



EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 1375 / AUGUST 2011

**PRECAUTIONARY
PRICE STICKINESS**

by James Costain
and Anton Nakov



EUROPEAN CENTRAL BANK

EUROSYSTEM



WORKING PAPER SERIES

NO 1375 / AUGUST 2011

PRECAUTIONARY PRICE STICKINESS

by James Costain¹
and Anton Nakov²

NOTE: This Working Paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

In 2011 all ECB publications feature a motif taken from the €100 banknote.

This paper can be downloaded without charge from <http://www.ecb.europa.eu> or from the Social Science Research Network electronic library at http://ssrn.com/abstract_id=1912484.

¹ Banco de España, C/Alcalá 48, 28014 Madrid, Spain; e-mail: james.costain@bde.es

² Banco de España, C/Alcalá 48, 28014 Madrid, Spain and European Central Bank, Kaiserstrasse 29, D-60311 Frankfurt am Main, Germany; e-mail: Anton.Nakov@ecb.europa.eu

© European Central Bank, 2011

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

All rights reserved.

Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

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	4
Non-technical summary	5
1 Introduction	7
1.1 Related literature	9
2 Sticky prices in partial equilibrium	9
2.1 The monopolistic competitor's decision	10
2.2 Alternative models of sticky prices	12
2.3 Deriving logit choice from a control cost function	12
2.4 Relation to information-theoretic models	14
3 General equilibrium	14
3.1 Households	14
3.2 Monetary policy and aggregate consistency	15
3.3 State variable	16
3.4 Detrending	17
4 Computation	18
4.1 Outline of algorithm	18
4.2 The discretized model	19
4.3 Step 1: steady state	21
4.4 Step 2: linearized dynamics	22
5 Results	23
5.1 Parameterization	23
5.2 Steady state microeconomic results	24
5.3 Effects of changes in monetary policy	25
5.4 Changing the risk of errors	27
5.5 Logit equilibrium versus control costs	28
6 Conclusions	28
References	29
Tables and Figures	33

Abstract

This paper proposes two models in which price stickiness arises endogenously even though firms are free to change their prices at zero physical cost. Firms are subject to idiosyncratic and aggregate shocks, and they also face a risk of making errors when they set their prices. In our first specification, firms are assumed to play a dynamic logit equilibrium, which implies that big mistakes are less likely than small ones. The second specification derives logit behavior from an assumption that precision is costly. The empirical implications of the two versions of our model are very similar. Since firms making sufficiently large errors choose to adjust, both versions generate a strong "selection effect" in response to a nominal shock that eliminates most of the monetary nonneutrality found in the Calvo model. Thus the model implies that money shocks have little impact on the real economy, as in Golosov and Lucas (2007), but fits microdata better than their specification.

Keywords: logit equilibrium, state-dependent pricing, (S,s) adjustment, near rationality, information-constrained pricing

JEL Codes: E31, D81, C72

Non-Technical Summary

Economic conditions change continually. A firm that attempts to maintain an optimal price in response to these changes faces at least two costly managerial challenges. First, it must repeatedly decide when to post new prices. Second, for each price update, it must choose what new price to post. Since both decisions are costly, managers may suffer errors or frictions along either margin. The most familiar models of nominal rigidity have studied frictions in the first decision, assuming that price adjustments can only occur intermittently, either with exogenous frequency (as in Calvo, 1983) or with endogenous frequency (e.g. Golosov and Lucas, 2007; Dotsey et al. 2009; Costain and Nakov 2011A, B). This paper instead explores the implications of frictions in the second decision. In other words, we assume that firms can adjust their prices costlessly in any period, but that whenever they adjust, their price choice is subject to errors. We study the implications of frictions along this margin both for microeconomic price adjustment data and for macroeconomic dynamics.

Our key assumption is that the probability of setting any given price is a smoothly increasing function of the expected present discounted value of setting that price. Our model nests full frictionless rationality as the limiting case in which the firm sets the optimal price with probability one every period. More specifically, we impose this assumption on our model in two different ways. In the first version of our model, we assume the probability distribution over price choices is a logit function. The general equilibrium of this version of the model is therefore a logit equilibrium (McKelvey and Palfrey 1995, 1998): the probability of each firm's choice is a logit function which depends on the value of each choice; moreover, the value of each choice is determined, in equilibrium, by the logit choice probabilities of other firms.

While this version of the model does not impose price stickiness directly, the risk of errors gives rise to a certain degree of endogenous price stickiness. Because they fear they may “tremble” when choosing a new price, firms may refrain from adjusting, on precautionary grounds, even when their current price is not exactly optimal. Whenever the firm's price is sufficiently close to the optimum, it prefers to “leave well enough alone”, thus avoiding the risk of making a costly mistake. Hence behavior has an (S,s) band structure, in which adjustment occurs only if the current price is sufficiently far from the optimum.

In a second version of our model, we derive logit behavior from an assumption that precision is costly. We show that the distribution of price adjustments takes logit form if decision costs are proportional to entropy reduction. Thus, this specification involves a cost of price adjustment, but it is not a “menu cost” in the sense of labor time devoted to the physical task of adjusting posted prices. Instead, the cost of price adjustment can be understood as a cost of managerial control, consistent with the evidence of Zbaracki *et al.* (2004). The welfare losses due to price stickiness are lower in this interpretation of the model, since a stickier price allows for a reduction in managerial costs.

The logit framework for modeling bounded rationality has been widely applied in experimental game theory, where it has very successfully explained play in a number of games where Nash equilibrium performs poorly, such as the centipede game and Bertrand competition games (McKelvey and Palfrey 1998; Anderson, Goeree, and Holt 2002). It has been much less frequently applied in other areas of economics; we are unaware of any application of logit equilibrium inside a dynamic general equilibrium macroeconomic model. Another possible reason why macroeconomists have so rarely considered error-prone choice is that errors imply heterogeneity; the computational simplicity of a representative agent model may be lost if agents differ

because of small, random mistakes. However, when applied to state-dependent pricing, this problem is less relevant, since it has long been argued that it is important to allow for heterogeneity in order to understand the dynamics of “sticky” adjustment models (see for example Caplin and Spulber 1987, Caballero 1992, and Golosov and Lucas 2007). Moreover, we have shown (Costain and Nakov 2008B) how distributional dynamics can be tractably characterized in general equilibrium, without relying on special functional forms or questionable numerical aggregation assumptions. The same numerical method we used in that paper (Reiter 2009) can be applied to a logit equilibrium model; in fact, the smoothness of the logit case makes it even easier to compute than the fully rational case. We therefore find that logit equilibrium opens the door to tractable models with implications both for macroeconomic and for microeconomic data.

Summarizing our main findings, both versions of our model are consistent with several “puzzling” stylized facts from micro price adjustment data, in spite of the fact that the model has only one free parameter to estimate. Our model implies that many large and small price changes coexist (see Fig. 2), in contrast to the implications of a fixed menu cost model (Midrigan, 2010; Klenow and Kryvtsov, 2008; Klenow and Malin, 2009). It also implies that the probability of price adjustment decreases rapidly over the first few months, and then remains essentially flat (Nakamura and Steinsson, 2008; Klenow and Malin, 2009). Third, we find that the standard deviation of price changes is approximately constant, independent of the time since last adjustment (Klenow and Malin, 2009). Most alternative frameworks, including the Calvo model, instead imply that price changes are increasing in the time since last adjustment. Fourth, “extreme” prices are more likely to have been recently set than are prices near the center of the distribution (Campbell and Eden, 2010). While a variety of explanations have been offered for some of these observations (including sales, economies of scope in price setting, and heterogeneity among price setters), our framework matches all these facts in a very simple way, with only one degree of freedom in the parameterization.

Finally, we calculate the effects of money supply shocks in our framework. Given the degree of rationality that best fits microdata, the effect of money shocks on consumption is roughly twice as large as in the Golosov-Lucas (2007) fixed menu cost setup. The effect is much weaker than in the Calvo model because of a selection effect: all the firms that require the largest price adjustments do in fact adjust. Thus, a model in which price adjustment is slowed down by mistakes fits microdata much better than a fixed menu cost model, but implies that the macroeconomy is relatively close to monetary neutrality.

1 Introduction¹

Economic conditions change continually. A firm that attempts to maintain an optimal price in response to these changes faces at least two costly managerial challenges. First, it must repeatedly decide when to post new prices. Second, for each price update, it must choose what new price to post. Since both decisions are costly, managers may suffer errors or frictions along either margin. The most familiar models of nominal rigidity have studied frictions in the first decision, assuming that price adjustments can only occur intermittently, either with exogenous frequency (as in Calvo, 1983) or with endogenous frequency (e.g. Golosov and Lucas, 2007; Dotsey et al. 2009; Costain and Nakov 2011A, B). This paper instead explores the implications of frictions in the second decision. In other words, we assume that firms can adjust their prices costlessly in any period, but that whenever they adjust, their price choice is subject to errors. We study the implications of frictions along this margin both for microeconomic price adjustment data and for macroeconomic dynamics.

Our key assumption is that the probability of setting any given price is a smoothly increasing function of the expected present discounted value of setting that price. Our model nests full frictionless rationality as the limiting case in which the firm sets the optimal price with probability one every period. More specifically, we impose our main assumption in two slightly different ways. In the first version of our model, we assume the probability distribution over price choices is a logit function. The general equilibrium of this version of the model is therefore a logit equilibrium (McKelvey and Palfrey 1995, 1998): the probability of each firm's choice is a logit function which depends on the value of each choice; moreover, the value of each choice is determined, in equilibrium, by the logit choice probabilities of other firms.

While this version of the model does not impose price stickiness directly, the risk of errors gives rise to a certain degree of endogenous price stickiness. Because they fear they may “tremble” when choosing a new price, firms may refrain from adjusting, on precautionary grounds, even when their current price is not exactly optimal. Whenever the firm's price is sufficiently close to the optimum, it prefers to “leave well enough alone”, thus avoiding the risk of making a costly mistake. Hence behavior has an (S,s) band structure, in which adjustment occurs only if the current price is sufficiently far from the optimum.

In a second version of our model, we derive logit behavior from an assumption that precision is costly. We show that the distribution of price adjustments takes logit form if decision costs are proportional to entropy reduction. Thus, this specification involves a cost of price adjustment, but it is not a “menu cost” in the sense of labor time devoted to the physical task of adjusting posted prices. Instead, the cost of price adjustment can be understood as a cost of managerial control, consistent with the evidence of Zbaracki *et al.* (2004).

The logit framework for modeling bounded rationality has been widely applied in experimental game theory, where it has successfully explained play in a number of games where Nash equilibrium performs poorly, such as the centipede game and Bertrand competition games (McKelvey and Palfrey 1998; Anderson, Goeree, and Holt 2002). It has been much less frequently

¹For their helpful comments, we thank Fernando Álvarez, Jordi Galí, Kevin Lansing, John Leahy, Bartosz Mackowiak, Filip Matejka, Antonella Tutino, Mirko Wiederholt, Jonathan Willis, and seminar participants at the Bank of Spain, the San Francisco Fed, the ECB, HEC Paris, the Federal Reserve Board, SNDE 2010, Zeuthen Macroeconomics 2010, the 11th CeNDEF workshop, SED 2010, CEF 2010, the 2010 Econometric Society World Congress, and ASSA 2011. Views expressed here are those of the authors and do not necessarily coincide with those of the Bank of Spain, the European Central Bank, the Eurosystem, or the Federal Reserve Board.

applied in other areas of economics; we are unaware of any application of logit equilibrium inside a dynamic general equilibrium macroeconomic model.² The absence of logit modeling in macroeconomics may be due, in part, to discomfort with the many potential degrees of freedom opened up by moving away from the benchmark of full rationality. However, since logit equilibrium is just a one-parameter generalization of fully rational choice, it actually imposes much of the discipline of rationality on the model.

Another possible reason why macroeconomists have so rarely considered error-prone choice is that errors imply heterogeneity; the computational simplicity of a representative agent model may be lost if agents differ because of small, random mistakes. However, when applied to state-dependent pricing, this problem is less relevant, since it has long been argued that it is important to allow for heterogeneity in order to understand the dynamics of “sticky” adjustment models (see for example Caballero 1992, and Golosov and Lucas 2007). Moreover, we have shown (Costain and Nakov 2008B) how distributional dynamics can be tractably characterized in general equilibrium, without relying on special functional forms or questionable numerical aggregation assumptions. The same numerical method we used in that paper (Reiter 2009) can be applied to a logit equilibrium model; in fact, the smoothness of the logit case makes it even easier to compute than the fully rational case. We therefore find that logit equilibrium opens the door to tractable models with implications both for macroeconomic and for microeconomic data.

Summarizing our main findings, both versions of our model are consistent with several “puzzling” stylized facts from micro price adjustment data. Our model implies that many large and small price changes coexist (see Fig. 2), in contrast to the implications of the standard fixed menu cost model (Midrigan, 2011; Klenow and Kryvtsov, 2008; Klenow and Malin, 2009). It also implies that the probability of price adjustment decreases rapidly over the first few months, and then remains essentially flat (Nakamura and Steinsson, 2008; Klenow and Malin, 2009). Third, we find that the standard deviation of price changes is approximately constant, independent of the time since last adjustment (Klenow and Malin, 2009). The Calvo model implies instead that price changes are increasing in the time since last adjustment. Fourth, “extreme” prices are more likely to have been recently set than are prices near the center of the distribution (Campbell and Eden, 2010). While a variety of explanations have been offered for some of these observations (including sales, economies of scope in price setting, and heterogeneity among price setters), our framework matches all these facts in a very simple way, using only one degree of freedom in the parameterization.

Finally, we calculate the effects of money supply shocks in our framework. Given the degree of rationality that best fits microdata, the effect of money shocks on consumption is similar to that in the Golosov-Lucas (2007) fixed menu cost setup. The impact on consumption is much weaker than the Calvo model implies because of a selection effect: all the firms that require the largest price adjustments do in fact adjust. Thus, a model in which price adjustment is slowed down by the risk of mistakes fits microdata better than a fixed menu cost model, but implies that the macroeconomy is relatively close to monetary neutrality.

