EUROPEAN CENTRAL BANK

EUROSYSTEM

ECB LAMFALUSSY FELLOWSHIP PROGRAMME

WORKING PAPER SERIES

NO 1197 / MAY 2010

INVESTORS WITH TOO MANY OPTIONS?

by Daniel Dorn

EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 1197 / MAY 2010

ECB LAMFALUSSY FELLOWSHIP PROGRAMME

INVESTORS WITH TOO MANY OPTIONS? [1]

by Daniel Dorn [2]

In 2010 all ECB publications feature a motif taken from the €500 banknote.

1 I thank James Choi and an anonymous ECB referee for comments. I also thank Charles Bergquist for his Perl programming help and Torsten LÄudecke for answering questions about the data from the Karlsruher Kapitalmarktdatenbank (KKMDB). This paper has been prepared by the author under the Lamfalussy Fellowship Program sponsored by the European Central Bank.

2 LeBow College of Business, Drexel University; 208 Academic Building; 101 North 33rd Street; Philadelphia, PA 19104; Email: dd79@drexel.edu

Lamfalussy Fellowships

This paper has been produced under the ECB Lamfalussy Fellowship programme. This programme was launched in 2003 in the context of the ECB-CFS Research Network on "Capital Markets and Financial Integration in Europe". It aims at stimulating high-quality research on the structure, integration and performance of the European financial system.

The Fellowship programme is named after Baron Alexandre Lamfalussy, the first President of the European Monetary Institute. Mr Lamfalussy is one of the leading central bankers of his time and one of the main supporters of a single capital market within the European Union.

Each year the programme sponsors five young scholars conducting a research project in the priority areas of the Network. The Lamfalussy Fellows and their projects are chosen by a selection committee composed of Eurosystem experts and academic scholars. Further information about the Network can be found at http://www.eu-financial-system.org and about the Fellowship programme under the menu point "fellowships".

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.europa.eu

Fax
+49 69 1344 6000

Information on all of the papers published in the ECB Working Paper Series can be found on the ECB's website, http://www.ecb.europa.eu/pub/scientific/wps/date/html/index.en.html

ISSN 1725-2806 (online)

CONTENTS

Abstract

During the last decade, markets for covered warrants (bank-issued options) have flourished in Europe and Asia. In these markets, investors often face a choice between many instruments that differ only slightly from each other. Based on retail trades in call options on the German DAX index, this paper documents substantial price dispersion across securities that are close substitutes. Moreover, investors generally fail to identify attractively priced options. The results suggest that the observed product proliferation imposes a substantial search cost on investors even though the products are homogenous and their pricing is well understood. The search cost is estimated to average 1% of the amount invested, the same order of magnitude as the average spread.

JEL codes: G11, G13, D83

Keywords: OTC derivatives, price dispersion, investor behavior, search costs

Non-Technical Summary

During the last decade, markets for covered warrants (also known as bank-issued options) have flourished in Europe and Asia. In recent years, trading volume in these over-the-counter derivatives has rivaled that in standardized contracts traded on traditional derivatives exchanges.

The over-the-counter market for bank-issued options exhibits two striking features documented here. First, the plethora of instruments offered: investors can typically choose between several bank-issued options that differ only slightly from each other. Second, prices differ substantially across instruments that are close substitutes. Price dispersion appears to be high both in absolute terms and relative to the price dispersion in comparable instruments traded on traditional derivatives exchanges.

These features prompt two questions. First, given that investors face a choice between several slightly different instruments, how do investors select among the available alternatives? Second, how do the options chosen perform relative to other options with similar characteristics that could have been bought, especially to those judged ex ante superior based on objective criteria such as spreads? Do trades made by more experienced or larger investors – who presumably have lower search costs or stronger incentives to search – perform relatively better?

To answer these questions, the study combines data on retail trades in DAX call options made at a large German discount broker and quotes for the population of DAX calls listed on the European Warrant Exchange (Euwax) during the period November 1999 to May 2000. The Euwax is one of the largest trading venues for bank-issued options worldwide.

The options actually chosen by the investors significantly underperform other options with a similar leverage, that is, a similar price elasticity with respect to changes in the DAX. The difference between the roundtrip returns of options bought and the corresponding returns of similar options not bought averages -0.7%; the average duration of a completed roundtrip is less than three days. The median return is also negative, suggesting that the sample investors would have been better off picking options randomly. Actual choices fare even worse when compared with similar options that are predicted to outperform based on information available to investors at the time of purchase, for example, because they have had lower spreads in the past. Options bought underperform similar options with the lowest spreads by about 1% per roundtrip, on average. To put this number into perspective: the average roundtrip return during this sample period in which the DAX appreciated by 22% was 1.6%. Consistent with a costly search conjecture, investors who submit small trades and novice option investors perform particularly poorly. A simple strategy of buying the option with the lowest spread

(and a similar leverage as the option actually bought) would have increased the absolute profit of inexperienced option investors by EUR 100 per trade, on average.

The standard explanation for the observed plethora of instruments is that banks compete to respond to the product preferences of rational investors. An alternative explanation is that low issuing costs and at most boundedly rational investors sustain an equilibrium with product proliferation and prices that are dispersed and above marginal cost. The paper documents substantial price dispersion and poor investor choices which are broadly consistent with the alternative view.

The study's results do not rule out that some investors benefit from the choice available in covered warrant markets. It is possible, though, that retail investors would be better off in an institutional setting found, for example, in Hong Kong, where higher listing and registration fees for covered warrants apparently lead to more concentrated trading in fewer listings. Analyzing the welfare implications of a listing fee increase is beyond the scope of the study, but perhaps an area for future research.

I Introduction

Conventional economic wisdom has consumers obtaining better outcomes when firms compete to offer them more choices. Moreover, rational consumers are thought to effortlessly use all publicly available information in their decision making.

This notion is suspect from at least two angles. First, a more complex choice set may be associated with higher search or psychological decision costs. One instance in the financial domain where more choice apparently fails to lead to superior decisions is reported by Huberman et al. (2004): participation in 401(k) retirement plans is negatively related to the number of investment options available, controlling for a host of investor and plan characteristics. Second, firms may strategically add complexity to the choice set to make it harder for consumers to find low-price offerings as pointed out by Carlin (2009). As a result, violations of the law of one price can occur – that is, prices can be dispersed and above marginal cost even when products are homogenous.

This paper examines choices of retail investors in the fast-growing market for covered warrants (also known as bank-issued derivatives). Over the past decade, these markets have grown rapidly in European and Asian countries, in part because they give retail investors easy access to option-like payoffs. The trading activity in covered warrants is remarkable, both in absolute terms and relative to traditional derivatives exchanges. Germany provides a good illustration because it is host to Eurex, one of the world's largest traditional derivatives exchanges, as well as to the European Warrant Exchange (Euwax), one of the world's largest exchanges for covered warrants. For 2007, Euwax reported aggregate premium trading volume of EUR 128 billion, or roughly 30% of Eurex volume. This is likely a gross underestimate of the overall trading volume in bank-issued derivatives because much of the trading in these instruments occurs on the proprietary platforms of the issuers (over 80% in the sample studied in this paper).

What makes the market for covered warrants a compelling case study of choice is not so much its size in terms of EUR trading volume, however, but the sheer number of different contracts that are typically traded. In November 2007, the number of different options listed on the Euwax crossed the threshold of 100,000 for the first time; in addition to these plain vanilla options, the Euwax listed more than 150,000 other more exotic derivatives and structured products. The number of bank-issued options dwarfs the number of stocks or options listed on traditional exchanges.

Retail investors who desire option-like payoffs clearly benefit from the ability to trade covered warrants, especially if they are shut out of traditional derivatives exchanges. It is much simpler for them to, say, buy a call option on the German DAX than to dynamically trade DAX stocks and the riskless asset so as to replicate the option's payoffs. It is less clear why they would benefit from having to choose between roughly 3,000 different DAX calls, as was the case on Euwax in November 2007, when many of the calls have highly similar payoff patterns. One explanation for the large number of different contracts is the standard argument that financial intermediaries respond to rational investors' desiring more complete markets. Another explanation is that a combination of low issuing costs and boundedly rational investors facing a costly search for the best option sustains an equilibrium with product proliferation and price dispersion.

The paper addresses two questions to discriminate between these two explanations:

First, given that investors face a choice between several slightly different instruments, how do investors select among the available alternatives? The choices of rational investors should be purely driven by criteria related to the objective value of the option as an investment. For example, a warrant with a lower spread should be preferred to a warrant with a higher spread, other attributes being equal. In practice, however, investors may also rely on other attributes or even choose to ignore relevant information which leads to decisions that are sub-optimal relative to the rational benchmark.

Second, how do the options chosen perform relative to other options with similar characteristics that could have been bought, especially to those judged ex ante superior based on objective criteria such as spreads? Do trades made by more experienced or larger investors – who presumably have lower search costs or stronger incentives to search – perform relatively better?

To answer these questions, the paper combines data on retail trades in DAX call options made at a large German discount broker and quotes for the population of DAX calls listed on the Euwax during the period November 1999 to May 2000. The DAX 30 Performance Index is the most popular underlying for warrants listed on the Euwax; Bartram and Fehle (2007) estimate that, in terms of Euwax trading volume in DAX options, the DAX would have ranked among the top five underlying assets on the Chicago Board Options Exchange (CBOE) during the sample period.

In a panel logit model of option choice, the decisions of investors are found to be influenced both by option attributes that should matter (such as spreads), and by attributes that should be irrelevant (such as the exchange or conversion ratio, that is, the number of units necessary to get a payoff of EUR 1 if the DAX finishes 1 point above the strike price at expiration), other things being equal. Preferences for some of the attributes are at odds with rational choices. For example, favoring options with lower conversion ratios leads investors to focus on options that are relatively more expensive to trade because these options have nominally lower prices and relative spreads generally decrease in price.

