

DISCUSSION PAPER SERIES

No. 9912

PRECAUTIONARY PRICE STICKINESS

James Costain and Anton Nakov

INTERNATIONAL MACROECONOMICS



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP9912.php

PRECAUTIONARY PRICE STICKINESS

James Costain, Banco de España
Anton Nakov, European Central Bank and CEPR

Discussion Paper No. 9912
March 2014

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INTERNATIONAL MACROECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: James Costain and Anton Nakov

ABSTRACT

Precautionary price stickiness*

This paper proposes a model in which retail prices are sticky even though firms can always change their prices at zero cost. Instead of imposing a "menu cost", we assume that more precise decisions are more costly. In equilibrium, firms optimally make some errors in price-setting, thus economizing on managerial time. Both the time cost of choice, and the resulting risk of errors, give firms an incentive to leave their prices unchanged until they perceive a sufficiently large deviation from the optimal price.

We show that pricing errors help explain several "puzzles" from microdata: (1) small and large price changes coexist; (2) the probability of price adjustment is largely independent of the time since last adjustment; (3) the size of the adjustment is largely independent of the time since last adjustment; (4) extreme prices are younger than prices near the center of the distribution; (5) the coefficient of variation of prices is greater than that of costs; (6) the standard deviation of price adjustments is largely independent of the inflation rate, and the fraction of price increases converges slowly towards 100% as inflation rises.

However, on the macroeconomic side, pricing errors do little to explain the real effects of monetary shocks. Since firms making sufficiently large errors always choose to adjust, a nominal shock generates a strong, inflationary "selection effect". Thus, like Golosov and Lucas (2007), we find that money shocks are almost neutral, but our model fits microdata better than their specification does.

JEL Classification: C72, D81 and E31

Keywords: (S,s), logit equilibrium, near-rational behavior, nominal rigidity, adjustment and state-dependent pricing

James Costain
División de Investigación
Servicio de Estudios
Banco de España
Calle Alcalá 50
28014 Madrid
SPAIN

Anton Nakov
European Central Bank
DG-Research
Frankfurt
GERMANY

Email: james.costain@bde.es

Email: anton.nakov@gmail.com

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=149411

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=160916

*For their helpful comments, we thank Fernando Álvarez, Jordi Galí, Kevin Lansing, John Leahy, Bartosz Mackowiak, Filip Matejka, Antonela Tutino, Mirko Wiederholt, Jonathan Willis, the editor and two anonymous referees, and seminar participants at the Bank of Spain, the San Francisco Fed, the ECB, HEC Paris, the Federal Reserve Board, the Simposio de Analisis Economico 2009, 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, or the Eurosystem.

Submitted 12 March 2014

1 Introduction

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 effects of frictions in the second decision. In other words, we assume that firms can adjust their prices costlessly at any time, and that they adjust if and only if it is optimal to do so, 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.

The key assumption of our model is that making precise decisions is costly. That is, while a firm could in principle adjust its price costlessly at any time, the price it sets would then exhibit maximal randomness. Therefore, in equilibrium, firms choose to pay a cost when they adjust prices, but it is not a “menu cost” in the sense of resources devoted to the physical task of adjusting posted prices. Instead, firms reduce their pricing errors by spending time on managerial decision making, consistent with the evidence of Zbaracki *et al.* (2004). Note that there are many possible ways of measuring precision; for concreteness, we measure it in terms of entropy. We show that if decision costs are proportional to entropy reduction, then the distribution of price adjustments takes the form of a multinomial logit. Our model nests full frictionless rationality as the limiting case in which the firm sets the optimal price with probability one every period. The general equilibrium of the model is 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.¹

The fact that precision is costly gives rise to price stickiness, for two reasons. First, 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.² Second, firms know that when they make an adjustment, they should devote some time to choosing which price to set. This managerial cost deters price adjustment, just like a menu cost would. For both these reasons, behavior has an (S,s) band structure, in which adjustment occurs only if the current price is sufficiently far from the optimum.

Summarizing our main findings, our model is consistent with several “puzzling” stylized facts from micro price adjustment data. It 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 over the first few months, and then remains essentially flat (Nakamura

¹Note, however, that since firms are infinitesimal, there are no strategic considerations in their price-setting decisions.

²In order to distinguish between the effects of precaution, *per se*, and the effects of costly decisions, *per se*, we also report an alternative specification that allows us to decompose the two. The alternative model assumes directly that prices are distributed according to a logit function, without subtracting off any managerial costs.

and Steinsson, 2008; Klenow and Malin, 2009). In contrast, contrary to what is observed, in the menu cost model the price adjustment hazard is initially strongly upward sloping. Third, we find that the absolute size 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). Fifth, log retail prices vary more than the log of the replacement cost of those retail goods (Eichenbaum, Jaimovich, and Rebelo, 2011). Sixth, we find that the standard deviation of price adjustments has little relation with the inflation rate, consistent with Gagnon (2009), whereas in the fixed menu cost model the fraction of price increases converges rapidly to 100%, causing the standard deviation of price adjustments to collapse. 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 study the effects of money supply shocks, using numerical methods for calculating dynamic general equilibrium in heterogeneous agent models. Given the degree of rationality that best fits microdata, the effect of money shocks on consumption in our model is similar to that in the Golosov-Lucas (2007) fixed menu cost setup. The impact on consumption is much weaker than it would be in the Calvo model because of a strong “selection effect”. That is, following a positive money growth shock there is big change in the distribution of adjusting firms: some just outside the upper (S,s) band that were contemplating a large price decrease are discouraged from adjusting, while others just inside the lower (S,s) band cross into the action region and choose a large price increase on average, making the aggregate price level quite flexible. Thus, a model of error-prone price adjustment fits microdata better than a fixed menu cost model, but implies that the macroeconomy is relatively close to monetary neutrality.

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 2008). Golosov and Lucas (2007) were the first to calculate the equilibrium effect of a money supply shock in a model that included large, persistent idiosyncratic shocks, which could thus be compared both to micro and macro facts in a quantitatively serious way.

However, the menu cost model presents some problems when confronted with microdata. To begin with, it seems contradicted by evidence that large price changes coexist with some very small ones; Klenow and Kryvtsov (2008) show 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; references include Nakamura and Steinsson (2008), Klenow and Malin (2009), and Campbell and Eden (2010). Recent work based on menu costs or related frameworks to address the details of these microdata include the stochastic menu cost model of Dotsey, King, and Wolman (2009); the model of menu costs and observation

costs of Álvarez, Lippi, and Paciello (2011); and the study of Argentinian data by Álvarez, Gonzalez-Rozada, Neumeyer, and Beraja (2011).

As an alternative, 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 framework for modeling bounded rationality, logit equilibrium, has been widely applied in experimental game theory. 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 applied in other areas of economics; we are unaware of any application of logit equilibrium inside a dynamic general equilibrium macroeconomic model.³ 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. But 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 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 issue 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 2011B) how the resulting distributional dynamics can be tractably characterized in general equilibrium. 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 easier to compute than the fully rational case. In contrast, this method and other state-of-the-art alternatives *cannot* be used to solve the dynamic general equilibrium of “rational inattention” models in the class proposed by Sims (2003), which have a much less tractable state space. We therefore find that logit equilibrium opens the door to tractable models that can be used to study both macroeconomic and microeconomic data.

Recent related papers from the “rational inattention” literature include Woodford (2009), Matejka (2010), and Matejka and McKay (2013), among others. What our framework shares with this literature is the presence of an entropy cost or an entropy constraint on the decision-maker. Under rational inattention, the cost is associated with the flow of information through the decision process (to be precise, costs depend on the mutual information between the exogenous shocks and the endogenous decisions). In contrast, our model is not explicitly based on Shannon’s information theory; we simply assume that precise decisions are costly. We impose a cost on the precision of each choice that is taken, and we use entropy as a convenient measure of precision. Thus, our setup might be thought of 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 then faces costs of *using* that information. The information-theoretic approach takes into account the costs of *receiving* that information too.

Intuitively, we agree that all stages of information processing are likely to imply important costs. The only reason we ignore the cost of *receiving* information is that this dramatically

³The 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.

reduces the dimension of the calculations required to solve our model. Since the rational inattention approach assumes the firm acts under uncertainty, it implies that the firm conditions on a prior over its possible productivity levels. In our model, the firm conditions on its true productivity level, but nonetheless finds it costly to set exactly the right price. The fact that a distribution of productivities is a much lower-dimensional object than a *distribution over priors* greatly simplifies computation in our setup. Moreover, once one knows that entropy reduction costs imply logit, one can simply impose a logit function directly rather than solving numerically for the form of the error distribution. These facts make DSGE modeling based on our approach feasible with contemporary numerical methods, whereas general equilibrium solution methods do not currently exist for rational inattention models (except in very special cases).

Given the similarity between our approach and that of the rational inattention literature, it is likely that the two may 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.

2 Sticky prices in partial equilibrium

In this section we describe the partial equilibrium decision of a monopolistically competitive firm that typically makes small errors when it adjusts its price. We show how a multinomial logit governing price probabilities can be derived from a cost function for error avoidance. Subsection 2.2 compares our setup with the Calvo and menu cost frameworks. We postpone discussion of general equilibrium until Section 3.

2.1 Deriving logit choice from a control cost function

There exists a continuum of monopolistically competitive firms, indexed by i . 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 $A_{it} \equiv \exp(a_{it})$:

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

The idiosyncratic shocks a_{it} are given by a time-invariant Markov process on a bounded set: $a_{it} \in \Gamma^a \subset [\underline{a}, \bar{a}]$, and they are *iid* across firms. Thus a_{it} is correlated with $a_{i,t-1}$, but it is uncorrelated with other firms' shocks.

