same period, the average probability of default for these bonds was about 0.5%, and the average recovery rate given default was 50%. Hence, the spread was more than eight times the expected loss from default as measured by the average loss over five years.[7]

Graph 2: The pricing of default risk

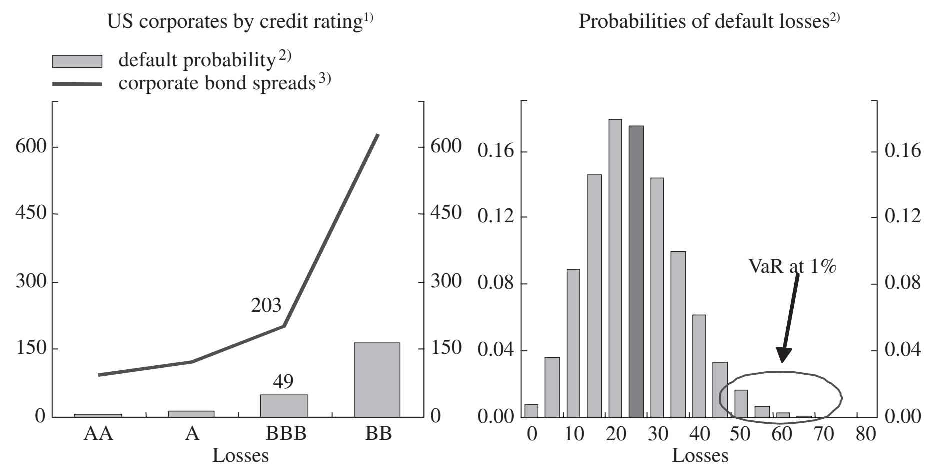

Sources: Bloomberg; Merrill Lynch; Standard & Poor's CreditPro; BIS calculations.
1) Average for the period 1998-2002; in basis points.
2) Probability of default within one year as calculated by Standard & Poor's.
3) Option-adjusted spread for US corporate bonds with five to seven years to maturity; for triple-B, all maturities.
4) Probability density on the vertical axis; losses in units of $1,000 on the horizontal axis; the darker bar represents the expected loss in US dollars for a portfolio of 1,000 US triple-B corporate bonds.

7 Indeed, there are investment strategies that attempt to arbitrage between spreads and expected default losses. The most prominent example of these strategies is the collateralised debt obligation (CDO), in which low-rated bonds are pooled together in a securitisation to create highly-rated securities. Elton et al. (2001) find that a significant portion of the spread can be accounted for by taxes.

... because default risk is hard to diversify
Even 1,000 names may not be enough

Corporate spreads are largely a compensation for bearing credit risk, and one reason why they are so wide is that actual losses from default can easily differ substantially from expected losses. Moreover, the risk of unexpected loss is evidently difficult to diversify away. To illustrate, consider a hypothetical portfolio worth a total of USD 10 million and divided equally among 1,000 different triple-B names.[8] Assume further that these names have identical default probabilities and independent default times (that is, defaults that are uncorrelated).[9] The right-hand panel of Graph 2 shows the probabilities of varying amounts of default losses for this portfolio given the triple-B default probability of 0.5% and recovery rate of 50%; the dark bar indicates an expected loss from default of $25,000. However, as the graph also shows, the probabilities of greater losses are significant. For example, 1% VaR represents a 1% probability that losses would exceed $50,000. As corporate bond portfolios go, one with 1,000 names is already unusually large, and yet our example shows that it could still be poorly diversified in that unexpected losses remain significant. By contrast, in the equity market, a portfolio with 50 different stocks can often be considered well-diversified.[10]

Correlations are higher for lower-rated names

It is important to understand the role a correlation in defaults would play in the risk of a corporate bond portfolio. Such a correlation would naturally limit the scope for diversification. In the extreme, a portfolio with 1,000 names but with 100% default correlation would have the risk profile of a portfolio with a single name. In practice, it is difficult to estimate default correlations with any precision. Market participants often assume that such correlations are significant for firms in the same industry, while for firms in different industries. correlations are small.[11] Correlations are also likely to be higher between low-rated names than between highly-rated names.[12] Such correlations are also likely to vary over time, increasing in precisely those periods when the benefits of diversification are most sought after. To estimate such correlations more accurately, some market participants rely on models that attempt to derive these correlations from the degree to which sharp downward movements in equity prices coincide between firms.

However, while such correlations limit the scope for diversification, they are not what makes corporate bond portfolios difficult to diversify. After all, equity returns tend to be much more highly correlated than default risk. And yet, as mentioned above, a small equity portfolio can be well diversified in that the idiosyncratic risk of individual stock returns is negligible, while a large corporate bond portfolio is likely to remain poorly diversified in that unexpected losses from default are significant.

---

[8] To keep things simple, we only account for the probability of default. In practice, losses can also arise from downgrades and wider spreads. Indeed, it is important to integrate credit and market risk into risk management. Duffie and Singleton (2003), for example, show how this might be done.

[9] We discuss the role of correlations below.

[10] Since the 1970s, the required number of stocks for diversification seems to have risen from about 20 to as many as 50, because idiosyncratic risk has increased. See, for example, Campbell et al. (2001).

[11] For example, in evaluating CDOs, Moody's assigns so-called diversity scores to the pool of collateral. These scores reflect the default correlations the rating agency sees, and the scores tend to differentiate mainly between correlations within an industry and correlations between firms in different industries.

[12] Zhou (1997) and Gersbach and Lipponer (2003), for example, show that credit losses are more highly correlated for debt with higher probabilities of default. This means that as credit quality declines over the cycle, default correlations would also rise.

Corporate bond returns are negatively skewed

The essential characteristic of credit risk that makes diversification so difficult is the asymmetry in the distribution of returns that this risk generates. In particular, the return distribution for a corporate bond portfolio is characterised by a rather long tail on the left, representing low probabilities of heavy losses from defaults or rating downgrades. In other words, the distribution is negatively skewed. By contrast, equity returns tend to show a much more symmetric distribution, in which the probabilities of large losses tend to be matched by the probabilities of large gains. It is the skewness in returns that presents the reserves manager with the challenge of diversifying a corporate bond portfolio.

4 Instruments in other currencies: do higher yields mean higher returns?

Do exchange rates move to offset yield differentials?

At present, most central banks manage their reserves by fixing their currency allocations, with a substantial portion devoted to US dollar-denominated highly-rated fixed income assets. Until recently, these assets have offered rather low yields. Can we gain by deviating from these currency allocations to tilt towards assets in currencies with higher yields? The hypothesis of uncovered interest rate parity suggests that on average there should be no gain: currencies with higher yields are likely to depreciate to an extent that the loss from the exchange rate offsets the gain from the yield differential. In its strict form – where the maturity of the instruments matches the investment horizon – the hypothesis is empirically found not to hold.[13] However, reserves are often placed in securities with maturities that exceed the investment horizon, and to our knowledge the uncovered interest rate hypothesis has not been tested for this case.

We look at bond returns instead of short-term yields

Do higher yields lead to higher returns once exchange rate movements are taken into account, particularly for longer-maturity instruments? For present purposes, we compare returns on government bonds denominated in euro (Deutsche Marks for the pre-euro period), pound sterling, Japanese yen and US dollars. We examine yields and returns for the period January 1994 to December 2002, calculating returns in US dollars. We fix the investment horizon at one year while comparing returns for securities with a five-year duration. If the uncovered interest rate parity hypothesis holds, yield differentials should have no effect on differential returns, because differences in yield should be offset by changes in the exchange rate.

The results are striking for the sample period considered. For yield differentials between the euro and dollar and between the pound and dollar, not only do we reject our version of the uncovered interest rate parity Hypothesis, but we also find that return differentials exceeded yield differentials by large amounts. As shown in Graph 3, a 10 basis point yield differential between euro and dollar bonds meant a 62 basis point differential in returns, while the same yield differential between sterling and dollar bonds led to a 32 basis point differential in

[13] The body of evidence against uncovered interest rate parity is quite large. One of the most careful tests is provided by Hansen and Hodrick (1980). More recent investigations of this issue include Flood and Rose (1999) and Brooks et al. (2001). The literature thus far relies on tests using maturities that match the holding periods, for example a one-year instrument for an investment horizon of one year.

Graph 3: Return differentials against yield differentials for 5-year bonds[1)]
(January 1994-December 2002; in annual percentage rates)

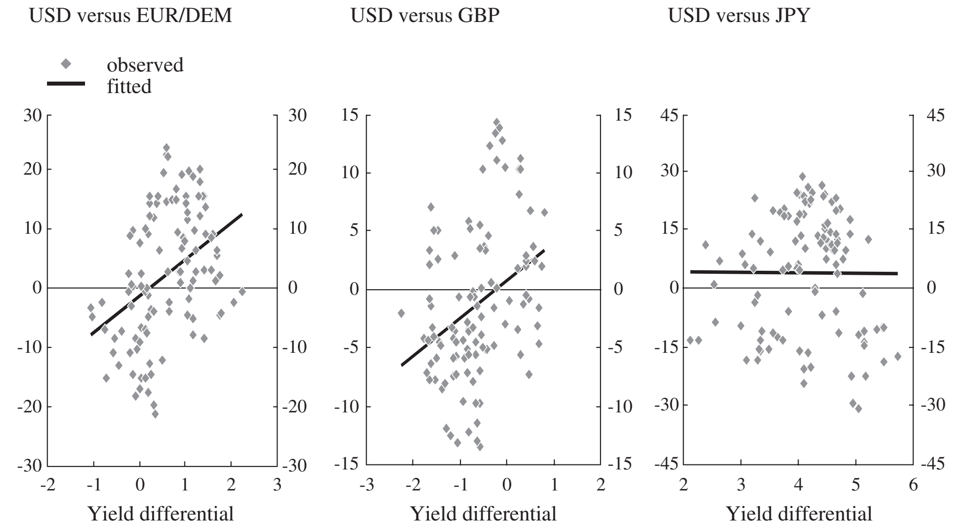

Sources: National data; BIS calculations.
[1] The return differential (shown on the vertical axis) is defined as the return on the US Treasury security minus that on the other currency government bond in US dollar terms. The yield differential (horizontal axis) is the difference between the nominal yields. The fitted line is based on an OLS estimate of $r^{US}_{t+1} - r^{J}_{t+1} = \alpha + \beta(y^{US}_t - y^{J}_t)$, where $r^{US}_{t+1} - r^{J}_{t+1}$ is the return differential between the US dollar and the other currency and $y^{US}_t - y^{J}_t$ is the corresponding yield differential.

returns. It happens that during this period the higher-yielding currency also tended to be the appreciating currency. Hence, exchange rate movements served to magnify the effect of yield differentials on returns.[14] Note, however, that this phenomenon did not extend to yield differentials between yen and dollar bonds. In this case, the outcome was roughly consistent with the hypothesis: exchange rate movements tended to just offset the yield differentials.

Higher-yielding currencies offer scope for enhancing returns

Our results suggest only that there may be some scope for enhancing returns by considering higher-yielding currencies. On one hand, yield differentials are generally not offset, and indeed may often be reinforced, by currency movements. On the other hand, the relationship does not seem to be reliable for all currencies and may not hold for all periods.

[14] For this sample period, conducting the test using 1-year government bonds, so that the maturity matches the investment horizon, leads to qualitatively similar but weaker results. The tendency of higher-yielding currencies to appreciate seems to be more strongly associated with long-term yields than with short-term ones.

Exchange rates can add volatility to bond returns

Another issue to consider in deviating from one's currency allocation is the benefit of diversification in reducing risk. As is well known, a low correlation between returns on different assets in a portfolio can reduce the volatility of returns of the overall portfolio. To what extent is this gain from diversification present in returns across currencies? In general, for the major currencies, fluctuations in exchange rates contribute more to the volatility of bond returns than movements in interest rates do. For example, over the 1994-2002 sample period, the volatility of returns in US dollar terms on a 2-year German government bond was two and a half times the volatility for a 2-year Treasury note. Although the correlation between returns is low between German government bonds and US Treasuries, the gain from diversification is limited by the fact that the return volatilities are so far apart. Note, however, that if the reserves manager calculates returns in local currency, there may be more scope for diversification, as here the difference in volatilities across foreign currencies would not be so pronounced.

5 Conclusion

The alternatives available to reserves managers who are seeking higher yields include extending their duration benchmark, investing in corporate bonds and shifting towards instruments in higher-yielding currencies. For each of these alternatives, we raise specific issues about either risk or return. In none of these cases do we try to resolve the issue. Rather, the intention here is limited to providing analyses that will allow a reserves manager to pose important questions in more focused ways.

For the alternative of extending the duration benchmark, we find that the critical risks have changed in a way that seems favourable to the reserves manager. In particular, we find that as yields on highly-rated government securities have declined, so have the relevant return volatilities for any given duration. This means that an unchanged VaR standard would allow the reserves manager to take advantage of the higher yields offered by longer durations. At the same time, however, the trade-off between risk and return also seems to have changed in a way that may not favour longer durations. One particular measure of this trade-off, the Sharpe ratio, seems to recommend durations not longer than two years. The question then becomes the appropriate standard for judging risk and return.

In the case of corporate bonds, we argue that the main challenge is one of diversification in the face of skewness in returns. Such skewness – representing the risk of small probabilities of large losses – makes corporate bond portfolios rather difficult to diversify. The good news is that this difficulty is reflected in corporate spreads that are much wider than would be implied by expected losses from default.

Finally, in the case of currency allocations, we find that, over a long sample period, exchange rates on average move in favour of the higher-yielding currencies, thus resulting in return differentials that magnify the yield differentials. Our analysis applies to the common case in which the instruments considered have longer maturities than the investment horizons. We find results that are stronger than the usual rejections of the hypothesis of uncovered interest rate parity, in which maturities and investment horizons are kept equal. Given our findings, the open question becomes the reliability of these results for a given currency pair and their robustness for different currency pairs.

References

Ait-Sahalia, Y. (1996): "Testing Continuous-time Models of the Spot Interest Rate", *Review of Financial Studies*, Vol. 9, pp. 385-426.

Borio, C. E. V. and R. N. McCauley (1996): "The Economics of Recent Bond Yield Volatility", *BIS Economic Papers*, no 45, July.

Brooks, R. J., H. Edison, M. S. Kumar and T. M. Slok (2001): "Exchange Rates and Capital Flows", *IMF Working Papers*, 01/190.

Campbell, J. Y., M. Lettau, B. G. Malkiel and Y. Xu (2001): "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk", *Journal of Financ*e, Vol. 56, pp. 1-143.

Cox, J., J. Ingersoll and S. Ross (1985): "A Theory of the Term Structure of Interest Rates", *Econometrica*, Vol. 53, pp. 385-407.

Duffie, D. and K. J. Singleton (2003): *Credit Risk: Pricing, Measurement and Management*, Princeton and Oxford: Princeton University Press, pp. 319-24.

Elton, E. J., M. J. Gruber, D. Agarwal and C. Mann (2001): "Explaining the Rate Spread on Corporate Bonds", *Journal of Finance*, Vol. 56, pp. 247-77.

Flood, R. and A. Rose (1999): "Understanding Exchange Rates without the Contrivance of Macroeconomics", *Economic Journal,* Vol. 109, pp. 660-72.

Gong, F. and E. M. Remolona (1997): "Two Factors along the Yield Curve", *Papers in Money, Macroeconomics and Finance, supplement to the Manchester school*, Vol. LXV, pp. 1-31.

Gersbach, H. and A. Lipponer (2003): "Firm Defaults and the Correlation Effect", *European Financial Management*, Vol. 9, pp. 361-77.

Hansen, L. P. and R. J. Hodrick (1980): "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis", *Journal of Political Economy*, Vol. 88, pp. 829-53.

Sharpe, W. F. (1966): "Mutual Fund Performance", *Journal of Business*, Vol. 39, pp. 119-38.

Zhou, C. (1997): "Default Correlation: An Analytical Result", FEDS paper 1997–27, Federal Reserve Board, May.

6

The risk of diversification

Peter Ferket and Machiel Zwanenburg

Abstract

Efficient management of foreign reserves is an important task of a central bank. Risk management contributes to this task by managing and controlling the exposure to financial risks. Its importance has increased over the past decade, as the management of foreign currency reserves has changed focus from a principal preservation objective to one of generating a total return. This paper addresses one important task of risk management: that of proposing a strategic asset allocation.

We present return and risk characteristics of asset classes which are of particular interest: long-term government bonds, global government bonds, investment-grade credits, high-yield bonds, and equities. The paper pays special attention to the international dimension of bond investing, addressing the potency of hedging currency risk and of active duration management. We provide a theoretical discussion of the risks involved as well as empirical results. Starting from the safest investment class, any diversification to another asset class introduces extra risk. This paper therefore addresses the risk of diversification.

1 Introduction

Efficient management of foreign reserves is an important task of a central bank. Risk management contributes to this task by managing and controlling the exposure to financial risks. Typically, risk management is involved in proposing a strategic asset allocation over different asset classes and establishing a framework for the day-to-day risk management of reserves. This paper addresses the first task of risk management, namely strategic asset allocation. Over the past decade, the management of foreign currency reserves has changed its focus from an objective of maintaining liquidity and principal preservation to one of generating a total return. The quantification of the trade-off between risk and expected return is one of the important problems of modern financial economics. For equity portfolios, the investor's portfolio selection problem was embedded in terms of expected return and variance of return by pioneers[1] such as Markowitz (1959), Sharpe (1964) and Ross (1976). In the private asset management industry, total return maximisation and risk diversification are always the main objectives of investments. In this paper we discuss a number of asset classes that are popular in the private asset management industry: long-term government bonds, global government bonds with both passive and active management variants, investment-grade credits, high-yield bonds, and equities. Besides a theoretical discussion of the risks involved in investing in these asset classes, we also present empirical risk and return results. We compare the return and risk characteristics of these asset classes with the characteristics of a cash benchmark that can be seen as the lowest risk investment alternative. As central banks must consider many factors in the currency allocation of their foreign reserves, such as the balance of trade and liability matching, this decision is beyond the scope of this paper. In our

[1] For a historic overview of quantitative models as applied to finance, see Campbell, Lo and MacKinlay (1997). For a less mathematical introduction to the concepts of risk and return, see Brealey and Meyers (1991).

presentation framework we restrict ourselves to two foreign reserves currencies: the US dollar and the euro. Looking for the total return enhancement of the foreign reserves portfolio, we leave the safe area of cash investments and we proceed in four dimensions. As we start from the safest investment class, any diversification into another asset class will bring extra risk into the portfolio[2]. This paper therefore addresses the risk of diversification. The risks taken into account are market, credit and liquidity risk.

The remainder of the paper is organised as follows. In the next section, we describe the data series and set-up used in our empirical analysis. The first dimension of extension that we consider is that of maturity. In the case of cash investments, we know the returns at any point in time. By investing in government bonds, investors are able to profit from higher yields owing to (normally) upward sloping yield curves. When investing in government bonds, we still know that our investments are safe, i.e. there is no default risk involved. But the return on our bond portfolio is no longer known in advance, i.e. we are exposed to market risk. We compare return and risk characteristics of investments in short-term bonds (term to maturity between one and three years) and all longer-term bonds (term to maturity greater than one year) with the cash benchmark in Chapter 3. We deliberately use these overlapping maturity bands in this paper because a maturity restriction for the investment universe (e.g. three years time-to-maturity) is often used in the private asset management industry. In the second dimension, we consider the diversification effects of international bond investment. Instead of investing only in the local government bond market of the foreign currency concerned, we consider a global investment approach in Chapter 4. The question of whether or not to hedge the currency risk is the most important topic that is discussed. Moving towards a global bond portfolio offers the possibility of introducing the third dimension: active duration management as an additional component to the passive global bond portfolio. This topic is discussed in Chapter 5. In Chapter 6 we step beyond government bonds and look at investment-grade credits, high-yield bonds, and equities. This paper does not address alternative investments such as emerging markets, private equity, and hedge funds. We end with a conclusion in Chapter 7.

2 Data and empirical set-up

2.1 Data

In this sub-section we successively treat cash, short-term and long-term government bonds, investment-grade credits, high yield and equity.

Cash

As a return on the cash benchmark we use the 3-month JP Morgan Cash Indices. These cash indices measure the total return performance of constant maturity euro-currency deposits. The indices serve as reliable performance benchmarks for a fund's cash component. Analyses in this paper are performed both for a US-based investor as well as for an EUR-based investor. We thus use both the USD and EUR cash indices.

[2] This holds for all risk measures satisfying the risk-free condition introduced by Artzner, Delbaen, Eber and Heath (1997).

Government bonds

In our analysis we will distinguish between short-term bonds and the government bond market as a whole. For the USD reserves we define the short-term bond market as all US Treasury bonds with a term to maturity between one and three years; the entire government bond market is defined by all US Treasuries with a term to maturity of longer than one year. We choose the JP Morgan US 1-3 year Government Bond Index to represent the short-term US Treasury market, and the JP Morgan US Government Bond Index to represent the entire US government bond market. We basically repeat this for EUR reserves, with the only difference that we now consider bonds denominated in EUR[3] that are issued by governments of countries participating in the euro. As a proxy for the global bond market, we take the JP Morgan Government Bond Index (GBI), which is a market-weighted index of the return of government bonds in the following countries: Australia, Belgium, Canada, Denmark, France, Germany, Italy, Japan, the Netherlands, Spain, Sweden, the United Kingdom and the United States.

The data source for the exchange rates, used in the calculation of hedged returns, is Reuters (6 o'clock GMT fixings).

Credits/High Yield/Equity

As a proxy for investment-grade credits, we use the Lehman Brothers US Credit Index. The credit index includes investment-grade bonds issued by corporations and non-corporate entities.

We use the Lehman US High-Yield Corporate Index as our data source for the high-yield market. The corporate index covers dollar-denominated, fixed-rate, below-investment-grade debt. For equities, we use the MSCI US Index.

Data period

All analyses are performed using monthly returns over the data period December 1987 through December 2002. We choose this data period out of convenience: JP Morgan publishes its returns starting December 1987. We believe that the period chosen is long enough to draw conclusions about the risk/return profile of different investment strategies. It includes periods of both rising and falling bond yields. We are aware that the results of a risk/return analysis of asset classes are strongly affected by the data period chosen. Typically, the return dimension is more sensitive to the choice of data period than the risk dimension. A historical period can be found for every asset class that shows superior return characteristics compared with any other asset class. It is explicitly not our goal to promote any specific asset class based on the empirical results. Our aim is to illustrate the changes in portfolio risk when diversifying foreign reserves investments over a number of asset classes that offer investors the possibility of return enhancement.

2.2 Empirical set-up

The risk and return characteristics of the different asset classes are compared based on different measures. There are a variety of risk measures that can be used to present the risk/return profile of an asset mix. The mean-variance approach, based on volatility or variance as risk measure, is the earliest method designed to solve the portfolio selection problem (Markowitz 1952, 1959). The principle of diversification is the foundation of this method,

[3] The euro was introduced in 1999. Since its introduction the behaviour of the different EMU bond markets has been quite similar, but this was not the case before 1999. Therefore, we use the German bond market as a proxy for the EMU bond market until December 1998.

and is still has wide application in risk management. However, there are some arguments against this approach. Controlling (minimising) the variance not only leads to low deviation from the expected return on the downside, but also on the upside. It may also bound the possible gains. Another disadvantage is that the use of variance as a risk measure is only advisable if returns are normally distributed. Although this might be true for returns on Treasuries, this is not the case for the asset categories that we consider in Section 6[4]. Empirically, the normal distribution roughly approximates the distribution of many variables.

In recent years, Value at Risk (VaR) has become a new benchmark for managing and controlling risk (Jorion 2003, Jorion 1997, RiskMetrics 1995). This risk measure overcomes the shortcomings of variance as a risk measure. VaR can be calculated using a multitude of methodologies, which include Monte Carlo simulation, historical simulation, and parametric approaches using BARRA multifactor models, RiskMetrics, or user-defined risk models. The methodology used in this article is presented in the box below.

Other risk and return characteristics used in this article are described below:

Risk/return characteristics

- *Average return:*
  Annualised average return
- *Volatility:*
  Annualised volatility of the returns
- *Sharpe ratio:*
  The Sharpe ratio (also known as the Reward-to-Volatility ratio) indicates the excess return (over cash) per unit of risk associated with the excess return. The higher the Sharpe ratio, the better the performance. Graphically, the Sharpe ratio is the slope of a line between the risk-free rate and the portfolio in the mean/volatility space. Actually, the task of finding the efficient portfolio in the Markowitz mean-variance framework with a risk-free asset is equal to maximising the Sharpe ratio of the portfolio. The Sharpe ratio does not refer to the market portfolio or any other benchmark. The implicit benchmark is the risk-free rate of return. This makes this measure particularly useful in our case. The excess return can be interpreted as a zero-investment strategy: it can be obtained by taking a long position in the portfolio and a short position in the risk-free rate, with the funds from the latter used to finance the purchase of the former.
  Note that the drawbacks of volatility as a risk measure, as described earlier in this section, also apply to the Sharpe ratio.
- *Probability of negative return*
  This gives the probability of a negative return over a one-year investment period. The probability provides the percentage of starting months for investments over the past 15 years that resulted in a negative return over the next period of one year.
- *Probability of outperformance*
  This gives the probability of outperforming cash over one-year and three-year investment periods. This probability is based on (overlapping) 12-month and 36-month returns respectively over the historical data period. The probability provides the percentage of starting months for investments over the past 15 years that resulted in a better return than cash over the next period of one or three years.
- *Value at Risk (VaR)*
  VaR summarises the expected maximum loss (or worst loss) over a target horizon at a given confidence level at a given moment in time. We present the VaR for a time horizon of one year and a confidence level of 95%[1]. Just as with the probability of negative return and the probability of outperformance, the VaR is based on overlapping yearly returns. It gives us the annual return that we would have surmounted in 95% of the starting months for investments over the past 15 years.

[1] All VaR results presented in this paper are calculated using a 95% confidence level. Although the absolute size of the results changes using a more conservative confidence level (e.g. 97.5%, 99%), conclusions based on VaR figures do not change.

[4] For example: the excess kurtosis of the monthly US high yield returns is 4.7, against 0 for a normally distributed return series. Using volatility as a risk measure would provide an incorrect picture of risk taken.

The risk characteristics that we present are fully based on the historical data. We deliberately do not present return and risk predictions for the asset classes for the coming years as we want to focus on the strategic asset allocation decision.

The different risk measures either present the risk over a one-year or a three-year horizon. The risk horizon chosen is subjective by definition. We choose a one-year horizon as this is standard practice in the private asset management business. Although central banks are long-lived agents, performance (and thus also risk) will often be evaluated on a shorter term (e.g. a one-year horizon). This is only logical as the long-term view can change into a short-term view quickly owing to unforeseen events. We forgo risk measures for a longer horizon as we believe that a historical data period of only 15 years does not suffice for such analyses.

Note that we do not take transaction costs into account. The investment strategy applied in this article is a buy-and-hold strategy, incurring only limited transaction costs. As a consequence we also forgo transaction costs on the currency forwards, used in Chapter 4 for hedging purposes. The currency market is one of the most liquid financial markets available, resulting in very low transaction costs. Taking transaction costs into account will only marginally alter the results[5].

Transaction costs are however included in the analyses presented in Chapter 5, as the active duration management strategy described has the potential to create transactions on a monthly basis.

In the next sections we show risk/return pictures for different combinations of asset classes. The vertical dimension of the pictures shows the average return of asset classes on a yearly basis, while the horizontal dimension (the risk dimension) shows the 95% empirical VaR as described above. We always show results for both the USD and the EUR reserves portfolio perspectives.

3 The maturity dimension

3.1 The risk of government bond investing

Putting reserves in a foreign deposit is often judged as the best guarantee of preserving the nominal value of your capital.[6] However, owing to the credit risk of the counterparty that is involved, this is not exactly the case. An alternative with less credit risk is investing in Treasury bills. Rates on deposits are normally higher than T-bill rates as the former are subject to more credit risk. As we do not want to make things easy for ourselves, we choose the higher-yielding deposits as our cash benchmark.

A straightforward way to enhance return without taking on too much risk is to invest foreign reserves in the local government bond market. For USD reserves this means investing in US Treasuries. For EUR reserves, the local government bond market contains all EUR-denominated bonds issued by governments of countries participating in the euro. By far the largest risk faced by an investor in government bonds is *market risk* or *interest rate risk*. As interest rates rise, the price of the bond will fall; conversely, as interest rates fall, the price of the bond will rise. For buy-and-hold investors who plan to hold the bond to maturity, the change in the bond's price is of no concern. However, for investors who may have to sell the bond before the maturity date (such as managers of foreign reserves), an increase in yields after the bond is purchased results in the realisation of a capital loss. In government bond

[5] The return of the different asset categories and investment strategies will diminish owing to the presence of transaction costs. The risk dimension is not affected.

[6] We do not consider inflation risk that will impact the real value of the capital in the future.

investment practice, the following first-order approximation for the relative price change of an (option-free) bond is often used:

$$\frac{\Delta P}{P} = -MD \times \Delta y, \tag{1}$$

where P is the price of the bond, MD is the modified duration of the bond and Δy the absolute yield change (see for example Fabozzi, Pitts and Dattatreya 1997, and Jorion 2003). Expression (1) tells us that the market risk, or interest rate risk, of the government bond depends on interest rate volatility and on the modified duration of the bond. The first factor is market-specific, while the second is bond-specific. The higher the modified duration, the greater the sensitivity of the bond's price to yield changes. So by using (modified) duration, investors can choose the amount of interest rate risk they wish to allocate to their portfolio.

In general, the extra risk of a long maturity bond versus a short-dated bond is compensated by a higher return. If we forgo interest rate changes, the return on a bond is equal to its yield. The most common relationship is a yield curve in which the longer the maturity, the higher the yield. That is, investors are rewarded for holding longer maturity Treasuries in the form of a higher potential yield. This risk premium is usually referred to as the "term premium". This shape is referred to as a normal or positively sloped yield curve.

The theory on market risk described above can be straightforwardly transferred to the case of a portfolio of government bonds. The modified duration of a government bond portfolio measures the sensitivity of the value of the portfolio to interest rate changes (in first order approximation). The portfolio's modified duration can easily be approximated by the market-weighted sum of the modified duration of the bonds in the portfolio. This approximation of market risk using the portfolio modified duration is only accurate for *parallel yield changes*. In practice, however, this is never the case. For example, changes in 2-year yields are usually different from changes in 10-year yields.

The closer together maturities are in terms of years, the better the approximation of parallel yield shifts becomes. For this reason it is common practice for bond investors to divide the bond market and bond portfolios into maturity buckets. An example of such a division is provided below:

1-3 years	3-5 years	5-7 years	7-10 years	> 10 years

The 1-3 year maturity segment contains all bonds (in the market or in the portfolio) with a time to maturity between one and three years. Using such a division, one can calculate the *partial duration* of a government bond portfolio for each maturity bucket. The portfolio's partial duration for a maturity bucket is calculated as the market-weighted sum of the modified duration of the bonds in the portfolio that fall into that specific maturity bucket.

The approach described above for the measurement of market risk is only one of many available approaches. It is very similar to the widely-used Key Rate Duration approach propagated by RiskMetrics (1995).[7]

Although market risk or interest rate risk is by far the most significant risk, it is not the only risk that government bond investors face. *Credit risk*, or *default risk*, refers to the risk that the issuer of a fixed-income security may not be able to make timely principal and coupon payments on the security – a situation known as a default. US Treasury securities are considered to be free of credit risk. The same holds for most EMU governments such as Germany and the

[7] See Jorion (1997) for an overview of the available approaches.

Figure 1: Risk (95% VaR)/return characteristics of the short-term

a) USD bond market and the USD bond market as a whole, Dec. 1987-Dec. 2002 (in %)

b) EUR bond market and the EUR bond market as a whole, Dec. 1987-Dec. 2002 (in %)

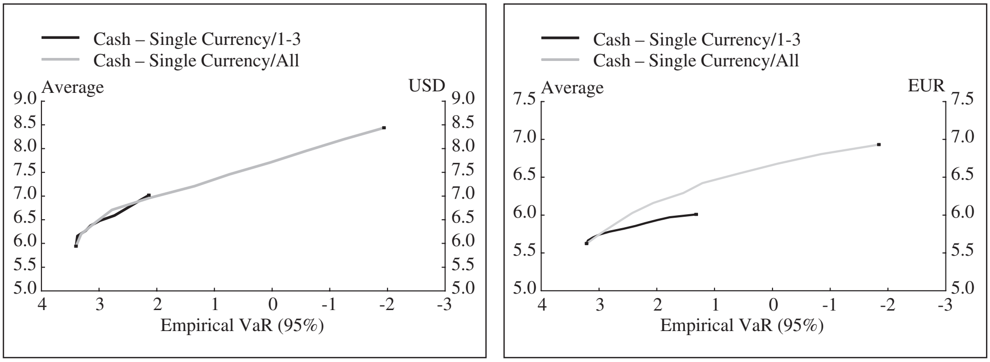

Netherlands. Some other EMU countries, such as Italy and Portugal, have a lower credit rating and hence are not considered to be free of default risk. The credit rating of these countries is sufficiently high to ensure that credit risk is much less important than market risk.

Especially for foreign reserves management, *liquidity risk* or *marketability risk* is an important issue. This risk involves the ease with which an issue can be sold at or near its true value, and is primarily measured by the spread between the bid and offer prices quoted by dealers. The US and EMU government bond markets are two of the largest and most liquid financial markets in the world. Bid/ask spreads are typically in the range of basis points and large volumes of government bonds can be traded within minutes. Therefore, liquidity risk is not a major source of risk for US and EMU government bond investments.

3.2 Empirical results

The risk/return profiles of the short-term and overall bond market for the USD and EUR markets are shown in Figure 1a and Figure 1b respectively. The ends of the lines represent the risk/return characteristics of a 100% investment in either the short-term bond market or the overall bond market. The connecting line[8] depicts the risk and return for combinations of investments in cash and in government bonds.

Both figures show that risk and return go hand in hand when investing in government bonds. Short-term bonds outperform cash in the long run, at the cost of an increase in the risk profile. The overall bond market outperforms the short-term bond market, but again at the cost of extra risk. This extra risk is the consequence of extra interest rate sensitivity. The duration of the short-term bond market is typically only 30% to 40% of the duration of the overall bond market. Therefore, the return volatility of a 100% allocation to the short-term bond market is similar to the return volatility of a 30-40% allocation to the overall bond market (assuming identical yield volatilities for both markets). Table 1 summarises the results for the different markets.

The positive Sharpe ratios indicate that an investor is compensated for the extra risk of bond markets over cash. It is up to each individual investor to decide if this compensation is

---

[8] The connected line is less smooth compared with standard mean-variance analyses because we plot *empirical* Value at Risk instead of volatility on the horizontal axis.

Table 1: Risk/return profile of single currency bond markets

Market	Cash	Single currency	Single currency
Maturity segment		1-3	all
Base currency USD			
Average return	5.9%	7.0%	8.4%
Volatility	0.6%	1.8%	4.4%
Sharpe ratio	n/a	0.61	0.55
VaR (95%)	3.4%	2.1%	-1.9%
Probability of negative return	0%	0%	8%
Probability of outperf. (1 yr)	n/a	66%	69%
Probability of outperf. (3 yrs)	n/a	79%	90%
Base currency EUR			
Average return	5.6%	6.0%	6.9%
Volatility	0.7%	1.8%	3.6%
Sharpe ratio	n/a	0.22	0.35
VaR (95%)	3.2%	1.3%	-1.8%
Probability of negative return	0%	1%	14%
Probability of outperf. (1 yr)	n/a	60%	66%
Probability of outperf. (3 yrs)	n/a	79%	77%

adequate. Although the risk of government bond investing in terms of volatility is more or less comparable for both USD and EUR scenarios, the probability of negative return differs. The probabilities in the EUR scenario are higher owing to the lower average returns.

The probability of outperformance depends on the investment horizon. In general the probability will increase with the length of the investment horizon, as negative excess returns are compensated by positive returns in the long run. This is illustrated by the figures in Table 1.

Figure 2a and Figure 2b confirm the added value of bonds over cash, showing the total return indices for the different investment possibilities (base value 100 in December 1987). The cumulative return depends in a positive way on the duration of the underlying position – as of course does the fluctuation of the return.

Figure 2: Total return indices for the

a) USD bond markets,
Dec.1987-Dec. 2002

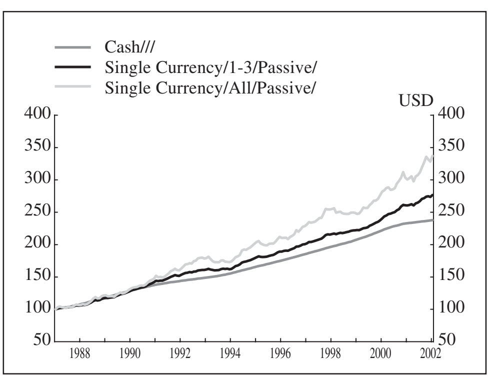

b) EUR bond markets,
Dec. 1987-Dec. 2002

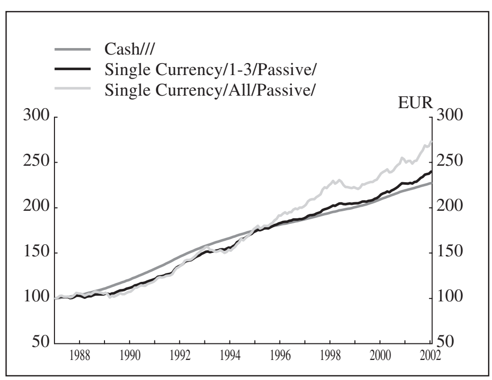

Figure 3: Rolling 12M returns for the

a) USD bond markets, Dec. 1987-Dec. 2002 (in %)

b) EUR bond markets, Dec. 1987-Dec. 2002 (in %)

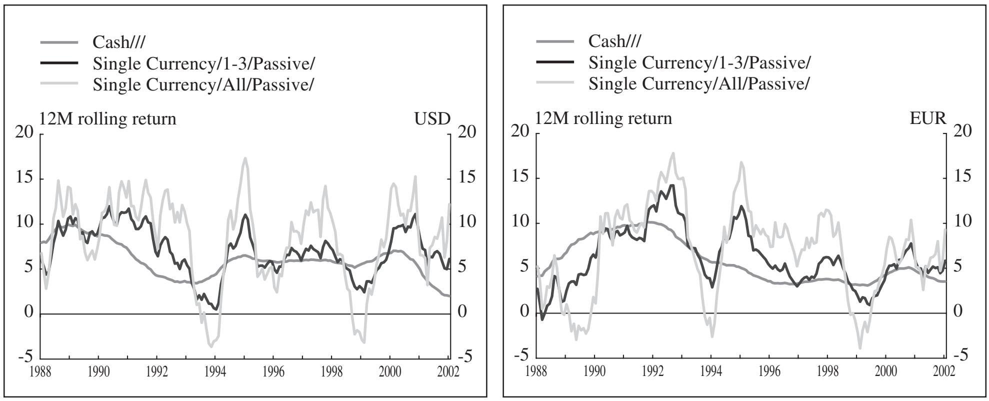

The extra risk of the bond positions becomes evident when we study the moving 12-month returns, shown in Figure 3a and Figure 3b. Moving return figures show an interesting picture of risk to investors as they indicate whether there have been consecutive months in which the average return falls below a specific level. For foreign reserves managers it may be interesting to see whether the rolling 12-month average return falls below the zero percent level (as an indication of capital preservation).

We can pinpoint periods with negative returns for both the USD and EUR markets; that is, periods where capital is not preserved over a 12-month horizon. During these periods, interest rates typically increased sharply. A risk scenario for government bond investing in the coming years is one of sharply rising interest rates. As history shows, such a scenario will result in negative 12-month returns on government bond portfolios. The same conclusion can be based on the probability of negative returns, presented in Table 1. Another interesting observation is the fact that the periods where returns increase and decrease overlap for both the short-term and overall markets. This is a direct result of expression (1): price changes are a result of yield changes and the duration. Apparently the yield changes in both markets are similar, whereas the duration of the overall market is evidently higher than the duration of the short-term market.

4 The international dimension

4.1 The risk of international bond investing

The most important rationale for international bond investing is diversification. A portfolio that is spread over multiple bond markets should have lower risk (or lower volatility of returns) than a portfolio invested in a single government bond market. The theoretical explanation for this is that the international bond markets are not perfectly correlated. There are a number of reasons for this:

- The dynamics of the business cycle, and the role of monetary policy in dealing with the business cycle, differ for each country (see for example Heathcote and Perri 2003 and Kose, Otrok and Whiteman 2003 for recent analyses of international business cycles);
- The role of buyers and sellers varies among government bond markets owing to institutional forces, fiscal policy and national traditions;

Table 2: Correlation coefficients between US and EMU government bond markets and international bond markets (based on monthly total returns converted into USD and EUR respectively)

Correlation coefficients with US		Correlation coefficients with EMU	
Australia	0.19	Australia	0.13
Canada	0.48	Canada	0.19
Denmark	0.31	Denmark	0.61
EMU	0.34	Japan	0.19
Japan	0.13	Sweden	0.23
Sweden	0.28	United Kingdom	0.34
United Kingdom	0.36	United States	0.23

- The trend of inflation and a country's tolerance of inflation vary from country to country;
- Geopolitical, foreign policy, and social developments also differ between countries.

As a consequence, price movements in different government bond markets often offset each other. This should reduce the overall return volatility of a government bond portfolio. Empirical results on the diversification effect of adding international bonds to a US bond portfolio over the period 1978-1995 are presented by Steward and Greshin (1997).

There is no reason from a theoretical point of view to suppose that certain countries or regions will structurally outperform or underperform in the future. There is scientific evidence of investment strategies that systematically beat the market, but all these strategies are active as opposed to passive. Nor is there any empirical support. With hindsight, it is not difficult to find what would have been the best allocation in a certain historical period, but the issue here is whether such an optimal allocation can be established in advance. Using a historically optimal portfolio for the future only provides an illusion of accuracy.

In the remainder of this section we look empirically at the correlation between international bond markets. In Table 2 we show the correlation coefficient between total returns converted into USD and EUR of the US and EMU government bond markets and other government bond markets.

For example, the correlation coefficient of 0.19 for Australia with the US applies to a situation where part of the USD reserves are invested in Australian government bonds. The return on this investment measured in the foreign reserves currency USD consists of two parts: a fixed-income market return (related to the Australian government bond market) and an FX return (related to changes in the AUD/USD exchange rate). For example, if Australian yields fall over the month, resulting in a positive market return of 1%, and the AUD depreciates against the dollar over the month by 2%, the net USD return on the Australian investment will be approximately minus 1%.

An immediate consequence of international bond investing is exposure to an important additional source of risk besides interest rate risk, namely *exchange rate risk* or *currency risk*. The market value of a non-US (or non-EMU) government bond will decrease if (a) the local yield increases and/or (b) the local currency depreciates against USD (or EUR). In a first-order approximation we can express the relative price change of the local government bond as follows (compare with expression (1)):

$$\frac{\Delta P}{P} = -MD \times \Delta y + \frac{\Delta S}{S} \qquad (2)$$

where S is the exchange rate and ΔS the absolute change in the exchange rate.

Table 3: Volatility of bond market total returns: in local currency, converted into USD and EUR
(in %)

Bond market	local currency	converted into USD	converted into EUR
Australia	5.5	11.0	14.3
Canada	5.5	8.1	13.3
Denmark	4.3	10.9	4.9
EMU	3.6	11.3	3.6
Japan	4.5	13.4	12.2
Sweden	4.7	11.0	8.6
United Kingdom	5.9	11.5	10.0
United States	4.4	4.4	11.1

As foreign exchange rates are typically more volatile than interest rates, currency risk will be the dominant factor in international bond investing. This is illustrated in Table 3, in which we present the annualised standard deviation of total returns for eight major bond markets. We compare the volatility of total returns in local currency with the volatility of total returns converted into USD and EUR respectively. We note that the volatility in the first column results from interest rate volatility alone, while the volatility in the other two columns results from both interest rate and FX volatility. The figures show that the additional exchange rate risk indeed significantly increases the volatility (and hence the risk) of international bond investments.

Another way of looking at the importance of currency returns is illustrated in Figure 4, where we show the influence of the currency factor on monthly global bond returns.

On average, 56% of monthly bond returns are the result of currency movements. Note that this percentage declines if we consider a time horizon longer than one month.

The risk of international bond investing can be reduced significantly by hedging the currency risk in the portfolio. For example, for a USD reserves portfolio we can hedge the JPY/USD currency risk of the JPY-denominated part of a government bond portfolio using JPY/USD currency forward contracts. To purchase a Japanese government bond, two separate transactions must be made. The bond must be purchased, and JPY must be purchased to pay for the bond. To hedge the JPY currency component of the bond, a simultaneous sale of the JPY is transacted through the forward currency market. In this way, the investor locks in a sale price for the JPY, thereby eliminating practically all exposure to currency volatility over a certain period in time. The common practice in currency hedging is to use rolling

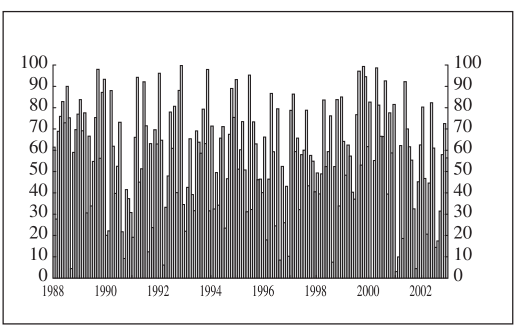

Figure 4: Influence of currency factor on monthly global bond returns, Jan. 1987-Jan. 2002
(in %)

Table 4: Correlation coefficients between US and EMU and international bond markets hedged to USD and EUR respectively (based on monthly total returns)

Correlation coefficients with US		Correlation coefficients with EMU	
Australia	0.59	Australia	0.44
Canada	0.70	Canada	0.50
Denmark	0.41	Denmark	0.72
EMU	0.58	Japan	0.34
Japan	0.25	Sweden	0.53
Sweden	0.35	United Kingdom	0.70
United Kingdom	0.51	United States	0.59

forward contracts of one month to one year in duration, which may be renewed at the expiration date. Rolling forward contracts are not perfect hedges, but the amount of residual foreign exchange exposure is generally quite small.

Transaction costs for this rolling hedge strategy are low as bid-ask spreads for forward contracts are small. A larger cost impact is the periodical settlement of the realised profits and losses from the hedge positions in the bond portfolio. An often-used approach is to run a (small) separate currency settlement account beside the bond portfolio.

As stated above, the argument for risk reduction through international bond investing is based on low correlation coefficients between international government bond market returns. An interesting question is how this argument holds out when we consistently hedge currency risks. The answer to this can be found in Table 4. Here we show the correlation coefficients between the major international government bond markets and the US and EMU government bond markets if the currency risk is consistently hedged.

It should come as no surprise that hedging the currency risk increases the correlation of international bond markets. All correlation coefficients in Table 4 are larger than in Table 2. Nevertheless, international bond markets are far from perfectly correlated, which gives room for risk reduction through diversification in an international hedged government bond portfolio. Grinold and Meese (2000) even suggest that both hedged and unhedged international assets should be considered as separate asset classes when conducting strategic asset allocation studies.

To complete the risk picture of international bond investing, we need to address the additional risk taken when currency risks are hedged. As an investor, we are exposed to *default risk* on the counterparty of the currency forward contracts. In practice the forward sale of a currency entails no money changing hands between the currency dealer and the investor. Therefore the nominal amount exposed to default risk is rather limited, and as a result credit risk in a hedged international government bond portfolio is negligible compared with market risk. The credit risk can be reduced even further by diversifying over currency dealers when initiating the currency hedges. The risk of a hedged international bond portfolio is by definition larger than that of an unhedged international bond portfolio, as more financial instruments are involved. For the currencies of the major bond markets, however, the OTC currency derivative market is very broad and deep and marketability is comparable with that of the government bonds themselves.

The best-known theory on exchange rate developments is the purchasing power parity theory (see for example Balassa 1965 and Froot and Rogoff 1995). The theoretical concept of long-term purchasing power parity means that over the long term, hedging will lead to the same return as leaving currency positions open.

Figure 5: Difference in annual returns between open and hedged investments in global bond markets
(in %)

a) (from a USD perspective) b) (from an EUR perspective)

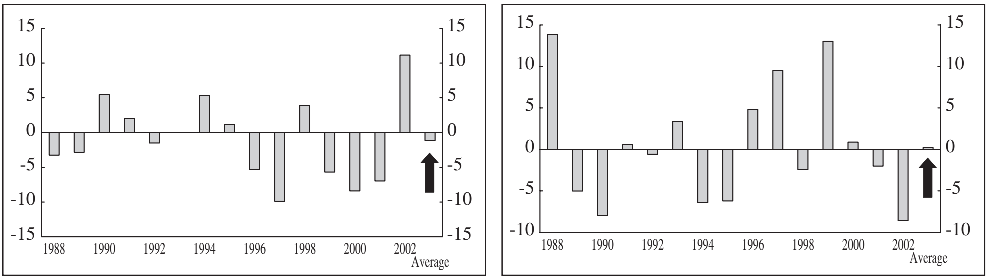

Empirical results further support the proposition that hedging currency risks does not lead to a higher or a lower return over the long term. Figure 5a and Figure 5b illustrate this. They show the differences in return between hedged and open investments for a USD and EUR investor in the overall global bond market. The short-term differences in return can be significant, but over the long term the difference in practice is zero.

4.2 *Empirical results*

Figure 6a and Figure 6b are an extension to Figure 1a and Figure 1b. We now present the risk/return profile of both single-currency (respectively USD and EUR) and global bond markets, both hedged and unhedged. The line between the 100% single currency and 100% global allocations correspond to portfolios partially allocated, for example, to the US Treasury market and the international bond market. More risk characteristics of global bond investing can be found in Table 5 below.

Figure 6: Risk/return characteristics

a) of USD and global bond markets, Dec. 1987-Dec. 2002 (in %)

b) of EUR and global bond markets, Dec. 1987-Dec. 2002 (in %)

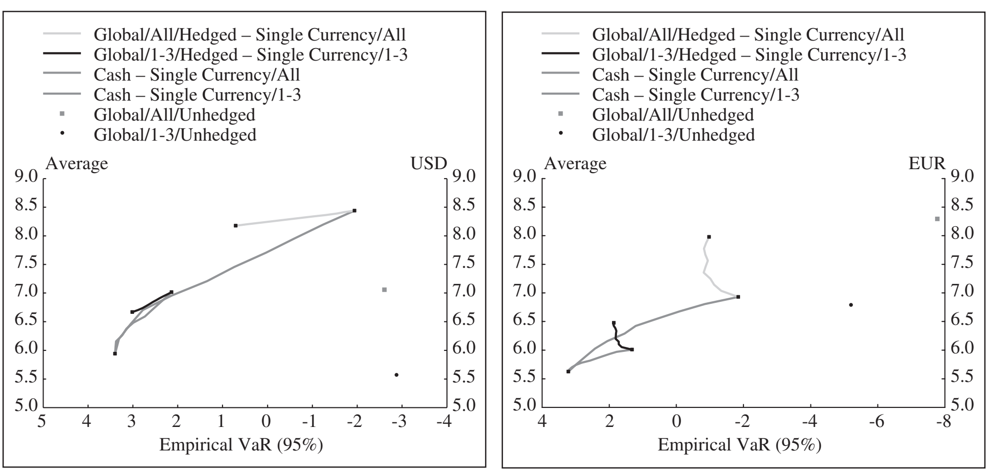

Table 5: Risk/return profile of global bond markets

Market	Cash	Single currency	Single currency	Global	Global	Global	Global
Maturity segment Currency		1-3	all	1-3 hedged	all hedged	1-3 unhedged	all unhedged
				Base currency USD			
Average return	5.9%	7.0%	8.4%	6.7%	8.2%	5.6%	7.1%
Volatility	0.6%	1.8%	4.4%	1.4%	3.3%	5.3%	6.2%
Sharpe ratio	n/a	0.61	0.55	0.57	0.66	-0.07	0.17
VaR (95%)	3.4%	2.1%	-1.9%	3.0%	0.7%	-2.9%	-2.6%
Probability of negative return	0%	0%	8%	0%	4%	18%	14%
Probability of outperf. (1 yr)	n/a	66%	69%	69%	72%	49%	49%
Probability of outperf. (3 yrs)	n/a	79%	89%	90%	89%	43%	56%
				Base currency EUR			
Average return	5.6%	6.0%	6.9%	6.5%	8.0%	6.8%	8.3%
Volatility	0.7%	1.8%	3.6%	1.5%	3.4%	6.5%	7.1%
Sharpe ratio	n/a	0.22	0.35	0.64	0.69	0.17	0.35
VaR (95%)	3.2%	1.3%	-1.8%	1.9%	-1.0%	-5.2%	-7.8%
Probability of negative return	0%	1%	14%	0%	7%	21%	17%
Probability of outperf. (1 yr)	n/a	60%	66%	69%	73.4%	57%	66%
Probability of outperf. (3 yrs)	n/a	79%	77%	92%	94%	58%	68%

Going from an investment in a single-currency bond market to an unhedged investment in global bond markets sharply increases risk. The increase in volatility and the probability of a negative return is staggering, without any compensation in the form of additional return. This is also visible in the decline of the Sharpe ratio and the probability of outperformance.

Hedging currency risk completely changes the picture. The diversification effect introduced by international bond investing is apparent for both base currencies. This is supported by all the risk measures calculated. The volatility of the (hedged) global bond market returns, and also the probability of a negative return, is considerably lower than their single currency counterparts.

Figure 7: Total return indices for the global bond markets, hedged and unhedged

a) USD, Dec. 1987-Dec. 2002 b) EUR, Dec. 1987-Dec. 2002

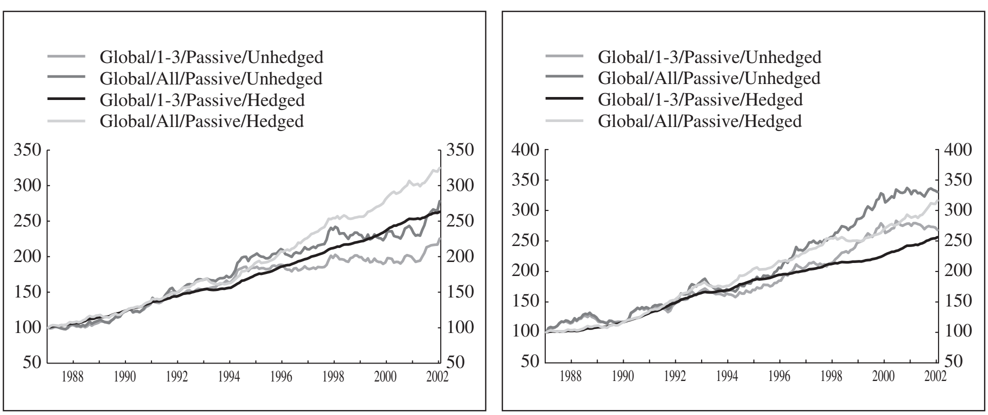

Figure 8: Rolling 12M returns for the global bond markets, hedged and unhedged
(in %)

a) USD, Dec. 1987-Dec. 2002 a) EUR, Dec. 1987-Dec. 2002

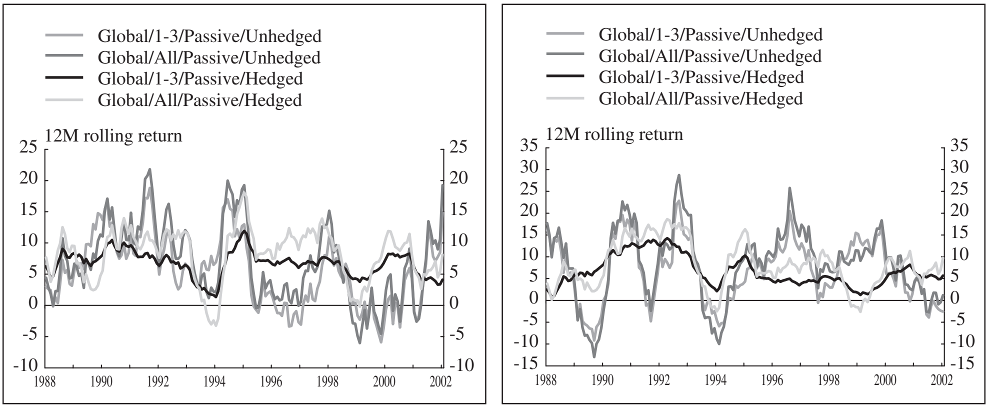

The path of the cumulative return of the global bond markets is depicted in Figure 7a and Figure 7b. The rolling 12-month returns for hedged and unhedged international bonds are shown in Figure 8a and Figure 8b.

We see a large discrepancy between the paths of the unhedged and hedged returns. The paths of the hedged returns are much smoother over time. The correlation between the two unhedged return series is more than 90%, versus less than 84% for the hedged series. This is due to the high volatility of exchange rates, which do have an impact on the unhedged returns and thus increase the interdependence of the unhedged return series. A negative effect of the additional currency risk in the unhedged portfolios is depicted in the rolling 12-month return figures. Owing to adverse currency movements, we see a large number of periods in which there is no capital preservation over a one-year horizon. This provides a strong argument for hedging currency risk in global bond investments, particularly for managers of foreign reserves. We also note that the one-year probability of outperforming cash is much higher for the hedged global bond return (72% and 73% for USD and EUR respectively) than the unhedged counterpart (49% and 66% respectively). The differences are even more apparent for the three-year probabilities.

5 Adding active duration management

5.1 Introduction

As stated in the previous chapter, an important rationale for international bond investing is risk reduction through diversification. From the return perspective, however, there are no strong fundamental arguments which demonstrate that international bond investing generates superior rates of return compared with single country (USD or EMU) bond investing. The fact that in the EUR (USD) risk return figure the average return for the global hedged portfolio is higher (lower) than for the single country portfolio is a pure coincidence. If we consider a different time frame (e.g. December 1992 through December 2002), the results would have been the other way around.

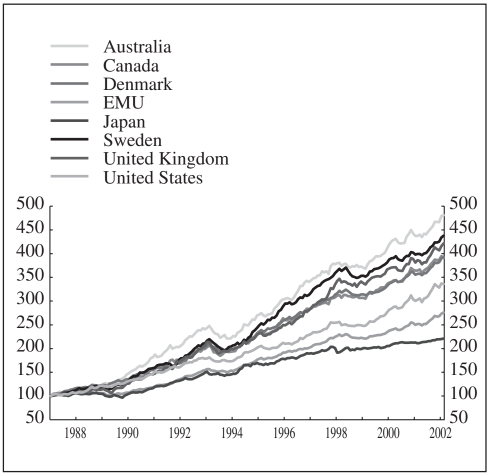

Figure 9: Total return indices of eight bond markets in local currency, Dec. 1987-Dec. 2002

So should we conclude that there is no return enhancement argument for international bond investing? The answer is no. An international bond portfolio offers attractive return enhancement features through successful *active management* (see for example Rosenberg 1997).

In Figure 9 we show the cumulative total returns of eight major bond markets. The chart shows the total return of the JP Morgan Government Bond Country Indices in local currency.

We show the correlation between local returns over the same data period in Table 6 below.

What we see is that there are periods when all bond markets rally and periods when (almost) all bond markets show negative returns. From the charts it is also clear that during any given period some bond markets perform much better (or much worse) than others. With hindsight, by actively switching between bond markets, one could have obtained a much higher average return for an actively managed international bond portfolio than for a passive international bond portfolio. In the next two sections we address two important questions. The first question is how one can *define* and *implement* an active international bond investing strategy. The second question is how much more risky an actively managed international bond portfolio is compared with a passive international bond portfolio. We choose the JP Morgan Government Bond Index as the *benchmark* passive global bond portfolio.

Table 6: Correlation coefficients between local bond markets (based on monthly total returns)

	Australia	Canada	Denmark	EMU	Japan	Sweden	UK	US
Australia	1.00							
Canada	0.64	1.00						
Denmark	0.35	0.40	1.00					
EMU	0.41	0.48	0.70	1.00				
Japan	0.25	0.29	0.11	0.33	1.00			
Sweden	0.43	0.38	0.55	0.55	0.21	1.00		
UK	0.44	0.55	0.56	0.68	0.24	0.53	1.00	
US	0.59	0.71	0.41	0.57	0.25	0.33	0.51	1.00

5.2 *Example of an active duration management strategy*

The main source of risk in a hedged international bond portfolio is market or interest rate risk. Modified duration is the measure for the exposure to interest rate risk. An international bond portfolio is exposed to market risk for all of the government bond markets that are part of the portfolio. The exposure of the portfolio to Japanese yield changes, for example, depends both on the modified duration of the part of the portfolio invested in Japanese government bonds and on the percentage weight of this part of the portfolio. For example, suppose that 25% is invested in Japanese government bonds and that the modified duration of this part of the portfolio is equal to four years. Then the partial duration of Japan in the international portfolio is equal to 1.0 year.

By overweighting and underweighting the partial duration of countries relative to the benchmark weights, one can profit from interest rate movements in the different countries. More interest rate sensitivity in a certain country can be achieved by increasing the market weight of the country in the portfolio (i.e. nominal overweight) or by increasing the duration of the part of the portfolio invested in the country (i.e. duration overweight).

There are various ways of taking active duration and country allocation decisions in an international government bond portfolio. A first distinction one can make is between fundamental and technical analysis. Fundamental analysis involves decisions taken on the basis of a country's economy, monetary and fiscal conditions, as well as its political situation. Technical analysis decisions, on the other hand, are taken based on historical price developments, and volume and sentiment indicators. Both fundamental and technical analysis can be approached in a qualitative or quantitative way. In a qualitative approach, investors assess the relative attractiveness of markets based on their interpretation of available information. A purely quantitative approach uses the available statistical material as input for a quantitative forecasting model which arrives at recommendations for markets by processing this data. It is common practice in international fixed-income investing to combine several or all of these types of market analysis to decide on the active partial duration positions in the portfolio.

In the next section we show the empirical results of an active management strategy that is purely based on the outcomes of a quantitative forecasting model. The model is Robeco's proprietary Duration and Country Allocation Model, which has been used in the investment process for global bond portfolios since 1995. Results of a similar active management strategy can be found in Ilmanen and Sayood (2002).

Before we turn to the empirical results, we will briefly describe the model itself. The duration model bases its forecasts for a bond market on six different indicators: a local bond market valuation indicator, a global bond market trend indicator, a local bond market mean reversion indicator, a local equity market trend indicator, a commodity price indicator and a seasonal indicator[9]. An important feature of the model is that all information for the indicators is available at the moment an active position is taken (i.e. no judgemental information about the future is used). These indicators have been carefully selected on the basis of economic interpretation and historical forecasting power. In each case, there is a clear economic relationship with the bond market that has been empirically demonstrated. Correlation between the indicators is low: each indicator reflects a different aspect of the bond market and thus adds value.

---

[9] This is closely related to the Halloween indicator used in equity investment strategies. Bouman et al. (2002) describe this calendar effect.

Table 7: Examples of scores and signals from the quantitative model

	US	CA	AU	EMU	UK	DK	SE	JP
Score	-0.02	0.11	-0.08	0.72	0.35	-0.08	0.26	-0.43
Signal	0	+	0	+	+	0	+	-

A normalised score (a so-called "z-score") is calculated for each country and each indicator. The average of the scores over the indicators provides the country score for each country. The country scores are translated into either a positive, a negative or a neutral signal. A negative country signal results if the country score is lower than the threshold -0.10; a positive country signal results if the country score is higher than the threshold +0.10. For country scores between -0.10 and +0.10, a neutral signal results. Table 7 illustrates this.

The table shows that the model generates recommendations for both active duration allocation and active country allocation. First, it supports the overall duration policy. Based on the above signals, we would want to maintain overall duration at higher than the benchmark level in our portfolio. After all, the model is positive for most government bond markets. Second, the model supports the country allocation. In this example, besides applying a higher duration, we would also want to underweight Japan versus the benchmark and overweight Canada, Sweden, the UK and EMU.

If we apply an active management strategy to an international bond portfolio, by definition this adds an additional source of risk. The reason is that on top of the static interest rate risk (and possibly currency risk) of the passive international bond portfolio, we are dynamically adding or removing interest rate risk at either the total portfolio or the individual country level. Active duration management will lead to returns that are different from those of the passive or benchmark portfolio. The larger the active duration bets (in terms of size of partial duration underweights and overweights), the greater this difference can be. The annualised volatility of the return differences between the active portfolio and the benchmark portfolio is known as the *tracking error*; this is the standard for the quantification of the *active risk* in an investment portfolio. Tracking errors for actively managed government bond portfolios are typically in the range of 1% to 2%. The total risk of the actively managed portfolio can be seen as a combination of the volatility of the benchmark portfolio and the tracking error. As the correlation between active returns and passive returns is less than perfect, the total risk (active plus passive) of an actively managed portfolio is usually only slightly higher than the total risk of the benchmark portfolio. It should be noted that an active strategy even has the potential to decrease the overall risk of a portfolio.

5.3 Empirical results for active management

In this section we show empirical risk/return results of an active management strategy compared with a passive international bond investing strategy. The active duration positions are based on the signals of the quantitative duration and country allocation model described in the previous section. The duration bet size in each country is equal to 6.25% of the duration of the bond market in that country (represented by the corresponding JP Morgan country index). For example, if the duration of the country index is equal to five years, then a plus-signal in the model for that country results in an overweight of 0.31 years of the partial duration of that country in the actively managed global bond portfolio. A similar approach is applied to minus signals from the model. With these bet sizes, the maximum overweight position (in case of plus signals for all eight countries) for the overall portfolio is approximately equal to 50% of

Figure 10: Risk/return characteristics of passive and active investment in global bond markets, (in %)

a) USD, Dec. 1987-Dec. 2002 b) EUR, Dec. 1987-Dec. 2002

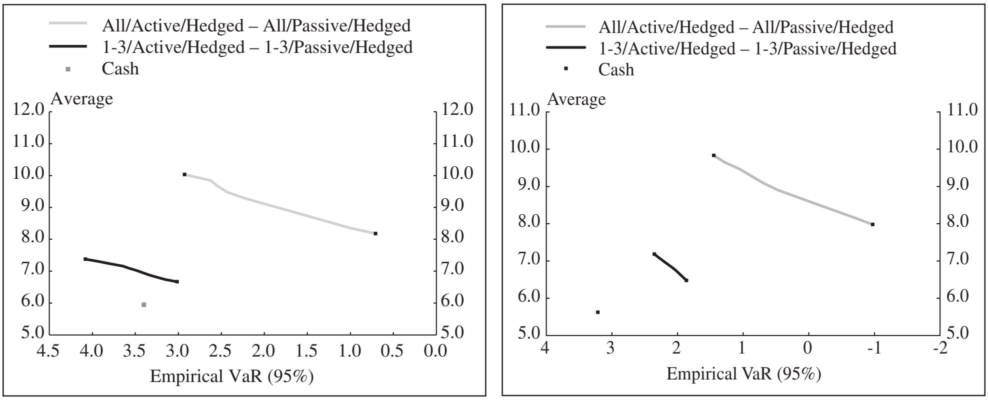

the duration of the passive global bond benchmark. For example, the overall duration of the actively managed portfolio will vary between 2.5 and 7.5 years compared to a benchmark duration of five years. Partial duration overweights can be implemented by physically trading bonds (replacing short-term bonds in the portfolio with long-term bonds). For reasons of convenience and to limit transaction costs, as well as to avoid unintended yield curve bets in the portfolio, it is recommended to implement active positions via bond futures.

In Figure 10a and Figure 10b we compare the risk/return profiles of passive global bond benchmarks and actively managed government bond portfolios. The results can be

Table 8: Risk/return profile of passive and active investment in global bond markets

Market	Cash	Single currency	Single currency	Global	Global	Global	Global
Maturity segment		1-3	all	1-3	all	1-3	all
Currency				hedged	hedged	hedged	hedged
Strategy				passive	passive	active	active
Base currency USD							
Average return	5.9%	7.0%	8.4%	6.7%	8.2%	7.4%	10.0%
Volatility	0.6%	1.8%	4.4%	1.4%	3.3%	1.4%	3.5%
Sharpe ratio	n/a	0.61	0.55	0.57	0.66	1.03	1.12
VaR (95%)	3.4%	2.1%	-1.9%	3.0%	0.7%	4.1%	2.9%
Probability of negative return	0%	0%	8%	0%	4%	0%	1%
Probability of outperf. (1 yr)	n/a	66%	69%	69%	72%	73%	78%
Probability of outperf. (3 yrs)	n/a	79%	90%	90%	89%	95%	100%
Base currency EUR							
Average return	5.6%	6.0%	6.9%	6.5%	8.0%	7.2%	9.8%
Volatility	0.7%	1.8%	3.6%	1.5%	3.4%	1.6%	3.6%
Sharpe ratio	n/a	0.22	0.35	0.64	0.69	1.09	1.14
VaR (95%)	3.2%	1.3%	-1.8%	1.9%	-1.0%	2.4%	1.4%
Probability of negative return	0%	1%	14%	0%	7%	0%	0%
Probability of outperf. (1 yr)	n/a	60%	66%	69%	73%	75%	79%
Probability of outperf. (3 yrs)	n/a	79%	77%	92%	94%	97%	100%

Figure 11: Total return indices for passive and active global bond market strategies

a) USD, Dec. 1987-Dec. 2002 b) EUR, Dec. 1987-Dec. 2002

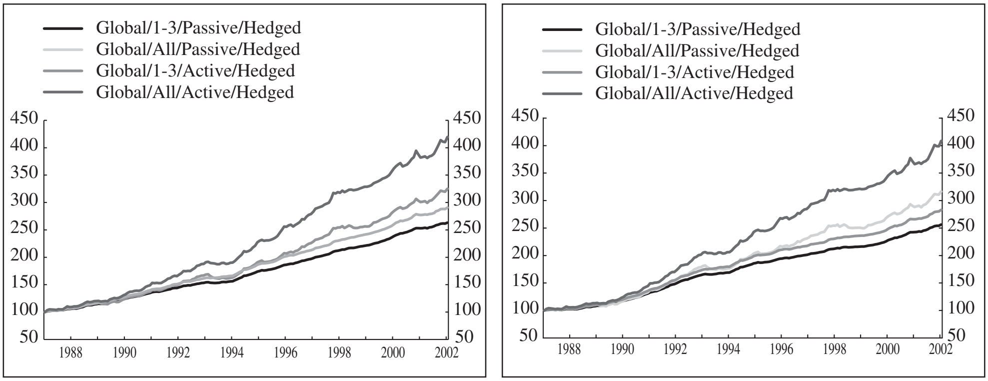

seen as theoretical back-test results of the quantitative duration and country allocation model.

Over the back-test period, the active management strategy apparently paid off. From Table 8 we can deduce that the annual return increases by 70 basis points (short-term bond market) and 180 basis points (overall bond market) respectively. Clearly these results depend on the quality of the quantitative forecasting model. From a risk perspective, we see that there is hardly any increase in volatility for the actively managed portfolio. Despite the fact that we introduce a substantial amount of active risk (overweights and underweights up to 50% of the duration of the benchmark), the absolute risk of the actively managed global bond portfolio, measured using VaR, is even lower than the absolute risk of the passively managed benchmark. The added value is evident when we examine the Sharpe ratios: these almost double when an active strategy is applied.

The paths of the cumulative return of the passive and active strategies are shown in Figure 11a and Figure 11b. The rolling 12-month returns are presented in Figure 12a and Figure 12b.

Figure 12: Rolling 12M returns on the passive and active global bond market strategies (in %)

a) USD, Dec. 1987-Dec. 2002 b) EUR, Dec. 1987-Dec. 2002

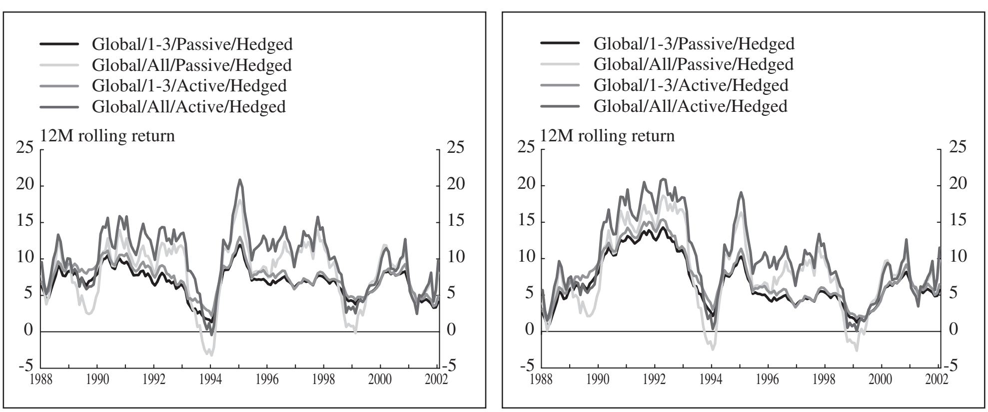

The return paths of the passive portfolios are clearly smoother than the return paths of the actively managed portfolios. However, in this case the active investor is rewarded for the extra risk taken by higher returns. The rolling 12-month return figures show us that active management is a way to avoid periods of 12 months with negative returns (the probability of negative returns is at most 1% for any of the active strategies).

The overall active bond strategy outperforms the cash market over a one-year horizon with a probability of 78% (versus 72% and 73% for the passive USD and EUR strategies respectively). Looking back, the overall active bond strategy has never underperformed cash over a three-year investment horizon.

6 Beyond government bonds

6.1 Investment-grade credits, high-yield bonds and equities

In searching for return, investors have increasingly turned to alternatives to government bonds. As stated before, US Treasuries are considered to be free of default risk. It is assumed that the US Treasury will always be able to make timely principal and interest payments on all its issued securities. Clearly this does not hold for all issuers of fixed-income securities. Investors in a bond issued by, for example, a corporation face the risk that this corporation at a future point in time might not be able (or willing) to fulfil its contractual obligations. Hence investors in this type of bond face a credit risk. As a consequence, these bonds are sold at a lower price to, and hence at a yield spread over, comparable US Treasury securities. This yield spread, which compensates the investor for this extra credit risk, is known as the *credit spread*.

Although credit spread is often purely attributed to the default risk of a credit, part of the spread can be attributed to other factors. The literature on the pricing of corporate bonds is extensive. Elton et al. (2001) argue that expected default can account for only a small part of the yield spread for investment-grade corporate over Treasury bonds, while state taxes (which are payable on corporate interest but not on Treasury interest) are relatively much more important. The remaining portion of the spread is closely related to the factors that we commonly accept as explaining risk premiums for common stocks. Huang and Huang (2003) reach a similar conclusion.

King et al. (2002) decompose the yield to maturity of a risky bond into three components: (1) the default-free yield to maturity; (2) the risk premium, which is the difference between the expected and default-free yield to maturity; and (3) the default premium, which is the difference between the promised yield to maturity and the expected yield to maturity. Their results suggest that cross-sectional differences in yield spreads at a given point in time are largely due to differences in the default premium portion of the yield spread. On the other hand, time variation in yield spreads is largely due to differences in the risk premium portion of the yield spread.

The different studies conclude that the credit spread consists of two components: a compensation for expected loss, which is determined by the expected default and the recovery rate, and a risk premium.

Figure 13 depicts the development of credit spreads over the last 14 years, whereas Figure 14 shows the historic annual one-year default rates of investment grade and high yield bonds. When we compare Figure 13 and Figure 14 we notice strong similarities between the two graphs: the years with peaks and troughs in annual default rates and spread levels coincide.

Figure 13: Credit spread per rating class, 1989-2002

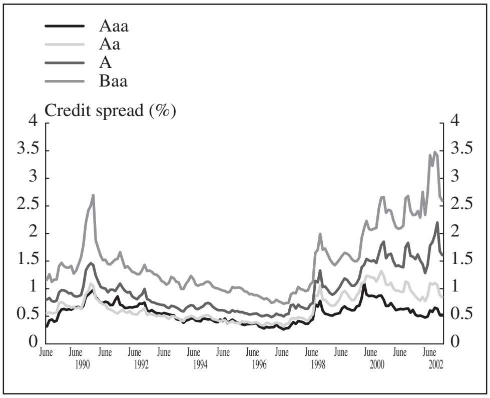

Source: Lehman Brothers.

Figure 14: Annual default rates per rating class, 1989-2002

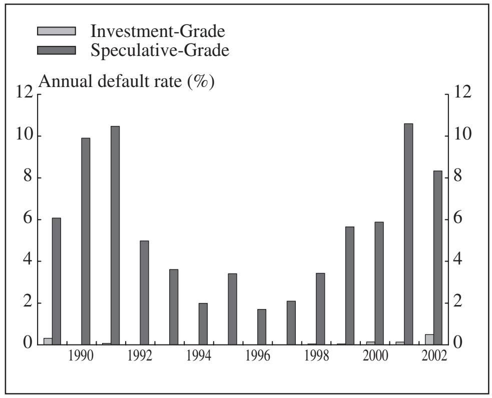

Source: Lehman Brothers.

Credit risk is gauged by quality ratings assigned by commercial rating companies such as Moody's Investor Service, Standard & Poor's, and Fitch Investors Service. The safer, higher-rated credit securities are known as investment-grade credits, whereas the riskier, lower-rated credit securities are known as high-yield bonds. For investment-grade credits, the probability of an actual default event within one year's time, for example, is negligible. Investors in investment-grade credits are more concerned with changes in perceived credit risk and/or the cost associated with a given level of credit risk. This is because the impact of a change in perceived credit risk or the spread demanded by the market for a given level of credit risk has an immediate impact on the security. The price of a corporate bond moves in an opposite direction to the change in credit spreads. As the credit spread widens, the price of the bond will fall; similarly, as the credit spread narrows, the price of the bond will rise. For buy-and-hold investors who plan to hold the bond to maturity, changes in the bond's price are of no concern. However, for investors who may have to sell the bond before the maturity date, a widening of the credit spread after buying the bond results in the realisation of a capital loss. The mechanics are similar to those of interest rates and government bonds: the higher the duration of the credit, the larger the impact of a change of the credit spread on the price of the bond. Therefore for credits, the market risk consists of two components: interest rate risk and credit spread risk. Changes in either of these two market variables have a direct impact on the price of the security. The variability of spreads, visible in Figure 13, is probably an important factor. The lower the credit rating, the higher the corresponding spread and spread volatility (and thus credit risk).

Below we show a scatter diagram of the total returns for the JP Morgan US Government Bond Index and the excess returns of the Lehman Brothers US Credit Index.

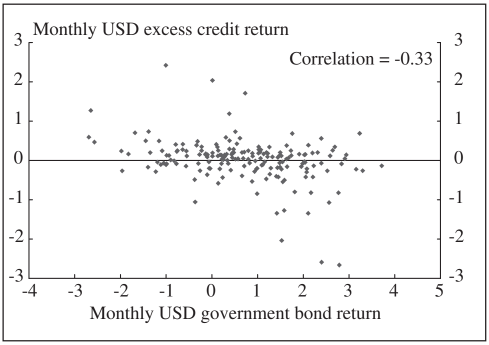

Figure 15: Scatter diagram of government bond returns and excess spread returns (monthly returns)
(in %)

The figure shows that months of very negative government bond returns (owing to interest rate hikes) typically also show a positive excess spread return (owing to decreasing credit spreads). Similarly, months of very negative excess spread returns typically show positive government bond returns. The correlation coefficient between the two return series appears to be negative (-0.33). Owing to this diversification effect, the investment-grade credit class may appear to be less risky (i.e. have a lower total return volatility) than the "safer" government bond class.

An important risk difference between investment-grade credits and government bonds is that the marketability of the former is much lower. Hence the liquidity risk associated with investing in credits is much higher than for government bonds. Bid/offer spreads for credits are in the order of tens of basis points, and reducing a position in a credit can take several business days.

For investments in the high-yield market, the risks are quite different. The default risk of high-yield bonds is much more apparent: the investor in a high-yield bond faces the real risk of the issuer being unable or unwilling to fulfil contractual payments in the coming years. For example, for Ba-rated bonds, Moody's reports an annual default probability of 1.2% as measured over the 1982-2001 period. For the lower-rated class of single-Bs, Moody's reports an annual default probability of 6.5% over the same period. As a result, the premium for default risk, the credit spread, is much higher for high-yield bonds than it is for investment-grade bonds. From a risk perspective, high-yield bond investing is more similar to equity investing than to government bond investing. As with equities, the market risk of a high-yield bond is typically a specific price risk, whereas for government bonds market risk arises from the combination of interest rate volatility and the duration of the bond. The marketability of high-yield bonds is much lower than it is for equities. Of all the asset classes considered in this paper, high-yield bonds have the highest liquidity risk. Leaving the territory of "safe" bond investing, it is only a small step to considering investments in equities. Over the past century, equity investing has proven to be riskier than bond investing, but also more rewarding. The extra return earned by investors in equities over bonds was nearly 7% on an annual basis, as stated by Dimson et al. (2002). As we shall see in the empirical results in Section 6.3, these long-term attractive return figures for equities cannot be guaranteed in the short term.