The options actually chosen by the investors significantly underperform other options with a similar leverage, that is, a similar price elasticity with respect to changes in the DAX. The difference between the roundtrip returns of options bought and the corresponding returns of similar options not bought averages -0.7%; the average duration of a completed roundtrip is less than three days. The median return is also negative, suggesting that the sample investors would have been better off picking options randomly. Actual choices fare even worse when compared with similar options that are predicted to outperform based on information available to investors at the time of purchase, for example, because they have had lower spreads in the past. Options bought underperform similar options with the lowest spreads by about 1% per roundtrip, on average. To put this number into perspective: the average roundtrip return during this sample period in which the DAX appreciated by 22% was 1.6%.

In short, there is substantial price dispersion in the market for covered warrants even for products such as homogeneous plain vanilla equity index options whose pricing is fairly well understood. Moreover, retail investors generally fail to identify options that represent good values. Consistent with a costly search conjecture, investors who submit small trades and novice option investors perform particularly poorly. The simple strategy of buying the option with the lowest spread (and a similar leverage as the option actually bought) would have increased the absolute profit of inexperienced option investors by EUR 100 per trade, on average.

Prior research on covered warrants and structured products has focused on the "average" pricing of these instruments as opposed to price dispersion or investor choice. Bartram and Fehle (2007) and Ter Horst and Veld (2008) report substantial price deviations between bank-issued options and comparable standardized options traded in Germany and the Netherlands. Wilkens et al. (2003), Muck (2006), Ter Horst and Veld (2008), and Henderson and Pearson (2009) report deviations between market and model prices for bank-issued derivatives traded in Germany, the Netherland, and in the U.S., especially at the time of issue. Bergstresser (2008) reports poor performance of a comprehensive global sample of structured equity products, especially before 2005. Bernard and Boyle (2008) report that the popularity of structured equity products with local caps and global floors is consistent with investors overweighing the probability of certain rare events that feature prominently in the selling prospectus. The tenor of these papers is that issuers use the complexity of structured equity products to extract rents from unsophisticated or biased retail investors. The paper by Poteshman and Serbin (2003) is one of the few to analyze actual trading decisions of retail option investors. They report that 2-3% of CBOE option exercises by retail investors can be classified as irrational.

The documented dispersion of prices for relatively homogenous equity index options is reminiscent of the dispersion of expense ratios among S&P 500 index funds reported by Hortaçsu and Syverson (2004). Hortaçsu and Syverson (2004) attribute the dispersion of expense ratios to search costs, nonportfolio differentiation (such as differential tax exposures of the fund) and switching costs. Elton et al. (2004) report that variation in nonportfolio attributes fails to explain much of the variation in S&P 500 index fund flows. In experiments with Harvard staff, Wharton MBAs, and Harvard college students, Choi et al. (2009) report that subjects overwhelmingly fail to identify the lowest cost alternative when faced with a choice of four S&P 500 index funds stripped of any nonportfolio differences.

It is possible, indeed likely, that part of the observed price dispersion is due to well-documented option price anomalies (relative to the Black and Scholes (1973) model). If differences in implied volatilities across the sample options were solely due to properties of the DAX return generating process, one would expect similar price dispersions in the market for bank-issued derivatives (Euwax) and in the market for exchange-traded derivatives (Eurex). Figure 1 illustrates that the standard deviation of Euwax implied volatilities is much larger than that of Eurex implied volatilities. One explanation, consistent with the other results presented in the paper, is that Eurex traders (predominantly institutions) have lower search costs and stronger incentives to search than Euwax traders (predominantly individuals).

The paper is organized as follows. The next section introduces the sample of warrants studied in the paper. Section III examines what makes retail investors choose a particular option from the pool of available options. Section IV evaluates the quality of the investors' choices. Section V reflects on the results in light of more recent developments in the market for covered warrants. Section VI concludes. The Appendix describes features of the market for covered warrants in Germany, compares option prices across different trade venues, and offers robustness checks for results discussed in Section IV.

II The Data

This paper focuses on bank-issued options, specifically on call options on the DAX 30 Performance Index traded on the Euwax between November 22, 1999 and May 31, 2000. The sample period is determined by the availability of Euwax quotes (available from November 22, 1999, onwards) and retail transaction records from one of the three largest German discount brokers (available until May 31, 2000).

DAX options are chosen because the DAX is the most common underlying asset for covered warrants in Germany. Glaser and Schmitz (2007) report that transactions in DAX options account for more than one third of all transactions in a sample of option trades made by a sample of German discount brokerage clients between 1997 and 2001. Bartram and Fehle (2007) estimate that, in terms of Euwax trading volume in DAX options, the DAX would have ranked among the top five underlying assets on the CBOE during the sample period. The population of DAX calls is identified using a register of securities listed on German stock exchanges obtained from the Karlsruher Kapitalmarktdatenbank (KKMDB) which makes data feeds from Deutsche Börse available to academics at cost. The register provides the security's local identification code (German cusip), the name of the issuer, the name of the underlying, and the expiration date. Frequently, the nature of the option (call or put) and the strike price can also be inferred. The cusip/ISIN can then be used to verify the static option features and obtain additional information from records maintained by the Börsen-Zeitung, the official publication of the German stock exchanges. A total of 668 DAX calls that are alive at some point during the sample period can thus be identified. The KKMDB also provides the daily trading volume separately for each option and German stock exchange.

Intraday Euwax quotes for 537 of these options, posted by the corresponding issuer in his role as market maker, are available from the Euwax website.[1] Despite Euwax's strong market position, some issuers (notably WestLB) choose to list their warrants on the Frankfurt Stock Exchange or one of the other regional exchanges, but not on Euwax.

This data can be merged with daily transaction records available for a large sample of German discount brokerage clients for the period January 1995 to May 2000. Dorn and Sengmueller (2009) provide a more detailed description of this sample which constitutes a non-trivial portion of the population of German discount brokerage customers. The focus of this paper is the intersection of the data sets described above: 14,687 trades made by 1,043 distinct investors in 223 different DAX calls between November 22, 1999, and May 31, 2000. The difference between the number of calls with Euwax quotes (537) and the number of calls traded by the sample investors (223) is not due to trading restrictions; in principle, the sample investors could trade all the DAX calls listed on the Euwax. The results in this paper suggest that this striking difference is due to both strong preferences for options with particular attributes and costly search.

Figure 1 illustrates the price dispersion in terms of Black-Scholes implied volatilities across

[1]The data was downloaded at
$http://www.kbl.boerse-stuttgart.de/archivselect/rfmstamm.php?archivselect=euwaxfixed$

the more than 300 active DAX calls on May 31, 2000. The interquartile range of implied volatilities is 24.2% to 30.7%. In contrast, the interquartile range of implied volatilities computed from the prices of 90 active DAX calls traded on Eurex on the same date is 21.4% to 24.3%.[2] The two distributions are not directly comparable partly because the range of Eurex options features differs from that of Euwax. For example, the range of Eurex strike prices (6,000-10,000) is smaller than that of Euwax strikes (5,800 to 11,000). However, the Euwax implied volatilities are based on ask quotes only whereas the Eurex implied volatilities are based on transaction prices which reflect both sides of the substantial spread. With these caveats in mind, prices of Euwax options appear to be higher, on average, and much more dispersed than prices of comparable Eurex options. Bartram and Fehle (2007) attribute the higher average Euwax implied volatilities to the fact that individual investors are unable to write options on the Euwax. Moreover, Eurex spreads of close to 5% (Bartram and Fehle (2007)) appear to be too large to allow Euwax-Eurex arbitrage by institutional traders.

The transaction records identify the trading venue as either a regular exchange such as the Euwax or the issuer's platform. In the sample, 85% of the trades are executed on issuer platforms and only 15% go to regular exchanges. Trading with the issuer directly enables investors to trade before and after exchange hours. Trading on a regular exchange during normal hours, however, is more transparent since detailed price information is available; moreover, the issuer's quoting behavior as an Euwax market maker is monitored by the exchange.

III How Do Investors Choose Among the Available Options?

This section presents a panel logit model of option choice. The model answers the question "given that an investor buys a call on the DAX, what makes him choose that particular call rather than one of the other available DAX calls?" The modelling strategy is similar to that of Grinblatt and Keloharju (2001) who explore, given an investors' decision to sell a stock from his portfolio, what makes the investor sell that particular stock rather than another stock in his portfolio.

For each purchase of a call, record the characteristics of the call bought *as well as* the characteristics of other available calls that day that the investor chooses not to buy. Given n calls on a particular day, an investor's decision to purchase one of these calls on a particular day generates n observations – one "bought" observation (dependent variable equal to one) and $n-1$ "not bought" observations (dependent variable equal to zero), each with its associated options characteristics.

[2]The underlying data were obtained from Deutsche Börse's data webshop.

This empirical strategy controls for changes in the sample investors' opportunity set and addresses the possibility that the issuers' choices of options characteristics reflect anticipated buyer preferences. For example, investment banks may forecast strong preferences for out-of-the-money calls and consequently issue more calls with relatively high strike prices; simply because of the greater availability, more investors will end up buying out-of-the-money calls even if they lack the forecasted preferences. Given the modelling strategy, such a change of the investment opportunity set will not be interpreted as a preference for low moneyness because the greater availability of out-of-the-money options generates both more "bought" and more "not bought" observations with low moneyness.

One can think of three categories of option characteristics that can systematically affect investor choices.

The first category consists of attributes that *should* systematically affect decisions, controlling for other option attributes: for example, the option's liquidity, its implied volatility, and the issuer's credit rating. Other things being equal, one would expect rational investors to choose an option with lower trading costs, lower implied volatility, and issued by a bank with a higher credit rating.

A second category consists of attributes for which the average investor may have a preference even though rationality does not imply a particular preference in the aggregate. This category includes, for example, the elasticity of the option's price with respect to the price of the underlying and the option's time to maturity. Of course, a preference for a higher elasticity or longer maturity may be dictated by particular circumstances and thus vary across individuals and over time. Such variation cannot be accommodated within the model studied in this section, but will be considered in Section IV.