²The logit choice function is probably the most standard econometric framework for discrete choice, and has been applied to a huge number of microeconomic contexts. But logit *equilibrium*, in which each player makes logit decisions, based on payoff values which depend on other players’ logit decisions, has to the best of our knowledge rarely been applied outside of experimental game theory.

1.1 Related literature

Early sticky price frameworks based on “menu costs” were studied by Barro (1972), Sheshinski and Weiss (1977), and Mankiw (1985). General equilibrium solutions of these models have only been attempted more recently, at first by ignoring idiosyncratic shocks (Dotsey, King, and Wolman 1999), or by strongly restricting the distribution of such shocks (Danziger 1999; Gertler and Leahy 2006). Golosov and Lucas (2007) were the first to include frequent large idiosyncratic shocks in a quantitative model of state-dependent pricing, and approximately calculated the resulting equilibrium dynamics. However, while menu costs have been an influential idea in macroeconomics, the fact that price changes come in a wide variety of sizes, including some very small ones, is hard for the menu cost framework to explain.

In particular, Klenow and Kryvtsov (2008) have shown that the distribution of price changes remains puzzling even if we allow for many sectors with different menu costs. As a possible explanation for the presence of small adjustments, Lach and Tsiddon (2007) and Midrigan (2010) proposed economies of scope in the pricing of multiple goods: a firm that pays to correct one large price misalignment might get to change other, less misaligned, prices on the same menu costlessly. An extensive empirical literature has recently taken advantage of scanner data to document other microeconomic facts about retail price adjustment, many of which are puzzling when viewed through the lenses of the Calvo model or the menu cost model; references include Nakamura and Steinsson (2008), Klenow and Malin (2009), and Campbell and Eden (2010).

Rather than assuming a menu cost, our model delivers price stickiness as the result of near-rational behavior. In this it is similar in spirit to Akerlof and Yellen (1985), who assume firms sometimes make mistakes if they are not very costly. Our setup is also closely related to the “rational inattention” literature (e.g. Sims, 2003; Mackowiak and Wiederholt, 2009). Whereas that literature imposes an entropy constraint on the flow of information from the economic environment to the decision-maker, we instead use changes in entropy as a measure of the cost of precise decisions. A number of game-theoretic papers have modeled near-rational behavior in the same way, including Stahl (1990) and Mattsson and Weibull (2002).

Technically, the only difference between using entropy to measure the cost of precision and using entropy to measure information flow is that in the former case, the firm’s decisions are a function of its true state (its productivity and price) at time t , whereas in the latter the firm’s decisions depend on its prior about the state of the world. One possible interpretation of our environment, which features full information but imperfect decisions, is that the decision-maker (e.g. the CEO of a large corporation) is fully rational and has complete information about the environment, but acts subject to an implementation constraint which prevents him from perfectly communicating and enforcing his decisions throughout the organization. Given this imperfect implementation, the decisionmaker sometimes rationally prefers “not to call a meeting” and leave prices as they are, as long as the firm is doing reasonably well.

2 Sticky prices in partial equilibrium

In subsection 2.1, we describe the partial equilibrium decision of a monopolistically competitive firm that sometimes makes small errors when it adjusts its price. Concretely, we assume the price probabilities are governed by a multinomial logit. Subsection 2.2 discusses how our framework differs from the Calvo and menu cost approaches. In subsection 2.3, we take a more structural approach, and show how the multinomial logit can be derived from a cost function for error avoidance. We postpone discussion of general equilibrium until Section 3.

2.1 The monopolistic competitor's decision

Suppose that each firm i produces output Y_{it} under a constant returns technology, with labor N_{it} as the only input, and faces idiosyncratic productivity shocks A_{it} :

$$Y_{it} = A_{it}N_{it}$$

The idiosyncratic shocks A_{it} are given by a time-invariant Markov process, *iid* across firms. Thus A_{it} is correlated with $A_{i,t-1}$ but is uncorrelated with other firms' shocks. For numerical purposes, we assume A_{it} is drawn from a finite grid of possible values $\Gamma^a \equiv \{a^1, a^2, \dots, a^{\#a}\}$.³

Firms are monopolistic competitors, facing the demand curve $Y_{it} = \vartheta_t P_{it}^{-\epsilon}$, where ϑ_t represents aggregate demand. A firm's only control variable is its price; that is, we assume firms must fulfill all demand at the price they set. They hire in competitive labor markets at wage rate W_t , so period t profits are

$$P_{it}Y_{it} - W_tN_{it} = \left(P_{it} - \frac{W_t}{A_{it}}\right) Y_{it} = \left(P_{it} - \frac{W_t}{A_{it}}\right) \vartheta_t P_{it}^{-\epsilon}$$

Likewise, a firm that produces with price P_{it} and productivity A_{it} at time t has some discounted present value, which we write as $V_t(P_{it}, A_{it})$. The time subscript on the value function denotes all dependence on aggregate conditions, such as aggregate shocks or deterministic trends.⁴

At each point in time, a firm must decide whether or not to adjust its price. To make this decision, it compares the value of maintaining its previous price with the value of choosing a new one. A firm that begins period t with some initial price \tilde{P}_{it} receives value $V_t(\tilde{P}_{it}, A_{it})$ if it chooses not to adjust its price. If it instead chooses to adjust, it faces a risk of error in the new price it sets. Note that since we regard decisions as error-prone, the firm's decision process determines a *distribution* across its possible actions, rather than picking out a single optimal value. We assume a distribution such that the probability of choosing any given price is a smoothly increasing function of the value of choosing that price. This is the key assumption of our model.

As is common in microeconometrics and experimental game theory, we assume the distribution of errors is given by a multinomial logit. In order to treat the logit function as a primitive of the model, we define its argument in units of labor time. That is, since the costs of decision-making are presumably related to the labor effort (in particular, managerial labor) required to calculate and communicate the chosen price, we divide the values in the logit function by the wage rate, to convert them to time units. Hence, the probability $\pi_t(P^j|A_{it})$ of choosing price $P^j \in \Gamma^P$ at time t , conditional on productivity A_{it} , is given by

$$\pi_t(P^j|A_{it}) \equiv \frac{\exp\left(\frac{V_t(P^j, A_{it})}{\kappa W_t}\right)}{\sum_{k=1}^{\#P} \exp\left(\frac{V_t(P^k, A_{it})}{\kappa W_t}\right)} \quad (1)$$

³Theoretically, our model would be well-defined with a continuum of possible values of productivity A_{it} and also a continuum of possible prices P_{it} . However, our numerical solution method requires us to approximate the continuous case by a finite grid of possible productivities and prices. Therefore, for notational convenience, we define the model on a discrete grid from the start.

⁴We could instead write $V_t(P_{it}, A_{it})$ as $V(P_{it}, A_{it}, \Omega_t)$, where Ω_t represents the aggregate state of the economy. For more concise notation we just write V with a time subscript.

For numerical purposes, we constrain the price to a finite discrete grid $\Gamma^P \equiv \{P^1, P^2, \dots, P^{\#P}\}$. The parameter κ in the logit function can be interpreted as representing the degree of noise in the decision; in the limit as $\kappa \rightarrow 0$ it converges to the policy function under full rationality, in which the optimal price is chosen with probability one.⁵

We will use the notation E_t^π to indicate an expectation taken under the logit probability (1). The firm's expected value, conditional on adjusting to a new price $P' \in \Gamma^P$, is then

$$E_t^\pi V_t(P', A_{it}) \equiv \sum_{j=1}^{\#P} \pi_t(P^j | A_{it}) V(P^j, A_{it}) = \sum_{j=1}^{\#P} \frac{\exp\left(\frac{V_t(P^j, A_{it})}{\kappa W_t}\right) V(P^j, A_{it})}{\sum_{k=1}^{\#P} \exp\left(\frac{V_t(P^k, A_{it})}{\kappa W_t}\right)} \quad (2)$$

The expected value of adjustment is

$$D_t(P_{it}, A_{it}) \equiv E_t^\pi V_t(P', A_{it}) - V_t(P_{it}, A_{it}) \quad (3)$$

We assume the firm adjusts its price if and only if the expected gain from adjustment is non-negative. That is, the probability of adjustment can be written as

$$\lambda(D_t(P_{it}, A_{it})) = \mathbf{1}(D_t(P_{it}, A_{it}) \geq 0) \quad (4)$$

where $\mathbf{1}(x)$ is an indicator function taking the value 1 if statement x is true, and zero otherwise.

We can now state the Bellman equation that governs a firm's value of producing at any given price P . The Bellman equation in this case is:

Bellman equation in partial equilibrium:

$$V_t(P, A) = \left(P - \frac{W_t}{A}\right) \vartheta_t P^{-\epsilon} + E_t \left\{ Q_{t,t+1} \left[(1 - \lambda(D_{t+1}(P, A'))) V_{t+1}(P, A') + \lambda(D_{t+1}(P, A')) E_t^\pi V_{t+1}(P', A') \right] \middle| A \right\} \quad (5)$$

where $Q_{t,t+1}$ is the firm's stochastic discount factor. Note that the aggregate price level is absent from the above expression; it is subsumed into ϑ_t , as we show in Section 3. On the left-hand side and in the current profits term, P refers to a given firm i 's price P_{it} at the time of production. In the expectation on the right, P represents the price $\tilde{P}_{i,t+1}$ at the beginning of period $t+1$, which is the same as P_{it} , and subsequently may (with probability λ) or may not (with probability $1 - \lambda$) be adjusted prior to time $t+1$ production.

We can simplify substantially by noticing that the value on the right-hand side of the equation is just the value of continuing without adjustment, plus the expected gains from adjustment, which we call G :

$$V_t(P, A) = \left(P - \frac{W_t}{A}\right) \vartheta_t P^{-\epsilon} + E_t \left\{ Q_{t,t+1} [V_{t+1}(P, A') + G_{t+1}(P, A')] \middle| A \right\} \quad (6)$$

where

$$G_{t+1}(P, A') \equiv \lambda(D_{t+1}(P, A')) D_{t+1}(P, A') = \mathbf{1}(D_{t+1}(P, A') \geq 0) D_{t+1}(P, A') \quad (7)$$

⁵Alternatively, logit models are often written in terms of the inverse parameter $\xi \equiv \kappa^{-1}$, which can be interpreted as a measure of rationality rather than a measure of noise.



2.2 Alternative models of sticky prices

To better interpret our results, we will compare simulations of our framework with simulations of two standard models of nominal rigidity: the Calvo model and the fixed menu cost (FMC) model. Both these models are consistent with Bellman equation (6) if we redefine the expected gains function G appropriately.

In the Calvo model, adjustment occurs with a constant, exogenous probability $\bar{\lambda}$, and conditional on adjustment, the firm sets the optimal price. This means $\lambda(D_{t+1}(P, A')) = \bar{\lambda}$ and $D_{t+1}(P, A') = V_{t+1}^*(A') - V_{t+1}(P, A')$, where

$$V_{t+1}^*(A') = \max_{P'} V_{t+1}(P', A'). \quad (8)$$

Therefore (7) is replaced by

$$G_{t+1}(P, A') \equiv \bar{\lambda} (V_{t+1}^*(A') - V_{t+1}(P, A')). \quad (9)$$

In the FMC model, the firm adjusts if and only if the gains from adjustment are at least as large as the menu cost α , which is a fixed, exogenous quantity of labor. If the firm adjusts, it pays the menu cost and sets the optimal price. So the probability of adjustment is $\lambda(D_{t+1}(P, A')) = \mathbf{1}(D_{t+1}(P, A') \geq \alpha W_{t+1})$, where $D_{t+1}(P, A') = V_{t+1}^*(A') - V_{t+1}(P, A')$. Therefore (7) is replaced by

$$G_{t+1}(P, A') \equiv \mathbf{1}(D_{t+1}(P, A') \geq \alpha W_{t+1}) (D_{t+1}(P, A') - \alpha W_{t+1}). \quad (10)$$

2.3 Deriving logit choice from a control cost function

The logit assumption (1) has the desirable property that the probability of choosing any given price is a smoothly increasing function of the value of that price. However, many other distributions have the same property. Is there any good reason for assuming a logit, other than its prominence in game theory and econometrics?

One way to derive a logit error distribution is to assume that managerial decision-making is costly. Thus, let us suppose that greater precision in the price choice (equivalently, a decrease in errors) requires greater managerial time. In particular, following Stahl (1990) and Mattsson and Weibull (2002), we will assume that the time cost of more precise choice is proportional to the reduction in the entropy of the choice variable, normalizing the cost of perfectly random decisions (a uniform distribution) to zero.⁶ Then the cost of choosing a price distribution $\vec{\pi} \equiv \{\pi^j\}_{j=1}^{\#P}$ is given by⁷

$$C(\vec{\pi}) = \kappa \left(\ln(\#P) + \sum_{j=1}^{\#P} \pi^j \ln(\pi^j) \right) \quad (11)$$

Here κ is the marginal cost of entropy reduction in units of labor time. This cost function is nonnegative and convex. It takes its maximum value, $\kappa \ln(\#P) > 0$, for any distribution that

⁶See also Marsili (1999), Baron *et al.* (2002), or Matejka and McKay (2011).

⁷Our normalization of the level of costs is equivalent to defining the cost function in terms of *relative entropy* (also known as Kullback-Leibler divergence). That is, by setting the cost of choosing a uniform distribution to zero, we are making the cost function proportional to the Kullback-Leibler divergence $D(p||q)$ between the price variable p and a random variable q that is uniform over the support of p . See Mattsson and Weibull (2002) for details.

places all probability on a single price $p \in \Gamma^P$. It takes its minimum value, zero, for a uniform distribution.⁸

While this version of our framework involves a cost of price adjustment, it should not be interpreted as the “menu cost” of the physical task of altering the posted price, but rather as a cost of managerial effort. We now show how to derive a logit distribution from this cost function. Consider a firm that has chosen to update its price at time t , and is now deciding which new price P^j to set, on the finite grid $\Gamma^P \equiv \{P^1, P^2, \dots, P^{\#P}\}$. The optimal price distribution $\vec{\pi}$ maximizes firm value, net of computational costs (which we convert to nominal terms by multiplying by the wage):

$$\max_{\pi^j} \sum_j \pi^j V_t(P^j, A) - \kappa W_t \left(\ln(\#P) + \sum_{j=1}^{\#P} \pi^j \ln(\pi^j) \right) \quad \text{s.t.} \quad \sum_j \pi^j = 1 \quad (12)$$

The first-order condition for π^j is

$$V^j - \kappa W_t (1 + \ln \pi^j) - \mu = 0,$$

where μ is the multiplier on the constraint. Some rearrangement yields:

$$\pi^j = \exp \left(\frac{V^j}{\kappa W_t} - 1 - \frac{\mu}{\kappa W_t} \right). \quad (13)$$

Since the probabilities sum to one, we have $\exp \left(1 + \frac{\mu}{\kappa W_t} \right) = \sum_j \exp \left(\frac{V^j}{\kappa W_t} \right)$. Therefore the optimal probabilities (13) reduce to the logit formula (1). Thus the noise parameter in the logit function corresponds to the marginal cost of entropy reduction in the control cost problem.