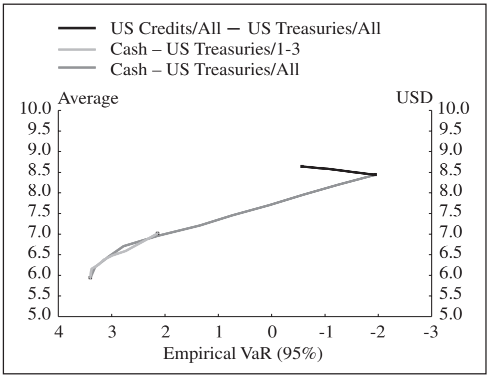

Figure 16: Risk/return characteristics of USD Treasuries and USD credits, Dec. 1987-Dec. 2002
(in %)

6.2 *Empirical results for investment grade credits*

Currently the European investment-grade credit market is only 40% of the size of its USD counterpart. The European credit market itself is still rather young. As data on the European credit market are only available since 1999, this section focuses on the USD credit market.

Figure 16 presents the risk/return profile of the USD government bond market and USD investment-grade credits, all from a USD perspective. A summary of risk/return characteristics is provided in Table 9.

We can see that the risk of investing in USD investment-grade credits is comparable to the risk of investing in USD Treasuries. The volatility of the investment-grade credit returns is even slightly lower than the volatility of the government bond returns. The same is true if we use Value at Risk as the risk measure, owing to the negative correlation between interest rate changes and credit spread changes. An important remark to be made at this point is that the returns presented here are the returns of a broadly diversified credit index. This index is only exposed to systematic credit risk, not to issue-specific risk. Any investment-grade credit portfolio will be exposed to issue-specific risk. This issue-specific risk can be seen as active risk in the credit portfolio compared with the credit index. Well-structured portfolios containing 50 to 100 securities can have reasonably low tracking errors of 0.5% to 0.3% per year.

Over the period under consideration, credits outperform Treasuries on both the risk and return dimensions.

Table 9: Risk/return profile of USD investment-grade credits

Market	Cash	Treasury	Treasury	Credits
Maturity segment		1-3	all	all
Currency		USD	USD	USD
Base currency USD				
Average return	5.9%	7.0%	8.4%	8.6%
Volatility	0.6%	1.8%	4.4%	4.2%
Sharpe ratio	n/a	0.61	0.55	0.63
VaR (95%)	3.4%	2.1%	-1.9%	-0.6%
Probability of negative return	0%	0%	8%	7%
Probability of outperf. (1 yr)	n/a	66%	69%	70%
Probability of outperf. (3 yrs)	n/a	79%	89%	83%

Figure 17: Total return indices for USD Treasuries and USD credits, Dec. 1987-Dec. 2002

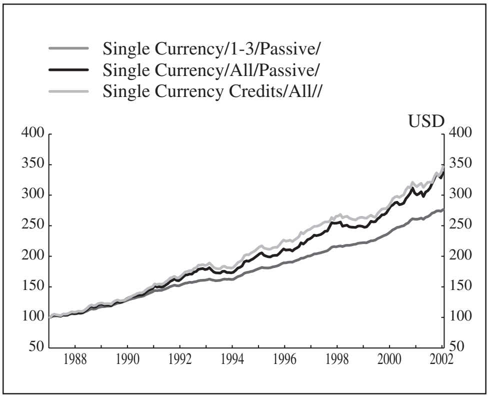

Figure 18: Rolling 12M returns on USD Treasuries and USD credits, Dec. 1987-Dec. 2002
(in %)

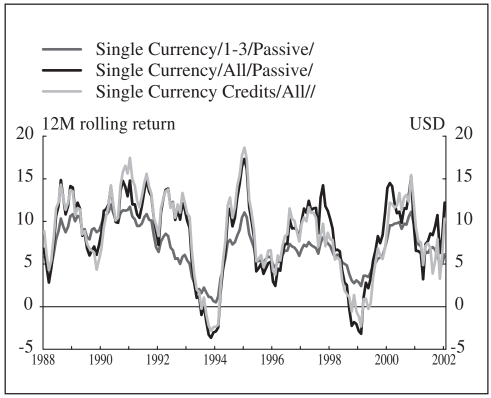

The cumulative returns are presented in Figure 17, and the 12-month rolling returns in Figure 18. The negative correlation between interest rate changes and credit spread changes results in a smoother total return index for investment grade credits than for government bonds. The 12-month rolling returns for credits show that the negative correlation between interest rate changes and credit spread changes is not strong enough to prevent negative returns over a one-year horizon. The past years have been tough ones for credit investors, with credits underperforming government bonds. Nevertheless, 12-month rolling credit returns have remained well on the positive side over the past few years.

6.3 Empirical results for high yield and equities

In this section we examine the risk/return profile of an investment in high yield and equities. In parallel with the previous section, we again focus on the USD viewpoint[10].

Figure 19 presents the risk/return profile of USD Treasuries, USD high yield and USD equity. Based on this, it is hard to make a case for high-yield investments. Risk measured as volatility of returns is nearly twice as large for high-yield assets as for government bonds.

Table 10: Risk/return profile of USD high yield and equity

Market	Cash	Treasury	High Yield	Equity
Maturity segment		all	all	all
Currency		USD	USD	USD
Base currency USD				
Average return	5.9%	8.4%	7.8%	11.6%
Volatility	0.6%	4.4%	7.8%	14.9%
Sharpe ratio	n/a	0.55	0.22	0.36
VaR (95%)	3.4%	-1.9%	-5.6%	-18.0%
Probability of negative return	0%	8%	23%	18%
Probability of outperf. (1 yr)	n/a	69%	51%	78%
Probability of outperf. (3 yrs)	n/a	89%	68%	86%

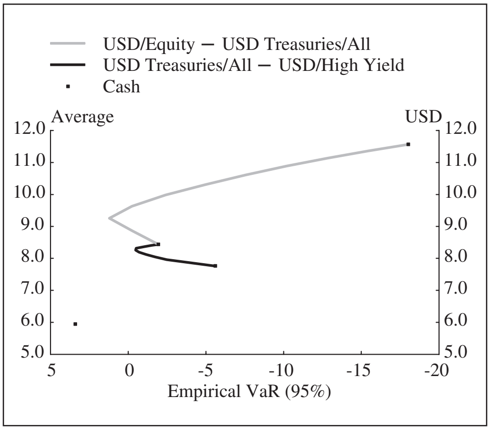

Figure 19: Risk/return characteristics of USD Treasuries, high yield and equity, Dec. 1987-Dec. 2002 (in %)

This extra risk was not however compensated for by extra return. Indeed, the opposite is true: the risk/return profile of high yield (and, to a lesser extent, equity as well) is far from impressive.

Note that the poor results for both asset classes are mainly due to the last three to four years, as can be concluded from Figure 20 (cumulative returns) and Figure 21 (12-month rolling returns). One would expect to be remunerated for the extra risk in both asset categories. In the long run we expect both high yield and equity to outperform Treasuries.

From a diversification point of view, high yield and equity are interesting asset classes. Table 11 presents the correlation between the asset classes. The correlation between high yield and Treasuries and equity and Treasuries is respectively only 0.11 and 0.08 (which can be compared with a correlation of 0.90 between credits and Treasuries). The high correlation between equity and high yield seems to imply that high yield behaves like a deleveraged position in equity. Figure 20 and Figure 21 tell us that this is only part of the story. Equity is in a league of its own. First of all, the volatility is nearly twice that of high yield (and more than three times that of Treasuries). Second, although this high volatility results in a relatively high probability of negative returns (18%), it does perform better than high yield on this risk measure. And third, there are periods where high yield outperforms equity in bull markets and vice versa in bear markets.

If capital preservation is the objective of the foreign reserves manager, equity is an unacceptable asset class: the lowest annual return over the fifteen year horizon is -27%. The

Table 11: Correlation coefficients between different US asset classes (based on monthly total returns)

	USD Treasuries	USD credits	USD high yield	USD equity
USD Treasuries	1.00			
USD credits	0.90	1.00		
USD high yield	0.11	0.43	1.00	
USD equity	0.08	0.27	0.49	1.00

[10] In this section we present results for the US market. We have shown that the risk/return profile of the hedged global bond market dominates that of a single currency bond market. The same principle holds for the equity market: in the long run, a hedged position in the global equity markets will dominate a position in a single currency equity market.

Figure 20: Total return indices for USD Treasuries, high yield and equity, Dec. 1987-Dec. 2002

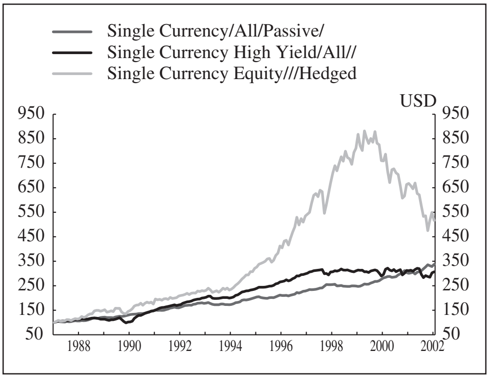

Figure 21: Rolling 12M returns on USD Treasuries, high yield and equity, Dec. 1987-Dec. 2002

(in %)

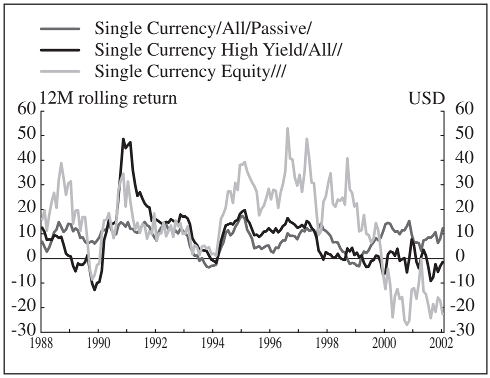

upside potential, on the other hand, is tempting: the highest annual return over the same period is 53%.

7 Conclusion

Over the past decade, the management of foreign currency reserves has changed its focus to an objective of maximising total returns. In this paper we have discussed several ways of enhancing return over a cash benchmark: single currency government bond investing, international government bond investing, investing in investment-grade credits and high-yield bonds, and equities. The cash benchmark has the lowest risk in terms of volatility of returns. Return enhancement is not possible without accepting a higher level of total risk. Diversification over asset classes in this case automatically leads to more risk. The assessment of an acceptable total risk level is an important first step to take, and can, for example, take the form of a maximum acceptable volatility of investment returns. The next step is to decide on the allocation of the available risk budget. One option is to invest all reserves in a relatively low-risk asset class such as short-term (term to maturity between one and three years) government bonds. Another option is to invest part of the reserves in a higher-risk asset class such as longer-term government bonds, investment-grade credits or equities, and then to invest the remainder in the cash benchmark. Both options can lead to comparable levels of absolute risk for the reserves portfolio. The final allocation over asset classes will be based on other decision variables, such as return expectations, available resources, limitations on operational matters, and liquidity requirements.

We have shown in this paper that international bond investing can be appealing both from a return perspective and a risk perspective. On the risk side, diversification over countries results in lower risk if currency risk is hedged. On the return side, we have seen that active duration management offers the opportunity to enhance portfolio returns with only a small amount of extra overall risk. Another interesting result shown in this paper is that the returns of an international bond portfolio with open currencies are much more volatile than the returns of its hedged alternative without an extra return compensation, even in the long run. This illustrates that extra risk is not automatically compensated for by extra return.

References

Artzner, P., F. Delbaen, J. M. Eber and D. Heath, 1997, "Thinking Coherently", *Risk* 10 (11), pp. 68-71.

Balassa, B., 1964, "The Purchasing-Power Parity Doctrine: A Reappraisal", *Journal of Political Economy*, 72 (6), pp. 584-96.

Brealey, R. A. and S. C. Myers, 1991, *Principles of Corporate Finance*, 4th edition, McGraw-Hill, Inc.

Bouman, S. and B. Jacobsen, 2002, "The Halloween Indicator, 'Sell in May and Go Away': Another Puzzle", *American Economic Review*, 92, pp. 1618-35.

Collin-Dufresne, P., R. S. Goldstein and J. Spencer Martin, 2001, "The Determinants of Credit Spread Changes", *Journal of Finance*, 56, pp. 2177-2207.

Dimson, E., P. Marsh and M. Staunton, 2002, *Triumph of the Optimists: 101 Years of Global Investment Returns*, Princeton University Press.

Elton, E. J., M. J. Gruber, D. Agrawal and C. Mann, 2001, "Explaining the Rate Spread on Corporate Bonds" *Journal of Finance*, 56, pp. 247-77.

Fabozzi, F. J., M. Pitts and R. E. Dattatreya, 1997, "Price Volatility Characteristics of Fixed Income Securities", in *The Handbook of Fixed Income Securities*, 5th edition, ed. by F. J. Fabozzi, McGraw-Hill, New York.

Froot, K. A. and K. Rogoff, 1995, "Perspectives on PPP and Long Run Real Exchange Rates", in *Handbook of International Economics*, Vol. 3, ed. by G. M. Grossman and K. Rogoff, Amsterdam: North Holland.

Grinold, R. C. and R. A. Meese, 2000, "Strategic Asset Allocation and International Investing", *Journal of Portfolio Management*, Fall 2000, pp. 53-60.

Heathcote, J. and F. Perri, 2003, "Why Has the U.S. Economy Become Less Correlated with the Rest of the World?", *American Economic Review, AEA Papers and Proceedings*, 93 (2), pp. 63-69.

Huang, J. and M. Huang, 2003, "How Much of the Corporate-Treasury Yield Spread is Due to Credit Risk?", Working paper.

Ilmanen, A. and R. Sayood, 2002, "Quantitative Forecasting Models and Active Diversification for International Bonds", *Journal of Fixed Income*, December 2002, pp. 40-51.

Jorion, P., 1997, *Value at Risk: The New Benchmark for Controlling Market Risk*, McGraw-Hill, New York, NY.

Jorion, P., 2003, *Financial Risk Manager Handbook*, John Wiley & Sons, Hoboken, New Jersey.

J. P. Morgan, 1995, *RiskMetrics – Technical Documentation*, Release 1-3, J. P. Morgan, New York, NY.

King, T. D. and K. Khang, 2002, "On the Cross-sectional and Time-series Relation between Firm Characteristics and Corporate Bond Yield Spreads", Working paper.

Kose, M. A., C. Otrok and C. H. Whiteman, 2003, "Understanding the Evolution of World Business Cycles", IMF Working Papers.

Markowitz, H., 1952, "Portfolio Selection", *Journal of Finance*, 7, pp. 77-91.

Markowitz, H., 1959, *Portfolio Selection: Efficient Diversification of Investment*, John Wiley & Sons, New York, NY.

Ross, S., 1976, "The Arbitrage Theory of Capital Asset Pricing", *Journal of Economic Theory*, 13, pp. 341-60.

Rosenberg, M. R., 1997, "International Fixed Income Investing: Theory and Practice", in *The Handbook of Fixed Income Securities*, 5th edition, ed. by F. J. Fabozzi, McGraw-Hill, New York.
Sharpe, W., 1964, "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk", *Journal of Finance*, 19, pp. 425-42.
Steward, C. B. and A. M. Greshin, 1997, "International Bond Investing and Portfolio Management", in *The Handbook of Fixed Income Securities*, 5th edition, ed. by F. J. Fabozzi, McGraw-Hill, New York.

The authors would like to thank Erik van Leeuwen and Jillert Blom from Robeco and an anonymous referee for their comments.

7

Currency reserve management by dual benchmark optimisation

Andreas Gintschel and Bernd Scherer

Abstract

We develop a framework for the currency allocation decision of central banks, and propose a dual benchmark optimisation problem in which the bank may attach different weights to a nominal wealth preservation benchmark and a liquidity benchmark. Furthermore, we allow for shifts in volatilities and correlations, and more generally deviations from normality in returns.

1 Currency allocation and central banks

Three problems dominate a central bank's management of currency reserves: determining the optimal level of reserves, the optimal mix of assets with different degrees of liquidity, and the optimal allocation among currencies. In this paper, we focus on the last problem.

Allocating among currencies is difficult because currencies are assets and numeraire at the same time. For example, the volatility of short-term US bonds, as an asset, depends on whether we view them from a European or a Japanese perspective, i.e. depending on the numeraire. The academic literature on currency allocation problems typically focuses on either one of the following optimisation problems:[1]

- **Wealth preservation approach.** Wealth held in foreign assets is volatile in terms of the numeraire (home) currency. Hence, positions are often diversified across many currencies to reduce risk.
- **Liquidity preservation approach.** To ensure international solvency, at least in the short run, central banks need sufficient reserves to cover payments for net imports and to service foreign debt payments. The closer the currency position is to the structure of these flows of payment, the less risky this is, even though the value of the position might be volatile in terms of the home currency.

Rather than simply following one approach in isolation, we model the currency allocation decision as a multi-objective optimisation problem, in which central banks may apply different weights to the various sub-problems and alternative risk regimes. We propose an objective function for central banks that captures the board's decision problem in an intuitive way. The board faces the usual trade-off between risk and return. Thus, we model the basic decision problem as mean-variance optimisation. However, we adjust the standard framework to take into account complications that arise for central banks, namely joint committee decision-making and a highly political environment. The board's problem is further exacerbated by uncertainty regarding the risk regime in which the bank operates.

The resulting framework is a highly practical solution for day-to-day management of reserves that takes into account political decision-making and the specifics of currency return distributions. The optimisation problem is set up in a simple, familiar and transparent framework, facilitating internal and external communication. Moreover, it can be readily implemented using standard statistical software packages, or appended to existing optimisation solutions. The framework makes explicit the trade-off between the two aims, wealth preservation and liquidity preservation. By varying the weight of these conflicting

[1] Currency return optimisation relative to various benchmarks has been addressed in Boorman and Ingves (2001), Ramaswamy (1999) and Scobie and Cagliesi (2000). Dual benchmark optimisation is covered in Scherer (2002).

aims in the optimisation problems, decision-makers can quantify the relative importance of these two aims. Thus, the framework can be interpreted as a management tool if the relative importance of the two aims is decided upon. The bank may also use the framework for decision support by investigating the consequences of attaching alternative weights to the aims.

The remainder of this paper is organised as follows. In the next section, we summarise the salient features of currency return data, and calculate the relevant covariance matrices and expected returns. In Section 3, we develop and apply a framework for allocating currencies in a central bank setting, taking into account the specific circumstances of committee decision-making. In Sections 4 and 5, we generalise the framework allowing state-dependent utility and non-normal returns.

2 Data

We use weekly return data (currency plus local cash return) from J. P. Morgan covering the period from January 1986 to December 2002 for the US dollar, UK pound, Japanese yen, Australian dollar and Swiss franc (all against the euro). Currency returns exhibit empirical features that are different from those of standard asset classes. These features have a potentially large impact on the risk of the currency portfolio. We do not provide a comprehensive review or argue for a preferred methodology. Instead, we focus on the two aspects of currency returns, time-varying second moments and non-normality, that are most relevant in the present context.

There are arguments that correlations break down in certain market conditions, i.e. when portfolio managers need them most. Here, we do not attempt to forecast the change in input parameters. Instead, we evaluate the diversifying properties of currencies in rivalling risk regimes. As policymakers and governing bodies are concerned mostly with short-term performance, central bankers often cannot afford to bet on average correlation or average volatility. Moreover, the highly political environment in which central bank policy is shaped suggests the need to focus on crises. Below, we develop a framework to tackle such issues formally. At this point, we merely document the risk regimes in returns and estimate the parameters for the optimisation program introduced later. To formalise the idea of alternative risk regimes, we distinguish between correlation and volatility during normal and hectic times.[2] We define unusual times according to the distance from the mean vector as given in

$$\left(\mathbf{R}_m - \bar{\boldsymbol{\mu}}\right)^T \bar{\boldsymbol{\Omega}}^{-1} \left(\mathbf{R}_m - \bar{\boldsymbol{\mu}}\right) = \mathbf{d}_m^T \bar{\boldsymbol{\Omega}}^{-1} \mathbf{d}_m , \tag{1.1}$$

where $\mathbf{d}_m$ is the distance vector at time m, $\mathbf{R}_m$ is a vector of currency return $\hat{\boldsymbol{\Omega}}$ observations for N currencies at time m, $\bar{\boldsymbol{\mu}}$ denotes a vector of average currency returns, and is the unconditional covariance matrix (over all $m = 1$, L, M observations).

For each cross-section of returns, we calculate (1.1) and compare it to the critical value of a $\chi^2_{0.95}(N)$. Defining an unusual observation as the outer 5% of a distribution (alternatively, it could be termed an outlier) for five return series, the critical distance is 11.07. In (1.1) the

[2] See Chow et al. (1999). Longin and Solnik (2001) propose an improved estimator for time-varying correlations.

return distance is weighted by the inverse of the covariance matrix. Thus, we take into account individual currencies' volatilities and correlations when defining hectic times. Outliers are not necessarily associated with down markets, although this is often the case in practice. We estimate covariance matrices

$$\mathbf{\Omega}_s = \boldsymbol{\sigma}_s \boldsymbol{\rho}_s \boldsymbol{\sigma}$$

where *s* = *normal, hectic*. Effectively, we split the covariance matrix into a high volatility regime and a normal volatility regime. We find that currency return correlations $\mathbf{r}_s$ are virtually unchanged across risk regimes. Thus, international currency allocation remains attractive for diversification purposes even in times of market crisis. More interesting results, from a risk management perspective, emerge for volatilities, which we report in Table 1 below:

Table 1: Volatility in normal and hectic times

$$\boldsymbol{\sigma}_{normal} = \begin{pmatrix} 9{,}3 &&&&& \\ & 6{,}6 &&& 0 & \\ && 9{,}8 &&& \\ &&& 11{,}9 && \\ & 0 &&& 3{,}6 & \\ &&&&& 0{,}4 \end{pmatrix} \boldsymbol{\sigma}_{hectic} = \begin{pmatrix} 20{,}6 &&&&& \\ & 14{,}6 &&& 0 & \\ && 20{,}1 &&& \\ &&& 26{,}4 && \\ & 0 &&& 7{,}8 & \\ &&&&& 0{,}4 \end{pmatrix}$$

By construction, currency risk on average increases substantially as we move from normal to hectic times. More interestingly, volatility roughly doubles across all currencies simultaneously. Thus, currency crises tend to be global, affecting all currencies at the same time. Therefore, singling out crises as special events is important for risk management, despite stable correlations.

Further investigation of currency returns reveals occasional skewness and substantial leptokurtosis. Thus, currency returns are generally not very well approximated by normal distributions. In later sections, we show how to adapt the basic optimisation framework to deal with substantial deviations from normality.

Finally, we need estimates of expected currency returns. We use the James-Stein (JS) estimator to reduce the error in the estimates of the means, reducing the tendency for corner solutions, i.e. when asset weights are determined by constraints. The JS estimator exploits the notion of pooling the information in all series improves the efficiency of the historical mean estimate for an individual return series. Effectively, the JS estimator shrinks the individual historical means towards the grand mean of all time series. In the case of complete shrinkage, we estimate means as equal to the grand mean, implying that the optimal allocation is the minimum variance portfolio. The JS estimator is

$$\boldsymbol{\mu} = \phi\bar{\mu}\mathbf{I} + (1-\phi)\bar{\boldsymbol{\mu}} \tag{1.2}$$

where denotes the grand mean of all series $\bar{\mu} = \sum_{n=1}^{N}\mu_n$, $\mathbf{I}$ represents a vector of ones and $\bar{\boldsymbol{\mu}}$ reflects the vector of historical means. The shrinkage factor is

$$\phi = \min\left[1, \frac{(N-2)}{M(\bar{\boldsymbol{\mu}} - \bar{\mu}\mathbf{I})^T \mathbf{\Omega}^{-1} (\bar{\boldsymbol{\mu}} - \bar{\mu}\mathbf{I})}\right].$$

If $(\bar{\boldsymbol{\mu}} - \bar{\mu}\mathrm{I})^T\, \Omega^{-1}\, (\bar{\boldsymbol{\mu}} - \bar{\mu}\mathrm{I})$ is small, either because the distance between the historical means and the grand mean is small or because the precision of the estimates is low, the JS estimator shades the individual estimates towards the grand mean. The same is true if there are many assets or few observations.

3 Multiple benchmarks and risk regimes

Committee decisions are notoriously difficult to model and predict; there is no generally accepted solution to multilateral bargaining models. In a pragmatic approach, we suggest modelling the committee decision as a single objective function similar to the standard expected utility function for individuals. However, we adapt the standard function to take into account the problems of joint decision-making.

Most central bank boards operate in a highly political environment and are under public scrutiny, which tends to intensify during economic crises. In these circumstances, decision-makers typically try to avoid being caught on the wrong foot, i.e. ending up with an outcome that is ex ante improbable or bad. Thus, we believe that a suitable objective function for central bank boards must contain a minmax element, safeguarding against the wrong policy objective, i.e. the wrong benchmark, during crises.

Moreover, committee members often disagree considerably about the probability of various scenarios. Without any further information about, or agreement on, the forecasting powers of the committee members, it is impossible to arrive at a common, committee-wide probability assessment. In such circumstances, a minmax approach can be seen as a compromise to focus on bad outcomes irrespective of subjective probabilities. Of course, such an approach is likely to lead to committee decisions that each member feels to be overly cautious. However, in political situations committee members do not have an incentive to oppose conservative policies. If, upon an individual committee member's intervention, a more aggressive policy were to be adopted, that committee member is likely to be singled out for blame in case of an unfortunate outcome, but will have to share the praise should the outcome be positive.

Formally, we model this situation as the central bank facing $s = 1\mathrm{L}\ S$ risk regimes, reflected in the associated currency covariance matrices $\boldsymbol{\Omega}_s$, and $b = 1\mathrm{L}\ B$ currency benchmarks. The central bank may invest in $n = 1\mathrm{L}\ N$, assets summarised by the vector of currency holdings, $\mathbf{w}$. The optimisation problem is

$$\max_{\mathbf{w},\mathbf{w}\geq 0}\left(\mathbf{w}^T\boldsymbol{\mu} - \lambda \max_{s,b}\left[\left(\mathbf{w}-\mathbf{w}_b\right)^T \boldsymbol{\Omega}_s\left(\mathbf{w}-\mathbf{w}_b\right)\right]\right)$$

where λ reflects the central bank's risk aversion. The risk aversion parameter depends on the joint preferences of the bank's board over the outcomes. Given the similarity, we suggest using the intuition from standard mean-variance utility about the magnitude of the risk aversion parameter. Traditionally, values between 1 and 30 are regarded as a reasonable range, for which we show calculations.

The optimisation program (1.4) provides maximum protection against the risk of adopting an investment strategy based on the wrong benchmark or the wrong risk regime. We can rewrite as a constrained quadratic program[3]

$$\max_{\mathbf{w},\sigma^2_{\max}} \left(\mu - \lambda\sigma^2_{\max}\right)$$
$$\left(\mathbf{w}-\mathbf{w}_b\right)^T \mathbf{\Omega}_s \left(\mathbf{w}-\mathbf{w}_b\right) \le \sigma^2_{\max}$$
$$\mathbf{w}^T\mathbf{I} = 1$$
$$\mathbf{w}^T\boldsymbol{\mu} = \mu$$
$$\mathbf{w} \ge 0$$

If there is only one risk regime and one benchmark, this program is a standard Markowitz problem. In the remainder of the paper, we consider two benchmarks, namely a liquidity and a wealth preservation target, and two risk regimes. As examples of liquidity and wealth benchmarks, we use

$$w_{wealth} = (0\ \ 0\ \ 0\ \ 0\ \ 0\ \ 1)^T$$
$$w_{liquidity} = (0{,}5\ \ 0\ \ 0{,}5\ \ 0\ \ 0\ \ 0)^T.$$

The wealth preservation benchmark consists of 100% euro cash, while the liquidity benchmark contains 50% US dollar and 50% Japanese yen. These figures are hypothetical examples, for which we investigate the effect of two benchmarks on portfolio choice. Starting with a portfolio of 100% US cash and gradually moving allocations into euro cash, we study the effect on benchmark relative risk. For example, a central bank with the above benchmark definitions and risk regimes may choose a 20% allocation in the US dollar. The relative riskiness varies between 1.87% under a wealth benchmark in normal times and 13.99% under a liquidity benchmark in hectic times, as we show in Figure 1. Board members might argue

Figure 1: Relative risks under alternative benchmarks and risk regimes
(in %)

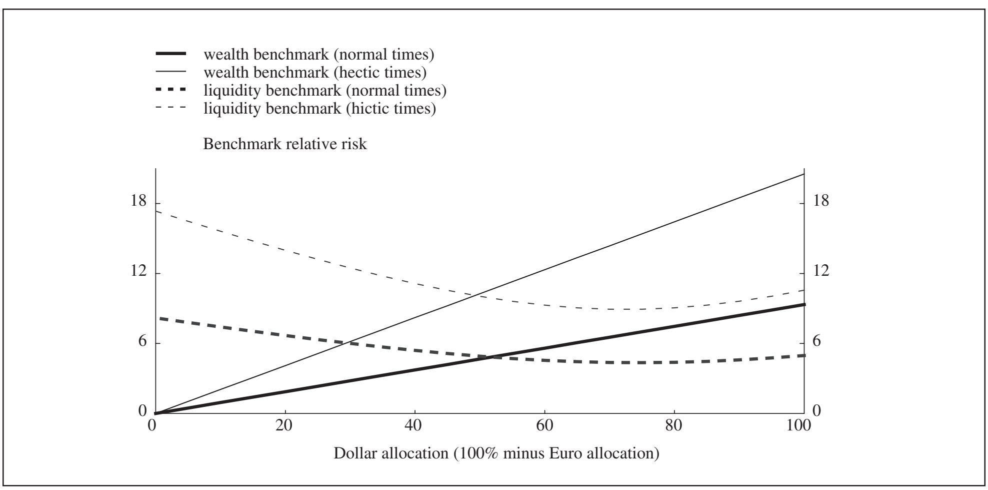

[3] All optimisations have been performed using NUOPT for S-Plus. See Scherer and Martin (2003) for an extensive treatment of portfolio optimisation problems in S-Plus.

about the riskiness of any given strategy depending on the weight a member puts on the two benchmarks and risk regimes.

Hence, any regrets at having decided to manage against the wrong benchmark in the wrong risk regime are potentially large. However, allocating 50% to US dollars and euro significantly narrows the range of potential outcomes to between 4.67% and 10.27%. Compared to a 20%-80% allocation, the worst case risk drops from 13.99% to 10.275%.

So far we have discussed risk sensitivity for a simple two currency portfolio; now we proceed to solve . The optimisation results are contained in Table 2 through Table 4 below:

Table 2: Optimal portfolios with dual benchmark and rivalling risk regimes
(in %)

Currency	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9	p_{10}
USD	0.00	25.61	19.07	12.09	5.13	0.66	0.02	0.00	0.00	0.00
Pound	0.00	0.00	6.95	15.16	23.35	32.78	44.05	55.63	67.22	100.00
Yen	0.00	23.91	24.65	24.24	23.83	21.90	17.67	13.05	8.43	0.00
AUSD	0.00	0.00	4.33	9.43	14.51	18.35	20.29	21.91	23.52	0.00
FRK	100.00	10.65	0.07	0.09	0.02	0.05	0.00	0.00	0.00	0.00
EUR	0.00	39.84	44.93	38.99	33.15	26.26	17.98	9.41	0.83	0.00
Risk	18.33	8.71	8.76	9.07	9.58	10.28	11.21	12.36	13.67	19.22
Return	-0.42	-0.13	0.17	0.46	0.76	1.05	1.35	1.64	1.94	2.23

Table 2 contains the allocations resulting from program (1.4) **for l = 2**. The columns of the table contain allocations for different target returns, which are reported in the last row. For comparison, we also report the allocations resulting from maximising (1.4) just for normal times (Table 3), and then only for hectic times (Table 4). In each case, we pretend there is only one possible risk regime, either normal or hectic. In all three tables, we use dual benchmarks for liquidity and wealth preservation aims.

Not surprisingly given the minimum regret preferences, the resulting allocations in Table 2 are heavily influenced by the worst-case risk scenario (Table 4). Hence, the solutions in which we use both regimes (Table 2) are virtually identical to the solutions in which we use only the hectic regimes (Table 4). Naturally, the risk under both risk regimes is also considerably higher than during normal times. Comparing these results with allocations in which we only use the covariance matrix for normal times (Table 3), the differences are probably smaller

Table 3. Optimal portfolios with dual benchmark and single risk regime (normal times)
(in %)

Currency	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9	p_{10}
USD	0.00	25.90	19.87	14.02	8.08	2.47	0.12	0.00	0.01	0.00
Pound	0.00	0.04	8.29	18.38	28.35	38.43	49.83	62.14	74.49	100.00
Yen	0.00	23.97	24.29	23.50	22.65	21.64	18.26	13.28	8.20	0.00
AUSD	0.00	0.06	2.84	5.78	8.94	11.89	13.54	14.28	14.97	0.00
FRK	100.00	10.95	0.10	0.01	0.26	0.00	0.00	0.00	0.00	0.00
EUR	0.00	39.08	44.62	38.31	31.72	25.57	18.25	10.31	2.34	0.00
Risk	8.88	4.11	4.14	4.28	4.53	4.85	5.25	5.75	6.32	7.85
Return	-0.42	-0.13	0.17	0.46	0.76	1.05	1.35	1.64	1.94	2.23

Table 4: Optimal portfolios with dual benchmark and single risk regime (hectic times) *(in %)*

Currency	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9	p_{10}
USD	0.00	25.61	19.07	12.10	5.14	0.03	0.01	0.00	0.00	0.00
Pound	0.00	0.00	6.95	15.15	23.36	32.46	44.04	55.63	67.22	100.00
Yen	0.00	23.91	24.65	24.25	23.82	22.29	17.68	13.05	8.42	0.00
AUSD	0.00	0.00	4.34	9.42	14.50	18.66	20.29	21.91	23.53	0.00
FRK	100.00	10.65	0.10	0.04	0.03	0.00	0.00	0.00	0.00	0.00
EUR	0.00	39.84	44.90	39.05	33.14	26.55	17.98	9.41	0.83	0.00
Risk	18.33	8.71	8.76	9.07	9.58	10.27	11.21	12.36	13.67	19.21
Return	-0.42	-0.13	0.17	0.46	0.76	1.05	1.35	1.64	1.94	2.23

than expected. However, since correlations hardly change and all volatilities increase by roughly the same factor, allocations do not differ much across volatility regimes. Only the Australian dollar attracts a larger weight because of the improved risk-return trade-off. The minimum risk portfolio remains unchanged under all optimisations, further supporting our observation that the relative risk is stable.

4 State-dependent preferences

Now we allow for varying risk preferences, effectively assigning different weights to various sub-problems.[4] State-dependent preferences capture the idea of central bank boards being exposed to different levels of political pressure and public scrutiny during crises. In a further generalisation, we also allow risk aversion coefficients to differ across benchmarks. The ratio of the risk aversion coefficients reflects the committee's views on the relative importance of both aims, liquidity and wealth preservation. The average of the two coefficients reflects the overall risk aversion. Thus, the general formulation of the problem is

$$\max_{\mathbf{w},\mathbf{w}\geq 0}\left(\min_{s,b}\left[\mathbf{w}^T\boldsymbol{\mu}-\lambda_{s,b}\left(\mathbf{w}-\mathbf{w}_b\right)^T\boldsymbol{\Omega}_s\left(\mathbf{w}-\mathbf{w}_b\right)\right]\right) \tag{1.5}$$

Equation (1.5) poses the problem of maximising the minimum risk-adjusted performance (utility) across the alternative benchmarks, risk regimes, and the associated risk aversion coefficients. Under mean variance preferences, this is equivalent to maximising the minimum utility. The resulting solution is Pareto optimal in the sense that it is impossible to increase utility any further without decreasing another sub-problem's utility below the minimum level. Equation (1.5) differs from the conventional treatment of multiple benchmark problems[5], which can be expressed

$$\max_{\mathbf{w},\mathbf{w}\geq 0}\left(\mathbf{w}^T\mathbf{R}-\lambda_1\left(\mathbf{w}-\mathbf{w}_1\right)^T\boldsymbol{\Omega}\left(\mathbf{w}-\mathbf{w}_1\right)-\lambda_2\left(\mathbf{w}-\mathbf{w}_2\right)^T\boldsymbol{\Omega}\left(\mathbf{w}-\mathbf{w}_2\right)-\mathrm{L}\right) \tag{1.6}$$

4 See Shectman (2000) on Pareto optimality and dual benchmark optimisation.
5 See Wang (1999) for the use of standard portfolio optimisers in multiple benchmark optimisation.

Table 5: Pareto optimality with dual benchmarks and dual risk regimes
(in %)

	$\lambda_2 = 1$	$\lambda_2 = 3$	$\lambda_2 = 10$	$\lambda_2 = 30$	
	0	7	25	35	USD
	53	32	17	9	Pound
$\lambda_1 = 1$	12	32	39	43	Yen
	21	20	10	6	AUSD
	0	0	0	0	FRK
	14	9	9	7	EUR
	0	8	23	33	USD
	29	20	11	6	Pound
$\lambda_1 = 3$	10	24	32	38	Yen
	13	12	7	4	AUSD
	0	0	0	0	FRK
	16	0	0	1	EUR
	0	8	20	29	USD
	15	11	6	3	Pound
$\lambda_1 = 10$	8	17	25	31	Yen
	8	7	4	2	AUSD
	0	0	0	0	FRK
	69	58	46	34	EUR
	0	6	15	23	USD
	9	6	3	2	Pound
$\lambda_1 = 30$	5	11	18	25	Yen
	5	4	2	1	AUSD
	0	0	0	0	FRK
	81	72	61	48	EUR

We examine the same benchmarks and risk regimes as in the previous section. The resulting allocations vary according to the risk aversion parameters attached to both benchmarks. For purposes of exposition, we restrict attention here to weighing benchmarks, keeping constant risk aversion across volatility regimes. The results are contained in Table 5. When decision-makers have low risk aversion ($\lambda_s = 1$), the solutions focus on assets with the highest returns. As the liquidity benchmark becomes more important (higher penalty term for relative risk), allocations approach this benchmark (assumed to be 50% USD and 50% yen). If both risk aversion parameters equal 30, i.e. an extremely risk-averse board, we arrive at an intermediate solution investing in the equally weighted benchmark. At the other extreme, if central bankers put an overwhelming emphasis on the avoidance of wealth benchmark-relative risk, we naturally come close to this benchmark (81% weighting in local cash). Obviously this is a solution few central banks would adopt. Interestingly, the case of extreme risk aversion with regard to both benchmarks is relatively close to the minimum variance solution. The optimal solution depends on the degree of risk aversion towards the benchmark relative risk. Hence, the proposed optimisation framework offers a straightforward technique to arrive at the optimal currency allocation if risk aversion differs with respect to benchmark-relative risk. To assess the importance of the rivalling risk regimes, we rerun the optimisation above under only one risk regime, namely normal times. The results can be found in Table 6. When risk is considerably low (normal times), it is optimal to allocate to the maximum return currency (99% in UK pounds for the used data set, time period, and currency perspective). However, for high risk aversion the results are almost the same as in Table 5. If risk aversion

Table 6: Pareto optimality with dual benchmarks and single (lnormal) risk regime *(in %)*

	$\lambda_2 = 1$	$\lambda_2 = 3$	$\lambda_2 = 10$	$\lambda_2 = 30$	
	0	0	20	35	USD
	99	52	27	17	Pound
$\lambda_1 = 1$	0	20	39	43	Yen
	0	28	14	8	AUSD
	0	0	0	0	FRK
	0	0	0	0	EUR
	0	0	8	25	USD
	82	85	44	27	Pound
$\lambda_1 = 3$	0	0	32	39	Yen
	2	15	16	9	AUSD
	0	0	0	0	FRK
	16	0	0	1	EUR
	0	0	5	21	USD
	28	42	32	18	Pound
$\lambda_1 = 10$	0	2	22	31	Yen
	1	8	10	6	AUSD
	0	0	0	0	FRK
	7	48	31	24	EUR
	0	0	6	18	USD
	16	23	18	11	Pound
$\lambda_1 = 30$	0	4	16	24	Yen
	1	5	6	3	AUSD
	0	0	0	0	FRK
	83	68	54	44	EUR

is sufficiently high, the level of risk does not matter. In general, there is a clear tendency to allocate more to the high return asset classes. All other results remain qualitatively the same.

5 Pareto optimality and non-normality

In this section, we incorporate non-normality of returns into the optimisation framework. Under non-symmetric return distributions, the mean-variance criterion is not appropriate. Instead of deriving the expected utility function in closed form, i.e. in terms of the distribution's moments, we calculate the expected utility directly using scenario optimisation. We assume that investors' preferences may be represented by a hyperbolic utility function $U_m^j = \frac{1}{1-\gamma_j} W_m^{1-\gamma_j}$, where $W_m = \sum_{n=1}^{N} w_n \left(1 + R_{mn}\right)$ is terminal wealth relative to the relevant benchmark in scenario m = 1L M, R_{mn} is the return of asset n in scenario m, and ã$_j$ is the constant coefficient of risk aversion for benchmark b. Under weak assumptions on the return distribution, problem (1.5) becomes

$$\max_{\mathbf{w},\mathbf{w}\geq 0}\left(\min_{j}\left[\frac{1}{m}\sum_{m=1}^{M} U_m^j\right]\right), \qquad (1.7)$$

which is equivalent to maximising the minimum expected utility. Bootstrapping, for example, scenarios m = 1L M, equation (1.7) is easily solved by standard numerical optimisation techniques.

Table 7: Pareto optimality under normality

(in %)

	$\gamma_2 = 3$	$\gamma_2 = 10$	$\gamma_2 = 30$	
	0	0	27	USD
	82	54	20	Pound
$\gamma_1 = 3$	0	22	42	Yen
	18	24	11	AUSD
	0	0	0	FRK
	0	0	0	EUR
	0	1	22	USD
	22	45	23	Pound
$\gamma_1 = 10$	1	19	36	Yen
	7	15	10	AUSD
	0	0	0	FRK
	61	21	9	EUR
	1	0	14	USD
	22	22	15	Pound
$\gamma_1 = 30$	1	10	24	Yen
	4	7	7	AUSD
	1	0	0	FRK
	72	61	40	EUR

Under the hyperbolic preferences, currencies exhibiting positive skewness (yen) are generally favoured relative to currencies that show negative skewness (UK pound). The results are shown in Tables 7 and 8. We see that the yen (positively skewed) is favoured versus the pound and Australian dollar (both negatively skewed) relative to the solution assuming normality. For extreme risk aversion, the two solutions more or less coincide.

Table 8: Pareto optimality under non-normality

(in %)

	$\gamma_2 = 3$	$\gamma_2 = 10$	$\gamma_2 = 30$	
	0	5	32	USD
	18	19	7	Pound
$\gamma_1 = 3$	71	63	55	Yen
	11	13	5	AUSD
	0	0	0	FRK
	0	0	0	EUR
	1	3	25	USD
	26	23	12	Pound
$\gamma_1 = 10$	31	43	47	Yen
	5	8	5	AUSD
	0	0	0	FRK
	38	23	11	EUR
	0	2	17	USD
	10	11	8	Pound
$\gamma_1 = 30$	12	21	32	Yen
	1	4	3	AUSD
	0	0	0	FRK
	76	62	41	EUR

6 Summary

We model the optimal currency allocation decision as a multiple benchmark optimisation problem. In contrast to traditional solutions, we show how a maximin approach can be successfully applied to the currency allocation problem, including alternative risk aversion, risk regimes, and benchmarks. We extend the above analysis to incorporate non-normality in return data using scenario optimisation as the most general form of portfolio optimisation. Our methodology is equally applicable to central banks of both developed and developing countries, because it is very general in specifying benchmarks, the relative importance of benchmark-relative risks, as well as risk regimes and distributional assumptions. The definition of these parameters determines the optimal solution for any given central bank.

References

Boorman J. and S. Ingwes (2001), *Issues in Reserve Adequacy Management*, IMF Working Paper

Chow G., Jacquier E., Kritzman M. and K. Lowry (1999), "Optimal Portfolios in Good Times and Bad", *Financial Analysts Journal*, 55, pp. 65-73

Longin F. and B. Solnik (2001), "Extreme Correlation of International Equity Markets", *Journal of Finance*, 56, pp. 649–76.

Ramaswamy S. (1999), *Reserve Currency Allocation: An Alternative Methodology*, BIS Working Paper No. 72

Scherer B. (2002), *Risk Budgeting and Portfolio Construction* (Riskwaters: London)

Scobie H. and G. Cagliesi (2000), *Reserve Management* (Riskwaters: London)

Shectman P. (2000), "Multiple Benchmarks and Multiple Sources of Risk", Working Paper, Northfields

Wang M. (1999), "Multiple-Benchmark and Multiple Portfolio Optimization", *Financial Analysts Journal*, 55, pp. 63-72

2 SPECIFICS OF RISK MEASUREMENT AND MANAGEMENT

8

Risk systems in central bank reserves management

Mark Dwyer and John Nugée

Abstract

For central banks, because of their reputation and emphasis on stability, risk control assumes a greater prominence than for many other institutions active in the financial sphere. At stake is not just the central bank's balance sheet, but also its reputation and its ability to command the respect and hence cooperation of the market. This paper considers risk in reserves management, both from a financial risk perspective, but also as part of the wider framework of central bank risk control and corporate governance, including legal risk, operational risk, reputational risk and business risk, as well as the more easily identifiable and quantifiable financial risks. This requires a "whole organisation" approach to risk management. The paper concludes that risk management in reserves management is an integral part of and must be consistent with the overall risk management of the central bank.

1 Introduction

Central banks are well known for their conservative, risk-averse approach to all aspects of their operations. This is a natural consequence of their primary function, which is the pursuit of financial stability. Although different central banks have differing objectives and emphases, these objectives generally revolve around financial stability; in particular, the stability of money and the stability of the financial system. For a central bank to succeed in delivering this stability, it is very important that it earns, and then retains, the confidence of other market operators as a secure, stable and reliable institution itself.

Because of this emphasis throughout the organisation on stability and reliability, risk control assumes a greater prominence for central banks than for many other institutions active in the financial sphere. At stake is not just the central bank's balance sheet, but also its reputation and its ability to continue to command the respect and cooperation of the market. Without this respect and reputation for reliability, a central bank may find it very difficult to continue to operate effectively, and may as a result find that it cannot succeed in fulfilling its overall objectives.

This paper considers risk management in the reserves management operation both from a narrow financial risk control perspective, and also in the context of the wider framework of risk control and corporate governance for the whole bank. This wider framework must cover all aspects of the central bank's operations, rather than merely its financial activities, and will therefore include legal risk, operational risk, reputational risk and business risk as well as the more easily identifiable and quantifiable financial risks. This requires a "whole organisation" approach to risk management.

The paper is constructed as follows. We start by examining the financial risks that a central bank faces in its reserves management operation, then briefly describe the traditional risk measures used, before examining the rise of Value-at-Risk techniques and their increasing use by central banks. We explore shortcomings in the traditional approach, and consider modern solutions to these concerns, all of which are extremely IT-dependent. We conclude this first section of the paper by looking at some of the operational issues connected with the large IT-based risk management solutions currently employed by central banks.

The second section of the paper broadens the discussion to examine non-financial risks. We consider how non-financial risks can be identified and controlled, and how they impact upon the reserves management operation. Finally, we bring the two parts of the paper together by introducing the concept of "whole enterprise" risk management. We conclude that, far from being a separate subject, risk management in reserves management is entirely consistent with the overall style of risk management in a central bank, which it forms an integral part of.

2 Financial risk in reserves management

As stated in the Introduction, central banks face a variety of risks in their operations, both financial and non-financial. Non-financial risks are best controlled by the adoption of appropriate management procedures and checks, and the naturally risk-averse nature of central bank senior management is well suited to these.

However, financial risks are too complex, and in many cases too fast-moving, to be left entirely to the natural conservatism of the central bank ethos. Moreover, the visibility of any financial losses means that the public is more likely to focus on them should they occur, while the quantifiable nature of the risks means that technological solutions to prevent this are more appropriate. For these reasons, central banks will typically expend considerable energies controlling these risks, with highly complex risk management systems and detailed risk control procedures.

This is particularly the case in the field of reserves management. In many areas of a central bank's operations, it has a privileged position, and can control the structure of an operation so as to minimise its own risk. For example, a central bank can control the risks it faces in its role as the operator of a payments system, by so constructing the rules concerning collateral for liquidity that it is never exposed. Similarly, a central bank can avoid daylight exposure to banks in its money market operations by demanding that they keep positive balances in their accounts at the central bank. And in its operations as the manager of public debt issues, a central bank can avoid underwriting risk (the risk that it is left with unsold stock) by requiring a panel of market makers or primary dealers to "cover the auction" – i.e. collectively bid for at least as much stock as is on issue.

In its reserves management operations, however, a central bank does not usually occupy a privileged position, and is therefore exposed to the same financial risks as other market participants. All market participants must assess for themselves the level of risk that they are able to bear, and the price they are prepared to pay to reduce those risks. For most institutions, where the end objective is the pursuit of financial return, this is an exercise in comparing risks (i.e. financial losses) and rewards (i.e. financial gains), which therefore makes the judgment mainly a financial one. If the rewards for an operation do not justify the risk, the institution will usually withdraw from that operation.

For a central bank, however, the end objective of its operations is usually the pursuit of a chosen policy (for example, maintaining the exchange rate at a certain level), rather than merely financial gain. In this case, it may be unwilling or unable to withdraw from an operation, even if a simple financial cost-benefit analysis suggests that it is unfavourable. The very act of holding net foreign exchange reserves entails an unavoidable currency risk, as the liabilities that fund the net reserves will be denominated in domestic currency. This risk cannot be reduced without reducing the holdings of the reserves themselves. Similarly, other risks (for example credit risk, liquidity risk, interest rate risk, currency risk and cross-country spread risks) constitute an integral part of holding reserve assets, and cannot be wholly avoided.

The 4 Stages of Risk Management	
	1: Ignored
	2: Calculated
	3: Avoided
	4: Controlled

Central bank responses to the financial risks in reserves management typically follow a well-travelled path, commonly-known by the acronym "ICAC". In the initial stage of reserves management, risks are largely *ignored*. They are neither quantified nor, in some cases, even recognised.

While this idyllic state of ignorance may suffice for the simplest operations (and certainly allows reserves managers to sleep more easily), it is not suitable for anything more complicated than the most basic reserves management, and sooner or later, senior management realise that risks cannot continue to be ignored. The second stage therefore is that risks are *calculated*, and quantified for the first time.

The third stage follows on from the second, usually fairly swiftly. When senior management become aware of the risks being run and see them quantified for the first time, their natural reaction is to seek to *avoid* them.

Finally, the central bank realises that avoiding all risks is neither sensible nor even possible, thus leading to the fourth stage, where risks are *controlled*. The control of risk is essential to preserve the central bank's efficacy and wider reputation; moreover, the *controlled use* of risk in reserves management is necessary for any central bank seeking to add excess return to its reserves.

Most central banks pass through these four stages of attitude to risk. They may do so at different speeds, and they may spend some time at each stage of their risk development. Few move straight to Stage 4, but equally, even fewer will find that a permanent decision to halt before they reach Stage 4 is in their best long-term interest. This leads to a general acceptance of the current best market practice among central bank reserves managers, which can be summarised as placing emphasis on risk control rather than risk avoidance, and stressing a formal framework within which this risk control operates.

3 Using risk

The previous section introduced not only the concept of the *control* of risk, but also of the *controlled use* of risk. To understand further how these two concepts interact, it is helpful to consider the options open to a reserves management portfolio manager when he or she is considering the portfolio under management.

Any trade that a portfolio manager completes on his or her portfolio will either increase or decrease the riskiness of the portfolio, and will also increase or decrease the expected return on that portfolio. This then leads to four possible outcomes of any trade (ignoring those which leave either or both of risk and return unchanged). These are shown diagrammatically below:

A: Return increases, risk decreases	B: Return and risk both increase
C: Return and risk both decrease	D: Return decreases, risk increases

With this framework, three types of trade can be identified. Trades of the first type are those that fall into Quadrant A, with increased expected return and reduced risk. These are known as *Portfolio improvement* trades; the portfolio after the trade is superior in a risk-return sense to its position before the trade. If the portfolio manager can find such a trade, it should always be done, provided it is compatible with the portfolio's other constraints.

Trades of the second type are those that fall into Quadrant D, with decreased expected return and increased risk. These are the exact opposite of trades in Quadrant A, and should never be done unless they are intended to fulfil a wider or more important policy decision.

Finally there are trades in Quadrants B and C, where risk and expected return are either both increased or both decreased. These trades are the essence of active management, and require judgment from the portfolio manager: the decision is whether the extra return justifies the extra risk, and so on.

In the 1980s the suitability of trades of types B and C was a debating point among central bank reserves managers. Some maintained that it was wrong for a central bank to seek to manage its reserves actively, and wrong for the reserves management operation to have the pursuit of return as one of its objectives. For the adherents of this approach, type B trades in particular, where extra risk is taken on in the pursuit or expectation of extra return, were deemed unsuitable for a central bank. Instead, they insisted that the role of central bank reserves managers was solely to focus on reducing risk. More recently, however, this debate appears to have been conclusively resolved in favour of the opposing approach, namely accepting the role of return enhancement in central bank reserves management, and few central banks these days stick to a pure "risk minimisation at all costs" strategy. Indeed, reserves management is one of the few areas of central bank operations where extra risk is voluntarily taken on in order to achieve a policy objective. It is the voluntary nature of the assumption of extra risk that highlights the importance of risk management for central banks.

4 Traditional risk control in reserves management

Traditionally, central bank risk control of fixed income portfolios was fairly simple. The tools at the disposal of management were straightforward, and the main question facing senior management was how to set appropriate numerical limits (i.e. setting an upper bound on the *amount* of risk being run) so as to be commensurate with their risk appetite. Typical measures used were nominal differences from a benchmark for currency and credit exposures, with additional duration measures used to control interest rate exposures.

For a central bank whose reserves management operation is characterised by infrequent deals and limited to simple instruments such as straight fixed income bonds, this approach is eminently suitable. There is little need for a highly complicated risk management structure and much advantage to be gained from simplicity: the simpler the risk system, the less likely it is to be misunderstood by portfolio managers and senior management. This simplicity also means that IT support for the portfolio managers can be kept to a minimum.

However, a more active central bank that includes more complex instruments such as derivatives in its reserves management operation will wish to consider more sophisticated measures of risk. Three of the main questions a central bank will need to consider are:

- the handling of complex positions, e.g. cross-market or cross-currency;
- the handling of non-linear risk;
- the frequency of risk measurement and analysis.

When faced with complex positions or non-linear risk, the traditional measures suffer from a number of drawbacks and limitations. Firstly, they are all static, whereas fund management takes place in a moving environment. This can to a certain extent be overcome with simulations and what-if analyses, but the quality of the information obtained from such exercises is very reliant on the quality of the forecasts fed into them. In addition, the assumption that the investor would hold his or her portfolio unchanged as various scenarios unwind has always been somewhat unrealistic.

Secondly, the risk measures are absolute, whereas markets move between calmer and more volatile phases. A position which is justifiable in calm markets might be too risky in more turbulent times. Traditional risk control methods, in which management for example lay down how much a portfolio may vary from a preset benchmark, struggle to respond adequately to varying market conditions, and the danger is that in order to avoid allowing too much risk in difficult markets, management set limits so tight that no worthwhile positions can be taken even in more favourable conditions.

Thirdly, traditional measures can fail to quantify fully the risks even in relatively simple portfolios. For example, a position short \$20 million 4-year bonds and long \$10 million 10-year bonds has no PV01 risk[1] (i.e. it will not gain or lose value on a general change in the level of the yield curve). However, it is nevertheless exposed to changes in the *slope* of the curve. Similarly, a position short \$20 million 4-year government bonds and long \$20 million 4-year bonds issued by another issuer (e.g. an agency) also has no PV01 risk, yet is not without risk as it is exposed to *spread* risk (i.e. the difference between yield levels on government bonds and on the other issuer's bonds).

Lastly, the traditional techniques struggle to handle adequately newer instruments such as derivatives. Even before the explosion of derivatives in the last ten years or so, such basic and well-established investments as callable bonds (i.e. bonds with an embedded option) posed problems for the more traditional measures such as duration. A fall in general yield levels which results in a callable bond being more likely to be called will markedly shorten the duration of any holding of that bond[2].

5 Consequences of complexity

As markets have developed, with new derivative constructions added to the underlying base instruments, the control and management of risk has become more complex. This is not only attributable to the development of liquid markets in derivative products, but also to the increase in the correlation relationship between asset classes and markets globally. For central banks as for other investors, these new instruments cannot be ignored, and in some markets, liquidity is migrating from basic cash instruments to the derivatives markets based on them. Government bonds are just one example of this, as is the swap market, now a massive global, liquid market that is considered to be essential to managing short-term interest rate exposures.

As central banks have adapted their reserves management techniques to include these new security types, they have to manage their risks in order to manage their portfolios effectively. The typical response to the inclusion of a new security type is to purchase or build a new valuation and risk system for this particular security. A specialist system purchase is often the

1 The PV01 of a bond holding, or "Price value of an 01", is the amount that the price of the bond will change for a 1 basis point (i.e. an "01") change in the bond's yield. The greater the PV01, the greater the sensitivity of the bond to yield changes, and thus the greater the risk in the bond holding.

2 For a more detailed discussion of this subject, see Nugée (2000).

response when a financial institution, irrespective of whether it is a central bank or not, is required for efficiency or portfolio performance reasons to change the investment process. This leads to the incremental building of risk systems, which has many consequences.

Methodology

It is natural, when adding an asset class, to seek the risk system that can optimally handle its particular characteristics. However, using several methodologies simultaneously over the reserves management operation considerably increases the complexity of calculating the exposure between different asset classes. Another consequence of modelling risk for each asset class separately is that there is no common reference or understanding of risk between portfolios or departments.

Data

Multiple risk systems usually result in different databases for each asset class. Even data common to several asset class exposures, such as the price of securities, interest and currency rates, have a different source and may be taken at a different time. Attaining consistency between data sets is a laborious process, and this makes the already complex task of evaluating institutional exposures even more complicated.

System integrity and security

Employing multiple systems using different methodologies and different data makes the task of ensuring the integrity of each system more time-consuming and difficult. Furthermore, ensuring system security is more difficult when risk system management is decentralised. For example, one of the more common reasons for financial institutions discovering unexpected exposures is the incomplete security-trade input to the risk system. This is more difficult to police and audit when key elements of risk system management are in discrete and separate systems.

System maintenance

Multiple risk systems can increase operational risk and the cost of system maintenance. Often, the cost of maintaining multiple systems is not transparent, having been lost in multiple department budgets. Additionally, the risk of key person dependency, which we discuss later in more detail, is more likely when complex department-specific system maintenance procedures have developed over time.

Institutional exposure

Among all the operational and financial risks associated with incremental risk system development, the most important for a central bank is probably the lack of an institutional or enterprise view of risk that such system structures provide. The last few years have seen a very substantial increase in the correlation of financial markets, whether at the security type, asset class or country level (Longin and Solnik, 2001). This increased level of dependence and correlation across markets and securities has been driven by improved communications, more sophisticated investment strategies and the vast increase in capital flows. A central bank evaluating risk at a security or asset class level is not taking into consideration the covariance relationships that exist, and consequently could be substantially overstating or understating the institutional risk.

The consequences of understating risk for a central bank are clear, and must be avoided. However, overstating risk can also have a significant impact on the investment performance of the reserves management operation. One of the responsibilities of central bank reserves managers is to manage the reserves efficiently within risk or position limits as defined by senior management. If the reserves management investment process is creating alpha, overstating institutional risk results in a sub-optimal utilisation of the risk budget allocated. This financial cost is often undetected without an institution-wide risk evaluation system, but can have a significant financial impact on the reserves managers' performance.

6 The rise of value-at-risk techniques

Value-at-Risk (VaR) analysis attempts to answer just one question: "How much am I likely to lose?". In answering this question, VaR brings together traditional measures of a portfolio's current positions, measurements of the state of the market's volatility, and allowances for correlations between markets and instruments, all of which enable VaR to assess statistically how much markets may move against the portfolio, and then mathematically what the consequences in terms of loss will be. The output from a VaR analysis is therefore phrased in terms of the likelihood that the loss on the portfolio over the coming period (day, week, etc.) will be less than a certain figure – for example "with 95% probability, the portfolio will not lose more than $\$x$ million in the next 24 hours".

VaR methods differ from the simpler approach of former years in two areas. Firstly, the measurement of market volatility is considerably more mathematical. Options markets have been in existence for a long time, but it was not until the Black-Scholes model for pricing options was developed that they entered the mainstream of portfolio management. Since then, the use of options has increased dramatically. As a result, not only have market players been able to observe volatility at work and establish a much better feel for it, but also the price discovery mechanism in the market has meant that it is much more precisely defined.

Secondly, VaR enables risk controllers to handle the relationship between different markets. The inclusion within the VaR process of cross-correlations between different markets, and the realisation of how one set of assets can influence another set, has enabled investors to consider their risk positions on a whole portfolio basis, even where their portfolios contain many asset classes. The use of ever more powerful computers has enabled cross-correlations between the huge number of instruments and markets in a complex portfolio to be formally included in the calculation, to produce one overall mathematically derived figure for the total VaR for a portfolio, no matter how large, diverse or complex.

It is this last factor which is mainly behind the current enthusiasm for VaR techniques. VaR seems to combine the ultimate in sophistication with the ultimate in simple outputs – a number, in dollars, which even a non-market-specialist can understand. Given the pressures on senior central bank management to understand fully the risks that the reserves are running, such a system has undoubted appeal. Furthermore, comparisons between the returns on and the VaR statistics for various markets or investments seemingly offer portfolio managers a method for ensuring the maximum return for each unit of risk: investments can be concentrated in markets with the most return or potential for return for each dollar at risk. In this way, risk measurement tools can be used to maximise portfolio efficiency.

7 The control of risk and the use of value-at-risk in practice

Historically, much of the VaR research and application development has been undertaken by the "sell side", i.e. investment banks and trading institutions, which are primarily concerned with absolute returns. For central bank reserves managers, however, portfolio management returns and risk parameters are most commonly defined by reference to a benchmark. In this approach, the important measures of risk and return are those of the actual portfolio relative to the benchmark.

The starting point for much risk control is measurement of ex post risk. Measures of ex post risk provide management with an analysis of how much risk they have been taking in their portfolios, and enable them to decide whether the observed results conform to the risk appetite of the central bank.

It is often difficult to understand the risk and return characteristics of a portfolio in isolation. Furthermore, evaluating the risk return characteristics of a portfolio relative to some well-recognised index, such as a government bond index, may also produce misleading comparisons. In many cases, a central bank's investment managers will be operating within investment guidelines that are different from those used when constructing the index. To overcome this problem, a benchmark should be defined based on and reflecting the investment guidelines of the central bank.

With a benchmark and active portfolio structure, relative risk measures can be used to clarify the risks that the managers are taking in excess of the risks as defined by and inherent in the benchmark. The most common relative risk measure is tracking error. Ex post tracking error is usually defined as the standard deviation of the difference in the return on a portfolio and the return on the benchmark. Ex post TE is a measure of the risk that the portfolio managers have taken, and the higher the TE, the higher the risk that has been taken relative to the benchmark portfolio.

While the measurement of ex post TE is important, it remains a measure of the difference in two returns streams that has already occurred, and is thus a statistical recording of what has happened rather than a control measure over what might happen in the future. Of more interest to reserves managers and risk controllers is an ex ante measure of the risk in a portfolio, measured now and looking forward. Slightly confusingly, the commonest measure of this is also called tracking error, and is known as ex ante TE to distinguish it from ex post TE.

Ex ante TE is an estimate of the expected difference in future returns streams. In other words, it is a forecast of the ex post TE which will be evident at the end of the forecasting period, given the current starting position of the portfolio. The key to successful risk control for reserves managers is to establish a relationship between ex ante TE (a forward-looking estimate) and ex post TE (a backward-looking statistic). Once this has been done, management can then control ex post TE by measuring and controlling ex ante TE.

There are several ways of modelling ex ante TE, and back-testing can be used to determine the most appropriate model for the central bank, given its own style of reserves management. One common approach makes use of VaR techniques. The statement that a portfolio has 1-day VaR at the 95% level of \$$x$ million means that the portfolio has a 95% probability of not losing more than \$$x$ million over the next day. This measure looks forward one day only, and is a one-sided measure; VaR statistics refer only to potential loss. To use this VaR to calculate the ex ante TE, we first have to allow for the length of the measurement period (typically a year), and then convert it into a two-sided measure. The box below illustrates this.

Using VaR to calculate ex ante TE

A portfolio of $100 million has 1-day VaR at the 95% level (relative to the benchmark) of $100,000.

Converting this first into annual VaR, we use the formula that the *n*-day VaR for a portfolio is equal to ($\sqrt{n}$) times the 1-day VaR (assuming a normal distribution and stationarity). Assuming 250 working days in a year, the annual VaR for this portfolio is = $100,000 times $\sqrt{250}$, or $1,580,000.

Second, we convert this one-sided measure into a two-sided measure. If a portfolio is expected to lose more than this sum relative to its benchmark 5% of the time, then (assuming symmetrical returns) it will also outperform its benchmark by at least this sum 5% of the time, whereas for the remaining 90% of the time the portfolio's relative return will be within this amount.

Finally, we turn to statistics and the normal distribution, which tells us that a portfolio will be within 1.65 standard deviations of its mean 90% of the time. $1,580,000 therefore represents 1.65 standard deviations, producing a standard deviation of about $960,000. This is the ex ante TE.

TEs are more usually expressed in basis point terms, and we therefore have an ex ante TE of 96 b.p.

As with all statistical processes, the above calculation, while apparently precise, does rely on several assumptions, in this particular case the symmetrical nature of returns, the use of the normal distribution, and correlations remaining unchanged in the future. These are significant assumptions, but nevertheless, in most cases, and assuming the portfolio maintains roughly the same "shape" during the review period, the observed ex post TE at the end of the period under review should correlate well with the calculated ex ante TE. When combined with past experience, therefore, senior management should be able to derive a satisfactory relationship between VaR levels and ex post observed TE. This then allows them to use the powerful VaR techniques, and, by setting a range for the VaR on their portfolios, to expect that the resulting ex post TE will also be within a predefined acceptable range.

By bringing these statistical techniques together in this way, the risk manager has the tools to determine the risk characteristics of a portfolio in both absolute and relative terms, and then to determine whether the current risk in the portfolio's position is commensurate with the risk appetite as expressed by senior management.

Nevertheless, VaR is not an answer to every risk management question, and those considering the use of VaR techniques in this way need to ensure that they understand not only the strengths of VaR as a risk control methodology, but also its weaknesses and dangers. These are accordingly explored in the next section.

8 Weaknesses and dangers in using VaR techniques

For all the mathematical precision and seeming simplicity of VaR techniques, they still need to be used with care. With regard to VaR techniques, the first caveat is that inevitably they remain predictions of the future based on the past. As with all statistical models, to fully understand the risk numbers themselves, the user must have some understanding of the assumptions underlying the statistical model. Fund managers are well aware that "past performance is no guide to future performance", but some of the other recipients of VaR statistics are perhaps persuaded by the seeming precision of the process to place undue faith on the results. Volatility and correlation measures lie at the heart of all VaR calculations, yet they can and do change, and major outside shocks to a market can make such changes abrupt at times.

It is in this area of effective communication that VaR can offer some significant advantages. The concept of Value at Risk, or potential capital loss, subject to some assumptions, can be clearly understood provided that the assumptions are comprehensible. However, without knowledge of the underlying statistical processes, the use of the VaR numbers by themselves can lead to difficulties.

Coupled with this is the uncertainty over the way historic volatilities should be used. Although the volatility numbers themselves can be precisely derived for each historical period, they can be quite variable from day to day, and an element of smoothing via moving averages is therefore usually applied. But how far back into history one should look, and how the data should be combined (i.e. whether some sort of weighting factor should be applied to give greater weight to more recent readings), are both subjective decisions. Too little past data, or too fast a decay factor in the weighting routine, leads to excessive emphasis on the most recent figures and a volatility input series for the VaR process that can significantly fluctuate. On the other hand, long moving averages with little decay in the weighting process lead to an input series that is slow to change and may be dangerously late in reacting to changes in the market.

A third factor in the use of VaR methods, especially when they are used by senior management who are perhaps less closely in tune with the markets, is that excessive reliance may be placed on them. The air of precision over the numbers can mislead the user into believing there is a similar precision in the process. However, in actual fact the process remains heavily based on statistical theory, and users should recall at all times that the VaR numbers generated by the process are entirely dependent on the statistical assumptions and the level of confidence chosen for the statistical part of the process.

The statistical assumptions are indeed significant, and (as we explore later in Section 9) there is no uniform agreement on the best way to handle the uncertainties. However, even if the statistics are sound, the layman's understanding of them is not always so good. A level of money at risk quoted with a confidence level of 95%, or two standard deviations, sounds fairly safe. Yet if this is rephrased as "once a month (i.e. one day in every 20 working days) you could lose at least this much", the perception of risk changes. Senior management can react by demanding to know what loss level will not be exceeded once a year, or even longer; the statistical process can generate these numbers easily enough by using three or even four standard deviation levels; nevertheless, the size of the figure for the maximum loss is correspondingly larger and may cause management to react by reducing positions further.

A considerable restraining factor in the growth of VaR techniques is the size of the IT investment required. The amount of data required to calculate and then put into use all the volatilities and cross-correlations is huge, and can consume computing power and IT resources on a commensurate scale. Although the emergence of commercially-available packages of volatility and cross-correlation data at least simplifies the data collection side of the IT challenge, using this data remains a very computer-intensive task, and even reserves managers of the largest central banks will find the task of building VaR models of their own from scratch rather daunting.

A more serious concern raised by VaR is that it encourages herding in the markets and may even promote instability. Faced with the same market information, and using similar risk

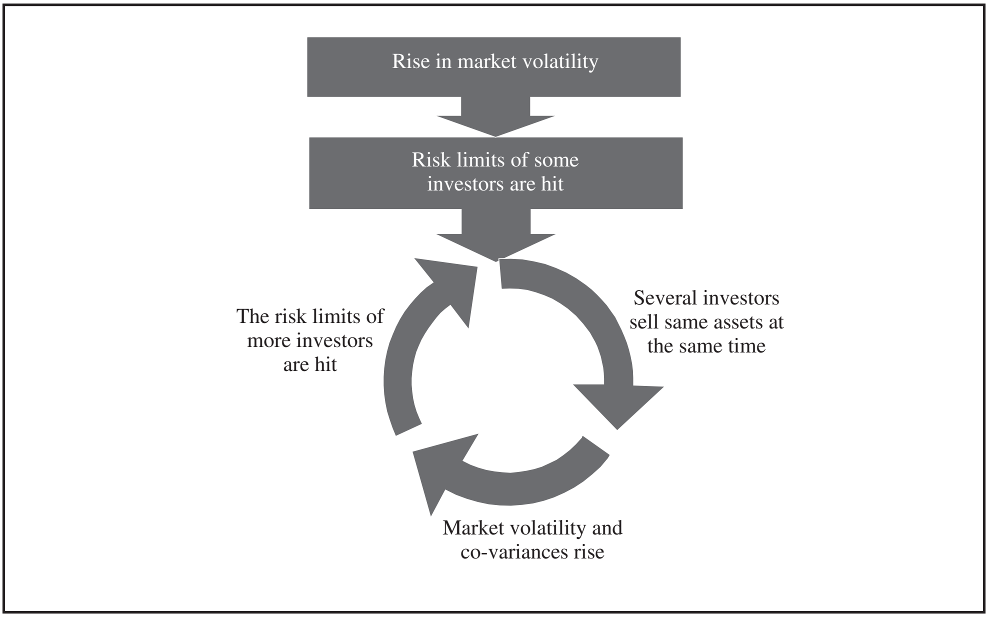

models, investors are drawn towards the same low risk portfolios. As volatilities rise, therefore, they will all probably seek to reduce their exposure at the same time, and even end up selling the same assets. This creates a spiral in which volatilities rise, causing more positions to be sold (Persaud, 2002).

Another issue that can arise is that, to a central bank, not all losses of the same financial size carry the same consequences. VaR techniques, which measure all losses in currency terms (and thus treat a loss of $*x* million identically however it is incurred), are unable to capture this difference, and may not fully reflect central bank preferences and risk appetites. For example, credit risk is often perceived more sensitively than market risk, and a loss of $*x* million through a credit default may be more damaging to the central bank's reputation than an identical loss through market price and yield movements. Another risk which central banks may be more than usually sensitive to is losses caused by the use of derivatives. Although derivatives are a standard part of the markets and of portfolio management, they retain an element of mystique which lends a more serious tone to losses incurred through their use. In contrast, currency risk, i.e. the risk of losses incurred through currency movements, is accepted more freely as inherent in central bank reserves management. These nuances are not easy to convey to the reader via VaR statistics.

Finally, there is the well-known effect known as the seatbelt syndrome. Just as a seatbelt can give a car driver a false sense of security and lead to faster and more dangerous driving, so VaR methods may lead central bank senior management to rely excessively on them and, perversely, to increase the level of risk in their portfolios. A risky portfolio does not become "safe" by the addition of a VaR statistic, and there is always the possibility of a three or even four standard deviation move in the market to throw a portfolio off balance and into heavy loss-making territory.

Indeed, perhaps the greatest challenge for promoters of VaR techniques is in advising users (including senior management) what they *cannot* do. In particular, no VaR statistic can guide a manager when the basic ground rules change in a market. For example, when the Asian crisis unfurled in 1997, some markets in Asian securities ceased to function at all, not only making a mockery of the VaR figures, but also locking investors into positions they could not sell at any price. In the last resort, using VaR statistics is not a substitute for assessing more basic risk such as whether markets will continue to function at all.

9 Calculating VaR in practice: different VaR methodologies

There are two main methodologies for analysing VaR risk: historical simulation and factor models.

The historical simulation approaches are based on the full valuation of the portfolio, and have the advantage that they capture the portfolio's full distribution. They elucidate all the market variables that determine the volatility of the portfolio, can be used on all security types in a consistent way, display the full impact of any non-linearity in the portfolio, and can be effectively used to analyse the "fat-tail" events.

The cost of this approach is a potentially substantial data management and processing task. The valuations and historical data must be accurate if the resultant analysis is to be useful. In fact, the data quality issue is essential to the successful implementation of historical simulation-based approaches. Poor quality position or historical data can lead to a large quantity of risk flags and exception reports. This will over time degrade the value of the system, as users will come to regard risk flags as data errors. Even more dangerous for the central bank are poor quality data that do not prompt sufficient risk flags. This can result in a significant understatement of risk and, more worryingly, an environment of false well-being that has been created by the risk system infrastructure.

An alternative to historical simulation is the factor model. This method utilises parametric approaches with a similar structure to that underlying the mean/variance optimisation tools. Parametric type models overcome many of the data issues by using statistical distribution assumptions to describe the security distribution. Furthermore, some users find factor-based models more easy to use as the factors "explain" the risks in an intuitive way. However, the data compression used in factor models can obscure exceptional events. Further parametric approaches tend to use different models for different security types, making aggregation and enterprise risk management more complex.

Central to the assumptions contained in the factor approach to VaR is the use of the normal curve for the distribution of market movements. Extensive analysis has shown that this assumption is probably only an approximation of real life, and VaR models based on it appear to underestimate the likelihood of extreme events (in statistical terms, the distribution of market movements is not normal but has "fat tails"). However, supporters of factor methods claim that the approximation is a good one, and the gains in simplicity of use outweigh any disadvantages[3].

3 For a more detailed discussion of the merits of historical simulation versus factor-based models, see Askari (2001).

In-house build or purchase

When a central bank decides to implement an enterprise risk management system, a decision on whether to build internally or purchase a commercial system needs to be made. Both options have advantages and disadvantages, and neither option is without complexity.

The advantage of building internally is that the system will be built specifically for the central bank. The development can be built with the constant and detailed involvement of all the major users. Furthermore, the systems operation procedures and output can be built to fit optimally with the bank's investment process.

However, there are a number of significant disadvantages to an enterprise risk system built in-house. Firstly, risk management systems are complex, and the larger the scope and complexity of the asset classes in the investment universe, the more complex the system. The cost of a system built in-house can be prohibitive, particularly when ongoing maintenance is considered. Secondly, the time needed to build such a system will be considerably longer than the time to implement a commercially available product, even with detailed customisation. The opportunity cost of such a delay needs to be factored into the decision-making process.

Risk systems developed in-house also have a tendency to create key personnel dependence, particularly after the first stage of development, when the department-specific functionality is added. By purchasing a risk system from a specialist, the central bank ensures that they have the latest available risk analytics without key person dependence. Furthermore, a commercially available system will be more scalable then a system developed in-house, and will have lower maintenance costs.

Unfortunately, purchasing a commercially available enterprise risk system is still a substantial project for a central bank. There are several enterprise risk systems available, and a detailed and thorough evaluation process is needed to determine the best fit.

Fundamentally, the historical simulation approach tends to offer advantages over factor-based models, where the investment universe is complex. However, as discussed above, a more important issue than the methodology is user understanding. Understanding the underlying assumptions of any risk model will ensure that the user does not have an unrealistic sense of control and security.

10 Reserves management within the wider central bank context

So far this paper has focused on the financial risks inherent in reserves management. A mathematical approach to their analysis has been outlined, and various statistically-based methods discussed. However, these risks, and the central bank's response to them, should not be seen wholly in isolation. The control of financial risks is both part of, and contributes to, the central bank's overall risk management. This overall task will also encompass non-financial risks, such as operational risk, legal risk, business risk and reputational risk.

One solution to these myriad risks is to adopt a separate risk control procedure and process for each one. In one sense there is merit in this, as the mere fact that one set of risks cannot be numerically quantified (for example reputational risk) should not deter the central bank from using mathematical methods for those risks that can be numerically quantified. However, as discussed above in Section 5, there are substantial disadvantages to incremental risk system development, making it worth considering whether a whole enterprise view of risk creates a more controlled and more effective investment process.

When selecting a whole enterprise risk management system, one of the key decisions relates to methodology. Different risk methodologies use different assumptions, and produce a different set of emphases. It is useful to discuss the underlying assumptions of a particular risk methodology in relation to a particular risk analysis, but little value is created by comparing the relative merits of one risk methodology to another in general terms. However, a prerequisite for a whole enterprise risk management system is that the risk methodology used in the system is appropriate for and consistent with the overall management framework and ethos of the central bank.

With whole enterprise risk management, the central bank has the means to minimise the risk associated with unexpected risk, creates a language of risk across departments, provides a tool for the Investment Committee to influence the investment process, and allows the institution to maximise the alpha from the allocated risk budget. Given these advantages, it is optimal for a central bank with multi-asset class exposure to give budget priority to purchasing or building a centralised enterprise risk management system. However, this should not necessarily be the only risk system in use.

As stated above, risk systems produce risk analytics subject to a statistical model. An effective risk model is one in which the assumptions are appropriate for the particular analysis. It may be that some of the specialist departments can add more value in the investment process by using a specialist risk management tool. These specialist risk management systems have some of the cost and maintenance problems discussed above; however, the operational, security, data and methodology issues are resolved by the enterprise system.

It is therefore perfectly consistent to advocate a whole enterprise risk management system while still permitting specialist risk systems within the central bank. In fact, it could be reasonably argued that, as risk systems only provide a risk estimate based on one set of particular statistical assumptions, it is appropriate to evaluate asset class risk using different models.

11 Conclusion

Financial risk is more quantifiable than the other types of risk; consequently, it is possible to use statistical models to analyse risk profiles. Some have argued that it is inconsistent to deal with financial risk alone with analytical rigour, as the same approach cannot be adopted for all risk types. This argument essentially puts consistency above effective management.

A substantial part of a central bank's financial risk is taken voluntarily, and consequently a substantial reputational risk is associated with the financial risk of a central bank. Furthermore, the non-financial risks are mainly controlled by the management and operational procedures. However, the complexity of financial risks makes operational procedures, on their own, of limited value.

Central banks should use all available tools to manage the institution effectively. If a central bank intends to use complex financial products in order to create alpha, it also needs to adopt the IT systems that have been built to deal with this complexity. In addition, the central bank should consider whole enterprise risk management. Security and product departmental approaches to risk can significantly overstate or understate the risk of the overall institution, as the covariance relationships are ignored.

Whether selecting a system from those available commercially or building a risk system internally, the decision is complex and requires detailed analysis. It is particularly important to recognise that risk models can only provide a statistical estimate of the institutional risk. It is more important that the users of a risk system understand the underlying methodology and consequently its limitations, than it is to have the most advanced statistical models.

In the final analysis, risk systems should be employed to complement the excellent management procedures and risk-aware culture of central banks. They are not a substitute for existing management practices, but instead tools which strengthen the natural caution which central banks bring to the task of managing their reserves. Used properly, they allow reserves managers to utilise complex financial products in a controlled environment, and thereby provide a more effective management of national reserves.

References

Longin, François & Bruno Solnik: "Extreme Correlation of International Equity Markets" – Journal of Finance, November 2001

Nugée, John: "Foreign Exchange Reserves Management" – Bank of England Centre for Central Banking Studies, October 2000

Persaud, Avinash: "The Folly of VaR: How Modern Risk Management Practices Are Creating Risk" – Reprint of Gresham College Lecture, December 2002

Askari: A Discussion of the Event Simulation Model versus Factor-based Models – published paper, September 2001

9

Corporate bonds in central bank reserves portfolios: a strategic asset allocation perspective

Roberts L. Grava

Abstract

As the discussion on generating excess returns from central bank reserves portfolios continues, consideration is increasingly given to allocations into investment-grade corporate bonds. This paper analyses the implications of adding corporate debt securities to government fixed-income portfolios, and finds that including exposure to corporates should result in long-term outperformance versus similar maturity government securities, and that this excess return is expected to come with minimal additional market risk. Using a combination of historical and simulation analysis, the author concludes that adding corporate risk is to an extent more attractive from a risk/reward standpoint than increasing portfolio duration. Furthermore, a long-term passive allocation to the corporate sector would likely outperform active management strategies, where institutions would add or remove corporate risk on a tactical basis. Finally, the ability of institutions to tolerate short-term underperformance relative to government sectors could increase long-term terminal wealth.