Let Ω_t be the aggregate state of the economy, which we will describe later when we discuss general equilibrium. Let $P(\Omega_t)$ be an aggregate price index. Each firm is a monopolistic competitor that chooses its nominal price P_{it} taking into account the demand curve $Y_{it} = \vartheta(\Omega_t)P_{it}^{-\epsilon}$, where $\vartheta(\Omega_t)/P(\Omega_t)$ represents real aggregate demand, and ϵ is the elasticity of substitution among product varieties. The firm's nominal price is its only control variable; this implies that it must fulfill all demand at the price it sets. The firm hires labor in a competitive market at the nominal wage rate $W(\Omega_t)$, so the real labor costs associated with its production are $W(\Omega_t)N_{it}/P(\Omega_t)$. Period t nominal profits are $(P_{it} - W(\Omega_t)/A_{it})\vartheta(\Omega_t)P_{it}^{-\epsilon}$.

In addition to these production costs, we assume that a firm incurs a cost, representing managerial time, whenever it resets its nominal price. Since we wish to allow for the possibility that firms may make errors, we think of the managerial decision process as defining a probability distribution across prices, instead of picking out a price deterministically. We assume that

avoiding errors is costly: that is, the management cost increases with the precision of the choice. To define the cost function, we first define the log real price of firm i as

$$p_{it} \equiv \log(P_{it}) - \log(P(\Omega_t)), \quad (1)$$

and we assume the manager allocates probabilities $\pi(p)$ to log real prices lying in a bounded set $p \in \Gamma^p \subseteq [\underline{p}, \bar{p}]$. We assume a grid sufficiently wide so that the real prices preferred at the minimal and maximal values of productivity, \underline{a} and \bar{a} , lie strictly inside $[\underline{p}, \bar{p}]$. It is not necessary at this point to specify whether Γ^p is a discrete or continuous set; in the former case, $\pi(p)$ should be interpreted as a set of discrete probabilities summing to one over the elements of Γ^p ; and in the latter, it should be interpreted as a probability density function on Γ^p .

Now, following Stahl (1990) and Mattsson and Weibull (2002), we define the cost function in terms of entropy. In particular, we assume that the cost of more precise choice is proportional to the reduction in the entropy of the choice variable, normalizing the cost of a perfectly random decision (a uniform distribution) to zero.⁴ That is, the time cost of choosing the probabilities π is given by $\kappa \mathcal{D}(\pi||u)$, where \mathcal{D} represents *Kullback-Leibler divergence*, also known as *relative entropy*. In the cost function, κ is a constant, representing the marginal cost of entropy reduction in units of labor time, and u represents a uniform distribution defined over the same support Γ^p as π . This cost function is nonnegative and convex, and is well-defined either for a discrete or continuous support Γ^p .⁵ The upper bound on this cost function is associated with any distribution that places all probability on a single price $p \in \Gamma^p$. The lower bound on costs is zero, associated with a uniform distribution.

Given this cost function, let us now consider the distribution of prices that would be chosen by the firm. Let $V(a, p, \Omega)$ indicate the nominal expected present discounted value of a firm that produces with log productivity a to sell at log real price p when the aggregate state is Ω .⁶ Now recall that we assume the manager allocates probabilities to maximize firm value, subject to the costs of management.⁷ Therefore, the optimal price distribution $\pi(p|a, \Omega)$ for a firm with log productivity a when the aggregate state is Ω solves the following maximization problem:

$$\tilde{V}(a, \Omega) = \max_{\pi(p): p \in \Gamma^p} \int_{\underline{p}}^{\bar{p}} V(p, a, \Omega) \pi(p) dp - \kappa W(\Omega) \int_{\underline{p}}^{\bar{p}} \pi(p) \log \frac{\pi(p)}{u(p)} dp \quad \text{s.t.} \quad \int_{\underline{p}}^{\bar{p}} \pi(p) dp = 1, \quad (2)$$

where we have converted computational costs to nominal terms, by multiplying by the wage. For each $p \in \Gamma^p$, $\pi(p|a, \Omega)$ solves the following first-order condition:

$$V(p, a, \Omega) - \kappa W(\Omega) \left(1 + \log \frac{\pi(p|a, \Omega)}{u(p)} \right) - \mu = 0,$$

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

$$\frac{\pi(p|a, \Omega)}{u(p)} = \exp \left(\frac{V(p, a, \Omega)}{\kappa W(\Omega)} - 1 - \frac{\mu}{\kappa W(\Omega)} \right). \quad (3)$$

⁴See also Marsili (1999), Baron *et al.* (2002), and Matejka and McKay (2013).

⁵Given two distributions π and τ with support on a bounded set $\mathcal{S} \subseteq [\underline{s}, \bar{s}]$, the Kullback-Leibler divergence of π with respect to τ is defined as $\mathcal{D}(\pi||\tau) \equiv \int_{\underline{s}}^{\bar{s}} \pi(s) \log \left(\frac{\pi(s)}{\tau(s)} \right) ds$. Note that we have not specified whether the support \mathcal{S} is a discrete or continuous set; in the former case, the integral with respect to s should be interpreted as a sum over a set of discrete points.

⁶We will see that it is convenient to write the value function in terms of the firm's *real* price p_{it} , but we could equivalently define it in terms of the nominal price P_{it} , since $P_{it} = \exp(p_{it})P(\Omega)$.

⁷We emphasize that assuming managers allocate probabilities across different prices is an "as if" assumption; it is not intended as a *literal* description of the firm's decision process. Instead, we regard this as a simple but disciplined way of modeling the fact that no decision maker can completely avoid errors.

Here $u(p) \equiv \bar{u}$ represents the density of a uniform, so it is a constant. Now, since the density π integrates to one, we have $\exp\left(1 + \frac{\mu}{\kappa W(\Omega)}\right) = \bar{u} \int_{\underline{p}}^{\bar{p}} \exp\left(\frac{V(p,a,\Omega)}{\kappa W(\Omega)}\right) dp$. Therefore (3) implies that the optimal probabilities are given by a logit formula:

$$\pi(p|a, \Omega) \equiv \frac{\exp\left(\frac{V(p,a,\Omega)}{\kappa W(\Omega)}\right)}{\int_{\underline{p}}^{\bar{p}} \exp\left(\frac{V(p',a,\Omega)}{\kappa W(\Omega)}\right) dp'}. \quad (4)$$

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.

We can now plug these probabilities into the objective function to solve for the value function analytically.⁸ Using (3) to obtain $V(p, a, \Omega) - \kappa W(\Omega) \log \frac{\pi(p|a,\Omega)}{u(p)} = \mu + \kappa W(\Omega)$, we have

$$\tilde{V}(a, \Omega) = \kappa W(\Omega) \log \left(\bar{u} \int_{\underline{p}}^{\bar{p}} \exp\left(\frac{V(p', a, \Omega)}{\kappa W(\Omega)}\right) dp' \right). \quad (5)$$

Here, $\tilde{V}(a, \Omega)$ represents the value of a firm that is choosing a new price; it includes the cost of management time devoted to the price decision. Since changing prices is costly and risky, it is not always optimal to adjust the current price. The expected gain from adjusting the price, conditional on the current log real price p , can be defined as

$$D(p, a, \Omega) \equiv \tilde{V}(a, \Omega) - V(p, a, \Omega). \quad (6)$$

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

$$\lambda(D(p, a, \Omega)) = \mathbf{1}(D(p, a, \Omega) \geq 0), \quad (7)$$

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 real price p . When considering the firm's value now, we must take into account the fact that it may choose to adjust its price next period. At that time, it will face a new log productivity level a' and a new aggregate state Ω' , and if it does *not* adjust its nominal price, its real log price will be altered by inflation, to $p' \equiv p - \log \frac{P(\Omega')}{P(\Omega)}$. Therefore, if the firm does *not* adjust, its value will be $V\left(p - \log \frac{P(\Omega')}{P(\Omega)}, a', \Omega'\right)$. It will adjust its nominal price with probability $\lambda\left(D\left(p - \log \frac{P(\Omega')}{P(\Omega)}, a', \Omega'\right)\right)$, and if it does so, its value will increase by $D\left(p - \log \frac{P(\Omega')}{P(\Omega)}, a', \Omega'\right)$. Therefore, we can write the firm's Bellman equation as follows:

$$\begin{aligned} V(p, a, \Omega) = & \left(P(\Omega) \exp(p) - \frac{W(\Omega)}{\exp(a)} \right) \vartheta(\Omega) (P(\Omega) \exp(p))^{-\epsilon} \\ & + E \left\{ Q(\Omega, \Omega') \left[V\left(p - \log \frac{P(\Omega')}{P(\Omega)}, a', \Omega'\right) + G\left(p - \log \frac{P(\Omega')}{P(\Omega)}, a', \Omega'\right) \right] \middle| a, \Omega \right\}, \end{aligned} \quad (8)$$

where $Q(\Omega, \Omega')$ is the stochastic discount factor, and G represents the expected gains from adjustment:

$$G(p', a', \Omega') \equiv \lambda(D(p', a', \Omega')) D(p', a', \Omega'). \quad (9)$$

⁸Anderson, de Palma, and Thisse (1992) solve a consumption allocation problem for a representative consumer with a utility function that includes an entropy term which can be interpreted as a desire for variety. Their problem is mathematically equivalent to the problem solved here, and they state an analytical solution of the form (5).

2.2 Alternative models of sticky prices

In our simulations below, we will report two specifications of our model. Our main specification, which we will label “ENT”, is defined by (2), (6), (7), (8), and (9), so logit choice is interpreted as optimal decision-making under entropy reduction costs. But we also consider a second specification, abbreviated as “PPS”, for “precautionary price stickiness”. This second specification instead imposes

$$D(p, a, \Omega) \equiv \int_p^{\bar{p}} V(p', a, \Omega) \pi(p'|a, \Omega) dp' - V(p, a, \Omega), \quad (10)$$

where the probabilities π are given by (4). Thus, this version simply assumes that the price distribution takes the form of a logit, without deriving this explicitly from control costs. The full PPS model thus consists of (4), (7), (8), and (9), and (10).