Taking logs in the first-order condition (13), we can calculate a simple analytical formula for the value of problem (12). Using $\ln \pi^j = \frac{V^j}{\kappa W_t} - \text{constant}$, the value of (12) equals:

$$\kappa W_t \ln \left(\frac{1}{\#P} \sum_{j=1}^{\#P} \exp \left(\frac{V_t(P^j, A)}{\kappa W_t} \right) \right)$$

Of course, if we interpret the logit choice distribution as the result of costly managerial time, these costs need to be subtracted out of the value function. The description of the firm’s problem in subsection 2.1 remains valid, except for the expected value of adjustment. Equation (3) is replaced by

$$D_t(P_{it}, A_{it}) \equiv E_t^\pi V_t(P', A_{it}) - W_t C(\vec{\pi}) - V_t(P_{it}, A_{it}) \quad (14)$$

$$= \kappa W_t \ln \left(\frac{1}{\#P} \sum_{j=1}^{\#P} \exp \left(\frac{V_t(P^j, A)}{\kappa W_t} \right) \right) - V_t(P_{it}, A_{it}). \quad (15)$$

The Bellman equation is then given by (6)-(7) as before.

Therefore, in our simulations below, we will report two specifications of our model. One specification, abbreviated as ‘PPS’, is defined by (3), (6), and (7), so that logit choice is simply interpreted as a decision process near to, but varying around, the optimal choice. The second specification, indicated by ‘ENT’, is defined by (14), (6), and (7), so logit choice is interpreted as the result of optimal decision-making constrained by entropy reduction costs.

⁸If π is uniform, then $\pi(p) = 1/\#P$ for all $p \in \Gamma^P$, which implies $\sum_{j \in \Gamma^P} \pi(p) \ln(\pi(p)) = -\ln(\#P)$.

2.4 Relation to information-theoretic models

Before leaving this section we discuss the relationship between our framework and the “rational inattention” papers like Sims (2003), Woodford (2009), and more recently Matejka and McKay (2011), that are explicitly based on information theory. Note that by imposing a cost function on error reduction we are emphasizing the managerial costs of evaluating information, choosing an optimal policy, and communicating it within the firm, but by treating (12) as a decision under full information we are ignoring the costs of *receiving* information. The information-theoretic approach instead takes account of all four of these stages of information processing.

Thus, our setup could be interpreted as the problem of a manager who costlessly absorbs the content of all the world’s newspapers (and other data sources) with her morning coffee, but who faces costs of *using* that information. The information-theoretic approach allows for costs of receiving the morning’s information too. Thus while the information-theoretic approach imposes a Shannon entropy constraint on the whole flow of information from reception to communicating choices, our approach imposes a Shannon entropy constraint only on the flow of information inside the manager’s head and then within her firm.

Intuitively, we believe all these stages of information flow imply important costs. Our only reason for ignoring the first stage of the information flow is that by doing so we dramatically reduce the dimension of the calculations required to solve our model. Since the rational inattention approach assumes the firm acts under uncertainty, it implies the firm conditions on a *prior* over its possible productivity levels (which is a high-dimensional object that complicates solution of the model). In our setup, the firm just conditions on its true productivity level. Moreover, once one knows that entropy reduction costs imply logit, one can simply impose a logit function directly (and then subtract off the implied costs) rather than explicitly solving for the form of the error distribution. These facts make our approach entirely tractable in a DSGE context, as we show in the next two sections.

Given the similarity between our approach and that of the rational inattention literature, it is likely that the two approaches will have many similar implications. Rational inattention may have policy-relevant empirical implications which our model does not capture; this is a relevant question for future research. But if the implications of the two approaches turn out to be essentially the same, our setup may be preferred for its greater tractability.

3 General equilibrium

We next embed this partial equilibrium framework into a dynamic New Keynesian general equilibrium model. For comparability, we use the same structure as Golosov and Lucas (2007). Besides the firms, there is a representative household and a central bank that sets the money supply.

3.1 Households

The household’s period utility function is

$$u(C_t) - x(N_t) + v(M_t/P_t);$$

payoffs are discounted by factor β per period. Consumption C_t is a Spence-Dixit-Stiglitz aggregate of differentiated products:

$$C_t = \left[\int_0^1 C_{it}^{\frac{\epsilon-1}{\epsilon}} di \right]^{\frac{\epsilon}{\epsilon-1}}.$$

N_t is labor supply, and M_t/P_t is real money balances. The household's period budget constraint is

$$\int_0^1 P_{it} C_{it} di + M_t + R_t^{-1} B_t = W_t N_t + M_{t-1} + T_t + B_{t-1} + Div_t,$$

where $\int_0^1 P_{it} C_{it} di$ is total nominal spending on the differentiated goods. B_t represents nominal bond holdings, with interest rate $R_t - 1$; T_t represents lump sum transfers received from the monetary authority, and Div_t represents dividend payments received from the firms. In this context, optimal allocation of consumption across the differentiated goods implies $C_{it} = (P_t/P_{it})^\epsilon C_t$, where P_t is the price index $P_t \equiv \left\{ \int_0^1 P_{it}^{1-\epsilon} di \right\}^{\frac{1}{1-\epsilon}}$.

3.2 Monetary policy and aggregate consistency

For simplicity, we assume the central bank follows an exogenous stochastic money growth rule:

$$M_t = \mu_t M_{t-1}, \quad (16)$$

where $\mu_t = \mu \exp(z_t)$, and z_t is AR(1):

$$z_t = \phi_z z_{t-1} + \epsilon_t^z. \quad (17)$$

Here $0 \leq \phi_z < 1$ and $\epsilon_t^z \sim i.i.d.N(0, \sigma_z^2)$ is a money growth shock. Thus the money supply trends upward by approximately factor $\mu \geq 1$ per period on average. Seigniorage revenues are paid to the household as a lump sum transfer, and the public budget is balanced each period. Therefore the public budget constraint is

$$M_t = M_{t-1} + T_t.$$

Bond market clearing is simply $B_t = 0$. Market clearing for good i implies the following demand and supply relations for firm i :

$$Y_{it} = A_{it} N_{it} = C_{it} = P_t^\epsilon C_t P_{it}^{-\epsilon}. \quad (18)$$

Also, total labor supply must equal total labor demand:

$$N_t = \int_0^1 \frac{C_{it}}{A_{it}} di = P_t^\epsilon C_t \int_0^1 P_{it}^{-\epsilon} a_{it}^{-1} di \equiv \Delta_t C_t. \quad (19)$$

The labor market clearing condition (19) also defines a weighted measure of price dispersion, $\Delta_t \equiv P_t^\epsilon \int_0^1 P_{it}^{-\epsilon} a_{it}^{-1} di$. which generalizes the dispersion measure in Yun (2005) to allow for heterogeneous productivity. An increase in Δ_t decreases the consumption goods produced per unit of labor, effectively acting like a negative shock to aggregate productivity.

Aggregate consistency also requires that the demand curve and the discount factor that appear in the firm's problem be consistent with the household's problem. In particular, to make

the firm's problem (6) consistent with the goods market clearing conditions (18), the aggregate demand shift factor must be

$$\vartheta_t = C_t P_t^\epsilon. \quad (20)$$

Also, we assume that the representative household owns the firms, so the stochastic discount factor in the firm's problem must be consistent with the household's Euler equation. This implies that the appropriate stochastic discount factor is

$$Q_{t,t+1} = \beta \frac{P_t u'(C_{t+1})}{P_{t+1} u'(C_t)}. \quad (21)$$

To write the firm's problem in general equilibrium, we simply plug (20) and (21) into the firm's problem (6). Then the value of producing with price $P_{it} = P$ and productivity $A_{it} = A$ is

Bellman equation in general equilibrium:

$$V_t(P, A) = \left(P - \frac{W_t}{A} \right) C_t P_t^\epsilon P^{-\epsilon} + \beta E_t \left\{ \frac{P_t u'(C_{t+1})}{P_{t+1} u'(C_t)} [V_{t+1}(P, A') + G_{t+1}(P, A')] \middle| A \right\}, \quad (22)$$

where $G_{t+1}(P, A')$ has the form described in equation (7). G in turn depends on the adjustment gain D , which takes the form (3) if we impose a logit choice distribution directly, or the form (14) if we derive logit choice by imposing entropy control costs.

3.3 State variable

At this point, we have spelled out all equilibrium conditions: household and monetary authority behavior has been described in this section, and the firms' decision was stated in Section 2. Thus can now identify the aggregate state variable Ω_t . Aggregate uncertainty in the model relates only to the money supply M_t . But since the growth rate of M_t is $AR(1)$ over time, the latest deviation in growth rates, z_t , is a state variable too. There is also a continuum of idiosyncratic productivity shocks A_{it} , $i \in [0, 1]$. Finally, since firms cannot instantly adjust their prices, they are state variables too. More precisely, the state includes the joint distribution of prices and productivity shocks at the beginning of the period, prior to adjustment.

We will use the notation \tilde{P}_{it} to refer to firm i 's price at the beginning of period t , prior to adjustment; this may of course differ from the price P_{it} at which it produces, because the price may be adjusted before production. Therefore we will distinguish the distribution of production prices and productivity at the time of production, which we write as $\Phi_t(P_{it}, A_{it})$, from the distribution of beginning-of-period prices and productivity, $\tilde{\Phi}_t(\tilde{P}_{it}, A_{it})$. Since beginning-of-period prices and productivities determine all equilibrium decisions at t , we can define the state at time t as $\Omega_t \equiv (M_t, z_t, \tilde{\Phi}_t)$.

It is helpful here to compare the dimension of the general equilibrium calculation our model requires with the calculation implied by a rational inattention model like Sims (2003). In our model, the firm's idiosyncratic state has two dimensions (P and A), and the value function is also contingent on the aggregate state Ω . Since the possible values of prices and productivities are constrained to a grid, the distribution $\tilde{\Phi}$ is an object of dimension $\#P\#a$, and thus Ω has dimension $\#P\#a + 2$. In contrast, if we were to attempt to compute a rational inattention model on the same grid, the firm's idiosyncratic state would have at least dimension $\#a$: one dimension for its price P and $\#a - 1$ dimensions for its prior over possible values of A . The true aggregate state of the economy Ω would then include at least M_t , z_t , and a *distribution*

over prices and priors. Moreover, in general equilibrium firms might hold nondegenerate priors over Ω itself, which would blow up the dimension of the problem exponentially again. In practice, therefore, some rational inattention models, like Mackowiak and Wiederholt (2011), have restricted attention to a linear-quadratic-Gaussian framework where all distributions collapse to a known form, whereas papers that have allowed for non-Gaussian shocks, like Tutino (2009) and Matejka (2010), have studied partial equilibrium problems only.

3.4 Detrending

So far we have written the value function and all prices in nominal terms, but we can rewrite the model in real terms by deflating all prices by the nominal price level $P_t \equiv \left\{ \int_0^1 P_{it}^{1-\epsilon} di \right\}^{\frac{1}{1-\epsilon}}$. Thus, define $m_t \equiv M_t/P_t$ and $w_t \equiv W_t/P_t$. Given the nominal distribution $\Phi_t(P_{it}, A_{it})$, let us denote by $\Psi_t(p_{it}, A_{it})$ the distribution over real transaction prices $p_{it} \equiv P_{it}/P_t$. Rewriting the definition of the price index in terms of these deflated prices, we have the following restriction:

$$\int_0^1 p_{it}^{1-\epsilon} di = 1.$$

Notice however that the beginning-of-period *real* price is not predetermined: if we define $\tilde{p}_{it} \equiv \tilde{P}_{it}/P_t$, then \tilde{p}_{it} is a jump variable, and so is the distribution of real beginning-of-period prices $\tilde{\Psi}_t(\tilde{p}_i, A_i)$. Therefore we cannot define the real state of the economy at the beginning of t in terms of the distribution $\tilde{\Psi}_t$.

To write the model in real terms, the level of the money supply, M_t , and the aggregate price level, P_t , must be irrelevant for determining real quantities; and we must condition on a real state variable that is predetermined at the beginning of period. Therefore, we define the real state at time t as $\Xi_t \equiv (z_t, \Psi_{t-1})$, where Ψ_{t-1} is the distribution of *lagged* prices and productivities. Note that the distribution Ψ_{t-1} , together with the shocks z_t , is sufficient to determine all equilibrium quantities at time t : in particular, it will determine the distributions $\tilde{\Psi}_t(\tilde{p}_i, A_i)$ and $\Psi_t(p_i, A_i)$. Therefore Ξ_t is a correct time t real state variable.