1 Introduction

As central banks continue to focus on generating excess return from their reserves portfolios, an energetic discussion on the merits and pitfalls of investing in corporate securities has emerged. Adding corporate risk to a reserves portfolio should result in long-term outperformance versus similar maturity government securities, and this excess return is expected to come with minimal additional portfolio market risk. Nevertheless, the inclusion of corporates in a reserves portfolio introduces a new set of risks, and raises a number of practical management issues.

This paper attempts to address some of these issues. We will discuss some of the arguments in favour of adding corporate risk to a reserves portfolio, especially in comparison with traditional means of excess return generation. We will explore active and passive investment strategies, and strategic and tactical allocations out of governments into corporates. We will investigate the historical and simulated performance of single-asset and mixed portfolios, and analyse the performance of these portfolios over different investment horizons.

We should note at the outset that we will only be investigating the characteristics of highly-rated investment-grade investments. Investments in lower-rated securities require a great deal of individual security analysis that may necessitate devoting an impractical amount of resources for most central banks. We focus instead on the risks and benefits of portfolios that aggregate a diversified collection of high-grade corporate securities, and analyse whether low-maintenance passive strategies involving investment-grade spread product can add greater value than more familiar means of excess return generation.

We find that when considering strategies for enhancing reserves portfolio returns, adding spread risk (to an extent) is more attractive from a risk/reward standpoint than increasing portfolio duration. When managing portfolios that allow exposure to corporates, a long-term passive allocation to credit sectors would likely outperform active management strategies where central banks would add or remove corporate exposure on a tactical basis. Finally, the ability for institutions to tolerate short-term underperformance relative to government sectors could increase long-term terminal wealth. All of the above leads us to conclude that the inclusion of corporate debt in a central bank reserves portfolio is a viable strategy that is at least worth examining.

2 Underlying assumptions

For the purposes of this study, we analyse relative performance vs. a benchmark portfolio 100% invested in 1-3 year US government bonds. The 1-3 year maturity sector is popular among central banks with non-cash benchmarks for several reasons: it provides expected returns well in excess of returns from cash-type assets; the sector generally has the highest Sharpe ratio among other maturities; and an analysis of historical data shows that there has been no 12-month period when 1-3 year US governments delivered a return less than zero. This last point can be viewed as a key consideration for central banks, which typically place a great emphasis on reputational risk. Our aim, then, is to examine whether an institution can expect to enhance expected returns from its reserves portfolio while maintaining an expected risk profile broadly in line with a 1-3 year government benchmark. Institutions managing against cash-type benchmarks should re-examine their attitude towards market risk in general before considering the addition of corporate bonds to their portfolio.

In generating the return data used in this study, we used monthly total return information from the Merrill Lynch US dollar 3-month T-bill, 1-3 year government, 3-5 year government, 1-3 year agency, 1-3 year AAA corporate, 1-3 year AA corporate, 1-3 year A corporate and 1-3 year BBB corporate indices, from March 1990 through June 2002. Annualised geometric mean return and standard deviation data are used as the annual return information presented in Table 1.

Table 1: Annual return and standard deviation, 1990 - 2002
(%)

	3M Bills	Govt 1-3	Agcy 1-3	AAA 1-3	AA 1-3	A 1-3	BBB 1-3	Govt 3-5
Return	5.08	6.83	7.01	7.41	7.58	7.54	7.24	7.99
Standard Deviation	*0.74*	*2.21*	*2.16*	*2.20*	*2.28*	*2.23*	*2.55*	*4.31*

Data source: Merrill Lynch Bond Indices, March 1990 - June 2002.

Unsurprisingly, the expected return from T-bills is significantly less than the expected annual return from 1-3 year government securities, although expected volatility is lower as well. Similarly, the 3-5 year government sector provides greater expected returns with greater expected risk. Analysing the 1-3 year credit sectors, expected volatility relative to governments rises only slightly as credit quality decreases, but the expected return increases steadily up to the AA/A sectors, whereupon expected returns decrease once more.

Table 2: Correlations, 1990 - 2002

	3M Bills	Govt 1-3	Agcy 1-3	AAA 1-3	AA 1-3	A 1-3	BBB 1-3	Govt 3-5
3M Bills	1.0000	0.5796	0.5793	0.5266	0.5154	0.4902	0.3731	0.4221
Govt 1-3	0.5796	1.0000	0.9870	0.9808	0.9762	0.9581	0.7847	0.9655
Agcy 1-3	0.5793	0.9870	1.0000	0.9760	0.9774	0.9576	0.7792	0.9440
AAA 1-3	0.5266	0.9808	0.9760	1.0000	0.9832	0.9678	0.8165	0.9502
AA 1-3	0.5154	0.9762	0.9774	0.9832	1.0000	0.9813	0.8426	0.9475
A 1-3	0.4902	0.9581	0.9576	0.9678	0.9813	1.0000	0.8661	0.9315
BBB 1-3	0.3731	0.7847	0.7792	0.8165	0.8426	0.8661	1.0000	0.7543
Govt 3-5	0.4221	0.9655	0.9440	0.9502	0.9475	0.9315	0.7543	1.0000

Clearly, the risk and return parameters for these asset classes are dominated by maturity considerations – the expected return differences are considerably more dramatic between T-bills and 1-3 year governments, for example, than between 1-3 year governments and 1-3 year AA corporates. This can also be seen in the derived correlation matrix shown in Table 2.

At first glance, Table 2 seems to show that there are greater diversification benefits from incorporating different maturity assets into the reserves portfolio, as the correlations between assets in the 1-3 year sector (with the exception of BBB corporates) are relatively higher. Indeed, the high correlations among similar maturity sectors might lead some analysts to conclude that diversifying into corporate risk does not provide any meaningful diversification. We will nevertheless devote the following sections to analysing both return and correlation data in investigating the strategic asset allocation consequences of maturity and credit deviations away from the 1-3 year government benchmark.

3 Compensation for assumed risk

A classic trade-off in finance is between return and risk. Searching for excess return should also entail the assumption of excess risk. Conversely, when assuming excess risk, an investor expects some compensation. Not all risks are compensated equally, however, and the following analysis will attempt to quantify the relative risk/return trade-offs from two types of risk familiar to fixed-income investors – duration risk and spread risk.

Central banks generally feel more comfortable assuming directional interest rate risk than undertaking the spread risk resulting from investments in corporate fixed income paper. Perhaps they feel that there is greater risk reduction potential from diversifying across maturity sectors; perhaps an abundance of economists in their institutions gives them the impression that they possess some unique insight into future interest rate movements unknown to the market in general. Certainly, reputational risk plays a part – central banks would be embarrassed to disclose that they had invested national reserves in corporate debt that had subsequently defaulted.

Table 3: Portfolios incorporating duration deviations
(%)

Portfolio Weights			Expected Annual Return	Expected Annual Std. Deviation
3M Bills	Govt 1-3	Govt 3-5		
50	50	-	5.97	1.35
40	60	-	6.14	1.52
30	70	-	6.31	1.69
20	80	-	6.48	1.86
10	90	-	6.66	2.04
-	100	-	6.83	2.21
-	90	10	6.95	2.41
-	80	20	7.07	2.61
-	70	30	7.19	2.82
-	60	40	7.31	3.02
-	50	50	7.43	3.23

Nevertheless, it is interesting to put the risk/reward trade-offs of these two types of return enhancement strategies into perspective. The following analysis examines the expected compensation for assuming directional risk (lengthening or shortening portfolio duration vs. the 1-3 year government benchmark) and assuming spread risk (replacing government bonds with similar maturity corporate paper). Utilising the return and correlation data from Tables 1 and 2, Table 3 examines the expected portfolio characteristics of several portfolios that implement duration deviations away from the 1-3 year government benchmarks. Similarly, Table 4 presents portfolio data for a series of portfolios that increasingly add spread risk which is generally duration-neutral.

Table 3, which describes the payoffs from assuming directional interest rate risk, exhibits the characteristics of a classic risk/return trade-off – an institution can increase expected return, but at the expense of increased expected volatility; or it can reduce risk, at the expense of decreasing expected returns. The actual amount of risk addition or reduction chosen in a strategic asset allocation exercise would depend on an individual institution's risk tolerance.

Table 4, however, examines the characteristics of portfolios that add spread risk relative to the benchmark (we cannot examine the effects of reducing spread risk here, since the

Table 4: Portfolios Incorporating Spread Risk Deviations
(%)

Portfolio Weights		Expected	Expected Annual
Govt 1-3	AA 1-3	Annual Return	Std. Deviation
100	-	6.83	2.21
90	10	6.91	2.22
80	20	6.98	2.22
70	30	7.06	2.22
60	40	7.13	2.23
50	50	7.21	2.23
40	60	7.28	2.24
30	70	7.36	2.25
20	80	7.43	2.26
10	90	7.51	2.27
-	100	7.58	2.28

Figure 1: Compensation for assumed risk

(%)

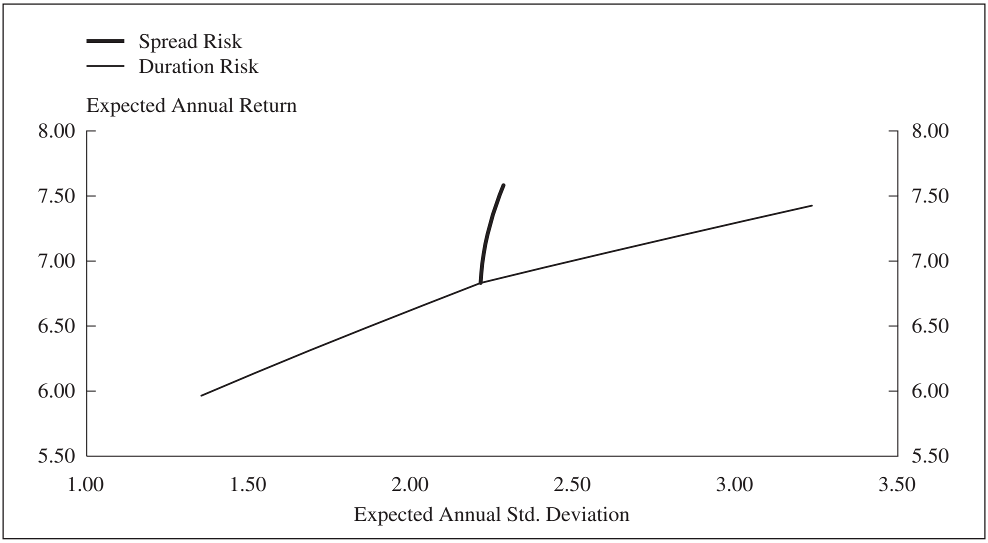

benchmark consists of 100% governments). It shows that increasing spread risk increases expected returns dramatically, but only marginally increases expected portfolio risks.

As can be seen in the Figure 1, when both of these portfolio sets are put together graphically, the image is striking:

The line charting expected return and volatility for the duration risk portfolios is positively sloped, which implies that increasing duration risk should add expected return. The line charting the characteristics of the spread risk portfolios, however, is much more positively sloped, and is nearly vertical. This means that increasing the expected portfolio return via assuming spread risk could be done with less expected portfolio volatility than by assuming duration risk. Conversely, for each basis point of additional portfolio risk, one is compensated considerably more for assuming spread risk than duration risk. Of course, there is a limit to the extent of excess return generation from investing in corporates. The upper point of the spread risk line is the maximum amount of expected excess return from AA corporate bonds without increasing portfolio leverage. Similarly, decreasing credit quality further would not necessarily increase expected returns, as can be seen in Table 1.

Duration, however, can be extended further than shown in the chart; holding a portfolio of 3-5 year government bonds yields an expected annual return of 7.99%, which exceeds the maximum expected return from 1-3 year spread sectors. The cost, of course, is much greater assumed market risk. The purpose of the preceding analysis, however, is to compare incremental excess returns to incremental excess risk; from this perspective, these results seem to make the argument for including high-quality investment-grade corporate bonds in a reserves portfolio quite compelling.

However, adding corporate paper to a reserves portfolio also adds different dimensions of risk – default risk, liquidity risk, concentration risk, event risk, etc. Most of these risks can be managed by holding a well-diversified portfolio of many corporate names, spread across industries. As an institution's portfolio more closely mirrors the characteristics of a broad corporate index, it should be able to deliver the excess return and risk diversification benefits of the asset classes analysed here. This necessitates holding a portfolio of more, rather than less, securities. While this might seem counterintuitive for central banks that lack the credit research resources of private sector asset managers, this is nevertheless an important point. It is precisely because most central banks do not have the resources to add value through specific corporate security selection, that an appropriate and attainable goal would be to construct portfolios that minimise tracking error against broad corporate indices. This also has the added benefit of lessening concentration risk: the inevitable negative surprises in specific names will have a lesser impact on the total portfolio, and consequently any possible reputational risk will be lessened, as well.

Default risk is a concern for any investor in corporate fixed income securities, but actual observed defaults for the types of investment-grade bonds that central banks generally consider are exceedingly small. Investments in high-yield speculative-grade securities require serious security-specific credit analysis, but central banks active in corporate bonds typically concern themselves with the highest ends of the credit quality spectrum. Table 5 summarises typical credit quality data by rating class.

Table 5: Descriptive credit statistics by rating category

(%)

	Aaa	Aa	A	Baa	Ba	B	Caa-C
Mean Annual Default Rate (1970-2001)	0.00	0.02	0.01	0.15	1.21	6.53	24.73
Mean Annual Credit Loss Rate (1982 - 2001)	0.00	0.02	0.01	0.12	0.88	4.15	16.43

Source: Moody's Investors Service.

By investing only in AAA – A securities, most investors would only feel the effects of default risk through changing spreads, as perceptions of credit quality change. The ratings agencies generally initiate ratings downgrades before debt becomes distressed, and this gives corporate fixed-income investors the opportunity to liquidate holdings (admittedly at a loss) before credit quality becomes completely unacceptable. Even holders of debt from once highly-rated firms such as Tyco, Enron or WorldCom would never have had to book credit losses if they had followed strict stop-loss policies, such as liquidating bonds when ratings or prices fall below a certain threshold. Clearly, liquidating positions at a market loss is painful, but it is less so than holding distressed paper originally bought at investment grade. Once again, diversification is key to ensuring that inevitable market losses from downgraded securities are offset by gains from better-performing paper.

Table 6 shows Moody's Investors Service data on 1-year average rating transition rates. For example, an A-rated bond would have a 2.41% probability of being upgraded to Aa over the

Table 6: 1-Year average rating transition rates

	Rating to:								
Rating from:	Aaa	Aa	A	Baa	Ba	B	Caa-C	Default	WR
Aaa	89.09	7.15	0.79	0.00	0.02	0.00	0.00	0.00	2.94
Aa	1.17	88.00	7.44	0.27	0.08	0.01	0.00	0.02	3.01
A	0.05	2.41	89.01	4.68	0.49	0.12	0.01	0.01	3.21
Baa	0.05	0.25	5.20	84.55	4.51	0.69	0.09	0.15	4.51
Ba	0.02	0.04	0.47	5.17	79.35	6.23	0.42	1.19	7.11
B	0.01	0.02	0.13	0.38	6.24	77.82	2.40	6.34	6.67
Caa-C	0.00	0.00	0.00	0.57	1.47	3.81	62.90	23.69	7.56

Source: Moody's Investors Service.

course of one year, and a 4.68% probability of being downgraded to Baa. These types of transitions are generally consistent across the ratings classes, and this suggests that a wise strategy in managing security-specific corporate risk is "avoiding losers", as opposed to "picking winners", as there should be more downward than upward rating migration in any given year.

4 Active or passive management?

For central banks considering whether to include corporate debt into their reserves portfolios, a natural question is whether to commit a fixed long-term allocation to spread product, or whether to utilise corporates as an opportunistic asset class – buying corporates when they are expected to outperform, and moving back to government securities when corporates are expected to underperform.

Many studies have addressed the question of whether active management strategies can outperform passive strategies over a longer horizon. Our goal here is to quantify the expected results from specific active strategies in portfolios that can include government bonds and corporates, and compare these findings with expected returns from holding long-term fixed allocations to corporate bonds.

Based on terminal wealth values utilising the data that underlie Tables 1 and 2, we calculate the annual outperformance of each credit sector against the 1-3 year government sector, which is presented in Table 7. In other words, holding a portfolio invested 100% in 1-3 year AA securities over the period March 1990 to June 2002 would have generated 75 basis points in average annual outperformance vs. a 1-3 year government portfolio.

Table 7: Annual outperformance vs. 1-3 Year governments, in bps

Agcy	AAA	AA	A	BBB
19	58	75	71	41

We compare these results to an active tactical strategy. For each asset class during the observed time period, every month the portfolio manager has a choice – to be 100% invested in governments, or 100% invested in the corporate or agency asset class. The portfolio manager would obviously want to be invested in governments when corporates are underperforming, and vice versa. As the manager can have no concrete knowledge of the coming month's performance in advance, it is reasonable to expect that in some months he or she will make the wrong choice by investing in the underperforming asset class. We model this by analysing the potential "success rates" of the portfolio manager. If the portfolio manager can deliver a 60% success rate, for example, he or she will have invested in the correct asset class 60% of the time, and chosen the underperforming asset class 40% of the time.

We iterated through the time series of observed monthly returns from 1990-2002. If we were analysing a 60% success rate, then each month for each asset class there was a 60% probability that the outperforming asset was chosen and a 40% chance that the underperforming asset was selected. The result was another series of monthly returns for each asset class' active management strategy, which incorporated a mix of government and non-government returns, calculated according to the methodology mentioned above. The terminal value for each asset class' active management strategy was compared to the terminal value of

Table 8: Annual outperformance from active management, in bps

Success Rate (%)	Agcy	AAA	AA	A	BBB
0	-45	-30	-27	-41	-132
5	-40	-24	-21	-33	-117
10	-34	-18	-14	-25	-102
15	-29	-12	-8	-18	-86
20	-23	-6	-1	-10	-71
25	-18	0	5	-3	-56
30	-13	5	12	5	-41
35	-7	12	18	13	-25
40	-2	17	25	20	-10
45	4	24	31	28	7
50	9	29	37	35	20
55	15	35	44	43	37
60	20	41	50	51	52
65	25	47	56	58	66
70	31	53	63	66	83
75	36	58	69	73	97
80	42	65	76	81	114
85	47	71	83	88	129
90	53	77	89	97	145
95	58	83	96	104	160
100	64	89	102	112	176

a portfolio 100% invested in 1-3 year governments during that time period, and average annual outperformance numbers were calculated. This process was then repeated a thousand times, resulting in mean outperformance numbers from active management strategies with varying success rates (see Table 8).

For example, with a 60% success rate, a portfolio manager should expect to outperform the 1-3 year government benchmark by 50 basis points annually, if actively choosing between governments and AA corporates. Interestingly, even "bad" portfolio managers, i.e. ones who choose the incorrect asset class more frequently than the correct one, can expect to outperform governments by actively investing in the AAA, AA or A sectors. Even a 30% success rate would generate 5, 12 and 5 basis point annual outperformance in the AAA, AA and A sectors, respectively.

However, these numbers do not compare favourably with the annual expected outperformance from simply holding the risky asset classes over the entire time period. For example, for an active management strategy in AAA – A corporates, in order to generate greater outperformance than a full-time allocation to the risky asset, a portfolio manager would have to demonstrate a success rate between 70–80%, which is quite unlikely to happen in the real world, especially over a multi-year time horizon.

Therefore, the data seem to suggest that central banks which include corporate paper in their reserves portfolios would be better served by committing a long-term allocation to the sector, instead of opportunistically attempting to increase and decrease their exposures. Unless, of course, they happen to employ portfolio managers with extraordinary amounts of prescience, in which case the opposite would apply.

5 Portfolio construction

Until now, we have concerned ourselves with the observed performance of portfolios consisting of only one or two of the analysed asset classes. We shall now review the performance of portfolios that can contain combinations of all of the analysed assets. Before we discuss concrete portfolios, however, we need to generate performance data to aid our analysis.

Since the return data we have analysed contain only 147 monthly observations, we will construct a much larger simulated set of monthly returns, using sampling with replacement methodology. From the set of actual monthly data, we draw a set of actual returns for three consecutive months, from a starting month chosen at random. We choose series of three monthly returns, instead of choosing single month returns randomly, in the hope of capturing some of the autocorrelation that is typically present in financial time series. We repeat this process until we have a simulated time series 60,000 months, or 5,000 years, long, from which we can draw performance data for monthly, quarterly, yearly, 5-year and 10-year horizons. Mean expected holding period outperformance (vs. 1-3 year governments) is presented in Table 9.

Table 9: Expected holding period outperformance vs. 1-3 yr. governments, in bps

Horizon	3M Bills	Govt 1-3	Agcy 1-3	AAA 1-3	AA 1-3	A 1-3	BBB 1-3	Govt 3-5
Month	-14	0	1	4	6	6	4	9
Quarter	-41	0	4	13	18	17	11	28
Year	-162	0	17	54	71	67	43	111
5 Years	-784	0	86	273	359	337	219	568
10 Years	-1,506	0	173	554	730	686	443	1,168

For example, a portfolio of 1-3 year A securities would be expected to outperform 1-3 year governments by 17 basis points over a one quarter holding period, and by 67 basis points over a one year holding period. T-bills are always expected to underperform 1-3 year government bonds on average.

However, since these outperformance calculations are only average numbers, we also examine the probability that an asset class could underperform relative to the benchmark over the holding period (i.e. the terminal value would be less than that of the benchmark at the end of the horizon). These data are presented in Table 10.

Table 10: Probability of underperformance at the end of the holding period *(%)*

Horizon	3M Bills	Govt 1-3	Agcy 1-3	AAA 1-3	AA 1-3	A 1-3	BBB 1-3	Govt 3-5
Month	65	0	41	30	24	26	29	42
Quarter	68	0	40	20	18	24	28	41
Year	82	0	29	10	6	13	37	30
5 Years	98	0	14	0	0	1	26	13
10 Years	100	0	5	0	0	0	18	5

As an example, a portfolio of AA corporates would be expected to outperform governments by an average of 6 basis points every month, but such a portfolio would be expected to underperform governments roughly one in every four months. In contrast, a T-bill portfolio would be expected to outperform 1-3 year governments 35% of the time; however, there were no 10-year holding period observations where the terminal value of the T-bill portfolio exceeded the terminal value of the 1-3 year government portfolio.

The simulated return set also yielded traditional volatility and covariance data which are used in the following portfolio optimisation exercise. Using the expected returns and volatility data from the simulated return set, and a mean-variance optimisation framework, we generated five optimal portfolios, each optimised to perform over a specific investment horizon. The "Month" portfolio was optimised using monthly data, whereas the "Quarter" portfolio was optimised using quarterly data, and so on for the other categories. The goal was to generate portfolios that would maximise outperformance vs. the government benchmark over the specified investment horizon, while being constrained to have expected portfolio volatility equal to the government benchmark. In other words, this was an attempt to maximise return without assuming excess portfolio risk[1]. The resulting portfolio weights, and expected horizon outperformance data, are presented in Table 11.

According to these results, over a one month horizon there is no combination of assets that would outperform 1-3 year governments with the same portfolio risk. However, over longer horizons the benchmark asset disappears from optimal portfolios entirely. Similarly, longer-maturity governments never appear in the optimal portfolios; corporate assets offer a better risk/return trade-off. Nor do agency securities ever appear in the optimal portfolios; while offering some yield enhancement over the benchmark governments, high-rated corporates offer greater income potential with favourable risk characteristics.

It is also interesting to note that while increasing credit risk results in greater expected returns versus governments, decreasing credit quality to BBB would not result in an optimal portfolio. AAA securities appear in some optimal portfolios, but the majority of optimal portfolio allocations go to the AA and A sectors. The increasing allocation to T-bills in the longer-horizon portfolios is also notable, although this still remains small.

Table 11: Portfolio weights and expected period outperformance (bps) vs. 1-3 yr. governments

Portfolio	Weights								Expected Period Out-performance
	3M Bills	Govt 1-3	Agcy 1-3	AAA 1-3	AA 1-3	A 1-3	BBB 1-3	Govt 3-5	
Month	-	100%	-	-	-	-	-	-	0
Quarter	1%	-	-	14%	44%	41%	1%	-	16
Year 1	%	-	-	14%	44%	42%	-	-	68
5 Years	4%	-	-	-	25%	70%	-	-	407
10 Years	8%	-	-	4%	29%	59%	-	-	999

[1] Conventional mean-variance optimisation uses a quadratic programming algorithm to minimise expected volatility for a given return target. We are interested in maximizing return while keeping expected risk constant, and therefore must use a more computationally intensive non-linear optimisation algorithm with non-linear constraints.

Table 12: Average holding period outperformance vs. 1-3 yr. governments, in bps

	Portfolio				
Period	"Month"	"Quarter"	"Year"	"5 Years"	"10 Years"
Month	0	6	6	5	4
Quarter	0	18	18	15	13
Year	0	75	75	64	56
5 Years	0	492	489	421	369
10 Years	0	1,379	1,370	1,183	1,045

Table 13: Probability of holding period underperformance
(%)

	Portfolio				
Period	"Month"	"Quarter"	"Year"	"5 Years"	"10 Years"
Month	0	22	22	26	28
Quarter	0	18	18	23	25
Year	0	6	5	12	13
5 Years	0	0	0	1	1
10 Years	0	0	0	0	0

6 Portfolio performance simulations

These portfolios were optimised under one set of historical data; we will now test their performance in a simulation using a similar, but different, set. Using the original set of 147 monthly observations, we construct a new 60,000 month simulated time series. We use the sampling with replacement methodology used in the previous section, but decay the data, so that the most recent data have a greater probability of appearing in the sample. Specifically, we use a decay factor of 0.9841, which corresponds to the most recent monthly returns being ten times more likely to appear in the sample than the oldest monthly returns. This should capture some of the negative aspects of holding corporate bonds recently, at a time when government yields are currently at historic lows, spreads are high, and investor enthusiasm for holding corporate names in the wake of various scandals has waned. Consequently, we would expect to see that the portfolios with greater allocations to lower-rated credit sectors perform less well than expected under the equally weighted data sets.

We compare the performance of all of the portfolios over the various investment horizons, so that we can see how, for example, a portfolio optimised with quarterly data would have performed over a longer holding period. The expected horizon outperformances (vs. governments) for all portfolios are presented in Table 12, and the probabilities of holding period underperformance are shown in Table 13.

Interestingly, the portfolios that had been optimised to ensure best performance over quarterly and yearly horizons using the equally-weighted data delivered the best performance over all horizons when subject to the return simulations utilising recently-weighted data. The heavy allocation to A corporate bonds in the "5 Years" and "10 Years" portfolios does not add as much value when returns to this sector are relatively more negative. Nevertheless, all of the portfolios still significantly outperform the government sector.

The results also highlight the weaknesses of traditional mean-variance portfolio optimisation. Portfolios designed to perform optimally under one set of return and covariance assumptions may perform quite differently under slightly different market conditions.

Allocations are highly sensitive to the input parameters, and different inputs would result in different allocations. Moreover, the optimisation process is not designed to produce portfolios that would perform well under varying market conditions. Consequently, mean-variance analysis should not be used as the sole input in a strategic asset allocation exercise. Scenario analysis, extra diversification and the use of common sense and market experience should also be considered when constructing portfolios.

The portfolio performance data above also suggest that observing the outperformance from investments in corporate bonds may take time. All of the non-benchmark portfolios are expected to underperform the benchmark over a monthly horizon 22–28% of the time, and 18–25% of the time over a quarterly horizon. There is a 5–13% probability that such portfolios could underperform during a 12 month holding period. It is only over the longer holding periods of 5 or 10 years that these portfolios would outperform with the greatest certainty.

Nevertheless, time is on the side of most central banks. Despite constant fears that central bank reserves might have to be liquidated during interventions, history shows that for the majority of central banks, reserves are de facto long-term assets. Consequently, the analysis of long-term performance is entirely appropriate.

As a result, the addition of duration-neutral corporate exposure into a fixed income portfolio traditionally populated with government paper may result in a portfolio that increases long-term expected returns in a very risk-efficient manner. Actual returns will clearly vary depending on market cycles, and the amount of accepted spread risk will be entirely dependent on each institution's risk tolerance and culture. Furthermore, the ability of institutions to tolerate periods of underperformance relative to governments will also influence actual allocations into corporate bonds.

7 Conclusions

The Bank of Latvia has allocated portions of its reserves to investment-grade corporate bonds for several years, and we have found through our experience that this is a strategy worth maintaining. We have experienced different market cycles, internal and external economic shocks, interventions and corporate downgrades, yet despite these negative factors, we have consistently been able to outperform benchmark portfolios consisting solely of government bonds.

Our practical experience also supports the conclusions from the preceding analyses. While we do engage in active management practices, we strive to maintain consistent diversified holdings of spread product throughout market cycles, to gain the benefits described above. Similarly, while we focus very much on short-term results, we recognise that we are managing towards long-term investment goals; each passing year seems to confirm the validity of this strategy.

Ultimately, it is up to each individual institution to mandate the amount, if any, of corporate bonds in a reserves portfolio. Our goal here has simply been to outline some of the positive aspects of diversifying away from government bonds, and to discuss some of the practical aspects involved in managing such portfolios.

References

de Beaufort, Roberto (2000) “The Changing Nature of Reserve Management Risks”, *Risk Management for Central Bankers*, Central Banking Publications

de Beaufort, Roberto, Benitez, Sandra and Palomino, Francisco (2002) “The Case for Reserve Managers to Invest in Corporate Debt”, *Central Banking*, Vol. XII, No. 4

Grava, Roberts (2002) “Fewer Government Bonds: Rethinking Reserve Portfolios”, *How Countries Manage Reserve Assets*, Central Banking Publications, 2003

Hickman, Kent, Hunter, Hugh, Byrd, John, Beck, John and Terpening, Will (2001) “Life Cycle Investing, Holding Periods, and Risk”, *Journal of Portfolio Management*, Winter 2001

Johnson-Calari, Jennifer (2000) “Risk Management at the World Bank: Global Liquidity Portfolios”, *Risk Management for Central Bankers*, Central Banking Publications

Keeley, Terrence (2002) “Reserve Management: A New Era”, *Central Banking*, Vol. XII, No. 3

Moody’s Investors Service (2002) “Default & Recovery Rates of Corporate Bond Issuers”, Moody’s Investors Service Global Credit Research, February 2002

Ramaswamy, Srichander (2002) “Managing Credit Risk in a Corporate Bond Portfolio”, *Journal of Portfolio Management*, Spring 2002

Zhou, Chunsheng (2001) “Credit Ratings and Corporate Defaults”, *Journal of Fixed Income*, December 2001

10

Setting counterparty credit limits for the reserves portfolio[1]

Srichander Ramaswamy

Abstract

This paper outlines a framework for establishing counterparty credit limits for the reserves portfolio of a central bank. Taking the view that agency ratings reflect the best credit assessment available for counterparties that a central bank deals with, the paper presents a simple methodology for measuring portfolio credit risk. Subsequently, it outlines how exposure limits for counterparties can be set up to facilitate communication between the middle office and front office, and to ensure that the overall level of portfolio credit risk conforms to the central bank's risk appetite.

1 Background

Foreign currency reserves held by central banks have grown from USD 1,639 billion in 1998 to USD 2,915 billion as at end-2003. Partly as a result of the increase in reserves levels, which has attracted greater public scrutiny, reserves managers have become increasingly conscious of the return on reserves. Other factors that have led to an increased focus on return include a gradual reduction in interest rate levels that has reduced interest income, and awareness that higher reserves holdings increase the cost of carry as opposed to investments in the domestic markets. As a result of the increased attention paid to returns, reserves managers have started debating the inclusion of lower credit quality products in the investment universe for reserves to improve return. At the short end of the yield curve, diversification into bank deposits, certificates of deposit and commercial paper at the expense of T-bills has already taken place. Over the medium to long end of the yield curve, some central banks are even considering investments in corporate bonds to boost returns.

Assuming that the diversification trend into lower quality credit will continue in the coming years, foreign currency reserves will be exposed to greater levels of credit risk than what reserves managers are traditionally accustomed to. This raises the question as to whether reserves managers have the right tools to measure and manage credit risk, and to set appropriate limits for counterparty credit exposure as well as monitor usage against these limits. One can safely argue that reserves managers are familiar with managing liquidity and market risk by establishing a control framework to ensure compliance with the established guidelines. However, in the area of credit risk management, there may still be scope for improvement in the existing procedures followed by reserves managers. As a consequence, to take on more credit risk, reserves managers will have to establish an appropriate framework for measuring credit risk and to set up a credit limit system to ensure that the foreign currency reserves are being managed prudently. This paper attempts to address some of these issues by providing a framework for implementing a counterparty credit risk limit system that is appropriate in a reserves management framework.

---

[1] The views expressed in this article are mine, and do not necessarily reflect those of the Bank for International Settlements.

2 Motivation and outline

Unlike financial institutions, central banks do not face any regulatory requirement to report potential credit-related losses against the available capital. Nor are central banks required to maximising shareholder value. Hence, the need to quantify credit risk in the reserves portfolio is difficult to justify either from a regulatory or a capital allocation perspective. Instead, the desire to measure and manage credit risk can only be motivated by the requirement to limit counterparty credit risk and to ensure that the overall level of portfolio credit risk conforms to the risk appetite of the central bank. To establish counterparty credit limits, an important prerequisite is the ability to measure credit risk both at the counterparty and portfolio level. To meet this requirement, it is first necessary to implement an internal credit risk model that facilitates the assessment and quantification of credit risk embedded in the reserves portfolio. This first step will in turn require an assessment to be made of the credit quality of counterparties against which credit limits exist.

In general, the counterparties of a central bank will be of a high credit quality, and most of these counterparties will be rated by the major rating agencies. As a result of this, the credit exposures for the reserves portfolio will be very different from those of financial firms. A large part of the credit exposure on the balance sheet of financial firms will arise from unrated counterparties. As a consequence, financial firms will be required to maintain an internal credit rating system to assess the creditworthiness of all the counterparties on their balance sheets. To do this, financial firms usually have a team of experienced credit analysts and related support staff. In the reserves management division of a central bank, however, it is sometimes difficult to retain a team of experienced credit analysts due to the mandatory job rotation policy adopted by many central banks.

Faced with the constraint that central bank staff have limited experience assessing the creditworthiness of counterparties, coupled with the fact that the credit counterparties of a central bank for its financial transactions are sovereigns or corporations that are rated by credit agencies, central banks may well prefer to rely primarily on agency ratings to measure credit risk and set appropriate limits. One may even go so far as to argue that for counterparties rated single-A or better, the agency rating reflects the best credit assessment of the counterparty, considering that information about such counterparties is mostly public. Without debating the merits of this view, and assuming that a central bank takes this route for assessing the credit risk exposure to various counterparties, independent credit assessment of counterparties thus becomes less of a requirement. The implementation of an internal credit risk model under this framework would then only require transaction data and specification of the credit risk model parameters for quantifying credit risk.

To a casual reader, the last statement might sound as if implementing an internal credit risk model for the reserves portfolio is a fairly straightforward task, as it only requires the relevant credit risk model parameters to be specified. On the contrary, implementing an internal credit risk model is quite a challenge in practice, because data required for quantifying credit risk are not easily available. For instance, sovereign defaults are scarce, and therefore recovery statistics for defaulted sovereign debt are primarily little more than educated guesses. When aggregating portfolio credit risk, for instance, the likelihood of joint defaults across counterparties has to be assessed. Assessing the likelihood of joint default between two sovereigns or between sovereign and non-sovereign counterparts is more art than science. In each of these cases, either educated guesses or several simplifying assumptions will have to be made.

The main focus of this paper is to develop a framework for establishing counterparty credit limits. To facilitate this, I will introduce a methodology for credit risk measurement that makes use of publicly available information, and I will provide suggestions for choosing the credit risk model parameters that will be required for quantifying credit risk. It is important to mention here that the suggested model parameter values are intended as indicative values that should be refined as additional data become available. The modelling approach used here makes the assumption that credit events can be modelled as a two-state process, meaning that the counterparty either remains in business or is bankrupt. Under this assumption, credit risk is modelled in the default mode framework. For a central bank that is primarily interested in prudently allocating counterparty credit limits, the default mode setting provides a good compromise between accuracy and simplicity. Moreover, if a significant part of credit risk in the reserves portfolio arises from exposures to bank deposits and commercial paper in the short end of the yield curve, then choosing the default mode framework would provide simplicity without compromising on the accuracy of the credit risk estimates.

It is useful to mention here that some central banks tend to differentiate between counterparty limits and issuer limits. Issuer limits can be thought of as the limits set for the purchase of debt issued by some corporate or sovereign entity. Counterparty limits, on the other hand, can be interpreted as the limits established for money market placements, which are primarily with financial institutions. In this paper I will not differentiate between the two, and the limit setting approach presented here will be valid for both.

The rest of the paper is organised as follows. Section 3 introduces the terminology required to understand credit risk modelling. Section 4 discusses counterparty credit risk by introducing the concept of expected and unexpected losses. Aggregating credit risk arising from various counterparties and the steps involved in estimating the likelihood of joint defaults are discussed in Section 5. Section 6 provides a detailed description of how credit risk can be managed by computing risk contribution arising from different counterparties and by setting exposure limits for counterparties. An optimisation framework for allocating credit risk limits for various counterparties is then presented in Section 7. Section 8 outlines a simulation framework for computing portfolio credit risk at some desired confidence level. Finally, Section 9 concludes.

3 Credit risk terminology

Credit risk, in broad terms, refers to the risk of a loss arising from the counterparty or obligor not being in a position to service the debt obligations. Counterparties or obligors are usually distinguished at the branch level of the legal entity which is responsible for servicing the debt obligations. The standard practice is to estimate the potential loss at some desired confidence level over a one-year time horizon. For short-term money market instruments, this insolvency risk encompasses all of the credit risk at the counterparty level. However, for instruments that have maturity greater than one year, such as bonds, additional sources of credit risk can arise. In particular, the mark-to-market loss of a bond resulting from widening credit spreads caused by a change in the market perception concerning the ability of the obligor to service the debt in the future is attributed to credit risk. In most cases, the change in the market perception regarding the obligor's ability to service debt obligations will be either preceded or succeeded by a change in the credit rating of the obligor. This process, usually described as credit migration, will result in a change in the credit spread of the obligor.

If credit migrations are taken into consideration for quantifying credit risk, then credit risk is said to be computed under the migration mode. If only obligor defaults are modelled, then credit risk is said to be computed under the default mode. In computing credit risk at the counterparty level, the following quantities have an important role to play:

- *Exposure amount:* The maximum positive value which could be lost in the event of a counterparty default before accounting for recovery costs. For money market deposits, loans and bond holdings, the exposure amount will be equal to the mark-to-market value of the instrument. For derivatives securities, such as interest rate swaps, futures and options, the exposure amount will be equal to the replacement cost of the instrument.
- *Default probability*: The probability that the counterparty will default on its contractual obligations to repay its debt. Since default probability is a function of the time horizon over which one measures the debt servicing ability, it is standard practice to assume a one-year time horizon to quantify this.
- *Recovery rates*: The extent to which the face value of an obligation can be recovered once the counterparty has defaulted. It should be mentioned here that, in the event of default, the credit market convention is to ask how much of the promised debt is lost rather than how much of it is recovered. This leads us to define the term loss-given-default (LGD) as 1 minus the recovery rate.
- *Credit migration*: Short of default, the extent to which the credit quality of the counterparty improves or deteriorates as expressed by a change in the default probability. This will impact the relative spread versus Libor at which the instrument trades.

For measuring and managing credit risk at the portfolio level, the following additional quantities play an important role:

- *Default correlation*: The strength of the default relationship between two counterparties. This is defined as the correlation between the default indicators for these two counterparties over some specified interval of time, typically one year. An important driver of default correlations is the asset return correlation between counterparties.
- *Risk contribution*: The incremental risk that the exposure of a single asset contributes to the total risk of the portfolio.
- *Concentration risk*: The extent to which a particular counterparty, industrial sector or country contributes to the overall credit risk of the portfolio. Managing this exposure is important to ensure that there is sufficient diversification of credit risk in the portfolio held.

For purposes of quantifying credit risk, two other measures are frequently used in the credit risk literature. These are given below.

- *Expected loss (EL)*: This is the average loss in market value of the asset resulting from credit-related events over the holding period (a one-year horizon is commonly used for the holding period).
- *Unexpected loss (UL)*: This is the estimated volatility of potential losses around the expected loss and is also expressed over a one-year horizon.

For reasons cited in Section 2, I will take the simpler route of credit risk quantification under the default mode for the reserves portfolio. Furthermore, I will only compute the first two moments of the credit loss distribution, namely expected loss and unexpected loss, in order to measure credit risk and to set counterparty credit risk limits. In Section 8 I will introduce a simulation framework to estimate the potential credit loss of the reserves portfolio at some desired confidence level.

4 Counterparty credit risk

The primary objective of developing a credit risk measurement framework for the reserves portfolio is to facilitate prudential risk management by establishing counterparty credit risk limits and identifying sources that give rise to concentration risk. While developing this framework for credit risk measurement, it is necessary to make some simplifying assumptions. One such assumption that is often made in estimating the potential loss arising from credit risk is to model the credit event as a two-state process – default or no default. In this simplified framework, one is interested in estimating the probability that the counterparty will default over a specified time interval. In the event of default, the potential face value of the debt that will have to be written off will also influence the loss estimates. The estimate of insolvency risk at the counterparty or obligor level in this framework will depend on three variables, namely: probability of default (PD), loss-given-default (LGD) and volatility of the LGD estimates, denoted σ_{LGD}.

4.1 Probability of default

The key determinant of the credit risk of a counterparty is the uncertainty regarding the counterparty's ability to service debt obligations as expressed through the default probability. Two possible routes exist for obtaining estimates of the default probability of the counterparty. The first is empirical in nature and requires the public credit-quality rating scheme to be used. The default probability is then inferred from the rating transition matrix given the counterparty's current credit rating. The second is based on the options theory framework, which uses the current estimates of the firm's assets, liabilities and asset return volatility, and is hence related to the dynamics of the underlying structure of the firm.

The advantage of the first method is that the default probability estimates will be more stable, so that the counterparty credit limits will also remain fairly stable if the nature of the business does not change. This is important for business continuity. The disadvantage, however, is that the default probability estimates will be less responsive to changes in the economic cycle, resulting in overestimation or underestimation of the likelihood of default over different time periods. Although greater responsiveness to changes in economic cycles may appear desirable, frequent changes in credit limits for counterparties will disrupt business planning. Hence, the approach followed here will be to assign default probabilities for different credit ratings on the basis of empirical estimates reported by rating agencies.

Table 1: One-year probability of default for various rating grades

Rating grade	Rating description	Standard & Poor's	Moody's	Recommended
1	Risk-free	-	-	0.00
2	Aaa / AAA	0.00	0.00	0.01
3	Aa1 / AA+	0.00	0.00	0.02
4	Aa2 / AA	0.00	0.00	0.03
5	Aa3 / AA-	0.03	0.08	0.04
6	A1 / A+	0.02	0.00	0.06
7	A2 / A	0.05	0.02	0.08
8	A3 / A-	0.05	0.00	0.10
9	Baa1 / BBB+	0.13	0.08	0.13
10	Baa2 / BBB	0.23	0.07	0.23
11	Baa3 / BBB-	0.37	0.46	0.46
12	Ba1 / BB+	0.48	0.67	0.67
13	Ba2 / BB	1.03	0.72	1.03
14	Ba3 / BB-	1.46	2.46	2.46
15	B1 / B+	3.25	3.97	3.97
16	B2 / B	9.37	8.41	9.37
17	B3 / B-	11.49	13.72	13.72
18	Caa-C / CCC	25.25	29.60	29.60

Table 1 shows 18 distinct rating grades and the default probabilities associated with them as estimated by Standard and Poor's and Moody's Investors Service.

Contrary to what one might expect, the empirical default probability estimates in Table 1 do not increase monotonically as one goes down the rating grades. Having this property will have the intuitive appeal that lower credit rating for a counterparty is associated with an increase in the counterparty's probability of default. The last column in Table 1 imposes this condition explicitly by modifying the rating agencies' empirical probability of default estimates. Specifically, for issuers with ratings between Aaa and Aa3, I have assumed that the default probability will increase by 1 basis point (one hundredth of a percentage point) for every one-notch downgrade. From grades Aa3 to A3, PD is assumed to increase by 2 basis points for every one-notch downgrade. For issuers rated Baa1 and lower, the PD values have been chosen to reflect the maximum of the estimates made by Moody's and Standard & Poor's.

4.2 Recovery rate statistics

The other variables that influence credit risk estimates are the expected LGD and its volatility. The expected LGD is bounded between 0 and 1, implying that in the best case scenario the counterparty holding the debt will recover all the money in the event of default, but will recover nothing in the worst-case scenario. In general, estimating the recovery value of defaulted securities is quite difficult, because the payments made to creditors could take the form of a combination of equity and derivatives securities, new debt or modifications to the terms of the surviving debt. Considering that there may be no market for some forms of payments, it may not be feasible to measure the recovery value. Moreover, the amount recovered could take several months or even years to materialise and could potentially also depend on the relative strength of the negotiating positions. As a result, estimating historical averages of amounts recovered from defaulted debt will require some simplifying assumptions to be made.

Moody's, for instance, proxy the recovery rate with the secondary market price of the defaulted instrument approximately one month after the time of default. The motivation for such a definition is that many investors may wish to trade out of defaulted debt. A separate investor clientele may then acquire these instruments and pursue the legal issues related to recovering money from them. In this context, Moody's recovery rate proxy can be interpreted as a transfer price between these two investor groups.

Depending on the data set and time period used for estimating recovery rate parameters, variations in reported recovery rate statistics tend to be quite large. Table 2 shows the typical range within which the parameters LGD and σ_{LGD} lie based on empirical estimates reported in the literature (Moody's Investors Service, 2002; Standard and Poor's, 1999; Altman et al., 2001).

Table 2: Empirical estimates for LGD statistics
in (%)

Security type	**LGD**	$\boldsymbol{\sigma}_{LGD}$
Senior secured bonds	27 to 35	23 to 28
Senior unsecured bonds	48 to 63	23 to 35
Subordinated bonds	65 to 74	20 to 30

If we were to use the above empirical estimates for guidance in choosing the recovery rate statistics for different instruments in the reserves portfolio, one could choose the LGD and σ_{LGD} values for different security types to be roughly in the middle of the ranges reported in Table 2. Since the differentiation on the σ_{LGD} values for different security types is relatively small if we follow this approach, one could choose σ_{LGD} for all security types to be 25%. Such a choice would result in the recovery rate statistics shown in Table 3. Supranational debt and debt issued by sovereign governments denominated in foreign currency can be treated as unsecured bonds, so that the LGD associated with such debt will be 55%. Money market deposits usually have the same priority in a bankruptcy proceeding as unsecured bonds, and hence have the same LGD value. It is important to emphasise here that the recovery rate statistics for different securities provided in Table 3 are only representative values that can be refined as and when improved estimates for these quantities become available.

Table 3: Recommended values for LGD and volatility of LGD
(in %)

Security type	LGD	σ_{LGD}
Money market deposits	55	25
Asset-backed securities	30	25
Unsecured bonds, bills and notes	55	25
Subordinated debt	70	25

4.3 Exposure amount

To quantify credit risk at the counterparty level, we need in addition to the above variables the exposure amount (EA) for the transaction in question. The EA is equal to the maximum amount that can be lost in the event of a counterparty defaulting before accounting for recovery costs. For instruments traded in the cash market, the EA is equal to the current market value of the transaction. It is assumed in this paper that exposure amounts are given in US dollars.

For derivatives transactions, the EA is equal to the positive part of the mark-to-market value of the transaction (also referred to as the replacement cost), plus an add-on. This add-on is meant to reflect the fact that, unlike cash instruments, a derivate transaction's exposure amount at the time of counterparty default could differ greatly from the current exposure amount. The add-on is usually computed on the basis of the average replacement cost of the instrument over a one-year time horizon. It is useful to note here that EA will usually be a small fraction of the notional amount of the derivatives transaction (typically in the range of 1-3% of the notional amount). If netting arrangements are present, this amount will be further reduced. In this paper, I will not discuss the details of how the add-on can be computed for derivatives transactions.

4.4 Expected and unexpected losses

In the previous sections I identified the important variables that influence credit risk at the security level. In this section I will focus on quantifying credit risk at the security level. In broad terms, risk is associated with the potential financial loss that can arise from holding the security, the exact magnitude of which is difficult to forecast. As a result, it is common to describe the potential loss in value using an appropriate probability distribution whose mean and standard deviation serve as useful measures for risk quantification. This practice is well known in the equities market, where investors focus on market risk that model variations in stock return. This leads us to quantifying the market risk measures through expected return and standard deviation of return. Under the assumption that equity returns are normally distributed, the realised return will lie within one standard deviation of the expected return with two-thirds probability.

Quantifying credit risk at the security level is similar in principle. Borrowing the principle underlying market risk quantification from the equities market, it has become common practice to quantify credit risk at the security level through the mean and standard deviation of the credit loss distribution. However, there is an important difference between the two risk measures. This pertains to the distribution of credit loss, which unlike market risk is far from normal. Hence, deviations from the expected loss by one standard deviation can occur more frequently than on one in three occasions. The credit market convention is to refer to the standard deviation of loss resulting from credit events as unexpected loss (UL), and the average loss as expected loss (EL). If the credit events modelled are only default and no default, then EL and UL will reflect the mean and the standard deviation of the credit loss distribution under the default mode.

I will not provide detailed derivation of the relevant formulas for EL and UL.[2] However, it is important to mention here that in deriving the expressions for EL and UL under the default mode, the simplifying assumption that recovery rate process and default process are independent is made. In practice, empirical results on the relationship between obligor defaults and recovery rates suggest that these two variables are negatively correlated. The relevant formulas for computing EL and UL under the assumption that default process and recovery process are independent are given below:

$$EL = EA \times PD \times LGD \tag{1}$$

$$UL = EA \times \sqrt{PD \times \sigma_{LGD}^2 + LGD^2 \times \sigma_{PD}^2} \tag{2}$$

In equation (2), σ_{PD}^2 is the variance of the default process. Since I assume that the default event is a two-state process, the default process will have a Bernoulli distribution. In this case, the variance of this process will be given by

$$\sigma_{PD}^2 = PD \times (1 - PD) \tag{3}$$

It should be noted here that if PD values reflect the default probability estimates over a one-year time horizon, then EL and UL will quantify credit risk over the same period. Considering that short-term money market instruments expire in less than one year, credit risk will be overestimated. However, if we take the view that money market instruments will be rolled over when they mature, then we can argue that EL and UL computed under the assumption that all instruments have at least one-year maturity provide reasonable estimates of the true credit risk under the default mode.

---

[2] Interested readers are referred to Ong (1999) for the derivations of the respective expressions.

5 Portfolio credit risk

The previous section primarily focused on identifying the elements of credit risk and quantifying credit risk in terms of expected and unexpected losses at the security level. A natural extension of this analysis will be to quantify credit risk at the reserves portfolio level. To do this, we will have to model the co-movement of defaults between two or more counterparties. This brings us to the topic of correlated credit events, which are fundamental to the modelling of portfolio credit risk. In practice, measuring these correlations directly is difficult, if not impossible. Standard techniques used to estimate them follow an indirect approach that utilises the correlation between the variables that drive credit events. The variable that is usually considered to drive credit events is the asset return of the firm. Since asset returns are not directly observable, the method used to estimate asset return correlation between different obligors is a much-debated topic.

In this section I will provide the relevant equations for computing portfolio credit risk under the default mode. I will also indicate how the default correlation between obligor pairs, which is required to compute portfolio credit risk, can be determined. A simple technique to derive approximate asset return correlations between obligor pairs is also presented in this section.

5.1 Quantifying portfolio credit risk

We have seen in Section 4 that credit risk at the security level can be quantified in terms of the mean and the standard deviation of the loss distribution resulting from counterparty defaults. Quantification of credit risk at the portfolio level follows a similar approach. Once again we are interested in the loss distribution of the portfolio arising from counterparty defaults. The mean and standard deviation of the portfolio loss distribution, denoted expected portfolio loss (EL_P) and unexpected portfolio loss (UL_P) are used to quantify portfolio credit risk. The relevant formulas to compute these quantities, assuming that *n* securities are held in the reserves portfolios, are given below.

$$EL_P = \sum_{i=1}^{n} EL_i \tag{4}$$

$$UL_P = \sqrt{\sum_{i=1}^{n}\sum_{k=1}^{n} \rho_{ik} \times UL_i \times UL_k} \tag{5}$$

In equation (5) the term ρ_{ik} denotes the default correlation between the *i*th and the *k*th security in the reserves portfolios. In practice, the use of default correlation will tend to overestimate the standard deviation (or the unexpected loss) of the credit loss distribution under the default mode if recovery rates across different counterparties are uncorrelated. Using loss correlation instead of default correlation will remove this bias (see Ramaswamy, 2003). However, if we assume that recovery rates across different counterparties are positively correlated, then loss correlation tends to approach default correlation values.

Considering that the default correlation between two securities that belong to the same counterparty will be equal to 1, we only need to estimate default correlation between counterparties in order to compute portfolio credit risk. Hence, I will only focus only on computing default correlation at the counterparty level rather than at the security level in the following section.

5.2 Estimating default correlation

Default correlation measures the strength of the default relationship between two counterparties. It answers the important question of how default by one counterparty can influence the contemporaneous default of other counterparties. An increase in default correlation between two counterparties will increase the unexpected loss of a two-bond portfolio, assuming all other parameters remain the same. Default correlation between two counterparties is formally defined as the correlation between the default indicators for these two counterparties over some specified interval of time, typically one year. Invoking the standard definition of correlation between random variables, we can derive the following relation for default correlation between the *i*th and *k*th counterparties:

$$\rho_{ik} = \frac{P(D_i \cdot D_k) - PD_i \cdot PD_k}{\sqrt{PD_i(1-PD_i)} \cdot \sqrt{PD_k(1-PD_k)}} \tag{6}$$

In equation (6), PD_i and PD_k are the one-year default probabilities of the *i*th and *k*th counterparty in the portfolio, while $P(D_i \cdot D_k)$ denotes the probability of joint default.

In spite of the apparently simple expression given above, estimating default correlation is difficult in practice because of the lack of joint default data. Hence, we need to be able to estimate the probability of joint defaults using an indirect approach. To do this, some simplifying assumptions have to be made, and variables identified that drive defaults to facilitate the indirect estimate of the probability of joint defaults. A standard assumption that is often made in practice is that default correlation between counterparties can be inferred from asset return correlation.

The intuition behind the use of asset returns of the counterparty to estimate the probability of joint defaults lies in Merton's structural model for default. In Merton's model, the firm's default is driven by changes in the asset value of the firm. As a result, the correlation between the asset returns of two obligors can be used to compute the default correlation between the two obligors. The way this works in practice is to use correlation between asset returns for two counterparties to estimate the probability of joint default of the counterparties. Using this information, one can easily compute the default correlation between any two securities using equation (6).

If we make the simplifying assumption that the joint distribution of asset returns is bivariate normal, then the probability of joint default of two counterparties *i* and *k* having asset return correlation ρ_{ik}^{α} is given by

$$P(D_i \cdot D_k) = \frac{1}{2\pi\sqrt{1-(\rho_{ik}^{\alpha})^2}} \int_{-\infty}^{D_i}\int_{-\infty}^{D_k} \exp\left(-\frac{(x^2 - 2\rho_{ik}^{\alpha}xy + y^2)}{2[1-(\rho_{ik}^{\alpha})^2]}\right) dx\,dy \tag{7}$$

Since the probability of default for each counterpart is known given its current rating, the limits for the integrals in equation (7) can be computed by solving the following expressions:

$$PD_i = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{D_i} \exp(-\tfrac{1}{2}z^2)\,dz \tag{8}$$

$$PD_k = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{D_k} \exp(-\tfrac{1}{2}z^2)\,dz \tag{9}$$

It is of interest to note here that the above integrals can be computed using any standard numerical package.

5.3 Estimating asset return correlation

To estimate the probability of joint defaults, we suggest using asset return correlation between counterparties, since they can be considered as drivers of defaults in Merton's structural model for bankruptcy. However, considering that asset returns themselves are not directly observable, further simplifications will be required to estimate asset return correlation across firms. One approach is to use factor models for estimating asset return correlation between firms. From the practitioner's perspective, it can be seen that the mechanics involved in the process of quantifying portfolio credit risk are susceptible to considerable model risk.

A possible route to obtaining the asset return correlation between different counterparties would be to buy this data from vendors such as Moody's KMV. However, the errors on these estimates for highly-rated sovereign counterparties may be high because some of the underlying assumptions involved in computing the asset return correlation may not hold for such counterparties. Since my intention in this paper is to propose a modelling framework for credit risk that makes use of public information, I will suggest a technique to derive approximate asset return correlation between counterparties using a factor model.

A standard practice followed in the literature is to assume that asset returns are generated by a set of common and idiosyncratic factors. Since idiosyncratic factors are firm-specific, asset return correlation only arises from exposure to common factors. Given any two counterparties, I will identify a set of common factors, and the exposure to these factors will contribute to the asset return correlation between the counterparties. The set of common factors for asset returns that I consider here are as follows:

- *Systematic risk component (SYS)*: The asset returns of all firms depend to a certain extent on the health of the global economy. The systematic risk component models the interdependencies of the firm's asset returns with this common factor.
- *Counterparty type component (CTY)*: This component captures the herding behaviour that is often observed in financial markets. For instance, trouble experienced by some large corporates can affect the stock valuations of all corporates. Similarly, a negative remark concerning a particular debt-issuing agency can have an impact on the asset prices of other agencies issuing debt. The counterparty type component captures this co-movement of asset returns. To identify this component of asset return correlation, I will distinguish between counterparty types consisting of sovereign institutions, supranational institutions, agencies and corporates.
- *Regional factor component (REG)*: This component of asset return correlation is meant to capture the regional bias present in the asset returns of different firms. For instance, the economic problems in Argentina could have an effect on corporate and sovereign borrowers in Latin America. To capture this component of asset return correlation, I will identify the following regional factors: Latin America, North America, Europe, Africa and Asia/Pacific. For supranational institutions, this component of asset return correlation will be set to 0 as it is difficult to motivate such a factor for supranationals.
- *Corporate sector component (COR)*: Empirical evidence suggests that corporates involved in similar business exhibit greater co-movement in asset returns. To capture this additional source for asset return correlation between corporates, I shall consider the following sub-classification of corporate counterparties (see Lehman Brothers 2001 for industries that fall under each classification):
  a. Basic and capital goods industries (BAC)
  b. Consumer cyclical and non-cyclical (CCN)
  c. Communications and technology (COT)
  d. Energy (ENE)

e. Transportation (TRA)
f. Utilities (UTL)
g. Banking (BNK)
h. Brokerage (BRO)
i. Financial companies (FIN)
j. Insurance (INS)
k. Asset-backed sector (ABS).

Assuming that asset return correlation between two counterparties *i* and *k* is the sum of the asset return correlation arising from each of the common factors identified above, the total asset return correlation between the counterparties is given by

$$\rho_{ik}^{\alpha} = \rho_{ik}^{SYS} + \rho_{ik}^{CTY} + \rho_{ik}^{REG} + \rho_{ik}^{COR} \tag{10}$$

In equation (10), the contribution of each of the common factors will have to be chosen in order to determine the asset return correlation between two counterparties. Although several intuitive choices are possible so that they reflect the actual estimates for asset return correlation, I provide some indicative values for these choices. Considering that asset return correlations between counterparties are typically in the range of 10-45%, I suggest the following representative values: 10% for the systematic risk component, 5% for the counterparty type component, 10% for the regional factor component, and 20% for the corporate sector component.

The advantage of taking this simple approach is that it is possible to pre-compute default correlation tables (as asset return correlation is limited to a set of eight discrete values) for various rating grades, so that simple look-up tables can be constructed to infer default correlation between counterparties. Table 4 shows the asset return correlations between a set of legal entities derived using the factor decomposition given by equation (10). It is interesting to note that the asset return correlations given in Table 4 turn out to be reasonable approximations in most cases to the asset return correlations computed using a much more rigorous approach and supplied by data vendors.

Table 4: Asset return correlation estimates between different legal entities

Legal entity	Daimler-Chrysler	General Dynamics	AT& T Corp	UBS AG	Citigroup Inc	JP Morgan Chase
Daimler-Chrysler	1.0					
General Dynamics	0.15	1.0				
AT&T Corp	0.15	0.25	1.0			
UBS AG	0.25	0.15	0.15	1.0		
Citigroup Inc	0.15	0.25	0.25	0.35	1.0	
JP Morgan Chase	0.15	0.25	0.25	0.35	0.45	1.0

6 Managing credit risk

The discussion so far has focused on measuring credit risk, which, although important, is not our end objective. The motivation for measuring credit risk is the intention to identify and mitigate risk concentrations in the portfolio. This brings us to the next topic, namely managing credit risk, which deals with ways to avoid risk concentrations through the prudent setting of credit limits for various counterparties. For the purposes of managing credit risk, proper tools are required that will allow us to do the following:

- Set exposure limits for various counterparties.
- View credit risk exposures across maturities, counterparties, rating grades, countries and geographical regions.
- Carry out "what-if" scenario analysis.
- Compute the risk contribution of individual trades to the portfolio credit risk.

In this section I will develop relevant formulas that will allow us to slice and dice the risk across different dimensions of the portfolio so that credit risk arising from different sources can be managed. This in turn will facilitate the setting-up and monitoring of credit risk limits.

6.1 Unexpected loss contribution

For the purposes of managing portfolio credit risk, one is usually interested in knowing how much incremental risk a single transaction will bring to the portfolio. In credit risk terminology, this is referred to as the unexpected loss contribution (ULC). The ULC is defined as the incremental risk that the exposure of a single transaction contributes to the portfolio's total risk. In mathematical terms, this is given by

$$ULC_i = UL_i \times \frac{\partial UL_P}{\partial UL_i} \tag{11}$$

By performing the actual differentiation, it is easy to show that the ULC arising from a single transaction is given by the following equation:

$$ULC_i = \frac{UL_i \sum_{k=1}^{n} UL_k \times \rho_{ik}}{UL_P} \tag{12}$$

Using equations (11) and (12), it can be clearly demonstrated that the following relation holds:

$$UL_P = \sum_{i=1}^{n} ULC_i \tag{13}$$

6.2 Disaggregating unexpected loss across maturities

For the purposes of monitoring and setting up credit risk limits, one may wish to disaggregate the unexpected loss across different maturity bands. For instance, we may consider disaggregating the ULC across the following maturity bands:

T_1 = Less than 1 month
T_2 = 1 month to less than 6 months
T_3 = 6 months to less than 1 year
T_4 = 1 year to less than 2 years
T_5 = 2 years to less than 3 years
T_6 = 3 years to less than 5 years
T_7 = 5 years to less than 7 years
T_8 = 7 years to less than 10 years
T_9 = 10 years or greater

Let us denote the ULC arising from insolvency risk of the *i*th instrument that matures in *t* years as $ULC_i(t)$. Using this notation, the ULC across the maturity band T_k arising from ins olvency risk will be given by

$$ULC(T_k) = \sum_{\substack{i=1 \\ t \in T_k}}^{n} ULC_i(t) \tag{14}$$

It is easy to show that the following relation also holds:

$$UL_P = \sum_{k=1}^{9} ULC(T_k) \tag{15}$$

6.3 Disaggregating unexpected loss across rating grades

Let us denote the ULC arising from insolvency and migration risk of the *i*th instrument whose rating grade is *r* as $ULC_i(r)$. Using this notation, the ULC across the rating grade R_k arising from insolvency risk will be given by

$$ULC(R_k) = \sum_{\substack{i=1 \\ r \in R_k}}^{n} ULC_i(r) \tag{16}$$

Since I have modelled 18 rating grades excluding the default state, it is easy to show that the following relation holds:

$$UL_P = \sum_{k=1}^{18} ULC(R_k) \tag{17}$$

Following the same indexing principle, one can easily derive relations for disaggregating the portfolio credit risk across other variables of the portfolio such as counterparties, legal entities, countries or geographical regions. To avoid repetition, I will skip the details. In the next section I will discuss how exposure limits can be set up for managing credit risk across various risk dimensions.

6.4 Setting up exposure limits

The ultimate objective of measuring credit risk is to be able to manage credit risk so that there is a reasonable diversification of this risk across the portfolio holdings. To achieve this objective, one can set up exposure limits for counterparties and subsequently monitor usage against these limits. Subsequently, it is important to translate the credit limits initially set up in terms of some suitable risk measure into measures that are meaningful to dealers. One such measure that serves this purpose is the loan equivalent exposure (LEE). This is defined as the amount of loan exposure (i.e. term loan or short-term money market deposit) to the counterparty that requires the same economic capital as the transaction in question. Essentially, this is meant to correct for different LGD characteristics of instruments, so that a global limit can be set for each counterpart against whom the total LEE of all transactions can be compared.

At the counterparty level, the LEE of the *i*th transaction is given by

$$LEE_i = EA_i \times \frac{UL_i^{transaction}}{UL_i^{deposit}} \qquad (18)$$

In equation (18), $UL_i^{transaction}$ is the unexpected loss associated with the transaction and $UL_i^{deposit}$ is the unexpected loss if the transaction with the counterparty is a money market deposit for the same exposure amount EA_i. If the counterparty happens to be a sovereign country, then the deposit is assumed to be in the domestic currency of the country. Since credit risk limits will be set in terms of LEE, dealers will be able to conduct a "what-if" analysis to check the loan equivalent exposure for a given nominal amount of the intended transaction, in order to verify whether credit limits are available for the transaction. For those transactions that have no risk contribution (for example, if the counterparty is the US Treasury), the LEE will be undefined, as both numerator and denominator will be 0. In this case, one can assume that the LEE is equal to the exposure amount of the transaction.

For the purposes of setting credit limits for counterparties, many institutions follow the practice of distinguishing counterparties at the branch level and at the legal entity level. The notion of a branch encompasses commercial bank branches, sovereign issuers, central bank counterparties, corporate borrowers and supranational institutions. A branch is assumed to have the following characteristics:

- A branch is the largest unit for which insolvency risk can be measured.
- A branch belongs to a legal entity and its insolvency risk is assumed to be the same as that of the legal entity.

A legal entity, on the other hand, is an incorporated entity that has an insolvency risk associated with it. Legal entities are assigned rating grades 1 to 18 to reflect the probability of default.

To control credit risk, LEE limits will have to be set up at the counterparty level. Additionally, one could also set up the loan exposure limits for counterparties across different maturity bands. Details of these bands are briefly described below.

6.4.1 Counterparty limit

Counterparty limits are set to control concentration risk to counterparties. These limits are set in terms of LEE. It is also common practice to set exposure limits at the legal entity level. In this case, monitoring the limit utilisation requires aggregating LEEs arising from all branches of the legal entity.

6.4.2 Limits across maturity buckets

LEE limits for each counterpart can be further set across various maturity bands. This limit may be required to ensure that a few long-dated instruments do not consume the exposure limit set for the counterparty. By redistributing the exposure limit set for the counterparty across specified maturity bands, any desired maturity exposure profile can be enforced for the counterparty.

7 An optimisation framework

In Section 6 I indicated an approach to setting up a credit risk limit system, which will make it easier to examine credit risk exposures across different risk dimensions. In establishing the appropriate credit limits for various counterparties, return considerations as well as the rating and the perceived importance of the counterparty will play a role. In this section I assume that qualitative and objective considerations have been used to determine existing business volume or exposure limits. Given such an initial exposure limit profile for various counterparties, I will present an optimisation framework that will facilitate the review of existing credit limits with the objective of improving the credit diversification of the reserves portfolio.

Before proceeding to discuss the formulation of the optimisation problem, we need to introduce some notations. In Section 6 the loan equivalent exposure of the *i*th transaction with a counterparty was denoted LEE_i . In order to be able to identify the counterparty associated with the transaction, I will denote the *i*th transaction with the *k*th counterparty as $LEE_k(i)$. Assuming that there are *m* transactions associated with the *k*th counterparty, the total loan equivalent exposure to the *k*th counterparty is given by

$$LEE_k = \sum_{i=1}^{m} LEE_k(i) \tag{19}$$

The unexpected loss for a short-term deposit with the *k*th counterparty is given by

$$UL_k^{deposit} = EA_k \times \sigma_k \tag{20}$$

In equation (20), σ_k is given by

$$\sigma_k = \sqrt{PD_k \times \sigma^2_{LGD,k} + LGD_k^2 \times \sigma^2_{PD,k}} \tag{21}$$

Using equation (20), the total loan equivalent exposure to the *k*th counterparty is given by

$$LEE_k = \frac{1}{\sigma_k} \sum_{i=1}^{m} UL_{k,i}^{transaction} \tag{22}$$

Considering that the unexpected loss associated with *m* transactions with a given counterparty is equal to the sum of the unexpected losses of each transaction (since default correlation in this case is 1), equation (22) can be written as

$$UL_k = \sigma_k \times LEE_k \tag{23}$$

In equation (23), UL_k is the unexpected loss arising from all transactions with the *k*th counterparty. It will be assumed here that there are exposures to *N* legal entities in the portfolio. I will now introduce one additional variable that is to be defined below, which will be required in the optimisation problem formulation:

$$\mu_k = PD_k \times LGD_k \tag{24}$$

In equation (24), the variable μ_k can be interpreted as the expected loss associated with a USD 1 short-term loan to the *k*th counterparty.

We are now in a position to formulate the optimisation problem for computing the risk exposure limits to various counterparties. If we denote the existing credit risk limits set in terms of LEEs to the *k*th counterparty as LEE_k^{now}, then the optimisation problem to compute the improved exposure limits LEE_k is given below:

$$\text{Minimise} \quad \sum_{i=1}^{N}\sum_{k=1}^{N} \sigma_i \times \sigma_k \times LEE_i \times LEE_k \times \rho_{ik}$$

Subject to the constraints

$$\sum_{k=1}^{N} LEE_k = \sum_{k=1}^{N} LEE_k^{now} \tag{25}$$

$$\sum_{k=1}^{N} \mu_k \times LEE_k \geq \sum_{k=1}^{N} \mu_k \times LEE_k^{now} \tag{26}$$

$$0.9 \times LEE_k^{now} \leq LEE_k \leq 1.1 \times LEE_k^{now} \tag{27}$$

One can easily recognise the objective function of the above optimisation problem as the square of the unexpected loss of the portfolio from counterparty credit risk. The constraint equation (25) ensures that there is no reduction in the total LEE limit. One can note that the objective function can be minimised at the expense of a reduction in expected earnings by increasing the limits for risk-free counterparties such as the US Treasury. To avoid this, I impose constraint (26). This constraint forces the expected loss of the optimal credit portfolio to be no lower than that of the existing portfolio. The inequality constraint (27) is meant to ensure that the recommended limit changes are marginal by restricting the limit changes to within ±10% of the existing limits for each counterpart.

Before concluding this section, I would like to point out that the optimisation problem formulated above is a quadratic programming problem that is easy to solve using standard software tools. In the next section I will present a simulation framework to evaluate credit risk under the default mode of the reserves portfolio at some desired confidence level.

8 Credit loss simulation

So far in this paper I have discussed how credit loss under the default mode can be quantified in terms of the mean and standard deviation of the loss distribution. Unlike market risk, which can be approximated by a normal distribution, the shape of the loss distribution arising from credit risk is closer to that of a beta distribution for some suitable parameter choices. Knowing the unexpected loss of the credit loss distribution, it is not possible to estimate the potential credit losses at some desired confidence level. To make an estimate of the potential credit-related losses at a given confidence level, it becomes necessary to perform a simulation so that the loss distribution can be generated. Given the simulated loss distribution, it is easy to compute the potential credit losses at any desired confidence level. In this section I will discuss how to perform such a simulation and how to determine the credit loss under the default mode at some desired level of confidence.

Briefly, the main computational steps involved in performing the Monte Carlo simulation to generate the credit loss distribution under the default mode are the following:

1. Simulate the correlated random numbers from a multivariate normal distribution function, which model the joint distribution of asset returns of the obligors in the portfolio.
2. Identify, for each simulation run, those counterparties for which the asset returns are greater than the default threshold D_i for the counterparty.

3. Compute, for each counterpart that is flagged as having defaulted in the simulation run, a random loss on default value by sampling from a beta distribution function that has a mean equal to the LGD and standard deviation equal to σ_{LGD}.
4. Sum the losses arising from each instrument held whose counterparty defaulted for a particular simulation run.

Repeating this simulation run many times and computing the portfolio's credit loss under each simulation run makes it possible to generate the distribution of portfolio credit loss under the default mode. It is to be noted here that the default threshold D_i for the *i*th counterparty can be determined using equation (8), given the probability of default of the counterparty.

Let us consider for the purpose of illustration that the simulated asset return for a particular counterparty is -2.5. Suppose the counterparty's default threshold D_i happens to be -2.2, then for the simulation run in question the counterparty will be deemed to be in default. Given that the counterparty has defaulted in the simulation run, we can try to determine the random loss on a specific security held by the counterparty by drawing a random number from a beta distribution with the following parameters:

$$\alpha = \frac{\mu^2(1-\mu)}{\sigma^2} - \mu \tag{28}$$

$$\beta = \frac{\alpha}{\mu} - \alpha \tag{29}$$

In the above equations, μ and σ correspond to the LGD and the volatility of LGD, both expressed as a fraction. It should be noted here that the random LGD being simulated should be the same in any given simulation run for all securities that belong to the same counterparty and have the same LGD values. Considering that in general one might hold securities from a particular counterparty with different LGD values, depending on the seniority in the debt structure, it is important to impose this condition in the simulated random losses. One way to impose this constraint is to sample from a uniform distribution over the interval 0 and 1, and then to compute the inverse function of the beta distribution function for the simulated uniform random variable. The parameters of the inverse beta distribution should correspond to the recovery rate statistics for the security in question.

Specifically, let $f(x, \alpha, \beta)$ denote the density function of the beta distribution and z the random value generated by sampling from the uniform distribution. To implement the condition that instruments with different seniority have recovery values consistent with their bankruptcy priority, we need to compute the following integral expression:

where
$$y = \int_0^z f(x, \alpha, \beta)\,dx \tag{30}$$

$$f(x,\alpha,\beta) = \begin{cases} \dfrac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\cdot\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}, & 0 < x < 1 \\ 0, & otherwise \end{cases} \tag{31}$$

In equations (30) and (31) the parameters α and β correspond to the recovery rate statistics for the instrument in question. The value y computed as in equation (30) will then correspond to the random LGD having a beta distribution.

To compute the portfolio credit loss L_k for any simulation run, we need to sum the losses across all securities in the reserves portfolio. Assuming that the number of simulation runs is equal to N, one can compute the expected and unexpected loss of the portfolio from the simulations as follows:

$$EL_P = \frac{1}{N}\sum_{k=1}^{N} \ell_k \quad (32)$$

$$UL_P = \sqrt{\frac{1}{N-1}\sum_{k=1}^{N} \left(\ell_k - EL_P\right)^2} \quad (33)$$

If we assume that the simulated losses are sorted in ascending order, then computing the credit values at risk at 99% confidence level over a one-year horizon boils down to picking the loss that corresponds to the 0.99N value of the simulated data. For instance, if the number of simulations is equal to 100,000, then the loss corresponding to the 99,000th value is the credit risk at 99% level of confidence.

9 Conclusions

In this paper I have outlined a simple framework for establishing counterparty credit limits. This required the development of an internal credit risk model under the default mode for the reserves portfolio. Arguing that the major difficulty involved in implementing an internal credit risk model arises from a lack of good estimates of the credit risk model parameters, this paper has provided some indicative choices for these parameters that are relevant in a reserves management framework. This paper has also provided suggestions for disaggregating credit risk across different dimensions. An optimisation approach for fine-tuning risk limits across different counterparties has also been presented. Finally, I have discussed how the distribution of the portfolio credit loss can be generated by performing a simulation, and how credit Value-at-Risk at some desired confidence level can be computed.

References

Altman, E., A. Resti and A. Sironi (2001): "Analyzing and Explaining Default Recovery Rates", a report submitted to The International Swaps & Derivatives Association.

Moody's Investors Service (2002): *Default and Recovery Rates of Corporate Bond Issuers.*

Ong, M. K. (1999): *Internal Credit Risk Models: Capital Allocation and Performance Measurement*, Risk Books.

Ramaswamy, S. (2003): *Managing Credit Risk in Corporate Bond Portfolios: A Practitioner's Guide*, John Wiley.

Standard & Poor's (1999): *Recovering your Money: Insights into Losses from Defaults.*

11

Multi-factor risk analysis of bond portfolios*

Lev Dynkin and Jay Hyman

Abstract

This paper describes the Lehman Brothers' Global Fixed Income Risk Model and other applications of interest in central banking. The paper discusses the importance of defining risk sensitivities in a manner that has intuitive meaning to portfolio managers and corresponds to the parameters managers use to make portfolio allocations, such as key rate durations and spread duration to various asset classes. The importance of reports with extensive diagnostics that accurately assess both systematic and security-specific risk is also discussed.

1 Introduction

In the early 1990s, central bank portfolios were composed almost exclusively of gold and short-term treasury securities (primarily US Treasuries). Ten years later, we see significant allocations to non-treasury asset classes, particularly US agency issues, top-rated asset-backed securities and holdings of supranational issues, with some banks also delving into credit and mortgage securities. Portfolio allocations to gold and short-term US Treasuries have become correspondingly smaller.

Several developments contributed to this trend (Figure 1). Currency interventions have become less frequent and less dramatic, reducing the liquidity demands on central bank reserves. The advent of the euro fuelled a deepening of European credit and securitised markets and created a potential rival to the US dollar's status as the world's predominant reserve currency. Equally important, the successful global campaigns against inflation in the 1990s eroded much of the need to hold gold and greatly depressed its value. Figure 2 documents the return advantage experienced by various categories of spread product relative to duration-matched US Treasuries.

For countries in which monetary policy is now set by the ECB, national central banks have reduced policy constraints on the allocation of reserve holdings, and may find themselves increasingly judged based on the performance of their investment portfolio. This has led

Figure 1: The Evolution of Reserve Management over the Past Ten Years

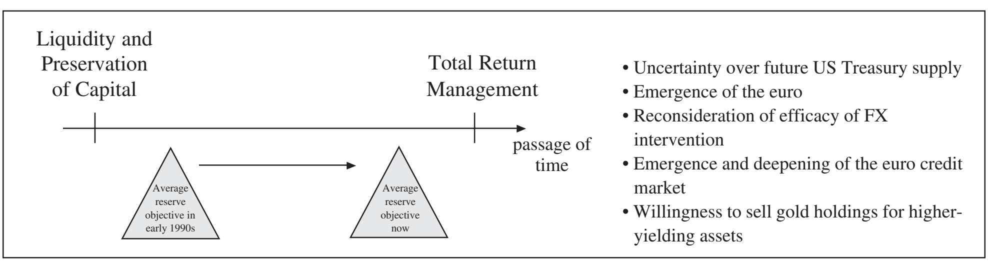

* Dynkin, Hyman and Wu (1999), Naldi, Chu and Wang (2002), and Berd and Naldi (2002) describe risk modelling at Lehman Brothers in detail. This paper draws material directly from these prior Lehman publications.

Figure 2: Monthly Total Return Over Duration-matched US Treasuries
August 1988-September 2002 (bps)

Asset Class	Average	Standard Deviation	Annualised Information Ratio [1)]
Agencies	3.3 bp	18.0 bp	0.64
Sovereigns (≥ A)	0.8	58.1	0.05
Supranationals (≥ Aa)	3.1	26.8	0.41
Asset-backeds (Aaa)	4.3	22.9	0.66
Corporates (≥ A)	1.4	45.0	0.11

Source: Carmel, Dynkin, Hyman and Phelps (2002).
1) The monthly information ratio is the average monthly excess return divided by the standard deviation of monthly excess returns. The annualised information ratio is the monthly information ratio multiplied by the square root of 12.

national central banks in the euro area to be particularly active in shifting portions of their dollar holdings from US Treasuries to high-grade dollar spread sectors.

Modern methods of portfolio risk assessment allow central banks to tap into the higher yields offered by high-grade spread product without impinging on their overarching responsibility to ensure the stability of the national currency and other macroeconomic objectives.[1] The current low yields on short-term US Treasuries have given additional impetus to the general movement by central banks toward diversification of their dollar reserves out of Treasury securities.