To better interpret our results, we will also compare our framework with two traditional models of nominal rigidity: the Calvo model and the fixed menu cost (FMC) model. Both of these models are consistent with Bellman equation (8) 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 $\lambda(D(p, a, \Omega)) = \bar{\lambda}$, and $D(p, a, \Omega) = V^*(a, \Omega) - V(p, a, \Omega)$, where

$$V^*(a, \Omega) = \max_{p'} V(p', a, \Omega). \quad (11)$$

Therefore in the Calvo model, (9) is replaced by

$$G(p, a, \Omega) \equiv \bar{\lambda} (V^*(a, \Omega) - V(p, a, \Omega)). \quad (12)$$

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 $\lambda(D(p, a, \Omega)) = \mathbf{1}(D(p, a, \Omega) \geq \alpha W(\Omega))$, where $D(p, a, \Omega) = V^*(a, \Omega) - V(p, a, \Omega)$. Therefore we define the FMC model by replacing (9) with

$$G(p, a, \Omega) \equiv \mathbf{1}(D(p, a, \Omega) \geq \alpha W(\Omega)) (D(p, a, \Omega) - \alpha W(\Omega)). \quad (13)$$

2.3 Distributional dynamics

As firms respond to productivity shocks, managing their prices according to (4) and (7), the distribution of prices and productivities evolves over time. We now state the equations governing the dynamics of the distribution.

We will use the notation \tilde{P}_{it} to refer to firm i 's nominal 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 beginning-of-period distribution of prices and productivity, $\tilde{\Phi}_t(\tilde{P}_{it}, a_{it})$, from distribution of prices and productivity at the time of production, $\Phi_t(P_{it}, a_{it})$. Besides keeping track of nominal prices P_{it} , it will also be helpful to track log real prices p_{it} , defined by (1). In analogy to the nominal distributions, we define $\tilde{\Psi}_t(\tilde{p}_{it}, a_{it})$ as the real beginning-of-period distribution, and $\Psi_t(p_{it}, a_{it})$ as the real distribution at the time of production. Finally, we also use lower-case letters to represent the

joint densities associated with these distributions, which we write as $\tilde{\phi}_t(\tilde{P}_{it}, a_{it})$, $\phi_t(P_{it}, a_{it})$, $\tilde{\psi}_t(\tilde{p}_{it}, a_{it})$, and $\psi_t(p_{it}, a_{it})$, respectively.⁹

Two stochastic processes drive the dynamics of the distribution. First, there is the Markov process for firm-specific productivity, which we can write in terms of the following *c.d.f.*:

$$S(a'|a) = \text{prob}(a_{i,t+1} \leq a'|a_{i,t} = a), \quad (14)$$

or in terms of the corresponding density function:

$$s(a'|a) = \frac{\partial}{\partial a'} S(a'|a). \quad (15)$$

Thus, suppose that the density of nominal prices and log productivities at the end of period $t-1$ is $\phi_{t-1}(P, a)$. This density is then affected by productivity shocks; the density at the beginning of t will therefore be

$$\tilde{\phi}_t(P, a') = \int s(a'|a) \phi_{t-1}(P, a) da. \quad (16)$$

But this equation conditions on a given nominal price P . Holding fixed a firm's nominal price, its real log price is changed by inflation, from $p_{i,t-1}$ to $\tilde{p}_{i,t} \equiv p_{i,t-1} - \log(P_t/P_{t-1})$. Therefore the density of real log prices and log productivities at the beginning of t is given by

$$\tilde{\psi}_t\left(p - \log \frac{P_t}{P_{t-1}}, a'\right) = \int s(a'|a) \psi_{t-1}(p, a) da, \quad (17)$$

and hence the cumulative distribution at the beginning of t , in real terms, is

$$\tilde{\Psi}_t(\tilde{p}, a') = \int^{\tilde{p}} \int^{a'} \left(\int s(b|a) \psi_{t-1}\left(q + \log \frac{P_t}{P_{t-1}}, a\right) da \right) db dq. \quad (18)$$

The second stochastic process that determines the dynamics is the process of real price updates, which we have defined in terms of a conditional density of logit form in (4). A firm with real log price \tilde{p} and log productivity a at the beginning of period t adjusts its price with probability $\lambda(D(\tilde{p}, a, \Omega_t))$, and its new real log price is distributed according to $\pi(p|a, \Omega_t)$. Therefore, if the density of firms at the beginning of t is $\tilde{\psi}_t(\tilde{p}, a)$, the density at the end of t is given by

$$\psi_t(p, a) = (1 - \lambda(D(p, a, \Omega_t))) \tilde{\psi}_t(p, a) + \int \lambda(D(\tilde{p}, a, \Omega_t)) \pi(p|a, \Omega_t) \tilde{\psi}_t(\tilde{p}, a) d\tilde{p}. \quad (19)$$

The cumulative distribution at the end of t is simply given by integrating up this density:

$$\Psi_t(p, a) = \int^p \int^a \psi_t(q, b) db dq. \quad (20)$$

3 General equilibrium

In order to be able to perform monetary policy experiments, 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.

⁹Our notation in this section assumes that all densities are well-defined on a continuous support, but we do not actually impose this assumption on the model. With slightly more sophisticated notation we could allow explicitly for distributions with mass points, or with discrete support.

3.1 Households

The household's period utility function is

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

where $u' > 0$, $u'' < 0$, $x' > 0$, $x'' \geq 0$, $v' > 0$, $v'' < 0$. 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

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

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^M represents lump sum transfers received from the monetary authority, and T_t^D 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 following price index:

$$P_t \equiv \left\{ \int_0^1 P_{it}^{1-\epsilon} di \right\}^{\frac{1}{1-\epsilon}}. \quad (21)$$

Optimal household labor supply implies the following first-order condition:

$$W_t u'(C_t) = P_t x'(N_t). \quad (22)$$

Intertemporal consumption is governed by an Euler equation,

$$R_t^{-1} = \beta \frac{P_t u'(C_{t+1})}{P_{t+1} u'(C_t)}, \quad (23)$$

and optimal money demand implies:

$$1 = \frac{v'(M_t/P_t)}{u'(C_t)} + \beta \frac{P_t u'(C_{t+1})}{P_{t+1} u'(C_t)}. \quad (24)$$

3.2 Monetary policy and aggregate consistency

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

$$M_t = \mu_t M_{t-1}, \quad (25)$$

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

$$z_t = \phi_z z_{t-1} + \epsilon_t^z. \quad (26)$$

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

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

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

$$Y_{it} = A_{it}N_{it} = C_{it} = P_t^\epsilon C_t P_{it}^{-\epsilon}. \quad (27)$$

Also, total labor supply must equal total labor demand:

$$N_t = \int_0^1 \frac{C_{it}}{A_{it}} di + K_t^\pi = P_t^\epsilon C_t \int_0^1 P_{it}^{-\epsilon} A_{it}^{-1} di + K_t^\pi \equiv \Delta_t C_t + K_t^\pi, \quad (28)$$

where K_t^π is the total time devoted to choosing which price to set by adjusting firms:

$$K_t^\pi = \kappa \int \int \lambda(D(\log \tilde{P} - \log P_t, a, \Omega_t)) \left(\int_{\underline{p}}^{\bar{p}} \pi(p|a, \Omega_t) \log \frac{\pi(p|a, \Omega_t)}{\bar{u}} dp \right) \tilde{\phi}_t(\tilde{P}, a) da d\tilde{P}. \quad (29)$$

The labor market clearing condition (28) also defines a weighted measure of price dispersion:

$$\Delta_t \equiv P_t^\epsilon \int_0^1 P_{it}^{-\epsilon} A_{it}^{-1} di, \quad (30)$$

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 (8) consistent with the goods market clearing conditions (27), the equilibrium demand shift factor must be

$$\vartheta(\Omega_t) = C(\Omega_t)P(\Omega_t)^\epsilon. \quad (31)$$

Finally, we assume that the representative household owns the firms. For consistency with the household's Euler equation, the stochastic discount factor in the firm's problem must be

$$Q(\Omega_t, \Omega_{t+1}) = \beta \frac{P(\Omega_t)u'(C(\Omega_{t+1}))}{P(\Omega_{t+1})u'(C(\Omega_t))}. \quad (32)$$

3.3 State variable

We now 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 we 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 $AR(1)$, the latest deviation in the money growth rate, 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, these are state variables too. Thus, the situation at the beginning of t depends on the joint distribution $\tilde{\Phi}_t(\tilde{P}_{it}, A_{it})$ of prices and productivity shocks, prior to adjustment. Therefore, we can define the time t state as $\Omega_t \equiv (M_t, z_t, \tilde{\Phi}_t)$.

This definition of the state variable is in nominal terms. But we can detrend the model, rewriting it in real terms, by deflating all prices by the nominal price level $P(\Omega_t)$. However, $P(\Omega_t)$ is a jump variable at t ; as we can see from (21) and (27), it aggregates prices at the end of the period, after adjustments have taken place:

$$P(\Omega_t) \equiv \left\{ \int \int P^{1-\epsilon} \phi_t(P, a) dP da \right\}^{\frac{1}{1-\epsilon}}. \quad (33)$$

But if $P(\Omega_t)$ is a jump variable, then so is the beginning-of-period real price $\tilde{p}_{it} \equiv \log \left(\tilde{P}_{it} / P(\Omega_t) \right)$, and the beginning-of-period real distribution $\tilde{\Psi}_t(\tilde{p}_{it}, a_{it})$ is too. Therefore we cannot define the real state of the economy at the beginning of t in terms of the distribution $\tilde{\Psi}_t$.