This also makes it possible to define a real value function v , meaning the nominal value function, divided by the current price level, depending on real variables only. That is,

$$V_t(P_{it}, A_{it}) = V(P_{it}, A_{it}, \Omega_t) = P_t v \left(\frac{P_{it}}{P_t}, A_{it}, \Xi_t \right) = P_t v_t(p_{it}, A_{it}).$$

Deflating in this way, the Bellman equation can be rewritten as follows:

Detrended Bellman equation, general equilibrium:

$$v_t(p, A) = \left(p - \frac{w_t}{A} \right) C_t p^{-\epsilon} + \beta E_t \left\{ \frac{u'(C_{t+1})}{u'(C_t)} [v_{t+1}(\pi_{t+1}^{-1} p, A') + g_{t+1}(\pi_{t+1}^{-1} p, A')] \middle| A \right\}, \quad (23)$$

where

$$\begin{aligned} g_{t+1}(\pi_{t+1}^{-1} p, A') &\equiv \lambda (w_{t+1}^{-1} d_{t+1}(\pi_{t+1}^{-1} p, A')) d_{t+1}(\pi_{t+1}^{-1} p, A'), \\ d_{t+1}(\pi_{t+1}^{-1} p, A') &\equiv E_{t+1}^{\pi} v_{t+1}(p', A') - v_{t+1}(\pi_{t+1}^{-1} p, A'). \end{aligned}$$

4 Computation

4.1 Outline of algorithm

Computing this model is challenging due to heterogeneity: at any time t , firms will face different idiosyncratic shocks A_{it} and will be stuck at different prices P_{it} . The reason for the popularity of the Calvo model is that even though firms have many different prices, up to a first-order approximation only the average price matters for equilibrium. Unfortunately, this property does not hold in general, and in the current context, we need to treat all equilibrium quantities explicitly as functions of the distribution of prices and productivity across the economy, and we must calculate the dynamics of this distribution over time.

We address this problem by implementing Reiter's (2009) solution method for dynamic general equilibrium models with heterogeneous agents and aggregate shocks. As a first step, Reiter's algorithm calculates the steady state general equilibrium that obtains in the absence of aggregate shocks. Idiosyncratic shocks are still active, but are assumed to have converged to their ergodic distribution, so an aggregate steady state means that $z = 0$, and Ψ , π , C , R , N , and w are all constant. To solve for this steady state, we will assume that real prices and productivities always lie on a fixed grid $\Gamma \equiv \Gamma^p \times \Gamma^a$, where $\Gamma^p \equiv \{p^1, p^2, \dots, p^{\#^p}\}$ and $\Gamma^a \equiv \{a^1, a^2, \dots, a^{\#^a}\}$ are logarithmically-spaced grids of possible values of p_{it} and A_{it} , respectively. We can then think of the steady state value function as a matrix \mathbf{V} of size $\#^p \times \#^a$ comprising the values $v^{jk} \equiv v(p^j, a^k)$ associated with the prices and productivities $(p^j, a^k) \in \Gamma$. Likewise, the price distribution can be viewed as a $\#^p \times \#^a$ matrix Ψ in which the row j , column k element Ψ^{jk} represents the fraction of firms in state (p^j, a^k) at the time of transactions. Given this discretized representation, we can calculate steady state general equilibrium by guessing the aggregate wage level, then solving the firm's problem by backwards induction on the grid Γ , then updating the conjectured wage, and iterating to convergence.

In a second step, Reiter's method constructs a linear approximation to the dynamics of the discretized model, by perturbing it around the steady state general equilibrium on a point-by-point basis. The method recognizes that the Bellman equation and the distributional dynamics can be interpreted as a large system of nonlinear first-order autonomous difference equations that define the aggregate dynamics. For example, away from steady state, the Bellman equation relates the $\#^p \times \#^a$ matrices \mathbf{V}_t and \mathbf{V}_{t+1} that represent the value function at times t and $t+1$. The row j , column k element of \mathbf{V}_t is $v_t^{jk} \equiv v_t(p^j, a^k) \equiv v(p^j, a^k, \Xi_t)$, for $(p^j, a^k) \in \Gamma$. Given this representation, we no longer need to think of the Bellman equation as a functional equation that defines $v(p, a, \Xi)$ for all possible idiosyncratic and aggregate states p , a , and Ξ ; instead, we simply treat it as a system of $\#^p \#^a$ expectational difference equations that determine the dynamics of the $\#^p \#^a$ variables v_t^{jk} . We linearize this large system of difference equations numerically, and then solve for the saddle-path stable solution of our linearized model using the QZ decomposition, following Klein (2000).

The beauty of Reiter's method is that it combines linearity and nonlinearity in a way appropriate for the model at hand. In the context of price setting, aggregate shocks are likely to be less relevant for individual firms' decisions than idiosyncratic shocks; Klenow and Kryvtsov (2008), Golosov and Lucas (2007), and Midrigan (2008) all argue that firms' prices are driven primarily by idiosyncratic shocks. To deal with these big firm-specific shocks, we treat functions of idiosyncratic states in a fully nonlinear way, by calculating them on a grid. But this grid-based solution can also be regarded as a large system of nonlinear equations, with equations

specific to each of the grid points. When we linearize each of these equations with respect to the aggregate dynamics, we recognize that aggregate changes are unlikely to affect individual value functions in a strongly nonlinear way. That is, we are implicitly assuming that aggregate shocks z_t and changes in the distribution Ψ_t have sufficiently smooth impacts on individual values that a linear treatment of these effects suffices. On the other hand, we need *not* start from any assumption of approximate aggregation like that required for the Krusell and Smith (1998) method, nor do we need to impose any particular functional form on the distribution Ψ .

Describing the distributional dynamics involves defining various matrices related to quantities on the grid Γ . From here on, we use bold face to identify matrices, and superscripts to identify notation related to grids. Matrices associated with the grid Γ are defined so that row j relates to the price $p^j \in \Gamma^p$, and column k relates to the productivity $a^k \in \Gamma^a$. Besides the value function matrix \mathbf{V}_t , we also define matrices \mathbf{D}_t , \mathbf{G}_t , and $\mathbf{\Lambda}_t$, to represent the functions d_t , g_t , and $\lambda(d_t/w_t)$ at points on the grid Γ . The distribution at the time of transactions is given by Ψ_t , with elements Ψ_t^{jk} representing the fraction of firms with real price $p_{it} \equiv P_{it}/P_t = p^j$ and productivity $A_{it} = a^k$ at the time of transactions. We also define the beginning-of-period distribution $\tilde{\Psi}_t$, with elements $\tilde{\Psi}_t^{jk}$ representing the fraction of firms with real price $\tilde{p}_{it} \equiv \tilde{P}_{it}/P_t = p^j$ and productivity $A_{it} = a^k$ at the beginning of the period. Shortly we will define the transition matrices that govern the relationships between all these objects.

4.2 The discretized model

In the discretized model, the value function \mathbf{V}_t is a matrix of size $\#^p \times \#^a$ with elements $v_t^{jk} \equiv v_t(p^j, a^k) \equiv v(p^j, a^k, \Xi_t)$ for $(p^j, a^k) \in \Gamma$. Other relevant $\#^p \times \#^a$ matrices include the adjustment values \mathbf{D}_t , the adjustment probabilities $\mathbf{\Lambda}_t$, and the expected gains \mathbf{G}_t , with (j, k) elements given by

$$d_t^{jk} \equiv d_t(p^j, a^k) \equiv E_t^\pi v_t(p, a^k) - v_t(p^j, a^k), \quad (24)$$

$$\lambda_t^{jk} \equiv \lambda(d_t^{jk}/w_t), \quad (25)$$

$$g_t^{jk} \equiv \lambda_t^{jk} d_t^{jk}. \quad (26)$$

Finally, we also define a matrix of logit probabilities $\mathbf{\Pi}_t$, which has its (j, k) element given by

$$\pi_t^{jk} = \pi_t(p^j | a^k) \equiv \frac{\exp(v_t^{jk}/(\kappa w_t))}{\sum_{n=1}^{\#^p} \exp(v_t^{jn}/(\kappa w_t))},$$

which is the probability of choosing real price p^j conditional on productivity a^k if the firm decides to adjust its price at time t .

We can now write the discrete Bellman equation and the discrete distributional dynamics in a precise way. First, consider how the beginning-of-period distribution $\tilde{\Psi}_t$ is derived from the lagged distribution Ψ_{t-1} . Idiosyncratic productivities A_i are driven by an exogenous Markov process, which can be defined in terms of a matrix \mathbf{S} of size $\#^a \times \#^a$. The row m , column k element of \mathbf{S} represents the probability

$$S^{mk} = \text{prob}(A_{it} = a^m | A_{i,t-1} = a^k).$$

Also, beginning-of-period real prices are, by definition, adjusted for inflation. Ignoring grids, the time $t-1$ real price $p_{i,t-1}$ would be deflated to $\tilde{p}_{it} \equiv p_{i,t-1}/\pi_t \equiv p_{i,t-1}P_{t-1}/P_t$ at the beginning

of t . To keep prices on the grid, we define a $\#^P \times \#^P$ Markov matrix \mathbf{R}_t in which the row m , column l element is

$$R_t^{ml} \equiv \text{prob}(\tilde{p}_{it} = p^m | p_{i,t-1} = p^l).$$

When the deflated price $p_{i,t-1}/\pi_t$ falls between two grid points, matrix \mathbf{R}_t must round up or down stochastically. Also, if $p_{i,t-1}/\pi_t$ lies outside the smallest and largest element of the grid, then \mathbf{R}_t must round up or down to keep prices on the grid.⁹ Therefore we construct \mathbf{R}_t according to

$$R_t^{ml} = \text{prob}(\tilde{p}_{it} = p^m | p_{i,t-1} = p^l, \pi_t) = \begin{cases} 1 & \text{if } \pi_t^{-1} p^l \leq p^1 = p^m \\ \frac{\pi_t^{-1} p^l - p^{n-1}}{p^n - p^{n-1}} & \text{if } p^1 < p^m = \min\{p \in \Gamma^p : p \geq \pi_t^{-1} p^l\} \\ \frac{p^{n+1} - \pi_t^{-1} p^l}{p^{n+1} - p^n} & \text{if } p^1 \leq p^m = \max\{p \in \Gamma^p : p < \pi_t^{-1} p^l\} \\ 1 & \text{if } \pi_t^{-1} p^l > p^{\#^p} = p^m \\ 0 & \text{otherwise} \end{cases} \quad (27)$$

Combining the adjustments of prices and productivities, we can calculate the beginning-of-period distribution $\tilde{\Psi}_t$ as a function of the lagged distribution of transaction prices Ψ_{t-1} :

$$\tilde{\Psi}_t = \mathbf{R}_t * \Psi_{t-1} * \mathbf{S}',$$

where $*$ represents ordinary matrix multiplication. The simplicity of this equation comes partly from the fact that the exogenous shocks to A_{it} are independent of the inflation adjustment that links \tilde{p}_{it} with $p_{i,t-1}$. Also, exogenous shocks are represented from left to right in the matrix Ψ_t , so that their transitions can be treated by right multiplication, while policies are represented vertically, so that transitions related to policies can be treated by left multiplication.

To calculate the effects of price adjustment on the distribution, let \mathbf{E}_{pp} and \mathbf{E}_{pa} be matrices of ones of size $\#P \times \#P$ and $\#P \times \#a$, respectively. Now suppose a firm has beginning-of- t price $\tilde{p}_{it} \equiv \tilde{P}_{it}/P_t = p^j \in \Gamma^p$ and productivity $A_{it} = a^k \in \Gamma^a$. This firm will adjust its production price with probability λ_t^{jk} , or will leave it unchanged ($p_{it} = \tilde{p}_{it} = p^j$) with probability $1 - \lambda_t^{jk}$. If adjustment occurs, the probabilities of choosing all possible prices are given by the matrix $\mathbf{\Pi}_t$. Therefore we can calculate distribution Ψ_t from $\tilde{\Psi}_t$ as follows:

$$\Psi_t = (\mathbf{E}_{pa} - \mathbf{\Lambda}_t) .* \tilde{\Psi}_t + \mathbf{\Pi}_t .* (\mathbf{E}_{pp} * (\mathbf{\Lambda}_t .* \tilde{\Psi}_t)), \quad (28)$$

where (as in MATLAB) the operator $.*$ represents element-by-element multiplication, and $*$ represents ordinary matrix multiplication.

The same transition matrices \mathbf{R} and \mathbf{S} show up when we write the Bellman equation in matrix form. Let \mathbf{U}_t be the $\#^P \times \#^a$ matrix of current payoffs, with elements

$$u_t^{jk} \equiv \left(p^j - \frac{w_t}{a^k}\right) (p^j)^{-\epsilon} C_t, \quad (29)$$

for $(p^j, a^k) \in \Gamma$. Then the Bellman equation is

⁹In other words, we assume that any nominal price that would have a real value less than p^1 after inflation is automatically adjusted upwards so that its real value is p^1 . This assumption is made for numerical purposes only, and has a negligible impact on the equilibrium as long as we choose a sufficiently wide grid Γ^p . If we were to compute examples with trend deflation, we would need to make an analogous adjustment to prevent real prices from exceeding the maximum grid point $p^{\#^p}$.

Dynamic general equilibrium Bellman equation, matrix version:

$$\mathbf{V}_t = \mathbf{U}_t + \beta E_t \left\{ \frac{u'(C_{t+1})}{u'(C_t)} [\mathbf{R}'_{t+1} * (\mathbf{V}_{t+1} + \mathbf{G}_{t+1}) * \mathbf{S}] \right\}. \quad (30)$$

The expectation E_t in the Bellman equation refers only to the effects of the time $t + 1$ aggregate shock z_{t+1} , because the shocks and dynamics of the idiosyncratic state $(p^j, a^k) \in \Gamma$ are completely described by the matrices \mathbf{R}'_{t+1} and \mathbf{S} . Note that since the Bellman equation iterates backwards in time, its transitions are represented by \mathbf{R}' and \mathbf{S} , whereas the distributional dynamics iterate forward in time and therefore contain \mathbf{R} and \mathbf{S}' .

While equilibrium seems to involve a complicated system of equations, the steady state is easy to solve because it reduces to a small scalar fixed-point problem, which is the first step of Reiter's (2009) method, to be discussed in the next subsection. The second step of the method, in which we linearize all equilibrium equations, is discussed in subsection 4.4.

4.3 Step 1: steady state

In the aggregate steady state, the shocks are zero, and the distribution takes some unchanging value Ψ , so the state of the economy is constant: $\Xi_t \equiv (z_t, \Psi_{t-1}) = (0, \Psi) \equiv \Xi$. We indicate the steady state of all equilibrium objects by dropping the time subscript t , so the steady state value function \mathbf{V} has elements $v^{jk} \equiv v(p^j, a^k, \Xi) \equiv v(p^j, a^k)$.