In addition to these long-term trends, the events of the past few years have further heightened interest in risk management for central bank portfolios. Yield volatility of the 2-year on-the-run US Treasury note, a mainstay of central bank reserves, is at record levels (see Figure 3). Realised returns of short-term US Treasuries have also been unusually volatile over the 1999-2002 period. Moreover, the current low yield on 2-year US Treasuries may not

Figure 3: Implied Volatility: 3-Month Option on the 2-Year On-the-Run Treasury Note, Jan 1, 1995 to Jan 1, 2004

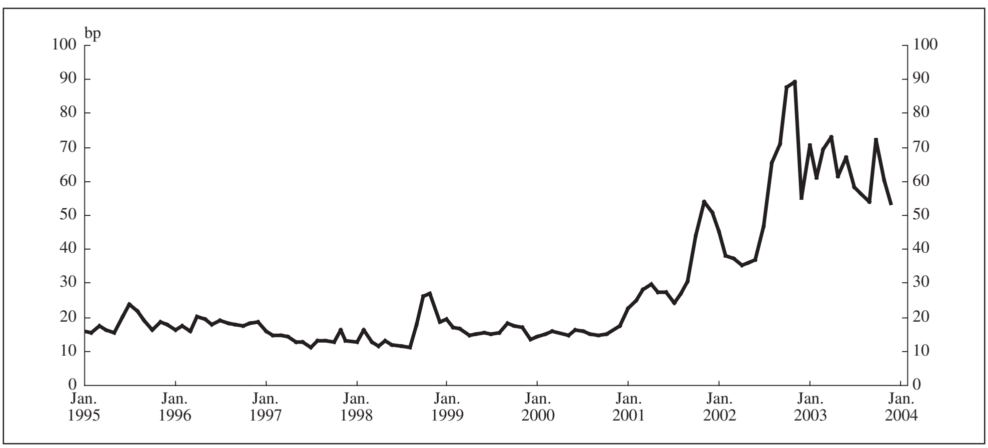

[1] Carmel, Dynkin, Hyman and Phelps (2002) discuss management of central bank reserves that include both US Treasury and non-Treasury holdings.

provide a carry return sufficient to cover even relatively modest capital losses that could occur in the event of a rate spike. This has led some conservative investors to be concerned about the possibility of negative realised returns for short-duration portfolios in the event that dollar interest rates rise in 2004 as forcefully as they fell in 2001.[2] Given the current level of yield volatilities for short-term US Treasury securities, such a scenario cannot be completely disregarded.

The Lehman brothers risk model

The increased prevalence of non-parallel shifts in the US Treasury curve and the movement away from Treasury-only investment policies have greatly increased both the complexity and the importance of risk modelling for central bank portfolios. The Lehman Brothers Risk Model for US dollar portfolios is a multi-factor model which projects portfolio risk over the next month by calculating values for a large number of observable portfolio characteristics (risk sensitivities), and then maps these characteristics into an estimated return volatility. Examples of the observable characteristics used as risk sensitivities in the model are: (i) distribution of interest rate exposures along the yield curve, (ii) portfolio allocations to various sectors and credit qualities, and (iii) a variety of factors concerning mortgages and callable or putable securities.

We feel the model has three important advantages:

- Risk sensitivities are defined exogenously so that they correspond tightly to the parameters managers use to make portfolio allocations (durations, spread durations, convexity, etc.). We then determine what change in the risk factors best fits the observed returns of individual securities in Lehman Indices.
- Intuitive, rather than independent, risk factors are used. For example, credit risk is measured by a set of spread durations, one for each credit tier, rather than by a set of independent "hidden" factors of market risk. Once again, this is more closely aligned to how investors think about markets and risks, and makes diagnostics more readily interpretable.
- The model is calibrated to individual bond returns, and the unexplained residuals are used to estimate security-specific risk. Most risk models ignore all sources of risk not accounted for by the systematic risk factors. The model under consideration goes one step further and estimates portfolio risk due to security-specific idiosyncratic sources. As events at Enron and WorldCom illustrate, firm-specific events such as accounting irregularities can generate substantial negative returns, which at times can be sizeable enough to affect (at a more moderated level) even the most diversified portfolio. This type of risk is estimated from historical residuals of bond returns after all risk factors have been applied. They should be uncorrelated to each other.

In all, the model has well over 100 factors. Many risk models exist with a small number of high-level factors. However, negative returns tend to cluster by industry or category on a detailed level, as recent events have emphasised (e.g. Californian utilities, US telecoms, South American sovereigns). Risk models with a large number of detailed factors are better able to pick up this clustering of risk and provide an accurate risk assessment.

The use of a large numbers of factors does however come at a cost: large amounts of detailed data are required, and the process can become inordinately complex. Fortunately, as a provider of fixed income indices for the past 30 years, Lehman has an extensive historical fixed-income

[2] See Dynkin and Carmel (2001).

database, enabling the implementation of models with large numbers of factors. Complexity is handled by organising factors into intuitive categories: risk due to term structure allocation, risk due to sector allocation, risk due to credit quality allocation, and so on.

Why a multi-factor model?

With the abundance of data available in today's marketplace, an asset manager might be tempted to evaluate the risk of each individual security directly using the security's own return history. The standard deviation of a security's return in the upcoming period can be projected to match its own past realised volatility; the correlation between any two securities can be set to their realised historical correlation. Despite the simplicity of this scheme, the multi-factor approach has several important advantages. First of all, the number of risk factors in the model is much smaller than the number of securities in a typical investment universe, and so the number of correlations and variances needed to calculate portfolio risk is greatly reduced. This in turn increases the speed of computation (which is becoming less important with gains in processing power) and, more importantly, improves the numerical stability of the calculations. A large covariance matrix of individual security volatilities and correlations is likely to cause numerical instability. This is especially true in the fixed income world, where returns of many securities are very highly correlated. Risk factors may also exhibit moderately high correlations with each other, but much less so than for individual securities[3].

A more fundamental problem with relying on individual security data is that not all securities can be modelled adequately in this way. For illiquid securities, pricing histories are either unavailable or unreliable; for new securities, histories do not exist. Even for liquid securities with reliable historical data, changes in security characteristics can make these data irrelevant to future results. For instance, a ratings upgrade of an issuer or simply shortening of time to maturity would make future returns less volatile than in the past. A change in interest rates can significantly alter the effective duration of a callable bond. As any bond ages, its duration shortens, making its price less sensitive to interest rates. A multi-factor model estimates the risk from owning a particular bond based not on the historical performance of that bond, but on historical returns of all bonds with characteristics similar to those currently pertaining to the bond.

Dynkin, Hyman and Wu (1999) describe multi-factor risk analysis and the Lehman model in detail. Much of the remainder of this article is a condensed version of this prior publication. The article is organised as follows. Section 2 presents basic risk model mathematics. Section 3 contains an overview of risk model applications. Section 4 looks at various reports and diagnostics that the risk model produces. Section 5 consists of an in-depth look at one particular risk model application: proxy portfolios. Section 6 describes calibration and model extensions. The article concludes with Section 7.

2 Basic risk model mathematics

This section provides the linear algebra behind the basic risk model equations. The risk model expresses security returns as a function of a set of systematic factors such as interest rate shifts at various points along the yield curve, industry and quality spread changes, etc. Let **x** be the

[3] Some practitioners insist on a set of risk factors that are uncorrelated to each other. We have found it more useful to select risk factors that are intuitively clear to investors, even at the expense of allowing positive correlations among the factors.

vector of realisations of the systematic factors for month t. The return of any bond *i* can be expressed as

$$r_i = time_i + \sum_{j=1}^{M} f_{ij} x_j + \varepsilon_i = time_i + \mathbf{f_i x} + \varepsilon_{i,} \tag{1}$$

where r_i is the bond's total return, $\mathbf{f_i}$ is a 1xM row vector of known factor loadings (sensitivities of the bond's return to a given risk factor), **x** is an Mx1 column vector of factor realisations, and $time_i$ is the deterministic time return for the next month for bond i. The random variable ε_i is the bond's non-systematic return. The non-systematic return ε_i is the portion of the bond's total return that is not explained by the factors.

Let $\boldsymbol{\varepsilon}$ be the Nx1 column vector of non-systematic returns. Let $\boldsymbol{\Omega}$ be the MxM covariance matrix of the random vector **x**, and let $\boldsymbol{\Gamma}$ be the sparse NxN covariance matrix of the random vector $\boldsymbol{\varepsilon}$. When idiosyncratic risk is modelled as issue-specific, all off-diagonal elements of $\boldsymbol{\Gamma}$ are zero. When idiosyncratic risk is modelled as issuer-specific, then element (i,j) of matrix $\boldsymbol{\Gamma}$ is zero if security i and j were issued by different firms.

The portfolio's squared tracking error (variance of portfolio return in excess of benchmark return) is

$$\sigma_{TE}^2 = (\mathbf{f_p} - \mathbf{f_b})\Omega(\mathbf{f_p} - \mathbf{f_b})^{\mathbf{T}} + (\mathbf{q_p} - \mathbf{q_b})\Gamma(\mathbf{q_p} - \mathbf{q_b})^{\mathbf{T}} \tag{2}$$

where

$\mathbf{f}_p$ is the 1xM vector of portfolio factor loadings,
$\mathbf{f}_b$ is the 1xM vector of benchmark factor loadings,
$\mathbf{q}_p$ is the 1xN vector of portfolio weights by issuer (or issue), and
$\mathbf{q}_b$ is the 1xN vector of benchmark weights.

The risk model's primary output is the portfolio's tracking error, which is the estimated standard deviation of the difference between portfolio and benchmark returns. The tracking error is calculated by simply taking the square root of σ_{TE}^2.

3 Overview of risk model applications

Estimation of portfolio return volatility and tracking error

The most fundamental application of the risk model is to estimate portfolio return volatility over the coming month under the assumption that portfolio allocations will remain unchanged in this period. When combined with an assumption about the distribution of portfolio returns (for example normality), this standard deviation can be used in monthly value-at-risk (VaR) calculations or to determine the probability that portfolio return will fall below any pre-specified level.

The risk model can be used to assess two forms of return volatility: (i) absolute return volatility, and (ii) volatility of returns in excess of the benchmark. The typical user of the Lehman Risk Model is an asset manager who is evaluated relative to a benchmark. For such a manager, risk is the standard deviation of portfolio return in excess of benchmark return or *tracking error*. The risk model assumes the presence of a benchmark and reports portfolio tracking error.

In instances where total return volatility, rather than tracking error, is the predominant concern, the user can set the benchmark to the 1-month US Treasury bill. For this benchmark, the tracking error and total return volatility are identical. It is not essential to set the benchmark to the 1-month US Treasury bill, since the risk model always reports total return volatility in addition to tracking error. However, investors interested in total return volatility may find the more detailed diagnostics more readily interpretable if the benchmark is present.

Determination of the sources of portfolio risk

In addition to providing portfolio tracking error and overall return volatility, the model provides a variety of diagnostics to determine the primary sources of portfolio risk. Such diagnostics are important for two reasons. First, blind reliance on model outputs is a dangerous practice in any situation involving financial modelling. By examining the sources of risk, investors can check whether the sources found by the model are in accord with their intuitive understanding of the portfolio's risks. If they are not, this may mean that the portfolio contains an unintended exposure. Second, these diagnostics are designed to be readily interpretable and provide insights as to how to reduce portfolio risk if it has risen to unacceptable levels.

Formation of risk-adjusted returns for performance evaluation

Suppose a manager generates a series of outstanding returns. How can it be determined whether these returns are due to picking undervalued assets and sectors, or to taking high-risk positions which happened to pay off? The sample standard deviation of the manager's realised returns could be calculated, but this approach often runs into the issue of small sample size. If the manager is lucky, a fundamentally risky strategy may generate a string of stable positive returns for several quarters. However, eventually the law of large numbers will make itself felt and a devastating negative realisation will probably occur.

The "fat-tailed" nature of most financial returns makes this problem particularly prevalent.[4] Fat-tailed returns are often a sign that volatility can change substantially through time. A fundamentally risky strategy may produce a sequence of stable returns during low volatility periods. However, once the market shifts into a period of high volatility, the same strategy may generate exceptionally volatile returns. This behaviour is typical of the standard deviation of realised returns, which is an ex post risk measure. Such measures include only the risks of events that have occurred, and ignore all risks that do not materialize.

For these reasons, when adjusting returns for risk, it is preferable to use an ex ante risk measure, which includes the effects of all adverse outcomes that could have been projected at the time of an investment decision. A risk model can be used to calculate the estimated volatility of each month's return. Then the realised return should be adjusted by its estimated volatility to measure the return achieved per unit of risk. Similarly, to judge risk-adjusted returns relative to a benchmark, one can divide each month's realised portfolio return in excess of the benchmark return by its risk model tracking error.

Risk budgeting

Consider a portfolio manager who wishes to limit portfolio risk. Without access to a risk model, the standard method for constraining risk is to calculate a variety of portfolio risk exposures (key rate durations and credit spread durations) and to make sure that each falls within an acceptable range. For example, a portfolio might be constrained to have its duration

[4] "Fat tails" refers to situations in which extreme events occur more often than would be predicted from a normal distribution.

within one year of benchmark duration, and its sector weights to be within 10% of benchmark allocations.

However, this is an inefficient mechanism for managing risk. Suppose each portfolio exposure is set to the maximum value allowed by the constraint. These exposures translate into an overall level of portfolio return risk. This level is unknown without the benefit of a risk model, but nonetheless exists. If the manager is willing to set one of the risk exposures more conservatively, and thereby reduce overall portfolio risk, some other portfolio exposure could probably be increased without exceeding allowable portfolio tracking error. A policy which constrains each portfolio risk exposure individually, however, does not allow the manager to trade off risk along one dimension for risk along another dimension.

With a risk model, the portfolio's *risk budget* can be set to a predetermined level indicating the maximum allowable tracking error. Suppose the manager perceives an unusual opportunity in the marketplace that cannot be taken advantage of without moving one of the portfolio's risk exposures outside its traditional range. The risk model will tell the manager how much to draw the portfolio's other risk exposures down in order to leave overall portfolio risk unchanged. A risk budgeting approach allows the manager to trade off risks along various dimensions, and greatly expands the manager's ability to implement profitable strategies without compromising overall portfolio risk.

Optimisation

The risk model contains an optimiser which ranks securities from a user-specified set of investable securities according to their marginal contribution to portfolio risk. If the portfolio manager is interested in purchasing a specific security from an investable pool, this security can be entered into the optimiser and matched with a security to be sold for the highest reduction in portfolio risk.

If the manager wishes to reduce the portfolio's tracking error or return volatility, the optimiser can be used to provide a list of the most effective possible trades, allowing the manager to choose the actual trades to be implemented.

Benchmarking

Investors generally wish to establish a benchmark both to evaluate managerial performance and to set the basic properties of the investment portfolio. When evaluating a potential benchmark, the risk model can add valuable insight into the total return risk that the portfolio would be subject to if it passively tracked the proposed benchmark.[5]

In addition to determining the benchmark for its overall holdings, a central bank may decide to assign a fraction of its holdings to be managed relative to a particular asset class (for example, US agencies or Aaa asset-backed securities). The risk model can be used to ensure that the return volatilities of the benchmarks for these sub-portfolios are at their desired levels.

[5] In many cases, the historical returns of a proposed benchmark may be studied directly. Nevertheless, risk model analysis can provide a more detailed breakdown of the sources of total return volatility, and can also react more quickly to changes in asset class characteristics such as MBS index duration.

Proxy portfolios

Consider a central bank that wishes to invest a portion of its assets in a certain market sector such as Aaa asset-backed securities. Clearly, the bank cannot simply buy each issue in the market in proportion to its market value. The resulting portfolio would be unwieldy and would require the bank to purchase a large number of issues that may not be available in the marketplace at reasonable bid-ask spreads. Instead, the bank can use the risk model optimiser to form a *proxy portfolio* by setting the benchmark equal to the index that mirrors the desired market segment. The user can determine the set of liquid issues available for purchase. The optimiser will recommend a series of trades to create a tracking portfolio with the fewest security holdings. Similarly, if a portfolio manager at a central bank is managing a portfolio relative to a particular benchmark, he or she can use this proxy portfolio technique to set up the core of the new portfolio.

Hedging

Suppose a manager is interested in using a set of liquid on-the-run Treasuries to hedge an exposure to a particular Treasury note over the coming month. The risk model optimiser could be used to determine the hedge by setting the benchmark as the Treasury note to be hedged and the pool of investable assets as the set of on-the-run Treasuries. The optimiser will then construct the desired hedge.

Generation of realistic scenarios

Scenario analysis is an important tool for understanding portfolio performance. Typically, the portfolio manager has a good grasp of the market parameters to be shifted in scenario analysis, but it is often less clear how to determine shift magnitudes. For example, suppose a portfolio manager is interested in determining portfolio performance in the event of an increase in credit spreads. With the risk model, the number of basis points that corresponds to a one standard deviation credit spread movement could be calculated.[6] The manager can then use these numbers to generate one, two, and three standard deviation scenarios for spread movements. Similarly, the correlations from the risk model covariance matrix can be used to generate the movements in other risk factors that are likely to coincide with a one standard deviation movement in a risk factor of interest. This is particularly popular for multi-currency portfolios.

4 Risk reports

The model produces a set of cross-sectional reports that compare the composition of the portfolio and the benchmark in as many ways as possible. These various comparative reports form the basis of the risk analysis by identifying structural differences between the two. By themselves, however, they fail to quantify the risk due to these mismatches. The model's anchor is therefore the tracking error report, which quantifies the risks associated with each cross-sectional comparison. Taken together, the various reports produced by the model provide a complete understanding of the risk of this portfolio versus its benchmark.

---

[6] Dynkin, Hyman and Wu (1999) provide a detailed exposition of this calculation.

The single most important output of the model is the overall annualised tracking error, defined as one standard deviation of the difference between the portfolio and benchmark returns. It is the random part of the difference between portfolio and benchmark returns. In simple terms, this means that the portfolio return over the next year will be within +/- the tracking error of the benchmark return with a probability of about 66%.[7] The model also reports the *beta* of the portfolio. Portfolio beta provides the sensitivity of portfolio return to shifts in benchmark returns. It is common to compare the beta produced by the risk model with the ratio of portfolio and benchmark durations. The duration ratio is often somewhat different from the risk model beta because the duration-based approach only considers term structure risk (and purely parallel shift risk at that), while the risk model includes the combined effects of all relevant forms of risk, along with the correlations among them.

Sources of systematic tracking error

What are the main sources of this tracking error? The model identifies market forces influencing *all* securities in a certain category as *systematic risk factors*. The tracking error is decomposed into components corresponding to different categories of risk[8] that correspond directly to the groups of risk factors. A detailed report of the differences in portfolio and benchmark exposures (sensitivities) to the relevant set of risk factors illustrates the origin of each component of systematic risk.

Sensitivities to risk factors are called *factor loadings*. They are expressed in units that depend on the definition of each particular risk factor. The factor loadings of a portfolio or an index are calculated as a market value-weighted average over all constituent securities. Differences between portfolio and benchmark factor loadings form a vector of *active portfolio exposures*. A comparison of the magnitudes of the different components of tracking error highlights the most significant mismatches.

Usually the largest component of portfolio risk is due to term structure. Risk factors associated with term structure movements are represented by a fixed set of points on the theoretical Treasury spot curve. Each of these risk factors exhibits a certain historical return volatility. The extent to which the portfolio and the benchmark returns are affected by this volatility is measured by factor loadings (exposures). These exposures are computed as percentages of the total present value of the portfolio and benchmark cash flows allocated to each point on the curve.[9] The risk that the portfolio will perform differently from the benchmark because of the term structure movements can be attributed to the differences in the portfolio and benchmark exposures to these risk factors and to their volatilities and correlations.

The tracking error is calculated from the vector of differences between portfolio and benchmark exposures. However, mismatches at different points are not treated equally. Exposures to factors with higher volatilities have a larger effect on tracking error. It should be noted that the risk caused by overweighting one segment of the yield curve can sometimes be

---

[7] This interpretation requires several simplifying assumptions. The 66% confidence interval assumes that returns are normally distributed. Furthermore, this presentation ignores differences in the expected returns of portfolio and benchmark (owing, for example, to a higher portfolio yield). Strictly speaking, the confidence interval should be drawn around the expected outperformance.

[8] The model also computes total return volatility for the portfolio and the benchmark. These items are labelled portfolio sigma and benchmark sigma, respectively.

[9] The details behind this calculation may be found in Appendix 2 of Dynkin, Hyman and Wu (1999). The new Lehman Brothers Risk model released in 2002 redefines term structure exposure in terms of option-adjusted key rate durations, rather than the allocation of static cash flows described here. See Section 6 and Naldi, Chu and Wang (2002).

offset by underweighting another, because adjacent points on the curve are highly correlated and almost always move together. To completely eliminate the tracking error due to term structure, exposures to each term structure risk factor need to be reduced to zero. To reduce term structure risk, it is most important to focus first on decreasing exposures at the long end of the curve, particularly those that are not offset by opposing positions at nearby points.

The tracking error due to sector exposures is explained by showing the sector allocations of the portfolio and the benchmark in two ways. In addition to reporting the percentage of market value allocated to each sector, it shows the contribution of each sector to the overall spread duration.[10] These contributions are computed as the product of the percentage allocations to a sector and the market-weighted average spread duration of the holdings in that sector. Contributions to spread duration (factor loadings) measure the sensitivity of return to systematic changes in particular sector spreads (risk factors), and are a better measure of risk than simple market allocations.

The analysis of credit quality risk follows the same approach. Portfolio and benchmark allocations to different credit rating levels are compared in terms of contributions to spread duration. The risk represented by tracking error due to quality corresponds to a systematic widening or tightening of spreads for a particular credit rating, uniformly across all industry groups.

The largest sources of systematic risk are often term structure, sector, and quality. The model provides similar reports for exposures to optionality factors (factors assessing callability and putability) and several mortgage factors. Dynkin, Hyman and Wu (1999) provide a detailed account of the risk factors and reports for these risk components, but they have only limited relevance for portfolios with minimal exposure to mortgages and callable bonds.

Sources of non-systematic tracking error

In addition to the various sources of systematic risk, portfolios have non-systematic tracking error, or special risk. This risk stems from portfolio concentrations in individual securities or issuers. As the portfolio is usually relatively small compared to the benchmark, each bond comprises a noticeable fraction. A negative credit event associated with a portfolio holding (i.e. a downgrade) will cause large losses in the portfolio, while hardly impacting the highly diversified benchmark. The Lehman Aggregate Index consisted of more than 6,700 securities as of December 2003, so that the largest US Treasury issue accounts for less than 1%, and most corporate issues contribute less than 0.01% of the index market value. Thus, any large position in a corporate issue represents a material difference between portfolio and benchmark exposures that must be considered in a full treatment of risk.

The magnitude of the return variance that the risk model associates with a mismatch in allocations to a particular issue is proportional to the square of the allocation difference and to the residual return variance estimated for the issue. This calculation is shown in schematic form in Figure 4. With the return variance based on the square of the market weight, it is dominated by the largest positions in the portfolio. The estimate of the monthly volatility of non-systematic return is based on bond characteristics such as sector, quality, duration, age, and amount outstanding.[11]

---

[10] Just as traditional duration can be defined as the sensitivity of bond price to a change in yield, spread duration is defined as the sensitivity of bond price to a change in spread. While this distinction is largely academic for bullet bonds, it can be significant for other securities, such as bonds with embedded options and floating-rate securities. The sensitivity to spread change is the correct measure of sector risk.

[11] See Appendix 3 of Dynkin, Hyman and Wu (1999) for details.

Figure 4: Calculation of variance due to special risk (issue-specific model)

	Portfolio Weights	Benchmark Weights	Contribution to Issue-specific Risk
Issue 1	w_{P_1}	w_{B_1}	$(w_{P_1} - w_{B_1})^2 \sigma_{\varepsilon_1}^2$
Issue 2	w_{P_2}	w_{B_2}	$(w_{P_2} - w_{B_2})^2 \sigma_{\varepsilon_2}^2$
…			
Issue *N-1*	$w_{P_{N-1}}$	$w_{B_{N-1}}$	$(w_{P_{N-1}} - w_{B_{N-1}})^2 \sigma_{\varepsilon_{N-1}}^2$
Issue *N*	w_{P_N}	w_{B_N}	$(w_{P_N} - w_{B_N})^2 \sigma_{\varepsilon_N}^2$
Total Issue-specific Risk			$\sum_{i=1}^{N} (w_{P_i} - w_{B_i})^2 \sigma_{\varepsilon_i}^2$

For bonds of similar maturities, the model tends to assign higher special risk volatilities to lower-rated issues. Thus, mismatches in low-quality bonds with long duration will be the largest contributors to non-systematic tracking error.

It is important to note that there are two separate decisions to be made in modelling non-systematic risk: (i) the correlation structure of non-systematic returns, and (ii) the estimation of the variance of the non-systematic component of returns. In the model of non-systematic return correlations, the non-systematic returns of securities issued by different firms are always assumed to be independent. However, in the estimation of the variance of non-systematic risk, it is essential to cluster securities into multiple-firm categories and to use information statistically derived from the entire category to estimate the variances of the non-systematic returns of each security in the category. This is discussed in more detail in Section 5.

This non-systematic risk calculation is carried out twice, using two different methods. In the *issuer-specific* calculation, the holdings of the portfolio and benchmark are not compared on a bond-by-bond basis, but are first aggregated into concentrations in individual issuers. This calculation is based on the assumption that spreads of bonds of the same issuer tend to move together. Therefore, matching the benchmark issuer allocations is sufficient. In the *issue-specific* calculation, each bond is considered an independent source of risk. This model recognises that large exposures to a single bond can incur more risk than a portfolio of all of an issuer's debt. In addition to credit events that affect an issuer as a whole, individual issues can be subject to various technical effects. For most portfolios, these two calculations produce very similar results. In certain circumstances, however, there can be significant differences. For instance, some large issuers use an index of all their outstanding debt as an internal performance benchmark. In the case of a single-issuer portfolio and benchmark, the issue-specific risk calculation will provide a much better measure of non-systematic risk. The reported non-systematic tracking error is the average of the results from the issuer-specific and issue-specific calculations.

Figure 5: Illustration of “Isolated” and “Cumulative” Calculations of Tracking Error Sub-components

a) Isolated Calculation of Tracking Error Components

b) Cumulative Calculation of Tracking Error Components

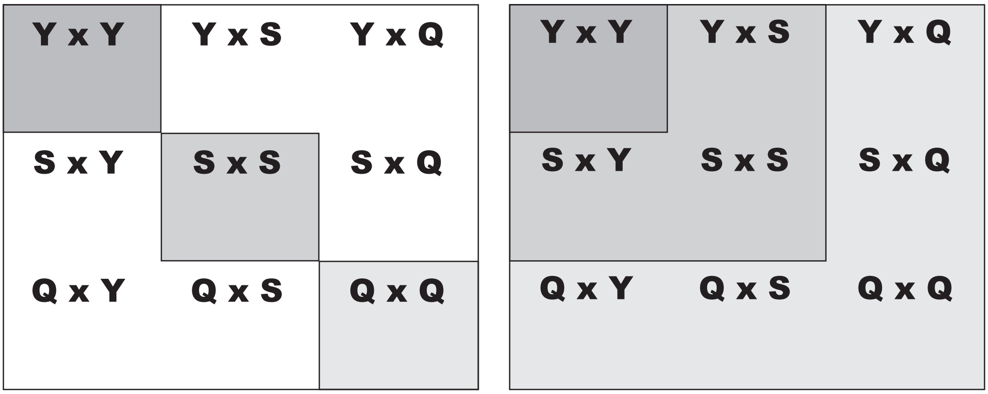

Combining components of tracking error

Given the origins of each component of tracking error, we can address the question of how these components combine to form the overall tracking error. The net result of the systematic and non-systematic sources of tracking error does not equal their sum. Rather, the squares of these two numbers (which represent variances) sum to the variance of the result. Next, we take the square root to obtain the overall tracking error.

The sub-components of tracking error due to sector, quality, etc. can be calculated in two different ways. We can estimate the *isolated tracking error* due to the effect of each group of related risk factors considered alone. The tracking error due to term structure, for example, reflects only the portfolio/benchmark mismatches in exposures along the yield curve, as well as the volatilities of each of these risk factors and the correlations among them. Similarly, the tracking error due to sector reflects only the mismatches in sector exposures, the volatilities of these risk factors, and the correlations among them. However, the correlations between the risk factors due to term structure and those due to sector do not participate in either of these calculations. Figure 5 depicts an idealised covariance matrix containing just three groups of risk factors relating to the yield curve (Y), sector spreads (S) and quality spreads (Q). Figure 5a illustrates how the covariance matrix is used to calculate the sub-components of tracking error in the isolated mode. The three shaded blocks represent the parts of the matrix that pertain to: movements of the various points along the yield curve and the correlations among them (Y × Y); movements of sector spreads and the correlations among them (S × S); and movements of quality spreads and the correlations among them (Q × Q). The portions of the matrix that are not shaded deal with the correlations among different sets of risk factors, and do not contribute to any of the isolated tracking errors.

A different way of subdividing tracking error computes the *cumulative tracking error*, which incrementally introduces one group of risk factors at a time to the tracking error calculation. As additional groups of risk factors are included, the calculation converges towards the total systematic tracking error, which is obtained by using the entire matrix. Figure 5b illustrates the rectangular section of the covariance matrix that is used at each stage of the calculation. The incremental tracking error due to sector reflects not only the effect of the $S \times S$ box in the diagram, but the $S \times Y$ and $Y \times S$ cross-terms as well. The partial tracking error due to sector therefore takes into account the correlations between sector risk and yield curve risk, and answers the question "Given the exposure to yield curve risk, how much more risk is introduced by the exposure to sector risk?"

The incremental approach has the intuitively pleasing property that the partial tracking errors add up to the total systematic tracking error. Of course, the order in which the various partial tracking errors are considered will affect the magnitude of the corresponding terms. Additionally, some of the partial tracking errors computed in this way may be negative. This reflects negative correlations among certain groups of risk factors.

The two methods used to subdivide tracking error into different components are complementary and serve different purposes. The isolated calculation is ideal for comparing the magnitudes of different types of risk in order to highlight the most significant exposures. The cumulative approach produces a set of tracking error sub-components that sum to the total systematic tracking error and reflect the effect of correlations among different groups of risk factors. The major drawback of the cumulative approach is that results are highly dependent on the order in which they are computed. The order currently used by the model was selected based on the significance of each type of risk; this may not be optimal for every portfolio/benchmark combination.

5 A detailed look at one application: proxy portfolios

Dynkin, Hyman and Wu (1999) consider a large number of detailed presentations of risk model applications. Owing to space constraints, we only consider one application in depth: proxy portfolios.

Investors often look for "index proxies" – portfolios with a small number of securities that nonetheless closely match their target indices. Proxies are used for two distinct purposes: for direct passive investment and for index analysis. Passive investment might be interesting for both passive portfolio managers and active managers with no particular view on the market at a given time. Proxy portfolios represent a practical method of matching index returns while containing transaction costs. In addition, the large number of securities in an index can pose difficulties for the application of computationally-intensive quantitative techniques. A portfolio can be analysed against an index proxy of a few securities using methods that would be impractical to apply to an index of several thousand securities. As long as the proxy is constructed to match the index along relevant risk dimensions, this approach can speed up many forms of analysis with just a small sacrifice in accuracy.

There are several approaches to the creation of index proxies. Quantitative techniques include stratified sampling or cell-matching, tracking error minimisation, and matching index scenario results. The replication of index returns can also be achieved using securities outside of indices, such as Treasury futures contracts[12] (an alternative way of obtaining index returns is entering into an index swap or buying an appropriately structured note). Regardless of the means used to build a proxy portfolio, the Lehman Brothers Risk Model can measure how well the proxy is likely to track the index.

In a simple cell-matching technique, a benchmark is profiled on an arbitrary grid, which reflects the risk dimensions along which a portfolio manager's allocation decisions are made. The index allocations to each cell are then matched by one or more representative liquid

[12] See Dynkin, Hyman and Lindner (2002).

Figure 6: Corporate Proxy – Tracking Error as a Function of Number of Bonds (the Effect of Diversification)

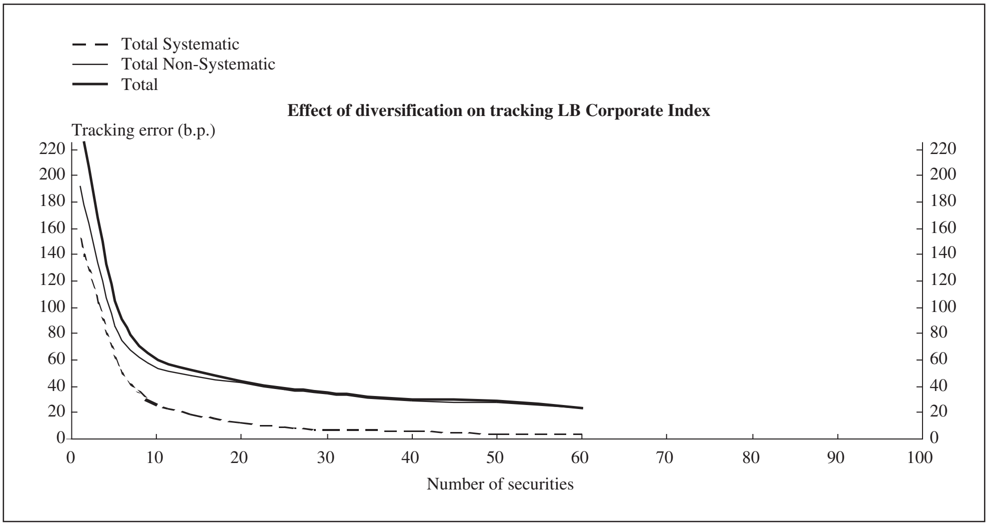

securities. The duration (and convexity) of each cell within the benchmark can be targeted when purchasing securities to fill the cell.

To create or fine-tune a proxy portfolio using the risk model, one starts by selecting a seed portfolio and an investable universe. Then the risk model optimiser recommends a sequence of transactions. As more bonds are added to the portfolio, risk decreases. The level of tracking achieved by a proxy portfolio depends on the number of bonds it contains. Figure 6 shows the annualised tracking errors achieved using this procedure, as a function of the number of bonds, in a proxy for the Lehman Brothers Corporate Bond Index. At first, adding more securities to the portfolio reduces tracking error quickly, but as the number of bonds grows, the process levels off. The difference in nature between systematic and non-systematic risk explains this phenomenon. As securities are added to the portfolio, systematic risk is quickly reduced. Once the corporate portfolio is sufficiently diverse to match index exposures to all industry groups and quality levels, non-systematic risk starts to dominate, and tracking error decreases at a slower rate.

Dynkin, Konstantinovsky and Phelps (2001) use the Lehman Risk Model to form proxy portfolios that track the Lehman Mortgage-Backed Securities Index. They study two investable universes. The first uses only current coupon TBAs (to-be-announced securities bought for forward settlement with only broad characteristics specified, not a specific issue). The advantage of this strategy is that each month the investor rolls out of the old TBAs into the new TBAs. As the strategy never entails taking delivery on actual mortgage pools, it enables funds to track the mortgage index even in situations in which their back offices are not set up to handle the complexities of mortgage payments and settlements. The second strategy does involve holding actual mortgages, but restricts the proxy portfolio to purchasing only large, newly-issued pools.[13] This second strategy necessitates a back office that can

[13] Each pool in this study is selected from a Lehman MBS generic (a group of similar pools with the same program, coupon and year of issuance) with at least $5 billion outstanding.

Figure 7: Realised Tracking Error (Annualised)
versus the Lehman MBS Index (bps)
Two-Year Overlapping Intervals, January 1994-May 2001

	TBAs Only		Large Pools Only	
Period	Tracking Error	Av. # of Holdings	Tracking Error	Av. # of Holdings
Jan. 94-Dec. 95	28.0	8	25.1	17
Jan. 95-Dec. 96	20.2	11	14.7	25
Jan. 96-Dec. 97	17.4	12	12.5	33
Jan. 97-Dec. 98	16.8	14	9.5	43
Jan. 98-Dec. 99	12.2	14	7.3	55
Jan. 99-May 01	18.3	11	5.1	69

Source: Dynkin, Konstantinovsky and Phelps (2001).

process actual mortgages, but only requires purchase of the most liquid pools during the most liquid part of their lifecycle (new issue). By replicating the MBS index using these proxy portfolio strategies, the investor can be confident that the portfolio will receive a return close to the MBS index without having to develop expertise in picking undervalued mortgage pools or investing resources in complicated prepayment models.

Figure 7 presents the realised tracking errors for these replication strategies. The tracking error for the "Large Pool Only" strategy steadily shrinks to 5.1 bp per year by the fifth year of the study (January 1999). The improvements in performance over time reflect a seasoning effect. The replication strategy starts in January 1994. It should be remembered that the large-pool strategy is restricted to purchasing only newly-issued pools. However, the actual index contains many pools that were issued several years earlier. Thus for the first few years, the large-pool proxy portfolio is unable to match the seasoning of the actual index. As time passes and the previously-purchased mortgages in the large pool strategy start to season, it is able to track the MBS index more closely.

By contrast, the TBA-only strategy can never hold seasoned pools. Therefore, there is no reason to expect its ability to track the index to improve over time. This is shown in Figure 7: replication results show no discernable trend from 1996 to 2001. The high level of tracking errors in the first two years of the study can be attributed to a lack of sufficiently long data history early in the study. The return history used by the mortgage component of the risk model starts in autumn 1993. In 1994 and 1995, the risk model had to compute mortgage risk based on only a few months of data.

A proxy portfolio for the Lehman Brothers Aggregate Index can be constructed by building proxies to track each of its major components, and combining them with the proper weights. This exercise clearly illustrates the benefits of diversification. The aggregate proxy in Figure 8 is obtained by combining the government, corporate and mortgage proxies shown in

Figure 8: Effect of Diversification – Tracking Error vs. Treasuries, Corporates, MBS, Aggregates

Index	No. of bonds in proxy	No. of bonds in index	Tracking Error (bp/year)
Treasury	6	165	13
Government	39	1,843	11
Corporate	51	4,380	26
Mortgage	19	606	15
Aggregate	115	6,928	10

the same figure. The tracking error achieved by the combination is smaller than that of any of its constituents. This is because the risks of the proxy portfolios are largely independent.

When using tracking error minimisation to design proxy portfolios, the choice of the "seed" portfolio and the investable universe should be considered carefully. The seed portfolio is the initial portfolio presented to the optimiser. Owing to the nature of the gradient search procedure, the path followed by the optimiser will depend on the initial portfolio. The closer the seed portfolio is to the benchmark, the better the achieved results will be. At the very least, it is advisable to choose a seed portfolio with duration near that of the benchmark. The investable universe, or bond swap pool, should be wide enough to offer the optimiser the freedom to match all risk factors. However, if the intention is actually to purchase the proxy, the investable universe should be limited to liquid securities.

It should be noted that the described methods for building proxy portfolios are not mutually exclusive, but can be used in conjunction with one another. A portfolio manager who seeks to build an investment portfolio that is largely passive to the index can use a combination of security picking, cell matching, and tracking error minimisation. By dividing the market into cells and choosing one or more preferred securities in each cell, the manager can create an investable universe of candidate bonds in which all sectors and credit qualities are represented. The tracking error minimisation procedure can then match index exposures to all risk factors, while choosing only securities that the manager would like to purchase.

6 Calibration and extensions

As per Equation (2), the square of the portfolio tracking error can be expressed as

$$\sigma_{TE}^2 = (\mathbf{f_p} - \mathbf{f_b})\Omega(\mathbf{f_p} - \mathbf{f_b})^{\mathbf{T}} + (\mathbf{q_p} - \mathbf{q_b})\Gamma(\mathbf{q_p} - \mathbf{q_b})^{\mathbf{T}}$$

where

$\mathbf{f}_p$ is the 1xM vector of portfolio factor loadings,
$\mathbf{f}_b$ is the 1xM vector of benchmark factor loadings,
$\mathbf{q}_p$ is the 1xN vector of portfolio weights by issuer (or issue),
$\mathbf{q}_b$ is the 1xN vector of benchmark weights,
$\mathbf{\Omega}$ is the MxM covariance matrix of the systematic risk factors, and
$\mathbf{\Gamma}$ is the NxN covariance matrix of idiosyncratic returns.

By design, the vector of factor loadings ($\mathbf{f}_p$ and $\mathbf{f}_b$) are observable characteristics that can be calculated using Lehman analytics (durations, spread durations, convexities, etc). Similarly, portfolio and benchmark weights ($\mathbf{q}_p$ and $\mathbf{q}_b$) are directly observable. The final inputs required to implement the risk model are estimates of the covariance matrices $\mathbf{\Omega}$ and $\mathbf{\Gamma}$. $\mathbf{\Omega}$ is the covariance matrix of $\mathbf{x}$, the vector of factor realisations. However, the vector of factor realisations ($\mathbf{x}$ in equation (1)) is not observable, even ex post. This complicates the estimation of the covariance matrix $\mathbf{\Omega}$.

To understand the procedure used to estimate $\mathbf{\Omega}$, it is helpful to recall from equation (1) that security returns are modelled as follows:

$$r_i = time_i + \mathbf{f_i x} + \varepsilon_{i,}$$

where

r_i is the bond's total return,
$time_i$ is the deterministic time return for the next month for bond i,
$\mathbf{f}_i$ is a 1xM row vector of known factor loadings for bond i,

x is an Mx1 column vector of factor realisations, and

ε_i is the portion of the bond's total return that is independent of model factors.

Omitting certain adjustments, the general approach used to estimate $\mathbf{\Omega}$ is as follows.[14] Each month, a series of cross-sectional regressions is run with realised bond returns r_i as the dependent variables and bond factor loadings f_{ij} as the independent variables. The estimated regression coefficient for the jth column of factor loadings in month t is the implied factor realisation x_{jt} for factor j in month t. This process is repeated each month and generates a time series of imputed factor realisations for each factor. The time-t factor realisations are then reassembled into a column vector $\mathbf{x} = [x_{1t},\ldots,x_{Mt}]^T$. The time series sample covariance matrix of the imputed realisations for **x** is used as an estimate of $\mathbf{\Omega}$.

In the monthly cross-sectional regressions, only the most relevant securities are used to estimate factor realisations for each set of factors. For instance, the cross-sectional regression used to estimate term structure factor realisations is restricted to Treasury notes and bonds. The dependent variable in this regression is Treasury returns in excess of time return, and the independent variables are the securities' term structure factor loadings. Similarly, the cross-sectional regression used to estimate credit factor realisations is restricted to corporate bonds. In this regression the dependent variable is the part of realised return in excess of time return which has not been explained by the term structure risk factors.[15] The independent variables are spread durations multiplied by dummy variables indicating the security's credit rating and sector.

Let ε_{it} be the residual return for bond i in month t. The time series variance of these residuals is calculated for each security. Securities are grouped into categories and the time series variances are then averaged inside each group. The diagonal element (i,i) of matrix $\mathbf{\Gamma}$ is set equal to the average variance for the group that security i belongs to.[16] Grouping of this form is essential to be able to estimate security-specific risk for new issues.

For calculations in which idiosyncratic risk is issue-specific, all off-diagonal elements of $\mathbf{\Gamma}$ are set to zero. When idiosyncratic risk is modelled as issuer-specific, element (i,j) of matrix $\mathbf{\Gamma}$ is set to zero if securities i and j were issued by different firms. Otherwise, the diagonal elements of $\mathbf{\Gamma}$ for securities i and j are calculated based on an assumption of perfect correlation.

Observable factors or observable factor loadings?

Most risk models determine factors first and then statistically estimate each security's loading based on these factors. In these models factors are established in one of two ways: (i) observable factors in which a set of observable time series are exogenously specified as risk factors (e.g. 2-year yield, Aa credit spread, etc); or (ii) statistically-estimated factors in which the factors are estimated from principal components decomposition of the covariance matrix of security returns. In either case, each security's return sensitivity (factor loading) to each factor is then estimated through regression.

---

[14] See Dynkin, Hyman and Wu (1999) and Naldi, Chu and Wang (2002) for more details.

[15] The value used on the left-hand side of the regression for security i is $r_i - time_i + \mathbf{f}_i\hat{\mathbf{x}}$ where the elements of $\hat{\mathbf{x}}$ are set equal to their estimated values for term structure factors. The elements corresponding to all other factors are set to zero.

[16] The actual procedure is somewhat more complicated. Inside each sector, a cross-sectional regression is estimated in which a set of security characteristics is used to explain the realised time series variances of non-systematic returns. The estimated regression parameters are then applied to an individual security's characteristics to estimate the security's non-systematic return variance. See Dynkin, Hyman and Wu (1999) for more details.

Lehman's risk model takes the opposite approach. Factor loadings are specified from observable (calculable) security characteristics. The factors and the factor covariance matrix are statistically estimated based on the exogenously-specified definitions of factor loadings. We feel that Lehman's approach has four advantages over more traditional models:

1. User-friendliness. Since factor loadings are defined exogenously, they can be set to correspond tightly to the parameters managers use to make portfolio allocations (durations, spread durations, etc).
2. Number of parameters to be estimated. Traditional models that specify factors first and then estimate factor loadings need to estimate statistically a tremendously large number of parameters. M denotes the number of factors in the model; assuming each security loads on each factor, the total number of parameters to be estimated is N times M.
   By contrast, Lehman's model, which exogenously specifies factor loadings and then imputes factor realisations from cross-sectional regressions, has N observations in each period to estimate M factor realisations. This approach only requires the estimation of M random variables in each period based on N observations.[17]
3. Ability to handle new issues and illiquid securities. Traditional models which statistically estimate each factor loading for each security inevitably face difficulties estimating factor loadings for new issues and for illiquid issues that trade infrequently and at large bid-ask spreads.
4. Robustness to changes in security characteristics. Traditional models which estimate factor loadings with time series regression encounter problems when a security's duration, spread duration, credit quality or other characteristics change. In this case, many or all of the observations in the time series regression used to estimate the security's factor loadings will be stale and will thus not reflect the security's current risk profile. When factor loadings are based on current observable (or calculable) security characteristics, they are never corrupted by stale data.

Modifying the covariance matrix

The risk model allows investors to specify options regarding the estimation of the covariance matrix. The new risk model allows the user to specify whether past data should be equally weighted or whether recent observations should be upweighted and distant observations downweighted. Specifically, the new risk model allows the user to impose weights on past data that decrease exponentially with a one-year half-life.[18]

Both approaches have merit. There are noticeable periods of low volatility and high volatility in fixed income markets. Similarly, various segments of the market undergo periods of high or low volatility. Another example is utilities which, prior to the California energy crisis, had displayed relatively low realised volatility. This argues for upweighting recent observations.

---

[17] Both approaches also estimate the M(M+1)/2 parameters of the factor covariance matrix. Since this is identical for both approaches, it is not figured in the above calculations. When the covariance matrix is based on estimated factor realisations rather than observed factor realisations, a correction should be implemented to adjust for estimation error. However, the traditional approach in which factor loadings are estimated is subject to an analogous problem. Equation (2) assumes that the factor loading vectors ($\mathbf{f}_p$ and $\mathbf{f}_b$) are observable. However, in the traditional approach, they are statistically estimated and contain estimation error. Traditional models cannot use equation (2) to estimate tracking error, but instead should modify it to adjust for estimation error in $\mathbf{f}_p$ and $\mathbf{f}_b$. As discussed above, the estimated values of $\mathbf{f}_p$ and $\mathbf{f}_b$ in traditional models are likely to contain greater amounts of estimation error than the imputed factor realisations in the Lehman model. We feel this is another advantage for the observable factor loading method.

[18] A one-year-old observation receives half as much weight as the most recent observation.

At times, financials can go from being a relatively quiet part of the fixed income market to exhibiting dramatically increased volatility, as in autumn 1998.

On the other hand, if exponentially decaying weights had used been used in late 1999, the previous period of major recession concerns (1990-1991) would have received close to zero weight in forming the risk model covariance matrix, giving the investor little indication of true portfolio risk as recession concerns reappeared in mid-2000. During the calm before a storm, upweighting of recent events can lead to excessively rosy estimates of risk exposure. Cautious investors may choose to use both approaches and emphasise the most conservative result.

Extensions

Up to this point, this paper has described the version of the Lehman Risk Model that was formally introduced in 1999.[19] In 2002, however, Lehman released a revised version of its risk model, which shares many of the features of the original model, but offers the following noteworthy innovations:

- Term Structure Factor Loadings: Option-adjusted key-rate durations replace expected cash flow arrival times as factor loadings for term structure exposure. Interest rate derivatives are included in the security set.
- Spread Exposures:
  - All static and zero-volatility spreads have been replaced by OAS-consistent measures of spread.
  - The cell-bucketing scheme for industry and quality spreads has been reconfigured.
  - Factor loadings involving optionality and mortgage prepayment have also been reconfigured.
  - Spreads are no longer assumed to shift in a parallel fashion.
  - Liquidity factors have been introduced.

A fuller exposition is provided in Naldi, Chu and Wang (2002). Most recently, Lehman has extended the model to cover all securities in the Lehman Global Aggregate Index. In addition to modelling exposures to term structure, swap spreads, and sector spreads in each currency, the model now measures currency risk as well. A detailed account of this Global Risk Model is in preparation.

7 Conclusion

Over the past decade, we have had the opportunity to work with several central banks and have found that they particularly value the following risk model features:

- **Easily Interpretable Factors.** Rather than generating factors based on principal component decompositions or other statistical methods, factors are designed to have readily understood financial meanings, aiding in interpretation and providing guidance for risk-reduction strategies.

  For example, consider an investor who is purchasing a bond outside the set of securities handled by a risk management system. If portfolio risks are presented in terms of factors with intuitive financial meanings, the manager can come to a reasonable assessment of the effects the new security will have on portfolio risk.

---

[19] Precursors to this model were in use at Lehman for many years prior to 1999.

On the other hand, suppose the risk is presented in terms of principal components decompositions and reports that the portfolio has an excessive exposure to the third principal component. This provides the manager with little guidance when considering the purchase of a security not handled by the risk management system. Does this new security increase or decrease portfolio exposure to the third principal component? Without input from the risk management system, a portfolio manager would be at a loss to understand the new security's potential impact on portfolio risk.

- **Extensive Diagnostics.** The model provides investors with a variety of reports detailing the portfolio's various risk exposures and itemising the amount of risk derived from each source.
- **Precision.** The model is based on a considerable amount of long-dated security-level data, allowing it to introduce a large number of systematic factors. This greatly enhances model accuracy. Even the largest portfolios contain many fewer issues than their benchmark indices. Thus, for large well-diversified portfolios, non-systematic risk invariably constitutes a substantial fraction of overall tracking error. The model captures security-specific, idiosyncratic risk as well as systematic risk, enabling more accurate tracking error assessment.
- **Ability to Respond to Changing Portfolio and Benchmark Characteristics.** The risk of individual securities is based on their current characteristics: durations, spreads, credit ratings, etc, not the security's past history. Thus if a security is downgraded or experiences a significant spread widening, the security's contribution to portfolio risk is based on the security's current, rather than historical, properties.

Central banking applications include estimating portfolio return volatility and tracking error, diagnosing the sources of portfolio risk, risk-adjusting returns for performance evaluation, and risk budgeting (imposing an overall return volatility or tracking error constraint on a portfolio, rather than constraining a large number of individual portfolio risk exposures).

The risk model allows investors to consider the magnitude of a spread change in terms of a number of standard deviations, and to translate this into an exact number of basis points that can be input into scenario analysis systems. For example, the covariance matrix from the risk model can be used to complete scenario definitions in a way consistent with history and with a certain set of portfolio manager's views.

One particularly useful feature of the risk model is the optimiser, which can quickly generate proxy portfolios containing a small number of assets that efficiently track a large index or market segment. It can moreover also be used for hedging. The optimiser additionally provides the user with recommendations for reducing overall portfolio risk if this has reached unacceptable levels.

References

Berd, A. and M. Naldi, 2002, *New Estimation Options for the Lehman Brothers Risk Model: Adjusting to Recent Credit Events,* Quantitative Credit Research Quarterly, Vol. 2002-Q3, Lehman Brothers Fixed Income Research publication.

Carmel, J. and L. Dynkin, 2001, *The Prospect of Negative Annual Returns for Short-duration Treasury Benchmarks,* Global Relative Value, 22 October 2001, Lehman Brothers Publication, Fixed Income Research publication.

Carmel, J., L. Dynkin, J. Hyman and B. Phelps, 2002, *Total Return Management of Central Bank Reserves,* Lehman Brothers Publication, Fixed Income Research publication.

Dynkin L., V. Konstantinovsky and B. Phelps, 2001, *Tradable Proxy Portfolios for the Lehman Brothers MBS Index,* Lehman Brothers Publication, Fixed Income Research publication.

Dynkin L., J. Hyman and P. Lindner, 2002, *Hedging and Replication of Fixed-income Portfolios,* Journal of Fixed Income, March 2002, pp. 43-63.

Dynkin L., J. Hyman and W. Wu, 1999, *The Lehman Brothers Multi-factor Risk Model,* Lehman Brothers Publication, Fixed Income Research publication.

Naldi, M., K. Chu and G. Wang, 2002, *The New Lehman Brothers Credit Risk Model,* Quantitative Credit Research Quarterly, Vol. 2002-Q2, Lehman Brothers Fixed Income Research publication.

12

Managing market risks: a balance sheet approach

Bert Boertje and Han van der Hoorn[1]

Abstract

A central bank's balance sheet largely reflects its responsibilities as a central bank. Monetary policy considerations have historically determined the composition of the balance sheet. We present a model which applies classical mean-variance optimisation techniques to a central bank's balance sheet, given constraints typical for central banks. It is based on concepts such as profits and potential losses. The model allows for an integrated approach towards market risks. We argue that the model helps to ensure financial independence, which is a necessary condition for central bank independence.

1 Introduction

Central banking is often seen as a mysterious or even mystical affair. Recently, a number of publications have appeared which deal with slightly less mysterious aspects of central banking such as asset and risk management, which are more common to the broader financial industry (IMF, 2001 and 2003, and Frowen, et al., 2000). We dwell on these aspects and will focus on concepts such as profits and potential losses, in order to present a framework for managing market risk.

For this purpose, the major responsibilities of central banks and their impact on the balance sheet will be discussed in Section 2. Section 3 then argues that, broadly speaking, the balance sheet composition determines both the profitability and the risk profile of a central bank. Furthermore, it is suggested that the necessity for central banks to be financially independent can be used as an ingredient to manage market risks. Section 4 discusses changes in the balance sheet composition and their impact on the relevant parameters. The paper concludes with a summary of the major advantages and limitations of the approach presented.

2 Central Bank responsibilities

The aim of this section is to introduce central bank responsibilities and how they relate to a typical central bank balance sheet. The balance sheet items play a central role in the optimisation model, which will be presented in sub-section 3.3 of this paper.

Central banks are entrusted with the issuance of banknotes, which are distributed to the public through the banking system. Consequently, changes in the amount of banknotes in circulation (a liability item on the balance sheet) imply an equal change in the demand for central bank credit facilities by the banks (an asset). Moreover, the central bank's credit facilities are used to implement monetary policy decisions by setting interest rates and/or controlling the volumes involved, and are therefore referred to as monetary policy operations. Implementing monetary policy is a key responsibility of central banks. Cash reserves held by commercial banks (a liability) are often among the instruments used to implement monetary policy. They aim at

[1] De Nederlandsche Bank NV, Financial Markets Department. This article has been written on a personal basis and the views expressed do not necessarily reflect those of De Nederlandsche Bank. Comments by Jos Heuvelman and Tom van Veen are gratefully appreciated.

creating a sufficiently large money market shortage, which is necessary to facilitate monetary policy operations, and/or to restrain the creation of money in a broad sense. By investing the means generated by issuing banknotes and by enforcing cash reserves in interest-bearing assets, central banks earn what is called 'monetary income' or 'seigniorage'.

Holding and managing gold and foreign assets reflects the international dimension of central bank responsibilities. Gold can be seen as the ultimate war chest: in periods of crisis, gold has proven to be an important payment device and is therefore widely considered to bolster a country's creditworthiness. Originally, central banks often controlled international capital inflows and outflows (which were almost fully related to trade flows). Foreign assets held by central banks functioned as a buffer in cases where imports exceeded exports. Since the collapse of Bretton Woods and the liberalisation of the international capital markets, however, the exchange rate regime under which the domestic currency operates largely determines the need for foreign asset holdings.

Under a fixed rate regime, the external value of the domestic currency is the primary target within the framework for monetary policy. A fixed exchange rate regime obliges a central bank to accommodate imbalances in international capital flows. Apart from a buffer function with respect to trade as well as international capital flows, foreign exchange holdings fulfil an important role in supporting confidence in the domestic currency in such a regime. An extreme case of a fixed rate regime is a currency board. Under a currency board, a central bank can only issue banknotes against the targeted foreign currency. Banknote issuance is therefore fully covered by foreign assets. In a floating rate regime, the external value of a currency is not a goal in itself; here, monetary policy focuses on the internal value of the currency. Nevertheless, foreign assets are needed to curb an unwanted currency depreciation and/or to dampen volatile currency movements. Europe's Exchange Rate Mechanism can be seen as a mixture of a floating and a fixed exchange rate regime.

Interventions to prevent a devaluation or depreciation of the domestic currency result in a lower amount of foreign assets and – depending on how the monetary instruments are designed – an increase in the demand for credit facilities or a decrease in cash reserves. Additionally, when preventing revaluation or appreciation of the domestic currency by interventions, higher currency holdings are offset by a decrease in credit facilities or an increase in cash reserves.

Different exchange rate regimes demand different investment styles. Under a fixed rate regime, the central bank will intervene more often than under a floating rate regime, and thus foreign assets will need to be invested in highly liquid assets with short duration. It is interesting, however, that despite these differences, a general model can be derived for balance sheet optimisation which applies to almost any central bank. This model will be presented in Section 3. Differences in exchange rate regime can be accounted for simply by changing one or a few parameters.

Table 1 shows a stylised balance sheet for a central bank that issues banknotes, holds gold and foreign assets and implements monetary policy. Capital and reserves act as a financial buffer to cover (unrealised) losses. Here, we use a broad definition of capital and reserves which may include revaluation accounts created out of unrealised profits on certain assets (notably gold). Capital and reserves (narrow definition) accumulate retained profits. The difference between capital and reserves, on the one hand, and revaluation accounts, on the other, is largely determined by accounting principles. Note, however, that central banks can usually retain only part of their profits. In our model, the distinction between both types of reserves is of minor importance and, consequently, the two are combined. Finally, losses may also stem from events outside the balance sheet. For a central bank entrusted with bank supervision, for example, there is always a (small) risk that the central bank will be held

Table 1: Stylised Balance Sheet

Assets		Liabilities	
Gold	G	Banknotes	B
Foreign assets	F	Cash reserves	C
Lending to financial sector	L	Capital and reserves	R
Other assets	O		

accountable should a commercial bank fail. A central bank with supervision responsibilities should therefore *ceteris paribus* hold more capital and reserves. Capital and reserves play an important role in the approach presented below with respect to managing market risks.

3 Balance sheet approach

The balance sheet approach presented here to manage market risks is essentially based on the need for central banks to be independent. Apart from the much-discussed question in the economic literature on how independent a central bank should be when implementing monetary policy (see, for example, Goodhart, 1990, Cukierman, 1992 or Alesina and Summers, 1993), central bank independence is also a question of financial independence. A central bank unable to cover its losses becomes dependent on alternative (government) resources and would therefore become vulnerable to external influences. One could argue that a central bank is, at least in theory, always capable of covering its losses. The fact that a central bank is a money-creating institution indeed allows for this. However, history has taught us that an insolvent central bank will ultimately lose its credibility and cease to exist.

Two dimensions can be distinguished with respect to financial independence: first of all, profit-generating capacity, and secondly, a central bank's ability to absorb losses. The first refers to the profit and loss account, while the second takes the financial buffers (revaluation accounts, capital and reserves) into account. Sub-section 3.1 starts with a general description of a central bank's expected return. Sub-section 3.2 then provides an introduction to market risk and potential losses. Finally, sub-section 3.3 presents the optimisation model.

3.1 Expected return and structural profit

Given the stylised balance sheet (see Table 1), a central bank's profitability on an annual basis can be defined as:

$$\begin{aligned} E(P) &= r_G G + r_F F + r_L L + r_O O - r_C C - OC \\ &= \mathbf{r}X - r_C C - OC, \end{aligned} \tag{3-1}$$

with $E(P)$ = expected profit,
X = vector of assets (G, F, L, O) (see Table 1, market values),
r = vector of corresponding expected total returns, r_i = expected return on asset i,
r_C = reimbursement on cash reserves C,
OC = operating costs.

In this simplified set-up, the expected profit is merely a function of the expected return on the assets on the one hand, and the expenses related to the reimbursement on the cash reserves and the operating costs on the other. Expected returns may contain an explicit view on the development of exchange rates and yield curves, but can also be obtained from simple assumptions, e.g. unchanged exchange rates and yield curves (an implicit market view). Historic returns are less suitable for predicting future returns on fixed income portfolios, especially given current low interest rates. The value of assets and liabilities is assumed to be endogenous, but may in reality also be affected by largely unpredictable exogenous inflows or outflows, e.g. from interventions.

From this simple set-up, two interesting observations can be made. The first observation is derived by assuming that gold does not generate any return ($r_G = 0$). This is a reasonable assumption, all the more since the return calculation on gold is often based – due to its illiquid nature – on valuation and accounting principles which deviate quite strongly from a total return approach. This means that any fluctuations in the gold price are usually accommodated in a revaluation account and, hence, do not contribute to the profit and loss account. Furthermore, it is reasonable to assume that in the long run, returns on domestic and foreign assets do not really differ ($r_L = r_F = r_O$). If one also assumes that cash reserves are interest-bearing with a return equal to the return on lending to the financial sector ($r_L = r_C$), one can derive the formula for the structural profit of a central bank. (Note that $G + F + L + O = B + C + R$, and thus $F + L + O - C = B + R - G$.)

$$E(P_s) = r_L(B + R - G) - OC, \qquad (3\text{-}2)$$

with $E(P_s)$ = (expected) structural profit.

Equation (3-2) shows that the structural profit of a central bank is driven by the interest rate level on the one hand, and the counterpart of non-interest-bearing liabilities not invested in gold on the other. This equation can also be used to draw a distinction between seigniorage or monetary income ($r_L B$) and other sources of structural profit. Note that r_L in equation (3-2) denotes a long-term forecast, while the horizon in equation (3-1) is typically one year. Consequently, the expected values may differ. In equation (3-2), it is common to derive structural profits from the current level of interest rates.

The second observation follows when the assumption that all returns are equal is relaxed. Under normal circumstances, yield curves are upward-sloping. Depending on the state of the economy, the long-run return on foreign assets might very well differ from the return on monetary policy operations. These observations could be accommodated by introducing additional reference rates. Here, only one is added, which represents the average return on assets not related to monetary policy (r_N). The formula for structural profit would then become:

$$E(P_s) = [\omega r_N + (1 - \omega) r_L](B + R - G) - OC, \qquad (3\text{-}3)$$

with ω = weight of assets not related to monetary policy, $0 \leq \omega < 1$.

Equation (3-3) can be used to identify an interesting dilemma for central banks. If in the long run the expected return on assets not related to monetary policy is higher than the expected return on lending to the financial sector, then there is an incentive to increase profitability at the expense of monetary policy operations. However, profit optimisation has its price, as will be discussed in the next sub-section.

3.2 Value-at-risk and potential loss

Just like any investor, a central bank bears market risks – notably foreign exchange risk and interest rate risk – as well as some other risks that are more difficult to quantify. Given our stylised balance sheet, whereby the central bank does not issue any debt other than banknotes, market risk arises only on the asset side of the balance sheet and may be quantified using a Value-at-Risk (VaR) measure. VaR measures the worst expected loss over a given horizon under normal market conditions at a given confidence level (Jorion, 2001). Over the years, dozens of methodologies for calculating VaR have been developed; common approaches include the asset-normal or delta-normal approach (also known as the variance-covariance approach), historical simulation and Monte Carlo simulation. For an overview, see Jorion (2001) or Manganelli and Engle (2001).

Here we propose to use the relatively straightforward asset-normal approach, which assumes that asset returns have a joint normal distribution. While it is widely known that returns are not distributed normally – in particular, they have fatter tails – two arguments justify our choice for simplicity. First, the goal is not to avoid 'embarrassing losses', as these can not be prevented by any risk management system (Wilson, 1998), but rather to compare alternative portfolios. Second, linear instruments without optionality usually predominate in central bank portfolios. For these portfolios and purposes, the normal distribution is sufficiently accurate. It is stressed that any method for calculating VaR may be chosen; the conclusions of this paper do not depend on the VaR methodology used. Consequently, VaR is calculated as in equation (3-4):

$$\text{VaR} = \rho\sqrt{t\mathbf{X}'\mathbf{\Sigma}\mathbf{X}}\,, \tag{3-4}$$

with ρ = scaling factor corresponding to a chosen confidence level,
$\mathbf{\Sigma}$ = covariance matrix of asset returns,
t = time horizon.

The covariance matrix may be estimated in many ways, using different historic observation periods and different decay factors. At DNB, we use a history of five years and monthly observations which are equally weighted. In this respect, our approach differs from the standard RiskMetrics™ approach, which applies a decay factor to historic returns (RiskMetrics, 1996). The covariance matrix should be derived from returns that are representative of the asset class. For fixed income portfolios, this implies that the (expected) composition of the asset class (by issuer as well as maturity) should be used to reconstruct historic time series of asset returns.

Value-at-Risk was developed for measuring market risk in trading portfolios with short horizons (the Basel Committee on Banking Supervision, 1996, prescribes a holding period of ten business days). This implies that expected returns over the holding period can be ignored. Here, we propose to apply the same technique to quantify market risk over a longer horizon – one year or even longer – and expected returns can no longer be ignored. We introduce the concept of potential loss (*PL*), defined as VaR less expected profit:

$$PL = \text{VaR} - E(P), \tag{3-5}$$

with *E(P)* as defined in equation (3-1).

In theory, *PL* may be negative and potential loss is thus actually a profit. However, given the substantial (and typically regarded as unavoidable) foreign exchange risk, *PL* will normally be positive and represent the maximum loss, given a certain confidence level. It is important to realise that VaR and *E*(*P*) are not independent. VaR is derived from potential adverse movements in risk factors, e.g. interest rates. A rise in interest rates will result in a drop in market value. However, at the same time, it will increase reinvestment return, thereby increasing $E(P)$. Accounting for this dependence will substantially increase the complexity of the model. For the sake of simplicity, we assume that VaR and $E(P)$ may be calculated independently, which will often result in a conservative estimate of potential loss.

The ability to absorb losses is determined by confronting the potential loss (PL) with the total amount of reserves (R), i.e. the sum of the revaluation accounts and capital and reserves. This confrontation can be seen as an important indicator with respect to risk tolerance. The indicator can be extended in principle to a period longer than one year[2] and therefore also permits, in combination with the calculation of structural profits, interpretations of a more strategic nature. The focus here is on explaining the main principles of the presented approach, and is therefore limited to a simple static model.

3.3 Managing market risk

By defining the expected profits and potential losses of a central bank, the main elements of a risk/return framework to manage and evaluate market risk can be identified. To optimise the composition of the balance sheet, the next step is to solve an optimisation model whereby the aim is to maximise profits. Given the nature of central banks, maximising profits may seem slightly odd, and of course it is possible to choose a different target (for instance risk minimisation). However, since the goal here is not to describe central bank behaviour but to develop a framework for managing market risks, aiming for profit maximisation is quite acceptable. Other ingredients for the optimisation model are constraints with respect to certain asset classes and a formal description of risk tolerance. The constraints stem from the responsibilities delegated to a central bank. Implementing monetary policy effectively could imply a minimum liquidity deficit ($L \geq L^{\min} \geq 0$). The exchange rate regime may require a certain amount of foreign assets ($F \geq F^{\min} \geq 0$). For example, as explained in Section 2, under a currency board regime the amount of foreign assets should at least be equal to the amount of banknotes in circulation ($F^{\min} = B$). Furthermore, to anticipate very exceptional circumstances, a minimum threshold for gold holdings can be introduced ($G \geq G^{\min} \geq 0$). By introducing the requirement that the total amount of reserves should be high enough to cover at least α times the potential losses (given a certain confidence level), a simple description of risk tolerance can be created ($\alpha\, PL \leq R$). A low risk tolerance is associated with a high value for α and vice versa. By changing α, the model may be used to find the efficient frontier:

$$\text{Maximise } E(P), \tag{3-6}$$

subject to $L \geq L^{\min}$,
$F \geq F^{\min}$,
$G \geq G^{\min}$,
$\alpha\, PL \leq R$,
$O \geq 0$ (no short selling),
$G + F + L + O = B + C + R$ (total assets equal total liabilities, B, C and R are known and assumed constant).

The model is the classical mean-variance optimisation model. Here, the objective function is linear in the decision variables, while one of the constraints ($\alpha\, PL \leq R$) is quadratic. The feasible region is convex, which implies that any local optimum will also be the global optimum. The model may be classified as a convex programming model, for which standard

[2] $PL^T = \sqrt{T}\,\text{VaR} - T\,E(P)$

optimisation techniques are available in many (statistical) software applications, including Excel, which applies an algorithm developed by Lasdon and Warren (1978).

Solving the model enables the optimal composition of the balance sheet to be determined, or more specific, the optimal asset composition ($G°, F°, L°, O°$) for a central bank. The model can be used to evaluate the characteristics of the market risks to which a central bank is exposed. As a first step, the total exposure by asset classes can be decomposed by calculating the market exposure for each single asset class. The sum of the potential losses on the individual market exposures is an estimate of the market risks which a central bank could be confronted with if diversification effects are ignored. Furthermore it can be noted that a – presumably major – proportion of the market risks taken cannot be avoided since they are directly related to the (monetary) obligations of a central bank. The framework presented above allows for a distinction between avoidable and unavoidable market risks. The latter can be calculated by taking into account only the minimum amounts for gold and foreign assets (G^{min} and F^{min}) and all other assets invested in the least risky asset, i.e. lending to the financial sector. Avoidable market risk is simply defined as the difference between actual risk (on optimal portfolio) and unavoidable risk. Central banks are increasingly transparent in their risk management and risk exposure. DNB, for instance, publishes its avoidable risk ignoring diversification between fixed income and equities in its annual report (see De Nederlandsche Bank, Annual Report 2002, p. 147).

While the techniques for solving model (3-6) have become relatively straightforward, the optimum should be treated with care. The model is very sensitive to estimation errors in expected returns and the covariance matrix. Typically, the model will recommend assets with high expected returns and low volatility and/or low correlations with other assets. Extreme solutions are particularly likely when the risk tolerance constraint ($\alpha PL \leq R$) is not binding, in which case the model would simply allocate all money to the asset with the highest expected return. Under normal circumstances, however, the risk tolerance constraint will be binding. Furthermore, small changes in parameters may lead to major changes in the optimal portfolio. As a result, mean-variance optimisation is sometimes dubbed "error maximisation" (see, for instance, Sharpe et al., 1999, pp. 176-77 or Michaud, 1989).

The model may be improved and the stability of results may be increased by introducing additional constraints, for instance on the maximum amount invested in each asset class. Another option is to increase the number of decision variables, even though this would enlarge the covariance matrix. Instead of treating each asset class as one homogenous asset, one could distinguish several investment categories (by issuer and by maturity) within each asset class. If fixed income assets are spread evenly over maturities, the likelihood of extreme solutions is reduced substantially. In any case, it is recommended that the robustness of the results should be verified.

Finally, the impact of changes in exogenous factors that may affect profits and the loss-absorbing capacity of a central bank can be analysed with the model. Some examples will be discussed in the next section to illustrate the model's main characteristics.

4 Impact of exogenous influences

For the central banks participating in the European System of Central Banks, the introduction of the euro implied a regime shift. The currency regime shifted from managed floating to floating, while the external value of the currency became a shared responsibility. In terms of the model, F^{min} and G^{min} decreased. Given that, under normal circumstances, at least one of the minimum allocation constraints on gold ($G \geq G^{min}$) and foreign assets ($F \geq F^{min}$) is binding – the risk on these assets is far higher than the risk on other assets and lending to the financial sector, and does not justify large allocations to both asset classes – the model predicts a shift from gold and/or foreign assets to other assets and lending to the financial sector. To put it another way, the model implies a move away from high-risk assets towards low-risk assets when responsibilities are shared. When looking at the major changes in the balance sheet of the European System of Central Banks, a major shift towards domestic assets has indeed taken place at the expense of foreign assets and gold.

Electronic payment devices, such as credit and debit cards, are becoming increasingly popular. So far these alternative payment devices have not led to a general decrease in the amount of banknotes outstanding. It is not hard to imagine, however, that such an impact might occur in the future. Given the model above, it is easy to hypothesise that when the amount of banknotes outstanding decreases, the amount of market risk which a central bank can bear comes under pressure, because a decrease in the amount of banknotes in circulation would lead to lower (structural) profits and as such to higher potential losses (see equation 3-2). This would imply – given an unchanged amount of reserves and a binding risk tolerance constraint ($\alpha\, PL \leq R$) – an adjustment in the amount of assets held in such a manner that market risks decrease. Generally speaking, this will be established by adjusting the amount of foreign assets and/or other financial assets downward, thereby increasing lending to the financial sector or decreasing cash reserves. As soon as the optimal amount of assets reaches the lower bounds imposed by (monetary) obligations, one would be forced to reconsider the attitude towards market risks or aim for higher financial buffers to ensure financial independence. Groeneveld and Visser (1997) argue that central banks have additional options to compensate for the loss of seigniorage, such as imposing reserve requirements on electronic money, or the charging of fees for services provided, e.g. bank supervision.

A similar observation can be made if interest rates decline and reach low levels, as this also implies lower (structural) profits. Since the late 1970s, interest rates in the major currency areas have declined to historically low levels. The model described above would under these circumstances imply a decrease in the amount of assets with relatively volatile return characteristics or higher financial buffers.

When assuming that a central bank is managing its market exposures in such a manner that the balance sheet composition could be called optimal, interventions in the foreign exchange market would always imply a deterioration in the risk return profile. Interventions to avoid a depreciation or devaluation would lead to a decrease in foreign assets below the optimal level, whereas preventing appreciation or revaluation would imply an accumulation of foreign assets beyond the optimal amount. These observations illustrate that the model is less suitable for evaluating market risks if, due to external circumstances (i.e. policy requirements), the balance sheet composition is relatively unstable. Nevertheless, the model could even under these circumstances serve as an anchor for evaluating profitability and potential losses, and thus promote risk awareness.

5 Concluding remarks

The model presented above offers an integrated approach towards market risks. It is based on an absolute risk parameter (potential loss), and can be used to analyse the financial strength of a central bank. These are important characteristics of the model, which are becoming ever more important. Central banks should manage their assets in an efficient way and will be held accountable for any losses if they occur. Moreover, low interest rates have made central banks more vulnerable to losses.

The model has its limitations. First of all, it should be mentioned that central banks usually have more responsibilities than the ones mentioned, for instance in respect of supervision and payment systems. These responsibilities might create a need for additional capital and reserves. Furthermore, like other financial institutions, central banks are confronted with other types of risk, such as credit and operational risk, which once again create a need for financial buffers. The model, on the other hand, only addresses market risks, and is fairly static in nature. More sophisticated insights into the financial strength of a central bank might be drawn from a proper asset/liability study, similar to the ones conducted by pension funds. Despite these shortcomings, the model might prove a powerful tool to promote risk awareness and evaluate financial strength within the central banking community.

References

Alesina, A. and L. H. Summers (1993), Central Bank Independence and Macroeconomic Performance, *Journal of Money*, *Credit and Banking* 25: 151-162.

Basel Committee on Banking Supervision (1996), Amendment to the Capital Accord to Incorporate Market Risks, BIS, Basel.

Cukierman, A. (1992), *Central Bank Strategy, Credibility and Independence: Theory and Evidence,* MIT Press, Cambridge, Mass.

De Nederlandsche Bank, *Annual Report 2002*, De Nederlandsche Bank, Amsterdam.

Frowen, S. F., R. Pringle and B. Weller (eds) (2000), *Risk Management for Central Bankers,* Central Banking Publications, London.

Goodhart, C. A. E. (1990), *The Evolution of Central Banks,* 2nd edition, MIT Press, Cambridge, Mass.

Groeneveld, J. M. and A. Visser (1997), Seigniorage, Electronic Money and Financial Independence of Central Banks, *BNL Quarterly Review,* 200: 69-88.

IMF (2001), Guidelines for Foreign Exchange Reserve Management, available at www.imf.org/external/np/mae/ferm/eng/index.htm.

IMF (2003), Guidelines for Foreign Exchange Reserve Management: Accompanying Document, available at www.imf.org/external/np/mae/ferm/2003/eng/index.htm.

Jorion, P. (2001), *Value at Risk: The New Benchmark for Managing Financial Risk,* 2nd edition, McGraw-Hill, New York.

Lasdon, L. S. and A. D. Warren (1978), Generalised Reduced Gradient Software for Linearly and Nonlinearly Constrained Problems, in H. G. Greenberg (ed.), *Design and Implementation of Optimization Software,* Sijthof and Noordhoff, Alphen aan den Rijn, The Netherlands.

Manganelli, S. and R. F. Engle (2001), Value at Risk Models in Finance, European Central Bank Working Paper 75.

Michaud, R. O. (1999), The Markowitz Optimization Enigma: Is 'Optimized' Optimal?, *Financial Analysts Journal,* 31-42.

RiskMatrics (1996), *Technical Document,* Morgan Guaranty Trust Company of New York.

Sharpe, W. T., G. J. Alexander and J. V. Bailey (1999), *Investments,* 6th edition, Prentice Hall, Upper Saddle River.

Wilson, T. C. (1998), Value at Risk, in C. Alexander (ed.), *Risk Management and Analysis Volume 1: Measuring and Modelling Financial Risk,* Wiley, New York.

13

Ex post risk attribution in a value-at-risk framework

Eugen Puschkarski

Abstract

In this article we will first describe a general procedure to decompose time-variation in Value-at-Risk from one reporting period to the next. Then, using standard methodology from the field of performance attribution, we analytically show how the new VaR Risk Attribution Model (RAM) ascribes these changes to an active trading factor, a market risk changes factor, a passive time decay factor and a resulting cross-product. With a slightly simplified version of the RAM, we subsequently demonstrate how risk-taking patterns can be detected in practice. We highlight the relevance of using a RAM for central banks and subsequently set the presented RAM into the context of existing risk attribution methods.

1 Introduction

Value-at-risk (VaR) has become the dominant risk measure for the sell-side of investment banking, and lately also for the so-called buy-side of investment management. One reason why VaR is widely used is that it can be aggregated across different portfolio hierarchies, investment product types and can be measured in a common denominator, namely as an amount of money.

Unlike its risk measure predecessor, volatility, VaR is also sensitive to moments higher than O(2), and can therefore also be applied to non-normal distributions. Financial data are usually negatively skewed and have a positive kurtosis, so that VaR has proven to be an appropriate risk measure for the financial community.

To use VaR as a risk controlling tool within the organisation and to comply with banking supervision guidelines, calculations need to be run periodically, applying the same risk parameters and risk methodology. The consistency of periodic risk calculations is vital for comparative risk analysis, and keeps the analysis from being simply ad hoc. With respect to VaR calculations, such static risk parameters usually include:

- the time decay factor for calculating volatilities and correlations;
- the length of the historical observation period from which volatilities and correlations are calculated, for example;
- the number of MC trials;
- distributional assumptions about risk factors;
- the VaR quantile (usually chosen as 95% or 99%);
- the method of time series gap filling;
- the time horizon (holding period) (usually one day to one month).

Additionally, many more are possible, depending on the VaR method used.

Unlike the mentioned static VaR parameters, risk managers update the market data set and the position loads usually at the same frequency as the periodic risk reports, to ensure that the reported VaR numbers reflect the most recent risk profile of the institution, as well as to give management valid and accurate information.

2 Risk attribution model (RAM)

VaR calculations can be represented as a very general function of the form

$$VaR_t(P_t, M_t, RD_t) \quad (1)$$

where t represents a time index, P denotes the financial positions for which the VaR is to be calculated, and M stands for the market data used in the VaR calculations.

The concrete market data M necessary for VaR calculations depend on the one hand largely on the VaR calculation method used, and on the other on the attributes of the financial instruments and markets in which these are traded. For example, the historical simulation and Monte Carlo methods need time series of risk factors such as stock prices, interest rates, volatilities and foreign exchange rates to calculate VaR. The variance-covariance method of calculating VaR requires, as its name suggests, a covariance matrix in order to calculate VaR analytically.

RD stands for a reference date which is used to calculate the time to maturity for fixed maturity financial products such as bonds, futures, swaps, options, etc.

The aim of this paper is to explain the difference in VaR from one report date (t-1) to the next (t).

$$VaR_{\Delta t} = VaR_t(P_t, M_t, RD_t) - VaR_{t-1}(P_{t-1}, M_{t-1}, RD_{t-1}) \quad (2)$$

When VaR calculations are periodically made, in addition to the static information about the VaR method used and the distributional assumptions applied, the information about P, M and RD is changed from report date (t-1) to report date (t) accordingly.

The following graph depicts the possible combinations of P, M and RD in three-dimensional space.