A third category consists of attributes that *should not* systematically affect rational decisions: for example, the option's conversion ratio and the time since issue. For each unit of the option held at expiry, an investor gets an amount in Euro equal to the product of the conversion ratio and the difference between the DAX level and the strike price (if the difference is positive). The most common conversion ratio is 1/100 but there are also ratios of 1/200 and 1/1000. Since there is no minimum number of units that need to be traded, the conversion ratio should be irrelevant for an investor's decision to purchase a particular option.[3]

[3]The maximum unit price of an option traded in the sample is EUR 27. Only 2 out of the sample purchases are for less than EUR 27. Option exercise is typically subject to a minimum of 100 units. However, since the underlying is a performance index that reflects reinvested dividends, investors should (and do) sell rather than

A Univariate Results

This subsection describes the option attributes in detail and summarizes their relation with the purchase decision. To obtain the univariate statistics reported in Table I, average the characteristics of options by investor and bought and not bought category. These investor averages then serve as the unit of observation for the statistics reported in Table I. Given the implicit assumption that observations are independent across investors, but not over time, the options bought significantly differ from the options not bought for each characteristic considered.

The first category of attributes are those that should matter for rational purchase decisions, other things being equal.

The spread, one measure of liquidity, is the most important component of trading costs. Moreover, it varies substantially across options. The relative spread for a given option and day is defined as

$$Spread = \frac{C_{\text{ask, open}} - C_{\text{bid, open}}}{\frac{C_{\text{ask, open}} + C_{\text{bid, open}}}{2}} \tag{1}$$

where $C_{\text{bid, open}}$ and $C_{\text{ask, open}}$ are the opening bid and ask quotes reported by Euwax. Because of the central importance of the spread for investor performance, the analysis in this paper focuses on DAX call options that list on Euwax and for which Euwax quotes are available (more than 80% of the universe). The relative spread of options bought is roughly twice the spread of options not bought, on average (1.3% versus 0.7%). This may seem surprising, at first glance, but can be explained by a preference for out-of-the-money calls that provide high leverage and tend to have lower prices (see below). Given that the absolute spread $C_{ask} - C_{bid}$ is commonly set at a fixed amount, such as two cents, a preference for out-of-the-money calls translates into buying options with higher relative spreads.

Trading volume is another measure of liquidity. The Karlsruher Kapitalmarktdatenbank (KKMDB) makes trading volume data available by option and date for each German stock exchange on which the sample option is traded (in general, however, this data does not include volume that is internalized by the issuer on his own platform). The volume share of a given option on a given day is the Euro trading volume in that option across all exchanges that day divided by the aggregate Euro trading volume across all options and exchanges that day. The average volume share of an option bought is 1.5% as opposed to 0.2% for options not bought. Trading volume should only be a consideration for investors submitting large orders since Euwax quotes are binding offers to trade a certain amount (up to EUR 3,000 or 10,000 units).

exercise before maturity.

Trading commissions, another element of transaction costs, are similar across different trading venues (that is, stock exchanges and issuers). One-way commissions average 0.9% of the traded amount and decline in trade size. Occasionally, issuers offer rebates for trades directed to their platforms during one-week promotions; one or more issuers run promotions in 5 out of the 27 weeks in the sample. Temporary rebates for trading fees range from a discount of 25% on trading commissions to a full reimbursement of trading commissions and spreads. Options with reduced fees represent 7.3% of the options bought, but only 5.6% of the options not bought.

Another measure for the attractiveness of an option is its implied volatility. The implied volatility for a given option and day is the standard deviation σ that solves the Black and Scholes (1973) equation

$$C_{\text{ask, open}} = SN(d_1) - Xe^{-rT}N(d_2), \text{ where} \tag{2}$$

$$d_1 = \frac{ln(S/X) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}, \text{ and} \tag{3}$$

$$d_2 = d_1 - \sigma\sqrt{T} \tag{4}$$

where S is the intraday value of the DAX 30 Performance Index at the time of the first reported Euwax ask quote. Since index changes are recorded every 15 seconds, this means that the maximum time difference between the option quote and the underlying quote is 7 seconds. DAX intraday data as well as certain static option features – security identification code, strike price X, time to maturity T, and issuer identity – come from the KKMDB. The static option features obtained from the KKMDB are checked against the option record in the Börsen-Zeitung, the official publication of the German stock exchanges. The records maintained by the Börsen-Zeitung also contain issue dates and exchange ratios. Datastream provides the daily German treasury yield curve at discrete points in time such as one week, one month, three months, and so on. The interest rate r for a given expiry date is the spot on the yield curve that is closest to that expiry date. The average implied volatility of the options bought is 28.9% as opposed to 33.4% for the options not bought.

An issuer's credit rating might also affect the attractiveness of an option since any payoff at expiry is conditional on the issuer being able to meet his obligation. Based on a Factiva search for the issuers' S&P counterparty risk rating one year before and during the sample period (Subject "Corporate/Industrial News"/"Funding/Capital"/"Corporate Credit Ratings"), all but one issuer fall into one of two rating categories during the sample in terms of their S&P counterparty risk rating: A (coded as 0), or AA (coded as 1). The one issuer without an S&P rating has a Fitch rating of "A" which is coded as a 0. The average rating score for options bought is 0.9, meaning that investors buy their options almost exclusively from AA-issuers; options not bought have an average rating score of 0.7.

The second category of attributes are those for which rational investors may, but need not, have a preference in aggregate.

An important aspect of an option investment for retail investors is the associated leverage. Leverage can be measured by the option's price elasticity with respect to changes in the DAX:

$$Elasticity = \frac{\partial C_{\text{ask, open}}}{\partial S} \cdot \frac{S * ratio}{C_{\text{ask, open}}} \tag{5}$$

where *ratio* is the conversion ratio. On average, the options bought have an elasticity of 10 (meaning that a 1% in increase in the DAX is associated with roughly a 10% increase in the price of an option bought) whereas the options not bought only have an elasticity of 6. Options with higher price elasticities are typically deeper out-of-the-money or have a shorter time to maturity. Because of their lower prices and because of fixed absolute spreads, higher-elasticity options tend to have higher relative spreads.

Options from certain issuers can not only be traded on Euwax (with the issuer as the market maker), but also directly with the issuer. One potential advantage of these platforms is that they are open outside the normal exchange trading hours from 9am to 5:30pm (for example, between 8am and 10:30pm). One potential disadvantage is that investors lack systematic information about platform quotes and how they compare to Euwax quotes. The transaction records identify the venue on which the orders are executed. During the sample period, about 85% of the trades go through the issuer's platform as opposed to an exchange. About 98% of the options bought could have been traded on issuer platforms as opposed to only 62% of options not bought.

The third category of attributes are those for which rational investors should not have a preference (unless such a preference is orthogonal to attributes that objectively affect the attractiveness of the option).

Investor favor options with lower exchange ratios (that is, lower prices). The average conversion ratio for options bought is 0.009 as opposed to 0.008 for options not bought.

Investors strongly favor options from familiar issuers. An investor is thought to be familiar with an issuer if he has already traded an option from the same issuer from which he buys the sample call. In more than 9 out of 10 cases, investors buy an option from a familiar issuer; in contrast, only half of all the options not bought are from a familiar issuer. Investors favor even more strongly familiar options – that is, instruments that they have already traded. In almost one third of the cases, investors buy a familiar instrument; in contrast, familiar instruments represent fewer than 1% of options not bought.

A related preference is that for options from issuers with a larger market share during the year before the sample period. The market share of a given issuer is computed as the number of trades in that issuer's options divided by the total number of option trades. Options bought are issued by a bank with a market share of 55%, on average, whereas options not bought are issued by a bank with a market share of only 17%, on average. An alternative measure of market share is the number of options listed by a particular issuer divided by the total number of options listed. Options bought are from issuers that account for 20% of listed options whereas options not bought are from issuers that account for 14% of listed options, on average.

B Multivariate results

Some of the univariate results presented above cannot be interpreted unambiguously. For example, investors tend to pick options with higher spreads, presumably because they prefer options with greater leverage which have higher spreads. It is possible, however, that investors prefer options with lower spreads conditional on leverage. Another example: investors might choose to become familiar with certain options because they are cheaper to trade, or because investors simply know about them through advertising. To distinguish between the different possibilities, a multivariate analysis is needed.

The panel logit specification is

$$P(\text{Option Purchased} = 1) = \Lambda(\beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3) \qquad (6)$$

where Λ is the logistic cumulative distribution function. X_1 is the vector of attributes that should matter for a rational purchase decision, that is, one would expect the coefficients β_1 to be significant in a particular direction: the relative spread lagged one trading day, a reduced fee dummy (indicating whether the broker or issuer offers a trading cost rebate for the option), the option's lagged volume share, lagged implied volatility, and a dummy indicating a "AA" counterparty credit rating by S&P (the "A"-issuer dummy is omitted). Lags are chosen for relative spread, volume share, and implied volatility to make sure that the attributes are available to the investor at the time of purchase and because current values are endogenous; for example, purchases on a given day will affect the volume share that day. X_2 is the vector of attributes for which rational investors may (but need not) have a preference: the log number of years to maturity, the option's price elasticity with respect to changes in the underlying, and a dummy variable indicating whether the option can be traded directly with the issuer. X_3 is the vector of attributes that should not matter to rational investors, that is, one would expect the coefficients to be indistinguishable from zero controlling for the other attributes: two dummies indicating exchange ratios of 1/200 and 1/1000, the log number of years since issue, the two measures of issuer market share, and the two familiarity dummies.

Without adjustments, the model's errors will be correlated across investors and over time. The primary source of error correlation in the cross section is that the number of options varies over time. Adding the inverse number of options available to trade on a given day – the unconditional probability of buying any available option that day – as a regressor addresses this problem. (Alternatively, one could add a full set of time dummies. The results using time dummies are similar to those reported below and thus omitted.) In addition, the coefficient standard errors are adjusted to allow same-investor errors to be arbitrarily correlated across time, using the adjustment recommended by Williams (2000). Observations with problematic quotes (option prices below intrinsic values or bid>ask or missing bid or ask) account for 8% of total observations and are dropped.