Long run monetary neutrality in steady state implies that the rate of nominal money growth equals the rate of inflation:

$$\mu = \pi.$$

Moreover, the Euler equation reduces to

$$\pi = \beta R.$$

Since the interest rate and inflation rate are observable, together they determine the required parameterization of β . The steady-state transition matrix \mathbf{R} is known, since it depends only on steady state inflation π .

We can then calculate general equilibrium as a one-dimensional root-finding problem: guessing the wage w , we have enough information to solve the Bellman equation and the distributional dynamics.¹⁰ Knowing the steady state aggregate distribution, we can construct the real price level, which must be one. Thus finding a value of w at which the real price level is one amounts to finding a steady state general equilibrium.

More precisely, for any w , we can calculate

$$C = \left(\frac{\chi}{w} \right)^{1/\gamma}, \quad (31)$$

and then construct the matrix \mathbf{U} with elements

$$u^{jk} \equiv \left(p^j - \frac{w}{a^k} \right) (p^j)^{-\epsilon} C. \quad (32)$$

¹⁰There are other, equivalent ways of describing the root-finding problem: for example, we could begin by guessing C . Guessing w is convenient since we know that in a representative-agent, flexible-price model, we have $w = \frac{\epsilon-1}{\epsilon}$. This suggests a good starting value for the heterogeneous-agent, sticky-price calculation.

We then find the fixed point of the value function:

$$\mathbf{V} = \mathbf{U} + \beta \mathbf{R}' * (\mathbf{V} + \mathbf{G}) * \mathbf{S}, \quad (33)$$

together with the logit probability function $\mathbf{\Pi}$, with elements

$$\pi^{jk} = \pi_t(p^j | a^k) \equiv \frac{\exp(v^{jk}/(\kappa w))}{\sum_{n=1}^{\#p} \exp(v^{jn}/(\kappa w))}.$$

We can then find the steady state distribution as the fixed point of

$$\mathbf{\Psi} = (\mathbf{E}_{pa} - \mathbf{\Lambda}) . * \tilde{\mathbf{\Psi}} + \mathbf{\Pi} . * (\mathbf{E}_{pp} * (\mathbf{\Lambda} . * \tilde{\mathbf{\Psi}})), \quad (34)$$

$$\tilde{\mathbf{\Psi}} = \mathbf{R} * \mathbf{\Psi} * \mathbf{S}'. \quad (35)$$

Finally, we check whether

$$1 = \sum_{j=1}^{\#p} \sum_{k=1}^{\#a} \Psi^{jk} (p^j)^{1-\epsilon} \equiv p(w). \quad (36)$$

If so, an equilibrium value of w has been found.

4.4 Step 2: linearized dynamics

Given the steady state, the general equilibrium dynamics can be calculated by linearization. To do so, we eliminate as many variables from the equation system as we can. For additional simplicity, we assume linear labor disutility, $x(N) = \chi N$. Thus the first-order condition for labor reduces to $\chi = w_t u'(C_t)$, so we don't actually need to solve for N_t in order to calculate the rest of the equilibrium.¹¹ We can then summarize the general equilibrium equation system in terms of the exogenous shock process z_t , the lagged distribution of idiosyncratic states $\mathbf{\Psi}_{t-1}$, which is the endogenous component of the time t aggregate state; and finally the endogenous jump variables including \mathbf{V}_t , $\mathbf{\Pi}_t$, C_t , and π_t . The equation system reduces to

$$z_t = \phi_z z_{t-1} + \epsilon_t^z, \quad (37)$$

$$M_t = \mu e^{z_t} M_{t-1}, \quad (38)$$

$$\mathbf{\Psi}_t = (\mathbf{E}_{pa} - \mathbf{\Lambda}_t) . * \tilde{\mathbf{\Psi}}_t + \mathbf{\Pi}_t . * (\mathbf{E}_{pp} * (\mathbf{\Lambda}_t . * \tilde{\mathbf{\Psi}}_t)), \quad (39)$$

$$\mathbf{V}_t = \mathbf{U}_t + \beta E_t \left\{ \frac{u'(C_{t+1})}{u'(C_t)} [\mathbf{R}'_{t+1} * (\mathbf{V}_{t+1} + \mathbf{G}_{t+1}) * \mathbf{S}] \right\}, \quad (40)$$

$$1 = \sum_{j=1}^{\#p} \sum_{k=1}^{\#a} \Psi_t^{jk} (p^j)^{1-\epsilon}. \quad (41)$$

If we now collapse all the endogenous variables into a single vector

$$\vec{X}_t \equiv (\text{vec}(\mathbf{\Psi}_{t-1})', \text{vec}(\mathbf{V}_t)', C_t, M_{t-1}, \pi_t)'$$

¹¹The assumption $x(N) = \xi N$ is not essential; the more general case with nonlinear labor disutility simply requires us to simulate a larger equation system that includes N_t .

then the whole set of expectational difference equations (37)-(41) governing the dynamic equilibrium becomes a first-order system of the following form:

$$E_t \mathcal{F} \left(\vec{X}_{t+1}, \vec{X}_t, z_{t+1}, z_t \right) = 0, \quad (42)$$

where E_t is an expectation conditional on z_t and all previous shocks.

To see that the variables in vector \vec{X}_t are in fact the only variables we need, note that given π_t and π_{t+1} we can construct \mathbf{R}_t and \mathbf{R}_{t+1} . Given \mathbf{R}_t , we can construct $\tilde{\Psi}_t = \mathbf{R}_t * \Psi_{t-1} * \mathbf{S}'$ from Ψ_{t-1} . Under linear labor disutility, we can calculate $w_t = \chi/u'(C_t)$, which gives us all the information needed to construct \mathbf{U}_t , with (j, k) element equal to $w_t^{jk} \equiv (p^j - \frac{w_t}{a^k}) (p^j)^{-\epsilon} C_t$. Finally, given \mathbf{V}_t and \mathbf{V}_{t+1} we can construct $\mathbf{\Pi}_t$, \mathbf{D}_t , and \mathbf{D}_{t+1} , and thus \mathbf{A}_t and \mathbf{G}_{t+1} . Therefore the variables in \vec{X}_t and z_t are indeed sufficient to evaluate the system (37)-(41).

Finally, if we linearize system \mathcal{F} numerically with respect to all its arguments to construct the Jacobian matrices $\mathcal{A} \equiv D_{\vec{X}_{t+1}} \mathcal{F}$, $\mathcal{B} \equiv D_{\vec{X}_t} \mathcal{F}$, $\mathcal{C} \equiv D_{z_{t+1}} \mathcal{F}$, and $\mathcal{D} \equiv D_{z_t} \mathcal{F}$, then we obtain the following first-order linear expectational difference equation system:

$$E_t \mathcal{A} \Delta \vec{X}_{t+1} + \mathcal{B} \Delta \vec{X}_t + E_t \mathcal{C} z_{t+1} + \mathcal{D} z_t = 0, \quad (43)$$

where Δ represents a deviation from steady state. This system has the form considered by Klein (2000), so we solve our model using his QZ decomposition method.¹²

5 Results

5.1 Parameterization

We calibrate our model to match the monthly frequency of price changes in US microdata. We set the steady state growth rate of money to 0%. This is consistent with the zero average price change in the AC Nielsen data of household product purchases (Midrigan, 2008), which we use to test our model's ability to replicate salient features of price adjustment behavior. The model, like the data, addresses "regular" price changes, excluding temporary "sales".

As in Costain and Nakov (2008A, B), we take most of our parameterization directly from Golosov and Lucas (2007). Thus we set the discount factor to $\beta = 1.04^{-1/12}$. Consumption utility is CRRA, $u(C) = \frac{1}{1-\gamma} C^{1-\gamma}$, with $\gamma = 2$. Labor disutility is linear, $x(N) = \chi N$, with $\chi = 6$. The elasticity of substitution in the consumption aggregator is $\epsilon = 7$. Finally, the utility of real money holdings is logarithmic, $v(m) = \nu \log(m)$, with $\nu = 1$.

We assume productivity is AR(1) in logs: $\log A_{it} = \rho \log A_{it-1} + \varepsilon_t^a$, where ε_t^a is a mean-zero, normal, *iid* shock. We take the autocorrelation parameter from Blundell and Bond (2000) who estimate it from a panel of 509 US manufacturing companies over 8 years, 1982-1989. Their preferred estimate is 0.565 on an annual basis, which implies ρ around 0.95 in monthly frequency.

The variance of log productivity is $\sigma_a^2 = (1-\rho^2)\sigma_\varepsilon^2$, where σ_ε^2 is the variance of the innovation ε_t^a . We set the standard deviation of log productivity to $\sigma_a = 0.06$, which is the standard deviation of "reference costs" estimated by Eichenbaum, Jaimovich, and Rebelo (2008). Given our grid-based approximation, this implies a maximum absolute log price change of 0.48, which

¹²Alternatively, the equation system can be rewritten in the form of Sims (2001). We chose to implement the Klein method because it is especially simple and transparent to program.

covers, with an extra margin of 7%, the maximum absolute log price change of 0.45 observed in the AC Nielsen dataset.¹³

The same technology and preferences are assumed for all adjustment models. Thus, for each adjustment model we are left with one free parameter to calibrate. In the PPS and ENT cases, this is the the logit noise parameter κ ; for the other specifications we must either set the Calvo adjustment rate $\bar{\lambda}$ or the menu cost α . In all three cases, we choose the parameter to match the 10% median monthly frequency of price changes estimated by Nakamura and Steinsson (2008). All three cases are well identified because κ and α both strongly affect the frequency of price changes (as does $\bar{\lambda}$, of course).

5.2 Steady state microeconomic results

Table 1 and Figure 1 summarize our main estimation results. The estimated noise parameter for the PPS specification is $\kappa = 0.0428$, implying a relatively small deviation from fully rational behavior (i.e. low probabilities of large errors). Table 1 reports the losses implied by the errors in our model: we report the mean loss suffered relative to a fully rational firm, as a percentage of average revenues of a fully rational firm. In the baseline estimate, firms lose half a percent of revenues due to imperfect rationality. Losses are of a similar order of magnitude in all the other specifications too, being largest in the Calvo case and smallest under fixed menu costs. The distribution of losses can also be seen in the last panel of Figure 1, which shows the distribution both before and after firms decide whether or not to adjust. In addition, Figure 1 illustrates other equilibrium objects, such as the logit probability function π (first panel), the adjustment gains function D (second panel), and the adjustment probability λ (third panel).¹⁴

The first column of Table 1 also shows the main statistics of price adjustment implied by our estimate of the PPS specification. With a single free parameter, we hit our single calibration target, the 10% monthly frequency of price changes estimated by Nakamura and Steinsson (2008) (the last three columns of the table show data from several sources). The remaining moments also appear quite consistent with the distribution of price adjustments in the data, which is illustrated in Figure 2 (where it is compared to the fixed menu cost and Calvo specifications). The histogram of price changes shows 51 equally-spaced bins representing log price changes from -0.5 to 0.5. The blue shaded bars in the figure represent the AC Nielsen data, and the black line represents the results of our estimated model; the PPS specification generates a widely dispersed and mildly bimodal distribution like that seen in the data. The standard deviation of price adjustments is matched relatively well by the PPS specification: 14.5% in our model versus 13.2% in the data. As in the data, half of the price adjustments in our baseline model are price increases. Kurtosis is somewhat lower in our model (2.6) than in the data (3.5), as can be seen from the relatively fat tails in the blue shaded data in the figure.

Notably, in Figure 2 we see that the PPS framework reproduces one of the puzzling observations in the context of the fixed menu cost model: the coexistence of large and small price changes. While typical price adjustments are large in our model, well over 10%, nonetheless

¹³The aggregate steady state is projected on a price-productivity grid Γ of 25 by 25 points. The 25 log productivity steps of 0.02 cover four standard deviations of productivity above and below the average. Likewise, 25 log price steps of 0.02 cover the distance from the minimum to the maximum flexible price associated with the maximum and minimum productivities, respectively. In calculations with trend inflation extra points are added in the price dimension to ensure that the ergodic distribution remains on the grid

¹⁴Figure 1a illustrates the PPS specification, while Figure 1b shows the same equilibrium objects under the ENT specification.

around 20% (10%) of all adjustments are less than 5% (2.5%) in absolute value. All these statistics are close to those the data. The Calvo and the fixed menu cost models imply more concentrated distributions,¹⁵ so both imply many absolute adjustments of less than 5%. However, the fixed menu cost model concentrates the distribution of adjustment on two sharp spikes around an empty center, with no adjustments smaller than 2.5%, whereas in the Calvo model more than half of all changes are that small. Like the menu cost model, the PPS framework also exhibits (S,s) behavior, as illustrated by the second-to-last panel of Figure 1; inside the (S,s) bands firms choose not to adjust because the expected value of adjustment is not high enough to justify the risk of an error. Nonetheless, since the actual size of price adjustment is determined stochastically, a wide range of price adjustments is observed, including many small ones.

Another striking finding relates to the behavior of the adjustment hazard, that is, the probability of a price change as a function of the time since last adjustment. Figure 3 illustrates this hazard, comparing two empirical data sets with our model and with the fixed menu cost and Calvo specifications. In our model, error-prone decisions imply that firms sometimes readjust quickly after making a change, when their decision turns out to have been a mistake. This accounts for the spike in the adjustment hazard at one month. The adjustment hazard remains mildly decreasing over the next few months, driven by small errors subsequently compounded by unfavorable shocks; it is thereafter largely flat. This pattern is quite consistent with microdata; it fails by construction in the Calvo model and also contrasts with the increasing hazard observed in the FMC model. Many studies have suggested that decreasing hazards in the data may be caused by heterogeneity in adjustment frequencies across different types of products. However, Nakamura and Steinsson (2008) find decreasing hazards even after controlling for heterogeneity (see the blue dotted line in the figure). Likewise, Campbell and Eden (2010) find that for a given product (as defined by the UPC code), recently-set prices are more likely to be adjusted than “older” prices.