Diagram 1

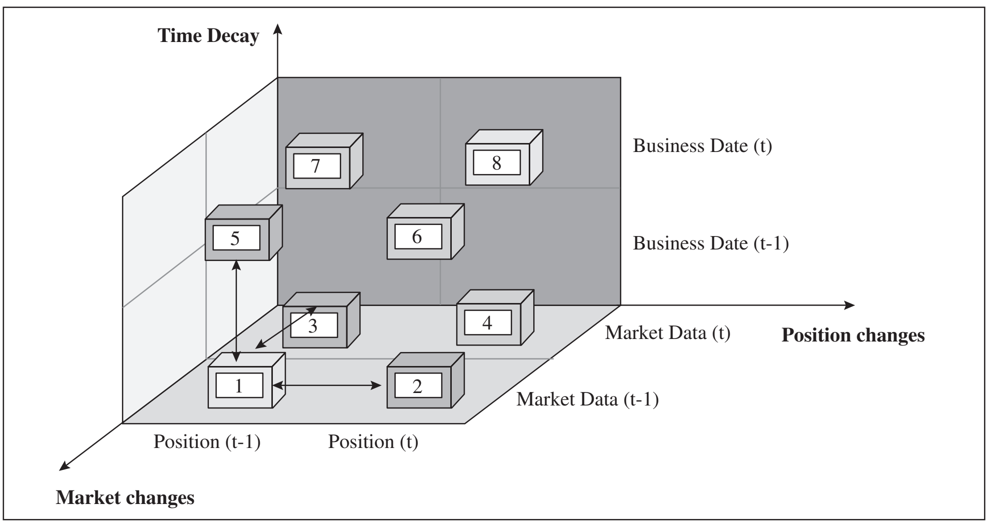

In the following we adapt the simple performance attribution method outlined in Brinson et al. (1986) and (1991) to the special requirements of a risk measurement framework. Unlike the Brinson model, the chosen benchmark against which the portfolio return is measured is not a separate benchmark but the lagged portfolio holdings.

In the diagram the different VaR outputs are symbolised by a cube (C). Let us say the VaR of report date (*t*-1) with positions held at (*t*-1) and market data collected up to time (*t*-1) is represented by C1, whereas C8 denotes the calculations made with position and market data at time (*t*). The difference between the VaR of C8 minus the VaR of C1 represents the total difference we now intend to explain with the RAM.

Total change in VaR from (*t*-1) to (*t*) = **C8 – C1**

This change in VaR in the course of time is often monitored by management to see whether a pre-set total VaR limit has been breached, and whether the risk appetite of the institution is still adequately being met.

2.1 Change in VaR owing to active trading

To calculate the effect of position changes owing to active trading from (*t*-1) to (*t*) on the total VaR difference, C2 is subtracted from C1. Here C2 represents a VaR scenario as it stands for the VaR which would have been measured at time (*t*-1) if the institution were to hold the positions of time (*t*) at time (*t*-1), *ceteris paribus*. This scenario can also be imagined as if the portfolio traders at time (*t*-1) instantly swap positions so that they hold the positions of time (*t*) at time (*t*-1).

Change in VaR owing to position changes from (*t*-1) to (*t*) = **C2 – C1**

2.2. Change in VaR owing to market risk factor changes

In addition, the framework allows the effect of changed market conditions owing to market risk factor fluctuations from (*t*-1) to (*t*) on the total VaR difference to be calculated. To do this, a further scenario (VaR C3) has to be calculated which is equal to the VaR of the positions held at time (*t*-1), but using the market data observed at time (*t*). The difference between this scenario and the initial VaR C1 can obviously be ascribed to the changes in market prices.

Change in VaR owing to market changes from (*t*-1) to (*t*) = **C3 – C1**

2.3 Change in VaR owing to time decay

Most financial products are of the fixed maturity type, which means that their basic cash flows accrue at certain points in time. This means that for every future VaR calculation the time span to the cash flows becomes shorter as the VaR is recorded through time. To determine this effect on VaR owing to time decay, another scenario (VaR C5) has to be calculated, where the reference date *RD* for calculating the time to maturity of the cash flows is set from (*t*-1) to (*t*). Again this change in the *RD* is done *ceteris paribus* for time (*t*-1). The difference between VaR C5 and VaR C1 can then be attributed to the time decay of fixed maturity-type financial instruments, which normally has a risk-decreasing effect on VaR.

Change in VaR owing to time decay from (*t*-1) to (*t*) = **C5 – C1**

It should be noted that changing the reference date *RD* without changing the date of financial positions is not possible in some “off-the-shelf” VaR applications. The author is for example aware of KVaR+ from Reuters, which does not allow the position date to be separated from the reference date. As a result, it is not possible to calculate how the VaR would change if all positions are one reporting period closer to their maturity. On the other hand, RiskManager™ from the RiskMetricsGroup offers the ability to separate *RD* from the positions date.

2.4 Change in VaR owing to the cross-product term, and an overview

The following overview shows how the different VaR components add up to the total difference if a cross-product term is added.

Table 1: RAM Overview

	Calculations
Position changes	C2 – C1
Market changes	C3 – C1
Time decay	C5 – C1
Cross Product	C8 – C5 – C2 – C3 + C1 + C1
Total VaR difference	C8 – C1

Laker (2000) provides a conceptual discussion about the cross-product term. Mathematically, the cross-product term ensures that the model is additive and represents intuitively the fact that in reality position changes, market changes and the decay of time to maturity happen simultaneously and cannot be separated from each other. The cross-product term can therefore be regarded as a "missing link" between the three factors.

The magnitude and sign of the cross-product term quantifies the extent to which the three factors influence each other by either amplifying or extenuating each other.

3 Risk attribution in practice

In the following section we present a slightly modified version of the RAM presented above, providing examples of risk-taking patterns that become more transparent with a RAM in place.

To facilitate risk communication in terms of VaR attribution to senior management, as well as to simplify the risk attribution calculations and lessen the requirements on the risk analysis software, the following shortcuts are made.

First, the change in VaR owing to a changed market risk environment is calculated exactly as in the full RAM with C3 – C1. Unlike the full RAM, the simplified version (sRAM) incorporates position changes, time decay and the cross-product all in the position change term. The reasoning behind this is simply to assume that where the change in VaR is not attributable to market risk changes, it is instead under the control of active risk management and is largely controlled by position-taking.
sRAM change in VaR owing to position changes: **C8 – C3**

Table 2: sRAM Overview

	Calculations
Position changes	C8 – C3
Market changes	C3 – C1
Total VaR difference	C8 – C1

With sRAM in place, it is much easier to keep the risk attribution synoptical. Additionally the simplified version requires only one VaR scenario (C3), and eliminates the need for C2, C5 or the calculation of the cross-product term. The scenario VaRs can be computationally rather costly with respect to duration of calculations when using the Monte Carlo method of calculating VaR. Additionally, as previously mentioned, the calculation of C5 is not always implemented in VaR software packages and might have to be customised. Of course, the simplification comes at the cost of losing risk drill-down options and more refined analyses.

The following chart shows the weekly change in VaR[1] of the foreign reserves (including gold holdings) of the Oesterreichische Nationalbank recorded over the course of approximately one year, and analysed with the sRAM.

Chart 1

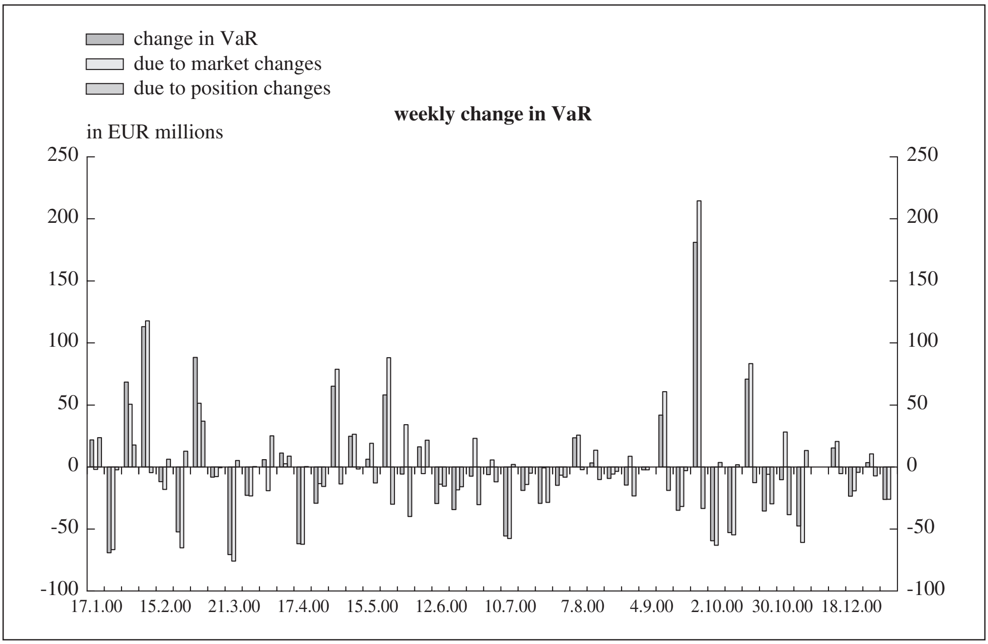

One of the first observations that we can make based on this overview is that the change in VaR is more driven by the change in market risk factors than the change in positions. What is even more interesting and provides insight into how market risk factor changes and corresponding reactions from traders interact becomes visible by ranking the market changes by magnitude.

[1] The VaRs shown use a 95% quantile and are calculated using the parametric approach of RiskManager 2.3 from the RiskMetricsGroup. The holding period is set to one month, as this equals the interval of reviewing the tactical benchmark.

Chart 2

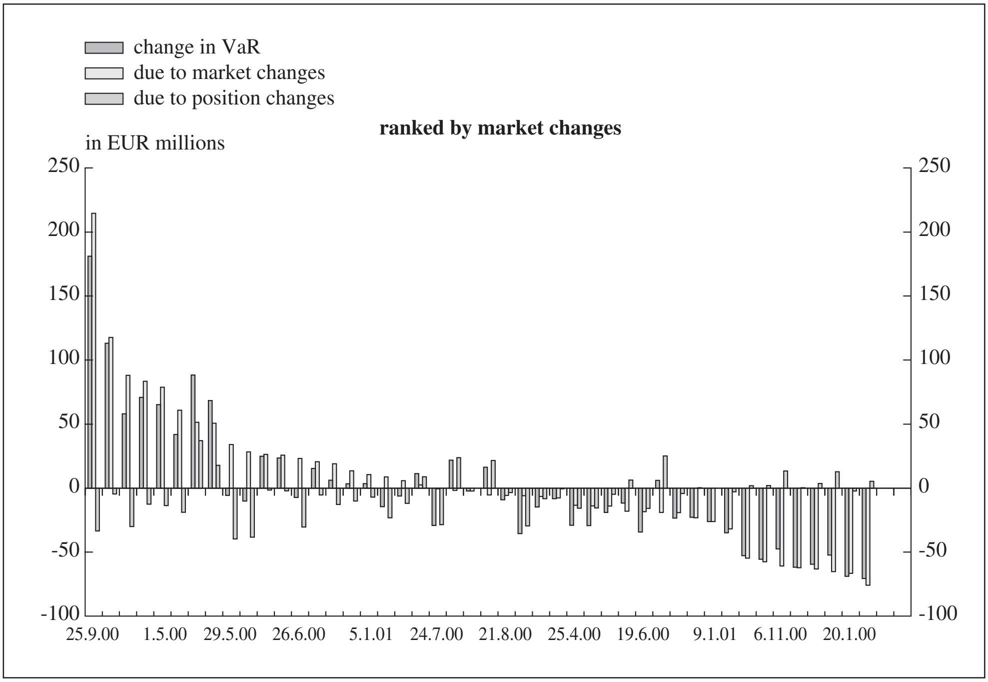

Chart 2 reveals a fairly consistent pattern of risk balancing by active trading on the part of the dealers. For large increases in market risk, the changes in positions are markedly risk-reducing, whereas for large reductions in market risk, traders exhibit a risk-taking propensity.[2]

This risk-balancing pattern of position changes is even more apparent in the following scatter plot, where VaR changes attributed to market changes are measured against VaR changes attributed to position changes.

2 At this point it should be noted that the long-term strategy of setting risk appetite is not determined by the traders but by senior management. So the bias in favour of risk reduction is rather due to the strategic risk policy.

Chart 3

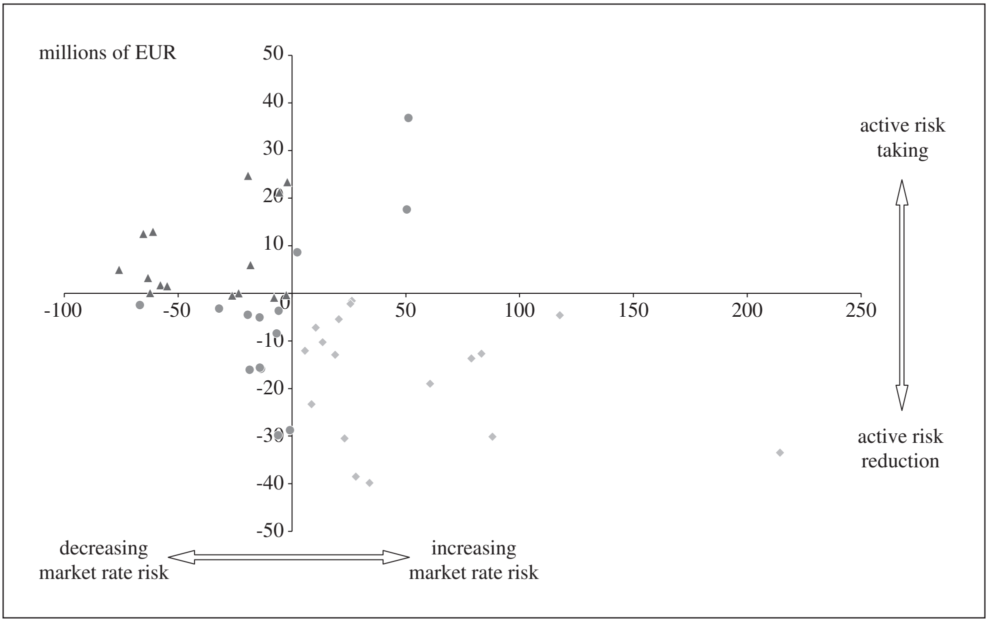

The diamonds represent the risk-reducing position changes which were accompanied by increasing market rate risk. The triangles show active risk-taking when the market environment was decreasing in risk.

4 Relevance of using an RAM for Central Banks

Historically, central banks mainly invested in fixed income instruments with set duration limits. Nowadays, the diversity of modern investment products calls for a comprehensive risk measure such as VaR that is capable of quantifying risk for fixed income products as well as equities, commodities and derivatives. As duration measures become increasingly meaningless in a risk-type diversified portfolio, VaR limits have recently been advocated as a suitable alternative. With VaR limits in place, an RAM helps monitor these limits.

Apart from detecting risk taking or risk reduction, the RAM is also capable of answering the question *why* set risk limits have been breached. A risk monitoring system, as presented in Culp et al. (1998), who use “soft” risk targets and/or “hard” risk limits within a set risk budget, will benefit from the detection of the specific cause that led to a limit overdraft. In case the change in market rate risk was the driving factor behind such a limit overdraft, one can interpret the breach as a passive one, which should lead to a discussion about the market environment and a possible readjustment of risk limits. For example, market risk generally increases rapidly in a falling market price environment. Automatically selling positions to reduce risk and stay within given risk limits can be dangerous without a detailed analysis of the reasons for the downturn. Using an RAM clearly separates the passive components of changes in market risk from active market risk. On the other hand, traders could be held directly responsible for a limit overdraft mainly driven by position changes. Active limit breaches should trigger a re-evaluation of the investment strategy and lead to a set of good and plausible arguments as to why taking on additional risk is a good strategy at this point of time.

5 Putting ex post risk attribution into perspective

Culp et al. (1999) highlight the fact that performance measurement and attribution are done ex post, rendering a backward-looking analysis of past returns. Risk analysis, on the other hand, is done ex ante, giving a forward-looking estimate of possible future returns. Risk attribution in the conventional sense is therefore also a forward looking breakdown of major risk drivers. For example, the usual categories for a VaR drill-down report include risk types (currency risk, interest rate risk, price risk, volatility risk, etc.), portfolio hierarchy, maturity of financial instruments, period of interest rates (1M, 3M, 1Y,….., 30Y+), and so on.

Moreover, risk attribution can also be seen relative to a benchmark showing the active risk taken by traders. This approach is presented in Mina (2002).

Unlike this conventional view of risk attribution, the RAM presented in this paper builds a link between the ex ante method of measuring risk and the ex post method of performance attribution. Basically, the ex ante risk forecasts of VaR are analysed ex post to detect the main drivers behind the updated VaR forecasts.

It is also possible to drill-down the risk attribution even further, giving insight into which market rates have most influenced the change in VaR owing to market changes. For example, a second-level risk attribution of market changes to risk types could provide management with a more detailed picture of the risk profile. Additionally, the position changes could be further broken up into asset types such as bonds, swaps, FRAs, futures, etc.

The question of how far to drill-down the risk attribution and in what periodicity the RAM should be used largely depends on the market risk of the traded instruments, whether the analysis is done for the traders or for management, and how capable the IT infrastructure is.

6 Conclusion

The presented method represents a new way of explaining time-variation of VaR, thereby giving management a useful tool for monitoring the risk-taking behaviour of traders. The analysis also complements existing methods of risk attribution, and aims to help in judging whether given risk budgets are appropriately used by traders. Moreover, it helps in the risk control process of monitoring VaR limits, which is becoming ever more important and needs to be easy to communicate to senior management.

References

Brinson, Gary P., L. Randolph Hood, and Gilbert L. Beebower, "Determinants of Portfolio Performance", *Financial Analysts Journal*, July-August 1986, pp. 39-44.

Brinson, Gary P., Singer, Brian D., and Gilbert L. Beebower, "Determinants of Portfolio Performance II: An Update", *Financial Analysts Journal*, May-June 1991, pp. 40-48.

Culp, Christopher L., Mensink, Ron, Neves, Andrea M. P., "Value at Risk for Asset Managers", *Derivatives Quarterly,* Vol. 5, No. 2, Winter 1998.

Culp, Christopher L., Mensink, Ron, "Measuring Risk for Asset Allocation, Performance Evaluation, and Risk Control: Different Problems, Different Solutions", *Journal of Performance Measurement,* Vol. 4, No. 1, Fall 1999.

Laker, Damien, "What Is this Thing Called 'Interaction'?, *Journal of Performance Measurement*, Vol. 5, No. 1, Fall 2000.

Mina, Jorge, "Risk Attribution for Asset Managers", *RiskMetrics Journal*, Vol. 3, No. 2, Winter 2002-2003, pp. 33-55.

14

Ruin theory revisited: stochastic models for operational risk

Paul Embrechts, Roger Kaufmann and Gennady Samorodnitsky

Abstract

The new Basel Capital Accord has opened up a discussion concerning the measurement of operational risk for banks. In our paper we do not take a stand on the issue of whether or not a quantitatively measured risk capital charge for operational risk is desirable; however, given that such measurement will come about, we review some of the tools which may be useful towards the statistical analysis of operational loss data. We also discuss the relevance of these tools for foreign reserves risk management of central banks.

1 Introduction

In [9], the following definition of operational risk is to be found: "The risk of losses resulting from inadequate or failed internal processes, people and systems or from external events." In its consultative document on the New Basel Capital Accord (also referred to as Basel II or the Accord), the Basel Committee for Banking Supervision continues its drive to increase market stability in the realms of market risk, credit risk and, most recently, operational risk. The approach is based on a three pillar concept where Pillar 1 corresponds to a Minimal Capital Requirement, Pillar 2 stands for a Supervisory Review Process and finally Pillar 3 concerns Market Discipline. Applied to credit and operational risk, within Pillar 1, quantitative modelling techniques play a fundamental role, especially for those banks opting for an advanced, internal measurement approach. It may well be discussed to what extent a capital charge for operational risk (estimated at about 12% of the current economic capital) is of importance; see Daníelson et al. [23], Goodhart [46] and Pezier [60, 61] for detailed, critical discussions on this and further issues underlying the Accord.

In our paper we start from the premise that a capital charge for operational risk will come about (eventually starting in 2007), and we examine some quantitative techniques which may eventually become useful in discussing the appropriateness of such a charge, especially for more detailed internal modelling. Independent of the final regulatory decision, the methods discussed in our paper have a wider range of applications within quantitative risk management for central banks, the financial (including insurance) industry and supervising authorities. In particular, we are convinced that these methods will play a role in the construction of quantitative tools for integrated risk management, including foreign reserves risk management for central banks. First of all, in an indirect way, central banks have a keen interest in the stability of the global banking system and as such do follow up onthe quality/diversity of tools used by the financial industry. Some of these tools used for the modelling of operational risk are discussed in the present paper. Secondly, whether for the demand on foreign reserves one adheres to the intervention model or an asset choice model (see Batten [10]), central banks face risk management decisions akin to commercial banks, albeit under different political and economical constraints. Finally, the role and function of central banks is no doubt under discussion (see Goodhart [45]), and therefore risk management issues which were less relevant some years ago may become important now. In particular, operational risk ought to be of great concern to any central bank. As discussed in Batten [10], a central bank typically confronts two types of economic phenomena – expected and unexpected – to which it makes policy responses. Faced with unanticipated economic but also external (e.g. catastrophic environmental) events to which it may wish to respond, it must hold additional reserves. Furthermore, as any institution, central banks face operational losses owing to system failure and fraud, for instance. How such losses impact on foreign reserve policy very much depends on the portfolio model chosen (Batten [10]).

Table 1: Types of operational risks (Crouhy et al. [19]).

1. **People risk:**	• Incompetence
	• Fraud
2. **Process risk:**	
A. Model risk	• Model/methodology error
	• Mark-to-model error
B. Transaction risk	• Execution error
	• Product complexity
	• Booking error
	• Settlement error
	• Documentation/contract risk
C. Operational control risk	• Exceeding limits
	• Security risks
	• Volume risk
3. **Technology risk:**	• System failure
	• Programming error
	• Information risk
	• Telecommunications failure

In Table 1, taken from Crouhy et al. [19], we have listed some typical types of operational risks. It is clear from this table that some risks are difficult to quantify (like incompetence under people risk), whereas others lend themselves much easier to quantification (as for instance execution error under transaction risk). As already alluded to above, most of the techniques discussed in this paper will have a bearing on the latter types of risk. In the terminology of Pezier [61], this corresponds to the ordinary operational risks. Clearly, the modelling of the latter type of risks is insufficient to base a full capital charge concept on.

The paper is organised as follows. In Section 2 we first look at some stylised facts of operational risk losses before formulating, in a mathematical form, the capital charge problem for operational risk (Pillar 1) in Section 3. In Section 4 we present a possible theory together with its limitations for analysing such losses, given that a sufficiently detailed loss database is available. We also discuss some of the mathematical research stemming from questions related to operational risk. Most of our discussion will use language close to Basel II and commercial banking. At several points, we shall highlight the relevance of the tools presented for risk management issues for central banks.

2 Data and preliminary stylised facts

Typically, operational risk losses are grouped in a number of categories (as in Table 1). In Pezier [61], these categories are further aggregated to the three levels of nominal, ordinary and exceptional operational risks. Within each category, losses are (or better, have to be) well-defined. Below we give an example of historical loss information for three different loss types. These losses correspond to transformed real data. As banks gather data, besides reporting current losses, an effort is made to build up databases going back about 10 years. The latter no doubt involves possible selection bias, a problem one will have to live with till more substantial data warehouses on operational risk become available. One possibility for the latter could be cross-bank pooling of loss data in order to find the main characteristics of the underlying loss distributions against which a particular bank's own loss experience can be calibrated. Such data pooling is well-known from non-life insurance or credit risk management. For Basel II, one needs to look very carefully into the economic desirability of such a pooling arrangement from a regulatory, risk management point of view. Whereas this would be most useful for the very rare large losses (exceptional losses), at the same time, such losses are often very specific to the institution and hence from that point of view make pooling somewhat questionable.

For obvious reasons, operational risk data are hard to come by. This is to some extent true for commercial banks, but considerably more so for central banks. One reason is no doubt the issue of confidentiality, another the relatively short period over which historical data have been consistently gathered. From the quantifiable real data we have seen in practice, we summarise below some of the stylised facts; these seem to be accepted throughout the industry for several operational risk categories. By way of example, in Figures 1, 2 and 3 we present loss information on three types of operational losses, which are for the purpose of this paper referred to as Types 1, 2 and 3. As stated above, these data correspond to modified real data. Figure 4 pools these losses across types. For these pooled losses, Figure 5 contains quarterly loss numbers. The stylised facts observed are:

- loss amounts very clearly show extremes, whereas
- loss occurrence times are definitely irregularly spaced in time, and also show (especially for Type 3, see also Figure 5) a tendency to increase over time. This non-stationarity can partly be explained by the already mentioned selection bias.

Figure 1: Operational risk losses, Type 1, $n = 162$

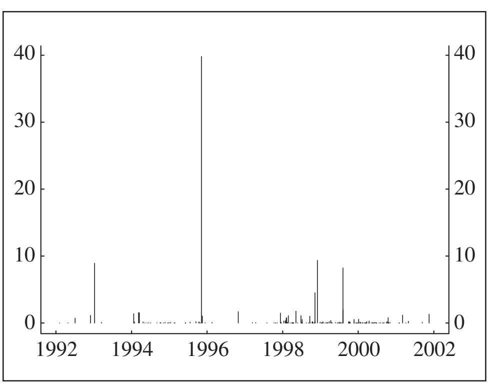

Figure 2: Operational risk losses, Type 2, $n = 80$

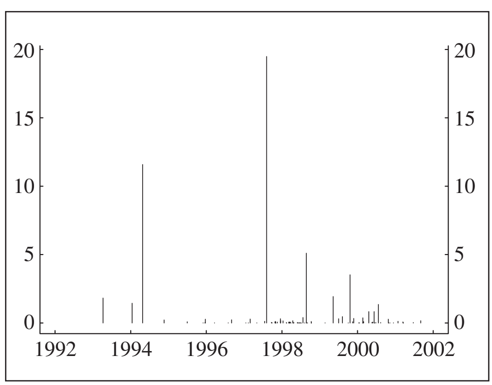

Figure 3: Operational risk losses, Type 3, $n = 175$

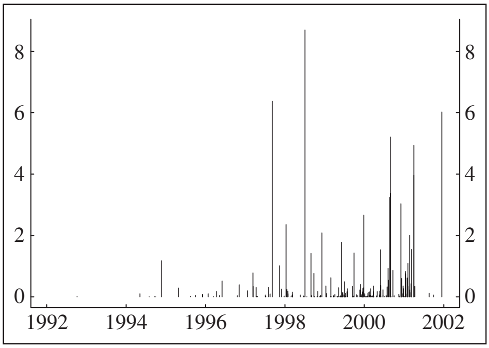

Figure 4: Pooled operational risk losses, $n = 417$

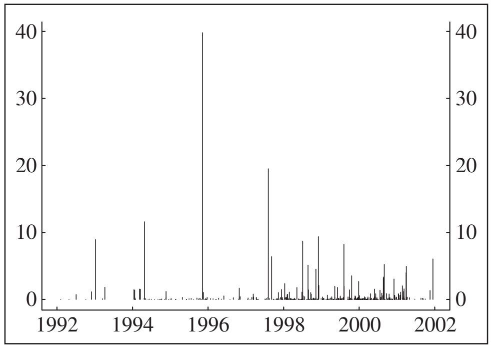

Figure 5: Quarterly loss numbers for the pooled operational risk losses

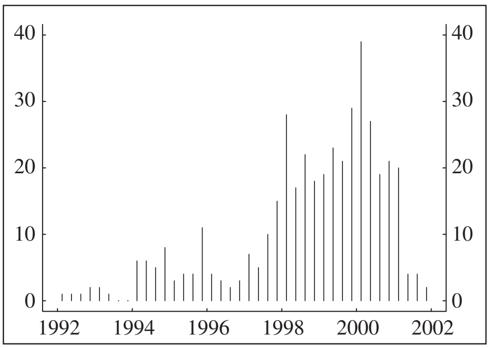

Figure 6: Fire insurance loss data, $n = 417$

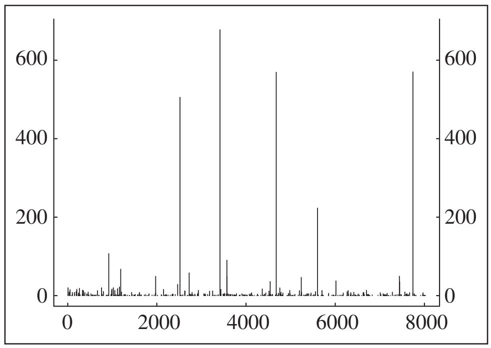

Any serious attempt of analytic modelling will at least have to take the above stylised facts into account. The analytic modelling referred to is not primarily aimed at calculating a risk-capital charge, but more at finding a sensible quantitative summary that goes beyond the purely descriptive. Similar approaches are well-known from the realm of reliability (see for instance Bedford and Cooke [11]), (non-life and re-)insurance (Hogg and Klugman [51]) and total quality control (as in Does et al. [26]).

In order to show some similarities with property insurance loss data, in Figure 6 we present $n = 417$ losses from a fire insurance loss database. For the full set of data, see Embrechts et al. [36], Figure 6.2.12.

Clearly, the large losses are a main concern, and hence, extreme value theory (EVT) can play a major role in analysing such data. Similar remarks have been made before concerning operational risk; see for instance Cruz [20] and Medova [55]. At this point, we would like to clarify a misconception which seems to persist in the literature; see for instance Pezier [61]. In no way will EVT be able to "predict" exceptional operational risk losses such as those present in the Barings case, for instance. Indeed, the introduction to Embrechts et al. [36] states very clearly that EVT is not a magical tool that can produce estimates out of thin air, but is instead one that tries to make the best use of whatever data exist on extreme phenomena. Moreover, and indeed equally important, EVT formulates very clearly under what conditions estimates on extreme events can be worked out. Especially with regard to exceptional losses

(Pezier [61]), there is very little that statistical theory, including EVT, can contribute. On the other hand, EVT is very useful when it comes to analysing extreme events such as catastrophic storms or floods, where these events occur within a specific physical or environmental model and where numerous observations on "normal" events exist; see Finkenstädt and Rootzén [40]. Clearly a case like Barings falls outside the range of EVTs applicability. On the other hand, when data on sufficient "normal" and a few extreme events within a well-defined class of risks exist, then EVT offers a very powerful statistical tool allowing one to extrapolate from the "normal" to the extreme. Numerous publications within financial risk management exemplify this; see for instance Embrechts [30]. Specific applications of EVT to risk management questions for central banks can be found in De Brandt and Hartman [25] and Hartman et al. [50]. Relevant problems where EVT technology could be applied are discussed in Borio et al. [14]. These papers mainly concern spillover of crises between financial markets, contagion, systemic risk and financial stability. An example of an EVT analysis related to the interest rate crisis of 2000 in Turkey involving interventions at the currency level (the lira) by the Turkish central bank is discussed in Gençay and Selçuk [44]. Embrechts [31] discusses the broader economic issues underlying the application of EVT to financial risk management.

In the next sections we concentrate on the calculation of an operational risk charge based on EVT and related actuarial techniques.

3 The problem

In order to investigate the kind of methodological problems one faces when trying to calculate a capital charge or reserve for (quantifiable) operational risks, we introduce some mathematical notation.

A typical operational risk database will consist of realisations of random variables

where $\left\{Y_k^{t,i} : t = 1,\dots,T, \quad i = 1,\dots,s \quad \text{and} \quad k = 1,\dots,N^{t,i}\right\}$

- T stands for the number of years ($T = 10$, say);
- s corresponds to the number of loss types (for instance $s = 6$), and
- $N^{t,i}$ is the (random) number of losses in year t of type i.

Note that in reality $Y_k^{t,i}$ is actually thinned from below, i.e.

$$Y_k^{t,i} = Y_k^{t,i} I_{\left\{Y_k^{t,i} \geq d^{t,i}\right\}}$$

where $d^{t,i}$ is some lower threshold below which losses are disregarded. Here $I_A(\omega) = 1$ whenever $\omega \in A$, and 0 otherwise. Hence, the total loss amount for year t becomes

$$L_t = \sum_{i=1}^{s} \sum_{k=1}^{N^{t,i}} Y_k^{t,i} = \sum_{i=1}^{s} L_{t,i}, \quad t = 1,\dots,T. \tag{1}$$

Denote by F_{L_t}, $F_{L_{t,i}}$ the distribution functions of L_t, $L_{t,i}$, $i = 1, \dots, s$. One of the capital charge measures discussed by the industry (Basel II) is the Value-at-Risk (VaR) at significance α (typically $0.001 \leq \alpha \leq 0.0025$ for operational risk losses) for next year's operational loss variable L_{T+1}. Hence

$$\text{OR-VaR}_{1-\alpha}^{T+1} = F_{L_{T+1}}^{\leftarrow}(1-\alpha),$$

where $F_{L_{T+1}}^{\leftarrow}$ denotes the (generalised) inverse of the distribution function $F_{L_{T+1}}$, also referred to as its quantile function. For a discussion of generalised inverses, see Embrechts et al. [36], p. 130. Figure 7 provides a graphical definition of this.

Figure 7: Calculation of operational risk VaR

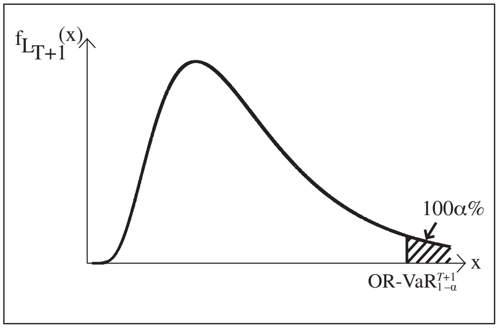

It is clear that, with any realistically available number T years worth of data, an in-sample estimation of VaR at this low significance level α is very difficult indeed. Moreover, at this aggregated loss level, across a wide range of operational risk types, any statistical theory (including EVT) will have difficulties in coming up with a scientifically sensible estimate. However, within quantitatively well-defined sub-categories, such as the examples in Figures 1-4, one could use EVT and come up with a model for the far tail of the loss distribution and calculate a possible out-of-sample tail fit. Based on these tail models, one could estimate VaR and risk measures that go beyond VaR, such as Conditional VaR (C-VaR)

$$\text{OR-C-VaR}_{1-\alpha,i}^{T+1} = E\left(L_{T+1,i} \mid L_{T+1,i} > \text{OR-VaR}_{1-\alpha,i}^{T+1}\right), \quad i = 1,\ldots,s,$$

or more sophisticated coherent risk measures; see Artzner et al. [2]. Furthermore, based on extreme value methodology, one could estimate a conditional loss distribution function for the operational risk category (categories) under investigation,

$$F_{T+1,u_i}^{i}(u_i + x) = P\left(L_{T+1,i} - u_i \le x \mid L_{T+1,i} > u_i\right), \quad x \ge 0, \quad i = 1,\ldots,s,$$

where u_i is typically a predetermined high loss level specific to loss category i. For instance one could take u_i = OR-VaR $_{1-\alpha,i}^{T+1}$. Section 4.1 provides more details on this.

We reiterate the need for extensive data modelling and pooling before risk measures of the above type can be calculated with a reasonable degree of accuracy. In the next section we offer some methodological building blocks which will be useful when more quantitative modelling of certain operational risk categories is demanded. The main benefit we see lies in a bank internal modelling, rather than providing a solution towards a capital charge calculation. As such, the methods we introduce have already been tested and made operational within a banking environment; see Ebnöther [28] and Ebnöther et al. [29].

4 Towards a theory

Since certain operational risk data are in many ways akin to insurance losses, it is clear that methods from the field of (non-life) insurance can play a fundamental role in their quantitative analysis. In this section we discuss some of these tools, also referred to as Insurance Analytics. For a discussion of the latter terminology, see Embrechts [32]. A further comparison with actuarial methodology can, for instance, be found in Duffy [27]. As mentioned in the Introduction, we have also made some references to EVT applications to specific risk management issues for central banks; the comments made below also apply to these.

4.1 Extreme value theory (EVT)

Going back to the fire insurance data (denoted $X_1,\dots,X_n$) in Figure 6, a standard EVT analysis is as follows:

– (EVT-1) Plot the empirical mean excess function

$$\hat{e}_n(u) = \frac{\sum_{k=1}^{n} (X_k - u)^+}{\sum_{k=1}^{n} I_{\{X_k > u\}}}$$

as a function of u and look for (almost) linear behaviour beyond some threshold value. For the fire insurance data, $\hat{e}_n(u)$ is plotted in Figure 8. A possible threshold choice is $u = 1$, i.e. for this case, a value low in the data.

– (EVT-2) Use the so-called Peaks over threshold (POT) method to fit an EVT model to the data above $u = 1$; plot the data (dots) and the fitted model (solid line) on log-log scale. Linearity supports Pareto-type power behaviour of the loss distribution $P(X_1 > x) = x^{-\alpha} h(x)$; see Figure 9. Here h is a so-called slowly varying function, i.e. for all $x > 0$, $\lim_{t\to\infty} \frac{h(tx)}{h(t)} = 1$. For $h \equiv c$, a constant, a log-log plot would be linear.

– (EVT-3) Estimate risk measures such as a 99% VaR – and 99% C-VaR – and calculate 95% confidence intervals around these risk measures; see Figure 9.

Figure 8: Empirical mean excess function for the fire loss data

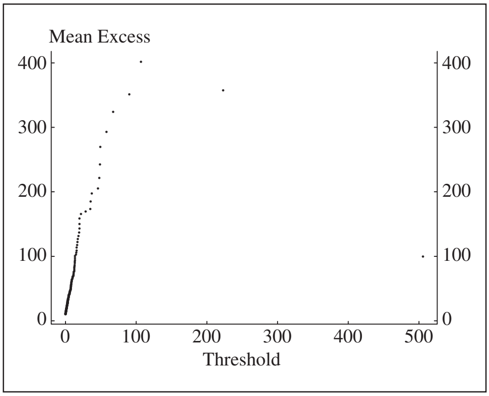

Figure 9: Empirical and fitted distribution tails on log–log scale, including estimates for VaR and C–VaR for the fire loss data.

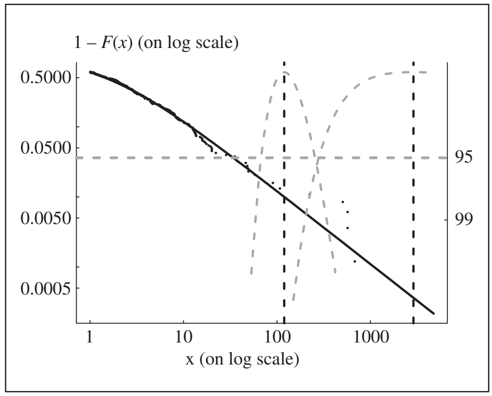

The estimates obtained are $\hat{\alpha} = 1.04$ with a corresponding 99% VaR value of 120 and an estimated 99% C-VaR of 2890 (note the huge difference). Figure 9 contains the so-called profile likelihood curves with maximal values in the estimated VaR and C-VaR. A 95% confidence interval around the 99% VaR 120 is given by (69, 255). The right vertical axis gives the confidence interval levels. The interval itself is obtained by cutting the profile likelihood curves at the 95% point. A similar construction (confidence interval) can be obtained for the C-VaR; owing to a value of $\hat{\alpha}$ (=1.04) close to 1, a very large 95% confidence interval is obtained which hence puts great uncertainty on the point estimates obtained. An α value less than 1 would correspond to an infinite mean model. A value between 1 and 2 yields an infinite variance, finite mean model. By providing these (very wide) confidence intervals in this case, EVT already warns the user that we are walking very close (or even too close) to the edge of the available data.

The software used, EVIS (Extreme Values in S-Plus) was developed by Alexander McNeil and can be downloaded from http://www.math.ethz.ch/~mcneil. It is no doubt a main strength of EVT that, under precise underlying model assumptions, confidence intervals for the risk measures under consideration can be worked out. The techniques used belong to the realm of maximum likelihood theory. We would however like to stress "under precise model assumptions". In Embrechts et al. [35] a simulation study by McNeil and Saladin [54] is reported which estimates, in the case of independent and identically distributed (iid) data, the number of observations needed in order to achieve a pre-assigned accuracy. For instance, in the iid case and a Pareto loss distribution with tail index 2 (a realistic assumption), in order to achieve a reasonable accuracy for the VaR at $\alpha = 0.001$, a minimum number of 100 observations above a 90% threshold u is needed (corresponding to an original 1,000 observations).[1]

[1] See also Embrechts et al. [36], pp. 194, 270 and 343 for the need to check conditions for the underlying data before an EVT analysis can be used. EVIS allows for several diagnostic checks on these conditions.

The basic result underlying the POT method is that the marked point process of excesses over a high threshold u, under fairly general (though very precise) conditions, can be well approximated by a compound Poisson process (see Figure 10):

$$\sum_{k=1}^{N(u)} Y_k \delta_{T_k}$$

where (Y_k) iid have a generalised Pareto distribution and $N(u)$ denotes the number of exceedances of u by (X_k). The exceedances of u form (in the limit) a homogeneous Poisson process and both are independent. See Leadbetter [52] for details. A consequence of the Poisson property is that inter-exceedance times of u are iid exponential. Hence such a model forms a good first guess. More advanced techniques can be introduced taking, for instance, non-stationarity and covariate modelling into account; see Embrechts [30], Chavez-Demoulin and Embrechts [16] and Coles [17] for a discussion of these techniques. The asymptotic independence between exceedance times and excesses makes likelihood fitting straightforward.

Figure 10: Stylised presentation of the POT method

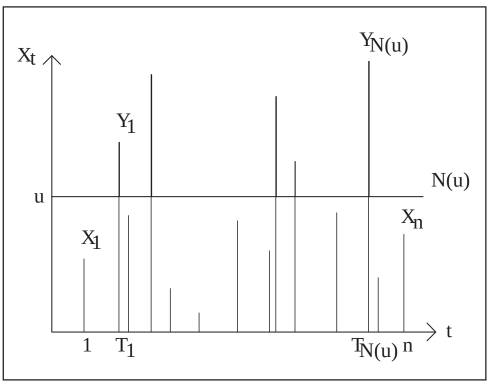

Turning to the mean excess plots for the operational risk data from Figures 1-3 (for the type-specific data) and Figure 4 (for the pooled data), we clearly see the typical increasing (nearly linear) trends indicating heavy-tailed, Pareto-type losses; see Figures 11-14 and compare them with Figure 8. As a first step, we can carry out the above extreme value analysis for the pooled data, though a refined analysis, taking non-stationarity into account, is no doubt necessary. Disregarding the possible non-stationarity of the data, one could be tempted to use the POT method and fit a generalised Pareto distribution to the pooled losses above $u = 0.4$, say. Estimates for the 99% VaR and the 99% C-VaR, including their 95% confidence intervals, are given in Figure 15. For the VaR we get a point estimate of 9.1, and a 95% confidence interval of (6.0,18.5). The 99% C-VaR beyond 9.1 is estimated as 25.9, and the lower limit for its 95% confidence interval is 11.7. Since, as in the fire insurance case, the tails are very heavy ($\hat{\alpha} = 1.63$), we get a very large estimate for the upper confidence limit for C-VaR.

Figure 11: Mean excess plot for operational risk losses, Type 1

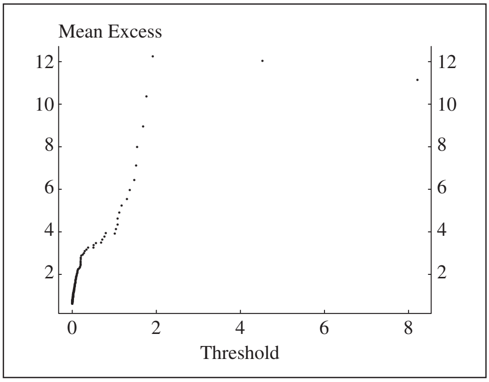

Figure 12: Mean excess plot for operational risk losses, Type 2

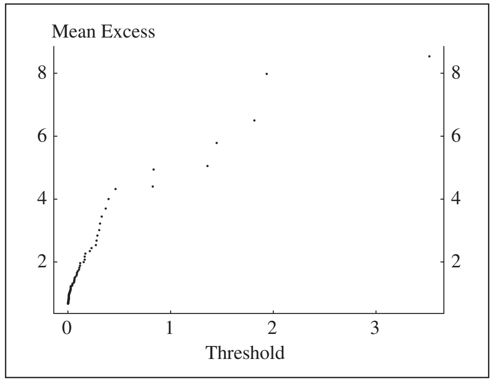

Figure 13: Mean excess plot for operational risk losses, Type 3

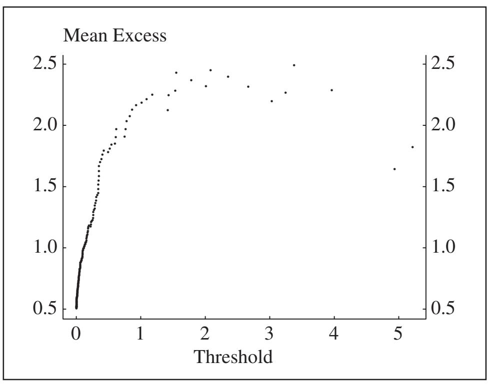

Figure 14: Mean excess plot for pooled operational risk losses

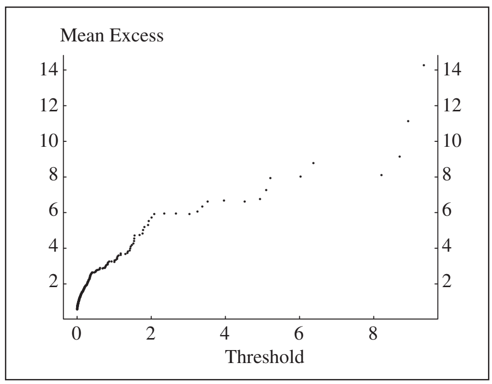

Figure 15: Empirical and fitted distribution tails for pooled operational losses on log-log scale, including estimates for VaR and C-VaR

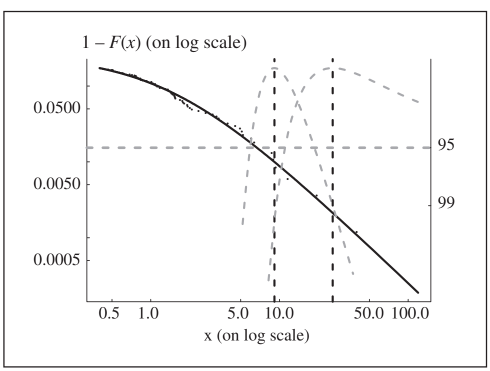

As already discussed before, the data in Figure 4 may contain a transition from more sparse data over the first half of the period under investigation to more frequent losses over the second half. It also seems that the early losses (in Figure 4 for instance) are not only more sparse, but also heavier. Again, this may be due to the way in which operational loss databases are built up. When gathering data for years some distance in the past, one only "remembers" the larger losses. Our EVT analysis should be adjusted for such a switch in size and/or intensity. Though Chavez-Demoulin and Embrechts [16] contains the relevant methodology, one should however realise that for such more advanced modelling, many more observations are needed.

In the next section, we make a more mathematical (actuarial) excursion in the realm of insurance risk theory. Risk theory provides models for random losses in a dynamic setting and yields techniques for the calculation of reserves given solvency constraints; see for instance Daykin et al. [24]. This theory has been applied in a variety of contexts and may also yield a relevant toolkit, especially in combination with EVT, if central banks were to manage their foreign reserves more actively and hence be more exposed to market, credit, and operational risk. In particular, ruin theory provides a longer-term view on solvency for a dynamic loss process. We discuss some of the key aspects of ruin theory with a specific application to operational risk below.

4.2. Ruin theory revisited

Given that (1) yields the total operational risk loss of s different sub-categories during a given year, it can be seen as resulting from a superposition of several (namely s) compound processes. So far, we are not aware of any studies which establish detailed features of individual processes or their interdependencies. Note that in Ebnöther [28] and Ebnöther et al. [29], conditions on the aggregated process are imposed: independence, or dependence through a common Poisson shock model. For the moment, we summarise (1) in a stylised way as follows:

$$L_t = \sum_{k=1}^{N(t)} Y_k ,$$

where $N(t)$ is the total number of losses over a time period $[0, t]$ across all s categories and the Y_k's are the individual losses. We drop the additional indices.

From an actuarial point of view, it would now be natural to consider an initial (risk) capital u and a premium rate $c > 0$ and define the cumulative risk process

$$C_t = u + ct - L_t , \quad t \geq 0. \tag{2}$$

In Figure 16 we have plotted such a risk process for the pooled operational risk losses shown in Figure 4. Again, the "regime switch" is clearly seen, splitting the time axis into roughly pre- and past-1998.

Figure 16: Risk process C_t with $u = 50$, $c = 28$ and the loss process from Figure 4

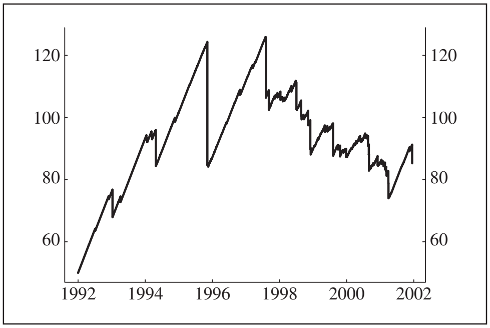

Given a small $\epsilon > 0$, for the process in (2), a risk capital u_ϵ can then be calculated putting the so-called ruin probability, i.e. the probability for the surplus process C_t to become negative, over a given time horizon $[\underline{T}, \overline{T}]$, equal to ϵ, small:

$$\Psi\left(u_\epsilon; \underline{T}, \overline{T}\right) = P\left(\inf_{\underline{T} \leq t \leq \overline{T}} (u_\epsilon + ct - L_t) < 0\right) = \epsilon\,. \tag{3}$$

For $c = 0$, $\underline{T} = T+$, $\overline{T} = T + 1$

$$u_\alpha = \text{OR–VaR}_{1-\alpha}^{T+1}.$$

The level of insolvency 0 is just chosen for mathematical convenience. One could, for instance, see c as a premium rate paid to an external insurer taking (part of) the operational risk losses or as a rate paid to (or accounted for by) a bank internal office. The rate c paid and the capital u_ϵ calculated would then be incorporated into the unit's overall risk capital.

Classical actuarial ruin theory concerns estimation of $\Psi(u; \underline{T}, \overline{T})$ in general and $\Psi(u, T) = \Psi(u; 0, T)$, $0 < T \leq \infty$ in particular, and for a wide range of processes. The standard assumption in the famous Cramér-Lundberg model is that $(N(t))$ is a homogeneous Poisson (λ) process, independent of the losses (Y_k) iid with distribution function G and mean $\mu < \infty$. Under the so-called net-profit condition (NPC), $c/\lambda > \mu$, one can show that, for "small claims" Y_k, a constant $R \in (0, \infty)$ (the so-called adjustment or Lundberg constant) and a constant $C \in (0, 1)$ exist, so that:

$$\Psi(u) = \Psi(u, \infty) < e^{-Ru}\,, \quad u \geq 0\,, \tag{4}$$

and

$$\lim_{u \to \infty} e^{Ru}\Psi(u) = C\,. \tag{5}$$

The small claims condition leading to the existence of $R > 0$ can be expressed in terms of $E(e^{RY_k})$ and typically holds for distribution functions with exponentially bounded tails. The constant C can be calculated explicitly; see for instance Grandell [47], Asmussen [3] and Rolski et al. [63] for details. For the case $\underline{T} = 0$, $\overline{T} = \infty$ and a process for which (4) holds, we can solve for u_ϵ in (3), obtaining

$$u_\epsilon = \frac{1}{R} \log \frac{1}{\epsilon} \,,$$

a quantity which can statistically be estimated given sufficient loss data.

For operational risk losses, the small claims condition underlying the so-called Cramér-Lundberg estimates (4) and (5) are typically not satisfied. Operational risk losses are heavy-tailed (power tail behaviour) as can be seen from Figures 11-14. Within the Cramér-Lundberg model, the infinite-horizon ($T = \infty$) ruin estimate for $\Psi(u) = \Psi(u, \infty)$ becomes (see Embrechts and Veraverbeke [39], Embrechts et al. [36]):

$$\Psi(u) \sim (\frac{c}{\lambda} - \mu)^{-1} \int_u^\infty (1 - G(x))\, dx \,, \quad u \to \infty \,. \tag{6}$$

Hence the ruin probability $\Psi(u)$ is determined by the tail of the loss distribution $1 - G(x)$ for x large, meaning that ruin (or a given limit excess) is caused by typically one (or a few) large claim(s). For a more detailed discussion on this "path leading to ruin", see Embrechts et al. [36], Section 8.3 and the references given there. The asymptotic estimate (6) holds under very general conditions of heavy tailedness, the simplest one being $1 - G(x) = x^{-\alpha} h(x)$ for h slowly varying and $\alpha > 1$. In this case (6) becomes

$$\Psi(u) \sim C\, u^{1-\alpha} h(u) \,, \quad u \to \infty \,, \tag{7}$$

where $C = [(\alpha - 1)\,(\frac{c}{\lambda} - \mu)]^{-1}$. Hence ruin decays polynomially (slow) as a function of the initial (risk) capital u. Also for lognormal claims, estimate (6) holds. In the actuarial literature, the former result was first proved by von Bahr [7], the latter by Thorin and Wikstad [65]. The final version for so-called subexponential claims is due to Embrechts and Veraverbeke [39]. In contrast to the small claims regime estimates (4) and (5), the heavy-tailed claims case (6) seems to be robust with respect to the underlying assumptions of the claims process. The numerical properties of (6) are however far less satisfactory.

Besides the classical Cramér-Lundberg model, an estimate similar to (6) also holds for the following processes:

– Replace the homogeneous Poisson process $(N(t))$ by a general renewal process; see Embrechts and Veraverbeke [39]. Here the claim inter-arrival times are still independent, but have a general distribution function, not necessarily an exponential one.
– Generalisations to risk processes with dependent inter-claim times, allowing for possible dependence between the arrival process and the claim sizes are discussed in Asmussen [3], Section IX.4. The generalisations contain the so-called Markov-modulated models as a special case; see also Asmussen et al. [6]. In these models, the underlying intensity model follows a finite state Markov chain, enabling for instance the modelling of underlying changes in the economy in general or the market in particular.
– Ruin estimates for risk processes perturbed by a diffusion, or by more general stochastic processes, are for instance to be found in Furrer [41], Schmidli [64] and Veraverbeke [66].
– A very general result of type (7) for the distribution of the ultimate supremum of a random walk with a negative drift is derived in Mikosch and Samorodnitsky [56]. Mathematically, these results are equivalent with ruin estimation for a related risk model.

For all of these models an estimate of type (7) holds. Invariably, the derivation is based on the so-called “one large claim heuristics”; see Asmussen [3], p. 264. These heuristics may eventually play an important role in the analysis of operational risk data.

As discussed above, as yet there is no clear stochastic model available for the general operational risk process (1). Consequently, it would be useful to find a way to obtain a broad class of risk processes for which (7) holds. A solution to this problem is presented in Embrechts and Samorodnitsky [38] through a combination of the “one large claim heuristics” and the notion of operational time (time-change). Below we restrict our attention to the infinite horizon case $\Psi(u)$. First of all, estimate (7) is not fine enough for accurate numerical approximations. It rather gives a benchmark estimate of ruin (insolvency), deliminating the heavy-tailed (“one claim causes ruin”) situation from the light-tailed estimates in (4) and (5) where most (small) claims contribute equally and ruin is remote, i.e. has an exponentially small probability. For a discussion on numerical ruin estimates of type (7), see Asmussen and Binswanger [4], and Asmussen et al. [5]; the keywords here are rare event simulation.

Suppose that we are able to estimate ruin over an infinite horizon for a general stochastic (loss) process (L_t), a special case of which is the classical Cramér-Lundberg total claim process in (2) or the risk processes listed above. Suppose now that, for this general loss process (L_t), we have a ruin estimate of form (7). From (L_t), more general risk processes can be constructed using the concept of time-change $(\Delta(t))$. The latter is a positive, increasing stochastic process that typically (but not exclusively) models economic or market activity. The more general process $(L_{\Delta(t)})$ is the one we are really interested in, since its added flexibility could allow us to model the stylised facts of operational risk data as discussed in Section 2. The step from the classical Cramér-Lundberg model (L_t) to the more general process $(L_{\Delta(t)})$ is akin to the step from the classical Black-Scholes-Merton model to more general stochastic volatility models. We can now look at this general time-changed process and define its corresponding infinite horizon ruin function:

$$\Psi_\Delta(u) = P\left(\sup_{t\geq 0}(L_{\Delta(t)} - ct) > u\right).$$

We then ask for conditions on the process parameters involved, as well as for conditions on $(\Delta(t))$, under which

$$\lim_{t\to\infty} \frac{\Psi(t)}{\Psi_\Delta(t)} = 1\,, \tag{8}$$

meaning that, asymptotically, ruin is of the same order of magnitude in the time-changed (more realistic) process as it is for the original (more stylised) process. These results can be interpreted as a kind of robustness characterisation for general risk processes so that the polynomial ruin probability estimate (7) holds. In Embrechts and Samorodnitsky [38], besides general results for (8) to hold, specific examples are discussed. Motivated by the example of transaction risk (see Table 1), Section 3 in the latter paper discusses the case of mixing through Markov chain switching models, also referred to as Markov-modulated or Markov renewal processes. In the context of operational risk, it is natural to consider a class of time-change processes $(\Delta(t))$ in which time runs at a different rate in different time intervals, depending on the state of a certain underlying Markov chain. The Markov chain stays a random amount of time in each state, with a distribution that depends on that state. Going back to the transaction risk case, one can think of the Markov chain states as resulting from an underlying market volume (intensity) index. These changes in volumes traded may for instance have an effect on back office errors. The results obtained in Embrechts and

Samorodnitsky [38] may be useful to characterise interesting classes of loss processes where ruin behaves as in (7). Recall from Figure 5 the fact that certain operational risk losses show periods of high (and low) intensity. Future dynamic models for sub-categories of operational risk losses will have to take these characteristics into account.The discussion above mainly aims to show that tools for such problems are at hand and await the availability of more detailed loss databases.

Some remarks should be made at this point. Within classical insurance risk theory, this result has been proven in full generality in Embrechts and Veraverbeke [39]. Several generalisations, including the case of Markov modulated processes, are discussed in Asmussen [3]. Within classical insurance risk theory, a full solution linking heavy-tailedness of the claim distribution to the long-tailedness of the corresponding ruin probability is discussed in Asmussen [3]. Alternative models leading to similar distributional conclusions can be found in the analysis of teletraffic data; see for instance Resnick and Samorodnitsky [62] and Finkenstädt and Rootzén [40]. Whereas the basic operational risk model in (1) may be of a more general nature than the ones discussed above, support seems to exist for the supposition that under fairly general conditions, the tail behaviour of $P(L_{T+1} > x)$ will be power-like. Furthermore, the notion of time-change may seem somewhat artificial. This technique has however been around in insurance mathematics for a long time and is used to transform a complicated loss process into a more standard one; see for instance Cramér [18] or Bühlmann [15]. Within finance, these techniques were introduced through the fundamental work of Olsen and Associates on Θ-time; see Dacorogna et al. [21]. Further work has been done by Ané and Geman [1], Geman et al. [42, 43] and more recently Barndorff-Nielsen and Shephard [8]; they use time-change techniques to transform a financial time series with randomness in the volatility to a standard Black-Scholes-Merton model. As stated above, the situation is somewhat akin to the relationship between a Brownian motion-based model (such as the Black-Scholes-Merton model) and more recent models based on general semi-martingales. It is a well-known result, see Monroe [57], that any semi-martingale can be written as a time-changed Brownian motion.

4.3 Further tools

In the previous section, we briefly discussed some (heavy-tailed) ruin type estimates which, in view of the data already available on operational risk, may become useful. From the realm of insurance, several further techniques may be used. Below we mention some of them without entering into details. Recall from (1) that a yearly operational risk variable will typically be of the form:

$$L = \sum_{k=1}^{N} Y_k \tag{9}$$

where N is some discrete random variable counting the total number of claims within a given period across all s loss classes, say, and Y_k denotes the kth claim. Insurance mathematics has numerous models of type (9) starting with the case where N is a random variable independent of the iid claims (Y_k) with distribution function G, say. In this case, the probability of a loss exceeding a certain level equals

$$P(L > x) = \sum_{k=1}^{\infty} P(N = k) \left(1 - G^{*k}(x)\right), \tag{10}$$

where G^{*k} denotes the kth convolution of G. Again, in the case that $1 - G(x) = x^{-\alpha}h(x)$, and the moment generating function of N is analytic in 1, it is shown in Embrechts et al. [36] that

$$P(L > x) \sim E(N)x^{-\alpha}h(x)\,, \quad x \to \infty\,.$$

Several procedures exist for numerically calculating (10) under a wide range of conditions. These include recursive methods such as the Panjer-Euler method for claim number distributions satisfying $P(N = k) = (a + \frac{b}{k})P(N = k - 1)$ for $k = 1, 2, \ldots$ (see Panjer [58]), and Fast Fourier Transform methods (see Bertram [12]). Grübel and Hermesmeier [48, 49] are excellent review papers containing further references. The actuarial literature contains numerous publications on the subject; good places to start are Panjer and Willmot [59] and Hogg and Klugman [51].

Finally, looking at (1), several aggregation operations are taking place, including the superposition of the different loss frequency processes $(N^{t,i})_{i=1,\ldots,s}$ and the aggregation of the different loss size variables $(Y_k^{t,i})_{k=1,\ldots,N^{t,i},\, i=1,\ldots,s}$. For the former, techniques from the theory of point processes are available; see for instance Daley and Vere-Jones [22]. The issue of dependence modelling within and across operational risk loss types will no doubt play a crucial role; copula techniques, as introduced in risk management in Embrechts et al. [37], can no doubt be used here.

5 Final comments

As stated in the Introduction, tools from the realm of insurance as discussed in this paper may well become relevant, conditional on the further development and implementation of quantitative operational risk measurement within the financial industry. Our paper aims at encouraging a better exchange of ideas between actuaries and risk managers. Even if one assumes full replicability of operational risk losses within the several operational risk sub-categories, their interdependence will make detailed modelling difficult. The theory presented in this paper is based on specific conditions and can be applied in cases where testing has shown that these underlying assumptions are indeed fulfilled. The ongoing discussions around Basel II will show at which level the tools presented will become useful. However, we strongly doubt that a full operational risk capital charge can be based solely on statistical modelling.

Some of the tools introduced and mainly exemplified through their application to the quantitative modelling of operational risk are no doubt useful well beyond in more general risk management. Further research will have to look more carefully at the risk management issues underlying the central bank landscape. In particular, the issue of market liquidity under extreme events, together with the more active management of foreign reserves and the longer time view will necessitate tools that complement the existing ones in risk management for commercial banks. Insurance analytics as presented in this paper, and discussed more in detail in Embrechts [33], will no doubt form part of this.

References

Ané, T. and Geman, H. [1] (2000). Order Flow, Transaction Clock and Normality of Asset Returns. *Journal of Finance*, 55, pp. 2259-84.

Artzner, P., Delbaen, F., Eber, J. M., and Heath, D. [2] (1999). Coherent Measures of Risk, *Mathematical Finance*, 9, pp. 203-28.

Asmussen, S. [3] (2000). *Ruin Probabilities*. World Scientific.

Asmussen, S. and Binswanger, K. [4] (1997). Simulation of Ruin Probabilities for Subexponential Claims. *Astin Bulletin*, 27(2), pp. 297-318.

Asmussen, S., Binswanger, K., and Højgaard, B. [5] (2000). Rare Events Simulation for Heavy-tailed Distributions. *Bernoulli*, 6(2), pp. 303-22.

Asmussen, S., Henriksen, L. F., and Klüppelberg, C. [6] (1994). Large Claims Approximations for Risk Processes in a Markovian Environment. *Stochastic Processes and their Applications*, 54, pp. 29-43.

von Bahr, B. [7] (1975). Asymptotic Ruin Probabilities when Exponential Moments Do Not Exist. *Scandinavian Actuarial Journal*, pp. 6-10.

Barndorff-Nielsen, O. E. and Shephard, N. [8] (2002). Financial Volatility and Lévy-based Models. Book project.

Basel Committee on Banking Supervision [9] (2001). Consultative Document. Overview of the New Basel Capital Accord.

Batten, D. S. [10] (1982). *Central Banks' Demand for Foreign Reserves under Fixed and Floating Exchange Rates*. Federal Reserve Bank of St. Louis.

Bedford, T. and Cooke, R. M. [11] (2001). *Probabilistic Risk Analysis: Foundations and Methods*. Cambridge University Press.

Bertram, J. [12] (1981). Numerische Berechnung von Gesamtschadenverteilungen. *Blätter der DGVM*, 15, pp. 175-94.

Blum, P., Dias, A., and Embrechts, P. [13] (2002). The ART of Dependence Modelling: The Latest Advances in Correlation Analysis. In: *Alternative Risk Strategies* (ed. M. Lane). Risk Waters Group, London, pp. 339-56.

Borio, C., Furfine, C., and Lowe, P. [14] (2001) Procyclicality of the Financial System and Financial Stability; Issues and Policy Options. BIS Papers No. 1.

Bühlmann, H. [15] (1970). *Mathematical Methods in Risk Theory*. Springer.

Chavez-Demoulin, V. and Embrechts, P. [16] (2004). Smooth Extremal Models in Finance and Insurance. *Journal of Risk and Insurance*, to appear.

Coles, S. [17] (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer.

Cramér, H. [18] (1930). *On the Mathematical Theory of Risk*. Skandia Jubilee Volume, Stockholm.

Crouhy, C., Galai, D., and Mark, R. [19] (2000). *Risk Management*. McGraw-Hill, New York.

Cruz, M. G. [20] (2002). *Modelling, Measuring and Hedging Operational Risk.* Wiley.

Dacorogna, M. M., Gençay, R., Müller, U. A., Olsen, R. B., and Pictet, O. V. [21] (2001). *An Introduction to High-frequency Finance*. Academic Press.

Daley, D. J. and Vere-Jones, D. [22] (1988). *An Introduction to the Theory of Point Processes*. Springer.

Daníelsson, J., Embrechts, P., Goodhart, C., Keating, C., Muennich, F., Renault, O., and Shin, H. S. [23] (2001). An Academic Response to Basel II. Financial Markets Group, London School of Economics, Special Paper No. 130.

Daykin, C. D., Pentikäinen, T., and Pesonen, M. [24] (1994). *Practical Risk Theory for Actuaries*. Chapman and Hall/CRC, London.

De Brandt, O. and Hartman, P. [25] (2000). Systemic Risk: A Survey. Working Paper No. 35, European Central Bank, Frankfurt.

Does, R. J. M. M., Roes, K. C. B., and Trip, A. [26] (1999). *Statistical Process Control in Industry*. Kluwer Academic.

Duffy, P. [27] (2002). Operational Risk. *The Actuary*, October 2002, pp. 24-25.

Ebnöther, S. [28] (2001). Quantitative Aspects in Operational Risk. Diploma thesis ETH Zürich.

Ebnöther, S., Vanini, P., McNeil, A. J., and Antolinez-Fehr, P. [29] (2003). Operational Risk: A Practitioner's View. *Journal of Risk*, 5(3), pp. 1-15.

Embrechts, P. [30] (ed.) (2000). *Extremes and Integrated Risk Management*. Risk Waters Group, London.

Embrechts, P. [31] (2003). Extremes in Economics and the Economics of Extremes. In: *Extreme Values in Finance, Telecommunications, and the Environment* (eds. B. Finkenstädt, H. Rootzén). Chapman and Hall/CRC, London, pp. 169-83.

Embrechts, P. [32] (2003). Insurance Analytics. Guest Editorial, *British Actuarial Journal*, 8(IV), pp. 639-41.

Embrechts, P. [33] (2003). Insurance Analytics: Actuarial Tools for Financial Risk Management. Opening Lecture at the XXXIV Astin Colloquium in Berlin.

Embrechts, P., Frey, R., and McNeil, A. J. [34] (2004). *Quantitative Risk Management: Concepts, Techniques and Tools*. Book manuscript, ETH Zürich.

Embrechts, P., Furrer, H., and Kaufmann, R. [35] (2003). Quantifying Regulatory Capital for Operational Risk. *Derivatives Use, Trading and Regulation*, 9(3), pp. 217-33.

Embrechts, P., Klüppelberg, C., and Mikosch, T. [36] (1997). *Modelling Extremal Events for Insurance and Finance*. Springer.

Embrechts, P., McNeil, A. J., and Straumann, D. [37] (2002). Correlation and Dependence in Risk Management: Properties and Pitfalls. In: *Risk Management: Value at Risk and Beyond* (ed. M. Dempster). Cambridge University Press, pp. 176-223.

Embrechts, P. and Samorodnitsky, G. [38] (2003). Ruin Problem and How Fast Stochastic Processes Mix. *The Annals of Applied Probability*, 13, pp. 1-36.

Embrechts, P. and Veraverbeke, N. [39] (1982). Estimates for the Probability of Ruin with Special Emphasis on the Possibility of Large Claims. *Insurance: Mathematics and Economics*, 1, pp. 55-72.

Finkenstädt, B. and Rootzén, H. [40] (eds.) (2003). *Extreme Values in Finance, Telecommunications, and the Environment*. Chapman and Hall/CRC, London.

Furrer, H. [41] (1998). Risk Processes Perturbed by α-stable Lévy Motion. *Scandinavian Actuarial Journal*, pp. 59-74.

Geman, H., Madan, D. B., and Yor, M. [42] (2001). Time Changes for Lévy Processes. *Mathematical Finance*, 11(1), pp. 79-96.

Geman, H., Madan, D. B., and Yor, M. [43] (2002). Stochastic Volatility, Jumps and Hidden Time Changes. *Finance and Stochastics*, 6(1), pp. 63-90.

Gençay, R. and Selçuk, F. [44] (2001). Overnight Borrowing, Interest Rates and Extreme Value Theory. Working Paper, Bilkent University.

Goodhart, C. [45] (1988). *The Evolution of Central Banks*. MIT Press.

Goodhart, C. [46] (2001). Operational Risk. Special Paper 131, Financial Markets Group, London School of Economics.

Grandell, J. [47] (1991). *Aspects of Risk Theory.* Springer.

Grübel, R. and Hermesmeier, R. [48] (1999). Computation of Compound Distributions I. Aliasing Errors and Exponential Tilting. *Astin Bulletin*, 29, pp. 197-214.

Grübel, R. and Hermesmeier, R. [49] (2000). Computation of Compound Distributions II. Discretization Errors and Richardson Extrapolation. *Astin Bulletin*, 30, pp. 309-31.

Hartman, P., Straetmans, S., and de Vries, C. G. [50] (2001). Asset Market Linkages in Crisis Periods. Working Paper No. 71, European Central Bank, Frankfurt.

Hogg, R. V. and Klugman, S. A. [51] (1984). *Loss Distributions.* Wiley.

Leadbetter, M. R. [52] (1991). On a Basis for 'Peaks over Threshold' Modeling. *Statistics and Probability Letters*, 12, pp. 357-62.

McNeil, A. J. [53] (1997). Estimating the Tails of Loss Severity Distributions Using Extreme Value Theory. *Astin Bulletin*, 27, pp. 117-37.

McNeil, A. J. and Saladin, T. [54] (1997). The Peaks over Thresholds Method for Estimating High Quantiles of Loss Distributions. *Proceedings of XXVIIth International Astin Colloquium*, Cairns, Australia, pp. 23-43.

Medova, E. [55] (2000). Measuring Risk by Extreme Values. *Risk*, November 2000, pp. 20-26.

Mikosch, T. and Samorodnitsky, G. [56] (2000). The Supremum of a Negative Drift Random Walk with Dependent Heavy-tailed Steps. *The Annals of Applied Probability*, 10, pp. 1025-64.

Monroe, I. [57] (1978). Processes that Can Be Embedded in Brownian Motion. *Annals of Probability*, 6(1), pp. 42-56.

Panjer, H. [58](1981). Recursive Evaluation of a Family of Compound Distributions. *Astin Bulletin*, 12, pp. 22-26.

Panjer, H. and Willmot, G. [59] (1992). *Insurance Risk Models*. Society of Actuaries, Schaumburg, Illinois.

Pezier, J. [60] (2002). A Constructive Review of Basel's Proposals on Operational Risk. Working Paper, ISMA Centre, University of Reading.

Pezier, J. [61] (2002). Operational Risk Management. Working Paper, ISMA Centre, University of Reading.

Resnick, S. and Samorodnitsky, G. [62] (2000). A Heavy Traffic Limit Theorem for Workload Processes with Heavy-tailed Service Requirements. *Management Science*, 46, pp. 1236-48.

Rolski, R, Schmidli, H., Teugels, J., and Schmidt, V. [63] (1999). *Stochastic Processes for Insurance and Finance*. Wiley.

Schmidli, H. [64] (1999). Perturbed Risk Processes: A Review. *Theory of Stochastic Processes*, 5, pp. 145-65.

Thorin, O. and Wikstad, N. [65] (1977). Calculation of Ruin Probabilities when the Claim Distribution Is Lognormal. *Astin Bulletin*, 9, pp. 231-46.

Veraverbeke, N. [66] (1993). Asymptotic Estimates for the Probability of Ruin in a Poisson Model with Diffusion. *Insurance: Mathematics and Economics*, 13, pp. 57-62.

3 CASE STUDIES

15

Risk management practices at the ECB

Ciarán Rogers

Abstract

This paper documents the approach to risk management at the ECB. It focuses mainly on the management of the foreign reserves and the attendant risks for the ECB. A simple balance sheet is presented, and it is demonstrated that the bulk of the ECB's financial risk comes from its holdings of foreign currency assets. The role of the risk management function at the ECB is elaborated upon, and its main responsibilities are summarised. The chosen organisational structure for the risk management area is also discussed. The paper concludes with an assessment of the effectiveness of the current set-up, highlighting the main challenges ahead.

1 Introduction

Since 1 January 1999, the European Central Bank (ECB) has taken its place on the central banking stage to form a triumvirate of key global central banks alongside the Federal Reserve and the Bank of Japan. The ECB is unique as it is a supranational central bank at the heart of the Eurosystem. The Eurosystem comprises the ECB and the 12 national central banks (NCBs) of the sovereign states that have agreed to transfer their monetary sovereignty to the European Community level. The bank is located in Frankfurt am Main, Germany, and was established together with the European System of Central Banks (ESCB) on 1 June 1998 in the run-up to the launch of the euro. The ECB has its own legal personality and acts as the hub within both the ESCB (comprising the 15 NCBs of the EU Member States) and the Eurosystem (comprising the ECB and the 12 NCBs of the EU Member States which have adopted the euro). The basic tasks of the Eurosystem are the definition and implementation of the monetary policy of the euro area, the conduct of foreign exchange operations, the holding and management of the official foreign reserves of the Member States, and the promotion of the smooth operation of payments systems. The Eurosystem is governed by the decision-making bodies of the ECB, namely the Governing Council and the Executive Board.

This paper focuses mainly on foreign reserves, but it also aims at giving an overview of all risk management activities at the ECB. It shows the broad outlines of the ECB balance sheet and the sensitivity of the bank to currency and interest rate movements. It also shows how the ECB chooses to manage these risks, and looks at the area in the bank charged with the monitoring, reporting and managing of risk. Later, the organisational set-up of the risk management area is discussed. The paper concludes with an assessment of the risk management framework at the ECB, pointing to the challenges ahead.

2 The ECB's financial position

It has been said that a central bank is an unconventional bank, but a bank nonetheless. The ECB's balance sheet reflects this. Consider the following simplistic central bank balance sheet:

Stylised central bank balance sheet

Assets	Liabilities
Monetary policy operations	Banknotes
Other financial assets	Capital

- *Monetary policy operations* typically contain collateralised lending to banks. This is remunerated at the operational target interest rate of the central bank.
- *Other financial assets* contain foreign exchange reserves including gold as well as possibly domestic financial assets not related to monetary policy. The return on these assets depends inter alia on the local yield curve developments and exchange rate movements.
- *Banknotes* are exogenously determined and normally follow an increasing trend over time.
- The *capital* item may, in addition to capital, also include provisions and any revaluation accounts.

Using these four categories to assess the ECB balance sheet, two main remarks can be made. First, the two most important items on the ECB balance sheet are banknotes and foreign currency reserves and their associated asset and liability counterparts. Together, they account for over 85% of the total balance sheet. The ECB itself does not print banknotes, relying instead on the NCBs. Nevertheless, it has been allocated a share of 8% of the total value of euro banknotes in circulation. This share represents a liability on the balance sheet. This share of the total euro banknote issue is in turn backed by claims on the NCBs. These claims, which bear interest, are on the asset side of the balance sheet. This combination of an interest-bearing asset with a non-interesting bearing liability creates seigniorage income for the ECB. This is the largest source of income for the ECB and carries minimal risk.

The other major balance sheet item is the foreign reserves. They derive from the transfer by the NCBs of the Eurosystem of foreign exchange assets (including gold) worth E 40 billion to the ECB as pooled reserves at the start of 1999.[1] The counterparty of these assets is a liability equivalent to the original euro value of the foreign exchange assets transferred by the NCBs, which is denominated in euro. While the return on the reserve assets is a function of exchange rate developments and yield curve movements in the currencies invested in, the associated liability is remunerated at the operational target interest rate of the ECB (85% of the main refinancing rate) regardless of the return on the foreign currency assets. This currency and interest rate mismatch creates a significant source of possible risk for the ECB.

The second important remark is that, unlike a standard central bank, the ECB for the time being does not directly execute monetary policy operations where central bank credit and collateral are exchanged. At present, these tasks are performed locally by the NCBs. As a direct result, the ECB does not hold assets consisting of lending to financial sector counterparties, and does it receive collateral in turn. The overall impact of this on the ECB

[1] "Pooled reserves" is the term used to describe the foreign currency reserves that are directly owned by the ECB and managed on an agency basis by the NCBs. Each NCB also continues to hold its own national foreign currency reserves separately from the pooled reserves and over which it has control, subject to ECB guidelines.

balance sheet is that the foreign currency reserves account for a larger percentage of assets than is the case in many central banks. In terms of the overall balance sheet, this increases exposure to the foreign assets held on the ECB's balance sheet. This contrasts with many other central bank balance sheets, and brings to the fore the need to manage risk.

In effect, of the four balance sheet items above, the bulk of the risk that the ECB faces comes from other financial assets. The remaining three items are exogenously determined or arise from the ECB's public mission, which is the formulation and implementation of monetary policy. In addition, most of these balance sheet items are sensitive to the ECB's main refinancing rate. The item that introduces most financial risk into the ECB's balance sheet is the foreign reserves. This is also arguably the area where the ECB has the most freedom to manage risk. To some degree, the ECB can decide which currencies and in which proportions to hold foreign reserve assets. Moreover, it can also decide how to manage the interest rate risk arising from investing in foreign assets. Unless the ECB is called upon to intervene, the management of the reserve assets is similar to the task faced by an asset manager in a large financial institution. The task of risk managers at the ECB is to inform decision-makers of the risks currently being taken by the institution and to arrive at situations where financial decisions are taken that are prudent and consistent with the risk appetite of the institution.

3 The role of risk management at the ECB

Within the Eurosystem, the ECB plays the role of decision-maker and coordinator. A clear example of this role is the risk management function. The Risk Management Division manages and monitors financial risks incurred by the ECB directly or indirectly by the 12 NCBs of the Eurosystem on behalf of the ECB.

The uniqueness of the Eurosystem demanded that particular consideration be given to an appropriate risk management structure. After some deliberation by the predecessor of the ECB, the European Monetary Institute (EMI), and in consultation with the NCBs, a decision was made to establish a centralised ECB risk management function with responsibility for the financial risk of all operations conducted by the ECB or the 12 NCBs of the Eurosystem on behalf of the ECB. The requirement to manage and monitor this risk centrally was of paramount importance in the organisational structure of the ECB.

The decision to manage risk centrally was also significant. The principle of decentralisation informs many of the organisational decisions at the ECB and in the Eurosystem. To the extent deemed possible and appropriate, the ECB shall have recourse to the NCBs to carry out operations which form part of the tasks of the Eurosystem. Risk management was judged not to be compatible with decentralisation. Risk management centralisation has the advantage of avoiding the inefficiencies that could arise from conflict and dispersion of responsibilities.

The charter of the ECB's Risk Management Division states that it shall:

- provide organisational support and propose policies and procedures on risk management for the range of operations conducted by the ECB or the 12 NCBs of the Eurosystem on behalf of the ECB.
- manage market risk and credit risk in relation to the ECB's investment operations, credit risk in relation to the monetary policy and foreign exchange operations of the ECB, and operational risk for all operations conducted by the ECB or the 12 NCBs of the Eurosystem on behalf of the ECB.

3.1 Responsibilities of the ECB's risk management division

The Risk Management Division contains two principal operational fields: risk management issues related to the implementation of Eurosystem monetary policy operations, and risk management issues related to the management of the foreign currency reserves and the ECB's own funds.

Risk management issues related to the implementation of Eurosystem monetary policy operations

In order to control the credit risk in its monetary policy operations (as well as in the intraday credit extended for payment systems purposes), the Eurosystem relies on high-quality collateral provided by its counterparties. Article 18 of the Statute of the European System of Central Banks and of the European Central Bank (the "Statute of the ESCB") requires that any credit provided by the Eurosystem shall be covered by adequate collateral. To this end, the Eurosystem has adopted an elaborate collateral policy, which is a key feature of its operational framework.

This policy seeks both to balance and reflect the very different practices applied by the NCBs prior to the establishment of monetary union, as well as the prevailing differences in the capital markets and in the legal and institutional structures of the EU Member States.

From today's perspective, it may seem self-evident that a central bank should require collateral to control the credit risk in its liquidity-providing operations. However, it is worthwhile recalling that at the time the Statute of the ESCB was drafted, it was not unusual for EU central banks to grant uncollateralised credits to their counterparties. Moreover, most EU central banks traditionally accepted only a very narrow range of collateral, mainly bonds issued by the respective national governments. As a consequence, there was only limited experience in most EU central banks of assessing the risk of individual assets. In particular, there was virtually no experience of assessing the rather complex risks related to cross-border transactions.