Table II reports the results of the regression. As one would expect, the propensity to buy an option decreases in the relative spread and increases in the option's volume share. Relative spreads, in particular, appear to be an important determinant of option choice; a one-standard deviation decrease in the relative spread (evaluated at the means of the other independent variables), is associated with a roughly 100% increase in the purchase probability. Perhaps surprisingly, variation in trading cost rebates and implied volatility are not associated with meaningful variation in the probability of purchase. Given an average holding period of 5 days, and the small credit rating differences between the issuers, it is not surprising that the credit rating dummy is insignificant.

Other things being equal, investors strongly favor options that can be traded directly with the issuer and options with a high price elasticity with respect to changes in the DAX. For example, a one-standard deviation increase in elasticity is associated with a roughly 60% increase in the probability of purchase, other things being equal.

Much of the explanatory power of the model, however, comes from variation in attributes that say little or nothing about the option's objective value as an investment. For example, investors are twice as likely to buy an option with a conversion ratio of 1/200 than an otherwise identical option with a conversion ratio of 1/100 (the only difference between these options is their nominal price). Other things being equal, investors appear to prefer younger options, that is, options that have been issued more recently. Such options may be more likely to capture the investors' attention since options are promoted through newspaper and direct mail advertising around the time they are issued. Moreover, investors prefer options issued by banks with a greater market share and options that they have traded before.

One straightforward interpretation of most of these results is that investors try to identify options with attractive features, for example, options that are cheap to trade. However, the sheer number of available options complicates this identification problem. Consequently, investors appear to rely on heuristics and tend to pick from what they know because of advertising or prior investment experience – similar to the mutual fund investors responsible for the flows observed by Sirri and Tufano (1998).

IV How Well Do Investors Choose?

Two broad results emerge from the empirical model of option choice presented in the last section. First, some option attributes that should matter do in fact matter. Second, some option attributes that are per se irrelevant for rational decision makers matter as well. Since the preferences for irrelevant attributes (such as conversion ratios) are at odds with the preferences for attributes that affect the option's objective value as an investment (such as relative spreads), investors likely make sub-optimal choices – that is, sub-optimal relative to those of a rational investor who uses past option attributes such as spreads to predict future performance. This section attempts to quantify the performance of the investors' option choices, both in absolute terms and relative to the available investment alternatives. The first subsection summarizes the actual performance of the investors. The second subsection estimates relative performance.

A Actual Performance

Actual performance is estimated as the roundtrip return

$$ret_{t,t+T}^{actual,broker} = \frac{P_{t+T}^{sell} - fee_{t+T}^{sell}}{P_t^{buy} + fee_t^{buy}} - 1 \quad (7)$$

where P_t^{buy} is the unit purchase price paid on day t, P_{t+T}^{sell} is the unit sales price received on day $t+T$ (with $T = 0, 1, 2, ...$), and fee_t^{buy} and fee_{t+T}^{sell} are the per unit purchase and sale commissions. P_t^{buy} and fee_t^{buy} are directly taken from the transaction records and adjusted for trading cost rebates where applicable. P_{t+T}^{sell} is the sale price recorded for the first sale following the option's purchase. If the first sale closes out the position established by the preceding purchase, which happens in two out of three cases, then fee_{t+T}^{sell} is directly taken from the transaction records. Otherwise, the selling commissions are calculated as if the investor had closed out the position. If the investor executes several sell orders on the same day instead of a single sell order, P_{t+T}^{sell} is the value-weighted average price received. Again, selling commissions are adjusted for rebates when necessary. If an investor does not sell an option during the sample period, which happens in 7% of the cases, P_{t+T}^{sell} is assumed to be the closing Euwax bid quoted on May 31, 2000, at the end of the last day of the sample period.

To examine actual performance further, two alternative return measures are calculated:

$$ret_{t,t+T}^{actual,Euwax} = \frac{P_{t+T}^{bid,close} - fee_{t+T}^{sell}}{P_t^{ask,open} + fee_t^{buy}} - 1 \quad (8)$$

$$ret_{t,t+T}^{actual,Euwax,\text{no costs}} = \frac{P_{t+T}^{bid,close}}{P_t^{bid,open}} - 1 \quad (9)$$

Equation 8 calculates the roundtrip return from buying at the opening Euwax ask quoted sometime after 9am on the day the investor actually bought the option (day t) and selling at the closing Euwax bid quoted sometime before 5:30pm on the day the investor actually sold the option for the first time after the purchase (day $t+T$, $T = 0, 1, 2, ...$). Unit trading commissions are calculated and adjusted as in Equation 7. To assess the impact of trading costs, Equation 9 calculates the corresponding roundtrip return before spreads and commissions. Observations with problematic Euwax quotes are excluded. Such observations include option prices below intrinsic values, bid>ask, missing bid or ask, and bid or ask not quoted within 20 minutes of market close or open.

Table III summarizes trade characteristics and actual trade performance as a function of the holding period. Three holding period categories are considered: one day (intraday), longer than one day but completed within the sample period, and not completed within the sample period.

Across all holding periods, actual roundtrip returns ($ret_{t,t+T}^{actual,broker}$) average 1.6%. Alternative roundtrip returns calculated from buying at the Euwax opening ask and selling at the Euwax closing bid on the actual days investors buy and sell ($ret_{t,t+T}^{actual,Euwax}$) are substantially lower, averaging only 0.1%. Alternative roundtrip returns before trading costs ($ret_{t,t+T}^{actual,Euwax,\text{no costs}}$) average 2.1%; trading costs represent 2% of the capital invested, on average.

The difference between actual returns and the corresponding returns calculated from Euwax opening ask and closing bid quotes could be due to order timing (the actual orders may not have been filled at the open or close) or due to traders obtaining better prices when trading with the issuer directly than when sending their orders to Euwax to be filled (most of the actual orders are not submitted to Euwax but directly to the issuer). The latter explanation seems unlikely given Euwax's best price principle – orders sent to the Euwax get filled at prices that are at least as good as or better than those of the market maker (that is, the issuer). Of course, it is possible that traders do get better execution from the issuer, for example, if it takes less time for their order to be routed to the issuer than to Euwax or if the order is larger than the maximum size for which Euwax requires market makers to honor posted quotes (EUR 3,000 or 10,000 units). Appendix B compares transaction prices across trading venues (regular exchanges and issuer platforms) and documents that option buyers tend to pay higher prices and option sellers tend to get lower prices when trading with the issuer than when trading on a regular exchange.

Order timing, driven by skill or luck, appears to be a better explanation for the discrepancy between actual returns and return estimates based on Euwax opening and closing quotes. During the sample period, Euwax trading closed at 5:30pm whereas trading with some issuers was possible until 10:30pm – a fact prominently advertised on the broker's web site. Given the positive market performance during the sample period, at least part of which was driven by information released after hours (such as earnings announcements or the performance of U.S. markets), waiting until after the Euwax close was likely profitable. To test the conjecture that such after-hours trading can explain the relatively high actual returns, one can compare the returns from Equation 8 with the hypothetical return from buying at the Euwax opening quote on day t (as before) and selling at the Euwax opening quote on day $t+T+1$ (to simulate selling after hours on day $t+T$):

$$ret_{t,t+T+1}^{actual,Euwax} = \frac{P_{t+T+1}^{bid,open} - fee_{t+T+1}^{sell}}{P_{t}^{ask,open} + fee_{t}^{buy}} - 1 \qquad (10)$$

In unreported results, average returns calculated according to Equation 10 are found to be statistically indistinguishable from average actual returns (calculated as in Equation 7). Thus, systematic selling after the Euwax close is one possible explanation for the discrepancy between actual returns and return estimates based on Euwax opening and closing quotes. The sample brokerage clients may have preferred to trade with the issuers after hours because they expected to get better prices (skill) or because it was more convenient for them to trade after coming home from work (luck). With the data at hand, it is difficult to pinpoint the cause for the difference in average returns.

Table III also shows that trade characteristics and trade performance differ substantially as a function of the holding period. More than half of all purchases are at least partly reversed intraday. Intraday purchases are roughly four times the size of purchases that are held longer, on average. Despite the high trading costs, intraday trades and other completed roundtrips are profitable which mainly reflects the 22% appreciation of the DAX index during the sample period. Hypothetical roundtrip returns on open positions at the end of the sample period are abysmal, largely because of the negative performance of the DAX at the end of the sample period.

B Performance Gap

One can compare the returns of the options actually bought with the returns available from the population of DAX calls, in general, and the subset of DAX calls that resembles the call actually chosen by the investor. A particularly interesting subset consists of DAX calls that are ex ante judged superior to the call actually chosen based on past option attributes.

Since the trading records are date-stamped, but not time-stamped, such a comparison requires assumptions about when exactly traders get their orders executed. Any particular set of assumptions implies little loss of generality as long as differences in the attractiveness of options persist over time. Such persistence can be conjectured in part because absolute spreads tend to remain fixed over the entire life of a given option. The results reported below will verify this conjecture.

The main trading strategy considered here assumes that the investors buys at the Euwax ask in effect at 9:20am (that is, twenty minutes after the market opens) on the day of the actual purchase and sells at the Euwax closing bid on the day of the actual sale. This trading strategy is feasible for orders that are small enough because the broker quickly relays orders to Euwax and because the Euwax market maker is obliged to honor publicly posted quotes up to a maximum of EUR 3,000 or 10,000 units. Less than 10% of the observations need to be dropped either because option prices are below intrinsic values, bid>ask, or bid/ask is missing. Another straightforward trading strategy consists of buying at the Euwax opening ask on the day of the actual purchase and selling at the Euwax closing bid on the day of the actual sale. This trading strategy, however, suffers from the fact that there is relatively more variation in the timing of the first quote after 9am across options. On average, the first quote is posted six minutes after 9am, with a standard deviation of four minutes. In contrast, the quotes in effect at 9:20am and 5:30pm are posted between one and two minutes before the time in question, on average. Of course, assuming 9:20am as the time of purchase is fairly arbitrary. Appendix C documents qualitatively similar results for two alternative trading strategies (buying at the Euwax open and buying at 9:40am instead of buying at 9:20am).