Figures 4 and 5 illustrate our model’s implications for two other stylized facts from microdata. In Figure 4, we see that in Klenow and Kryvtsov’s (2008) microdata, the average size of the price adjustment is essentially invariant with the time since last price adjustment. The fixed menu cost model is strongly consistent with these data; the precautionary price stickiness model also does fairly well, except at a horizon of one month. The Calvo model is strongly rejected on this score, since it implies that the size of the price adjustment increases with the time since last adjustment. Figure 5 illustrates Campbell and Eden’s (2010) finding that for any given product, prices far from the time average are likely to have been set recently. The Calvo model is again inconsistent with this observation. Both the fixed menu cost and precautionary price stickiness models generate a U-shaped relation between the percentage of “young” prices (those set within the last month) and the deviation from the average price; the former understates the relationship whereas the latter overstates it.

5.3 Effects of changes in monetary policy

Figure 6 investigates the effects of large changes in steady state inflation, compared with empirical observations from Mexican data reported by Gagnon (2007) for periods with annual inflation rates of 5%, 29%, and 63%. The Calvo model is not shown, since the initial Calvo specification

¹⁵To understand this result, note that conditional on a given productivity process, the PPS and ENT models generate greater price variability than the Calvo and FMC models do, since price adjustments in the PPS and ENT cases partially reflect errors, instead of being reactions to productivity shocks only.

with adjustment rate $\bar{\lambda} = 0.1$ is incompatible with the rise in inflation shown in the figure (it generates negative firm values at all inflation rates shown in the figure except for the lowest one). The first panel shows how the probability of adjustment rises with the steady state inflation rate. The FMC, PPS, and ENT specifications all generate increases in the adjustment frequency very similar to that observed in the data (the level of the adjustment frequency is higher in the data than in the models, but this simply reflects the fact that we are comparing US model calibrations with Mexican data).

The second panel shows how the standard deviation of price adjustments changes with the inflation rate. The Mexican data show a small, nonmonotonic change in the standard deviation as inflation rises. The PPS and ENT models both generate a standard deviation fairly close to that in the data, and there is little change in this standard deviation as inflation rises. The FMC model instead has a much smaller standard deviation of price adjustments; moreover, this standard deviation drops dramatically (by more than half) as inflation rises to the levels observed in Gagnon's data. Finally, the third panel reports the fraction of price adjustments that are increases; all three models somewhat exaggerate the rise in the percentage of price increases, as compared with Gagnon's data.

Figures 7-9 illustrate the implications of our estimated model for macroeconomic dynamics. The figures show impulse responses of inflation and consumption to money supply shocks ϵ_t^z , comparing the PPS and ENT specifications with the Calvo and FMC models. The AR(1) persistence is set to $\phi_z = 0.8$ monthly, implying quarterly persistence of 0.5. Under the baseline parameterization, both the PPS and ENT versions of our model imply real effects of monetary shocks much smaller than those in the Calvo case. In fact, the consumption response for the ENT specification is almost identical to that associated with FMC; the initial impact on consumption under PPS is similar, though in this case the response is more persistent. The same conclusions can be drawn from Table 2, which quantifies the real effects of money shocks across models following Golosov and Lucas (2007). The table reports the percentage of output variation that can be explained by money shocks alone (assuming the money shock is sufficiently variable to explain 100% of inflation variation); it also reports the slope coefficient of a "Phillips curve" regression of output on inflation. Under both of these measures, the ENT and FMC models have imply nearly identical degrees of monetary nonneutrality, whereas PPS implies real effects roughly twice as strong. Under the Calvo specification, money shocks have real effects roughly four times stronger than they do in the PPS case.

The explanation for the weak response of consumption in our model is the same one Golosov and Lucas (2007) identified in the FMC context: a strong "selection" effect. In our model, as in the case of FMC, adjustment occurs if and only if it is sufficiently valuable. Therefore all the most important price misalignments are eliminated after a money shock, which makes money approximately neutral. To measure this effect precisely, define the productivity-specific optimal price level $p_t^{*k} \equiv \operatorname{argmax}_p v_t(p, a^k)$, and also $x_t^{*jk} \equiv \log(p_t^{*k}/p^j)$, the desired log price adjustment of a firm at time t with productivity a^k and real price p^j . The actual log price adjustment of such a firm (call it i) can thus be decomposed as $x_{it} = x_t^{*jk} + \epsilon_{it}$, where ϵ_{it} is an error, in logs. We can then write the average desired adjustment as $\bar{x}_t^* = \sum_{j,k} x_t^{*jk} \tilde{\Psi}_t^{jk}$, and write the fraction of firms adjusting as $\bar{\lambda}_t = \sum_{j,k} \lambda_t^{jk} \tilde{\Psi}_t^{jk}$, and write the average log error as $\bar{\epsilon}_t = \sum_{j,k,l} \pi_t^{lk} \log(p^l/p_t^{*k}) \lambda_t^{jk} \tilde{\Psi}_t^{jk}$.

Then inflation can be written as

$$\pi_t = \sum_{j,k} x_t^{*jk} \lambda_t^{jk} \tilde{\Psi}_t^{jk} + \bar{\epsilon}_t. \quad (44)$$

To a first-order approximation, we can decompose the deviation in inflation at time t as

$$\Delta\pi_t = \log \Pi_t - \log \Pi = \bar{\lambda} \Delta \bar{x}_t^* + \bar{x}^* \Delta \bar{\lambda}_t + \Delta \sum_{j,k} x_t^{jk} (\lambda_t^{jk} - \bar{\lambda}_t) \tilde{\Psi}_t^{jk} + \Delta \bar{\epsilon}_t, \quad (45)$$

where terms without time subscripts represent steady states, and Δ represents a change relative to steady state.¹⁶

The first term, $\mathcal{I}_t \equiv \bar{\lambda} \Delta \bar{x}_t^*$, can be called the “intensive margin”; it is the part of inflation due to changes in the average desired adjustment, holding fixed the fraction of firms adjusting. The second term, $\mathcal{E}_t \equiv \bar{x}^* \Delta \bar{\lambda}_t$, can be called the “extensive margin”; it is the part of inflation due to changes in the fraction of firms adjusting, assuming the average desired adjustment among those who adjust equals the steady-state average in the whole population. The third effect, $\mathcal{S}_t \equiv \Delta \sum_{j,k} x_t^{jk} (\lambda_t^{jk} - \bar{\lambda}_t) \tilde{\Psi}_t^{jk}$, represents the part of inflation due to redistributing adjustment opportunities from firms desiring small price adjustments to firms desiring large adjustments, while fixing the total number adjusting. Finally, there is a fourth term, $\Delta \bar{\epsilon}_t$, which is the change in the average log error.

The inflation impulse responses associated with all the specifications considered are decomposed in Figure 9. By construction, all variation in inflation in the Calvo model is caused by the intensive margin, since the probability of adjustment never varies in that specification. The PPS, ENT, and FMC specifications also show a rise in inflation due to the intensive margin, since the money supply shock increases the average desired price adjustment. But in these cases there is a much more important rise in inflation that goes through the selection effect. The money supply shock leaves some firms with an especially low real price; some of these firms then decide to adjust their prices, which is the selection effect. The larger price increase in these specifications explains the small consumption impact seen in Figure 7.

Note finally that under all specifications the extensive margin effect is negligible. This occurs whenever we start from an inflation rate near zero: changing the fraction adjusting while holding the mean adjustment fixed must have a negligible impact since the mean adjustment must be roughly zero when inflation is roughly zero. Likewise, the change in the average log error (not shown) has a negligible impact on inflation. Instead, it is the reallocation of adjustment opportunities *to those that desire a large price increase*, which is classified as the selection effect, that really matters for the inflation response.

5.4 Changing the risk of errors

Tables 1-2 and Figure 8 also show how the behavior of the PPS model varies with the parameter κ that controls the amount of noise in the decision process. In Table 1, cutting noise in half causes the frequency of price adjustment to rise from 10% to 12.3% per month, and the average price change becomes smaller. This makes sense— with greater rationality, price adjustment is less risky, so firms become willing to adjust even when their prices are not so far out of line. Moreover, the elimination of errors also reduces the size of price changes directly by cutting out one component of the price adjustment process.

¹⁶See Costain and Nakov (2011B) for further discussion of this decomposition.

Likewise, in Figure 8, cutting noise in half increases the initial spike in inflation, so that the real effect of the nominal shock is smaller than in the baseline calibration. In other words, as rationality increases the PPS model eventually converges to a fully flexible setup in which money is essentially neutral.

5.5 Logit equilibrium versus control costs

Finally, it is interesting to contrast the behavior of the baseline logit equilibrium version of our model, which we called ‘PPS’, with the version based on control costs, called ‘ENT’. The main conclusion is that the implications of the two versions of the model are very similar, though calibrating the ENT specification results in somewhat lower frictions than PPS.

In terms of microeconomic implications, Table 1 and Figure 2 show that ENT implies a more concentrated distribution, with smaller adjustments, than PPS. While PPS and ENT are both quite consistent with the histogram of price adjustments, the ENT specification is somewhat more sharply bimodal than the data, whereas PPS is more weakly bimodal and more spread out. Figures 3-5 show that PPS and ENT have qualitatively similar implications for several other micro facts. In terms of impulse responses, Figures 7-9 show that ENT is even closer to monetary neutrality than PPS, behaving almost identically to the fixed menu cost model.

These findings can be understood by noting that we must recalibrate the noise parameter κ when going from PPS to the ENT setup. In the PPS specification, adjustment occurs if and only if the value of adjustment D is nonnegative; we calibrate $\kappa = 0.0478$ in order to match a 10% monthly adjustment frequency. Moving to the ENT specification implies that firms are less likely to adjust, *ceteris paribus*: adjustment does not occur unless the value of adjustment (in labor units) exceeds the control cost $C(\vec{\pi})$. Therefore matching the same 10% adjustment frequency requires a substantially lower degree of noise; the required calibration for the ENT specification is $\kappa = 0.0050$.

Since the ENT model combines a substantially lower noise parameter with a nonzero cost of adjustment, it behaves more like the FMC model than PPS does. This is especially true in terms of impulse responses. The lower noise parameter also explains why ENT exhibits smaller price adjustments than PPS. Nonetheless, the presence of errors means that ENT generates a much more realistic distribution of price changes than FMC does, with a smoother histogram that includes some very small adjustments. The presence of errors also implies that ENT exhibits price adjustment hazard that declines (slightly) with the time since last adjustment, in contrast with the FMC model. The lower noise in the ENT calibration also explains why the mean loss reported in Table 1 is lower for ENT than it is for PPS, in spite of the fact that no adjustment costs are deducted from the value function under the PPS specification.

6 Conclusions

We have analyzed the pricing behavior of near-rational firms subject to idiosyncratic and aggregate shocks which can adjust their prices at any time, but may make mistakes. We model error-prone behavior by assuming firms play a dynamic logit equilibrium. Prices are therefore endogenously sticky: when a firm’s current price is sufficiently close to the optimum, it prefers to leave well enough alone, avoiding the risk of accidentally choosing a worse price.

This way of modeling price stickiness is consistent with several observations from microdata that are hard to explain in most existing frameworks. Even though the decision to adjust prices

has an (S,s) structure, nonetheless many small price changes coexist with larger ones. Moreover, the price adjustment hazard exhibits negative duration dependence; the size of price changes is largely independent of the time elapsed since the last adjustment; and extreme prices tend to have been set recently.

When we estimate the logit rationality parameter, our model matches both the frequency of adjustment and the size distribution of price changes quite successfully. This is especially true for the PPS specification. Since our setup guarantees that firms making sufficiently large errors will choose to adjust, it generates a "selection effect" in response to nominal shocks that largely mitigates the strong real effects of money shocks found in the Calvo model.

While the Calvo model offers one way to explain the presence of small price changes, it fails to generate many other prominent micro facts. Our model performs considerably better than Calvo in matching microdata, but because of its strong selection effect, it generates far less stickiness of the aggregate price level than the Calvo model does. Thus, resolving the "puzzle" of small price adjustments need not, by itself, alter Golosov and Lucas' conclusion that money shocks have small real effects.

As we hinted in the introduction, thinking of nominal rigidity as a near-rational phenomenon makes it natural to consider mistakes both in the *size* and *timing* of price adjustment. We are studying a model of this type in ongoing work (Costain and Nakov 2011C). It would also be interesting to apply logit decision-making in many macroeconomic contexts other than price setting.

References

- [1] Akerlof, George, and Janet Yellen (1985), "A near-rational model of the business cycle with wage and price inertia." *Quarterly Journal of Economics* 100 (Supplement), pp. 823-38.
- [2] Álvarez, Luis (2008), "What do micro price data tell us on the validity of the New Keynesian Phillips curve?" *Economics: The Open-Access, Open-Assessment E-Journal* 2 (19).
- [3] Anderson, Simon; Jacob Goeree, and Charles Holt (2002), "The logit equilibrium: a perspective on intuitive behavioral anomalies." *Southern Economic Journal* 69 (1), pp. 21-47.
- [4] Baron, Richard, Jaques Durieu, Hans Haller, and Philippe Solal (2002), "Control costs and potential functions for spatial games." *International Journal of Game Theory* 31, pp. 541-561.
- [5] Barro, Robert (1972), "A theory of monopolistic price adjustment." *Review of Economic Studies* 39 (1), pp. 17-26.
- [6] Caballero, Ricardo, and Eduardo Engel (1993), "Microeconomic rigidities and aggregate price dynamics." *European Economic Review* 37 (4), pp. 697-711.
- [7] Caballero, Ricardo, and Eduardo Engel (1999), "Explaining investment dynamics in U.S. manufacturing: a generalized (S,s) approach." *Econometrica* 67 (4), pp. 783-862.
- [8] Caballero, Ricardo, and Eduardo Engel (2007), "Price stickiness in Ss models: new interpretations of old results." *Journal of Monetary Economics* 54, pp. 100-21.