From the outset, the Eurosystem opted for a policy of accepting a broad range of high-quality assets as collateral. This was seen as necessary in order to ensure sufficient availability of collateral. A broad range of collateral was also seen as conducive to equal treatment of counterparties, promoting the development of markets for private paper, and as a means of ensuring that the NCBs could continue with established practices. It appears that the Eurosystem accepts a wider range of asset types as collateral in its monetary policy operations than any other central bank in the world. This fact in itself poses a major challenge for the ECB's Risk Management Division.

The Eurosystem distinguishes between two categories of assets eligible for its monetary policy operations, which it categorises as tier one and tier two. However, extensive work has taken place on this issue and it is likely that in the future a proposal to move towards a single list of collateral will emerge.

Tier one consists of marketable debt securities that fulfil certain eligibility criteria established by the ECB. The criteria endeavour to ensure that tier one assets consist of high-quality marketable securities with high liquidity that can easily be traded on a cross-border basis.

Tier two consists of additional assets (currently including, inter alia, non-marketable securities and equities) which are of particular importance to the national financial markets or banking systems. The precise criteria for tier two assets are proposed by each NCB according to certain minimum standards established by the ECB. However, tier two assets must at least meet the same high credit standards as tier one assets.

The risk managers assist in fine-tuning the protection provided by the collateral by calculating various haircuts to be applied to the monetary policy repos. This ensures that adequate collateral is available at all times. In 2004, tier one assets where classified into four liquidity categories and each assigned a specific haircut schedule. Market and liquidity risks are the main risks addressed in developing these risk control measures.

Risk management issues related to the management of the foreign currency reserves and the ECB's own funds

The foreign currency reserves are invested in high-quality and liquid assets, and are split among the 12 NCBs of the Eurosystem. The size of each portfolio is based on the capital key of each NCB.[2] The ECB has a two-tiered benchmark system for the foreign currency reserves with a strategic benchmark for each relevant currency, reflecting the long-term investment objectives of the ECB. This is supplemented by a tactical benchmark set by the ECB's portfolio managers and organisationally separated from the Risk Management Division. NCB portfolio managers, who in practice manage the actual portfolios, are permitted to take positions within a predefined band around this tactical benchmark, which is, in turn, restricted to a band around the strategic benchmark. There is a credit policy and limits system that restricts the instruments, countries, issuers and counterparties that NCBs may use. The Risk Management Division is responsible for analysing and reporting the performance of both the benchmarks and the NCB portfolio managers. A similar although less elaborate structure exists for the ECB's own funds (the counterpart of capital and reserves), which is managed by a dedicated Portfolio Management Division within the ECB.

In terms of investment operations, the risk management area has three key areas of responsibility: risk management measures and compliance, analysis and reporting of investment performance, and strategic asset allocation.

3.1.1 Risk management measures and compliance

The ECB's exposure to market, credit and liquidity risk is measured and compliance with agreed limits is checked on a daily basis. Market risk exposure is measured using modified duration and Value-at-Risk (VaR). Credit and settlement exposure are also calculated daily and checked against a range of counterparty and asset class limits. All eligible counterparties must meet minimum credit criteria and satisfy certain operational constraints. Rating and other credit information are monitored together with ongoing trading volumes with portfolio managers to ensure that the counterparty complies with the credit requirements and has a good business relationship with the portfolio managers. In addition, the liquidity profile of the investments is monitored daily. Liquidity limits are in place to ensure that adequate amounts of assets are held in cash or highly liquid securities should the need to intervene arise. Clear procedures are in place for reporting and dealing with any breaches of the limits.

---

[2] The capital key is the percentage that each NCB of the Eurosystem had to contribute to the initial capital of the ECB. The weighting in the key is equal to the sum of 50% of the share of its respective Member State in the population of the Community and 50% of the share of its respective Member State in the euro area's GDP.

There are also daily procedures to verify the accuracy of the prices used to mark to market positions, as well as a mechanism to ensure that all prices traded at by portfolio managers reflect the market price range for that day.

3.1.2 Analysis and reporting of investment performance

The measurement and analysis of the performance of all the relevant investment portfolios forms a large part of the risk management work. The current framework has 24 separate portfolios in the foreign reserve management area alone. Monthly performance is calculated for all portfolios. In addition, more detailed and analytical semi-annual and annual performance reports are prepared.

Performance attribution techniques are also used to determine the source of performance. The source of performance relative to the investment benchmarks provides the portfolio managers with useful feedback. Aggregating this data over longer periods also allows more detailed conclusions to be drawn. This analysis of overall performance as well as the relative performance of portfolios versus the relevant benchmark can also be useful for policy-makers in deciding on the investment framework.

3.1.3 Strategic asset allocation

Strategic asset allocation is the process of dividing investments among different kinds of assets to optimise the risk-reward trade-off based on the institution's goals. The key decision influencing the return of the investment portfolios is the choice of strategic benchmarks.

Strategic asset allocation for the own funds involves the selection of an optimal allocation of assets. For the foreign exchange reserves, this involves a two-step process. The first step is to determine the optimal currency mix of the foreign currency reserves. The second step is to propose an optimal asset allocation for each currency portfolio. All of the above are reflected in the strategic benchmarks. At the moment, work is also underway to explore the possibility and implications of combining these two steps into a joint optimisation exercise (currency and interest rate benchmark). This could facilitate greater exploitation of diversification benefits.

The strategic benchmarks are designed to reflect the long-term risk return preferences of the ECB. Given the importance of selecting the correct investment benchmark, considerable resources are devoted to the development of the methodology used in the asset allocation process. Recent developments include the construction of a purpose-built econometric model which can be used to derive the expected returns. The normal mean variance optimisation technique is also supplemented by alternative optimisation techniques designed to enhance the robustness of the analysis. The use of forward-looking returns is a considerable improvement on using historical data as input for the optimisation exercise. Given the importance of the asset allocation methodology to the ECB's investment decisions, the process is continuously reviewed with the aim of seeking improvements or refinements.

4 Organisational set-up

Originally, in line with the recommendations of bank supervisory guidelines, the risk management function was organisationally separate from the operational area of the ECB. The Head of Risk Management reported directly to a member of the Executive Board. The ultimate decision-making body of the ECB is the Governing Council, consisting of the six Executive Board members of the ECB and the governors of the 12 NCBs of the Eurosystem.

Overall policy is decided by the Governing Council, but much of the operational authority for managing risk, such as the approval of counterparties and the tactical benchmarks, is delegated to the Executive Board. Following an organisational review, it was decided to place the risk management area within the operations area for administrative reporting. However, for issues relating to the establishment of strategic benchmarks, reporting on performance and risk, non-compliance and the assessment of eligibility of investments and counterparties, a direct reporting line to the Executive Board member responsible for the operations area was maintained. For other purposes, the Head of Risk Management reports via the Head of Operations to the relevant board member. This set-up attempts to fuse organisational efficiency with the unique and independent nature of much of the reporting and assessment provided by a risk management division.

To facilitate communication with the NCBs, the ECB regularly consults and receives feedback from the NCBs on various operational aspects including risk management. The main vehicle for this is the Market Operations Committee (MOC). This is a Eurosystem committee comprising ECB members and representatives from each of the 12 NCBs of the Eurosystem. Its focus is on the investment operations of the ECB and issues arising from the implementation of monetary policy. This committee is extensively consulted and can make comments and observations on papers that are sent via the Executive Board to the Governing Council. The committee is supported by three working groups: the Working Group on Foreign Reserve Assets WGFRA, the Working Group on Monetary and Exchange Rate Policy and Procedures (WGME) and the Monitoring Working Group (MWG).

Another important committee worth mentioning is the Asset and Liability Committee (ALCO), which is an internal ECB committee. Its mandate is to consider the impact of various individual investment, accounting and risk policies on the overall asset and liabilities positions of the ECB. In particular, it assesses the impact of investment activity on the ECB's balance sheet and profit and loss account. It also makes projections of income and relates these to annual budget needs. It is an important forum for the formulation of ECB policy and plays a key role in the ECB's own corporate planning. It also facilitates interaction between the accounting, operational and risk management areas of the ECB.

5 Risk management support

Two key support functions that the ECB risk management draws heavily on are IT support and legal services. Both are crucial to delivering sound risk management practices.

For its investment operations, the ECB uses an IT system that captures all trading activity in real-time. The system also calculates performance on a daily basis and has a comprehensive limit functionality to ensure that the portfolio managers behave in accordance with their mandates. An additional system designed specifically for risk management requirements is currently being developed to supplement the core system. It will provide additional analytical tools, improved database facilities and enhanced reporting for risk management and performance.

The system support for the monitoring of risks associated with monetary policy operations is not standardised to the same extent. Risk policy is defined at the ECB level but implemented at the NCB level. The collateral management systems may differ between the NCBs. Nevertheless, there are centrally maintained databases of eligible assets and counterparties for monetary policy operations. The NCBs provide the asset and counterparty updates, while the database maintenance takes place at the ECB. The Eligible Assets Database contains, inter alia, details of the credit rating of assets and the appropriate risk control measures. The ECB regularly collects data on the use of eligible assets and the various participants in operations. This material is collated and analysed and, on an annual basis, the Governing Council is presented with a report on the use of underlying eligible assets and the participation of counterparties. The collation of this information is not fully automated yet, but plans are being made to improve automation as well as to further standardise systems used for collateral management at the NCB level.

Finally, it is also important to remark that all ECB operations are supported by adequate legal documentation. Eurosystem operations are conducted under a contractual or regulatory framework to which counterparties are legally bound. In those euro area countries in which there is a contractual framework, counterparties enter into comprehensive, customised contractual documentation with the local national central bank. In addition, ECB foreign exchange and foreign reserve management operations are conducted under standard industry-wide repo and swap master agreements, and a proprietary master netting agreement, all of which have been executed with the ECB's counterparties. From a credit risk perspective, a primary motivation for executing such contractual documentation is to protect the ECB's interests in the event of counterparty insolvency and, in particular, to activate the close-out netting provisions in these agreements.

6 Conclusion

To date, the risk management framework established by the ECB has functioned well. The ECB has incurred no credit losses, has had no liquidity problems and has quite successfully managed market risk in accordance with its expressed risk return preferences, generating substantial returns at the same time. The management of currency risk, however, has proven more problematic. The ECB performs an annual exercise to assess the proportions of reserve currencies that minimise exchange rate risk relative to the euro. Currency risk is not managed on a dynamic basis, and although such active management would arguably be inconsistent with the ECB's membership of the central banking community, this nevertheless represents the largest risk factor for the ECB balance sheet. Further work is needed to improve management of this risk while respecting the parameters that the policy goals of the ECB impose on the bank's investment operations.

The main challenge facing the ECB in the near future is the enlargement of the euro area. In 2004 ten new member states will join the EU[3]. Ultimately, these countries will also participate in the euro and the Eurosystem will expand to at least 22 members. The exact design of the reserves management framework in the expanded Eurosystem has yet to be decided. Nevertheless, it is highly likely that enlargement will increase the burden of analysis, monitoring and coordination as well as the absolute size of foreign reserve assets involved and the variety of monetary policy instruments.

Risk management in the ECB must evolve in order to meet these changes. It is important to the ECB that the risk management structures in place are flexible enough to evolve in response to these changes, but also sufficiently resilient to be able to track all relevant risks effectively. As the assets of the central bank represent assets for the citizens of the euro area, the optimal management of these risks can enhance the economic wealth of the euro area.

[3] The ten acceding countries – Cyprus, the Czech Republic, Estonia, Hungary, Latvia, Lithuania, Malta, Poland, Slovakia and Slovenia – will join the EU on 1 May 2004.

Owing to its position at the centre of the Eurosystem, the ECB is extremely aware of the importance of managing, monitoring and controlling risk to ensure the success of the Eurosystem. The ECB has invested considerable resources in its risk management structures to ensure that a complete view of the risks facing the Eurosystem is available at all times. Knowing and managing this risk is critical for the credibility of the Eurosystem.

The challenge for the ECB is to build on its successes to date and to meet and respond proactively to the challenges that will emerge in the future. Its risk management philosophy is not the avoidance of all risks at all costs, but instead a careful and analytical approach ensuring that the full risk implications of decisions are understood at the different levels of the organisation. It is vital for the success of the ECB that this philosophy should continue to inform its decisions and business practices.

16

Management of currency distribution and duration

Karel Bauer, Michal Koblas, Ladislav Mochan, Jan Schmidt

Abstract

The paper discusses the approaches Česká národní banka (CNB) uses to define the currency composition of its foreign exchange reserves and the duration of its FX reserves portfolios. Currency composition and duration are the parameters that have the greatest impact on foreign reserves returns. Because the risks connected with these parameters are of different importance (99% of VaR of FX reserves is due to currency risk), these parameters are treated separately. Currency composition is based on the following assumptions: only USD and EUR are considered; CNB does not hold any significant FX liabilities; and from the return point of view, there is no preference for any of the currencies. Two approaches are then examined: optimisation of the composition with respect to the koruna value of FX reserves, and the liability match approach. Duration is set to match the no-negative-return requirement with the use of historical data analysis. Static and dynamic methods, taking into account the actual level of interest rates, are also discussed.

1 Introduction

Česká národní banka's (CNB) objective in managing foreign exchange reserves is to achieve maximum and stable returns subject to the liquidity restrictions and the limits on market and credit risks. This paper sets out to show how CNB implements this objective on the two most significant components of the return – currency composition and duration.

CNB does not use a single approach tailored to find the optimal set-up for both parameters at once, e.g. Value-at-Risk (VaR), as the risks connected with these parameters are of different scales. Measuring by VaR, it was found that over 99% of the risk was accounted for by currency risk. In addition, as explained later, VaR is used only as a complementary method to other risk measures used in CNB.

2 Currency composition

The massive capital inflows to the Czech Republic from abroad in the period 1998–2003 resulted in a significant increase in the country's foreign reserves. Apart from the regular interventions during that period, CNB, as an agent of the government, purchased about €6 billion in privatisation proceeds. The reserves became an outgrowth of monetary policy and today they account for 90% of CNB's balance sheet. On one hand, this means that the foreign exchange reserves are practically the only source of income for Česká národní banka. However, on the other hand they represent a high risk of financial losses.

Prior to any analysis seeking to set the currency composition, three assumptions have been made:

- no currencies other than USD and EUR would be considered (although this two-currencies-only approach may be sub-optimal, the costs connected with an introduction of additional currencies would reduce the diversification effect to a great extent);

- there is no explicit foreign exchange liability that would serve as an anchor;
- there are no economic essentials that would back holding either of these two currencies from the return point of view. The question of maximising returns is transformed into a question of minimising losses.

The entire process can then be broken down into two steps.

In the first approach, we consider the relative weights of the reserve currencies. We optimise their expected koruna value based on statistical models. This approach is convenient as long as the assets are managed with an infinite horizon, in which their koruna value is of primary importance.

The second step consists in identifying the future outflows (scenarios) in absolute terms. The optimisation assumes the adoption of a currency structure consistent with that of the future liabilities. In doing so, we minimise the exchange rate risk arising from fluctuations in the rates of reserve currencies. The koruna value of the portfolios is not considered.

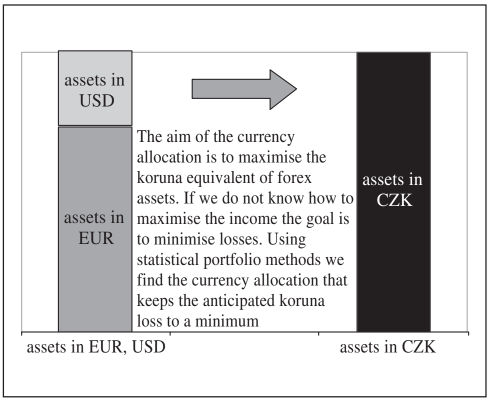

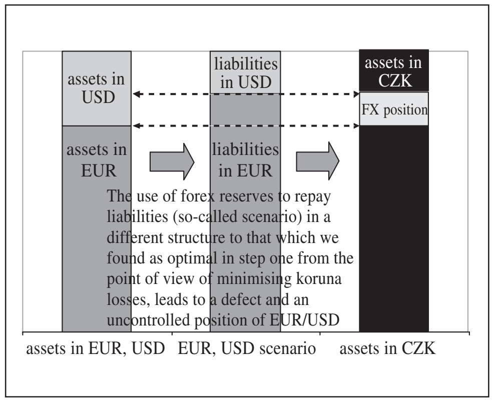

The two approaches are not consistent. Prioritising the latter one may create a koruna loss-bearing position, which may be observed when we evaluate the difference between the koruna value of the optimal portfolios of the first and second approaches.

2.1 Optimisation of the currency allocation with regard to the koruna value of the foreign exchange reserves

Seeking the optimal currency allocation in the first step, we minimise losses upon the re-evaluation of FX reserves in CNB's accounting. It is analogous to the principle of immunisation (i.e. the comparison of assets and liabilities). The liabilities side of the CNB balance sheet is mainly composed of koruna liabilities (OMOs, money in circulation, capital). The alternative is to maximise profit. Since we have found no economic reason to support! – with respect to return – the holding of one currency at the expense of the others, the statistical application of the strategy of maximisation of expected return has been rejected. Therefore the focus is on minimising losses either in terms of minimum VaR or minimum volatility of return (assuming long-term average returns do not vary across different currency compositions, and assuming volatility varies across different currency compositions), as illustrated on the chart bellow.

Chart 1: Probability distribution of two portfolios with different currency structures and identical average return

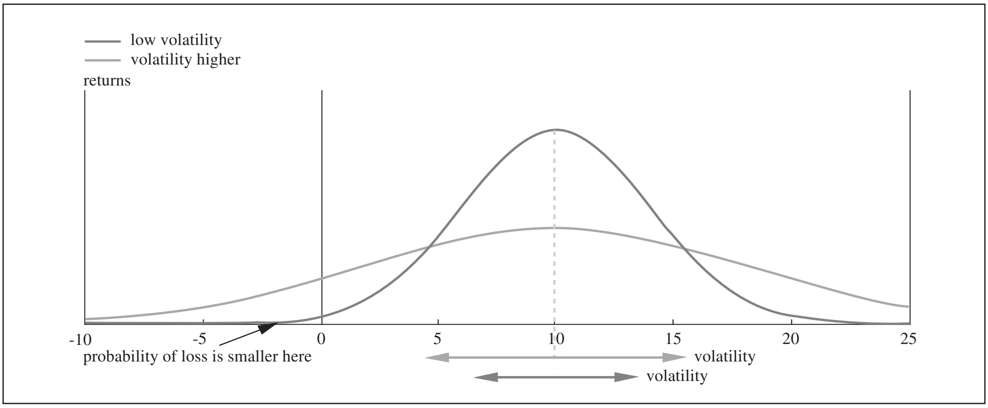

2.1.1 Minimisation of expected losses

This approach consists of the application of traditional VaR. We use the historical time series of the benchmark indices to model the distribution of probability of returns. Then, a particular worst-case quantile is found for each distribution. The calculations are repeated for portfolios with different currency composition, and a portfolio with the lowest VaR is found. This portfolio is assumed to be optimal.

This approach has the following characteristic. A one-year horizon has been used at the level of 1% confidence. The distribution of the returns has been normalised for the reasons explained below.

Since we are looking for the optimal portfolio rather than the accurate value of its VaR, we feel comfortable using the normal distribution prior to others. We are allowed to do so thanks to the following assumption. A portfolio which is VaR-optimal for normal distribution of returns is assumed to be optimal as well for any other symmetrical, semi-monotone distribution with a finite second momentum. This allows us to say that we have found the optimal portfolio not only for the normal distribution of returns, but for any other distribution

Chart 2: VaR dependent on dollar representation

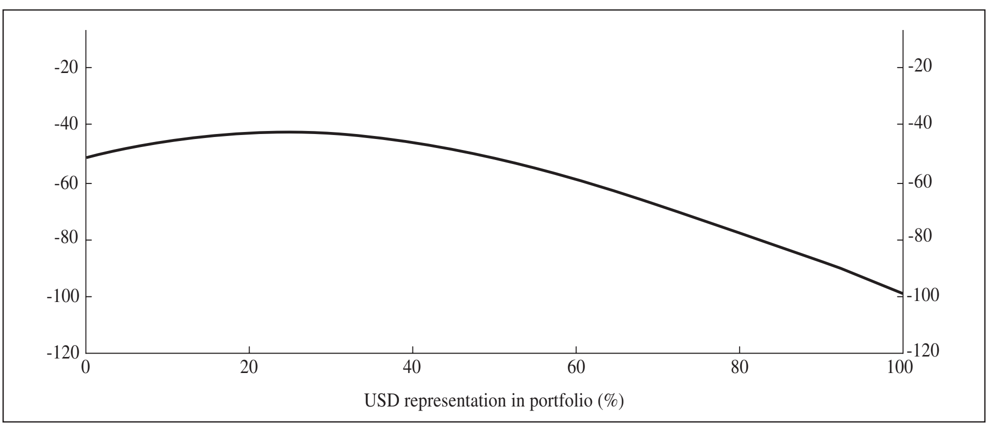

satisfying the assumptions above. It does not allow exact calculation of the values of the extreme-case losses, but it allows us to find portfolios, which generate them. We use VaR as a method that complements others. Rather than as an applied analysis, we use it for illustrative purposes.

The best koruna value of VaR was obtained for a portfolio with a USD proportion of 20-25%, the rest being allocated in EUR.[1]

2.1.2 Minimisation of return volatility

This strategy consists in finding a currency structure that will have the lowest volatility of anticipated returns. The economic rationale behind this choice is the anticipation of the long-term efficiency of the EUR/USD FX markets in relation to the development of the interest rate differential. It may be assumed that the currency composition of reserves does not have a significant impact on their long-term accumulated returns. If this is the case, then a portfolio with the lowest volatility of returns should be selected as the one with the lowest risk of short-term losses.

Chart 3: Average annual income at today's CNB benchmarks over the previous 18 years (9 years CZK) measured in different currencies

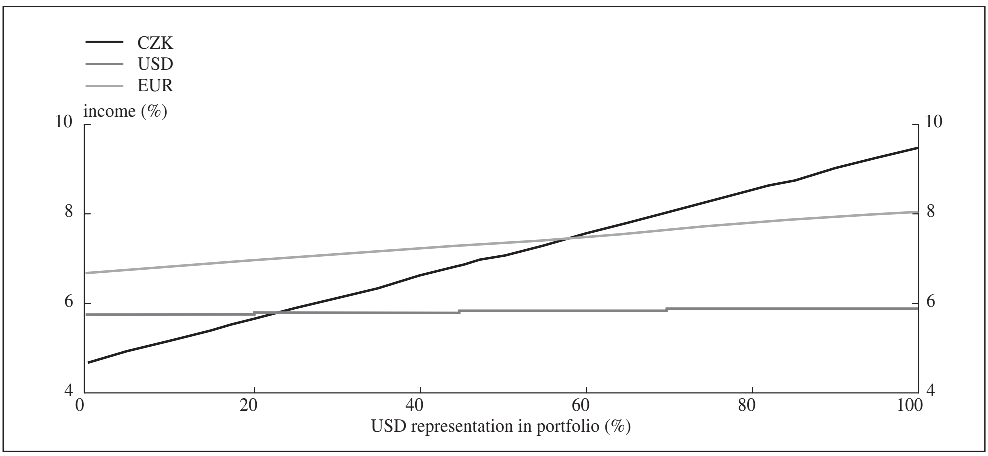

Whether it is realistic here to anticipate long-term, identical returns for portfolios with different currency compositions remains a theoretical question. In this respect, the attached graphs are only illustrative. They show the average annual returns for portfolios with different dollar representations, where returns are measured gradually in all three reference currencies. Whereas the returns seem to be even for the dollar, this is not the case for the euro or the koruna. The question remains whether the situation would balance out if a longer horizon were selected.

[1] A more precise figure depends on the portfolio's duration, although this is of second-order importance.

Chart 4: Volatility of annual returns in CZK

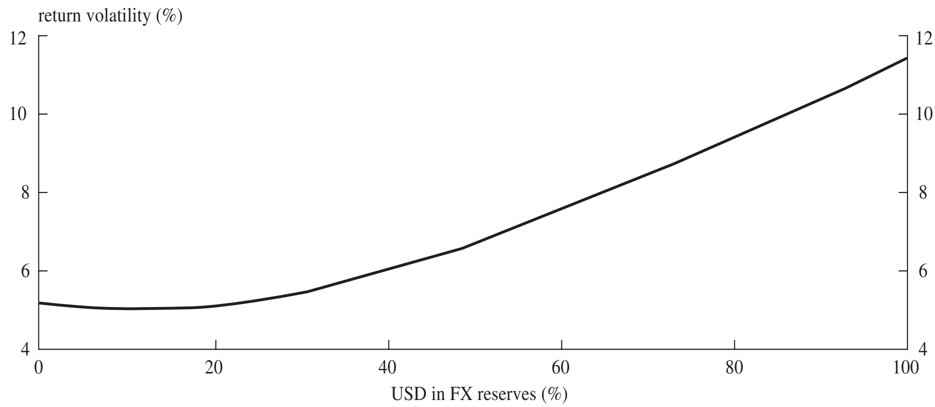

It may be concluded therefore that the method of minimum volatility is essentially the method of VaR, where there is a limitation of the degree of freedom in the original portfolio model from average return and volatility to just volatility. When looking for a portfolio that is ideal in this sense, the returns from portfolios with different currency compositions are modelled gradually. The portfolio with the lowest volatility of annual return measured monthly is selected as the most appropriate one.

Between 1994 and 2002 we measured the returns of portfolios based on the fusion of the euro and dollar benchmarks used for foreign exchange reserves. This fusion is made up of various USD and EUR ratios; here we are looking for a ratio that produces the lowest volatility of returns. We see, for example, that when measuring returns in Czech koruna and monitoring return over a period of one year, the lowest volatility was recorded by a portfolio consisting of a 14% USD benchmark and 86% euro benchmark.

2.2 Liabilities match

This approach tends to minimise EUR/USD FX rate risk in the event that the reserves are used to repay some liabilities (immunisation against liabilities). The optimal currency allocation then corresponds to the minimum currency risk[2], as illustrated in the chart above. The scenarios are chosen so as to reflect the function of FX reserves. The most important aspect is not the highest koruna foreign exchange equivalent, but to own the currency in which the liability will be paid, thereby avoiding undertaking conversion risk. For each of the scenarios, a currency structure exists for foreign exchange reserves with the lowest currency risk. It is easy to assume that the “suitability” of this approach depends on whether the scenario in

[2] If, for example, we had a currency structure of assets of €70 billion and USD 30 billion, and future liabilities of €50 billion and USD 50 billion, our current position is long in euro by 20 billion and short in dollars by 20 billion. These 20 billion are the source of currency risk which, at the moment of repayment of the liability, will manifest itself in the form of a profit or loss resulting from a change in the value of €20 billion measured in dollars accumulated over the whole period, from now to the repayment of the liability. The optimal currency allocation is therefore €50 billion and USD 50 billion.

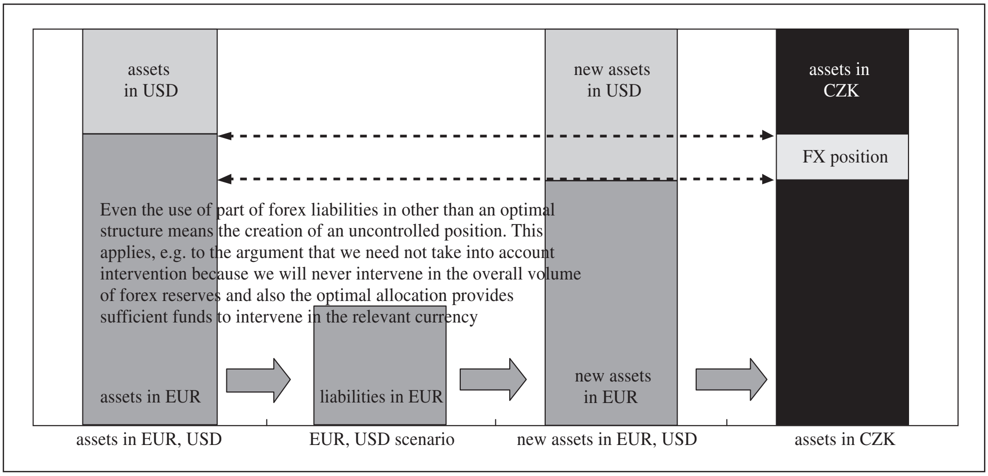

question actually occurs and whether its impact on the business of Česká národní banka is not prioritised. It should be pointed out that the following scenarios are rather complementary than alternatives, e.g. the need to finance the balance of payments does not exclude the possible need for interventions.

2.2.1 *Scenario: Financing the balance of payments*

Let us assume that the function of foreign exchange reserves is to cover the state's foreign exchange payments. If a FX market exists, this function is fulfilled by interventions. Conducting the interventions on the domestic market would imply a 100% holding of reserves in EUR. If the koruna market were to crash, forex reserves would be the sole source of foreign exchange for trading purposes. The aim is therefore to maximise the purchasing power of the reserves. That means maximising the amount of reserves measured in a currency basket that corresponds to the structure of foreign trade, with no option of profit repatriation and capital outflow. The long-term import average is 77% in EUR and 23% in USD[3], which would accordingly set the currency composition of the basket in the same way.

[3] It would be technically more correct to work with the difference between import and export, which is, however, 20% EUR and 80% USD. This approach produces a corresponding result, although the given figures are much less stable and do not correspond to the generally applied methodology, while the current FX reserves are equivalent to the trade balance for around five years.

2.2.2 Scenario: Entry of the Czech Republic into EMU

Thirdly, let us assume a scenario whereby the Czech Republic joins EMU. In this situation we will consider the currency risk implied by flows and arrangements associated with EMU entry. The liabilities are identified as follows:

- **Size of the future major liabilities**

  After the Czech Republic's accession to EMU, there will also be a re-denomination of the CNB balance sheet – koruna assets and liabilities will become euro assets and liabilities, and the originally foreign exchange reserves in euro will become reserves in the domestic currency. There are two possible strategies. The first one assumes that from now up until joining EMU, CNB will maintain the non-euro part of its reserves at the level it will hold after EMU entry as its future non-euro open position assets. The second possible approach is, conversely, to maintain the euro part of the reserves at a level corresponding to the size of the euro liabilities "payable" upon EMU entry. Both approaches require the target reference volume of future liabilities to be determined and/or estimated. As it turns out, a link can be established between them which allows a solution to be found which is consistent with both approaches.

- **Euro liabilities**

  Most of CNB's current passive balance is denominated in koruna. After EMU entry, this would be re-denominated into euro. Two major items of the passive balance are the currency base and repo operations. The latter item represents the monetary operations that withdraw liquidity from the market. These two items account for 90% of the balance, and both will become EUR-denominated liabilities after EMU entry. Unlike the currency base, the repo operations may, soon after EMU entry, be changed into other liabilities. If the domestic commercial banks transfer their funds from CNB to any foreign counterparts within TARGET, the CNB repo operations liability would thus turn CNB's large-scale liability towards TARGET.[4]

  In order to hedge the currency risk associated with this liability, we shall target the volume of the euro part of the FX reserves to the volume of the current domestic repo operations. Their amount outstanding is roughly equal to the sum of

- the net size of any possible CNB FX market operations during the EMU accession period,
- the ultimate size of the repo operations at the end of this period.

  That is because CNB's FX market operations are fully financed from domestic market repo operations as a result of the Bank's monetary policy. It assumes that the more EUR-denominated assets are spent on FX market interventions during this period, the less repo operations are needed to finance the FX reserves, and consequently the smaller the final size of the Bank's repo operations (and vice versa).

---

[4] Conversion of liabilities vis-à-vis domestic banks to liabilities vis-à-vis the NCBs and ECB from TARGET. Domestic commercial banks currently deposit excess liquidity at CNB. On joining EMU, we assume that (i) CNB will cease to offer this option (within the framework of ECB monetary policy implementation, liquidity will – on the contrary – be provided), and (ii) commercial banks will seek more appropriate investment opportunities (this should be facilitated by the fact that the euro market generally has a shortage of liquidity). For domestic banks to be able to realise the latter, they will be compelled to transfer their funds via their correspondent (currently CNB) to the account of a foreign bank (whose correspondent bank is one of the NCBs). However, TARGET works in such a way that funds are not transferred between the two NCBs, but the payer incurs a liability vis-à-vis the beneficiary (the CNB vis-à-vis one of the NCBs) and the beneficiary a claim against the payer (the NCB against the CNB). Hence the payment is not an "expense", but merely a liability.

Options for repaying TARGET liabilities vis-à-vis NCBs: a liability can be repaid by a set-off. This means that if CNB had a liability vis-à-vis an NCB, it would only be set off in a situation where that NCB makes a payment in favour of CNB. This can be achieved, for example, by CNB selling part of its EUR-denominated securities portfolio to buyers whose correspondent bank is the relevant NCB.

The sum of the EUR-denominated assets needed to fund the above liabilities (FX market operations *plus* the ultimate size of the repo operations upon EMU entry) can thus be determined from the outstanding amount of the current repo operations. In recent times, these represent funding of over two-thirds of the FX reserves. With the euro part of the FX reserves being targeted at the level of our future TARGET liability to the NCBs and any outstanding FX market operations during the EMU accession period, the currency structure requirement would be to hold an equivalent of around €16 billion (CZK 500 billion).

- **Non-euro liabilities**

Targeting the non-euro part of the FX reserves at the size of the future non-euro liabilities assumes that the size of these liabilities can be confidently predicted. Beside marginal outflows, the creation of the post-EMU entry FX reserves turns out to be the only expected large-scale liability. The size of the post-EMU entry FX reserves will be subject to future debate; however, even without this knowledge, we can still link arising liabilities to those on the EUR-denominated side.

- **The link between euro and non-euro liabilities**

A link between targeting EUR at the volume of TARGET liabilities and targeting the USD part of the reserves at the monetary base level can be detected. Targeting the non-euro part of the reserves at the amount of the monetary base is consistent with targeting the euro part at the amount of repo operations. With each additional euro by which the reserves are increased (for instance as a result of CNB interventions), the volume of repo operations sterilising the domestic currency rises. It follows that holding this additional euro in this currency covers the exchange rate risk of the future liability from repo operations. If at some point in the future they are settled against each other, an amount roughly corresponding to the size of the monetary base will remain in the FX reserves, which would be denominated in USD.

Using the "liabilities match" approach, however, results in the following:

- The currency structure derived from the requirement for a balanced position as outlined does not solve exchange rate risk from the re-denomination of foreign exchange assets in domestic currency.
- As the results are different, it is not clear-cut to find a compromise (suitable combination) that would fit all scenarios to a certain extent.
- Each scenario can have a different probability, although conversely, each scenario has a different relevance. For example, if CNB suffers an exchange rate loss[5] realised in its future liabilities towards the NCBs in TARGET, this would represent a matter of accounting; if however the CNB as a result of an exchange rate loss found its power of intervention reduced, in an extreme case its very credibility as a central bank could be threatened.
- All considered quantities will change in time (the structure and size of foreign trade will change according to the development of the economy and exchange rate; while the amount of repo operations, just like the currency base, will change according to the size of future interventions and the development of the exchange rate and money in circulation). It will therefore be necessary to adapt the currency structure to these changes.
- All considered quantities may change as a result of sudden movements on both sides. The currency allocation should therefore be decided by mean values within relatively broad intervals.

---

[5] Owing to the accumulated appreciation of the domestic currency after the capital inflows in 1999-2003, and before its possible outflows after the EMU entry.

3 Principles for setting out benchmark duration

As CNB has only minor FX liabilities, the duration of the FX reserves has no "natural" anchor. Therefore, for the purpose of setting up the reserves duration, CNB focuses on achieving the maximum return at the lowest possible risk. Once the set of permissible instruments represents the highest rating, the return is primarily a function of duration. Generally, we can say that the higher the duration, the greater the return, as measured by historical returns. In simpler terms, the maximisation of return therefore means the maximisation of the ratio of the bond component in reserves. The interest rate risk limit can be expressed as the goal to retain the absolute value of the relevant portfolio in the appropriate currency, i.e. not to incur any negative return for the period in question, which is twelve months in CNB's case.

Choosing components for the benchmark is rather arbitrary. In our case it is an index created from components from the money market and bond market. For example, for the EUR benchmark the money market components are the 1-month Libid, the 3-month FIXBIS and the T-bill Merrill Lynch EGB0 index. The bond market is represented by an index made up of bullet bonds from Germany, France and Italy[6] with maturity periods of between one and ten years; each country moreover has a one-third presence in the index. This structure takes into account our desire to cover, if possible, the entire euro yield curve and to a certain extent our opinion on the required liquidity.[7]

3.1 The static approach

For the sake of liquidity constraints, another limitation was introduced. The average maturity of components from the money market is set at three months and furthermore there is a fixed ratio of the FIXBIS component of 10% in the benchmark. The model then has three assets: the 3-month Libid will approximate the money market part, the bond part will approximate the appropriate bond index, and the FIXBIS component has a fixed weight. The goal is to find the minimum ratio of the Libid component so that the overall return from the portfolio for each rolling 12-month period is not negative. Data from 1992 were used.[8]

$R_{3m,i}$ represents the return from the Libid component, $R_{f,i}$ the return from FIXBIS, and $R_{b,i}$ the return from the bond index for the i rolling 12-month period, w_{3m} the weight of the Libid component, w_f the weight of the FIXBIS and w_b the weight of the bond component. The return for the benchmark R_i for the i period is thus given by the equation

$$R_i = R_{3m,i} \cdot w_{3m} + R_{f,i} \cdot w_f + R_{b,i} \cdot w_b \, .$$

[6] These indices have been chosen because of their liquidity; other ones could also substitute for them in this example as well.

[7] With respect to the size of the reserves, we assume that it is sufficient to keep a minimum of 25% of the assets in money market instruments. This decision is embodied in the fixed part of the benchmark's composition (e.g. on the US side the benchmark includes 3-month FIXBIS, and 6-month T-bills, together amounting to a fixed 25% of the benchmark).

[8] The reason for using a ten year old series is that BBG does not provide data on the desired instruments before 1992.

$\overline{R}$ is a matrix of returns of type (n,3), where n is the number of 12-month periods, and w the vector of the weights of individual components. Taking into account the defined limits, it arrives at a simple problem of linear programming:

$$\min w_{3m}$$
$$\overline{R} \otimes w \geq \Theta,$$
$$w_{3m}, w_f, w_b \geq 0,$$
$$w_3 + w_f + w_b = 1.$$

It is also known that $w_f = 0,1$; this means that $w_b = (1 - w_{3m} - w_f) = 0,9 - w_{3m}$. On condition that $R_{3m,i} \geq 0$ (always fulfilled due to the properties of Libid) and $R_{b,i} \leq 0$ (because only these cases interest us), the problem is simplified to looking for the minimum w_{3m} that fulfils

$$w_{3m} \geq \frac{0,1 \cdot R_{f,i} + 0,9 \cdot R_{b,i}}{R_{b,i} - R_{3m,i}} \quad \text{for all } i = 1,..,n.$$

Using this procedure, we find a benchmark for which a following property is guaranteed. Even in the event of a repetition of the historical scenarios over the past ten years, the benchmark will not set down a loss for any 12-month period. Such a benchmark, however, will have a fixed structure and therefore also a fixed duration.

One of the consequences of the above-mentioned procedure is that the duration fixed in this way may in the future only be reduced (upon condition that data are considered from the same initial point). This means that the system is static and ignores current developments on the financial markets. Nevertheless, the current size of short-term interest rates significantly affects the amount of interest rate risk that can be undertaken in order to fulfil the requirement for a non-negative anticipated return in each 12-month period.

3.2 The dynamic approach

Considering the influence of the current size of interest rates on the benchmark's duration, the ratio of one of the components in the benchmark would have to become a function of these short-term rates. This approach is based on the assumption that the higher the short-term interest rates, the greater the coverage of any capital losses from the bond component by the return from the money market components. We have documented this relationship together with our partners from Bear Stearns. The result of the above-mentioned assumption is that w_{3m} should be an inverse function of short-term interest rates; we thus define $w_{3m,i} = \gamma / R_{3m,i}$. The task is to find a minimum γ that fulfils

$$\gamma \geq \frac{(0,1 \cdot R_{f,i} + 0,9 \cdot R_{b,i}) \cdot R_{3m,i}}{R_{b,i} - R_{3m,i}} \quad \text{for all } i = 1,.., n.$$

This dynamic approach, as we will demonstrated, provides higher return under the same risk limitation conditions. The static approach is also potentially more risky if interest rates fall. Conditions may arise that do not fall within the framework for historical data used to define the static setting w_{3m}; consequently, the static setting w_{3m} under these new conditions may not be optimal. The dynamic approach, on the other hand, is capable of reacting to this to a certain extent.

The weights of other components at the short end of our benchmark, i.e. 1-month Libid and T-bill components, are then arbitrarily defined. Their aggregate is equal to $w_{3m,}$ so that the average maturity of all components of the money market, including FIXBIS, is equal to 0.25 of one year. The following chart compares the respective returns on benchmarks where the setting is defined by a static or dynamic approach.

Chart 5: Return on benchmarks

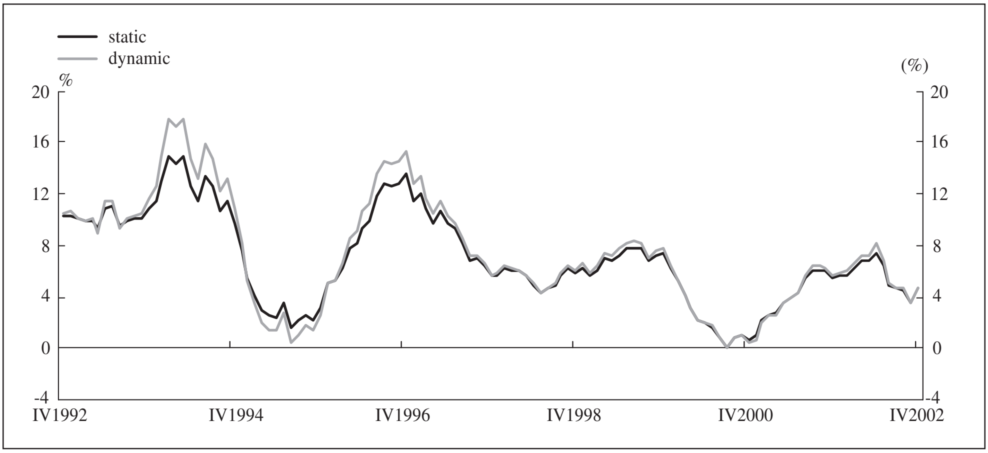

It can be seen that both approaches ensure that the benchmarks always have non-negative returns; however, the dynamic benchmark achieves this goal with a higher return. The average annual extra return of the dynamic benchmark over its static equivalent amounts to 83 basis points for the period 1992-2002.

The following chart shows how the benchmark's duration alters depending on the development of the 3-month Libid. It can be seen how the duration of the dynamic benchmark rises in relation to increases in the 3-month Libid.

Chart 6: Duration versus LIBID 3M

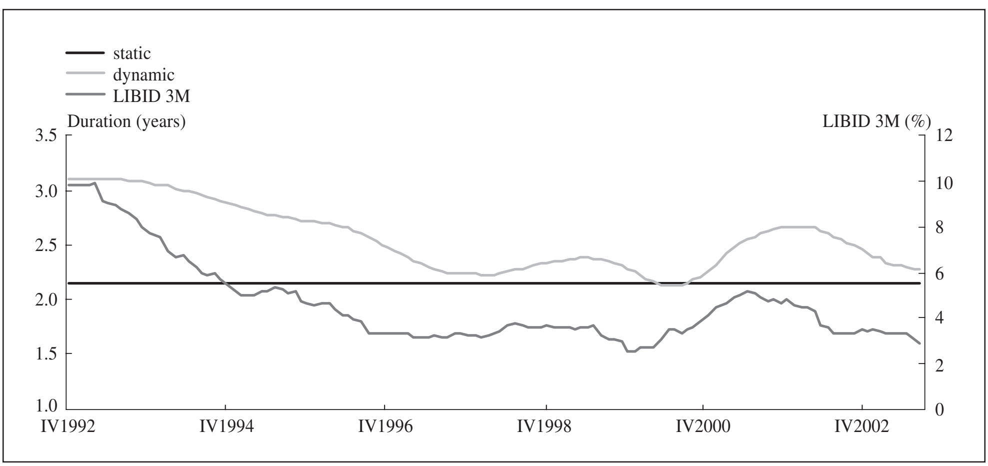

For an approximate mathematical expression of the relation between the change in duration and the change in the 3-month Libor rate, we proceed from the relation for the benchmark's duration:

$$D = w_{3m,i} \cdot D_{3m} + 0{,}1 \cdot D_f + (0{,}9 - w_{3m,i}) \cdot D_b \,,$$

where D_{3m}, D_f, D_b are the duration of individual components of the benchmark. After inserting the values for the duration and the relation for $w_{3m,i,}$, we get

$$D = w_{3m,i} \cdot 0{,}250 + 0{,}1 \cdot 0{,}250 + (0{,}9 - w_{3m,i}) \cdot 3{,}847 = 3{,}487 - 3{,}597 \frac{\gamma}{R_{3m,i}} \,,$$

The change of the duration $\Delta D=D^{(2)}-D^{(1)}$ for the change of the return from $R^{(1)}_{3m,i}$ to $R^{(2)}_{3m,i}$, is then equal to

$$\Delta D = 3{,}597 \cdot \gamma \cdot \frac{R^{(2)}_{3m,i} - R^{(1)}_{3m,i}}{R^{(1)}_{3m,i} \cdot R^{(2)}_{3m,i}} \,,$$

This formula shows that the size of the change in duration depends not only on the size of the change in return of the 3-month Libid, but also on the absolute amount of the return. After inserting the value *g*, arrived at through the calculation described above, this relation is shown on the following graph.

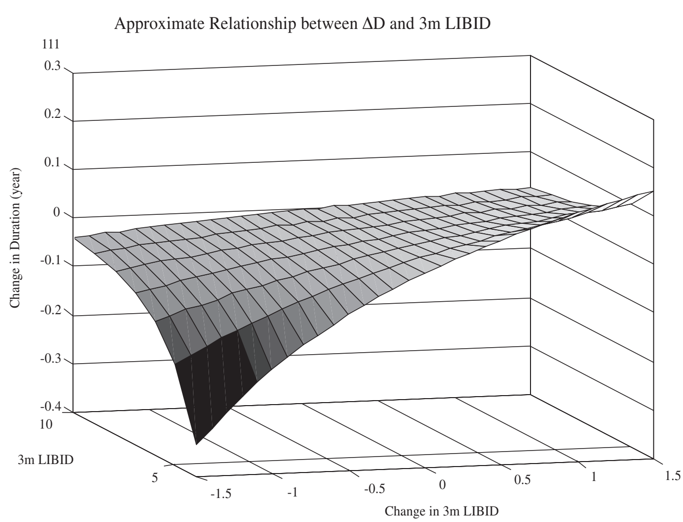

The question that arises is how safe is the dynamic benchmark in comparison with the static benchmark in the sense of non-negative return over the 12-month period. Here, a decisive role is played by the return from the bond index, which is the most volatile and also the only one of all the benchmark's components to record negative return over the 12-month period. Focusing on the behaviour of the return from the bond index, the question is how large the change in the return from the bond index would have to be for the overall return of the benchmark to be zero. The following graph shows the maximum permissible change for both benchmark settings in comparison with the actual behaviour of return of the bond index over 12 months.

Chart 7: Rermissible versus actual returns of bond index

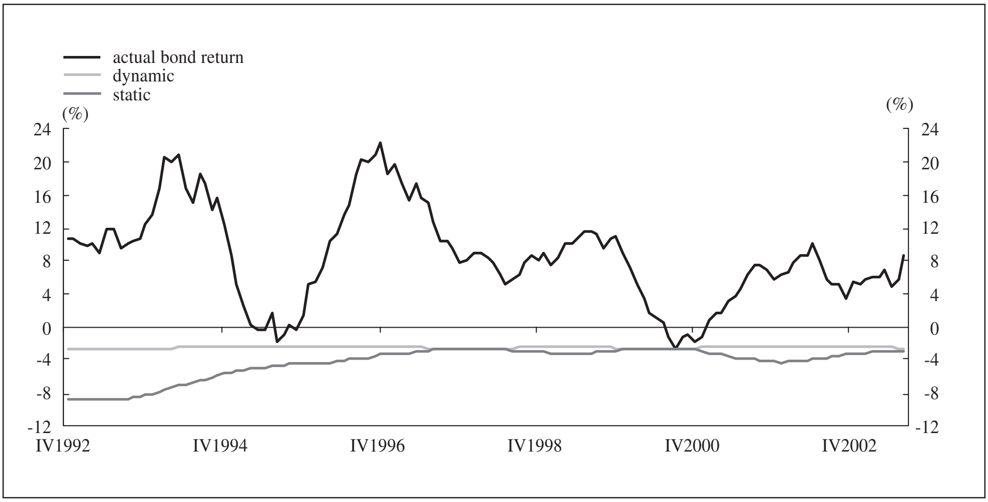

We can conclude from the graph that both approaches for benchmark settings are sufficiently robust to withstand changes in the index. It can even be said that the static benchmark is perhaps too robust, as the historical data do not justify such a large reserve.

There is another, perhaps more natural, viewpoint from which we can examine the robustness of both benchmarks. One may ask how large the horizontal shift of the yield curve would have to be for the benchmark to result in a zero return. To answer with a simplified but for illustrative purposes satisfactory example, let us assume that the return curve is flat, and let us replace the sector 1-10y, which represents the bond index, with a 5-year bond. We then find what change (over a 12-month period) in the 5-year bond return, ΔY_{5y}, would make the portfolio overall return decrease to zero.

The equation for the overall return is:

$$R = w_{3m,i} \cdot R_{3m,i} + 0{,}1 \cdot R_f + (0{,}9 - w_{3m,i}) \cdot R_{5y}$$

The relation for the modified duration can be used: $\Delta P_{5y} = -D \cdot (\Delta Y_{5y}) \cdot P_{5y}$, where P_{5y}, or Y_{5y}, or D is the price, or return, or the modified duration of the 5-year bond. The change in return caused by a change in the 5-year return, ΔY_{5y}, is indicated as r_d, i.e. $r_d = \Delta P_{5y}/P_{5y}$. Then, upon the assumption that the return curve is flat, it can be:

$$R_{5y} = R_{3m} + r_d,$$

After insertion,

$$R_{5y} = R_{3m} - D\Delta Y_{5y.}$$

Now we seek a value of ΔY_{5y} for which the following holds true

$$w_{3m} \cdot R_{3m} + 0{,}1 \cdot R_f + (0{,}9 - w_{3m}) \cdot (R_{3m} - D \cdot \Delta Y_{5y}) = 0\,.$$

After modifying this for the static case, we get

$$\Delta Y_{5y} = \frac{R_{3m}}{D \cdot (0{,}9 - w_{3m})},$$

and for the dynamic

$$\Delta Y_{5y} = \frac{R_{3m}}{D \cdot (0{,}9 - \frac{\gamma}{R_{3m}})}\,.$$

The following graph shows the maximum permissible and actually monitored ΔY_{5y} for the static benchmark and dynamic benchmark. The *static* or *dynamic* lines demonstrate how approximately large the change would have to be in the yield of a 5-year bond over a 12-month period for the benchmarks to record zero or negative return in a 12-month period.

The graph implies that both methods are sufficiently robust to withstand possible shifts in the yield curve. For the last data, there is a reserve in the form of a shift of up to 150 basis points over twelve months. The border was exceeded slightly in 1995 and 2000 as a result of simplified assumptions, particularly concerning the flat yield curve and approximation of the return from the bond index. In reality, this border is reached, as Chart 4 shows, as it is set.

Chart 8: Permissible versus actual shift in 5-yr yields

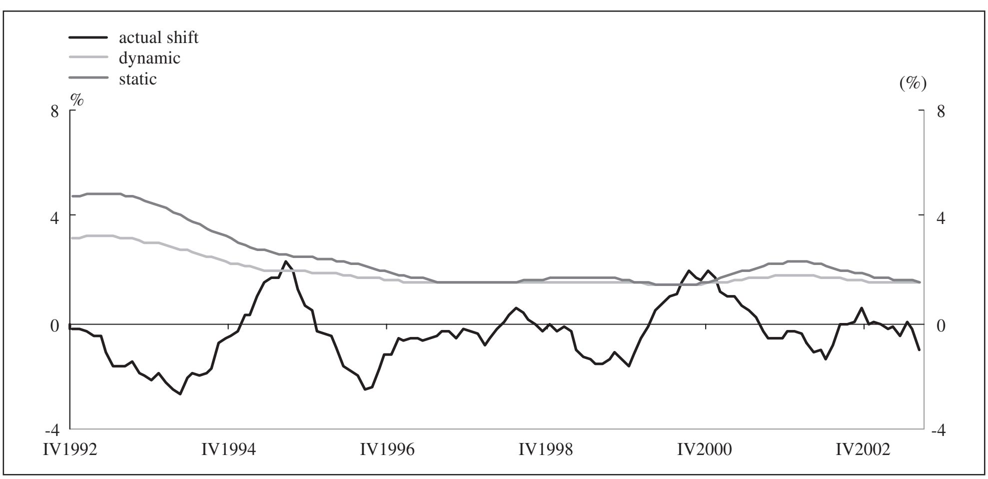

4 Conclusions

4.1 Currency allocation

Discussion regarding the currency allocation was focused on the two following points of view.

The first mainly regards CNB's accounting in the domestic currency. It examines the impact of the currency composition on the koruna value of the FX reserves. It was concluded that holding no more than 25% FX assets in USD and the rest in EUR would fit the requirements for the lowest koruna volatility of returns. A positive volume of USD in FX reserves is in line with one of the decisions to have own FX reserves within the future EMU participation. The "koruna accounting" approach however serves only as a guideline, as it examines the accounting performance of the reserves on the run. The ambiguous consistency of its assumptions with the current economic theory was also pointed out during our internal discussion.

The final decision about the currency structure thus takes place when the assets and future flows are finally settled against future liabilities upon EMU entry, whereby the accumulated profits/losses are to be realised (i.e. the "liabilities match" approach). Here, the central bank's aim is to have the lowest possible exchange rate risk. A loss arising from a speculation-motivated currency composition will lead to the position being closed, whereas the same loss arising from the same allocation, but motivated by the hedging of future liabilities, will lead to a continuation of the position. For this reason, the "liabilities match" approach was prioritised.

The following key flows and liabilities were considered:

- There will be EUR-denominated flows consisting of CNB's FX market operations during the EMU accession period. These include interventions either in favour or against the domestic currency, and any client operations on the FX market.
- The major EUR-denominated liability is expected to arise from the domestic repo operations.[9] These are used to sterilise the domestic currency after any central bank FX market operations. These liabilities would turn into EUR-denominated liability towards either domestic commercial banks or towards TARGET.

  As explained in the text, the sum of the two latter *future* liabilities shall be roughly equal to the outstanding amount of the *current* domestic repo operations. This therefore represents the projected net size of the two major future EUR-denominated liabilities.
- There will also be a future liability arising from the need to hold a sufficient amount of non-EUR currencies for CNB's calls after EMU entry.

It was concluded that all the above liabilities would be secured towards currency risk to the same extent as if the value of the dollar part of the reserves were to be linked to the USD value of the koruna monetary base. The resulting currency structure of the foreign reserves, with the dollar part maintained at the level of the monetary base, and the remainder in euro, has been therefore approved as the optimal one.

4.2 Benchmark duration

The way to set appropriate duration differs from the above approach. It is rather based on statistical analysis of historical data. Moreover, it is believed that economically higher duration means higher expected return. The natural inability to define an exact limit of

[9] Maintained by CNB to sterilise the domestic currency, as a result of CNB's monetary policy.

acceptable loss for the use of VaR has been partially compensated by an approach that allows the bank to react actively to interest rate changes. The dynamic benchmark provides the possibility of changing the benchmark structure and thereby of changing the duration of the benchmark. Extremely important and relevant in this regard is the fact that the dynamic benchmark results in a higher return than the static benchmark while fulfilling the same conditions that are applied on the return, i.e. that there should not be any negative return from the benchmark over the 12-month period.

The frequency at which the weights of individual components of the dynamic benchmark are reset depends on the volatility of short-term interest rates. During periods of low volatility, there is obviously no need to alter the setting of the benchmark. Alternatively, a certain threshold at which the benchmark will be changed can be set, e.g. in the form of a minimum change in the duration of the benchmark when a change takes place.

Functions other than the above-mentioned $w_{3m}= \gamma/R_{3m}$, naturally provide slightly different results. If, for example $w_b= \gamma_2 * R_{3m}$, (the constant γ_2 is calculated analogously), then the subsequent benchmark would have a significantly more volatile duration than a benchmark calculated by means of the first function. It is clearly also possible to alter the weights of all the components, not just two of them. For example, the weights could be based upon returns from the previous period. We could even consider altering the number of components in the benchmark by, for example, including the index of bonds of longer maturity. But the calculations would remain analogous to a large extent.

17

Foreign reserves risk management in Hong Kong

Clement Ho

Abstract

This paper aims to elaborate the framework and the risk management practices of the reserves management in Hong Kong. It sets out the details of the monitoring and control process of major risk parameters such as market, credit, liquidity and operational risk. It also describes the investment process in terms of establishing the investment benchmark as the long-term strategic asset allocation and the use of risk in tactical deviations for shorter to medium-term investment decisions. The investment benchmark is derived from a risk-based global CAPM equilibrium model anchored in the modern portfolio theory of efficient markets, and the optimisation process takes into account the investment objectives and constraints for managing the reserves in Hong Kong. Finally, the framework of investment performance measurement will also be discussed.

1 Introduction

Today, risk management of foreign exchange reserves has become increasingly challenging for central bankers. This is partly due to more common market recognition that a well-established risk management framework is the most critical element of successful reserves management. If it is conducted properly and effectively, we will not only meet various investment objectives, but also strengthen public confidence in the government's monetary policy. Responding to these challenges, central bankers are now more aware of the importance of adopting best market practices in building their risk management regimes. With the development of information technology, they are applying more sophisticated risk management tools to measure and control their risk exposure.

For the risk management of foreign reserves in Hong Kong, we have strengthened our diagnostic and analytical capabilities during the past few years to keep pace with the changes both in the marketplace and in the investment approach of the Exchange Fund. In 1999, a new investment benchmark and tactical deviation limits were implemented, which better defined the strategic asset allocation and the use of risk in the investment process. To measure and control the market risks inherent in the portfolios under different adverse market conditions, we have adopted the risk measures of "Value-at-Risk" (VaR) and scenario stress testing.

In the context of risk management, performance attribution is designed to test the optimality of asset allocation and distribution of active risk at the macro level, and to identify manager skills at the portfolio level. In the Hong Kong Monetary Authority (HKMA), the performance attribution analysis is divided into two main categories of active management effects, namely the overall asset allocation effect and the portfolio active return. Performance of the externally and internally managed portfolios is periodically evaluated against respective investment benchmarks. Performance attribution analyses are also undertaken against the well-tested attribution models in bond and equity markets to achieve a better manager mix of different styles and skills.

This paper is structured to define the framework and the application of risk management in the reserves management of Hong Kong. Section 1 provides a brief history and the objectives of the investment management of the Exchange Fund assets, division of duties and delegation and control of authority. Section 2 sets out the details of the risk management framework including the monitoring of major risk parameters such as market, credit and operational risk. Section 3 describes the investment process in terms of deriving the strategic asset allocation and the use of risk in tactical deviations from the benchmark portfolios. Section 4 extends the risk management framework to describe the framework of investment performance analysis.

2 Developing a sound governance and institutional framework for reserves management in Hong Kong

The history of the exchange fund of Hong Kong

The Exchange Fund of Hong Kong (the Fund) was established in 1935 by the Currency Ordinance (later renamed the Exchange Fund Ordinance). The purpose of the Fund was to provide backing for note issue and was then held in gold, silver and British pounds. The role of the Fund was further expanded in 1976 to include the management of fiscal reserves for the government. From that year onwards, the government began to transfer its fiscal reserves to the Exchange Fund. In addition, the merger of the ex-Land Fund assets into the Exchange Fund in November 1998 further boosted the strength of the reserves. As at December 2002, the foreign currency assets of the Fund amounted to US$ 111.9 billion, the fifth largest in the world.

The linked exchange rate regime in Hong Kong

Hong Kong's Linked Exchange Rate System was established in 1983. It is in essence a Currency Board Arrangement, which requires both the stock and the flow of the Monetary Base to be fully backed by foreign reserves. Any change in the size of the Monetary Base has to be matched by a corresponding change in the foreign reserves. In Hong Kong, the Monetary Base comprises the following components:

a. Certificates of Indebtedness (as backing for banknotes) and coins issued;
b. the sum of balances of banks' clearing accounts (Aggregate Balance) maintained with the HKMA for the purpose of clearing and settling transactions between the banks themselves, and also between the banks and the HKMA; and
c. the outstanding amount of Exchange Fund Bills and Notes.

Under the Currency Board system, the stability of the Hong Kong dollar exchange rate is maintained through an automatic interest rate adjustment mechanism. When there is a decrease in demand for Hong Kong dollar assets and the Hong Kong dollar exchange rate weakens to the convertibility rate, the HKMA stands ready to purchase Hong Kong dollars from banks, leading to a contraction of the Monetary Base. Interest rates then rise, creating the monetary conditions conducive to capital inflows so as to maintain exchange rate stability. Conversely, if there is an increase in the demand for Hong Kong dollar assets, leading to a strengthening of the exchange rate, banks may purchase Hong Kong dollars from the HKMA. The Monetary Base correspondingly expands, exerting downward pressure on interest rates and so discouraging continued inflows.

Statutory role of the Exchange Fund

The statutory role of the Exchange Fund, as defined in the Exchange Fund Ordinance, is to safeguard the exchange value of the currency of Hong Kong. Its functions were extended with the enactment of the Exchange Fund (Amendment) Ordinance in 1992 by introducing a secondary role of promoting the stability and integrity of the monetary and financial systems, and enhancing Hong Kong's role as an international financial centre. Another important objective is to maintain the long-term purchasing power of the reserve assets for the benefit of Hong Kong as a whole.

The day-to-day management of the Fund is vested in the HKMA, which was established on 1 April 1993 by merging the Office of the Exchange Fund with the Office of the Commissioner of Banking. It was formed to ensure that the central banking functions of maintaining monetary and banking stability can be performed with a high degree of professionalism and continuity, in the lead up to 1997 and beyond, in order to command the confidence from the people of Hong Kong and the international financial community.

Delegation of investment management

The authority for the investments held by the Exchange Fund rests with the Financial Secretary of the Hong Kong Special Administrative Region, who consults the Exchange Fund Advisory Committee (EFAC) to establish the long-term strategic direction of the Fund. The day-to-day management of the Fund is performed by the Reserves Management Department of the HKMA. In carrying out its responsibilities, the Reserves Management Department operates under the delegated authority of the Financial Secretary and within the investment and credit exposure policies approved by the EFAC.

The delegation of investment authority for managing the Fund is achieved through a three-level framework, covering the basic types of investment decision:

- Long-term strategic decisions controlled by the EFAC, in which the Financial Secretary is the ex officio chairman;
- Medium-term market decisions delegated to senior management down to the Executive Director level of the HKMA; and
- Short-term trading decisions delegated to portfolio manager levels.

Chart 1: Investment decision-making structure and delegation of authority for the management of the Exchange Fund

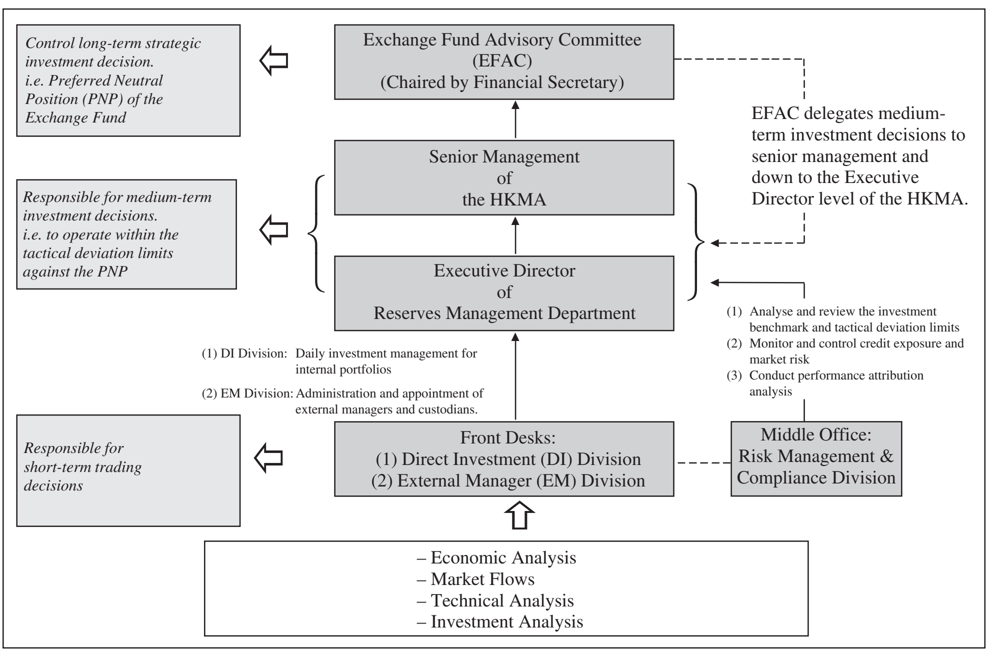

This three-level framework of investment decisions is reviewed every year to examine the system of delegation, limits and reporting schedules for the management of the Exchange Fund assets. Chart 1 below summarises the investment decision-making structure and delegation of authority for the management of the Exchange Fund assets.

Transparency and accountability

Hong Kong's Legislative Council defines and limits the HKMA's powers and responsibilities. This is achieved by setting out the Monetary Authority's accountability, through the Financial Secretary and the Hong Kong SAR government, to the Legislative Council under Article 64 of the Basic Law. Furthermore, the HKMA recognises a broader responsibility to the community of Hong Kong, and a duty to promote understanding of our role and objectives, and to keep ourselves informed of community concerns and to be open to public debate.

To this end, the HKMA pursues a policy of transparency and accountability. This policy has two main objectives:

(i) to keep the financial industry and the general public as fully informed of the HKMA's work as possible; and

(ii) to ensure that the HKMA is in touch with, and responsive to, the community that it serves.

Over the past several years, the HKMA has pursued various measures to be in line with best market practices and international standards with the aim of increasing transparency and disclosure requirements. Annex A details the major milestones and steps that the HKMA has taken over the past decade to enhance transparency and disclosure standards.

3 Risk management framework

At a macro level, risk management is used to control risks arising from investment activities of the Exchange Fund in order to achieve investment objectives and to support policy requirements. At a micro level, the primary purpose of risk management is to ensure that the Exchange Fund is invested within the acceptable limits and that the investments will not result in excessive losses. In developing its risk management framework, the HKMA draws upon the best market practices and international standards in the marketplace. For day-to-day investment operations, the framework serves to fulfil the following two main functions:

– to identify and assess the financial and operational risks; and
– to restrict these risks within acceptable parameters and levels.

The risk management of the Exchange Fund begins at the board level. The EFAC defines and controls the overall investment strategy and risk appetite by the Exchange Fund Investment Policy (the Investment Policy) and the Exchange Fund Credit Exposure Policy (the Credit Policy). The Investment Policy specifies the strategic asset allocation by means of the Preferred Neutral Position (PNP) framework (Section 3 on the Strategic Allocation of Exchange Fund Assets will elaborate in detail the derivation of the PNP and the allocation of risk for tactical positioning). In addition, reports on the current asset allocation of the Fund and compliance with the investment policy are submitted to EFAC members for their review at regular monthly meetings.

The Risk Management and Compliance Division of the HKMA is responsible for day-to-day monitoring and operational control of six major types of financial and operational risks, namely:

– Market Risk – Credit Risk – Liquidity Risk – Currency Risk – Settlement Risk – Operational Risk.

Market risk

Market risk is defined as losses due to adverse changes in the value of financial assets. The control of market risk requires timely and accurate information about the current value of the Fund's assets and liabilities, as well as off-balance-sheet positions.

For the measurement of market risk exposure, a weekly asset allocation report is prepared to reflect the mark-to-market value of the Exchange Fund. The report also includes the maturity profile by asset market and currency type. A weekly asset allocation meeting, with the participation of senior executives and reserves management staff, is held to review the short to medium-term investment strategies for the overall Exchange Fund and the allocations to both internal and external managers.

The overall market exposure of the Fund is effectively controlled by the Exchange Fund's Preferred Neutral Position (PNP) and tactical deviation limits for each asset class, duration and currency type. This ensures that the exposure of the fund is in line with the investment policy established by the EFAC.

In addition to the weekly asset allocation report, market risk is measured by Value-at-Risk (VaR). Using a well-tested risk management model, both absolute and relative (to benchmark) VaR are monitored regularly based on a 95% confidence interval for a 1-month horizon. The relative VaR is analogous to the measurement of tracking error against benchmark return. The above VaR calculation takes into account the price volatility of all asset class, market segmentation and correlation across markets given a specified time horizon and decay factor. In addition, VaR in both US dollar and percentage terms is monitored to ensure that the Fund is not unduly exposed to market risk at any point in time.

VaR measurement quantifies the maximum potential loss in a portfolio under normal market conditions and a given confidence level. Stress testing will also be performed to quantify the price behaviour of the Exchange Fund under extreme market conditions. The test evaluates the price impact of a simultaneous recurrence of the worst equity, bond and currency market crashes in recent history on the Exchange Fund.

Pros and cons of using VaR

The benefits of VaR vis-à-vis the traditional position controls are seen in various aspects. Firstly, VaR directly measures risk in terms of potential portfolio loss instead of notional amount of positions where the inherent risk, i.e. the expected loss, is not quantifiable and clear to the reserves managers. Secondly, VaR represents comparable risk measures and limits across different markets. Thirdly, VaR can facilitate the aggregation of risks across markets or asset classes. This provides flexible ways to dissect and to identify the portfolio effects at both overall fund and individual market levels. The final advantage of using VaR is the reality check that enables the comparison of risks and the actual results.

The shortfall of VaR is that it only measures the potential losses under normal market conditions. However, there are times when market movements can be extremely volatile owing to unanticipated political or economic events. Another possible approach involves the use of extreme value theory (EVT), which augments the standard VaR with an estimate of the expected shortfall based on stochastic simulations. This method better estimates the tail of the distribution for the purpose of managing risk because it is always the tail of the distribution that wipes out the overall portfolio performance. Nonetheless, the static mean-variance framework is powerful and easy to compute, although VaR is not a perfect solution for price risk control and measurement. The dynamic extension to new risk management tools for measuring market risk such as EVT is conceptually important and empirically relevant but not widely applied yet in the marketplace. Unless there are tremendous increases in simulation technology and data availability, the actual industry-wide application of these new risk management tools for reserves management will take some time.

In order to complement the shortcomings of VaR, we perform stress testing to quantify the market risks inherent in the portfolios for extreme market conditions based on a reasonably long period of time. Examples of extreme market conditions include the stock market crash in 1987 and the currency crisis in 1994. Stress testing is employed to test the impact of a simultaneous occurrence of the worst equity, bond and currency market crashes in the past 20 years on the Exchange Fund's assets. The results of VaR and stress testing are submitted to the EFAC for review on a regular basis.

Credit risk

Credit risk is associated with the possibility that a counterparty will not fulfil its obligation to repay. For the Exchange Fund, credit risk is controlled by establishing an authorised list of counterparties and issuers based upon their individual credit ratings and the analysis of respective financial ratios. A rule-based method is used to determine the credit limit for each counterparty and issuer. Regular monitoring of their financial positions and our exposure to them is required to avoid unduly large credit exposure.

With the delegation from the EFAC, the Credit Review Committee (CRC) was established to serve the following functions:

- Monitor credit risks set by the Credit Policy;
- Review the Credit Policy to meet any new investment requirements and to control counterparty or issuer credit risks; and
- Determine and evaluate the credit limits for counterparties including banks and securities houses, issuers and countries.

The Risk Management and Compliance Division is responsible for the credit analysis and assessment based on market-available information. The control mechanism is segregated into two major areas:

- credit assessment by standardising the analytical framework in terms of economic and financial performance, debt ratio, liquidity, capital, asset size and rating assessment by international credit agencies; and
- credit line allocation between internal and external managers through the CRC.

A risk-weighted exposure (RWE) limit is set for each authorised counterparty. The limit restricts the risk-weighted credit exposure to a counterparty arising from all kinds of investment activities. Different risk weighting is assigned to each allowable investment instrument depending on its market risk or price volatility, and the investment horizon.

For internally managed funds, an online pre-deal exposure checking system has been installed. Dealers have to check the limit availability before they deal with any counterparty. For externally managed funds, it is more practical to appoint the custodians to perform the investment and credit compliance checking functions based upon the limits allocated to the external managers. The counterparty and issuer credit limits for external investment managers are specified in the investment guideline of each mandate (the Investment Guideline), which is approved by the CRC for the credit limit allocation. Based on the Investment Guideline, the custodian manager will perform compliance checking for each transaction and cumulative positions. The External Manager Division will report breaches of credit limits to the CRC for sanctioning and other actions if deemed necessary.

Liquidity risk

For the Exchange Fund, the primary liquidity needs arises from Currency Board operations. Under the Convertibility Undertaking (CU) arrangement, the HKMA undertakes to sell USD/HKD at a rate of 7.8000. Any Authorized Institution[1] which maintains an account with the HKMA and has access to the domestic Real Time Gross Settlement System can buy USD/HKD at the CU rate.

Another primary need for liquidity arises from the Hong Kong Government's withdrawal of its fiscal reserves placed with the Exchange Fund. The Fund is required to provide cash on a timely basis for the repayment of principal and the payment of interest due on transfers from the fiscal reserves and other interest-bearing liabilities of the Fund.

To ensure that the Fund maintains sufficient liquidity, the EFAC has stipulated liquidity limits as follows:

- A minimum percentage of foreign currency assets that can be converted into cash within two and five working days;
- A maximum percentage of the Fund's assets that is placed for fixed-term deposits and a maximum percentage for deposits over a 1-month horizon; and
- An issue concentration limit, i.e. the holding amount of a particular bond issue as a percentage of particular issuance size or issuer's programme size. This is to ensure efficient liquidation without disrupting the markets.

Currency risk

Currency risk refers to the possibility of losses arising from adverse foreign exchange movements. The overall currency risk for the Exchange Fund is determined by combining the benchmark currency mix and corresponding deviation limits. The PNP of the currency mix strategically dictates the long-term currency exposure for the overall Exchange Fund's assets. This takes into account the long-term return and risk profile of currency as an independent asset class together with all investment constraints such as the maintenance of adequate US dollar assets to provide full backing for the HK dollar Monetary Base under the Currency Board system.

The Backing Portfolio can only be invested in liquid US dollar assets. The currency exposures of other portfolios are subject to their underlying investment benchmarks and permissible deviation limits.

Settlement risk

Settlement risk refers to the scenario that the counterparty fails to make payment after we have delivered our payment. In order to mitigate settlement risk, the HKMA has standardised settlement instructions with delivery versus payment (for bonds and equities), and encouraged foreign exchange transactions through the established real-time gross settlement (RTGS) platforms.

For externally managed assets, the HKMA will only use reputable and reliable master custodians to centralise and control the settlement workflow by portfolio types to improve

[1] As defined in the Banking Ordinance, Authorized Institutions are those which can conduct banking business in Hong Kong.

operational efficiency and mitigate settlement risks. The major criteria of selecting a master custodian to achieve higher quality and efficient custodian services are:

- Financial strength and credit rating;
- Reliability, flexibility and user-friendliness of the technological platform;
- The level and quality of local support; and
- Expertise and track record in custody business.

Operational risk

Operational risk refers to unexpected losses resulting from inadequate controls, human fraud and system errors. In the HKMA, there is a proper segregation of duties between settlement staff and dealing functions. This can achieve better control by avoiding potential outright fraud or collusion among staff.

The HKMA also conducts contingency drills on a regular basis to ensure that the dealing and back office processing functions can operate smoothly in the back-up sites even if the main office is not accessible owing to system breakdown, power failure or disaster.

4 Strategic allocation of assets

Objectives of the Exchange Fund of Hong Kong

The strategic asset allocation of the Exchange Fund should at all times be able to meet the objectives of the Fund, which are defined as follows:

- to preserve capital;
- to ensure that the entire Monetary Base will be at all times fully backed by highly liquid short-term US dollar-denominated securities;
- to ensure sufficient liquidity for the purpose of maintaining monetary and financial stability; and
- subject to the above, to achieve an investment return that will preserve the long-term purchasing power of the assets.

Portfolio segregation

The fund is divided into two portfolios: the Backing Portfolio and the Investment Portfolio. The Backing Portfolio was established for the operations of the Currency Board Account which linked the exchange value of the Hong Kong dollar to the US dollar. Under the Currency Board arrangement, Hong Kong's Monetary Base is at all times fully backed by the assets of the Backing Portfolio, which are composed of primarily short-term, highly liquid US dollar interest-bearing securities.

The Investment Portfolio is predominantly invested in OECD bond and equity markets to preserve the long-term purchasing power of the Fund's value for the future generations of Hong Kong. The Investment Portfolio is further divided into an internally managed and an externally managed investment portfolio.

Construction of the strategic investment benchmark

The strategic asset allocation of the Exchange Fund is defined by the Preferred Neutral Positions (PNP), which is approved by the EFAC. The PNP represents a benchmark for the long-term asset allocation mix of the Exchange Fund with respect to (a) currency distribution, (b) allocation of asset classes, and (c) the maturity profile of the assets. In addition, the EFAC will approve tactical deviation limits against the PNP for the HKMA to operate for short to medium-term investments of the Exchange Fund assets.

The PNP is derived from a global capital asset pricing model based on the Markowitz framework. This market equilibrium approach uses a global market capitalisation-weighted portfolio as a neutral reference point. The approach is able to blend investor views, if any, with global market equilibrium expected returns to determine the optimal asset allocation mix, which is a set of deviations from the global market capitalisation-weighted portfolio. Given no specific investor views on markets, such a market equilibrium optimisation model applies global market equilibrium returns. This is better than the traditional mean variance approach of using a complete set of expected returns on all assets during the optimisation process. Furthermore, under the traditional approach, a small shift in the expected return for a market can cause drastic swings in the weight of an asset market at optimal level, and hence give rise to extreme portfolio weightings.

Determination of covariance matrix and permissible markets

As a first step to determine the strategic asset allocation, we evaluate the full set of allowable asset classes based on our specific objectives and constraints. Available market indices are used to represent global asset markets including equities, bond, cash and currency. Different time periods and methods are considered to estimate a long-term asset market volatility and correlation matrix. Such a risk structure is assumed to hold constant to derive expected portfolio volatility. In addition, various asset allocation factors including currency hedging decision and equity risk premium are considered during the process.

The investments of the Exchange Fund are expected to operate in well-established markets with sufficient breadth and depth to ensure that transactions can efficiently be absorbed in the market without undue effects on liquidity and the availability of funds. This is particularly the case for central banks including the HKMA, since most transactions are typically in decent size. We assess the liquidity of each market and instrument based on the following factors:

(i) bid/offer spread in both normal and crisis conditions.
(ii) dealing size in both normal and crisis situations.
(iii) total portfolio holdings as a percentage of daily market turnover.
(iv) availability of repo market for each instrument type.

To quantify the investment objective and constraints and to transform them into "critical investment performance criteria"

For the optimisation process, the investment objectives and constraints are to be quantified:

- To preserve capital, it is required to set a high probability level that the returns in the short, medium and long term will not be negative;
- To ensure that the entire monetary base is at all times fully backed by liquid US dollar assets, it is required to set a minimum percentage of total portfolio holdings in liquid US dollar assets;
- To preserve the long-term purchasing power of assets, it is required to set a reasonably high probability that the expected benchmark portfolio return shall exceed the domestic inflation rate; and

– To measure the downside risk below the long-term domestic inflation level, it is required to set the annualised shortfall in return against the long-term domestic inflation rate at an acceptably low probability level.

To determine the strategic asset allocation by testing against the "critical investment performance criteria"

Institutional investors will typically set a target portfolio risk for the investment benchmark. Instead of using a specific portfolio risk target, we would test the optimal asset allocation against a set of "critical investment performance criteria", which effectively reflects the control of overall portfolio risk parameters in order to meet investment objectives and constraints.

Although the PNP represents an optimal asset mix for long-term investment, it is typically reviewed on an annual basis to reflect the dynamic changes in market conditions and investment constraints and requirements. The investment benchmark, or PNP, in terms of currency and asset distribution (as at December 2002) is as follows:

Bonds	80%
Equities	20%
Currencies	80% (USD bloc)
	15% (European bloc)
	5% (JPY)

Use of indices in reserves management

The construction of the strategic investment benchmark for the Exchange Fund of Hong Kong uses indices extensively. Factors such as liquidity, volatility, the cost of replication and representation of market are considered in the selection of the appropriate indices for the market allocations at total fund level and for each type of investment manager at individual portfolio level.