Baseline Results

Column (1) of Table IV summarizes the difference between the return of the option actually bought (option i) on day t and sold on day $t+T$ and the average return across the set S of available options

$$eret^{i}_{t,t+T} = ret^{i,actual,Euwax}_{t,t+T} - \frac{1}{||S||}\sum_{j\in S} ret^{j,actual,Euwax}_{t,t+T} \qquad (11)$$

where $ret^{i,actual,Euwax}_{t,t+T}$ is the return from buying option i at the Euwax ask in effect at 9:20am on day t and selling at the closing Euwax bid on day $t+T$, $T = 0, 1, 2, ...$, as similarly defined in Equation 8. An option is deemed available if it could have been bought and sold on the dates the actual option was bought and sold.

Across all observations, the average excess return of -0.1% is statistically indistinguishable from zero assuming that all observations are independent. This result mostly reflects the performance of intraday trades. Completed roundtrips with holding periods longer than one day average significant excess returns of 2.3% and incomplete roundtrips average significant excess returns of -12.6%. The superior excess performance of completed roundtrips with longer holding periods and the poor excess performance of the incomplete roundtrips can be explained by the performance of the DAX during the sample period coupled with the investors' preferences for options with a higher price elasticity. Completed roundtrips tend to happen earlier during the sample period during which the DAX rose; in contrast, incomplete roundtrips tend to happen after the market peak in March 2000 and hence reflect the subsequent negative performance. The investors' preference for higher-elasticity options amplifies the performance of the options actually bought relative to that of the other options.

Column (2) of Table IV summarizes a similar excess return measure which compares the performance of the option actually bought with the average return across the set of options whose price elasticity is within 10% of the actual price elasticity (for example, if the option bought has a price elasticity of 10, the set of comparable options is restricted to those with price elasticities between 9 and 11). Across all observations, the options bought by the investors underperform other options with a similar price elasticity by a statistically significant 0.7%. This translates into an average loss of EUR 20 per transaction relative to investing in an equally-weighted portfolio of similarly price-elastic options.

Matching options by their price elasticity is not only important because traders have a preference for higher-elasticity options; option performance crucially depends on this attribute during the sample period. This is illustrated by the statistics in Column (3) and (4) which summarize excess returns of options chosen over available options with higher elasticities (Column (3)) and lower elasticities (Column (4)). Due to the DAX performance during the sample period, options chosen tend to outperform options with lower elasticities and tend to underperform options with higher elasticities.

Columns (5) reports performance gap statistics by trader, that is, assuming that observations are independent across traders but not over time. Across all his option purchases, a trader loses an entire percentage point for a given purchase relative to similar options he could have bought, on average. This translates into an average loss of EUR 120 per trader relative to investing in an equally-weighted portfolio of similarly price-elastic options during the sample period.[4]

Results Conditioned on Trade Size and Trader Experience

The observed product proliferation and price dispersion should motivate traders to search for a good deal. If high search costs were to blame for the poor selection, one would expect traders with lower search costs or greater search incentives to make better choices. Higher past option trading activity of a trader, measured as the number of option trades made during the year prior to the sample period, proxies for greater experience (associated with lower search costs) and higher expected future trading activity (associated with stronger search incentives). Traders who submit larger orders should also have stronger incentives to search. Partitioning the sample along order size is also useful to evaluate a story that makes the opposite prediction. Euwax requires market makers to honor posted quotes only up to a volume of EUR 3,000 or 10,000 units. For investors who submit larger orders, the posted quotes thus may not be a good indication for the prices at which they trade. One might thus conjecture that observed inferior selection is concentrated in large trades.

Ex ante, what should traders look for, that is, what constitutes a good deal? Given that most purchases are reversed within days if not hours, a low spread is probably the most important aspect of a good deal. A purchase strategy based on spreads can be easily implemented because no calculation is required and because spreads are persistent; a low-spread option before the time of purchase will likely be a low-spread option at the time of purchase and sale as well. Given that option quotes are not updated continuously, traders may be able to identify temporarily mispriced options by comparing their implied volatilities against those of

[4]The number of traders reported in Column (5) of Table IV is smaller than the number of traders/clusters reported in Table II primarily because of different quote data requirements. The calculation of implied volatilities and spreads for Table II only requires that the Euwax market maker post one valid quote sometime during the previous trading day. The performance gap requires both a valid quote before 9:20am on the day of purchase and a valid quote between 5:10pm-5:30pm on the day of the first subsequent sale.

similar options at a given point in time. Such a strategy based on low implied volatility is more difficult to implement, although the sample broker, Euwax, and independent financial data providers (such as OnVista) maintain electronic databases that allow investors to rank options by their current implied volatilities.

Table V examines the performance gap for small trades and large trades separately. The maximum Euwax depth requirement, EUR 3,000 or 10,000 units (as chosen by the market maker), is an obvious threshold that divides the purchase sample roughly in half. The average large trade (EUR 18,000) is more than ten times the size of the average small trade (EUR 1,500). Given the average elasticity of options purchased (10), the average trade of EUR 10,000 generates profits and losses of a similar magnitude as stock positions of EUR 100,000.

Panel A of Table V shows the performance gap between the trader's choice and the option with the lowest spread among all the options with a similar price elasticity at the time of the trade. Consistent with the search cost conjecture, the performance gap of small trades (-1.3%) is significantly larger than that of large trades (-0.7%). These performance gaps translate into an average dollar loss of EUR 15 per small trade and EUR 100 per large trade.

Panel B of Table V shows the performance gap between the trader's choice and the option with the lowest implied volatility among all the options with a similar price elasticity at the time of the trade. Again, the performance gap of small trades (-4.5%) is significantly larger than that of large trades (-2.3%). The dollar losses relative to buying the option with the lowest implied volatility among all the options with a similar leverage are substantial; 50 EUR for small trades and 170 EUR for large trades, on average.

Panel C of Table V considers a variation of the spread-based trading strategy. Instead of conditioning on the lowest spread at the time of the actual purchase, it selects the option for which the lowest spread was quoted at the close of the last trading day. Although the performance gap estimates are not as precise, the average small trade (-1.8%) and large trade (-0.9%) gaps are quantitatively similar to the corresponding estimates in Panel A. Panel D considers a similar variation of Panel B by picking the option with the lowest implied volatility at the close of the last trading day. Again, the performance gap is significantly larger for small trades than for large trades, although the magnitude is smaller than that reported in Panel B.

Table VI examines the performance gap for trades as a function of the experience of the trader. Quartile 1 consists of trades made by the least experienced traders with an average of 3 option trades during the year prior to the sample period; Quartile 4 consists of trades by the most experienced traders with an average of 176 option trades under their belt. Although experience and trade size are positively related, experience is not simply a proxy for trade size; the median trade size in Quartile 4 is smaller than that in Quartile 3. The number of trades increases in the quartile number indicating that investors judged as more experience by their higher number of past option trades also trade more actively during the sample period.

The results in Table VI are again broadly consistent with search costs varying across investors. For example, Panel A of Table VI reports that options purchased by the least experienced traders exhibit twice the performance gap (-1.4%) as options purchased by the most experienced traders (-0.7%) – relative to the option with the lowest spread at the time of purchase. The average dollar losses relative to buying the option with the lowest spread among all the options with a similar leverage are again substantial. They range from around EUR 40 per trade for the most experienced investors to more than EUR 100 per trade for the relatively inexperienced investors in Quartile 2.

The performance gap between the options purchased and the population of available options, especially low spread or low implied volatility options, appears highly statistically significant and economically meaningful. The average performance gap is of the same order of magnitude as the average roundtrip return during a sample period in which the underlying appreciated by 22%. The money that appears to be left on the table is of the same order of magnitude as the spread, on average.

V Looking Ahead

The reader might be tempted to dismiss the results presented in the previous sections because the underlying sample is a bit dated and could reflect a market in its growing pains. Over time, bank-issued options may have become cheaper to trade and option investors may have learned to identify higher value investments. It is difficult to examine the learning conjecture without high-quality data on investor behavior. The conjecture that option trading has become cheaper over time (presumably as more issuers have discovered this market segment and stimulated competition) can be checked more easily.

For recent years, Datastream provides a searchable database of bank-issued DAX options that includes the options' security identifier (ISIN) for most issues. This, in turn, means that most of the Datastream option data can be merged with Euwax quote data, when available. To examine the conjecture that option spreads have declined over time, Table VII reports summary statistics of Euwax spreads on three trading days: 15 December 1999, 12 December 2007, and 10 December 2008. The two first dates were chosen such that the prior one-month performance of the index was roughly equal for both dates. (Given that spreads tend to be constant in absolute terms, poor performance will be associated with higher relative spreads. The opposite is true for positive performance.) Average relative spreads are not significantly different in 1999 and 2007 (1.3% in 1999 versus 1.5% in 2007); spread dispersion is slightly higher in 2007. Absolute spreads are significantly higher in 2007 than in 1999: 2 cents versus 1.7 cents. At the same time, the number of available investment alternatives increases more than tenfold from about 250 to almost 3,000. The statistics for December 2008 presumably reflect the higher market volatility and poor returns leading up to the snapshot date which makes for a difficult comparison. Although the number of plain vanilla DAX warrants declined substantially between 2007 and 2008, the total number of warrants with the DAX as their underlying was relatively steady at about 6,500.

The dramatic increase in the number of different DAX calls suggests that, if anything, the search problem for the best option has become more difficult over time. Moreover, the results are not consistent with the conjecture that spreads have trended lower over time.

VI Conclusion

Bank-issued option markets have grown rapidly during the last decade in many Asian and European countries. In some countries such as Germany, the trading volume in bank-issued options rivals that in options traded on traditional derivatives exchanges. Investors often face a choice between dozens of bank-issued options that differ only slightly in their attributes.

The standard explanation for the observed plethora of choices is that banks compete to respond to the product preferences of rational investors. An alternative explanation is that low issuing costs and at most boundedly rational investors sustain an equilibrium with product proliferation and prices that are dispersed and above marginal cost. The paper documents substantial price dispersion and poor investor choices which are broadly consistent with the alternative view.

The paper's results do not rule out that some investors benefit from the choice available in covered warrant markets. It is possible, though, that retail investors would be better off in an institutional setting found, for example, in Hong Kong, where higher listing and registration fees for covered warrants apparently lead to more concentrated trading in fewer listings. Analyzing the welfare implications of a listing fee increase is beyond the scope of this study, but perhaps an area for future research.