Instead, to define equilibrium in real terms, we must condition on a real state variable that is predetermined at the beginning of the period. Therefore, we define the time t real state as $\Xi_t \equiv (z_t, \Psi_{t-1})$, where Ψ_{t-1} is the distribution of *lagged* end-of-period prices and productivities. We will now show that the distribution Ψ_{t-1} , together with the shocks z_t , suffices to determine all real quantities at time t . In particular, it will determine the distributions $\tilde{\Psi}_t(\tilde{p}_{it}, a_{it})$ and $\Psi_t(p_{it}, a_{it})$. Therefore Ξ_t is a correct real state variable for time t .

3.4 Equilibrium definition

To show that we can define a general equilibrium for this economy in terms of the real state Ξ_t , we will rewrite the equations in terms of inflation, defined as

$$i(\Xi_t, \Xi_{t-1}) \equiv \log \left(\frac{P(\Omega_t)}{P(\Omega_{t-1})} \right). \quad (34)$$

Substituting inflation in place the nominal price level $P(\Omega)$ will allow us to eliminate all reference to nominal variables in the model.

Thus, we define the real value function v , deflating the nominal value function by the current price level. That is,

$$v(p, a, \Xi_t) \equiv \frac{V(p, a, \Omega_t)}{P(\Omega_t)}$$

Deflating the value function and related functions in this way, and defining the real wage $w(\Xi_t) \equiv W(\Omega_t) / P(\Omega_t)$, we can rewrite the Bellman equation in detrended form:

$$v(p, a, \Xi_t) = \left(\exp(p) - \frac{w(\Xi_t)}{\exp(a)} \right) \frac{C(\Xi_t)}{\exp(\epsilon p)} \quad (35)$$

$$+ \beta E \left\{ \frac{u'(C(\Xi_{t+1}))}{u'(C(\Xi_t))} \left[v(p - i(\Xi_t, \Xi_{t+1}), a', \Xi_{t+1}) + g(p - i(\Xi_t, \Xi_{t+1}), a', \Xi_{t+1}) \right] \middle| a, \Xi_t \right\},$$

where E represents an expectation over a' and Ξ_{t+1} , conditional on a and Ξ . This equation includes the real expected gains function g , which we define as follows for any beginning-of-period log real price \tilde{p} :

$$g(\tilde{p}, a, \Xi) \equiv \mathbf{1}(d(\tilde{p}, a, \Xi) \geq 0) d(\tilde{p}, a, \Xi), \quad (36)$$

$$d(\tilde{p}, a, \Xi) \equiv \tilde{v}(a, \Xi) - v(\tilde{p}, a, \Xi), \quad (37)$$

$$\tilde{v}(a, \Xi) \equiv \kappa w(\Xi) \log \left(\bar{u} \int_{\underline{p}}^{\bar{p}} \exp \left(\frac{v(p', a, \Xi)}{\kappa w(\Xi)} \right) dp' \right). \quad (38)$$

The density of new real prices associated with these value functions is

$$\pi(p|a, \Xi) \equiv \frac{\exp \left(\frac{v(p, a, \Xi)}{\kappa w(\Xi)} \right)}{\int_{\underline{p}}^{\bar{p}} \exp \left(\frac{v(p', a, \Xi)}{\kappa w(\Xi)} \right) dp'} = \bar{u} \frac{\exp \left(\frac{v(p, a, \Xi)}{\kappa w(\Xi)} \right)}{\exp \left(\frac{\tilde{v}(a, \Xi)}{\kappa w(\Xi)} \right)} \quad (39)$$

The probability densities then evolve as follows:

$$\tilde{\psi}_t(\tilde{p}, a') = \int s(a'|a) \psi_{t-1}(\tilde{p} + i(\Xi_t, \Xi_{t-1}), a) da \quad (40)$$

$$\psi_t(p, a) = \mathbf{1}(d(p, a, \Xi_t) < 0) \tilde{\psi}_t(p, a) + \int \mathbf{1}(d(\tilde{p}, a, \Xi_t) \geq 0) \pi(p|a, \Xi_t) \tilde{\psi}_t(\tilde{p}, a) d\tilde{p}. \quad (41)$$

These densities can be integrated up to give the following cumulative distributions:

$$\tilde{\Psi}_t(\tilde{p}, a) = \int^{\tilde{p}} \int^a \tilde{\psi}_t(q, b) db dq, \quad (42)$$

$$\Psi_t(p, a) = \int^p \int^a \psi_t(q, b) db dq. \quad (43)$$

Also, aggregate variables must satisfy the representative household's first-order conditions. For simplicity, from here on we will focus on the special case of linear labor disutility, $x(N) = \chi N$, which allows us to calculate the equilibrium without actually solving for N .¹⁰ In this case, the first-order condition on labor supply is

$$w(\Xi_t) u'(C(\Xi_t)) = \chi. \quad (44)$$

The Euler equation for intertemporal consumption and the money demand equation can be combined (by eliminating the nominal interest rate) to give:

$$1 - \frac{v'(m(\Xi_t))}{u'(C(\Xi_t))} = \beta E \left(i(\Xi_t, \Xi_{t+1}) \frac{u'(C(\Xi_{t+1}))}{u'(C(\Xi_t))} \middle| \Xi_t \right). \quad (45)$$

where $m(\Xi_t) \equiv M_t/P(\Omega_t)$ is the real money supply. Its growth rate must be consistent with the growth rate of the nominal money supply, and inflation, which implies:

$$\frac{\mu \exp(z_t)}{\exp(i(\Xi_{t-1}, \Xi_t))} = \frac{m(\Xi_t)}{m(\Xi_{t-1})}. \quad (46)$$

Finally, the distribution of real prices must be consistent at all times with the definition of the aggregate price level, (21), which implies the following identity:

$$\int \int \exp((1 - \epsilon)p) \psi_t(p, a) da dp = 1. \quad (47)$$

Note that equations (35)-(46) serve to determine the value functions v , g , d , and \tilde{v} , the densities and cumulative distributions π , $\tilde{\psi}$, ψ , $\tilde{\Psi}$, and Ψ , and the functions w , C , and m . Finally, given the last equation (47), we have one more condition to determine the function i . Therefore the number of unknowns matches the number of equations.

We can now define equilibrium. We will use the notation Ξ_t to denote the pair (z_t, Ψ_{t-1}) , where z_t is a scalar and Ψ_{t-1} is a cumulative distribution function on the set $[\underline{p}, \bar{p}] \times [\underline{a}, \bar{a}]$.

¹⁰For the general case of a nonlinear $x(N)$, the equilibrium definition would require three more equations, corresponding to (28), (29), and (30), to determine labor $N(\Xi)$, decision costs $K^\pi(\Xi)$, and price dispersion $\Delta(\Xi)$.

Definition. A *real general equilibrium* consists of value functions $v(p, a, \Xi)$, $\tilde{v}(a, \Xi)$, $g(p, a, \Xi)$, and $d(p, a, \Xi)$, a price policy function $\pi(p|a, \Xi)$, consumption demand and money demand functions $C(\Xi)$ and $m(\Xi)$, and wage and inflation functions $w(\Xi)$ and $i(\Xi, \Xi')$ that satisfy the following conditions.

- (1.) Taking as given the functions w , C , and i , let a_t be governed by the Markov process (14)-(15), and let $\Xi_t \equiv (z_t, \Psi_{t-1})$ be governed by (26) and (40)-(43). Then the value functions v , g , d , and \tilde{v} satisfy (35), (36), (37), and (38).
- (2.) Taking as given the wage function w and the value function v , the price policy π is determined by (39).
- (3.) Let $\Xi_t \equiv (z_t, \Psi_{t-1})$ be governed by (26) and (40)-(43). Then the functions w , C , m , and i satisfy (44), (45), (46), and (47).

Note that computing equilibrium requires us to track the evolution of the distribution Ψ of prices and productivities across firms. We compute the model following the algorithm of Reiter (2009), as described in the appendix.

4 Results

4.1 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 (2011A, 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 (2011). Given our grid-based approximation, this implies a maximum absolute log price change of 0.48, which covers, with an extra margin of 7%, the maximum absolute log price change of 0.45 observed in the AC Nielsen dataset.¹¹

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

¹¹The 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.