The use of benchmark indices serves two major purposes. On an ex ante basis, we use a market index to make asset allocation and to choose investment styles. Such asset allocation is effected through the construction of a passively managed portfolio that replicates the selected market index. On an ex post basis, the benchmark index is used for performance attribution analysis to measure the performance of investment managers on a risk-adjusted basis. The selection of appropriate indices is critical because it determines not only the investment style but also the level of risks for the Fund.

For the bond allocation, we need to be particularly aware that the conventional global government bond indices are on a pure market capitalisation-weighted basis. This implies that any reserves fund using such indices will be increasingly exposed to countries with rising government debt. On the one hand, it can reduce rebalancing costs since country weightings will logically be adjusted relative to market development in the countries concerned. On the other hand, this may not satisfy the investment need from the perspectives of diversification of credit risk and expected investment performance for markets with rising country debts.

Risk budgeting and the use of active risk

In the HKMA, risk budgeting is the allocation of portfolio active risks, i.e. tracking error limit, to different asset classes, markets or investment managers in order to generate active return. The determination of tracking error limit over the PNP is approved by the EFAC and reviewed on a regular basis. The use of active risk is derived from two aspects, which include the deviation from the neutral allocation to different asset classes and markets, and the allocation of risks to active mandates over their individual portfolio benchmark.

A risk-based approach using risk contribution and marginal volatility of underlying composite asset markets is used to calculate the permissible tactical deviations for each market. This method is essentially a VaR approach used by institutional investors to control the risks arising from trading positions. All things being equal, the most volatile asset will logically be given the least permissible deviation.

In addition, active risks to each asset class and market will be allocated with respect to their expected returns. In this regard, we allocate risk to active mandates with respect to their investment performance relative to other styles and passive mandates.

5 Performance measurement framework

The objective of performance attribution analysis is to test the optimality of asset allocation and distribution of active risks. At the overall fund level, the attribution of active return is analysed at two broad levels, namely the overall fund allocation effect and portfolio active return. For the latter, we employ market-tested performance analytical models to evaluate whether, where and how the fund manager has added value. We measure the out/under-performance of a portfolio relative to its benchmarks, and decompose the active return into various market attributes in corresponding bond and equity market models.

The overall fund allocation effect measures the impact on portfolio return owing to the overweight or underweight in different asset markets against the benchmark-neutral allocations. The overall fund allocation effect can be further divided into four contributory factors:

(i) Asset allocation effect – the selection of good performing asset classes (e.g. by allocating more funds into fixed income instead of equity in 2002).
(ii) Country weighting effect – the selection of outperforming countries within a particular asset class (e.g. for equities, allocating more funds to Hong Kong instead of the US market in 2002).
(iii) Currency effect – the selection of a better-performing currency against the benchmark allocation (e.g. euro versus the USD in 2002).
(iv) Benchmark adjustment effect – the return difference owing to different market indices chosen at individual portfolio and overall fund level (e.g. the return difference between Lehman Brothers MBS Index assigned for an MBS type bond manager versus JP Morgan Global Government Bond Index as the benchmark for bond allocation at the total fund level).

Portfolio active effect is the out/under-performance at portfolio levels relative to individual benchmarks. Since 1999, the HKMA has begun to use market-based multi-factor models to analyse the active returns of fixed income. We have further extended the analysis into active equity portfolios since mid-2002. The monthly attribution analysis is based on the daily portfolio positions. The results are submitted to management to monitor and review the performance of individual fund managers. In evaluating an investment manager, we typically rely on a long-term horizon of three to five years' portfolio performance instead of a short-term investment record.

Market-based three-factor model for analysing bond portfolios

For fixed-income attribution analysis, the HKMA uses a market-based three-factor model which principally captures the term-structure risks: yield curve parallel shift, slope and curvature changes. The model has several distinct advantages in that firstly, it corresponds directly to the types of term-structure decisions for investment in bonds. Secondly, the factors are pre-specified, open with published formulas, and relatively easy to understand. Thirdly, the factors used in this model are consistent over time, markets and cycles. However, the explanatory power of this three-factor model sometimes may not be as high as that of a principal component approach. Nonetheless, the latter alternative is inconsistent with our commonly recognised equilibrium models that imply continuous term-structure shifts in maturity. This kind of principal component approach applies esoteric factors that have no clear economic interpretation. In addition, different bond markets or time horizons may use different factors which make the performance results difficult to evaluate.

Multi-factor model for equity portfolios

For equity analysis, we also use a multi-factor model which decomposes active returns into a number of fundamental attributes in three dimensions: size, style and momentum. Size can be measured by the market capitalisation of the company. Style can be measured by price/earnings ratio, book/price ratio and dividend yield. Momentum factors include net earnings revisions, price reversals and earnings torpedo (susceptibility of stocks with high future expected earnings to negative surprises). Apart from the fundamental attributes, the historic beta and the industry classifications are also included as factors in the model.

On top of the above factoring model analysis, we maintain regular return-based regression analysis in the following parameters for the performance of individual investment managers:

– Alpha (α): the ability of the fund manager to select securities that out-perform the market;
– Down-market beta (β): measures the portfolio's exposure to the benchmark in a down market;
– Gamma (γ): the fund manager's skill in managing market exposure against the benchmark at times of down market and up market; and
– Information Ratio: measures the risk-reward trade-off given the investment manager's added value and excess volatility over the benchmark.

6 Conclusion

The case study presented in this paper illustrates the risk management approach and the evolution that has taken place in the area of reserves management in Hong Kong. Similar to many central banks, the basic traditional objectives of reserves management in Hong Kong are essentially security, liquidity and investment return. In addition, the reserves management strategies are used to support our own specific policy environment, particularly the Linked Exchange Rate Regime adopted in 1983.

The establishment of a sound governance and institutional framework aims at accomplishing a wider use of sophisticated risk management techniques and hence an efficient reserves management. In the HKMA, it is evident that the internal governance structure for reserves management reflects a clear division of responsibilities together with a well-defined delegation of authority from the EFAC down to the Reserves Management Department of the HKMA. The daily monitoring and operational risk control arising from the investment of Exchange Fund assets together with performance attribution analysis are performed by the Risk Management & Compliance Division in the HKMA.

We face different kinds of risks in managing the Exchange Fund. For market risk, we measure and control these risks by means of Value-at-Risk and stress testing. We also prudently control our credit risk exposure by establishing credit limits and constantly reviewing our counterparties. Operational risks can be mitigated by a proper segregation of duties and performing contingency drills.

To control the market risk of the Exchange Fund, the EFAC approves the adoption of the PNP which dictates the long-term strategic asset allocation mix and hence the market exposure. In addition, the EFAC will approve the tactical deviation limits for the HKMA to operate for short to medium-term investments. The PNP is derived from a global CAPM model based on modern portfolio theory of efficient markets and is well-tested and widely accepted by the international investment community. The optimisation process has taken into account the objectives and constraints such as the liquidity requirement of the Exchange Fund. For the tactical deviation limits, we adopt a risk budgeting concept to allocate active risks systematically across allowable markets and among managers to facilitate asset allocation.

The objective of performance attribution analysis is to test the optimality of asset allocation and the distribution of active risks. The analysis is used to evaluate the performance at the overall fund and individual portfolio levels against investment benchmarks. In addition, it can help facilitate the asset reallocation to accomplish enhanced portfolio performance on a risk-adjusted basis.

References

Fabozzi, Frank (1998): Handbook of Portfolio Management, pp 27-30, 57-74.

Hong Kong Monetary Authority (2002): Annual Report 2002, pp 78-83.

Hong Kong Monetary Authority (1997): Money and Banking in Hong Kong (Volume 2), pp 96-100.

Litterman, Bob (2003): Modern Investment Management, pp 36-109.

Pearson, Neil (2002): Risk Budgeting, pp 3-11.

Appendix A

Measures to increase transparency of disclosure

Prior to 1992, the accounts of the Exchange Fund were confidential. However, in keeping up with the commitment to greater transparency in disclosure, the government started publishing the accounts of the Exchange Fund annually in 1992. Biannual accounts have been published since June 1995, and headline figures for the foreign exchange reserves have also been released monthly since January 1997.

In recent years, we have done a great deal to advance the policy of transparency and accessibility. The major milestones in enhancing the transparency of disclosure are displayed in Table 1 below. This applies to technical matters such as the increased transparency of the Currency Board system, the daily publication of the size and composition of the Monetary Base, and the virtually real-time publication of the Aggregate Balance of the Banking System. In addition, we extend the information available for general public consumption, whether in print, through the press, on the Internet, or in our programmes of educational briefings, seminars and large-scale exhibitions.

Table 1: Major Milestones in enhancing the transparency of disclosure

Year	Measures
1992	Started publishing the accounts of the Exchange Fund annually
June 1995	Started publishing the accounts of the Exchange Fund biannually
Sept 1995	Started publishing quarterly figures on foreign currency assets
Jan 1997	Started monthly disclosure of foreign exchange reserves
Jan 1999	Started publishing monthly International Reserves in accordance with IMF's SDDS Started publishing monthly data on Analytical Accounts of the Central Bank in accordance with IMF's SDDS Started publishing monthly Abridged Exchange Fund Balance Sheet and Currency Board Accounts
Apr 2000	Started publishing Template on International Reserves and Foreign Currency Liquidity in accordance with IMF's SDDS

The considerations of market sensitivity, commercial confidentiality and statutory restrictions on disclosure of confidential information will, of course, place some limits on the amount of material we can make public. Indeed, given the increasing transparency and the extremely limited discretion over the Currency Board arrangements, which draws its credibility from a strict rule-based system, a certain constructive ambiguity is also necessary if we are to have the flexibility to deal decisively with sudden crises. Within these limits, we are committed to developing our policy of transparency, and to increasing the quality and comprehensibility of the materials we release for public consumption.

18

Performance attribution analysis – a homemade solution

Alojz Simicak and Michal Zajac

Abstract

The paper provides information on the performance attribution analysis used in the reserves management of Národná banka Slovenska (NBS). The approach chosen was tailor-made to fit the individual requirements of NBS. The performance attribution described in this paper has been designed as a tool to facilitate the investment process based on the analysis of past relative and absolute returns. The first part of the paper addresses the motivation for performance analysis in the reserves management of NBS. Attention is paid to the clear and understandable definition of performance attributes measured by the tool. The second part is dedicated to a technical description of the calculations. The paper also points out the variability and customisable character of performance analysis outcomes. Finally, the possible limitations of the system are discussed in the conclusion.

1 Introduction

Národná banka Slovenska (NBS) is, as the central bank of the Slovak Republic, responsible by law for the management of the country's foreign exchange reserves. The main role of the Bank is to act as the monetary authority, and reserves management forms an essential part of its activities. There are two main reasons for this. The first is associated with the role of the Bank as a guarantor of price stability, while the second is a practical one – the profitability of the Bank is closely linked to the reserves management and its returns.

Like any other institution managing public funds, NBS strongly focuses on safety and liquidity aspects. Nevertheless, the pressure for increased efficiency and higher returns is growing.

In light of these requirements, there are natural limitations to efforts to achieve higher returns. We acknowledge that basically every ambition to improve profitability is accompanied by an increase in some associated risks, and we have therefore decided to provide a more analytical background and support for the investment process, as well as to provide some tools to facilitate the decision-making process within reserves management.

Designing a tool that enables past returns to be analysed was one of the steps we took to enhance the chance of higher future returns. The reasoning behind this is simple – without critical assessment of the past, it is impossible to gauge properly what might happen in the future.

Performance attribution as a tool was developed to fit the specific conditions and needs of reserves management in NBS, as per the following characteristics:

Characteristics:	
• Fixed income	the targeted portfolio consists of purely fixed income products
• Single currency	the portfolio has no room to deviate into other currencies
• Daily basis	the tool provides outcomes on a daily basis
• Security-based	the analysis is based on the behaviour of single securities against the yield curve, and not, unlike other similar tools, on the behaviour of market indices or yield curve sectors
• Including transaction costs	the tool reflects trading, including the bid/ask spreads
• Performance measurement against a customised benchmark	the establishment of a benchmark enables relative performance and benchmark to be also the subject of analysis
• No cash flows	no cash flow in/out of a portfolio during a measurement period.

2 Motivation

The bold ambition of performance attribution is to promote the increase of returns within the constraints imposed by low risk tolerance on the part of the investor. To achieve this objective, information produced by the performance attribution analysis should enable:

the Bank's management
- to recognise the sources of income from reserves management.
- to understand better the dependence upon market conditions and the magnitude of their effect.
- to assess/monitor the quality of the internal reserves management team by means of relative performance decomposition.

the risk management unit
- to quantify the sources of out/underperformance vs. benchmarks.
- to risk-weight the components of the absolute and relative return.
- to identify expertly, comment and/or promote preferred investment styles.
- to reveal weak spots in the benchmarks.

Portfolio managers
- to receive meaningful feedback that can be used in the investment process.

3 Description of the attributes

There are a wide range of ways to structure the attributes of performance. The criteria to be used for the segmentation of return attributes depends purely on the specific conditions and preferences of each respective investor.

NBS tried to achieve a relatively simple, understandable and unambiguous structure of performance attributes to support the practical applicability of analysis.

The performance of a portfolio typical for a fixed income investor such as NBS can generally be driven by attributes from three groups:

- Income generated by the securities held
- Income generated by cash management
- Other income sources (securities lending, special vs. general repo).

Income generated by the securities held

Central banks belong to the traditional type of investors who prefer investing mainly in "real" securities, instead of creating synthetic exposure. The main part of their income, by nature, is generated by the securities held. Splitting this income into three attribute categories enables the principal income drivers of the portfolio to be understood better, as well as the associated risks.

The first attribute category, "Time", is the basic element of any fixed income portfolio. As the name suggests, "time" covers any income (positive or negative) produced by the passage of time, i.e. it counts the change in value of a portfolio over time, provided other parameters (yield curve, spreads) remain constant.

Accrued interest is the most apparent and easily understandable attribute. There are two other attributes associated with the passage of time. Tracking the amortisation of premium/discount of a fixed income instrument helps to catch the "time" attribute for instruments with coupon higher/lower than par yields. The last attribute of this category is known as the roll-down effect. In case of a positive yield curve, this is the additional income resulting from yields decreasing towards maturity.

The second category of attributes is "Curve". This category includes the income generated by changes in the yield curve – i.e. changes in its level, shape and steepness. The magnitude of this attribute is directly linked to the interest rate exposure inherent in the portfolio.

The *third category of attributes in this group is "Spread"*. The last source from which fixed income portfolio can derive income is investing in spread product vis-à-vis the benchmark curve. This income can be earned either by holding a security lying above the benchmark curve, or by widening or narrowing of spread measured as the distance between the yield of a particular security and the appropriate benchmark curve.

Income generated by cash management

Cash may be considered as an alternative investment or a residual appendix of a securities' portfolio. Nevertheless, cash bears a return that must be included in the performance measurement. Time deposits and reverse repos are the main vehicles used to generate income from cash.

Other income sources

There is a third group of income from a fixed income portfolio– securities lending and special vs. general repo transactions. These vehicles have two similarities:

– it is difficult to incorporate their return into a benchmark, and
– the importance and volume of the additional income they generate depends on the attractiveness of the securities held in the portfolio.

These transactions are quite popular among central banks, especially because of their perceived low risk. The magnitude of returns changes over time quite significantly and depends on various factors. It is therefore basically impossible for benchmarks to measure this income attribute in a reasonable manner.

Structuring the attributes

Having divided the attributes into basic groups, we obtained the following vertical structure:

Time
Curve
Spread
Cash
Other income

Asset classes

The structure of performance attributes described so far does not distinguish between different types of securities held. It is therefore logical to add this information to the analysis, because the performance of different asset classes could be relevant. In the case of NBS we have created a split into two categories – governmental and non-governmental debt. This creates a more complex picture:

	Government	Non-Government
Time		
Curve		
Spread		
Cash		
Other income		

Absolute versus relative performance

Benchmarking is an industry standard for central banks managing foreign exchange reserves. It enables their management to assess the competitiveness of their internal investment teams. Therefore, besides the analysis of absolute returns, analysis of performance vs. benchmark can add a very important insight. Whereas the absolute performance attribution aims to establish how the portfolio achieved the return, the relative performance attribution looks for the answer to a different question: namely, why the portfolio outperformed/underperformed the benchmark.

	Portfolio		Benchmark		Out/underperformance	
	Gov	Non-Gov	Gov	Non-Gov	Gov	Non-Gov
Time						
Curve						
Spread						
Cash						
Other income						

4 Calculating performance attribution

There are two approaches used for calculating performance attribution. On one hand, there are simple models based on complete decomposition of return without any residual variable. On the other hand, many institutions use very complex and sophisticated procedures. These systems typically use a residual variable for calculating an unexplained proportion of returns of the decomposition procedure.

The calculation of performance attribution presented here is based on a complete decomposition of daily returns. In this section we focus only on decomposition of the return generated by securities. We use graphical and numerical tools to describe the nature of performance attributions.

The formulas used in this section to compute attributes will use the notations given below.

$\mathbf{R}$ return on security,
$\mathbf{R}_{\mathbf{Curve}}$ return from change in yield curve,
$\mathbf{R}_{\mathbf{Time}}$ return from change in time,
$\mathbf{R}_{\mathbf{Spread}}$ return form change in spread,
$\mathbf{y}_{\mathbf{(t)}}$ yield to maturity (YTM) in time t,
$\mathbf{y}_{\mathbf{(t)\ Curve\ (t)}}$ yield to maturity in time t on yield curve in time t,
$\mathbf{s}_{\mathbf{(t)}}$ spread of security,
$\mathbf{P}\ (\mathbf{y}_{\mathbf{(t)}})$ function to calculate price of security (Y(t) is an argument)

The daily returns of a single security can be written as the sum of several return attributes. Finally, performance attributes are determined by solving the following equations.

$$R = R_{Curve} + R_{Time} + R_{Spread} \quad (1)$$
$$s_{(t)} = y_{(t)} - y_{(t)\ Curve\ (t)} \quad (2)$$
$$R_{Curve} = P(y_{(t-1)\ Curve\ (t-1)}) - P(y_{(t-1)\ Curve\ (t)}) \quad (3)$$
$$R_{Time} = P(y_{(t-1)\ Curve\ (t)}) - P(y_{(t)\ Curve\ (t)}) \quad (4)$$
$$R_{Spread} = [P(y_{(t-1)}) - P(y_{(t-1)\ Curve\ (t-1)})] - [P(y_{(t)}) - P(y_{(t)\ Curve\ (t)})] \quad (5)$$

To demonstrate this numerically and graphically, we use an example based on holding a five-year bond, BUNDESOBL #139:

	Current Date	Previous Date
Clean Price	103.74	102.67
Value Date	17 July 2002	16 July 2002
Duration	4/03	4/02
Coupon	4%	
Maturity Date	16 February 2002	
Type Of Accrued	Act/Act	

Note: Although BUNDESOBL#139 was the part of the euro benchmark yield curve, we assume that there was a significant spread between the yield of this security and the interpolated yield based on the benchmark curve.

Chart 1: Graphical interpretation of the yields needed for further calculations

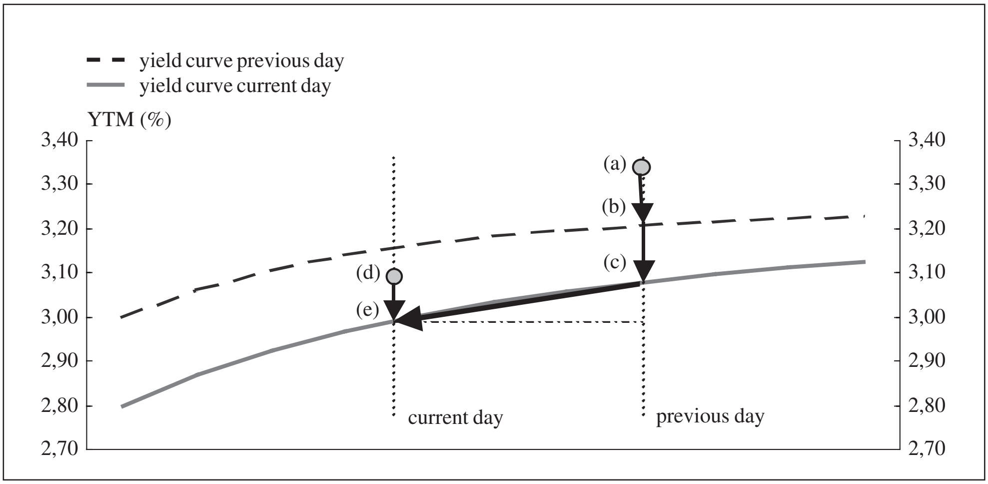

	Current Date (%)	Previous Date (%)
YTM of security	3.10 (d)	3.35 (a)
YTM of benchmark curve	3.00 (e)	3.20 (b)
		3.08 1)

1) The yield to maturity of 3.08% is gathered from the benchmark curve of the current date but calculation of the security price is based on the conditions of the previous date (settlement date is equal to previous date). This situation is depicted in the next exhibit.

Using equations (1) to (5) we calculate the profit or loss of the security in the given time interval for each attribute of return. The daily measure of decomposition is recommended as it allows easy handling of possible cash flows generated by security (trading, coupons, etc.).

The performance attributions calculated as yield and price change (in b.p.) are as follows:

	Yield	Price
Time	8	36
Curve	12	52
Spread	5	21

The success of performance attribution is subject to accurate input data such as prices of securities and curve construction. The main input is the benchmark yield curve. Each factor of return is determined by this curve. We recommend using a constant maturity yield curve. If the curve consists of certain securities, then these securities have zero spread against the customised benchmark curve.

The procedure of computing performance attributions was designed in the widely used Microsoft Excel worksheet environment. To evaluate the standard parameters of bonds, dynamic linking libraries (*.dll) that had been developed by an external firm were used.

Absolute and relative returns

There are many ways to interpret the results of performance attributions. If a report shows only the results of one portfolio, then this report is based on absolute returns. If the returns of two portfolios are compared (in many cases, portfolio against benchmark), then the report is on relative returns. In this section we show three sample reports, the first and second of which are based on absolute returns, while the third depicts the relative return of the portfolio against the benchmark.

1. A time series of absolute returns describes market circumstances and their relation to return on the portfolio. The performance of the portfolio (626 basis points (b.p.)) was primarily generated by the "time" attribute (393 b.p., 63% of the total return). Thanks to a falling yield curve, the portfolio produced its second largest income (192 b.p., or 31%). Owing to narrowing or non-widening of spreads, the portfolio recorded a return "spread" of 11 b.p., or 2%. Other factors of return are not significant in this example.

Chart 2: Sample time series of absolute returns
(in b.p.)

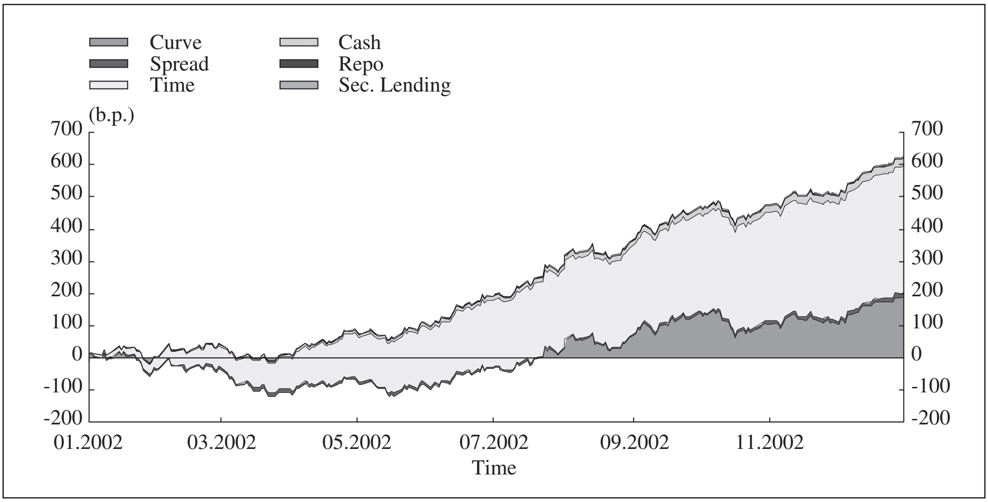

2. As the returns of the portfolio are related to the magnitude of risk exposure, it is useful to see the attribution of respective credit classes (in our case, government and non-government bonds). The system purely distinguishes income generated by securities held.

Table 1: Return on government and Non-government curities
(in b.p.)

	Gov	Non-Gov	Total
Time	283	110	393
Curve	134	58	192
Spread	4	6	11
Cash		24	24
Repo		2	2
Sec. Lending		4	4
			626

The split of performance attributes based on duration, i.e. of interest rate exposure, is visible in the next exhibit:

Table 2: Decomposition of return based on maturity
(in b.p.)

	Time	**Curve**	**Spread**	**Total**
less then 0.5 year	73	3	-1	75
0.5 – 1.5 years	29	10	0	39
1.5 – 2.5 years	94	35	2	130
2.5 – 3.5 years	42	26	3	70
3.5 – 4.5 years	82	59	2	142
greater than 4.5 years	74	60	5	139
	393	192	11	596

3. Here we compare the returns of the portfolio against the returns of an appropriate benchmark. The goal is to extend analysis of the return and provide precise information about sources of relative return. The total outperformance of the portfolio (27 b.p.) was primarily generated by proper cash management (23 b.p.). There are also 6 b.p. from other sources of income, namely repo and securities lending operations. It can therefore be concluded that the portfolio manager preferred investing cash against holding bond positions because of an inverted yield curve.

Table 3: Relative returns of portfolio against an appropriate benchmark
(in b.p.)

	Portfolio		Benchmark		Outperformance		Total
	Gov	**Non-Gov**	**Gov**	**Non-Gov**	**Gov**	**Non-Gov**	
Time	283	110	349	63	-66	47	-19
Curve	134	58	133	57	1	1	2
Spread	4	6	-3	-2	7	8	15
Cash	24		1		23		23
Repo	2		0		2		2
Sec. Lending	4		0		4		4
	625		598		27		

5 How can data be transformed into information?

The crucial part of the performance attribution analysis is the use of the outcomes. The numbers generated by the system only form the basis for further discussion, which is dedicated to improving the quality of the investment process.

To illustrate the usefulness of this tool, some theoretical examples and situations are presented in the following paragraphs.

The first group of situations relates to the interpretation of the absolute returns (i.e. returns as a percentage of a single portfolio). Let us imagine that the "curve" attribution of the performance is strongly dominant. Several conclusions can be drawn from this fact. As the "curve" attribution relates to the change in the level and shape of the curve, it is not rational to expect similar returns in the future and simply to extrapolate the return into the future. Furthermore, as the yield curve lowers, the "time" attribution in the following periods will be lower too as it directly depends upon the level of interest rates. Given the change in yields and the performance related to it, decision-makers may gain better understanding of the magnitude of interest rate risk built into the portfolio. It is true that the "curve" attribution can have a positive as well as a negative sign.

Another situation that could be imagined is that the portfolio consists of government and non-government securities. The management and risk managers will seek to answer the question whether it was right to include non-governments (i.e. credits) from the return point of view. Performance attribution can answer this, providing a number that indicates how much of return relates to the credits. The answer can be structured in several ways, as a single number representing the portion of the performance of credits in the total performance, as a breakdown alongside the yield curve based on sectors of the yield curve, or as a number signalling how much return relates to the credit risk of the portfolio.

The analysis of absolute returns is performed mainly for bank management and risk management purposes. Simply expressed, a reliable guide for estimating future returns is desirable. On the other hand, the analysis of relative performance against a predefined benchmark can provide all involved parties – management, risk and portfolio managers – with helpful insight and additional information.

Management generally seek to obtain information on the skills of the in-house portfolio management team. Performance attribution breaks down the "value added" into several parts. For example, a volatile relative attribution coming from "curve" would indicate an extensive use of duration bets in portfolio management. On the other hand, a steady contribution from "credit" implies careful stock selection with yield spread vs. the benchmark yield curve. Conversely, a very steady contribution from a particular investment strategy could suggest to a risk manager that there might be some weak points either in the selected benchmark or in the rules and limits applied.

Finally, portfolio managers are continually comparing their efforts dedicated to different investment tactics with their performance. Using the numerical outcomes from the performance attribution, they can react quickly to cut or build up a particular position in a more justifiable way.

Performance attribution is defined as an analytical process or technique that evaluates the decisions that contribute to return and excess return. These days, decomposition provides us with very precise data, yet making statements is more or less intuitive. The main challenge and future development issue regarding the decomposition model is connecting attributes with risk factors. This approach provides risk-weighted performance attribution.

6 Limitations of the performance attribution

The usefulness of this type of system is limited by the accuracy of price/yields supplied by market information providers. Especially when entering less liquid market segments, one has to keep this in mind and pay greater attention to proper data collection.

Data management is a further limitation. The capacity of spreadsheet calculators cannot cope with an extensive analysis of securities across several asset classes. A powerful database machine based on Structured Query Language standard should be employed.

The largest and most important limitation of performance analysis is however not technical, but human. A clear explanation of calculated outcomes and an understanding of the economic value of these are the primary conditions for full exploitation of the system's potential.

7 Conclusion

The users of performance attribution analysis receive a huge amount of information. The task of making reasonable conclusions rests on these users and their professional judgement. For example, the proportion of the "curve" effect on overall performance is a good indicator of interest rate exposure, but the decision as to whether to increase or decrease duration based on this indicator depends on a subjective assessment of inherited information. As there are no industry norms or standards, the results of performance analysis are not intended to assess or to serve as a form of traffic light. The analysis does however provide useful guidance.

A professional approach and sound assessment of the results of performance attribution analysis are therefore necessary. Performance analysis has been developed to facilitate communication among persons involved in the investment process, and to provide them with useful information. The system is now in its second year of use at NBS and is being updated continuously.

Performance attribution for fixed income portfolios in Central Bank of Brazil international reserves management

Antonio Francisco de Almeida da Silva Junior[1]

Abstract

There is no standard format for performance attribution reports, and each methodology is developed to capture some specific issues of the investment process. However, there are some generic theoretical features to implement a performance attribution system. This paper discusses some details of the performance attribution framework used by Central Bank of Brazil in its international reserves operations in fixed income portfolios, and highlights its features. It reaches the following conclusion: i) performance attribution reports contribute to the transparency of the investment process, since it is possible to analyse the impacts of each decision; ii) it is not strictly an accurate result of a mathematical procedure, since there are assumptions underlying each approach; iii) it contributes for returns to be risk adjusted. Furthermore, as a topic for further studies, this paper suggests a link between performance attribution and performance evaluation reports.

1 Introduction

Countries need international reserves to trade with other countries. Furthermore, international reserves are necessary to conduct an exchange rate policy, to avoid speculative currency attacks and to pay foreign debt. The level of reserves depends on the nature of the exchange rate policy, the international trade volume, and the cost of holding reserves among other factors.

Central banks are responsible for managing international reserves. Some might imagine that this task is similar to that of an asset manager, but in fact this is not the case, as central banks are very conservative institutions and need a high level of liquidity. There are few papers on central bank tasks in the management of their portfolios. Eichengreen and Mathieson (2000) point out that the management of international reserves remains one of the more understudied aspects of the international monetary system.

Central banks' management of reserves has become a matter of concern in recent years. In 1994, several Latin American countries held a seminar in Colombia entitled "Management of International Reserves by Central Banks", supported by CEMLA (the Latin American Monetary Studies Centre). One of the key conclusions reached at this seminar was that the trend among central banks is towards a more performance-driven approach.[2]

Therefore, the construction of a framework for reserves management is a key goal for central banks. This framework must clearly identify the investment objectives, the benchmark and the desired risk level. It is also very important to have risk and performance reports that allow the investment process to be evaluated. This kind of report also helps make the process

---

[1] The author works at the International Reserves Operations Department in the Risk Analysis Team at the Central Bank of Brazil. The views expressed here are those of the author and do not necessarily represent the views of the Central Bank of Brazil. The author is also very grateful to the International Reserves Operations Department Team for invaluable comments.

[2] See Goldenstein (1994) and Lichtbuer (1994).

more transparent, since it is possible to evaluate whether investment objectives have been attained. Effective reserves management is a useful tool that provides more credibility for central banks, since it allows the cost of holding reserves to be reduced.

This paper discusses some features of performance attribution that are designed to identify the sources of return in the investment process. Performance attribution is one step of performance report. It contributes to better knowledge of the risk return profile, and is a powerful tool that can help managers understand the decision standards.

2 Investment process oriented by a benchmark portfolio in the Central Bank of Brazil

The Central Bank of Brazil's investment policy of international reserves management is a formal framework that defines a reference portfolio, or benchmark, with the strategic investment objectives. International reserves are invested in a currency hedging strategy. The currency distribution replicates the profile of Brazil's external debt in the short run. Liquidity and profitability are other fundamental variables that the investment framework takes into account.

The Central Bank of Brazil is currently using an active portfolio management strategy oriented by a benchmark portfolio[3]. There are eight portfolios: three for money market operations in dollar, euro and yen; two government bond portfolios (in dollar and euro), one in gold, one portfolio for foreign currency operations, and finally one for Brazilian government bonds.

Furthermore, a small part of international reserves is shared between six external managers. For the Central Bank of Brazil, the aim of the external management programme is to "capture" the investment process expertise of managers in the international markets. Therefore, they have periodic meetings with Central Bank of Brazil staff, present periodic reports and run a training programme for the Bank.

Portfolio management oriented by a benchmark was chosen because it provides a quantitative description of portfolio management objectives, given that it is possible to define the currency profile and the average maturity or modified duration level of the reserves. This facilitates adherence to plan constraints and allows the Bank to verify whether managers are within the guideline limits. It is a vehicle to delegate risk mandates effectively and measure comparative performance since managers have strategic and tactic risk levels, which can be evaluated in terms of risk and return of benchmark and effective portfolios. This permits more meaningful performance evaluation and attribution, is a fair and effective discriminator of management skill, and allows gaps and overlaps in overall asset allocation to be identified.

There are some drawbacks in using portfolio management oriented by a benchmark. One of them is the fact that benchmark indices with constant duration are rebalanced without taking into account transaction costs. Generating virtual transactions for the benchmark portfolio can indeed solve this problem. Another drawback of the benchmark approach is that there are some operations that can be performed in the effective portfolio, such as swaps, repurchase agreements, securities lending etc. that usually cannot be implemented in the benchmark portfolio. In fact, Elton and Gruber (1995) state that one way to compare portfolios is to

[3] Libid is used for money market operations as a benchmark in US dollar, euro and yen deposits. Libid is the Libor minus 1/8 and is an average of money market deposit rates. Chase JP Morgan 1 to 3-year indices are used as a benchmark for US and European government bond portfolios. The EMBI+ Brazil from Chase J.P. Morgan was used as the benchmark for Brazil's bonds portfolio.

examine the return earned by alternative portfolios of the same risk; ideally, therefore, effective and benchmark portfolios should have the same risk level.

Managers can deviate from the benchmark portfolio, and this deviation is the key feature of active management. Deviations from the benchmark portfolio are however constrained by some guidelines that limit managers' room to act. The guidelines used by the Central Bank of Brazil include among other constraints a maximum differential market risk allowance (2% of differential VaR on an annual basis), limits on the amount of individual issues in the portfolio (liquidity risk constraints), and a maximum amount of investment in securities of different credit risk levels (supranationals and agencies).

The total differential return of a portfolio is the difference between effective and benchmark portfolios returns, and corresponds to the sum of each "return factor" owing to deviations from benchmark decisions in a mark-to-market basis, which is the result from active management operations. Sources of differential return include foreign currency profile deviations, duration and yield curve deviations, credit risk deviation, etc.

In the investment process there are strategic monthly meetings, where managers and the research team report past results to the strategic committee, briefly review the economic variables, explain the risk profile and suggest the strategy for the next month. The committee then discusses the proposal and defines the strategy for the next month.

3 The role of performance attribution

The Risk and Performance Report is presented on a quarterly basis to the Central Bank of Brazil's board of directors. These meetings are used as a management tool, which enables evaluation of the investment process and provides a forum to discuss sources of return. The performance report of the Central Bank of Brazil's international reserves management is presented in three complementary blocks. The first is a return calculation report, which measure profits or losses (P&L) incurred by trades in the portfolio management activity.

Although it seems very simple, there are different approaches of return calculations that may be classified in internal rate of return measures (IRR) and time-weighted calculations. The internal rate of return is the single discount rate that, when applied to a bond's expected future cash flows, equates those future cash flows to the price at which the bond is currently traded in the market. There are four yield measures: yield to maturity; yield to call; yield to average life; cash flow or discounted cash flow yield (Douglas, 1988). Otherwise, time-weighted return is a calculation of the investment return generated by a manager over specific time periods that are geometrically linked or compounded. Time-weighted returns calculations isolate a manager's specific performance from investor timing and the direction of cash flows, which allow for a consistent measure of a manager's performance relative to other managers and to alternative investments and indices (Clarfeld, 1998).

Different approaches to return calculations drive investors to different conclusions about the performance of their portfolios. Therefore, the Association for Investment, Management and Research (AIMR) has introduced standards for performance presentation, which recommend using the time-weighted approach for performance calculation, evaluation and attribution (see AIMR (2001)).

The Central Bank of Brazil uses time-weighted returns for performance reports. Sales and purchases, accrued interest, amortisation and foreign currency operations are taken into account at the time they occur in both the effective and benchmark portfolios, and approximations and simplifications in the return calculus due to operational constraints are kept to a minimum.

The second block on performance reports is performance evaluation, which tries to define whether traders are under or outperforming the "market portfolio" or the benchmark on a consistent statistic basis. There are several indices for this task, such as the Sharpe ratio, the Treynor index, Jensen's alpha etc.

The third block is a performance attribution report, which is designed to identify the sources of differential returns in a portfolio with many assets and currencies. Bodie et al. (1999) state that rather than having a focus on return adjusted by risk indices, which is the task of performance evaluation reports, the main goal of investors and managers is to identify the decisions that have resulted in underperforming or outperforming the benchmark. Therefore, underperforming or outperforming the benchmark would be a result of market and asset choices and timing. Even though good past performance is not a guarantee of good future results, performance attribution reports can be used as an additional tool by managers in the investment strategy process. In fact, performance attribution and performance evaluation analysis are complementary tools to the investment process, and Central Bank of Brazil takes them both into account.

Performance attribution reports split global performance into discrete components so as to identify sources of differential return. These are only one tool enabling better understanding of what has happened in the portfolio. Such reports need to have three key features: reliability, accuracy and relevance. Reliability is related to results that are true and not a consequence of information manipulation. Accuracy is related to the approach and methodology used in the calculations, while relevance is associated with the identification of sources of return that are relevant in the investment strategy process.

Performance attribution calculations using a benchmark approach for active management need to take into account the guidelines for a fair identification of return sources. Therefore, if the benchmark has only Treasuries, for instance, and traders are allowed to invest in agencies and supranationals, which have different credit risk levels, the performance attribution report must identify this source of differential return. For this reason, performance attribution and performance evaluation are complementary tasks.

4 The building block approach for performance attribution

There is no standard for performance attribution reports, and each financial institution builds its own reports based on its investment strategy process. In general, reports attempt to measure how much of the P&L can be attributed to foreign currency, market allocation and security selection (Bodie et al., 1999, Reilly and Brown, 2000). The fact that performance attribution should be appropriate to the investment process poses a challenge to any generic performance measurement solution (Thompson, 1997). Dynkin et al. (1998) state that the explanatory power of performance attribution depends on how well the selected variables set reflects the actual views on allocation decisions specific to the portfolio.

This section presents some considerations regarding the implementation of a performance attribution system for fixed income investments. This approach is termed here the building block approach, which takes into account returns in the securities level and identifies the differential return relative to benchmark sources.

Bodie et al. (1999) discuss a typical performance attribution system with three components: (1) global allocation in equities and fixed income (2) market allocation in each sector; and (3) asset selection in each sector. Considering a benchmark and an effective portfolio, the differential return may be calculated by the following equation:

$$r_P - r_B = \sum_{i=1}^{n} \left(w_P^i . r_P^i - w_B^i . r_B^i \right) \tag{1}$$

where

r_P, r_B	are the effective and benchmark portfolio returns,
w_P^i, w_B^i	are the effective and benchmark portfolio weights in each security or sector, and
r_P^i, r_B^i	are the effective and benchmark portfolio returns in each security or sector.

The equation above may be written in such a way so as to show how decisions in market allocation and security selection are taken:

$$allocation = \left(w_P^i - w_B^i \right) \cdot r_B^i \tag{2}$$

$$selection = w_P^i . \left(r_P^i - r_B^i \right) \tag{3}$$

$$allocation + selection = w_P^i \cdot r_P^i - w_B^i . r_B^i \tag{4}$$

The allocation term measures the impact of choosing to be over weighted or under weighted in a specific sector or market versus the benchmark, whereas the selection term shows the contribution to total performance of selecting a security in an asset class.

Equations (1) to (4) represent a typical approach to investments in stocks, where the wisdom of decisions in the period in question can be identified. However, this approach is not good enough for bond investments. Dynkin et al. (1998) present Lehman Brothers' performance attribution approach for fixed income portfolios, and they use the above equations to generate their reports. The Lehman Brothers' approach calculates yield curve allocation based on a breakdown of performance owing to curve positioning using equation (4) for each duration sector. This approach does not split portfolio duration effect and the curve positioning effect, as they are presented together.

The duration effect is the differential return owing to a global duration contribution to return, so if the effective portfolio has a longer duration than the benchmark, a positive parallel shift in the yield curve, for instance, will result in a negative differential return. Otherwise, even with an effective portfolio with a global duration equal to the benchmark one, a change in the yield curve that is not parallel could result in a differential return not equal to zero. This is what is termed here a 'curve effect' owing to bets in changes of the yield curve shape.

For fixed income portfolios, it is necessary to consider the modified duration differential between effective and benchmark portfolios and capture the duration effect. An alternative is to split duration and allocation effects adopting the following equations used by Robeco Asset Management (Breukelen, 2000):

$$duration = \left(\frac{d_P^{tot}}{d_B^{tot}} - 1 \right) \cdot \left(r_B^{tot} - r_{mm} \right) \tag{5}$$

$$allocation = \sum_{i=1}^{n} \left(w_P^i . d_P^i - w_B^i . d_B^i . \frac{d_P^{tot}}{d_B^{tot}} \right) \cdot \left(\frac{r_B^i - r_{mm}}{d_B^i} - \frac{r_B^{tot} - r_{mm}}{d_B^{tot}} \right) \tag{6}$$

$$selection = \sum_{i=1}^{n} \left(\frac{r_P^i - r_{mm}}{d_P^i} - \frac{r_B^i - r_{mm}}{d_B^i} \right) \cdot w_P^i . d_P^{tot} \tag{7}$$

where, for each country or sector *i*:

w^i_P *and* w^i_B are the portfolio and benchmark weights,

d^i_P *and* d^i_B are the portfolio and benchmark modified duration,

r^i_P *and* r^i_B are the portfolio and benchmark return, and

r_{mm} is the money market return.

Equation (5) attributes the excess returns relative to the benchmark owing to distribution along the yield curve together with duration effect. In other words, this approach does not identify the curve positioning effect. This problem could be solved for an "*i*" sector by:

$$curve = \sum_{k=1}^{n} \left(w_P^k - w_B^k \right) . \left(r_{B\_adj}^k - r_{BCD} \right) \tag{8}$$

$$r_{B\_adj}^k = \frac{d_P}{d_B^k} \cdot \left(r_B^k - r_{mm} \right) + r_{mm} \tag{9}$$

$$r_{BCD} = \frac{d_P}{d_B} \left(r_B - r_{mm} \right) + r_{mm} \tag{10}$$

where index *k* refers to a duration bucket, and

$r^k_{B\_adj}$ is the benchmark bucket return adjusted for the portfolio duration, while

r_{BCD} is the benchmark return adjusted for the portfolio duration.

Equations (9) and (10) are based in the assumption of a linear relationship between return and duration, and so returns are corrected by portfolio duration. This correction is also implicit in equations (5) to (7). Although this is an acceptable assumption, the correction using money market as a way to define the linear relationship is not a very good one. It would be better to use returns of assets near the portfolio duration, which would reduce the error of a linear relationship assumption.

Furthermore, for curve effect calculation it is necessary to have the *k*-bucket return of the yield curve. One alternative is to use the return of the assets that are in the effective portfolio, and then split them into duration buckets to use equations (8) to (10). Although this is a good alternative, the fact is that the effective portfolio is not static (there are purchases and sales), and does not allow bucket returns to be used, as *k*-bucket return is affected by trade operations. Figure 1 shows what could happen with bucket return, using assets return to calculate them.

So, an alternative is using the zero coupon bucket returns to calculate the curve effect. The zero coupon returns are also shown in Figure 1. As the calculations are based on the bucket return differentials, zero coupon curves can be used as an approach to curve effect calculations. So, these are adjusted using a linear relationship assumption according to the following equation:

$$r_{adj} = r_{d0} + \left(rz_{d\_adj} - rz_{d0} \right) \tag{11}$$

where

rz_{d_adj} is the return of the zero coupon at adjusted duration,
rz_{d0} is the return of the zero coupon at the base duration,
r_{d0} is the return of the portfolio (or asset) at the original duration, and
r_{adj} is the return of the portfolio (or asset) adjusted for the new duration.

Figure 1: Bucket return for curve effect calculation

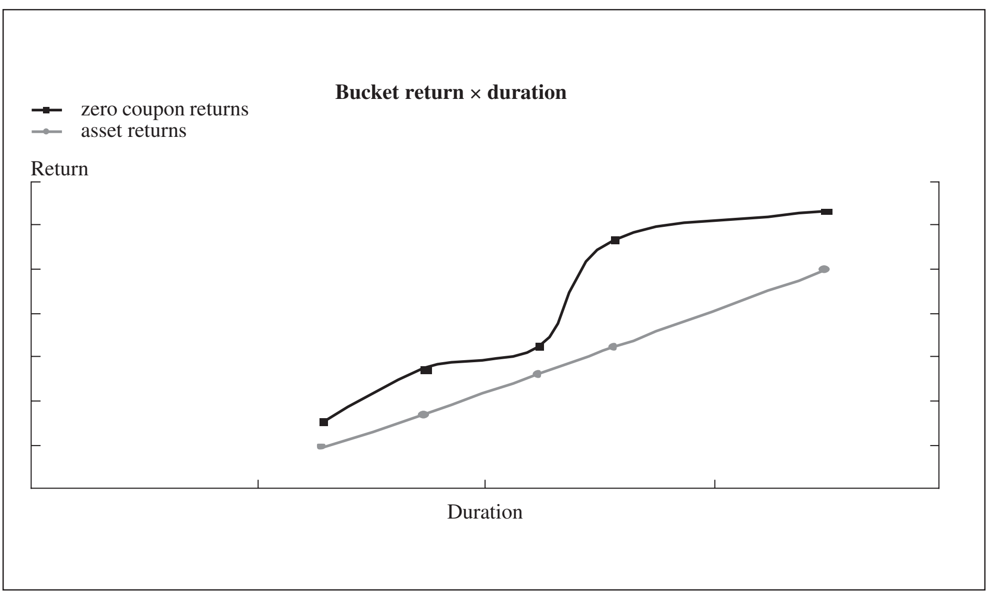

Note that the term in brackets in equation (11) is the zero coupon return difference. This approach considers that zero coupon return and a non-zero coupon return curves are parallel. In fact, the Central Bank of Brazil has implemented return adjustment using equation (11), as it produces better results than equations (9) and (10).

Another effect to be segregated is the credit effect. If the benchmark has only Treasuries and the portfolio has bonds with different credit risk, such as agencies and supranationals, there is a credit risk effect in portfolio performance compared with the benchmark portfolio. It is to be expected that the portfolio return will be higher than the benchmark return owing to credit risk. In fact, this could not happen in the short run owing to spread changes.

One could solve this problem using a benchmark with spread products, but it is first necessary to develop an internal benchmark with specific guidelines. Another way is to use two portfolios, the first with Treasuries and the second with spread products, each with its own benchmark. However, this would require building an internal portfolio for spread products owing to specific guidelines. There are few indices or benchmark portfolios for agencies and supranationals available in the market, and these often cannot be adjusted for internal constraints.

The third alternative is a type of virtual portfolio with a composite of the Treasuries index and a naive investment strategy in agencies and supranationals with a constant equal distribution of investments in the same assets of the real portfolio in the period. In fact, this is only an approximation approach that can be used if there are operational limits that constrain the construction of a benchmark portfolio for spread products with the specific guidelines. The credit contribution to the portfolio is calculated by the following equation:

$$credit = \left(x_A . r'_A + x_S . r'_S + x_B . r'_B\right) - r'_B \tag{12}$$

where

x_A	is the average of weights in agencies in the effective portfolio,
x_S	is the average of weights in supranationals in the effective portfolio,
x_B	is the average of weights in Treasuries in the effective portfolio,
r'_A	is the agencies return of a naive portfolio with the real portfolio duration,
r'_S	is the supranationals return of a naive portfolio with the real portfolio duration, and
r'_B	is the benchmark return corrected by the effective portfolio duration.

Figure 2 provides a brief view of the building block approach discussed in this paper. Point P is the effective portfolio at duration *dp* and return *rp*. Benchmark return is *rb* at duration *db*. So, using equation (11), benchmark return is corrected from B to B' at effective portfolio duration. The overall duration contribution for differential return is *rb'-rb*. Spread products return and duration are presented at point A, and their duration adjusted return is *r'a*. Therefore, point C is a weighted result of point A' and B', and equation (12) shows credit contribution to differential return (segment CB'). Thus, segment PC represents allocation, repo, selection and curve effects. Allocation effect is calculated using equation (6), and curve effect using equation (8). Moreover, as the benchmark portfolio does not have repo operations, repo contribution is calculated based on operations. Selection effect is calculated as the difference between the total differential return and allocation, duration, repo, credit and curve effect.

Figure 2: Building block approach view

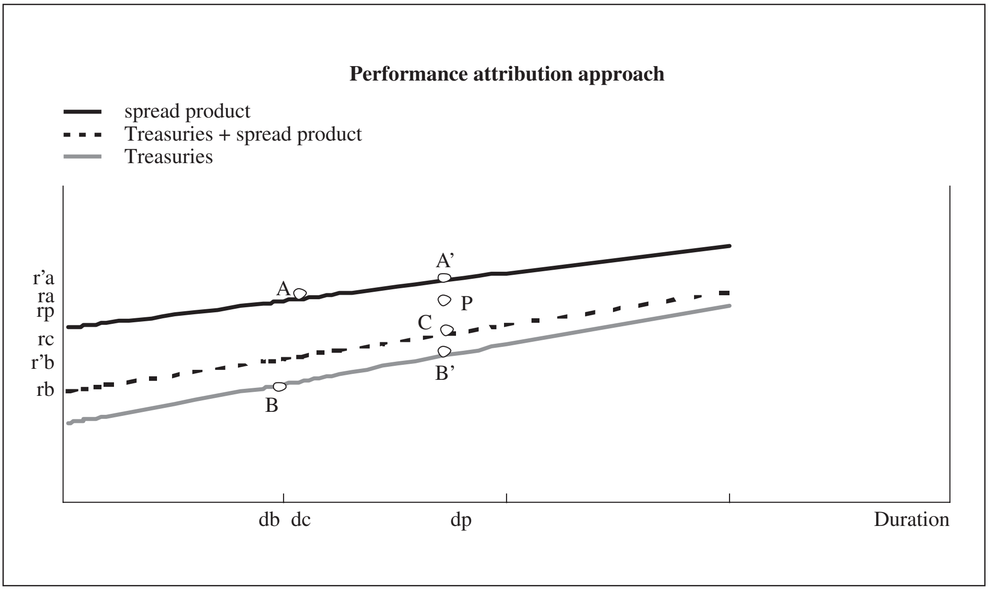

Credit effect can be divided into two contributions. The first one is a cushion effect, as there is an intrinsic spread between agencies and supranationals and Treasuries. Therefore, to measure this contribution, average spread is used. The second effect is spread variations during the period.

5 Specific issues in performance attribution

This section discusses some issues related to performance attribution. First of all, it is important to recognise that performance attribution is not strictly an accurate result of a mathematical process. It is rather an attempt to explain the sources of differential return that has mathematical restrictions, as it is not possible to identify exactly each contribution.

Dynkin et al. (1998) suggest that a model which assumes a static portfolio over the reporting period and does not account for intra-period trading and cash flows in and out of the portfolio constrains the ability to analyse actively traded portfolios. This problem is also discussed by Ramaswamy (2000), who states that many of the existing software platforms do performance attribution on the basis of securities held at the beginning and end of the month. In fact, performance attribution in a dynamic portfolio cannot exactly identify each source of return. Since deals during the intra-period do not permit an exact calculation, the performance attribution approach is a sort of "differential return sharing".

The second feature is the use of performance attribution in complementary analysis for performance evaluation. Figure 3 shows the accumulated differential return between effective and benchmark portfolios, the duration differential and the evolution of the zero coupon yield in a sample period of a simulation bond portfolio. The figure is divided into four periods. In the first one there was a positive duration differential and a yield increase with a negative differential return. In the second period there was a positive duration differential, and yield was relatively stable, as was the differential return. In the third period there was a slight

Table 1: Performance attribution in a sample period

	Period					
	1	2	3	4	Average	Std Deviation
Duration	-7.72	0.85	3.01	12.37	2.13	8.25
Other Effects	-4.74	4.54	-0.71	0.38	-0.13	3.81

duration differential, a decrease in yields and a smooth increase in differential return. In the last period the duration differential was increased and, as yield decreased, differential return increased. Table 1 shows performance attribution results for this duration strategy.

Looking only at duration contribution numbers in Table 1 and in Figure 3, it can be inferred that variability in duration performance results is a function of duration strategy changes and yield movements. It is therefore not easy to make any kind of assumption about the probability distribution function of duration contribution numbers, since it depends upon changes in the investment strategy. This is not a good feature, as it is very important to understand if there is any kind of consistency in duration results (and even in other contributions). Generally this kind of analysis is not available in performance attribution studies, since consistency analysis is a field of performance evaluation. In fact, performance attribution analysis is restricted to answering what occurred with past return as a complementary analysis.

Table 1 shows that the duration effect numbers have a high standard deviation compared with the mean. It shows that the positive results may not necessarily be associated with skill at beating the market. The use of performance attribution as a way of evaluating results as an integrated step in performance evaluation is suggested here for further study.

An alternative to analysing Figure 3 and Table 1 is evaluating the returns of benchmark and effective portfolios in a plot similar to that of Figure 4. If managers have timing ability, they should have duration shorter than the benchmark in a scenario of yield tightening, and a

Figure 3: Accumulated differential return evolution in a sample period

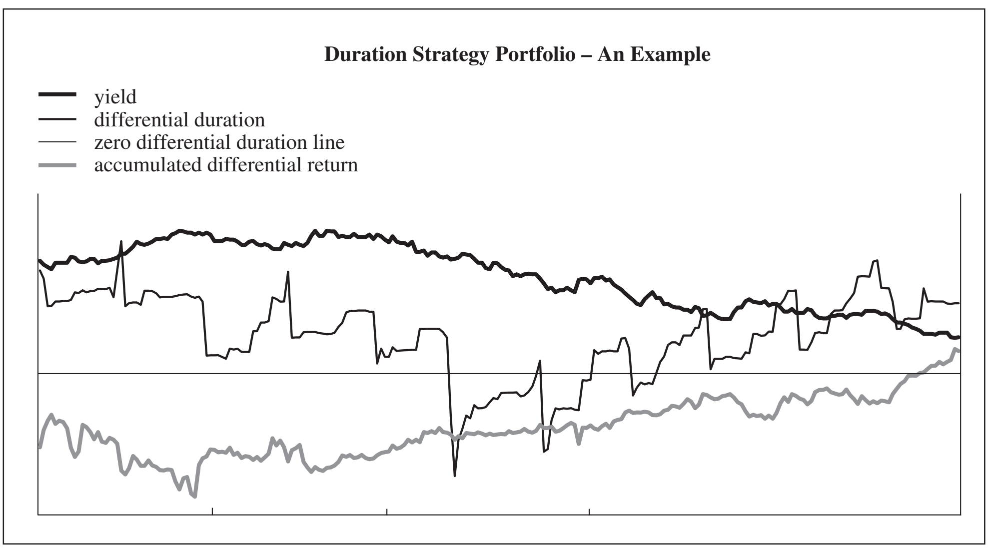

Figure 4: Timing analysis

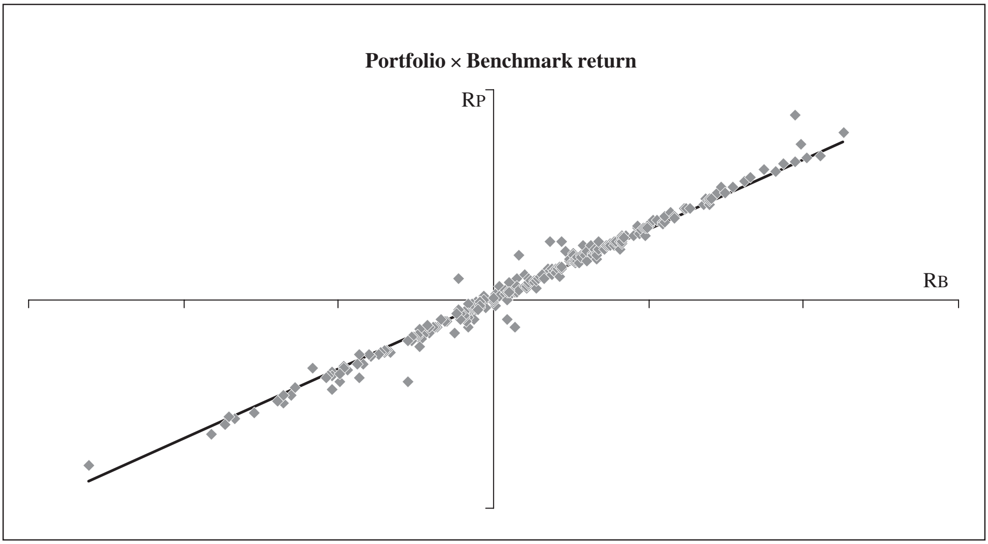

longer duration than the benchmark in a scenario of yield easing[4]. Timing ability should give results above the line in Figure 4 that would result in a curvature in the scattered points. In fact, on closer inspection Figure 4 does not reveal timing ability in the duration strategy, since there is no curvature. Treynor and Mazuy (1966) present one procedure to test the existence of curvature with the regression of equation (13). Results of the regression from equation (13) with data from Figure 4 show that the "c" coefficient is not different from zero, which means that there is no timing skill with this portfolio.

$$(r_{Pt} - r_{Ft}) = a + b.(r_{Bt} - r_{Ft}) + c.(r_{Bt} - r_{Ft})^2 + e \qquad (13)$$

where

r_{Pt}	is the return on portfolio in period t
r_{Bt}	is the return on the benchmark in period t
r_{Ft}	is the riskless asset return
e	is the residual return, and
a, b, c	are constants.

The third issue to discuss here is the fact that each performance attribution approach must be developed taking into account the investment process. Several alternative approaches exist to deal with this task. Table 2 reports the performance attribution of various international reserves external managers and from the internal portfolio of the Central Bank of Brazil. Based on this, it can be concluded that each manager has his or her own performance attribution approach, and that these approaches probably reflect their individual investment processes.

[4] Elton and Gruber (1995) suggest that a measure of a manager's timing ability is to compare directly the portfolio return and the market return.

Table 2: External asset manager and Central Bank of Brazil performance attribution reports

Manager A Differential Return	Manager B Differential Return	Manager C Differential Return	Manager D Differential Return
Foreign Exchange	Currency	Modified Duration Allocation	Country Allocation
Asset Allocation	Duration	Country Allocation	Currency Selection
Duration	Curve	Yield Curve Allocation	Curve Selection
Yield	Credit Exposure	Bond Selection	Credit Allocation
Stock Selection / Yield Curve	Accrual Conventions / Others	Foreign Currency Allocation Trading	Transaction Cost
			Security Selection

Manager E Differential Return	Manager F Differential Return	Central Bank of Brazil Differential Return
Duration	Pricing, Currency & Future Diff.	Forex
Country Allocation	Yield Adv. (Benchmark Roll Yield)	Differential Return
Govt. Bond Selection	Duration (Parallel Shift)	Duration
Credit Selection	Yield Curve (Reshaping)	Allocation
Currencies	Spread (Rolling Yield Effect)	Credit
	Total Mkt Spread (Market Effect)	Curve
	Spread Advantage	Selection
	Currency Allocation	Repo

It is interesting to note briefly that although numbers are not provided here, the pattern of the performance attribution results for all external managers as well as the Central Bank of Brazil is similar to that of Table 1, i.e. variability is high when compared to the mean. Some possible explanations for this pattern include market fluctuations, strategy changes, valuation sources and model features, as previously discussed. As has already been pointed out, data variability is a question of analysis. Since an active management strategy is based on the assumption that managers possess the requisite skills to beat the market, one could expect that data variability would not be so high compared with the average.

It could also be argued that when analysing the split between various investment decisions, collective decisions should be taken into consideration. If a manager wants to invest in spread products, owing to liquidity restrictions, he or she could make curve allocation bets. Therefore, splitting performance could prove problematic.

Table 3: Performance attribution – Portfolio 1

Two-week basis

	Duration	Credit	Curve	Selection	Repo
1st quarter	1.20	13.74	0.59	3.98	0.63
2nd quarter	2.54	9.00	-0.45	0.06	0.15
3rd quarter	-3.69	-0.49	1.67	-0.70	0.24
4th quarter	-2.33	6.90	-1.84	-0.40	0.42

Monthly basis

	Duration	Credit	Curve	Selection	Repo
1st quarter	-0.39	13.00	0.53	6.96	0.63
2nd quarter	2.50	9.00	-0.46	0.58	0.15
3rd quarter	0.00	0.22	1.86	-2.53	0.24
4th quarter	-2.90	7.52	-1.19	-0.80	0.43

Quarterly basis

	Duration	Credit	Curve	Selection	Repo
1st quarter	-1.88	9.50	0.56	11.21	0.63
2nd quarter	2.44	9.20	-0.56	0.23	0.15
3rd quarter	-2.49	0.92	1.93	-3.30	0.24
4th quarter	-0.05	7.48	-1.48	-3.59	0.43

Another issue for analysis is the frequency of performance attribution calculations. Tables 3 and 4 present results that were calculated on a two-week, monthly and quarterly basis over a sample period for two simulation portfolios. There is an influence of frequency calculation in performance attribution results.

Table 4: Performance attribution – Portfolio 2

Two-week bais

	Duration	Credit	Curve	Selection
1st quarter	-14.01	0.00	-0.61	10.69
2nd quarter	5.54	0.00	-0.43	3.99
3rd quarter	5.31	0.05	-0.48	0.36
4th quarter	7.03	0.16	0.63	-2.11

Monthly basis

	Duration	Credit	Curve	Selection
1st quarter	-12.05	0.00	-0.52	9.62
2nd quarter	5.27	0.00	-0.28	4.17
3rd quarter	4.74	0.15	-0.57	0.99
4th quarter	6.26	0.14	0.56	-1.23

Quarterly basis

	Duration	Credit	Curve	Selection
1st quarter	-13.01	0.13	-0.38	8.10
2nd quarter	7.20	0.12	-0.13	3.20
3rd quarter	4.40	0.14	-0.50	1.88
4th quarter	5.20	0.13	0.59	0.30

Moreover, the way numbers are compounded is another source of problems, since differential returns (performance attribution numbers) may not be compounded on a geometric or arithmetic basis, and so some calculus approximation is necessary, such as a normalisation procedure. An alternative to compound results would be to define differential returns not as a subtraction, but as a division operation.

Finally, it is important to highlight the fact that, in general, selection effect is calculated as a difference result, so it comprises the results of the skill in selecting assets and all the errors owing to modelling approximations.

6 Conclusions

Performance attribution is a useful tool in the investment process as a way of analysing return numbers and the sources of good or bad performance. Otherwise, past results are not a guarantee of future ones. Performance attribution reports contribute to the transparency of the investment process, as it is possible to analyse the impacts of each decision.

Although it represents a useful tool in the investment process, there is no standard theoretical approach, and the process is not the result of a strictly mathematical procedure. There are assumptions underlying each approach and approximations in the calculations, and each procedure must be related to the investment process.

Performance attribution could also be used for risk-adjusted returns, since it allows the sources of return to be identified, and could provide useful information to correct differential return. In credit risk sources of return, for example, there are two sub-sources. The first one is the spread cushion, which is not a fair result against the benchmark, i.e. it is not necessarily a result of active management, since agencies and supranationals have a higher yield than Treasuries. Nevertheless, spread products results could be the consequence of active management, as managers change positions in agencies and supranationals in order to take advantages on spreads movements.

Finally, this paper suggests as a topic for further studies a link between performance attribution and performance evaluation as a way of analysing numbers on a statistical basis searching for long run performance patterns. Therefore, it can be identified whether results comes from a good standard of decision or not. In fact, the question to be answered is whether active management really aggregates results in the investment strategy, or if results are only the consequence of a random process or a consequence of a choice for more risk.

References

AIMR (2001); "AIMR Performance Presentation Standards"; Association for Investment Management and Research – AIMR

Bodie, Z., Kane A. and Marcus, A. (1999); *Investments*; McGraw-Hill

Breukelen, G. (2000); "Fixed Income Attribution"; *Journal of Performance Measurement*; Summer; pp. 61-68

Clarfeld, R. A. (1998); "Time-weighted Returns"; *Journal of Accountancy*; February

Dynkin, L., Hyman, J. and Vankudre, P. (1998); "Attribution of Portfolio Performance Relative to an Index"; Lehman Brothers Fixed Income Research; Working Paper

Douglas, L. (1988); *Yield Curve Analysis: The Fundamentals of Risk and Return*; New York Institute of Finance

Elton, E. and Gruber, M. (1995); *Modern Portfolio Theory and Investment Analysis*; John Wiley

Eichengreen, B. and Mathieson, D. (2000); "The Currency Composition of Foreign Exchange Reserves: Retrospect and Prospect"; IMF Working Paper WP/00/131.

Goldenstein, A. (1994); "La Administración de Activos Internationales em la Banca Central"; Reuniones y Seminarios – CEMLA.

Lichtbuer, P. (1994); "Overview of Reserve Management by Central Banks"; Reuniones y Seminarios – CEMLA.

Ramaswamy, S. (2000); "Risk and Performance Attribution Report for Fixed Income Portfolio Mandates"; BIS Banking Papers Issue 2; Bank for International Settlements.

Reilly, F. and Brown, K. (2000); *Investment Analysis and Portfolio Management*; Dryden Press

Thompson, I. (1997); "How to Successfully Develop and Implement a Performance System"; *Journal of Performance Measurement*; Winter

Treynor, Jack and Mazuy, M. (1965); "Can Mutual Funds Outguess the Market?" *Harvard Business Review*, 43, No. 1, Jan./Feb, pp. 63-75.

20

Management of the international reserve liquidity portfolio

D. Delgado Ruiz, P. Martínez Somoza, E. Osorio Yánez,
R. Pabón Chwoschtschinsky

Abstract

The purpose of this paper is to offer a methodology related to the management of Banco Central de Venezuela's liquidity portfolio of its international reserves. The model establishes the desired size of the liquidity portfolio given the liquidity risk that the institution is willing to tolerate, and considers the trade-off between the amount of liquid funds and the probability of a liquidity crisis. This methodology is useful to measure and control the liquidity risk associated with BCV's daily holdings of funds[1].

1 Introduction

The purpose of this article is to describe the methodology developed by Banco Central de Venezuela, (BCV) with regard to the management of its liquidity portfolio. This approach enables us to establish the desired size of the liquidity portfolio and to measure and control the liquidity risk associated with daily holdings of funds.

In assessing the liquidity risk framework, it is important to bear in mind some characteristics of the Venezuelan economy. Firstly, the country's main source of inflows comes from oil exports, which necessarily brings volatility to Venezuela's international reserves. This volatility is exacerbated by the fact that these revenues represent the highest proportion of inflows to the international reserves (85% on average for the last four years). Secondly, as the dollars provided by oil have to be sold to the central bank, BCV cannot avoid selling them back to the economy in order to import the goods and services not produced internally. By doing so, BCV plays the unusual additional role of being the main source of foreign exchange for the economy.

Likewise, since the beginning of the 1980s BCV has attempted to attain price stability by reinforcing its monetary policy, with the same objective being assigned to the country's foreign exchange policy. BCV has accordingly been very active in the domestic forex market through various quite different regimes, ranging from exchange controls (during three different periods) to managed systems and free floating.

Within these institutional arrangements, BCV has been given the mandate to centralise and to manage the country's international reserves. The main goals for the reserves are to:

- Ensure continuity of international payments
- Support foreign exchange policy
- Reduce vulnerability towards external shocks.

Taken together, these objectives demonstrate that BCV places a larger than usual emphasis on liquidity in its reserves management, a feature further reinforced by a current internal constraint related to capital preservation, whereby assets can only be sold at prices higher than the purchased price.

[1] We would like to thank Guadalupe Asprino, Begoña Arróspide, Germán Utreras, Eduardo Zambrano, Claudio Rocco, Roberto de Beaufort and Ronald van der Wouden for helpful discussions.

The document is organised in three sections. The first part refers to the management of the liquidity portfolio, in which the methodology for measuring the liquidity risk and establishing the number and size of liquid tranches is described. The second part contains some numerical examples in order to illustrate the concrete implementation of the system. Finally, the third part draws some conclusions.

2 Management of the liquidity portfolio

2.1 Theoretical background

In this study, the term "liquidity" refers to the immediate availability of funds without the need to sell financial instruments under inappropriate conditions. This aspect is particularly relevant for BCV's reserves management, as BCV is not permitted to sell instruments at a loss. The trading of negotiable instruments should be motivated by return maximisation criteria, and not by liquidity concerns.

Therefore, the liquidity portfolio is only composed of short maturity instruments, and the liquidity risk refers to the probability of a liquidity crisis occurring, whereby the liquid portfolio would be unable to meet the outflow requirements.

The probability of a liquidity crisis diminishes as the amount of liquid funds increases. However, there is a cost to maintaining excess liquidity, as long as a positive relationship exists between the return of an instrument and its maturity (positive sloped yield curves). The goal is to find a systematic rule in order to determine appropriate liquidity levels for BCV's international reserves. These liquidity levels are represented by tranches, which all together integrate the liquidity portfolio.

In this context, the developed model establishes different alternatives to choose the desired size of the liquidity portfolio given the liquidity risk that the institution is willing to tolerate, considering the trade-off between return and maturity.

It does so by setting the elements of a dynamic system in which the hard currency demand, represented by net daily changes in international reserves, is treated as a random variable. The sequence, sign and size of daily observations are used to estimate the probability that a liquidity crisis will occur.

In contrast, a static analysis, such as a histogram of daily changes in international reserves, would not be appropriate to establish the required size of liquid funds, because it would ignore the time path of possible combinations of inflows and outflows.

A histogram calculated using data referring to periods longer than one day (for example monthly data) could potentially be misleading, as the resulting percentile might not necessarily reflect intra-month changes. Even though it satisfies at a certain confidence level the net demand of reserves for the month, some daily net changes in reserves can be greater than the amount defined. Consequently, the probability of a crisis would then be underestimated.

A percentile derived from a daily histogram provides information about the amount of liquid funds that has to be maintained every day. However, this constant amount does not provide any information about how and when to replenish and restructure the liquidity tranche. In fact, the only possible way to include such information is to create a system that considers the dynamic nature of daily inflows, outflows and transfers between liquidity tranches.

The defined system incorporates the fact that small daily negative changes in the reserves can lead to a liquidity crisis, if such changes occur consecutively over several days. Conversely, significant negative changes in international reserves can be followed in subsequent days by positive changes and transfers between liquidity tranches, thereby avoiding a liquidity crisis.

Another drawback to a static analysis such as a histogram approach is that it does not provide information regarding the likelihood of a liquidity crisis during a predefined horizon different than the one used to calculate the histogram. Conversely, in our model the estimation of the conditional probability answers this issue by determining the likelihood of a crisis occurring in a given period of time (the next seven days in our example).

The daily net change in reserves is simulated using a Monte Carlo simulation technique. This process enables us to determine how the fitted hard currency demand dynamically affects the daily level of the liquid tranches, which in turn allows us to quantify the number of unsatisfied requests for funds.

The Monte Carlo simulation is applied to replicate the empirical data, as several theoretical distributions were fitted but none of them were statistically appropriate.[2]

In this study, the liquidity portfolio is comprised of overnight and one week deposits.[3] A target level is set for both types of deposit, and a specific rule is defined for the transfer of funds between them in order to accomplish the mentioned target. Daily holdings of liquid tranches can temporally deviate from the target level owing to net demand on the international reserves.

A crisis will occur when the liquidity portfolio cannot satisfy the outflows of reserves. As BCV must guarantee the continuity of international payments, the crisis is considered to occur when it is necessary to find funds outside the system. The probability of a liquidity crisis therefore indicates that it would be necessary to find funds outside the liquidity system in order to solve the liquidity crisis the same day, for example by selling negotiable instruments of its optimisation portfolio, borrowing. etc.

Hence, the methodology is a toolkit that allows us to control the liquidity risk by selecting the level of liquidity targets and/or number of tranches for a chosen level of liquidity risk that the institution is willing to accept. A higher level of liquidity target and/or number of tranches implies a lesser probability that a liquidity problem will occur.

To guarantee that the parameters of the probabilistic distribution are constant through time, the daily net change in international reserves needs to follow a stationary process. Moreover, if this condition is satisfied, the results of the system will be validated regardless of the level of international reserves.

To verify the stationarity of the process, Dickey-Fuller[4] and Phillips-Perron[5] tests were applied. These are unit root tests used to establish the order of integration of the daily changes in the international reserves. Empirical values calculated from the daily data were compared to the theoretical ones; the results support the constant parameter hypothesis.

---

[2] Extreme Value, Logistic, Weibull, Uniform, Triangular, Normal were used; however, the series did not satisfy the Kolgomorov-Smirnov and Chi-square goodness-of-fit tests, performed to find the best fit between theoretical and empirical distributions. Since the data contain negative numbers, some distributions were omitted.

[3] It is possible to include additional liquidity tranches, for example a one-month deposit or other instruments (negotiable), depending on the requirements.

[4] Dickey, D. and W. Fuller. 1979. "Distribution of the Estimators for Autoregressive Time Series with a Unit Root". *Journal of the American Statistical Association*, 74, pp. 427-31.

[5] Phillips, P. and P. Perron. 1988. "Testing for a Unit Root in the Time Series Regression". *Biometrika*, 75, pp. 335-46.

2.2. Liquidity risk model

The parameters of the model are the desired liquidity levels and the probability of a liquidity crisis. The mentioned targets are initially set at an arbitrary amount for each tranche. These initial levels, together with the simulated demand and the defined rules for rebalancing the tranches, result in a number of cases in which initial funds were not sufficient to satisfy currency requests. This procedure is run until the amount of liquid funds to start with generates a tolerable probability of a liquidity crisis.

2.2.1 Probability of a liquidity crisis

This section describes the liquidity risk model's mathematical formulation, considering its daily and weekly dynamic.

For each period, the amount of the overnight tranche is equal to the sum of its previous level, ON_{t-1} and the net change in the international reserves (net request), ΔR_t, whenever this sum is less than or equal to the liquidity target level of the overnight, ON^*.

$$ON_t = \begin{cases} ON_{t-1} + \Delta R_t & \text{if} \quad ON_{t-1} + \Delta R_t \leq ON^* \\ ON^* & \text{otherwise} \end{cases} \tag{1}$$

When the sum is greater than the target level for overnight, this position would be equal to the target level. The surplus, EX_{ON_t}, would be transferred to a one-week deposit tranche.[6]

$$EX_{ON_t} = \begin{cases} ON_{t-1} + \Delta R_t - ON^* & \text{if} \quad ON_{t-1} + \Delta R_t > ON^* \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

On the other hand, transfers take place from the daily overnight tranche to support the one-week tranche whenever there is a surplus in overnight funds in relation to the target. The system only considers the amount that is necessary to complete the target level of the one-week deposits, W^*.

$$W_t = \begin{cases} W_{t-1} + EX_{ON_t} & \text{if} \quad W_{t-1} + EX_{ON_t} \leq W^* \\ W^* & \text{otherwise} \end{cases} \tag{3}$$

[6] For reasons of simplicity, it is assumed that all the deposits mature on the same day at the end of the week.

In addition, it is possible to register a surplus in the one-week deposit, EX_{Wt}, with respect to its desired amount. The surplus is transferred to the optimisation portfolio (Opt_t)as it occurs and leaves the system.

$$EX_{W_t} = \begin{cases} W_{t-1} + EX_{ON_T} - W^* & if \quad W_{t-1} + EX_{ON_T} > W^* \\ \\ 0 & otherwise \end{cases} \tag{4}$$

$$Opt_t = Opt_{t-1} + EX_{wt} \tag{5}$$

Every week the overnight tranche can be fed with the transfers, $Transf_W$, from the one-week deposits, W_t, so that the capacity of the overnight funds approaches the target.

$$Tranf_W = \begin{cases} ON_t - (ON_{t-1} + \Delta R) & if \quad W_t \geq ON^* - ON_{t-1} + \Delta R \geq 0 \\ W_t & if \quad ON^* - (ON_{t-1} + \Delta R) > W_t \\ 0 & if \quad ON^* - (ON_{t-1} + \Delta R) < 0 \end{cases} \tag{6}$$

$$ON_t = ON_{t-1} + \Delta R + Transf_{wt} \tag{7}$$

With a sufficient number of iterations of the simulation procedure described above, it is possible to calculate the daily probability of a crisis arising, i.e. the probability that the overnight and one-week deposits are not enough to satisfy the net outflows faced by the monetary authority.

The probability of a liquidity crisis is calculated by the empirical frequency of this event occurring. This result is based on De Moivre-Laplace's theorem: the weak law of large numbers shows the convergence of the empirical frequency to the probability.[7] This theory requires the series of events to be independent.[8]

$$\text{Prob (liquidity crisis)} = \frac{n^\circ \ liquidity\ crisis}{n^\circ \ observations}$$

[7] Goldfarb, B. and C. Pardoux. 1993. *Introduction à la Méthode Statistique*. Gestion. Économie. Dunod. p. 149.

[8] In case the daily changes in the international reserves are not independent from each other, the probability of a liquidity crisis can be estimated considering the autoregressive structure of the variable.

Figure 1: below illustrates the proposed liquidity risk system.

Figure 1: Liquidity risk system

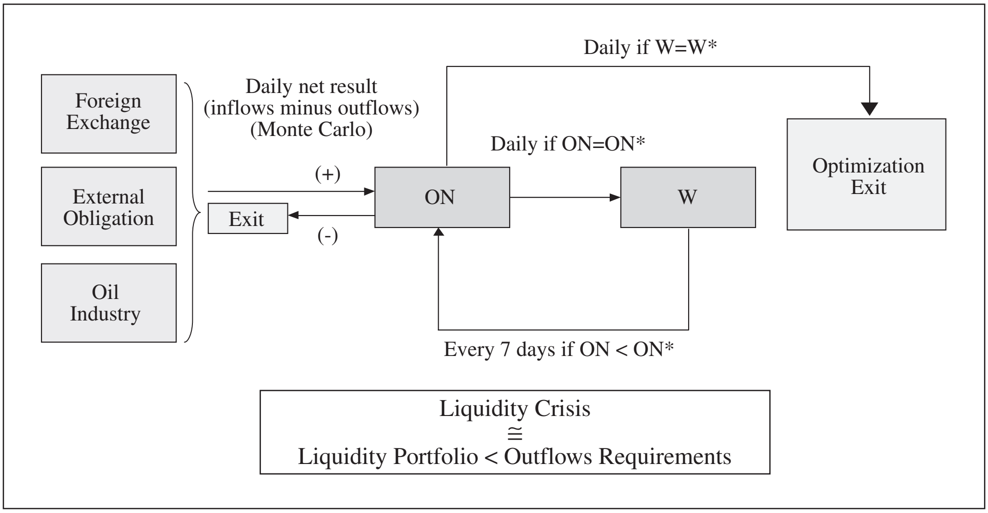

When there is a liquidity crisis in the model, the funds destined for overnight and one-week deposits are reset to the desired levels, *ON** and *W**. This allows for the control of the number of times that the targets face a liquidity crisis as well as the necessity to sell negotiable instruments to cover the deficit and re-establish the target levels.

It is important to consider that a negative relationship exists between the probability of a liquidity problem and the level of the liquidity targets, as Figure 2 shows:

Figure 2: Determining the liquidity risk

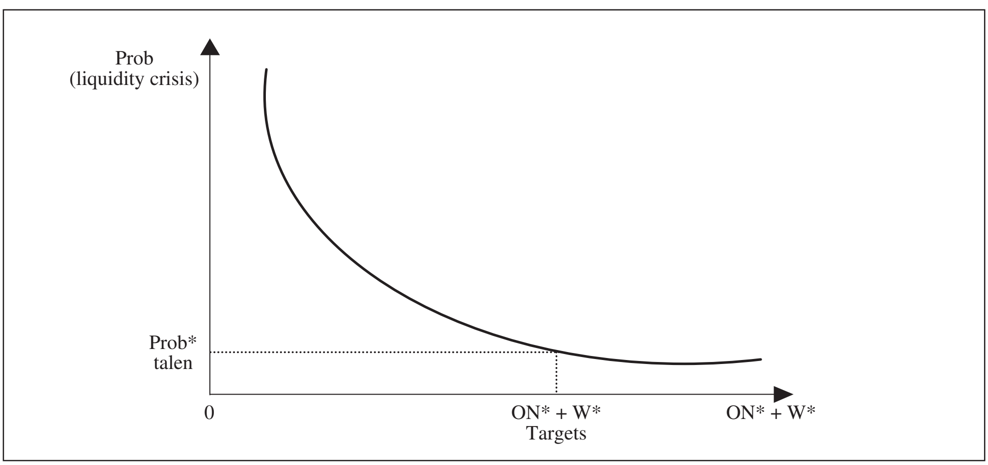

In this sense, the liquidity risk framework provides the tools to establish the desired liquidity amount that should be kept in the portfolio with regard to a specific level of liquidity risk.