Carlin (2009) conjectures that professional advice can help uninformed investors better navigate the menu of choices, unless issuers raise industry complexity or offer advisors incentives to share in industry profits. During the last ten years, the complexity in covered warrant markets has substantially increased as banks have offered a rapidly increasing number of DAX-based derivatives. Moreover, much of the financial advice in Germany appears to be dispensed by employees affiliated with the issuers. It is unclear, however, whether one should expect much improvement over the observed choices even with an impartial advice channel. In principle, the sample investors have access to up-to-date information on option attributes provided by the sample broker (a discount broker), the options exchange Euwax, and third-party providers such as OnVista.

The paper does not address why retail investors trade or hold options in the first place. Options may just provide a simple means of placing leveraged bets on the market. Alternatively, the observed trades may be part of a sophisticated trading strategy that involves assets held outside the observed account. The interpretation of the paper's results, however, is likely unaffected by the actual motivation behind the trades.

References

Baden-Württembergische Wertpapierbörse (2001). Jahresbericht 2000 der Baden-Württembergischen Wertpapierbörse. EuWax, Stuttgart (Germany).

Baden-Württembergische Wertpapierbörse (2008a). Euwax-report 2007, Stuttgart, Germany.

Baden-Württembergische Wertpapierbörse (2008b). Gebührenordnung (fee schedule), Stuttgart, Germany.

Bartram, S. M. and Fehle, F. R. (2007). Competition without fungibility: Evidence from alternative market structures for derivatives, *Journal of Banking and Finance* **31**(3): 659–677.

Bergstresser, D. B. (2008). The retail market for structured notes: Issuance patterns and performance, 1995-2008, Working Paper.

Bernard, C. and Boyle, P. (2008). Structured investment products with caps and floors, Working Paper.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities, *Journal of Political Economy* **81**: 637–659.

Bundesministerium der Justiz (2005). Vermögensanlagen-Verkaufsprospektgebührenverordnung, *Bundesgesetzblatt* **I**: 1873.

Carlin, B. I. (2009). Strategic price complexity in retail financial markets, *Journal of Financial Economics* **91**(3): 278–287.

Choi, J. J., Laibson, D. and Madrian, B. C. (2009). Why does the law of one price fail? an experiment on index mutual funds, *Review of Financial Studies* . Forthcoming.

Dorn, D. and Sengmueller, P. (2009). Trading as entertainment?, *Management Science* **55**(4): 591–603.

Elton, E. J., Gruber, M. J. and Busse, J. A. (2004). Are investors rational? choices among index funds, *Journal of Finance* **59**(1): 261–288.

Glaser, M. and Schmitz, P. (2007). Privatanleger Am Optionsscheinmarkt (Retail Investors and the Market for Covered Warrants), *Zeitschrift für Bankrecht und Bankwirtschaft* **19**: 214–230.

Grinblatt, M. and Keloharju, M. (2001). What makes investors trade?, *Journal of Finance* **56**: 589–616.

Henderson, B. J. and Pearson, N. D. (2009). The dark side of financial innovation, Working Paper.

Hortaçsu, A. and Syverson, C. (2004). Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P Index funds, *Quarterly Journal of Economics* **119**: 403–456.

Huberman, G., Jiang, W. and Iyengar, S. (2004). *How Much Choice is Too Much?: Contributions to 401(K) Retirement Plans, Chapter in Pension Design and Structure: New Lessons from Behavioral Finance*, Oxford University Press, pp. 83–96.

Muck, M. (2006). Where should you buy your options? the pricing of exchange traded certificates and OTC derivatives in Germany, *Journal of Derivatives* **14**: 82–96.

Poteshman, A. M. and Serbin, V. (2003). Clearly irrational financial market behavior: Evidence from the early exercise of exchange traded stock options, *Journal of Finance* pp. 37–70.

Sirri, E. R. and Tufano, P. (1998). Costly search and mutual fund flows, *Journal of Finance* **53**(5): 1589–1622.

Ter Horst, J. and Veld, C. (2008). An empirical analysis of the pricing of bank issued options versus exchange options, *European Financial Management* **14**(2): 288–314.

White, H. (1980). A heteroskedasticity-consistent covariance estimator and direct test for heteroskedasticity, *Econometrica* **48**: 817–838.

Wilkens, S., Erner, C. and Röder, K. (2003). The pricing of structured products in Germany, *Journal of Derivatives* **11**: 55–69.

Williams, R. (2000). A Note on Robust Variance Estimation for Cluster-Correlated Data, *Biometrics* **56**: 645–646.

Figure 1: Distribution of Implied Volatilities of DAX Calls as of May 31, 2000

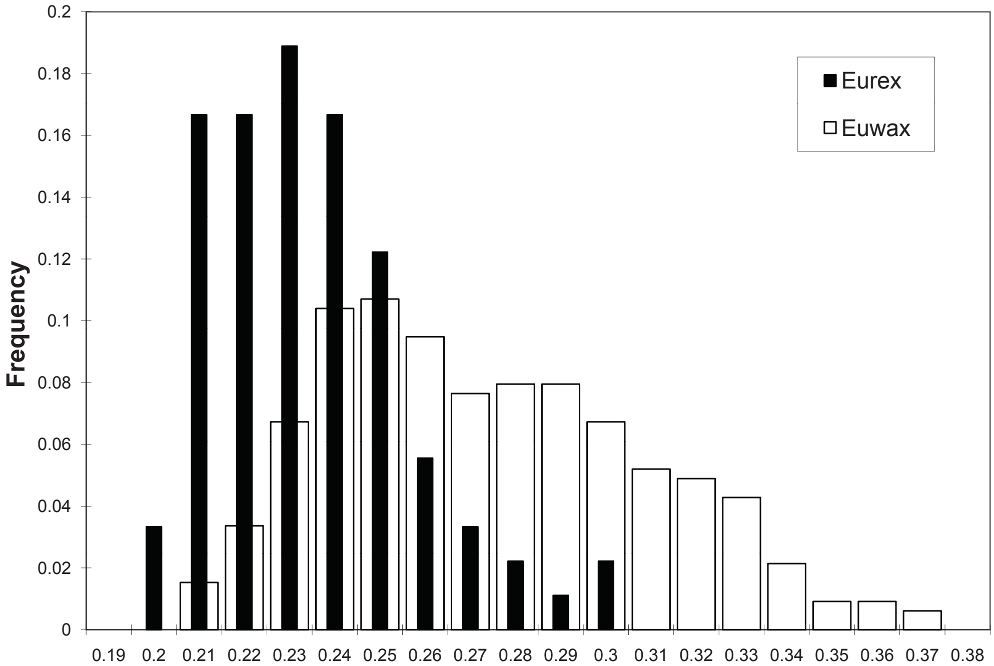

Table I: Attributes of Options Bought and Options Not Bought

Every time an option is purchased during the sample period, it is classified as an "Option Bought". All the other options that could have been bought instead are classified as "Options Not Bought." Attributes of "Options Bought" and "Options Not Bought" are averaged for each of the 886 investors. These investor-means form the unit of observation underlying the statistics reported below. "Lagged" means that the corresponding statistics are calculated using values lagged one trading day. "Relative Spread" is the difference of the Euwax opening ask and quotes relative to the quote midpoint. "Volume Share" is the sum of the EUR trading volume in the option across all German stock exchanges divided by the total EUR trading volume across all options and stock exchanges that day. The "Reduced Fees Dummy" is one if the purchase is eligible for at least a partial reimbursement of trading commissions and zero otherwise. "Implied Volatility" is the Black-Scholes volatility implied by the Euwax opening ask quote. The "AA Credit Rating Dummy" is one if the issuer has a "AA" long-term counterparty credit rating from S&P and zero otherwise. "Ln Years To Maturity" is the natural logarithm of the years remaining to maturity. "Price Elasticity" is the Black-Scholes implied elasticity of the option price with respect to the DAX. The "Issuer Platform Dummy" is one if the option can be traded directly with the issuer and zero otherwise. For each unit of the option held at expiry, the issuer pays the amount equal to the product of the "Conversion Ratio" and the difference between the DAX level and the strike price (if the difference is positive). "Ln Years Since Issue" is the natural logarithm of the years since issue. The "Issuer Market Share" is the volume share of the issuer at the sample broker between November 1998 and October 1999. The "Issuer Number Share" is the number of options offered by the issuer divided by the total number of options available. The "Issue Familiar Dummy" is one if the investor has traded the option before and zero otherwise. The "Issuer Familiar Dummy" is one if the investor has traded an option from the same issuer before and zero otherwise. All differences in means are significant at conventional levels of significance assuming that observations are independent across investors and time.

	Options Bought			Options Not Bought		
	Mean	Std	Median	Mean	Std	Median
Lagged Relative Spread	1.4%	5.8%	0.6%	0.7%	0.7%	0.4%
Lagged Volume Share	1.5%	1.8%	0.9%	0.2%	0.1%	0.2%
Reduced Fees Dummy	7.4%	19.5%	0.0%	5.8%	13.0%	0.0%
Lagged Implied Volatility	29%	8%	28%	33%	3%	33%
"AA" Credit Rating Dummy	0.9	0.3	1.0	0.7	0.0	0.7
Ln Years To Maturity	-1.1	0.9	-1.1	-0.8	0.1	-0.8
Lagged Price Elasticity	9.6	4.9	8.9	5.8	1.1	5.6
Issuer Platform Dummy	98%	14%	100%	62%	5%	61%
Conversion Ratio	0.008	0.002	0.009	0.009	0.000	0.009
Ln Years Since Maturity	-1.0	0.9	-0.9	-1.4	0.2	-1.4
Issuer Market Share	55%	23%	68%	17%	2%	17%
Issuer Number Share	20%	4%	21%	14%	1%	14%
Issue Familiar Dummy	32%	35%	25%	0%	1%	0%
Issuer Familiar Dummy	60%	38%	60%	48%	24%	47%

Table II: Panel Logit Model of Option Choice

The dependent variable is the "Option Bought" dummy which is one if the option was indeed bought and zero otherwise. All variables except the conversion ratio are defined as in Table I. The three conversion ratios (1/100, 1/200, and 1/1,000) are modeled as two dummy variables that are one if the conversion ratios are 1/200 and 1/1,000, respectively. The inverse of the number of active issues on a given day, the unconditional probability of buying any particular option that day, is an additional control variable (not reported). All reported standard errors are robust to heteroskedasticity and allow for clustering of errors across same-investor observations (see White (1980) and Williams (2000)). ***/**/* indicate that the coefficient estimates are significantly different from zero at the 1%/5%/10% level.