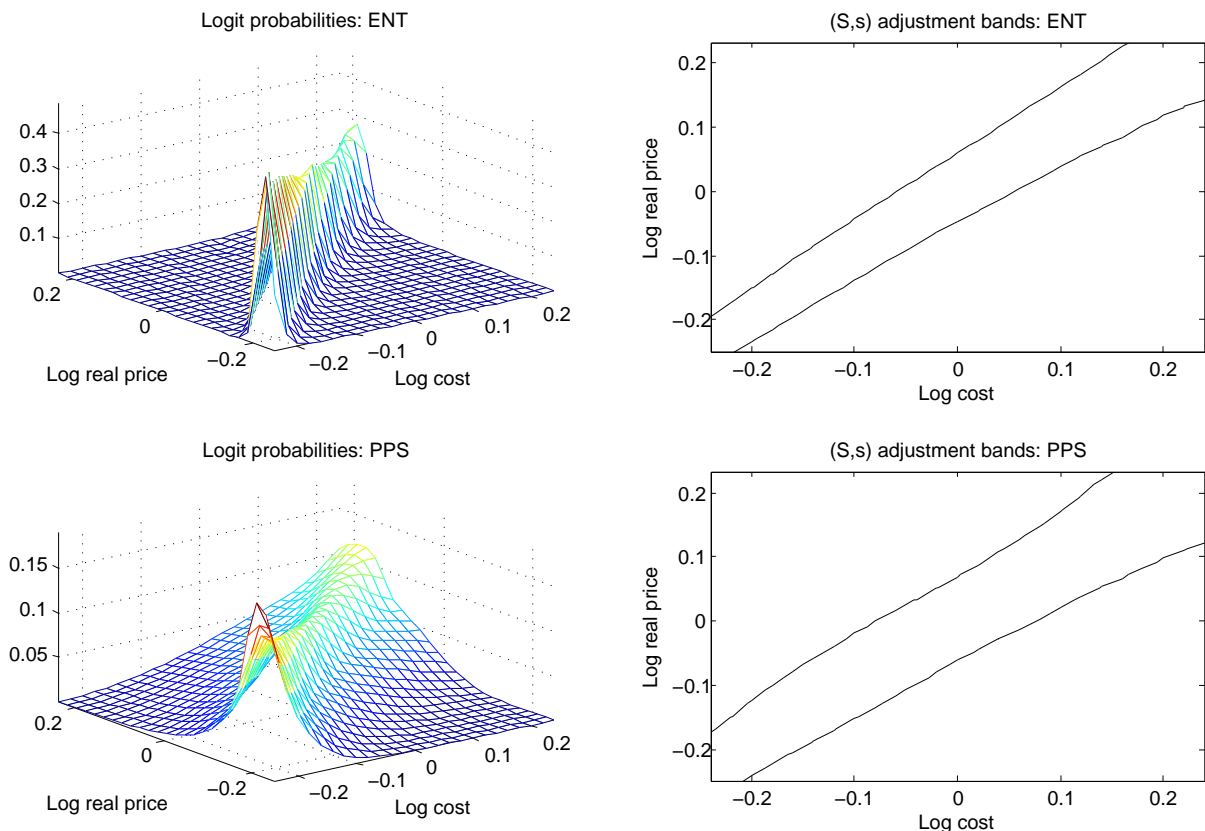


Fig. 1. Pricing and timing policy functions: ENT and PPS

Note: Left panels show the logit probabilities π of choosing different prices, conditional on costs. Right panels show the boundaries of the inaction region associated with adjustment indicator λ . Top panels: ENT specification. Bottom panels: PPS specification.

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).

4.2 Steady state microeconomic results

Figure 1 illustrates the two main decisions a firm must make: whether or not to adjust its current price, and which new price to set, conditional on adjustment. The top panels refer to the ENT specification. The axes of each panel in the figure show the logarithm of the firm's cost ($1/A_{it}$) and the log of its real price; both variables are centered around their nonstochastic steady-state values, so that the origin of each graph represents productivity $A_{it} = 1$ and real price $P_{it}/P_t = 1$. The top left panel illustrates the distribution π of prices, conditional on productivity, when a firm chooses to adjust (the vertical dimension in the graph represents probability). By assumption, more precise price decisions are more costly; indeed, under the assumed cost function perfect precision (allocating probability one to one particular price, and

zero to all others) is infinitely costly on the margin. Therefore, firms set prices with some randomness, according to the logit formula (4). Firms with higher costs (lower a_{it}) are better off setting higher prices, so the peak of the probability distribution roughly follows the 45° line in the graph. But random errors generate a bell-shaped distribution around this modal choice, as would be seen by slicing through the distribution along some constant value of costs.

The top right panel of Fig. 1 illustrates the decision whether or not to adjust the current price, under the ENT specification, as summarized by the (S,s) adjustment bands. The bands delimit the region of inaction that surrounds the 45° line: when its cost and price lie within the bands, a firm leaves its current price unchanged; outside the bands, it adjusts. There are two reasons why a firm might prefer to leave its current price unchanged. First, choosing a new price is costly; a firm that chooses to readjust knows that it must devote some time to choosing the right price, and that it must devote more time if it desires greater accuracy. This decision cost gives an incentive not to adjust the price every period, just as a traditional “menu cost” would. Second, the fact that decisions are error-prone means they are risky: when a firm chooses to readjust, it risks setting the wrong price, implying a loss of profits. For both these reasons, a firm is better off leaving its price unchanged if it is currently sufficiently close to the optimum. The inaction region gets wider as cost increases, since higher costs are associated with higher prices and therefore lower sales. Lower sales make profits less sensitive to the price, implying a wider inaction band.

The bottom two panels of Figure 1 illustrate the distribution of prices and the adjustment decision under the PPS specification. The only difference, relative to the baseline setup shown in the top panels, is that the conditional price distribution π is imposed exogenously, instead of being derived from a problem of decision cost minimization. Under the PPS framework, the logit price distribution (4) is motivated by its two desirable properties: (a) more profitable prices are more likely, and (b) frictionlessly optimal behavior is nested as a limiting case. Note that in this case, there is no physical cost or decision cost associated with price adjustment; nonetheless, as in the ENT specification, the adjustment decision is characterized by an inaction region. Failure to adjust arises in this case only because of the risk of setting a suboptimal price, which is the aspect of nominal rigidity we have called “precautionary price stickiness”.

The shape of the policy functions in the ENT and PPS specifications is similar. The main difference is quantitative: errors are much larger in the PPS specification than in the ENT specification, reflected in a wider distribution of prices for any level of costs, as can be seen by comparing the two left panels of Fig. 1. This quantitative finding is a result of our calibration strategy, which chooses the noise parameter κ to match the 10% monthly frequency of adjustment in our data. Since the PPS specification models price rigidity *only* on the basis of the cost of potential errors, these errors need to be relatively large in order to ensure that firms adjust only once in ten months. The decision cost in the ENT specification acts an additional brake on adjustment, so the same adjustment frequency can be matched with a much lower level of noise. Thus, as we see in Table 1, the noise parameter drops from $\kappa = 0.0428$ in the PPS model to $\kappa = 0.0050$ in the ENT model.

Table 1 also shows that the errors implied by both specifications are small, on average, in terms of their economic value. The implied losses are shown in the first row of the table: we report the mean revenue loss suffered relative to a fully rational firm, as a percentage of the average revenues of a fully rational firm. In the baseline (ENT) estimate, firms lose a third of a percent of revenues due to imperfect rationality. Losses are of a similar order of magnitude in all the other specifications too; conditional on a 10% adjustment frequency, losses are largest in the Calvo model and smallest under fixed menu costs. The distribution of losses is displayed in

Table 1. Model-simulated statistics and evidence (zero trend inflation)

| | Model PPS | | | ENT | | Calvo | | FMC | | Evidence | | |
|-------------------------------------|-----------------|-------------------|-----------------------------|----------------|-------------------|-----------------------------|------|------|------|----------|------|----|
| | $\kappa=0.0428$ | $2 \times \kappa$ | $\frac{1}{2} \times \kappa$ | $\kappa=0.005$ | $2 \times \kappa$ | $\frac{1}{2} \times \kappa$ | | | | MAC | NS | KK |
| Mean loss (% flex. revenue) | 0.55 | 0.70 | 0.40 | 0.34 | 0.45 | 0.24 | 0.61 | 0.31 | | | | |
| Freq. of price changes | 10 | 7.4 | 12.3 | 10 | 7.6 | 13.3 | 10 | 10 | 20.5 | 10 | 21.9 | |
| Mean absolute price change | 11.9 | 13.5 | 10.1 | 6.5 | 7.9 | 5.4 | 2.8 | 5.5 | 10.4 | | 11.3 | |
| Std of price changes | 14.5 | 16.5 | 12.3 | 7.3 | 9.1 | 6.0 | 3.7 | 5.6 | 13.2 | | | |
| Kurtosis of price changes | 2.6 | 2.5 | 2.7 | 2.3 | 2.5 | 1.9 | 4.2 | 1.2 | 3.5 | | | |
| Percent of price increases | 50 | 50.2 | 49.7 | 50 | 49.9 | 50 | 48 | 50.7 | 50 | 66 | 56 | |
| % of abs price changes $\leq 5\%$ | 19.5 | 17.1 | 23.2 | 32.2 | 23.7 | 44.7 | 83.6 | 42.4 | 25 | | 44 | |
| % of abs price changes $\leq 2.5\%$ | 9.4 | 8.4 | 10.9 | 10.1 | 8.3 | 12.5 | 55 | 0 | 11 | | | |
| Std(p)/Std(a) | 1.04 | 1.01 | 1.05 | 0.97 | 0.95 | 0.97 | 0.71 | 0.94 | | | | |

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).

First 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.

Last row reports ratio of coefficient of variation of real prices to coefficient of variation of productivity.