- [9] Calvo, Guillermo (1983), “Staggered prices in a utility-maximizing framework.” *Journal of Monetary Economics* 12, pp. 383-98.
- [10] Campbell, Jeffrey, and Benjamin Eden (2010), “Rigid prices: evidence from U.S. scanner data.” FRB Chicago WP2005-08 (revised July 2010).
- [11] Caplin, Andrew, and Daniel Spulber (1987), “Menu costs and the neutrality of money.” *Quarterly Journal of Economics* 102 (4), pp. 703-25.
- [12] Caplin, Andrew, and John Leahy (1991), “State-dependent pricing and the dynamics of money and output.” *Quarterly Journal of Economics* 106 (3), pp. 683-708.
- [13] Carvalho, Carlos (2006), “Heterogeneity in price stickiness and the real effects of monetary shocks.” *Berkeley Electronic Press Frontiers of Macroeconomics* 2 (1), Article 1.
- [14] Costain, James, and Anton Nakov (2011A), “Price adjustments in a general model of state-dependent pricing.” *Journal of Money, Credit and Banking* 43 (2-3), pp. 385-406.
- [15] Costain, James, and Anton Nakov (2011B), “Distributional dynamics under smoothly state-dependent pricing.” European Central Bank Working Paper 1333.
- [16] Costain, James, and Anton Nakov (2011C), “Logit price dynamics.” Manuscript, Bank of Spain.
- [17] Cover, Thomas, and Joy Thomas (2006), *Elements of Information Theory*, 2nd ed. Wiley Interscience.
- [18] Danziger, Leif (1999), “A dynamic economy with costly price adjustments.” *American Economic Review* 89 (4), pp. 878-901.
- [19] Dotsey, Michael; Robert King, and Alexander Wolman (1999), “State-dependent pricing and the general equilibrium dynamics of money and output.” *Quarterly Journal of Economics* 114 (2), pp. 655-90.
- [20] Dotsey, Michael, and Robert King (2005), “Implications of state-dependent pricing for dynamic macroeconomic modeling.” *Journal of Monetary Economics* 52 (1), pp. 213-42.
- [21] Eichenbaum, M.; N. Jaimovich, and S. Rebelo (2008), “Reference prices and nominal rigidities.” CEPR Discussion Paper 6709.
- [22] Gagnon, Etienne (2007), “Price setting under low and high inflation: evidence from Mexico.” *Quarterly Journal of Economics* 124, pp. 1221-1263.
- [23] Gertler, Mark, and John Leahy (2008), “A Phillips curve with an Ss foundation.” *Journal of Political Economy* 116 (3), pp. 533-72.
- [24] Golosov, Mikhail, and Robert E. Lucas, Jr. (2007), “Menu costs and Phillips curves.” *Journal of Political Economy* 115 (2), pp. 171-99.
- [25] Kashyap, Anil (1995), “Sticky prices: new evidence from retail catalogs.” *Quarterly Journal of Economics* 110 (1), pp. 45-74.

- [26] Klenow, Peter, and Oleksiy Kryvtsov (2008), “State-dependent or time-dependent pricing: does it matter for recent US inflation?”, *Quarterly Journal of Economics* 123, pp. 863-904.
- [27] Klenow, Peter, and Benjamin Malin (2009), “Microeconomic evidence on price setting.” Forthcoming, *Handbook of Monetary Economics*, edited by Benjamin Friedman and Michael Woodford.
- [28] Mackowiak, Bartosz, and Mirko Wiederholt (2011), “Business cycle dynamics under rational inattention.” ECB Working Paper 1331.
- [29] Mankiw, N. Gregory (1985), “Small menu costs and large business cycles: a macroeconomic model of monopoly.” *Quarterly Journal of Economics* 100 (2), pp. 529-37.
- [30] Marsili, Matteo (1999), “On the multinomial logit model.” *Physica A* 269, pp. 9-15.
- [31] Matejka, Filip (2010), “Rationally inattentive seller: sales and discrete pricing.” CERGE-EI Working Paper wp408.
- [32] Matejka, Filip, and Alisdair McKay (2011), “Rational inattention to discrete choices: a new foundation for the multinomial logit model.” Manuscript, CERGE-EI and Boston University.
- [33] Mattsson, Lars-Göran, and Jörgen Weibull (2002), “Probabilistic choice and procedurally bounded rationality.” *Games and Economic Behavior* 41, pp. 61-78.
- [34] McKelvey, Richard, and Thomas Palfrey (1995), “Quantal response equilibrium for normal form games.” *Games and Economic Behavior* 10, pp. 6-38.
- [35] McKelvey, Richard, and Thomas Palfrey (1998), “Quantal response equilibrium for extensive form games.” *Experimental Economics* 1, pp. 9-41.
- [36] Midrigan, Virgiliu (2011), “Menu Costs, Multi-Product Firms and Aggregate Fluctuations.” Forthcoming, *Econometrica*.
- [37] Nakamura, Emi, and Jón Steinsson (2008), “Five facts about prices: a reevaluation of menu cost models.” *Quarterly Journal of Economics* 123 (4), pp. 1415-1464.
- [38] Reiter, Michael (2009), “Solving heterogeneous-agent models by projection and perturbation.” *Journal of Economic Dynamics and Control* 33 (3), pp. 649-65.
- [39] Rotemberg, Julio (1982), “Monopolistic price adjustment and aggregate output.” *Review of Economic Studies* 49, pp. 517-31.
- [40] Sheshinski, Eytan, and Yoram Weiss (1977), “Inflation and costs of price adjustment.” *Review of Economic Studies* 44 (2), pp. 287-303.
- [41] Sims, Christopher (1998), “Stickiness.” *Carnegie-Rochester Conference Series on Public Policy* 49, pp. 317-56.
- [42] Sims, Christopher (2003), “Implications of rational inattention.” *Journal of Monetary Economics* 50, pp. 665-90.

- [43] Stahl, Dale (1990), “Entropy control costs and entropic equilibrium.” *International Journal of Game Theory* 19, pp. 129-138.
- [44] Tutino, Antonella (2009), “The rigidity of choice: lifetime savings under information-processing constraints.” Manuscript, Federal Reserve Bank of Dallas.
- [45] Woodford, Michael (2008), “Information-constrained state-dependent pricing.” *Journal of Monetary Economics* 56, pp. S100-S124.
- [46] Zbaracki, Mark; Mark Ritson; Daniel Levy; Shantanu Dutta; and Mark Bergen (2004), “Managerial and customer costs of price adjustment: direct evidence from industrial markets.” *Review of Economics and Statistics* 86 (2), pp. 514-33.

Table 1. Model-Simulated Statistics and Evidence (zero trend inflation)

	Model PPS			Model ENT		Calvo		MC		Evidence		
	$\kappa = 0.0428$	$2 \times \kappa$	$\frac{1}{2} \times \kappa$	$2 \times \kappa$	$\kappa = 0.0050$	10	10	10	10	MAC	NS	KK
Freq. of price changes	10	7.4	12.3	10	10	10	10	10	10	20.5	10	21.9
Mean absolute price change	11.9	13.5	10.1	6.5	6.5	2.8	5.5	10.4	11.3			
Std of price changes	14.5	16.5	12.3	7.3	7.3	3.7	5.6	13.2				
Kurtosis of price changes	2.6	2.5	2.7	2.3	2.3	4.2	1.2	3.5				
Percent of price increases	50	50.2	49.7	50	50	48	50.7	50	66	56		
% of abs price changes $\leq 5\%$	19.5	17.1	23.2	32.2	32.2	83.6	42.4	25	44			
% of abs price changes $\leq 2.5\%$	9.4	8.4	10.9	10.1	10.1	55	0	11				
Mean loss (% flex. revenue)	0.55	0.70	0.40	0.34	0.34	0.61	0.31					

Note: All statistics refer to regular consumer price changes excluding sales, and are stated in percent. The last three columns reproduce the statistics reported by Midrigan (2011) for AC Nielsen (MAC),

Nakamura and Steinsson (2008) (NS), and Klenow and Kryvtsov (2008) (KK).

Last row reports the mean difference between profits under flexible and rigid prices, as a percentage of the mean revenues of a flexible-price firm, for each model of rigidity.

Table 2. Variance decomposition and Phillips curves

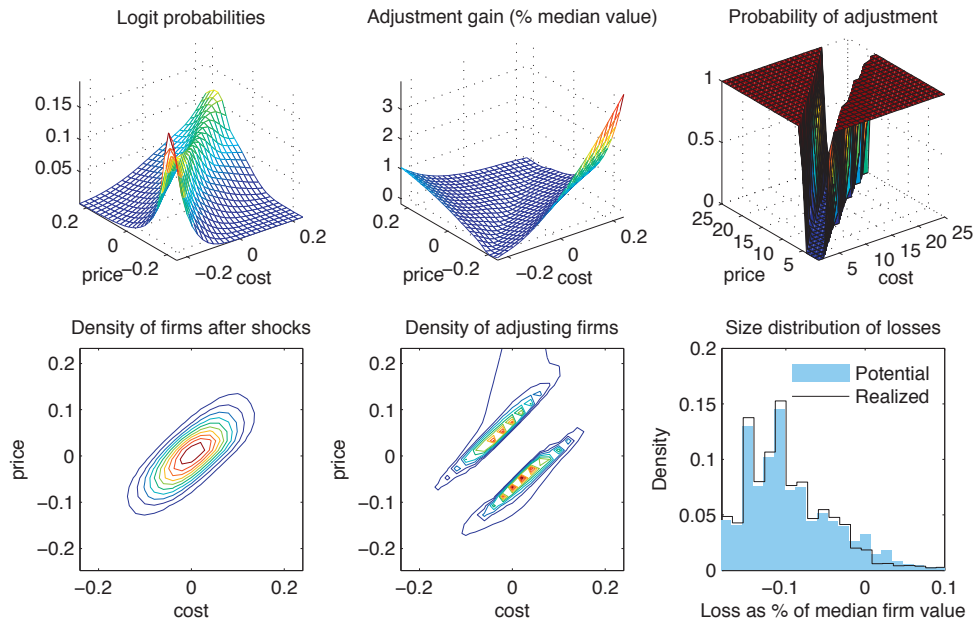
<i>Correlated money growth shock</i> ($\phi_z = 0.8$)	Data	Model PPS $\kappa = 0.0428$	$\frac{1}{2} \times \kappa$	Model ENT $\kappa = 0.0050$	Calvo	FMC
Frequency of non-zero price changes (%)		10	12.3	10	10	10
Std of money shock (x100)		0.153	0.192	0.134	0.331	0.122
Std of quarterly inflation (x100)	0.246	0.246	0.246	0.246	0.246	0.246
% explained by μ shock alone		100	100	100	100	100
Std of quarterly output growth (x100)	0.510	0.310	0.460	0.174	1.08	0.195
% explained by μ shock alone		60.7	90.3	43.2	212	38.3
Slope coeff. of Phillips curve*		0.273	0.429	0.191	1.100	0.149
Standard error		0.006	0.009	0.008	0.070	0.012
R ² of regression		0.987	0.987	0.949	0.892	0.832

*The "slope coefficients" are 2SLS estimates of the effect of inflation on consumption.

First stage: $\pi_t^q = \alpha_1 + \alpha_2 \mu_t^q + \epsilon_t$; second stage: $c_t^q = \beta_1 + \beta_2 \hat{\pi}_t^q + \epsilon_t$, where the instrument

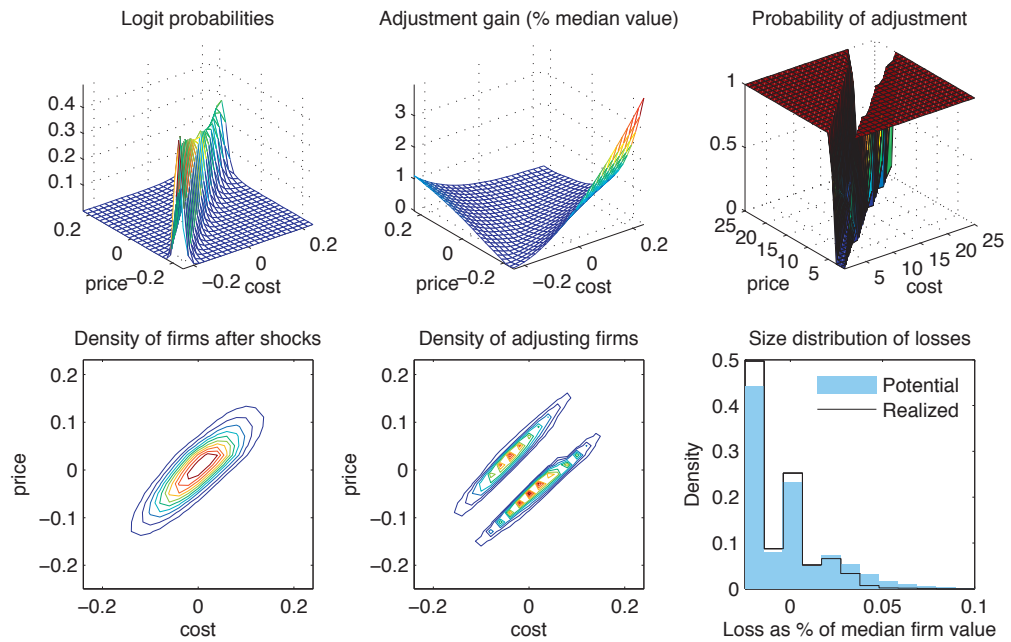
μ_t^q is the exogenous growth rate of the money supply and the superscript q indicates quarterly averages.

FIGURES FOR “PRECAUTIONARY PRICE STICKINESS”



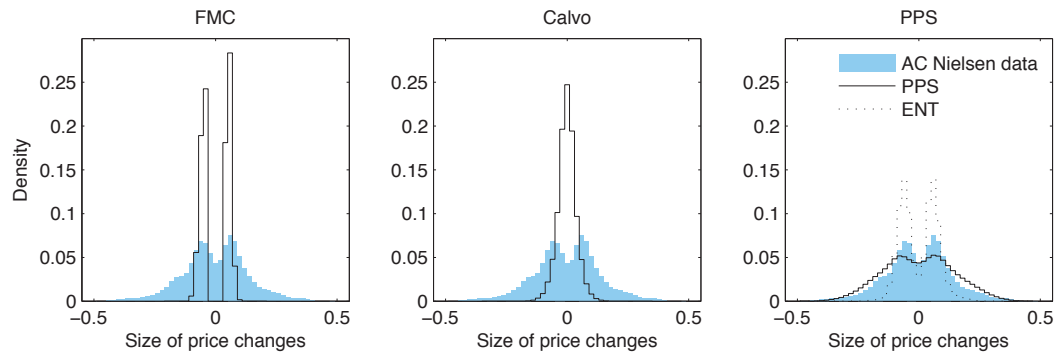
- 1 Fig. 1a. Properties of the steady state distribution: PPS specification
- 2 Note: first five panels show real prices p and costs $1/a$ as log deviations from mean.
- 3 First panel: Probability π of choosing real price p , conditional on cost $1/a$.
- 4 Second panel: Adjustment value D , expressed as fraction of median firm value.
- 5 Third panel: Adjustment probability λ .
- 6 Fourth panel: Contour plot of beginning-of-period density of firms, $\tilde{\Psi}$.
- 7 Fifth panel: Contour plot of density of adjusting firms.
- 8 Sixth panel: Densities of adjustment values before adjustment (D , blue shading), and after adjustment (solid line).