Dependent Variable	Option Bought
Constant	-12.864***
	(0.440)
Lagged Relative Spread	-26.200***
	(4.461)
Lagged Volume Share	5.677***
	(1.169)
Reduced Fees Dummy	-0.021
	(0.054)
Lagged Implied Volatility	-1.652**
	(0.831)
"AA" Credit Rating Dummy	-0.012
	(0.323)
Ln Years To Maturity	0.043
	(0.102)
Lagged Price Elasticity	0.146***
	(0.018)
Issuer Platform Dummy	2.235***
	(0.268)
Conversion Ratio = 1/200	0.631***
	(0.145)
Conversion Ratio = 1/1,000	0.470
	(0.690)
Ln Years Since Maturity	-0.130***
	(0.044)
Issuer Market Share	1.804***
	(0.301)
Issuer Number Share	-0.917
	(1.555)
Issue Familiar	3.638***
	(0.150)
Issuer Familiar	1.756***
	(0.183)
Ancillary Statistics	
Number of Observations	1,848,388
Number of Clusters	886
Other Control Variables	1/Number of Issues
Pseudo-R^2	39.9%

Table III: Actual Performance

Actual performance is computed three ways. $ret_{t,t+T}^{actual,broker}$ is the roundtrip return based on actual transaction prices and fees from the transaction records. A roundtrip is defined as a purchase on day t followed by the first sale of the same option on day $t+T$, $T = 0, 1, 2, ...$ $ret_{t,t+T}^{actual,Euwax}$ is the roundtrip return assuming that the investor buys at the Euwax opening ask on the actual day of the purchase and sells at the Euwax closing bid on the actual day of the sale and that the investor pays the same fees he actually paid. $ret_{t,t+T}^{actual,Euwax,\text{no costs}}$ is the roundtrip return assuming that the investor buys at the Euwax opening bid on the actual day of the purchase and sells at the Euwax closing bid on the actual day of the sale and that the investor pays no trading commissions. If the roundtrip is incomplete, each of the three return measures are computed assuming that the investor sells at Euwax closing bid on the last day of the sample period (May 31, 2000).

	All		Completed, Intraday		Completed, Not Intraday		Not Completed	
Number of observations	4,175		2,294		1,608		273	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median
Holding period [# days]	5	0	0	0	6	3	44	36
Trade size [EUR]	11,236	3,328	17,013	5,252	4,448	2,340	3,555	1,760
$ret_{t,t+T}^{actual,broker}$	1.6%	1.6%	1.3%	1.2%	8.3%	5.6%	-35.7%	-35.6%
$ret_{t,t+T}^{actual,EUWAX}$	0.1%	-1.4%	-0.9%	-1.9%	7.7%	3.0%	-36.8%	-35.3%
$ret_{t,t+T}^{actual,EUWAX,\text{no costs}}$	2.1%	0.3%	0.7%	-0.8%	10.3%	5.0%	-34.5%	-33.0%

Table IV: Performance Gap – Baseline Results

The performance gap is the difference between the return of the option actually bought and the return of an equally weighted portfolio of a specific subset of options that could have been bought as well (including the option bought). The performance gap is calculated for each roundtrip trade in the sample. To make the returns comparable across options, it is assumed that all options are bought shortly after the Euwax market open at the ask quote in effect at 9:20am on the actual day of purchase (day t) and sold at the Euwax closing bid (at 5:30pm) on the actual day of the first sale following the purchase (day $t+T$). If no sale follows the purchase, it is assumed that the investor sells at the Euwax closing bid on the last day of the sample period (May 31, 2000). Column (1) reports excess return statistics for the subset of all available options. The subset of options underlying Columns (2) and (5) contain options whose price elasticity is within 10% of the price elasticity of the option actually bought. The subset of options considered in Columns (3) and (4) are those with strictly higher and strictly lower elasticities than the option actually bought. In Column (5), the performance gap is first averaged across all transactions for a given trader, and second across all traders. ***/**/* indicate that the performance gap is significantly different from zero at the 1%/5%/10% level.

	(1)	(2)	(3)	(4)	(5)
Subset filter	None	Matched Elasticity	Higher Elasticity	Lower Elasticity	Matched Elasticity
Unit of observation	Trade	Trade	Trade	Trade	Trader
Panel A: All Observations					
Number of Observations	4,231	4,231	4,130	4,230	758
Mean Return	-0.1%	-0.7%***	-2.9%***	0.6%**	-1.0%***
Std of Returns	16.7%	5.0%	27.3%	18.2%	5.9%
Panel B: Intraday Roundtrips					
Number of Observations	2,327	2,327	2,275	2,327	
Mean Return	-0.2%	-0.5%***	-2.1%***	0.2%	
Std of Returns	10.3%	3.4%	26.7%	11.3%	
Panel C: Roundtrips Completed in $\geq$ 1 Day					
Number of Observations	1,631	1,631	1,584	1,630	
Mean Return	2.3%***	-0.5%***	-6.1%***	4.2%***	
Std of Returns	21.7%	6.0%	29.2%	23.3%	
Panel D: Incomplete Roundtrips					
Number of Observations	273	273	271	273	
Mean Return	-12.6%***	-2.9%***	8.5%***	-18.0%***	
Std of Returns	20.9%	8.2%	14.9%	20.2%	

Table V: Performance Gap Conditioned on Trade Size

Column (1) reports performance gap statistics for small trades (those with a volume of up to EUR 3,000 and involving up to 10,000 units). Column (2) reports the corresponding statistics for large trades (those with either a volume exceeding EUR 3,000 or involving more than 10,000 units). The performance gap is defined as the difference between the return of the option actually bought and the return of the option whose leverage is within 10% of that of the option bought, but that also has the lowest spread at 9:20am on the actual day of purchase (Panel A), the lowest implied volatility at 9:20am on the actual day of purchase (Panel B), the lowest spread at the close of trading on the day prior to the purchase (Panel C), and the lowest implied volatility at the close of trading on the day prior to the purchase (Panel D). In Columns (1) and (2), ***/**/* indicate that the performance gap is significantly different from zero at the 1%/5%/10% level. In Column (3), ***/**/* indicate that the average performance gap for small trades is significantly different from the average performance gap for large trades at the 1%/5%/10% level.

	(1)	(2)	(3)
	Trade size		(2)-(1)
	Small	Large	
Mean trade size [EUR]	1,487	18,257	
Median trade size [EUR]	1,455	8,000	
Panel A: Relative to option with lowest spread at time of trade			
Number of Observations	1,986	2,245	
Mean Return	-1.3%***	-0.7%***	0.6%***
Std of Returns	9.0%	4.4%	
Panel B: Relative to option with lowest IV at time of trade			
Number of Observations	1,986	2,245	
Mean Return	-4.5%***	-2.1%***	2.4%***
Std of Returns	27.6%	18.8%	
Panel C: Relative to option with lowest spread at close of t-1			
Number of Observations	1,986	2,245	
Mean Return	-1.8%***	-0.9%***	0.9%
Std of Returns	22.9%	14.9%	
Panel D: Relative to option with lowest IV at close of t-1			
Number of Observations	1,986	2,245	
Mean Return	-1.7%***	-0.2%	1.5%***
Std of Returns	23.4%	6.0%	

Table VI: Performance Gap Conditioned on Option Trader Experience

Column (1) reports performance gap statistics for trades made by the quartile of the least experienced option traders. Experience is measured as the number of option trades between November 1998 and October 1999, that is, the year before the sample period. Columns (2)-(4) report the corresponding statistics for trades made by traders in experience quartiles 2-4. The performance gap is defined as the difference between the return of the option actually bought and the return of the option whose leverage is within 10% of that of the option bought, but that also has the lowest spread at 9:20am on the actual day of purchase (Panel A), the lowest implied volatility at 9:20am on the actual day of purchase (Panel B), the lowest spread at the close of trading on the day prior to the purchase (Panel C), and the lowest implied volatility at the close of trading on the day prior to the purchase (Panel D). In Columns (1)-(4), ***/**/* indicate that the performance gap is significantly different from zero at the 1%/5%/10% level. In Column (5), ***/**/* indicate that the average performance gap for trades made by the least experienced traders is significantly different from the average performance gap for trades made by the most experienced traders at the 1%/5%/10% level.

	(1)	(2)	(3)	(4)	(5)
		Option trading experience			(4)-(1)
	Lowest			Highest	
Number of traders	189	190	189	190	
Avg # option trades	3	18	47	176	
Median # option trades	2	17	45	124	
Mean trade size [EUR]	5,353	6,535	7,337	16,022	
Median trade size [EUR]	2,776	2,845	4,355	3,302	
Panel A: Relative to option with lowest spread at time of trade					
Number of Observations	592	620	975	2,044	
Mean Return	-1.4%***	-1.2%***	-1.1%***	-0.7%***	0.7%**
Std of Returns	8.9%	7.4%	7.4%	5.7%	
Panel B: Relative to option with lowest IV at time of trade					
Number of Observations	592	620	975	2,044	
Mean Return	-7.0%***	-5.3%***	-2.5%***	-1.9%***	5.1%***
Std of Returns	42.4%	26.3%	19.2%	14.6%	
Panel C: Relative to option with lowest spread at close of t-1					
Number of Observations	592	620	975	2,044	
Mean Return	-4.1%**	-1.6%*	-1.1%**	-0.6%***	3.5%***
Std of Returns	39.5%	20.8%	16.2%	6.7%	
Panel D: Relative to option with lowest IV at close of t-1					
Number of Observations	592	620	975	2,044	
Mean Return	-3.3%**	-1.6%*	-0.2%	-0.3%**	3.0%***
Std of Returns	36.2%	20.3%	7.5%	6.2%	

Table VII: Spreads During and After the Sample Period

The statistics below summarize DAX call option spreads quoted on Euwax on 15 December 1999, 12 December 2007, and 10 December 2008. The absolute spread is the difference between the opening Euwax ask and bid quotes. The relative spread is the absolute spread divided by the quote midpoint.