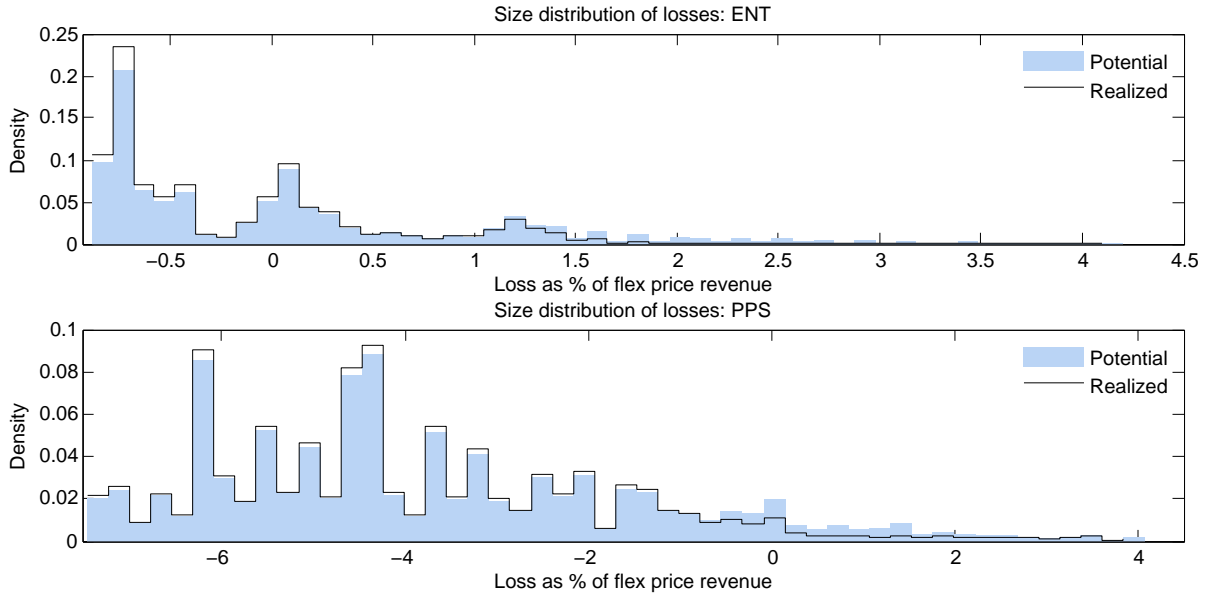


Fig. 2. Size distribution of losses: ENT and PPS

Note: Panels show the distribution of the potential (shaded) and realized (stairs) loss D as a percentage of average monthly revenue of a flexible-price firm. Negative losses represent gains from inaction.

Top panel: ENT specification. Bottom panel: PPS specification.

Figure 2, for the ENT and PPS models, both before (shaded blue area) and after (black line) firms decide whether or not to adjust. Comparing the distributions before and after adjustment, we see that some probability mass is shifted from the thin right tail of the distribution, where losses are positive, to the fat left tail, where losses are negative (that is, firms on the left are strictly better off not adjusting). Prior to adjustment, we see some losses as large as 4% of monthly revenues. After adjustment, almost all losses are below 2% in the ENT specification, and most of the distribution lies between 0% and -1%. In the PPS model, even after adjustment, there are still some 4% losses, but most losses lie between -2% and -6% of monthly revenues. That is, due to the large errors in the PPS calibration, some firms suffer substantial losses even *after* adjustment, but most firms are substantially better off *avoiding* the risk of adjusting.

Next, Table 1 reports price adjustment statistics for the various specifications we compare. By construction, a single free parameter allows us to 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). Thus we report calibrations of ENT, PPS, FMC, and the Calvo model that all match this adjustment frequency. The resulting size distribution of price adjustments is summarized in the middle of the table; the distribution is somewhat too concentrated under the ENT specification, and is more spread out under PPS. This distribution is also illustrated in Figure 3, which shows a histogram of price changes with 49 equally-spaced bins representing log price changes from -0.48 to 0.48. The blue shaded bars

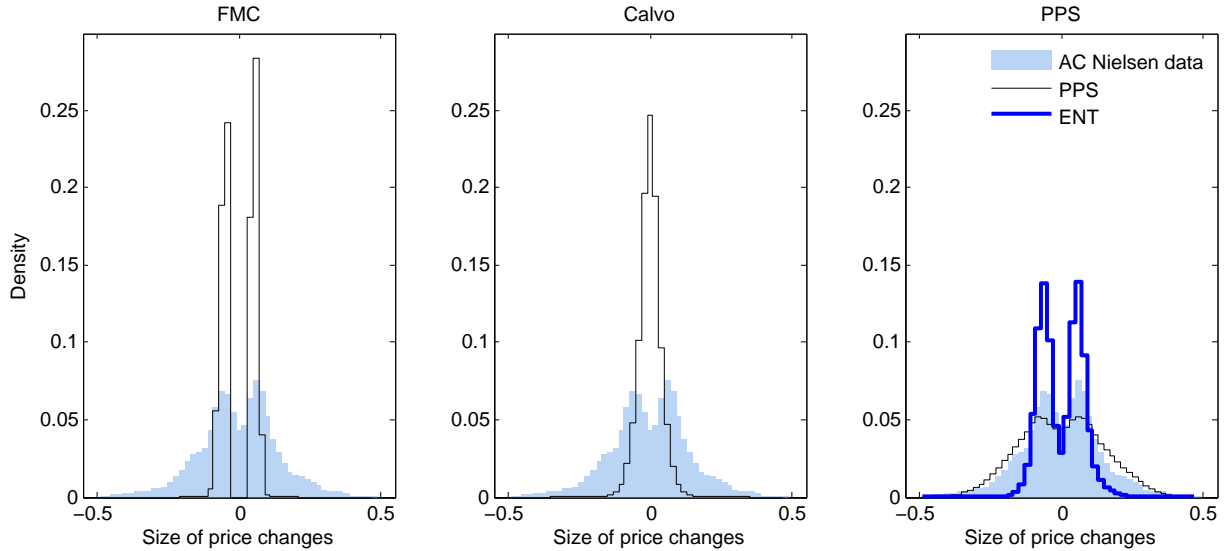


Fig. 3. Distribution of nonzero price adjustments: comparing models

Note: Shaded blue area: nonzero price adjustments in AC Nielsen data of Midrigan (2011), controlling for sales. Black lines: model simulations under fixed menu costs (left), Calvo model (middle), and PPS (right, thin) and ENT (right, thick blue line) specifications of our model.

in the figure represent the AC Nielsen data, and the solid lines represent the results of our calibrated models. The ENT and PPS specifications generate a bimodal distribution, like the empirical distribution, but the standard deviation of price adjustments is only 7.3% in the ENT model, versus 13.2% in the data. The standard deviation doubles under the PPS specification, to 14.5%, allowing this model to match the empirical histogram quite well. Kurtosis is substantially higher in the data (3.5) than it is in our model (2.3 for ENT, 2.6 for PPS), as can be seen from the relatively fat tails in the blue shaded data in the figure. As in the data, half of the price adjustments in the model are price increases.

The histograms in Figure 3 also illustrate the coexistence of many large and small price changes in the data, a pattern which is hard for the FMC model to explain. While many price changes are quite large in the ENT and PPS models, nonetheless twenty to thirty percent of all adjustments are less than 5% in absolute value, and around one tenth of the adjustments are smaller than 2.5%. These numbers are close to those observed in the AC Nielsen data. Conditional on the same productivity process and adjustment frequency, the Calvo and FMC models imply much more concentrated distributions than our model does; this is an obvious consequence of the errors that our setup adds on top of the fundamental price variation. But besides this difference in variability, each specification implies a different shape for the distribution of price changes. The histogram of price adjustments from the FMC model consists of two sharp spikes, representing price increases and decreases by firms concentrated near the lower and upper ad-

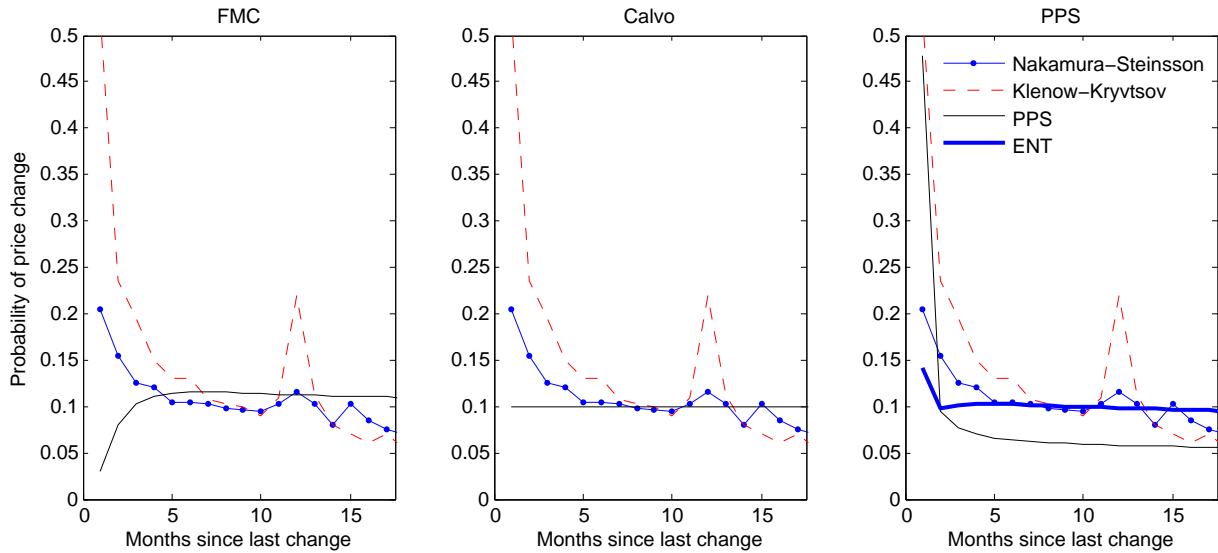


Fig. 4. Price adjustment hazard: comparing models

Note: Price adjustment hazard as function of time since last adjustment. Red dashed line: data of Nakamura and Steinsson (2008); blue line with dots: Klenow and Kryvtsov (2008). Black lines: model simulations under fixed menu costs (left), Calvo model (middle), and PPS (right, thin) and ENT (right, thick blue line) specifications of our model.

justment bands, respectively. Hence there is little variation in the size of the changes; while over 40% of all adjustments are less than 5% in absolute value, there are none at all less than 2.5% in absolute value. In contrast, the Calvo model generates a smoothly unimodal distribution, since there is no relation between the firm's current price and its adjustment hazard. Since the level of noise is very low in our calibration of the ENT model, it generates a distribution similar to that of the FMC model, but smoother. Most price changes observed in this specification are jumps from the (S,s) bands to new prices near the optimum. Thus the modes of the distribution are similar to those of the FMC model; but the distribution is smoother since prices are chosen with error, including some very small adjustments. Given the much higher noise in our calibration of the PPS model, it generates a histogram that is much smoother still, though it also displays some bimodality associated with the (S,s) bands.