FIGURES FOR “PRECAUTIONARY PRICE STICKINESS”



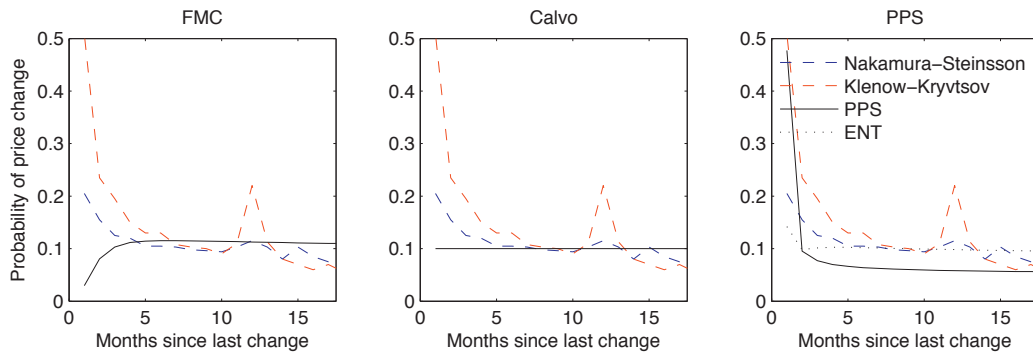
- 1 Fig. 1b. Properties of the steady state distribution: ENT specification
- 2 Note: first five panels show real prices p and costs $1/a$ as log deviations from mean.
- 3 First panel: Probability π of choosing real price p , conditional on cost $1/a$.
- 4 Second panel: Adjustment value D , expressed as fraction of median firm value.
- 5 Third panel: Adjustment probability λ .
- 6 Fourth panel: Contour plot of beginning-of-period density of firms, $\tilde{\Psi}$.
- 7 Fifth panel: Contour plot of density of adjusting firms.
- 8 Sixth panel: Densities of adjustment values before adjustment (D , blue shading), and after adjustment (solid line).

FIGURES FOR “PRECAUTIONARY PRICE STICKINESS”



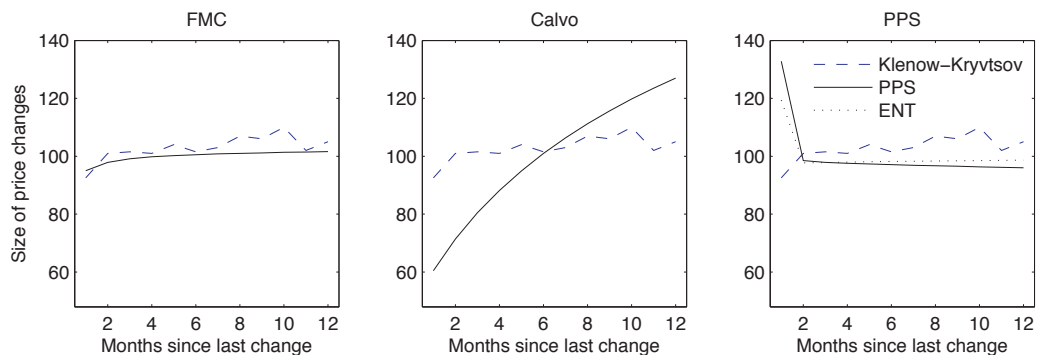
- 1 Fig. 2. Distribution of nonzero price adjustments: comparing models
- 2 Shaded blue area: Nonzero price adjustments in AC Nielsen data of Midrigan (2011), controlling for sales.
- 3 Black lines: model simulations under fixed menu costs (left), Calvo model (middle), and 'PPS' (right, solid) and 'ENT'
- 4 (right, dotted) specifications of our model.

FIGURES FOR “PRECAUTIONARY PRICE STICKINESS”



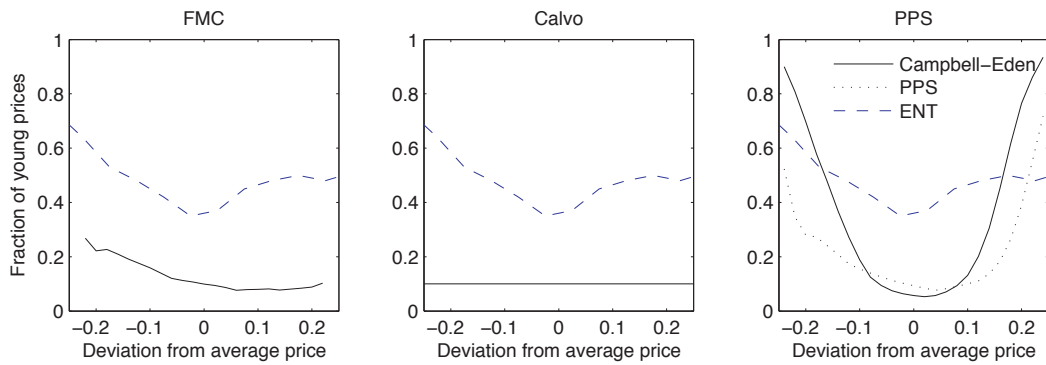
- 1 Fig. 3. Price adjustment hazard: comparing models
- 2 Red and blue dotted lines: price adjustment hazard as function of time since last adjustment, in data of Nakamura and
- 3 Steinsson (2008) and Klenow and Kryvtsov (2008).
- 4 Black lines: model simulations under fixed menu costs (left), Calvo model (middle), and 'PPS' (right, solid) and 'ENT'
- 5 (right, dotted) specifications of our model.

FIGURES FOR “PRECAUTIONARY PRICE STICKINESS”



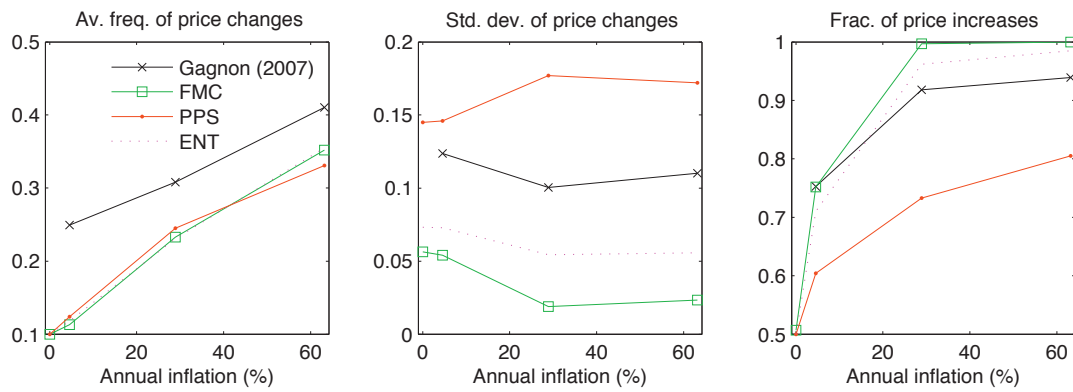
- 1 Fig. 4. Mean adjustment and price duration: comparing models
- 2 Blue dotted line: average absolute nonzero price change, in percent, as function of time since last adjustment, in data
- 3 of Klenow and Kryvtsov (2008).
- 4 Black lines: model simulations under fixed menu costs (left), Calvo model (middle), and 'PPS' (right, solid) and 'ENT'
- 5 (right, dotted) specifications of our model.

FIGURES FOR “PRECAUTIONARY PRICE STICKINESS”



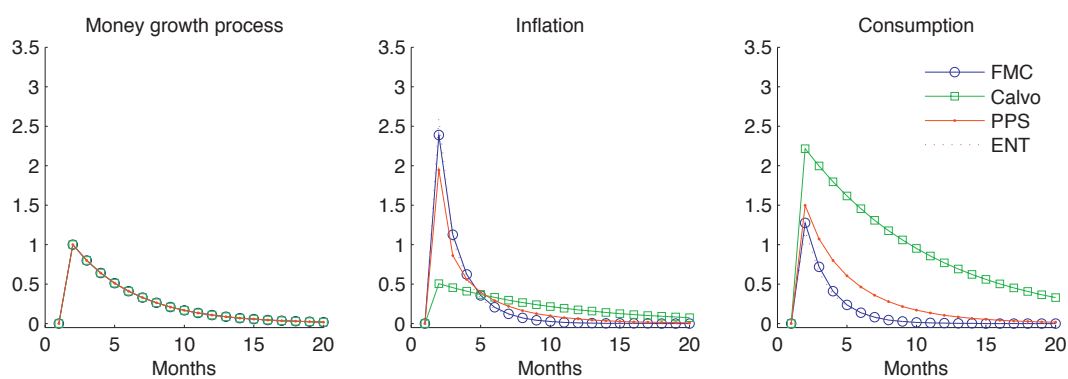
- 1 Fig. 5. Extreme prices tend to be young: comparing models
- 2 Blue dotted line: fraction of prices set no more than two months previously, as function of log deviation from product-
- 3 specific mean price, in data of Campbell and Eden (2010).
- 4 Black lines: model simulations under fixed menu costs (left), Calvo model (middle), and 'PPS' (right, solid) and 'ENT'
- 5 (right, dotted) specifications of our model.

FIGURES FOR “PRECAUTIONARY PRICE STICKINESS”



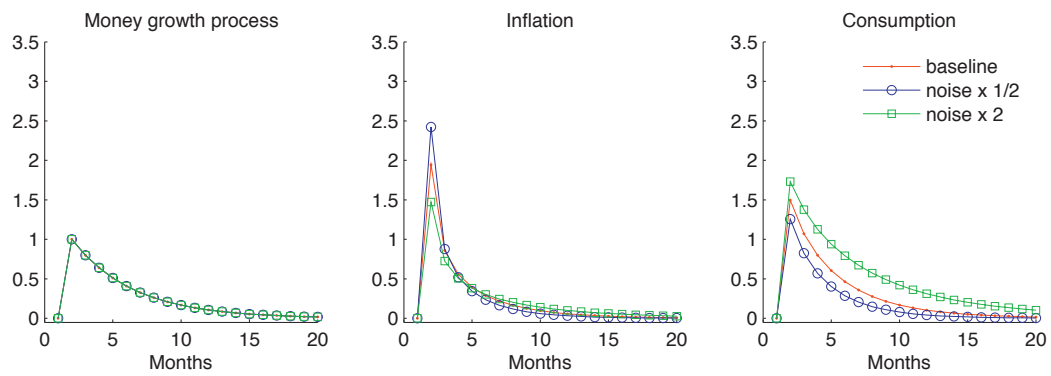
- 1 Fig. 6. Effects of trend money growth: comparing models
- 2 Black line with checks: Mexican data of Gagnon (2008), at annual inflation rates of 5%, 29%, and 63%.
- 3 Other lines: simulations of FMC (green with circles), PPS (red with dots), ENT (dotted pink).
- 4 First panel: average frequency of price adjustments, as function of inflation rate.
- 5 Second panel: standard deviation of nonzero price adjustments, as function of inflation rate.
- 6 Third panel: fraction of nonzero price adjustments which are increases, as function of inflation rate.

FIGURES FOR “PRECAUTIONARY PRICE STICKINESS”



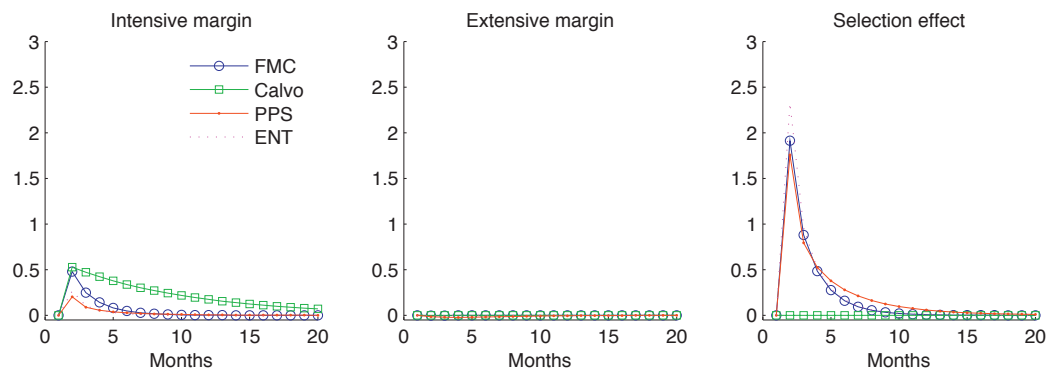
- 1 Fig. 7. Impulse responses to money growth shock: comparing models
- 2 Simulated responses of inflation and consumption to a money growth shock, monthly autocorrelation 0.8.
- 3 Green squares: Calvo model; blue circles: fixed menu costs; red line: 'PPS' model; dotte pink: 'ENT' model.

FIGURES FOR “PRECAUTIONARY PRICE STICKINESS”



- 1 Fig. 8. Impulse responses to money growth shock: effects of noise parameter
- 2 Simulated responses of inflation and consumption to a money growth shock, monthly autocorrelation 0.8.
- 3 Red line: 'PPS' model with baseline noise parameter, $\kappa = 0.0428$; green squares: 'PPS' model with higher noise,
- 4 $\kappa = 0.0856$; blue circles: 'PPS' model with lower noise, $\kappa = 0.0214$.

FIGURES FOR “PRECAUTIONARY PRICE STICKINESS”



- 1 Fig. 9. Inflation decomposition after money shock: comparing models
- 2 Intensive, extensive, and selection margins of inflation response to a money growth shock, monthly autocorrelation 0.8.
- 3 Green squares: Calvo model; blue circles: fixed menu costs; red line: 'PPS' model; dotted pink: 'ENT' model.