Date	# DAX call options	Spread Type	Mean	Spread Std	Median
December 15, 1999	238	Relative	1.32%	7.51%	0.27%
		Absolute	0.017	0.005	0.020
December 12, 2007	2,974	Relative	1.49%	8.54%	0.28%
		Absolute	0.020	0.090	0.020
December 10, 2008	1,674	Relative	47.76%	69.46%	7.66%
		Absolute	0.033	0.125	0.020

A The Market for Covered Warrants in Germany

Covered warrants are non-standardized derivatives written by large banks such as Citigroup, Commerzbank, and Deutsche Bank. The payoffs of these derivatives depend on the price movements in the underlying asset – individual stocks, stock indices, currencies, commodities, and bonds. Issuers state in the prospectus that they continuously hedge their exposure by taking positions in the underlying assets. This does not mean, however, that covered warrants are free from issuer default risk as reported by Bartram and Fehle (2007). The prospectus also states that covered warrants are senior unsecured obligations of the issuer and thus subject to issuer default risk. The attribute "non-standardized" refers to the fact that the banks writing the warrants have full discretion regarding the characteristics of the derivatives – in the case of options, for example, underlying, strike price, issue date, maturity date, and earliest exercise date.

Investors can trade these warrants either on regular exchanges, where the issuer typically serves as the market maker, or directly through the issuer as many issuers provide an electronic link between retail brokers and their own trading platforms. Retail warrant investors are effectively short-sale constrained – none of the major retail brokers allows his retail clients to write options even if the clients maintain a long position in the underlying.

The most important German exchange for covered warrants, and one of the largest of its kind in the world (see Glaser and Schmitz (2007)), is the European Warrant Exchange (Euwax) based in Stuttgart. Its market share of covered warrant trading in Germany averaged 60% in 1999 and increased to 80% by the end of 2000 (see Baden-Württembergische Wertpapierbörse (2001)). In 2007, the Euwax listed roughly 250,000 different covered warrants and reported a trading volume of EUR 128 billion (trading volume here refers to the price paid for the warrant, not the notional; see Baden-Württembergische Wertpapierbörse (2008a)). It is difficult to assess the magnitude of the volume executed on the issuers' proprietary trading platforms because this volume is not publicly reported and not included in the Euwax trading volume reported above. In the sample considered in this paper, more than 80% of the trading volume is handled by the issuer directly which suggests that Euwax trading volume is a conservative estimate of the total trading volume in covered warrants. To put these numbers in perspective: During the same period, the EuRex (one of the world's largest derivatives exchanges based in Frankfurt) reported a premium volume of EUR 455 billion in roughly 15,000 different options (author's estimates from the Eurex website $http : //www.eurex.de$). Hence, if the sample ratio of proprietary platform to exchange volume were representative of the aggregate ratio, the trading activity in covered warrants would be similar to that in traditional exchange-traded options. During the sample period, retail investors at major retail brokers could not participate in the trading of standardized options contracts that takes place on the Eurex. Bartram and Fehle (2007) provide a detailed comparison of the Euwax and the Eurex.

The focus of this paper is on the simplest type of covered warrants: plain-vanilla options. The issuing process of these options is quick and cheap which partly explains the plethora of instruments listed. A selling prospectus, which can be used to issue entire series of options, can be approved within 10 business days by the German securities regulator BAFIN at a cost of EUR 1,000 (Bundesministerium der Justiz (2005)). The issuer further bears the cost of publishing the prospectus as well as exchange listing fees. Euwax listing fees are EUR 500 per option (Baden-Württembergische Wertpapierbörse (2008b)), although Euwax caps listing fees by issuer such that issuers with more than 80 new listings in a calendar year have zero marginal listing fees; large banks such as Deutsche Bank issued hundreds of covered warrants during 1999. As part of the Euwax listing, the issuer commits to making a liquid market in his issues (or to hiring another firm to do so). Market maker quotes on Euwax are subject to fixed minimum depth requirements (EUR 3,000 or 10,000 units). The issuer's commitment to quote spreads below a stated maximum is also published by Euwax, but varies across issuers. In most cases, the maximum absolute spread corresponds to roughly one Euro per unit. This maximum is hypothetical during the sample period given that almost all options have absolute spreads of less than 4 cents.

B Prices on Euwax versus Issuer Platforms

One can compare trades that are executed on issuer platforms in a given option and on a given day with trades that are executed on regular exchanges such as the Euwax in the same option and on the same day. For the sample of DAX calls considered in this paper, prices received from the issuer appear to be higher than those received on regular exchanges, but the difference is only marginally significant (as reported in Panel A of Table VIII). These comparisons are based on few observations; there are only 72 instances in which a DAX call is bought both on an issuer platform and on a traditional exchange on the same day and 110 observations in which a DAX call is sold the same day on both types of trading venues. Enlarging the sample to include all options traded by clients at the sample broker as well as all trades beginning on January 2, 1995, yields 4,533 option-days on which an option is bought and 3,510 option-days on which an option is sold on both types of trading venues on the same day. For the larger sample of observations, average purchase prices are significantly higher and average sale prices are significantly lower on issuer platforms than the corresponding purchase and sale prices on regular exchanges. Panel B documents that the difference between issuer and exchange prices averages 1.6% for option purchases and -0.9% for option sales; option buyers on issuer platforms pay significantly higher prices and option sellers on issuer platforms get significantly lower prices than their peers who trade on regular exchanges on the same day. Absent time-stamped or tick data for both trading venues, however, it is difficult to discriminate between different explanations – issuers quote larger spreads, issuers handle more trades that require quote adjustments (e.g., larger trades), smarter investors trade on regular exchanges without affecting the adverse selection component of the spread, statistical fluke – for the observed price discrepancies.

C Alternative Trading Strategies

To compute the performance gap between options purchased and the population of available options (or subsets thereof), the main part of the paper assumes that the sample investors buy shortly after the open (9:20am) on the day the option was actually purchased, and sell at the close (5:30pm) on the day the option was actually sold for the first time. To examine the robustness of these assumptions, this section considers two alternative trading strategies: buying directly at the open and buying at 9:40am, instead of 9:20am. Panel A of Table IX suggests that the performance gap (relative to the universe of available options with similar leverage) indeed varies over time, but is significantly negative under all three sets of assumptions. Moreover, the performance gap relative to the option with the lowest spread is significantly negative and quantitatively similar across all trading strategies (see Panel B). Similarly, the performance gap for small trades is significantly worse than that for large trades across all trading strategies (see Panel C). In sum, it appears that changes in the assumptions about when orders are executed have little effect on the estimated performance gap.

Table VIII: Comparison of Prices on Issuer Platforms and Regular Exchanges

Each day that a warrant is bought both on an issuer platform and on a regular exchange, one can calculate the difference between the average issuer prices and the average exchange prices and divide the difference by the average of the averages. Positive differences indicate that transaction prices on the issuer platform tend to be higher than transaction prices on regular exchanges. The corresponding observations can be generated for sales. Panel A reports statistics for the sample studied in the paper. Panel B reports statistics for the full data set of warrant trades. ***/**/* indicate that the average of the relative differences across all purchase or sale observations are significantly different from zero at the 1%/5%/10% level.

Panel A: Trades of DAX calls, November 1999 - May 2000				
	Nobs	Mean	Std	Median
Issuer-Exchange (Purchases)	72	-0.1%	11.8%	0.0%
Issuer-Exchange (Sales)	110	1.7%*	10.0%	0.7%
Panel B: Trades of all bank-issued warrants, January 1995 - May 2000				
	Nobs	Mean	Std	Median
Issuer-Exchange (Purchases)	4,533	1.6%***	10.0%	1.1%
Issuer-Exchange (Sales)	3,510	-0.9%***	12.0%	0.0%

Table IX: Performance Gaps Under Alternative Trading Strategies

Column (1) assumes that options are bought at the Euwax opening ask on the actual day of purchase, Column (2) assumes that they are bought at the Euwax ask in effect at 9:20am, and Column (3) assumes that they are bought at the Euwax ask in effect at 9:40am. For all three trading strategies, options are assumed to be sold at the Euwax closing bid (at 5:30pm) on the actual day of the first sale following the purchase. If no sale follows the purchase, it is assumed that the investor sells at the Euwax closing bid on the last day of the sample period (May 31, 2000). Panel A reports the performance gap for the subset of options whose price elasticity is within 10% of the price elasticity of the option actually bought. The subset of options considered in Panel B are similarly price-elastic options that have the lowest spread at the time of purchase (for Columns (2) and (3)) and the lowest spread at the close of the prior trading day (for Column (1)). Panel C reports the difference between the performance gap for large trades and the performance gap for small trades. A positive value indicates that the returns of options bought in large volumes compare more favorably to those of similarly price-elastic options than those of options bought in small volumes. In Panels A and B, ***/**/* indicate that the performance gap is significantly different from zero at the 1%/5%/10% level. In Panel C, ***/**/* indicate that the difference in performance gaps between large and small trades is significantly different from zero at the 1%/5%/10% level.

	(1)	(2)	(3)
Trader buys at	Open	9:20	9:40
Panel A: Matched elasticity			
Number of Observations	4,253	4,231	4,328
Mean Return	-0.4%***	-0.7%***	-0.1%*
Std of Returns	4.8%	5.0%	3.9%
Panel B: Matched elasticity, lowest spread			
Number of Observations	4,253	4,231	4,328
Mean Return	-0.7%***	-1.0%***	-0.7%***
Std of Returns	7.4%	6.9%	7.4%
Panel C: Performance gap (Large - small)			
Mean	0.5%***	0.6%***	0.5%***