Another interesting finding relates to the behavior of the adjustment hazard, that is, the probability of a price change as a function of the time since the last change. Figure 4 illustrates this hazard, comparing two empirical data sets with our model and with the FMC and Calvo specifications. In our model, error-prone decisions imply that firms often readjust quickly after making a change, if their chosen price turns out to have been a mistake. This accounts for the spike in the adjustment hazard at one month. Some firms that have made smaller errors are subsequently hit by productivity shocks that push them across the adjustment thresholds; hence

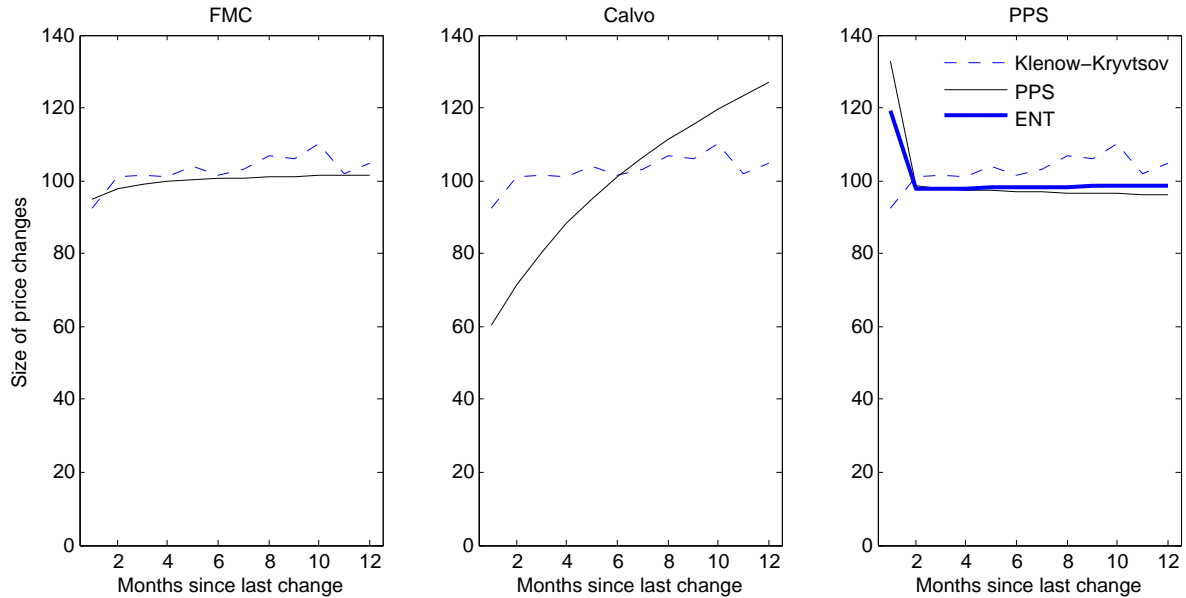


Fig. 5. Mean adjustment and price duration: comparing models

Note: Blue dashed line: average absolute nonzero price change, as function of time since last adjustment, in data of Klenow and Kryvtsov (2008).

Normalized: 100 represents the average price change in the data or the simulation.

Black lines: model simulations under fixed menu costs (left), Calvo model (middle), and PPS (right, thin) and ENT (right, thick) specifications of our model.

the adjustment hazard in the PPS model continues to decrease over the next few months, until it eventually flattens out. This pattern is quite consistent with microdata; it fails by construction in the Calvo model. It also contrasts with the behavior of the FMC model, where the expected distance from the optimal price increases with the time since last adjustment, resulting in an increasing hazard. Many studies have suggested that the 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 line with dots 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.

Figure 5 explores a related stylized fact: in Klenow and Kryvtsov’s (2008) microdata, the average absolute price adjustment is approximately invariant with the time since last price adjustment. The fixed menu cost model matches these data remarkably well. The fact that almost all price adjustments occur close to the (S,s) bands in the FMC model (regardless of the time since last adjustment) means that the size of the change cannot vary much with the time since last adjustment. The PPS and ENT specifications also match this observation fairly well, for the same reason, except at a horizon of one month, where the immediate correction of

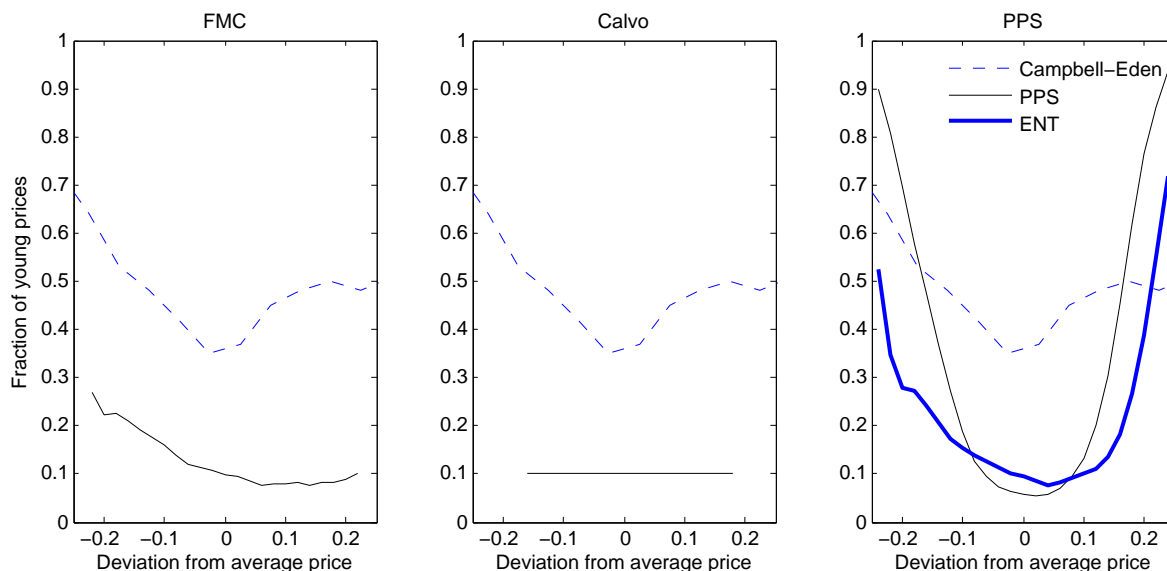


Fig. 6. Extreme prices tend to be young: comparing models

Note: Blue dashed line: fraction of prices set in the previous month, as function of log deviation from product-specific mean price, in data of Campbell and Eden (2010).

Black lines: model simulations under fixed menu costs (left), Calvo model (middle), and PPS (right, thin) and ENT (right, thick blue line) specifications of our model.

recent mistakes implies somewhat larger adjustments. In contrast, the Calvo model is strongly rejected on this score; it implies that the size of the price adjustment increases with the age of the previous price, since the distance between the price and the optimal price tends to increase with the time since the previous adjustment.

Figure 6 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 inconsistent with this observation: at a zero inflation rate, the fraction of young prices conditional on *any* price is exactly equal to the Calvo adjustment parameter.¹² Both the fixed menu cost and ENT 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. In the FMC model, this occurs because exceptional prices always reflect responses to exceptional cost shocks; since costs are mean-reverting, these exceptional prices are likely to be reverted soon thereafter. The intuition in the ENT and PPS models is similar: exceptional prices *either* reflect responses to exceptional cost shocks (and are therefore likely to revert in response to mean reversion of the shocks) *or* represent exceptionally large errors, which the firm is likely to correct soon (as we already

¹²In the Campbell and Eden data, and in our simulations, “young” prices are those set in the previous month. Note that under the Calvo model, in the absence of inflation, the distribution of prices is just a weighted average of past distributions of *new* prices. In steady state, the distribution of *new* prices is always the same, and therefore also equals the distribution of all prices. This property does not hold with nonzero inflation.

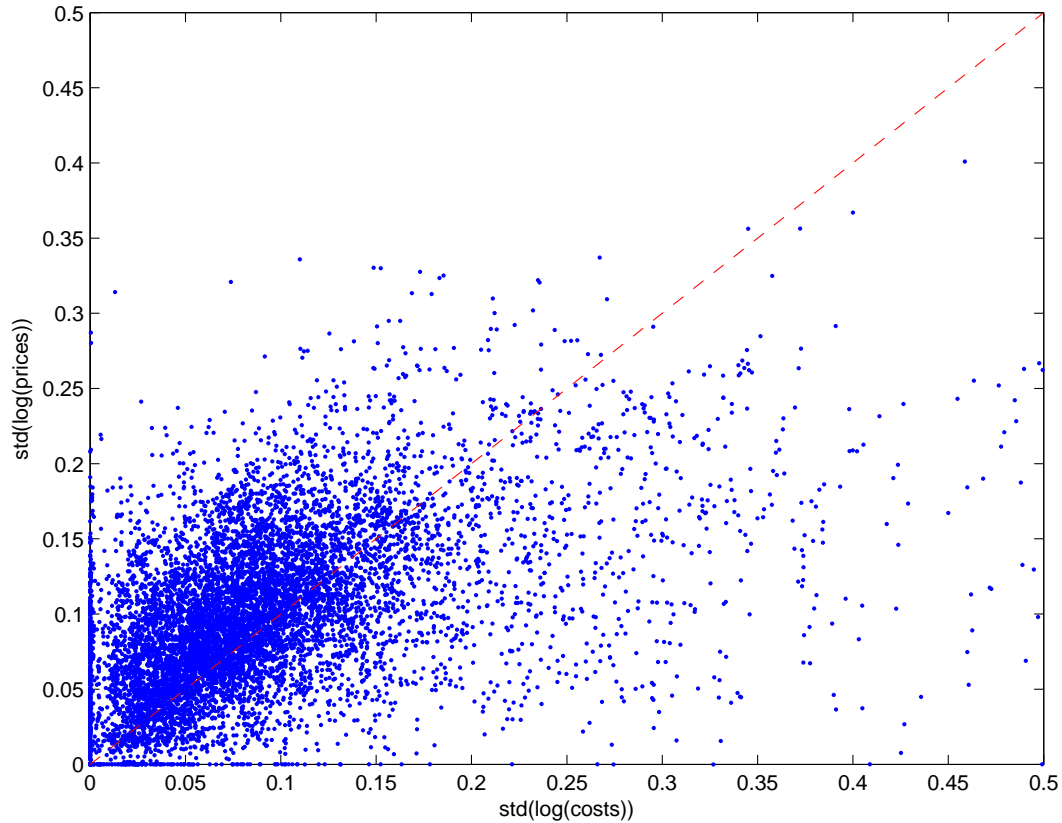


Fig. 7. Coefficients of variation of prices and costs

Note: Standard deviations of log prices and log costs for different retail products, as defined by UPC code.