2.2.2 Conditional probability of a liquidity crisis

For monitoring reasons, it could be useful to estimate the probability of a crisis situation for a given amount of overnight funds (x) during a particular day and within a determined period (p).[9] The daily probability provides an indication of how many liquidity crises can occur, but it does not provide information the timing during the mentioned period (p) when the crisis is going to happen. This information is instead provided by the conditional probability of a liquidity crisis.

The same frequentist probability theory approach mentioned earlier is applied to the determination of the conditional probability. The monitory horizon is established according to current needs on a weekly basis.

$$\Pr ob\,(weekly\ liquidity\ crisis\,/\,ON = x) = \frac{\Pr ob\left(crisis\,(weekly\,horizon) \bigcap ON = x\right)}{\Pr ob\left(ON = x\right)} \qquad (9)$$

The horizon is calculated by dividing the number of liquidity crises within a certain horizon by the number of observations, given a certain overnight amount:

$$\Pr ob\,(weekly\ liquidity\ crisis\,/\,ON = x) = \frac{n^{o}\ crisis\,(weekly\ \ horizon) \bigcap ON = x}{n^{o}\ observations\ \ ON = x} \qquad (10)$$

In this step of the process, the generation of a larger amount of data using a Monte Carlo simulation gains relevance because of the decrease in the number of observations.

3 Numerical examples

3.1 Estimation of the probability of a liquidity crisis

In this section the previously developed methodology is illustrated by a hypothetical numerical example in order to show its practical implementation. Tables 1 and 2 show the dynamic nature of the liquidity portfolio, structured by the overnight and one-week deposit tranches, and how they interact by transferring resources between each other. The overnight tranche can access the one-week deposit tranche once a week.

The first step is to simulate the distribution of the daily changes in international reserves representing the country's net demand of funds, using the Monte Carlo method.[10]

The next step is to set the initial target levels for overnight and one-week deposits. As an example, Table 1 shows these targets at 1,000 and 300 units of hard currency respectively. We start with this level of liquidity on Day 1. On Day 2, the net change in the international reserves is +306. As the overnight position equals the target level, the system transfers this amount to the one-week tranche (outflow to one-week deposits on Day 2, Table 1), and then to the optimisation portfolio, owing to the fact that the position in the second tranche satisfies its target level.

---

[9] As it is unlikely to find the exact amount of overnight within the sample, we use a range: [x-k; x+k].

[10] In our example, 16,000 observations were simulated.

On Day 6, for example, the net demand for hard currency is satisfied by the overnight tranche, implying that the remaining amount, 953, is less than its target level. On Day 7, the net change of reserves is +9 units of hard currency, but is not enough to reach the desired level of overnight. As a result, part of the one-week deposits that are maturing that day of the week, 38 units, are transferred to tranche 1 in order to re-establish the overnight target level (inflow from one-week deposits in Table 1, and outflow to overnight in Table 2).

On Day 8, the one-week tranche amount once again reaches its desired level thanks to the resources received from the overnight tranche (outflow to one-week deposits in Table 1 and inflow from overnight in Table 2). The surplus in the one-week deposits is +16 units of hard currency, which is sent to the optimisation portfolio.

An example of a liquidity crisis takes place on Day 63 as shown in Table 1, when the amount of overnight is insufficient by 734 units of hard currency to meet the liquidity requirements of the economy. It should be noted that in Table 2 the amount of the one-week tranche is zero, owing to the transfers previously made to the overnight tranche. Each time a liquidity crisis occurs, the initial values of the overnight and one-week deposits are reset to their target levels. We estimate the probability of a liquidity crisis by dividing the number of crises by the total number of observations, as indicated in equation (8).

This simulation procedure is performed several times by setting new target levels of overnight and one-week deposits in order to obtain several alternatives to choose from. This therefore establishes the parameters of the model, which are the desired target liquidity levels and the resulting crisis probability, according to the liquidity risk (confidence level) that the institution is willing to face.

Table 1: Temporal sequence of overnight positions (Units of Hard Currency)

Day	Initial overnight amount	Net change in reserves (simulation)	Inflows from 1 week deposits	Outflows to 1 week deposits	Final overnight amount
1	**1000**	0	0	0	1000
2	1000	306	0	306	1000
3	1000	8	0	8	1000
4	1000	53	0	53	1000
5	1000	7	0	7	1000
6	1000	-47	0	0	953
7	953	**9**	**38**	**0**	**1000**
8	1000	54	0	54	1000
9	1000	-26	0	0	974
10	974	-19	0	0	955
-	-	-	-	-	-
-	-	-	-	-	-
-	-	-	-	-	-
62	**436**	**-268**	**0**	**0**	**167**
63	167	-901	0	0	-734
64	1000	-19	0	0	981
65	981	51	0	32	1000
66	1000	-63	0	0	937
67	937	-74	0	0	864
68	864	-48	0	0	816
-	-	-	-	-	-
-	-	-	-	-	-
-	-	-	-	-	-
16000	784	0	0	0	785

Table 2: Temporal sequence of one-week desposits (Units of Hard Currency)

Day	Initial 1 week deposits amount	Inflows from overnight	Outflows to overnight	Outflows to optimisation	Final 1 week deposits
1	300	0	0	0	300
2	300	306	0	306	300
3	300	8	0	8	300
4	300	53	0	53	300
5	300	7	0	7	300
6	300	0	0	0	300
7	**300**	**0**	**38**	**0**	**262**
8	262	54	0	16	300
9	300	0	0	0	300
10	300	0	0	0	300
-	-	-	-	-	-
-	-	-	-	-	-
-	-	-	-	-	-
62	**0**	**0**	**0**	**0**	**0**
63	0	0	0	0	0
64	300	0	0	0	300
65	300	32	0	32	300
66	300	0	0	0	300
67	300	0	0	0	300
68	300	0	0	0	300
-	-	-	-	-	-
-	-	-	-	-	-
-	-	-	-	-	-
16000	132	0	0	0	132

As for the weekly conditional probability of a liquidity crisis, we search for the given amount of overnight in the total observations, and then calculate the frequency of a crisis within the next seven days.

3.2 Liquidity crisis histogram

In addition to estimating the conditional probability, it is also useful for monitoring reasons to consider the potential size of the liquidity crisis. This can be found in Table 1 (Final Overnight Amount column), corresponding to negative values. These values represent, in the simulation, the amount in hard currency unattended by* the monetary authority's liquidity portfolio, for a specific and hypothetical liquidity target level. The following histogram reveals the size of the most frequent and extreme liquidity crisis.

Figure 3: Histogram of liquidity crisis

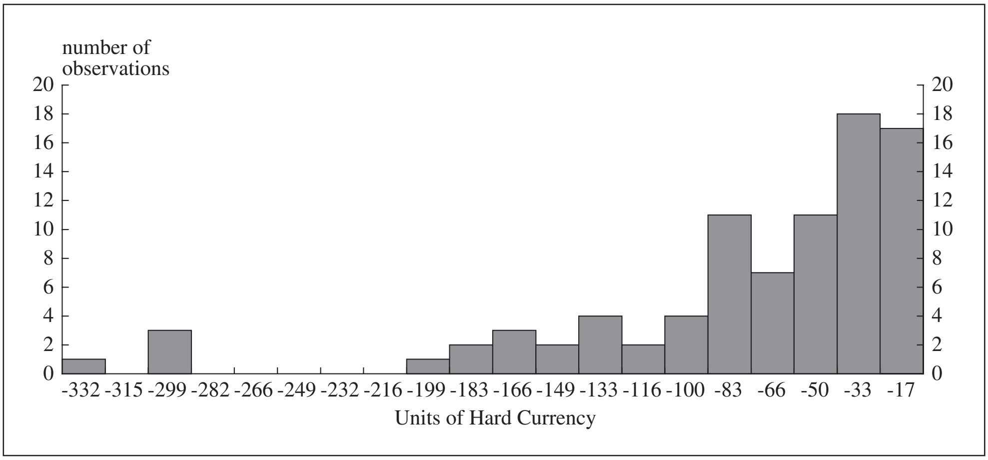

To solve the liquidity problem, the central bank would need to use the instruments that are maturing and/or sell tradable instruments.

4 Conclusions

The model establishes the desired size of the liquidity portfolio given the liquidity risk that the institution is willing to tolerate, considering the dynamic nature of flows and the trade-off between the amount of liquid funds and the probability of a liquidity crisis.

The framework is a simulation tool that determines the probability of a liquidity crisis for the desired size and/or daily holdings of the international reserve portfolio.

By changing the target level of liquidity, it is possible to have several sets of alternatives in order to select the desired size of the liquidity portfolio, which is consistent with the liquidity risk profile of the central bank.

The model has a transparent rule to re-establish the target level of liquidity, based on the transfers between liquid tranches.

The estimation of the probability of a liquidity crisis working with the daily changes in international reserves, is a reliable estimation because the series follows a stationary process*, and because the Monte Carlo simulation technique generates a large sample of values all over its distribution.

The model is flexible because different scenarios of the daily change in the distribution of the international reserves can be incorporated into the system. The scenarios can include a structural change in the behaviour of the economy. It is also possible to modify or add different rules of interaction between liquidity tranches.

References

Dickey, D. and W. Fuller, 1979. "Distribution of the Estimators for Autoregressive Time Series with a Unit Root", *Journal of the American Statistical Association,* 74.

Erlang, A. K., 1909. " The Theory of Probabilities and Telephone Conversations", *Nyt Tidsskrift for Matematik B*, Vol. 20.

Erlang, A. K., 1917. "Solution of Some Problems in the Theory of Probabilities of Significance in Automatic Telephone Exchange", Elektrotkeknikeren, Vol. 13.

Goldfarb, B. and C. Pardoux, 1993. *Introduction à la Méthode Statistique.* Gestion. Économie. Dunod.

Phillips, P. and P. Perron, 1988. "Testing for a Unit Root in the Time Series Regression". Biometrika, 75, pp. 335-46.

Taha, H., 1997. *Operations Research: An Introduction.* 6ª Ed. Prentice-Hall.

Jackel, P., 2002. *Monte Carlo Methods in Finance*. Wiley Finance.

Makridakis, S., S. Wheelwright and R. Hyndman, 1998. *Forecasting: Methods and Applications*. John Wiley & Sons.

21

Determining neutral duration in the Bank of Israel's dollar portfolio[1]

Janet Assouline

Abstract

The holding period return on the reserves portfolio is determined to a great extent by the parameters of the benchmark portfolio. This underlines the importance of the way the benchmark's parameters are established, as well as their periodic review in order to bring them into line with changing market conditions, as necessary. Duration is one of the salient parameters of the benchmark portfolio, and exerts considerable influence over holding period return. In 1998 the Bank of Israel (BoI) undertook to research what the desired duration of the benchmark portfolio should be. In this study a method of determining the duration of the BoI's benchmark dollar portfolio using the shortfall approach was presented. According to this method, the target duration was determined in such a way that at a given probability, the annual holding period return on the portfolio would not fall below a minimum desired threshold. The portfolio manager was required to set three parameters to determine the target duration of the portfolio which reflect his/her preferences: 1) the most appropriate minimum threshold level (in this study, the minimum threshold was defined as half the return on a risk-free asset); 2) the confidence level at which the portfolio yields a higher return than the minimum threshold level; and 3) the criterion at which the minimum threshold is met over the sample period as a whole. After choosing the parameters and running the model accordingly, the results showed that according to the preferences of the portfolio manager, the duration of the BoI's dollar portfolio should be 16 months.

1 Introduction

The composition of the Bank of Israel's reserves portfolio is based mainly on that of the benchmark portfolio as regards duration, currency composition, asset composition, and distribution along the curve. Positions taken against the benchmark represent only a small share of the portfolio's composition. As a result, the holding period return on the reserves portfolio is determined to a great extent by the parameters of the benchmark portfolio. This underlines the importance of the way the benchmark's parameters are established, as well as their periodic review, bringing them into line with changing market conditions, as needed.

One of the salient parameters of the reserves portfolio, and one which exerts considerable influence over holding period return, is duration. For several years before the study described here, the neutral duration[2] of the dollar portfolio was 12 months. In 1998 the Foreign Currency Department of the Bank of Israel undertook to research what the desired neutral duration of the dollar portfolio should be using the shortfall approach. According to this method, neutral

---

[1] This paper is based on a study undertaken in the Foreign Currency Department of the Bank of Israel in 1998. The findings were presented to the Foreign Currency Committee for decision-making.

[2] Neutral duration is the risk-neutral position of the portfolio relative to changes in interest rates.

duration is determined in such a way that, at a given probability, the holding period return on the portfolio will not fall below a minimum level determined by the portfolio owner or manager[3].

The shortfall approach and its implementation are described in Sections 2 and 3 of this paper. Section 4 focuses on determining the desired minimum threshold. In Sections 5, 6 and 7 the results of the model implementing the shortfall approach are presented. Section 8 examines the model for an out-of-sample period, and in Section 9 the results of the model are applied to the Bank of Israel's benchmark dollar portfolio. Section 10 contains a backward-looking analysis to this paper, written at a later date. It discusses the actual application of the results of the model, as well as the results that would have been obtained had the model been run over a longer period which includes the last five years.

2 The shortfall approach

Duration is a measure of the sensitivity of fixed-income assets and portfolios to changes in interest rates. Fluctuations in interest rates create capital gains or losses, which are reflected in the holding period return of the asset or portfolio. The longer the duration of the asset or portfolio, the greater its capital gains in a declining interest rate environment, but the greater its capital losses in a rising interest rate environment.

In general, a rational, risk-averse portfolio manager views interest rate risk in an asymmetric manner: he or she would like to fully benefit from capital gains in a declining interest rate environment, but to limit capital losses in a rising interest rate environment. Since he or she is unable to do this, the portfolio manager aspires to structure his or her portfolio such that he or she will suffer from capital losses only up to a certain level, even though this is at the expense of reducing potential capital gains. The level beyond which the portfolio manager is not prepared to forgo any more return is called the "minimum threshold". Section 4 discusses how this is established.

In order to ensure a minimum level of return at a given probability, the expected distribution of future returns must first be established. Once this has been determined, it is possible to set the duration of the portfolio such that at a certain confidence level – e.g. 95% – the return on the portfolio for a given period will not fall below the minimum threshold. In the present study, the period is one year, since portfolio managers tend to look more closely at this period, as it displays less volatility than shorter periods and, possibly, also as a result of the end-of-year reporting requirement. This approach to setting duration is called the shortfall approach.

3 The model

a. The database

Securities: the model is based on the weekly yields to maturity of on-the-run US dollar Treasury notes. In order to calculate the yields for those terms to maturity that do not exist in the on-the-run market, a linear interpolation of the yields that exist in the market was undertaken.

The period of the study is June 1984-June 1998, weekly observations.

[3] In general the portfolio owner defines the risk profile and sets the investment parameters accordingly. However, in the case of a portfolio comprising a country's foreign currency reserves, the country is really the portfolio owner and cannot (realistically) determine the portfolio's parameters. In this case, the portfolio manager, i.e. the central bank, must define the country's risk profile in place of and on behalf of the country. For this reason, the term used here is portfolio manager rather portfolio owner.

b. *Explanation of the model*

The problem is to ascertain the best way of estimating the distribution of holding period returns for the ensuing year on Treasury notes of various durations, so that the shortfall criterion can be used when choosing a neutral duration. The annual holding period return, one year forward, at a given point in time is a function of the level of yields to maturity at that point and the path of future changes in yields to maturity during the year. As the level of yields to maturity at each point is known, only the changes in yields to maturity during the ensuing year need to be estimated. The path of future changes in yields to maturity was estimated on the basis of the distribution of past observations, assuming that the distribution of changes in the year ahead is the same as that of past changes. The question that now must be answered is how many past observations should be used to estimate most effectively the future distribution of changes in yields to maturity. While there are statistical methods for examining this problem, the empirical approach of trial and error was chosen (see Appendix). Empirical examination indicates that the best estimate is obtained from observations of *one year retroactively*.

c. *Estimating the distribution of annual holding period returns*

This section describes the process of estimating the distribution of annual holding period returns on securities (Treasury notes) of various durations at different points in time, across the sample period. To make matters clearer, the reference is to a specific point in time, i.e. the estimation of the distribution at time *t*, and this process is repeated at monthly intervals throughout the sample period. The stages for estimating the distribution of the annual holding period return on securities of varying durations across the yield curve at time *t* are as follows:

1. First, the mean and standard deviation of changes in yields to maturity on the securities at a specific point on the yield curve were calculated, using weekly data for one year retroactively. The working hypothesis was that the distribution of changes in yields to maturity is normal, so that it is possible to define this distribution using the two parameters which determine a normal distribution (a sample of different periods and securities of varying durations showed that this hypothesis was reasonable).
2. On the basis of this distribution of changes in yields to maturity, successive random samples of 52 observations (equal in number to the weeks in a year) were taken, representing the weekly changes in the yields. This gave the path of changes in yields to maturity of specific securities for one year forward from time *t*.
3. On the basis of the random path of yields to maturity in the sample, the holding period return for one year forward was calculated. This represents one random observation from the hypothetical distribution of annual holding period returns on securities with a given duration, at time *t*.
4. To obtain the entire distribution, a Monte Carlo simulation was then used to take numerous samples, in accordance with the procedure described above. This produced a distribution of annual holding period returns on securities, representing a specific point on the yield curve at time *t*.
5. To estimate the distribution of holding period returns on different securities spread over the *entire* yield curve, the above procedure was applied for the various securities at each of the points on the yield curve at intervals of one month to maturity.
6. The distribution of annual holding period returns on various securities across the *entire* yield curve was obtained, each security having a different duration. From all these distributions, those that maintain the shortfall criterion were chose, i.e. those for which at least 95% of the area of each distribution was above the minimum threshold.

7. From all the distributions that maintain the shortfall criterion, the one in which the securities had the longest duration was chosen[4] and defined as 'neutral' at time *t*.

To obtain a time series of neutral durations, the entire estimation process was implemented for different points at one-month intervals over the sample period (June 1985-June 1998). The series obtained was therefore dependent on the two parameters which had to be determined in advance: the *minimum threshold* and the *confidence level* at which it is certain that the holding period return will not fall below the minimum threshold. The next section addresses possible ways of determining the minimum threshold.

4 Determining the minimum threshold

The minimum threshold is the minimum annual holding period return which a portfolio manager is prepared to accept. Three types of minimum thresholds are defined in this paper:

a. *Fixed minimum threshold*: this threshold does not change over time and is not dependent on the level of interest rates in the market (e.g. 1% or 0%).

b. *Threshold dependent on the yield to maturity on a risk-free asset*: this threshold is defined as the yield to maturity on a short-term security (which can be treated as a risk-free asset), *less* a fixed spread (for example, the yield to maturity on a risk-free asset less 2 percentage points). The rationale behind this kind of threshold is to ensure a minimum holding period return that reflects interest rates in the market. Naturally, the minimum threshold will be lower than the return on the risk-free asset in order to make it possible to obtain a higher return than that on a risk-free asset at a time of declining interest rates.

c. *Threshold linked relatively to the yield to maturity on a risk-free asset:* this threshold is defined as a specific percentage of the yield to maturity on a short-term or risk-free asset (e.g. 50% of the yield to maturity on a risk-free asset). The rationale behind this minimum threshold is that the return which is forfeited when interest rates are low is smaller, relative to the yield on the risk-free asset, than that which is forfeited when interest rates are high. Thus, a threshold of this kind provides better protection in times of lower interest rates than the previous one, and at times of higher interest rates, gives the portfolio manager more flexibility when setting neutral duration.

In this paper we define seven possible minimum thresholds, calculating a neutral duration over the entire sample period for each one:

Threshold 1: Yield to maturity on 3-month Treasury bills *less* 2 percentage points.
Threshold 2: Yield to maturity on 3-month Treasury bills *less* 3 percentage points.
Threshold 3: 50% of the yield to maturity on 3-month Treasury bills.
Threshold 4: 65% of the yield to maturity on 3-month Treasury bills.
Threshold 5: Fixed threshold of 2%.
Threshold 6: Fixed threshold of 1%.
Threshold 7: Fixed threshold of 0%.

The first two thresholds are dependent on the yield to maturity of the risk-free asset (type b). The next two are linked relatively to the yield to maturity on a risk-free asset (type c), while the last three are fixed thresholds (type a).

---

[4] The choice of the distribution of securities with the longest duration stems from the fact that over time a portfolio with a longer duration will yield a higher holding period return than one with a shorter duration (however, the return on the longer duration portfolio will be more volatile than that on a shorter duration portfolio, and in some short periods the longer duration portfolio may yield a lower return). Section 7 tests this claim empirically for the sample period.

Figure 1: Neutral durations: Time series derived from threshold 3
50% of the yield to maturity on a 3-month Treasury-bill

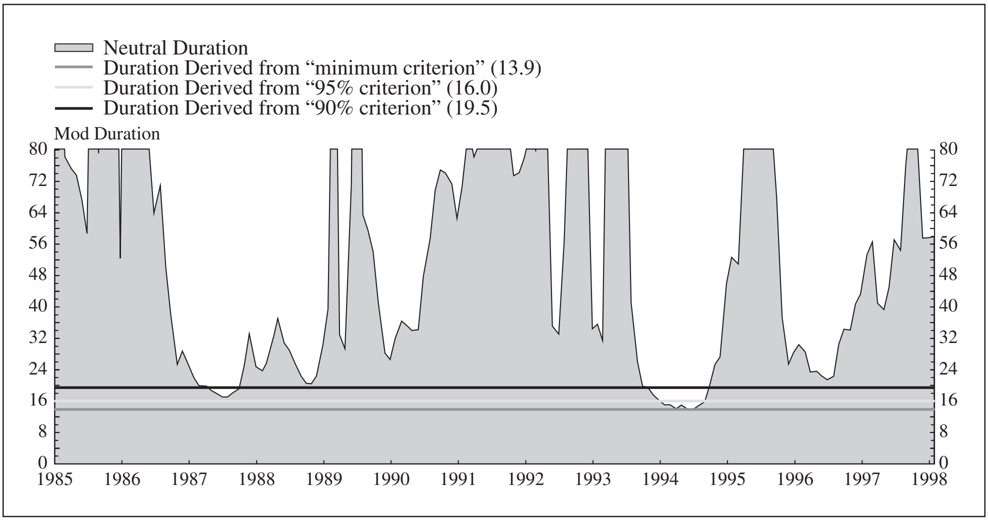

5 Results

For every minimum threshold, a calculation of neutral durations was carried out based on the model, for every point in time in the sample period. The results of the model are shown in Figure 1[5]:

The line in Figure 1 which borders the coloured area represents (at each point in time) the neutral duration calculated for the shortfall criterion at a 95% confidence level, and the area beneath it defines the range of the durations which are lower than the neutral duration and meet the shortfall criterion. As the graph shows, the series is very volatile, varying widely from month to month. Since the duration of the benchmark portfolio should remain stable over time and not change frequently, it was impossible to determine directly from this data what the benchmark duration of the dollar reserves portfolio should be. Consequently, it was necessary to define another criterion for choosing a neutral duration so that it would be stable over a relatively long time. Three different possibilities were examined:

1. *The minimum neutral duration criterion*: the duration determined by this criterion would be the lowest limit of the series of neutral durations throughout the sample period. In this manner, by setting the neutral duration of the dollar portfolio at the beginning of the sample period and leaving it unchanged for 13 years, the annual returns on the portfolio would be above the minimum threshold for 95% of the moving annual returns.
2. *The 95% neutral duration criterion*: this criterion is slightly less stringent than the previous one. The duration of the portfolio determined according to it will be below 95% of the observations in the series of neutral durations throughout the sample period.
3. *The 90% neutral duration criterion*: the duration of the portfolio to be determined according to it will be below 90% of all the observations in the series of neutral durations throughout the sample period.

[5] From this point onward only those graphs related to Minimum Threshold 3 (the threshold which was ultimately chosen) will be displayed. The process of choosing the appropriate threshold will be discussed in Section 9 below.

Figure 2: Neutral durations and their rates of return calculated using minimum tresholds and different criteria, June 1985-June 1998

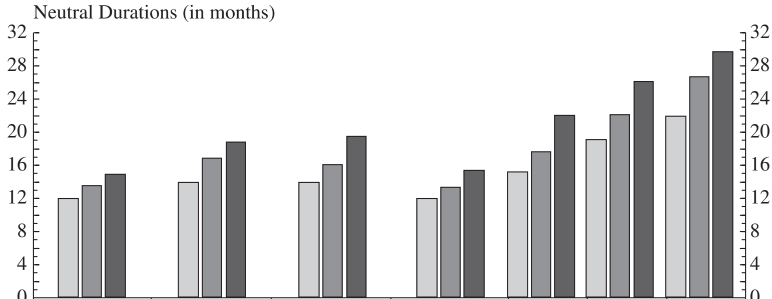

	3m. Treasury Bill less 2%	3m. Treasury Bill less 2%	50% of 3m. Treasury Bill	65% of 3m. Treasury Bill	Fixed 2%	Fixed 1%	Fixed 0%
□ Minimum Criterion	11.9	13.9	13.9	11.9	15.1	19.0	21.9
Rate of Return	6.8	6.9	6.9	6.8	6.9	7.3	7.5
■ 95% Criterion	13.5	16.8	16.0	13.2	17.5	22.1	26.6
Rate of Return	6.8	7.1	7.0	6.8	7.2	7.6	8.0
■ 90% Criterion	14.9	18.7	19.5	15.3	21.9	26.0	29.6
Rate of Return	6.9	7.3	7.3	6.9	7.5	8.1	8.0

Minimum Thresholds

The decision as to which criterion to use to determine the target duration of the dollar portfolio depends on the portfolio manager's set of risk/return preferences. The results according to the three criteria – minimum, 95%, and 90% – were calculated for each of the minimum thresholds. After examining the results, the minimum threshold and criterion most appropriate to the Bank of Israel's preferences were chosen (see Section 9 below).

Figure 1 above shows the time series of neutral durations derived from the minimum threshold. The horizontal lines represent the duration calculated on the basis of the different criteria.

Figure 2 shows the neutral durations and their holding period rates of return obtained using different minimum thresholds and different criteria. As the figure shows, the minimum target duration obtained is 12 months and the maximum is about 30 months. After examining additional data in the next sections, an attempt is made to determine the neutral duration for the Bank of Israel's dollar portfolio.

6 Retroactive examination of meeting the shortfall duration

Since the calculations for choosing the neutral duration are made ex ante on the basis of the distribution of annual holding period returns, which is estimated from past observations, a retroactive examination was undertaken to ascertain whether the actual (ex post) holding period return was above the minimum threshold in at least 95% of the cases throughout the sample period. This examination was implemented for each of the durations derived from the

Table 1: Percentage of cases in which the moving annual holding period returns were below the minimum threshold, June 1985-June 1998

90% Criterion		95% Criterion		Minimum Criterion		
w/out 1994	entire period	w/out 1994	entire period	w/out 1994	entire period	Type of Minimum Threshold
1.6%	2.8%	0%	0.7%	0%	0%	3-month Treasury bills less 2%
0%	0.7%	0%	0%	0%	0%	3-month Treasury bills less 3%
0.8%	6.9%	0%	4.1%	0%	2.1%	50% of 3-month Treasury bills
1.6%	8.3%	0.8%	6.2%	0%	3.4%	65% of 3-month Treasury bills
0.8%	8.3%	0%	7.6%	0%	5.5%	Fixed threshold of 2%
0.8%	7.6%	0%	5.5%	0%	4.8%	Fixed threshold of 1%
0%	5.5%	0%	4.8%	0%	2.1%	Fixed threshold of 0%

various minimum thresholds, and for each kind of criterion. The results may be seen in Figure 3[5] and in Table 1 below.

As Table 1 and Figure 3 show, the actual annual returns over the entire period were below the relevant thresholds in close to or even less than 5% of observations (i.e. they met the 95% confidence level) for two categories – the minimum and 95% criteria. The results for the 90% criterion were less successful.

Excluding the effects of 1994, the 95% confidence level held in all cases. The "minimum criterion" met the shortfall criterion for the entire period and for all thresholds. The thresholds defined as 3-month Treasury bills *less* 2 percentage points and Treasury bills *less* 3 percentage points, performed the best – the return on the portfolio was below the minimum threshold in less than 3% of cases.

Figure 3: Moving annual rates of return of T-notes with varying durations derived from threshold 3 and different criteria

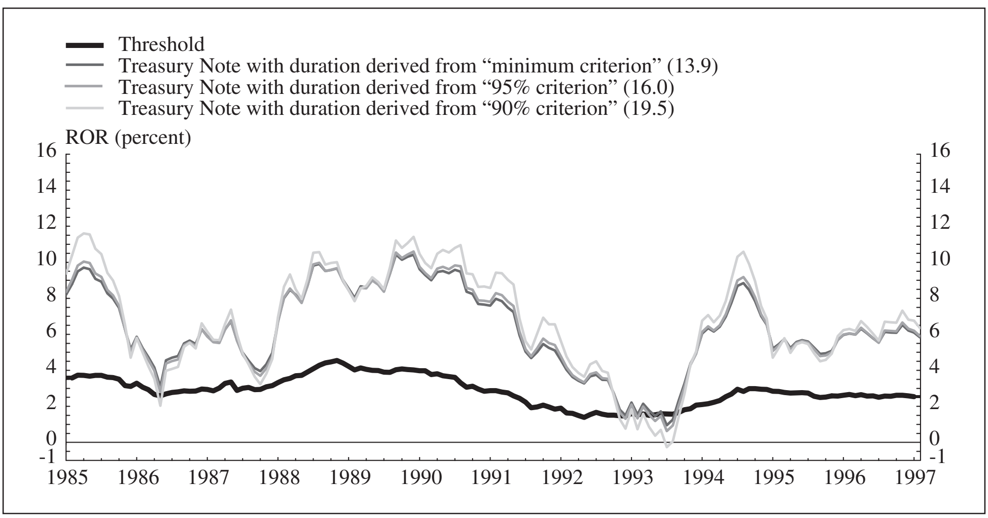

7 Comparison between actual holding period returns for portfolios of different durations

Another important point worth examining before deciding on a neutral duration for the dollar portfolio is the actual holding period return on Treasury notes (T-notes) of varying durations over periods of several years. Table 2 shows statistical calculations on the annual return differences between T-notes of varying durations and 12-month duration T-notes (recall that 12 months was the neutral duration of the dollar portfolio at the time of this study).

As Table 2 shows, the longer the duration of the T-note, the higher its average annual return across the sample period. Thus, for example, the average return on a 16-month T-note was 13 basis points higher, and the average return on a 21-month note was 49 basis points higher than that on a 12-month duration T-note, throughout the sample period. Similarly, the cumulative return over the sample period for a 16-month T-note was 2.7 percentage points higher and that on a 21-month T-note was 10.5 percentage points higher than that on a 12-month T-note, for the sample period.

Table 2: Moving annual return differences between treasury notes of varying durations and 12-month duration Treasury notes (basis points) (Period 1985-1997)

24m-12m	21m-12m	18m-12m	16m-12m	
75	49	41	13	Average
-255	-198	-133	-93	Minimum
402	285	202	105	Maximum
12.0%	10.5%	6.5%	2.7%	Cumulative Return

8 Examination of the model over the long term (out of sample)

The model presented here is based entirely on historical data over the sample period (June 1984-July 1998). It is sometimes customary, after the results of a model based on a specific period are received, to apply the same model to other, out of sample, periods in order to test its reliability. This makes it possible to check whether the results also meet all the criteria set for them in periods that are exogenous to the calculation of the model itself.

For this reason, the period from 1954 to 1998 was examined in order to test the validity of the results obtained previously for the long term, too. This test is important because, while the sample period is characterised by a long-term trend of declining interest rates, the period beginning in 1954 includes a long-term trend of rising interest rates as well as periods of sharp interest rate fluctuations. The yields to maturity over the period reviewed are shown in Figure 4.

Figure 4: Yields to maturity on on-the-run treasury notes
(Monthly Averages)

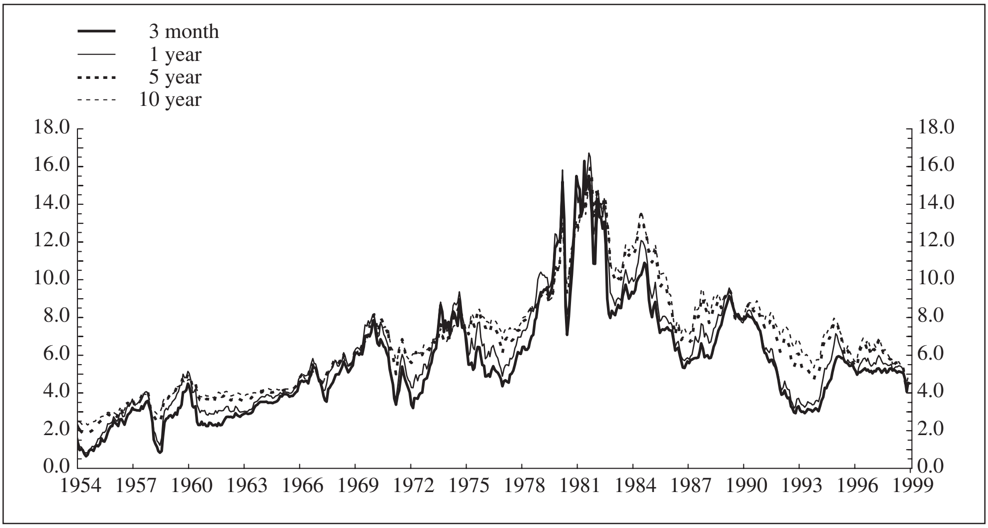

The test sought to examine whether the neutral duration derived from the model met the shortfall criterion in a period that was out of sample and differed from the sample period regarding market trends, the level of interest rates and their volatility. For this purpose, the moving annual holding period returns were calculated from 1954 (on the basis of monthly data) on T-notes of varying durations, and compared with the appropriate minimum threshold levels. The results can be seen in Figure 5[5] and in Table 3.

Table 3 and Figure 5 show that the results of this test do not differ significantly from those obtained in a similar test for the sample period, and that the percentage of cases in which the holding period returns fell below the minimum threshold was relatively low, with the exception of the 2% fixed threshold.

This test provided a good indication of what might happen in the future, as it incorporated a very long-term view of interest rate levels and varying trends over time. The results showed that out of sample, the test meets the shortfall criterion very well, providing additional confidence in its ability to meet this criterion in the future.

Table 3: Percentage of cases in which moving annual holding period returns fell below the minimum threshold, January 1954-December 1998

90% Criterion	95% Criterion	Minimum Criterion	Type of Minimum Threshold
2.8%	1.5%	0.8%	3-month Treasury bills *less* 2%
3.6%	2.5%	0.6%	3-month Treasury bills *less* 3%
6.2%	3.8%	2.1%	50% of 3-month Treasury bills
6.0%	3.4%	1.7%	65% of 3-month Treasury bills
11.2%	9.3%	8.1%	Fixed threshold of 2%
7.0%	5.5%	4.3%	Fixed threshold of 1%
4.9%	4.0%	2.3%	Fixed threshold of 0%

Figure 5: Moving annual rates of return of T-notes with varying durations derived from threshold 3 and different criteria (1954-1998)

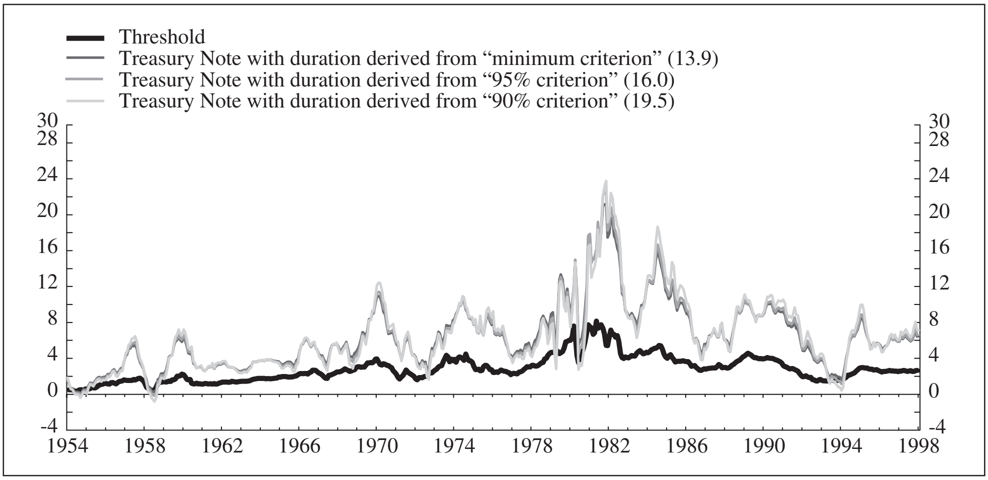

9 Choosing the minimum threshold level for the Bank of Israel' dollar portfolio

To choose the appropriate minimum threshold from among the seven possibilities, a process of elimination was used. First, the various minimum threshold levels were rejected since they did not suit the Bank's risk/return preferences. The fixed minimum threshold does not change with the market, and may even be described as "detached from reality". When the level of interest rates is high, for example, the minimum threshold remains low and will not express a preference for a minimum return on the portfolio. When interest rates are low, the minimum threshold may actually be higher than the return that could be realised under prevailing market conditions. Consequently, it would not seem desirable to prefer a fixed minimum threshold.

In deciding between a threshold with a yield set as a fixed spread below that on 3-month T-bills and a threshold with a yield set as a percentage of that on 3-month T-bills, the second option was preferred, on the grounds that it provides better protection against low interest rates than the former, and also offers the portfolio manager a greater degree of flexibility at times of high interest rates. In addition, the latter minimum threshold cannot be negative, which is sometimes the "red line" below which a portfolio manager is not prepared to go.

After choosing a threshold set as a percentage of the yield on a risk-free asset, it was necessary to decide between the two possibilities presented in this paper:

1. A minimum threshold of half the yield to maturity on 3-month T-bills; or
2. A minimum threshold of 65% of the yield to maturity on 3-month T-bills.

A scientific justification for choosing one option over the other cannot be given because, as stated, the choice depends on the risk/return preferences of the portfolio manager. It is possible, however, to examine the incremental return obtained by setting duration according to each of these two minimum thresholds.

Figure 6 shows the incremental return on an asset at a duration derived by using the first minimum threshold as compared with that derived from using the second, in accordance with the three criteria defined in Section 5. As the figure shows, the average annual return obtained

Figure 6: Incremental return on a duration derived using the threshold of 50% of the yield-to-maturity on 3-m T-bills as compared with that derived using the threshold of 65% of the yield-to-maturity on 3-m T-bills

(in accordance with the three criteria for choosing over time, basis point)

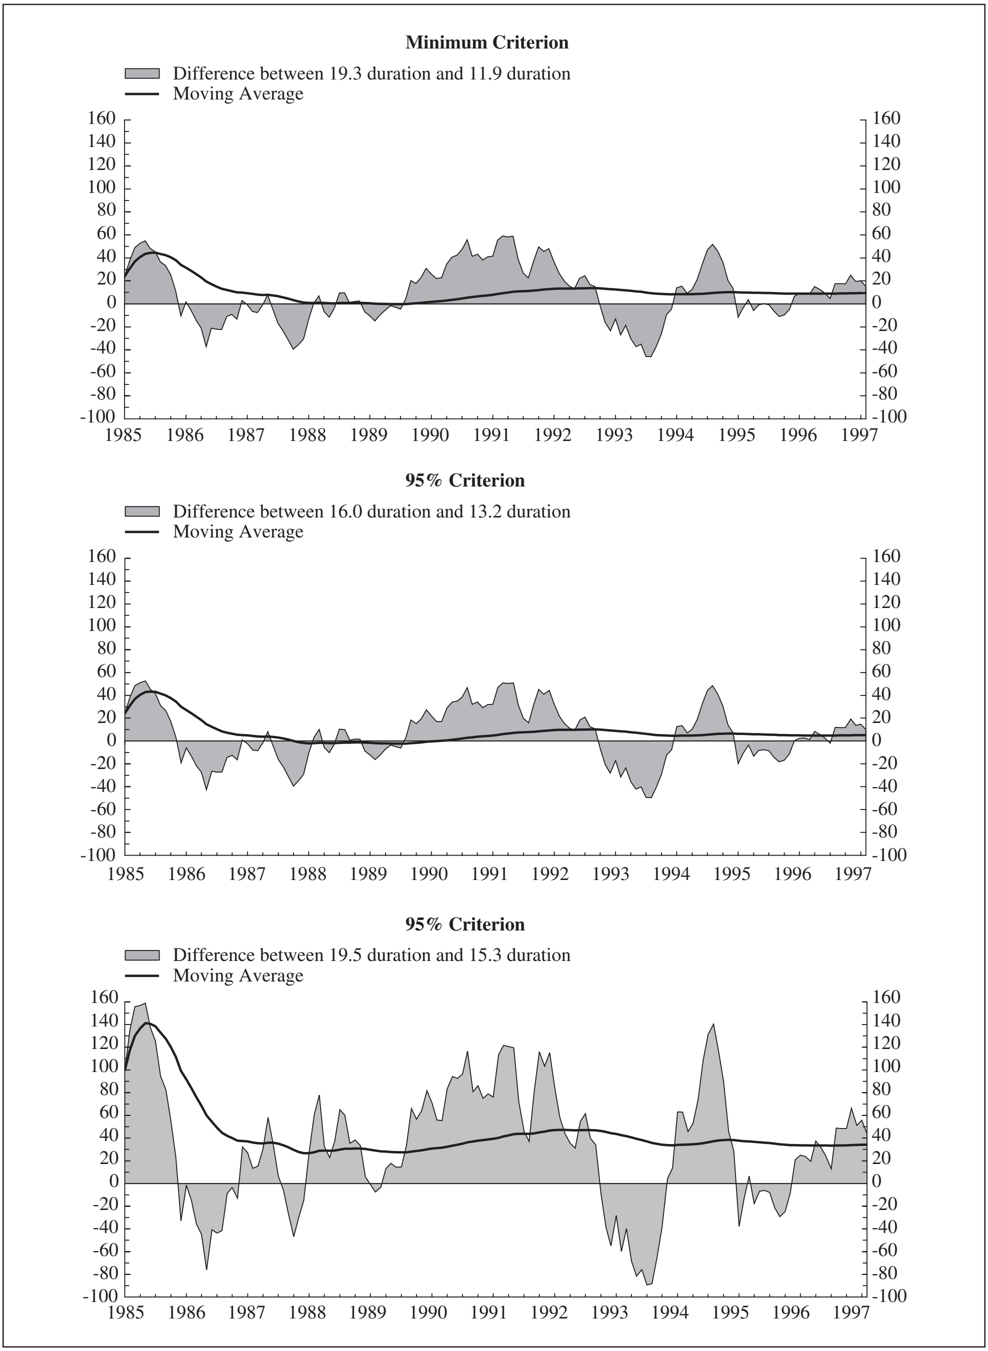

between June 1985 and June 1997 is higher when using the first rather than the second threshold for all the criteria. Therefore, if the preferences of the portfolio manager are not known and it is necessary to choose solely on the basis of returns, the first threshold, i.e., *half the yield to maturity on 3-month T-bills,* should be chosen.

At this stage, all that remains is to choose one of the three already identified criteria for "meeting the minimum threshold throughout the period".

Of these three criteria, the 95% one was chosen, as it was consistent with the confidence level of the shortfall criterion, where 95% was also chosen.

Thus, according to the model and the two decisions, the desired neutral duration for the dollar portfolio was 16 months.

On the basis of this study, a recommendation was made to the Foreign Currency Committee to increase the duration of the benchmark dollar portfolio from 12 to 16 months. This recommendation was accepted at the beginning of 1999.

10 Practical experience

The research presented in this paper was undertaken five years ago. On the basis of the results obtained, the benchmark duration of the dollar portfolio was changed from 12 to 16 months. Today it is worth examining whether a) the portfolio has withstood the shortfall criterion during the past five years, i.e. did the portfolio with a 16-month duration earn an annual return above the minimum threshold (half of the yield to maturity of the 3-month T-bill) in at least 95% of the observations; and b) if the study were undertaken today, including the last five years in the sample period, what neutral duration would the model have suggested?

These questions are even more relevant in light of recent exceptional market trends: sharp falls in interest rates during the last three years, especially at the short end of the yield curve, as a result of fast and intensive moves by the Fed in response to the bursting of the high-tech bubble and the economic downturn which followed, bringing rates down to their lowest levels in 50 years (Figure 7).

Figure 7: Yields to maturity on on-the-run Treasury notes
(in percentages)

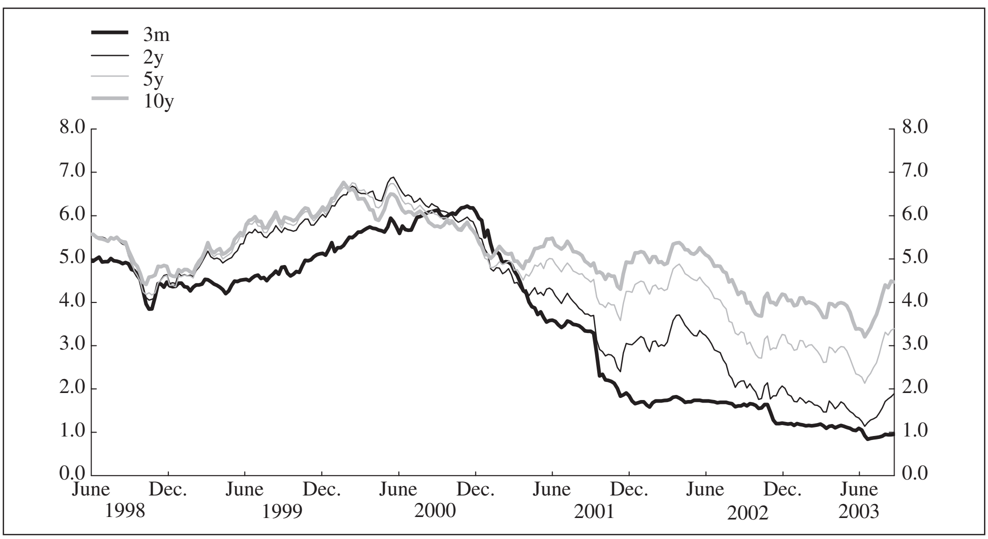

When rates are so low, a small upward movement in yields to maturity can result in capital losses. This, together with the low underlying yield, could cause the holding period return to fall below the minimum threshold, in which case it would not meet the shortfall criterion.

To answer the above questions, the following tests were conducted:

1. A historical test was run, based on data from the last five years (i.e. from July 1998 to August 2003), which examined whether the 16-month duration portfolio actually met the shortfall criterion.
2. The model was then run over a longer sample period, incorporating the last five years of data, to establish what the neutral duration of the portfolio should have been.

Methods and results

1. For the first test, a bullet portfolio was constructed with a 16-month duration by taking the two Treasury notes with duration closest to 16 months (one longer, one shorter) and weighting them accordingly. The daily holding period return was calculated for this portfolio. Then, for every observation the historical annual return was calculated and compared with the minimum threshold relevant to that observation. As Figure 8 shows, the annual return of the portfolio was greater than the threshold at every point in time during the entire period (July 1998-August 2003). Thus, the Bank of Israel's dollar benchmark, set at 16 months, met the shortfall criterion. Furthermore, the 16-month duration portfolio earned a return of 5.27%, in annual terms, for the period, while the Treasury portfolio with a 12-month duration earned only 5.06% – a difference of 21 basis points. In retrospect we can see that the move to a longer duration provided extra return while still maintaining the shortfall criterion.

Figure 8: Moving annual rates of return of T-note portfolio with 16-m. duration vs. threshold (June 1998-August 2003)

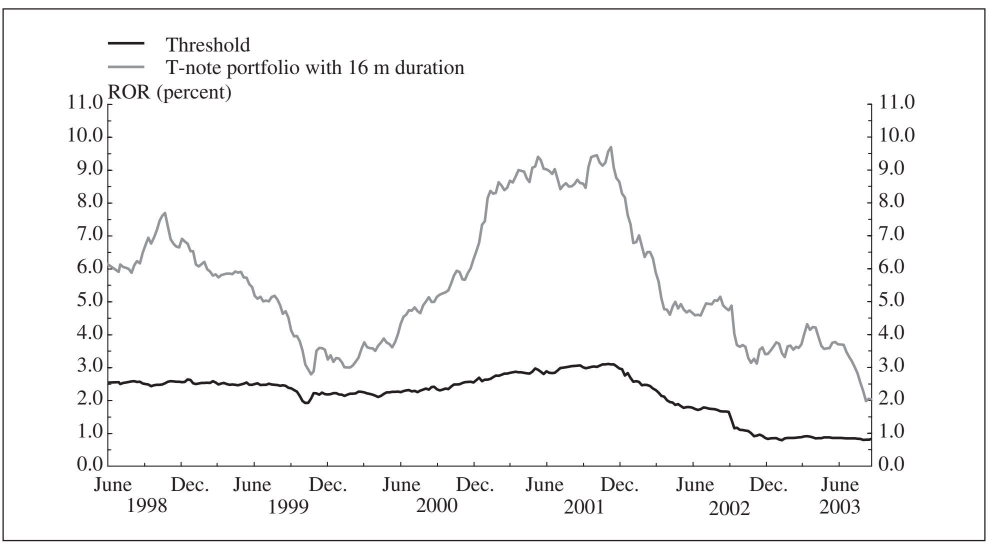

Figure 9: Neutral durations: Time series derived from 50% of 3-m. t-bills threshold
(as at June of each year)

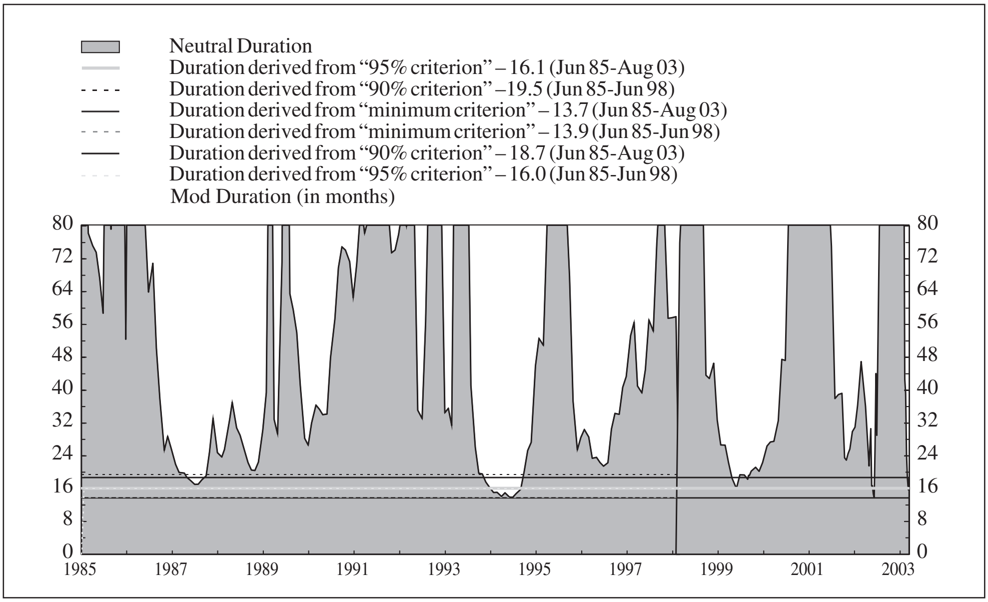

2. For the second test, the original sample period was extended to include the additional five years of data. The extended period, from June 1985 to August 2003, tests what the neutral duration of the portfolio should have been, leaving all other parameters of the model unchanged. Using the 95% confidence level, the model suggests a duration of 16.1 months, versus 16 months for the original sample period. In view of market trends in the last three years, the stability of the result was surprising, appearing to indicate that neutral duration, as derived by the model and in keeping with the criterion of a long-term target, is very stable over time and not greatly affected by cyclical changes in yields.[6] These results can be seen in Figure 9, which presents a time series of neutral durations derived by the model for each observation, as well as the neutral duration derived on the basis of different criteria: minimum, 95%, 90%, extended sample period and the period from June 1985-August 2003 (horizontal lines), and the original sample period (dashed horizontal lines).

In conclusion, it is necessary to explain the choice of the neutral portfolio, i.e. the "sanctity" of the neutral duration decision. Although the neutral duration derived by the model is based on the portfolio manager's set of risk/return preferences, if the manager believes that market conditions have changed so as to expose the portfolio to undue risks not accounted for by the model, he or she may alter the neutral duration. For precisely this reason, the Bank of Israel decided at the end of 2002 to temporarily shorten the neutral duration of the dollar portfolio from 16 to 11 months. This decision was taken against the background of a long-term decline in interest rates, bringing them to historical lows. It was estimated that the low level of interest

6 Note, however, that the last three years of the extended sample period were marked by falling yields, leading to capital gains (as opposed to capital losses). As the model is based on historical yields to maturity (one year back), the distribution of holding period returns which the model estimated was positively biased. Hence, if a period of rising yields is just beginning, running the model again to include this period might generate different results.

rates could cause the portfolio's return to drop below the minimum threshold as rates began rising and capital gains turned to losses. While running the model over the recent period did not suggest setting neutral duration below 16 months, it still did not negate the risk of earning a return below the minimum threshold – since the choice of a 16-month duration was based on the "95% criterion" for meeting the minimum threshold over time. According to this criterion, 95% of all the observations in the time series of durations derived by the model (in the sample period) were above 16 months duration, while 5% were below. Therefore, in addition to the 5% probability which the model allows for not meeting the minimum threshold (i.e. the 95% confidence level), there is another risk whereby the 16-month duration portfolio would not meet the minimum threshold. The decision to shorten temporarily the neutral duration of the dollar portfolio was based on the assessment that in the near term the probability that this would occur had increased significantly.

11 Conclusion

In this study a method of determining the target duration of the Bank of Israel's dollar portfolio using the shortfall approach was presented. According to this method, the target duration was determined in such a way that at a given probability, the holding period return on the dollar portfolio would not fall below a minimum threshold. The portfolio manager was required to set three parameters for deciding the target duration of the dollar portfolio. These parameters, which reflect the portfolio manager's preferences and risk aversion, are as follows:

1. The most appropriate minimum threshold level. In this paper seven different possibilities were presented, from which the minimum threshold defined as half the return on a risk-free asset was chosen.
2. The confidence level at which the portfolio yields a higher return than the minimum threshold level. In this paper, the confidence level was set at 95%.
3. The criterion at which the minimum threshold is met over the period as a whole. In this paper three possibilities were presented, and the "95% criterion".

After choosing the parameters and running the model accordingly, our results showed that according to the preferences of the portfolio manager, the duration of the Bank of Israel's dollar portfolio should be extended from 12 to 16 months. This recommendation was accepted by the Foreign Currency Committee.

On the basis of the research presented in this paper and the results obtained, which was carried out five years ago, the benchmark duration of the dollar portfolio was changed from 12 to 16 months. It was interesting to examine ex post a) whether the portfolio had withstood the shortfall criterion during the past five years, and b) if the study were undertaken today, including the last five years in the sample period, what neutral duration would the model have suggested? Closer examination of these two questions shows firstly that the annual return of the portfolio was greater than the threshold at every point in time during the period July 1998-August 2003. In other words, the Bank of Israel's dollar benchmark, set at 16 months, met the shortfall criterion. Secondly, when the sample period was extended to include an additional five years of data, using the 95% confidence level, the model suggests a duration of 16.1 months - versus 16 months for the original sample period. In view of market trends in the last three years, the stability of the result was surprising, and appears to indicate that neutral duration, as derived by the model and in keeping with the criterion of a long-term target, is very stable over time and not greatly affected by cyclical changes in yields.

Finally, yields to maturity on individual Treasury securities were used to estimate the distribution of holding period returns on assets of varying durations. In reality, however, asset portfolios usually consist of a wide range of securities which have different durations distributed throughout the portfolio's benchmark duration. These distributions affect the portfolio's holding period return, especially when the shape of the yield curve changes in a non-parallel manner. Thus, if the model had been run with different portfolios comprising a large number of securities, rather than a single security portfolio, the results obtained would have differed slightly from those presented here. The issue of choosing a neutral asset distribution in addition to a neutral duration is beyond the scope of this study, however.

References

Butler, Kirt C. and Dale L. Domain; "Risk, Diversification and the Investment Horizon"; *Journal of Portfolio Management,* Spring 1991

Langetieg, Terence C., Martin L. Leibowitz, and Stanley Kogelman; "Duration Targeting and the Management of Multiperiod Returns"; *Financial Analysts Journal, September/October 1990*

Leibowitz, Martin L. and Roy D. Henriksson; "Portfolio Optimization with Shortfall Constraints: A Confidence-limit Approach to Managing Downside Risk"; *Financial Analysts Journal,* March/April 1989

Leibowitz, Martin L., Stanley Kogelman and Thomas E. Klaffky; "A Shortfall Approach to Duration Management"; Salomon Brothers Inc., April 1990

Leibowitz, Martin L. and Stanley Kogelman; "Assets Allocation under Shortfall Constraints"; Salomon Brothers, *Journal of Portfolio Management,* Winter 1991

Leibowitz, Martin L. and Terence C. Langetieg; "Shortfall Risk and the Assets Allocation Decision: A Simulation Analysis of Stock and Bond Risk Profiles"; Salomon Brothers, *Journal of Portfolio Management,* Fall 1989

Piterman, Silvia; "Management of Interest Rate Risk by Central Banks"; Salomon Brothers 1994, Central Bank Seminar

Thomas, R. L. (1997); "Modern Econometrics"

Appendix

Determining the range of historical observations according to which the distribution of annual holding period returns are measured

The objective here is to examine the range of historical observations to be used as a basis for estimating the distribution of the annual holding period returns, in such a way that the distribution obtained reflects reality as closely as possible.

For the purpose of this examination, the distribution of the annual holding period returns was calculated simultaneously for different ranges of historical observations (as described in Section 2).

- Range of one year retroactively
- Range of two years retroactively
- Range of three years retroactively
- Range of 74 retroactive observations, where the most recent are given a higher weight than those that are furthest away.

As stated, the distribution of annual holding period returns was calculated for each of these ranges and compared with those actually obtained ex post (i.e. a comparison of the forecast with reality). This comparison was based on a sample period of 13 years, at monthly intervals.

The mean and standard deviation were calculated from the distribution for each point in time, and a range defining 95% and 99% of the distribution was defined accordingly. These ranges were then compared ex post with the annual holding period returns actually obtained. The *fewer* the instances throughout the period in which the actual holding period return was not in the range defined by the distribution, the better the distribution. *The figures below show quite clearly that the distribution calculated on the basis of* ***one year*** *retroactively provides the best reflection of reality.* It can also be seen that the period between 1985 and 1988 had the worst predictive ability out of all the alternatives presented. During this period autocorrelation appears to have been low, making it more difficult to estimate future returns on the basis of the past alone. Discounting this period, as it is relatively distant and since the market conditions that prevailed then do not resemble those present at the time of this study, it would seem that the number of times in which the actual holding period returns are not within the distribution range, when choosing a forecast for one year retroactively, is very low.

Table 4: Comparison of actual holding period returns with distributions based on different ranges of observations; percentage of cases in which actual yields were outside the range of the distribution with a 99 % confidence level

June 85-July 97	3m duration	12m duration	21m duration
Range of one year	1%	25%	16%
Range of two years	14%	30%	21%
Range of three years	24%	29%	23%
74 weighted observations	19%	29%	27%
June 89-July 97			
Range of one year	0%	7%	2%
Range of two years	12%	23%	18%
Range of three years	17%	16%	15%
74 weighted observations	11%	24%	26%

Actual annual RoR of 12m. tnotes vs. expected – Based on 3 years history

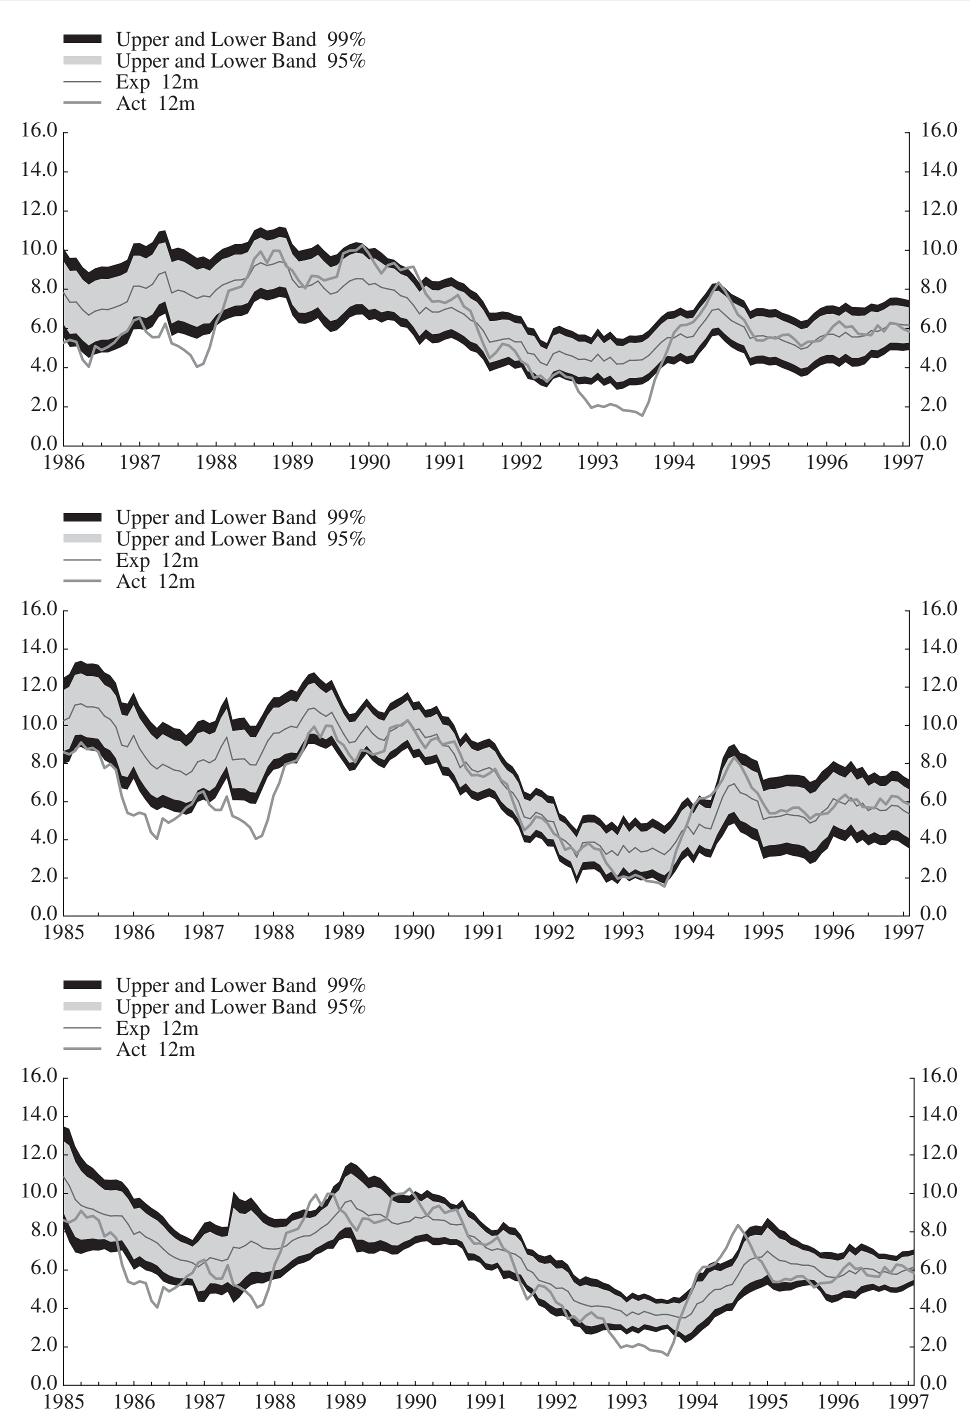

List of contributors

Janet Assouline is Senior Economist in the Foreign Department at the Bank of Israel. She has a degree in Economics and a Masters in Statistics and Operations Research from the Hebrew University. She has been at the Bank of Israel for the past 20 years.
E-mail address: janet_as@bankisrael.gov.il

Antonio Francisco de Almeida da Silva Junior is an Advisor in the Risk Analysis Team, International Reserves Operations Department of the Central Bank of Brazil. He has a degree in Engineering, an MSc in Process Control and another in Finance from the Universidade Federal da Bahia. His doctorate studies are in the field of exchange rate optimal stochastic intervention control at the Instituto Tecnológico de Aeronáutica.
E-mail: antonio.francisco@bcb.gov.br

Karel Bauer is Deputy Executive Director of the Risk Management and Transaction Support Department at the Czech National Bank, which he joined in 1992. Before joining the risk management area, he worked in the Front Office as a Portfolio Manager of USD part foreign exchange reserves. He graduated from the Faculty of Nuclear Science in Physical Engineering in 1989.
E-mail: karel.bauer@cnb.cz

Carlos Bernadell has been the Head of the Risk Management Division at the European Central Bank since 1998. The Division is responsible for the management of market risk in the ECB's investment operations and credit risk in the investment, monetary policy and foreign exchange operations of the ECB. Before joining the ECB, Carlos Bernadell was Director of Market Analysis and Control at La Caixa de Ahorros y Pensiones de Barcelona and Senior Team Leader in the Auditing/Corporate Finance Department at Coopers & Lybrand. He holds a degree and postgraduate diploma in Economics and Finance from the University of Barcelona.
E-mail: carlos.bernadell@ecb.int

Bert Boertje has been Deputy Departmental Director in the Financial Markets Department at De Nederlandsche Bank for the last 12 years. He started out as a Fixed Income Portfolio Manager before heading up the Equity Investment Section. For the last three years he has been managing the Risk Management Unit.
E-mail: b.boertje@dnb.nl

Pierre Cardon, CFA is responsible for third-party asset management at the Bank for International Settlements as well as acting as Secretary to the BIS Pension Fund Committee. He also advises central banks and official institutions on reserve management issues and is an Advisory Committee member of the European Pension Fund Investment Forum. He has over 20 years' experience in asset management and risk management. He obtained a Masters in Economics from the University of Louvain-la-Neuve and a MBA from the University of Chicago.
E-mail: pierre.cardon@bis.org

Stijn Claessens is Professor of International Finance at the University of Amsterdam. He holds a Ph.D. in Business Economics from the Wharton School of the University of Pennsylvania. From 1987 to 2001, Stijn Claessens worked at the World Bank, most recently

as Lead Economist in the Financial Sector Strategy and Policy Group. His policy and research interests are in external finance and in the domestic financial sector issues. He has published extensively in these areas and is the editor of several books.
E-mail: stijn.claessens@uva.nl

Joachim Coche is Senior Economist at the European Central Bank in Frankfurt. He has a Masters and Doctorate in Economics from the University of Osnabrück. Within the Risk Management Division he is responsible for developing asset allocation strategies for the ECB's foreign reserves and domestic assets. Before joining the ECB in 2000, he worked for DZ Bank, Frankfurt, and advised institutional clients on their investment process.
E-mail: joachim.coche@ecb.int

David Delgado Ruiz is Financial Analyst in the Risk Control Division at the Banco Central de Venezuela. His main responsibilities are performance and liquidity risk measurement of the Bank's international reserve portfolios. Before joining the Risk Control Division, he worked in the Economics Department, building econometric models in order to forecast economic variables. He holds a degree in Economics from Universidad Central de Venezuela and is currently studying for a Masters in Operational Research at Universidad Central de Venezuela.
E-mail: dadelad@bcv.org.ve

Francis X. Diebold is W.P. Carey Professor of Economics and Professor of Finance and Statistics at the University of Pennsylvania and its Wharton School, and Faculty Research Associate at the National Bureau of Economic Research in Cambridge, Massachusetts. He is a leader in financial modelling, risk management and forecasting, and publishes widely and serves on numerous editorial and advisory boards. Diebold received his BS from the Wharton School in 1981 and his Ph.D. in 1986. He also worked as an economist under Paul Volker and Alan Greenspan at the Board of Governors of the Federal Reserve System in Washington DC, before returning to the University of Pennsylvania in 1989.
E-mail: fdiebold@wharton.upenn.edu

Mark Dwyer is a Senior Risk Consultant at DST International, a UK-based provider of asset and risk management systems for institutional investors.
E-mail: mark.dwyer@dstintl.com

Lev Dynkin is Managing Director at Lehman Brothers in New York. He is Head of the Quantitative Portfolio Strategies Group in Fixed Income Research. The Group focuses on issues of portfolio optimisation, investment style and benchmark customisation for global institutional investors.
E-mail address: ldynkin@lehman.com

Paul Embrechts is Professor of Mathematics at the ETH Zurich and Centennial Professor of Finance at the London School of Economics. His main research areas are the modelling of extrernal events in insurance and finance, as well as quantitative risk management.
E-mail: embrechts@math.ethz.ch

Peter Ferket is Head of the Quantitative Research Department of Robeco Asset Management in the Netherlands. He joined the IRIS research joint venture of Robeco and Rabobank in

1997 and participated in many risk management projects for Robeco and Rabobank International. Since 2000 he has been responsible for co-ordinating all quantitative research efforts of the European business units of the Robeco Group. Peter Ferket holds a Ph.D. in Scientific Computing from the Eindhoven University of Technology (1996).
E-mail: p.j.j.ferket@robeco.nl

Stephen J. Fisher, Managing Director, is a Client Portfolio Manager in the International Fixed Income Group based in Singapore. Initially an employee since 1992, Stephen Fisher rejoined in September 2000 after a period working with Merrill Lynch Investment Managers. Prior to this, he was Regional Head of Capital Markets Research for Asia-Pacific from 1996 to 1998. Before this, he worked in the Capital Markets Research group in New York. Stephen Fisher's expertise is in fixed income, equities, asset allocation and derivatives. He obtained an MS in Finance from the University of Rochester and a B.Ec. from the University of Sydney, Australia. He also earned a Ph.D. in Finance from the W.E. Simon Graduate School of Business Administration, University of Rochester.
E-mail: stephen.j.fisher@jpmorganfleming.com

Andreas Gintschel Ph.D., Assistant Vice-President Investment Solutions and Overlay Management Group at DeAM, Frankfurt. Andreas Gintschel joined DeAM in 2002 as a quantitative analyst. Prior to that he was an Assistant Professor of Finance at Goizueta Business School, Emory University, from 2001 to 2002. He has a Ph.D. in Finance from Atlanta and an MS in Applied Economics from the University of Rochester, as well as a Master's degree ("Diplom Kaufmann") in Business Administration from the University of Trier, where he has also been a Lecturer in Capital Markets.
E-mail: andreas.gintschel@db.com

Roberts L. Grava is a Member of the Executive Board and Head of the Market Operations Department at the Bank of Latvia. Since 1994, he has been running the Department currently responsible for reserves management, portfolio risk management, investment and risk analysis, foreign currency interventions and domestic monetary policy operations, national debt management (together with the Finance Ministry and State Treasury), and payments and settlements.
E-mail: roberts.grava@bank.lv

Clement Ho, CFA joined the Hong Kong Monetary Authority in 1995. From 1995 to 2001, he was responsible for the management of internally managed bond portfolios, and from September 2001 for the Risk Management & Compliance Division. Before joining the Hong Kong Monetary Authority, he worked at Chase, Fuji International Finance and HSBC Markets. He graduated from the Chinese University of Hong Kong in 1984.
E-mail: clement_kc_ho@hkma.gov.hk

Han van der Hoorn is a Senior Risk Manager in the Financial Markets Department at De Nederlandsche Bank. He holds degrees in Econometrics (Erasmus University) and Economics (University of Amsterdam). He started his career at Coopers & Lybrand and joined De Nederlandsche Bank in 1996, where he first worked as an Economist in the Supervision Department and as a Fixed Income Portfolio Manager.
E-mail: h.h.j.van.der.hoorn@dnb.nl

Jay Hyman is a Senior Vice-President at Lehman Brothers, where he has worked for ten years in the Quantitative Portfolio Strategies Group. Current areas of interest include risk budgeting, performance attribution and optimal portfolio structuring. He holds a Ph.D. in Electrical Engineering from Columbia University.
E-mail: jhyman@lehman.com

Roger Kaufmann was awarded his Degree in Mathematics by the ETH Zurich in 1999. From 1999 until 2003, whilst working on his Ph.D. thesis entitled “Long-Term Risk Management”, he was a researcher at the Financial and Insurance Mathematics Institute RiskLab. He will receive his Ph.D. in Mathematics from the ETH Zurich in summer 2004.
E-mail: kaufmann@math.ethz.ch

Michal Koblas is a Financial Risk Analyst in the Risk Management and Trade Support Department of the Czech National Bank, which he joined in 2002. He has a Masters in Applied Mathematics from Prague Technical University and undertook a Ph.D. level course in Economics at CERGE-El, Prague.
E-mail: michal.koblas@cnb.cz

Jerome L. Kreuser, Executive Director and Founder. In his current role at The RisKontrol Group, he continues to develop and apply advanced techniques for asset/liability risk management for sovereigns, (re)insurance companies, pension funds, etc. He has also served as an “International Reserves Management Advisor” for the IMF. Prior to that, he worked at the World Bank for 24 years and was an Adjunct Professor at George Washington University, USA. He has held several international workshops on asset/liability modelling and applications of optimisation in economics and finance. He received his doctorate from the University of Wisconsin in Mathematical Programming.
E-mail: kreuser@riskontroller.com

Min Ce Lie, Vice-President, has been a quantitative research analyst in the International Fixed Income Group at JP Morgan Fleming since1999. Prior to this, he was a member of the Strategic Investment Advisory Group, where he served as an investment adviser to pension plans, foundations, endowments and central banks. He also worked on issues related to asset allocation, asset-liability modelling, risk management and portfolio construction. Min Ce Lie obtained a BS in Mathematics and Economics from Duke University and a MS in Computational Finance from Carnegie Mellon University.
E-mail: min.lie@jpmorganfleming.com

Simone Manganelli is a Senior Economist in the General Economic Research Division of the European Central Bank, where he began work in September 2000. His main research interests lie in the areas of financial econometrics, risk management and asset allocation. Simone Manganelli received his B.Sc. from Bocconi University in 1994 and his Ph.D. from the University of California, San Diego in 2000, under the supervision of Robert Engle and Clive Granger.
E-mail: simone.manganelli@ecb.int

Pedro José Martínez Somoza is a Financial Analyst at the Banco Central de Venezuela**.** His main responsibilities are portfolio performance measurement of the BCV’s international reserves and the Investment Fund for Macroeconomic Stabilization (FIEM). Prior to this, he

was an Economist involved in the supervision and evaluation of the quality of survey data, the periodical elaboration of quality control reports, the elaboration of definitive information tables with the outputs of the survey and the validation of all the data collected nationally for adjusting national accounts. He has a degree in Economics from the Universidad Central de Venezuela.
E-mail: pmartine@bcv.org.ve

Ladislav Mochan is a Financial Risk Analyst in the Risk Management Division of the Czech National Bank, which he joined in 1999. He graduated from the Faculty of Nuclear Science and Physical Engineering, Czech Technical University, in 1999 and from the University of Economics in Prague in 2000.
E-mail: ladislav.mochan@cnb.cz

John Nugée is Head of the Official Institutions Group at State Street Global Advisors. From 1977 to 2000 he worked for the Bank of England and the Hong Kong Monetary Authority, predominantly in the markets and reserves management areas. From 1992-96 he was Executive Director of Reserves Management at the HKMA and from 1996-2000 Chief Manager of the reserves at the Bank of England. From 1997-2000 he was also a member of the Board of Directors of the European Investment Bank. He is an experienced lecturer on markets, investment and central banking reserves management themes.
E-mail: john_nugee@ssga.com

Eneira Osorio Yánez is Head of the Risk Control Division at the Banco Central de Venezuela, reporting to the Head of the International Investment Department. The Division co-ordinates activities relating to the BCV's portfolio performance measuring, risk management and the development of new quantitative methodologies in order to improve and keep these processes updated. Before joining the Investment Department, she was a Senior Statistician in the Research Analysis Division, providing statistical support to different projects at the department level, more specifically in aspects related to complex sampling techniques. She has a degree in Statistics from the Universidad Central de Venezuela and a Masters in Applied Economics from the University of Michigan at Ann Arbor. She was also an Assistant Professor at the Statistic and Actuarial Science School between 1987-1991.
E-mail: eosorio@bcv.org.ve

Reinaldo Alex Pabón Chwoschtschinsky is a Financial Specialist in the Risk Control Division at the Banco Central de Venezuela. His main responsibilities are measuring the market risk of the Banco Central de Venezuela's (BCV) international reserves portfolio and the Macroeconomic Stabilization Fund (FIEM). He holds a degree in Economics from the Université Paris – Dauphine (France) as well as a Masters in Economics (Finance option) from the same university, and is currently working on his Ph.D. in Economics at the Universidad Católica Andrés Bello.
E-mail: rpabon@bcv.org.ve

Eugen Puschkarski studied International Business Administration at the University of Vienna and in New York. He focused his education on modern portfolio theory, risk management and term structure modelling. Upon graduating in 1998 he joined the CCEFM doctoral programme, which he completed in 2001. He is currently finishing his Ph.D. in "Applied Prescriptive Investment Decision Theory". He is also working in the Treasury

Department of the Oesterreichische Nationalbank (OeNB) as Market Risk Manager and Financial Products Specialist.
E-mail: eugen.puschkarski@oenb.at

Bluford H. Putnam is President of Bayesian Edge Technology & Solutions Ltd., a firm specialising in financial risk management research, integrating practical market experience and economic theory with state-of-the-art dynamic quantitative methods. He has worked in the financial markets in New York and London since 1976 and has published numerous books and articles on international finance.
E-mail: blu@bayes1.com

Srichander Ramaswamy is Head of Investment Analysis at the Bank for International Settlements and Adjunct Professor of Banking and Finance at the University of Lausanne. He holds degrees in Electrical Engineering and Aerospace Engineering.
E-mail: srichander.ramaswamy@bis.org

Eli M. Remolona runs the Financial Markets area of the Bank for International Settlements (BIS). His group monitors the markets, undertakes research and publishes the *BIS Quarterly Review*. Before joining the BIS in 1999, he was a Research Officer at the Federal Reserve Bank of New York. He has a Ph.D. in Economics from Stanford University.
E-mail: eli.remolona@bis.org

Ciarán Rogers joined the ECB in 1998 in the Risk Management Division. He helped develop the risk management framework at the ECB. From 2000 onwards, he led the team responsible for the risk management of investment operations. In March 2004, he became Head of the Liquidity Management Section in the Directorate Operations. He holds degrees from University College Dublin.
E-mail: ciaran.rogers@ecb.int

Gennady Samorodnitsky is a Professor and Director of Graduate Studies at the School of Operations Research and Industrial Engineering, Cornell University. He has a Ph.D. in Statistics from Technion (1986). He is the author of close to 100 published research papers and a monograph “Stable non-Gaussian Processes” (with M. Taqqu). He is the Associate Editor of a number of journals in Applied Probability.
E-mail: gennady@orie.cornell.edu

Bernd Scherer, Managing Director, heads the Investment Solutions Group at Deutsche Asset Management in Frankfurt. He has written several books on quantitative methods in investment management and lectures in finance at the University of Augsburg.
E-mail: bernd.scherer@db.com

Jan Schmidt is Executive Director of the Risk Management and Transaction Support Department at the Czech National Bank, which he joined in 1991. Before working in risk management, he worked in the Front Office as a Portfolio Manager of USD part foreign exchange reserves. He graduated from the Faculty of Nuclear Science in Physical Engineering in 1989.
E-mail: jan.schmidt@cnb.cz

Martijn A. Schrijvers is an Economist and Portfolio Manager at the Financial Markets Department of De Nederlandsche Bank. He manages the euro-fixed income portfolio and analyses developments in the financial markets. In his previous position he was involved in advising on and the implementation of monetary policy. In 2003 he was seconded to the Financial Markets area of the Bank for International Settlements for six months. He has Masters degrees in General Economics and Development Economics from the Erasmus University of Rotterdam.
E-mail: m.a.schrijvers@dnb.nl

Alojz Simicak has been a Risk Manager at the National Bank of Slovakia for three years. He is responsible for developing and maintaining the asset allocation system.
E-mail: alojz.simicak@nbs.sk

Michael Zajac is Head of the Risk Management Unit at the National Bank of Slovakia. The department is responsible for asset allocation, risk management and performance measurement. Prior to this, he worked for ten years in foreign debt management.
E-mail: michal.zajac@nbs.sk

Machiel Zwanenburg is a Risk Manager for the Robeco Asset Management Unit of Robeco in the Netherlands. Until July 2003 he worked in the Quantitative Research Department of the same firm, with a focus on risk management issues. He graduated with a Master of Arts in Econometrics from the Erasmus University Rotterdam in 1998 and a Master of Science in Economics from the London School of Economics and Political Science in 1999.
E-mail: m.zwanenburg@robeco.nl