Data source: Dominick's.

mentioned in the context of Fig. 4).

Another microeconomic fact is addressed in Table 1 and Figure 7. Eichenbaum, Jaimovich and Rebelo (2011) demonstrate, using data on retail prices and restocking prices for the same products at the same retail sellers, that prices are more variable than costs. More precisely, they show that the coefficient of variation of prices is roughly 15% greater than the coefficient of variation of costs, regardless of whether they compare “weekly” prices and costs, or “reference” prices and costs, and regardless of whether or not they condition on adjustment. We illustrate the same fact in Fig. 7, which shows a scatter plot across products from the Dominick's data set, where we have information on costs as well as prices. The coefficient of variation of costs is shown on the horizontal axis, with the coefficient of variation of prices on the vertical axis. There is great heterogeneity across products; for some products either the price or the cost is never observed to change. Nonetheless, the bulk of the data points lie slightly above the 45° line, meaning that for typical products, prices vary more than costs.

In the last line of Table 1 we see that the Calvo model is strongly inconsistent with this observation. Because firms cannot control the timing of their adjustments in the Calvo model, they price very conservatively. When costs are exceptionally high, Calvo firms set a low markup (they choose a price much lower than a firm with flexible prices would set), because they expect costs to revert to the mean, so they worry about getting stuck with an excessively high price. Likewise, Calvo firms set a high markup when costs are unusually low. There is also a tendency to price “conservatively” in the FMC model, but it is much weaker, since under menu costs firms readjust whenever their prices get too far out of line with costs. Thus the ratio of standard deviations of log prices relative to log costs rises from 0.71 in the Calvo model to 0.94 in the FMC model. Price variability increases in the PPS model, for an obvious reason: the presence of pricing errors. Thus in the PPS case the ratio of the volatilities of prices and costs is 1.04. The results for the ENT model are intermediate between FMC and PPS, with a ratio of 0.97. In other words, the extra price volatility induced by the small errors in the ENT specification is insufficient to offset the “conservativeness” induced by nominal rigidity, so the volatility ratio is less than one, but higher than that in the FMC case.

4.3 Effects of changes in monetary policy

Figure 8 investigates the effects of large changes in steady state inflation, compared with empirical observations from Mexican data reported by Gagnon (2009) for periods with annual inflation rates of 5%, 29%, and 63%. The Calvo model is not shown, since attempting to match these data with the Calvo model would be pointless.¹³ 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 that there is little change in the standard deviation, which fluctuates around 11% as inflation rises. The standard deviation is also fairly constant in the ENT and PPS models; it is around 7% and falling slightly in the ENT specification, and around 16% and rising for PPS. In contrast, in the FMC model the standard deviation of price adjustments deviation drops dramatically (by almost two-thirds) as inflation rises to the levels observed in Gagnon’s data. The reason is that as inflation increases, the fraction of price adjustments that are increases converges to almost 100%, as we see in the third panel of the figure. The histogram of price adjustments goes from being bimodal to having a single sharp peak, and the remaining variability of price changes is just 2%. Our error-prone specifications are more consistent with the data on this score, because the fraction of increases rises much more slowly towards 100%, and therefore the standard deviation of price changes collapses more slowly than it does in the FMC model.

Figures 9-11 illustrate the quantitative implications of our model for short-run 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

¹³The problem is not only that the Calvo model fails, by construction, to match the variation in the adjustment rate as inflation rises. In addition, at 29% and 63% inflation rates, our Calvo calibration with adjustment rate $\bar{\lambda} = 0.1$ generates *negative* firm values, even for firms that have just adjusted to their preferred prices.

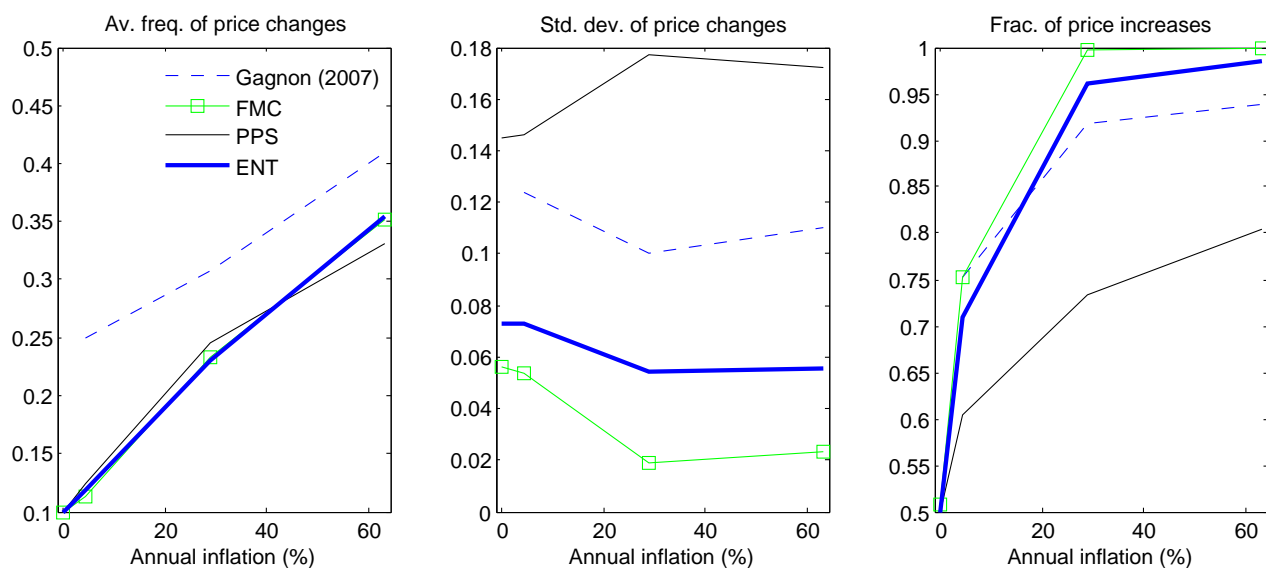


Fig. 8. Effects of trend money growth: comparing models

Note: Blue dashed line: Mexican data of Gagnon (2009), at annual inflation rates of 5%, 29%, and 63%.

Other lines: simulations of FMC (green with squares), PPS (thin black line), ENT (thick blue line).

First panel: average frequency of price adjustments, as function of inflation rate.

Second panel: standard deviation of nonzero price adjustments, as function of inflation rate.

Third panel: fraction of nonzero price adjustments which are increases, as function of inflation rate.

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 overall consumption response is considerably larger because it is more persistent. The same conclusions can be drawn from Table 2, which compares the real effects of money shocks across models following the methodology of Golosov and Lucas (2007). The table reports the percentage of output variation that can be explained by money shocks alone (under the assumption that the money shock suffices 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 imply nearly identical degrees of monetary nonneutrality, while PPS implies real effects roughly twice as strong as ENT. Under the Calvo specification, the real effects of money shocks are almost four times stronger than they are in the PPS case.

The reason for the weak response of consumption in the ENT specification is the same one Golosov and Lucas (2007) identified in the FMC context: a strong “selection” effect. In the FMC model, 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. The ENT model behaves much like the FMC model, with some small errors

Table 2. Variance decomposition and Phillips curves

| <i>Correlated money growth shock</i> ($\phi_z = 0.8$) | Data | Model PPS | ENT | Calvo | FMC |
|------------------------------------------------------------|-----------------------|--------------------------|----------------------|-------------------------------------|-------|
| % of non-zero price changes | $\kappa=0.0428$ 10 | $2 \times \kappa$ 7.4 | $\kappa=0.005$ 10 | $\frac{1}{2} \times \kappa$ 13.3 | 10 |
| Size of money shock (std $\times 100$) | 0.153 | 0.192 | 0.120 | 0.116 | 0.122 |
| Quarterly inflation (std $\times 100$) | 0.246 | 0.246 | 0.246 | 0.246 | 0.246 |
| % explained by μ shock alone | 100 | 100 | 100 | 100 | 100 |
| Quarterly output growth (std $\times 100$) | 0.31 | 0.46 | 0.174 | 0.336 | 1.08 |
| % explained by μ shock alone | 60.7 | 90.3 | 34.1 | 65.8 | 212 |
| Slope coeff. of Phillips curve* | 0.273 | 0.429 | 0.136 | 0.294 | 1.100 |
| | | | | 0.122 | 0.149 |

*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: $c_t^q = \beta_1 + \beta_2 \pi_t^q + \epsilon_t$, where the instrument

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

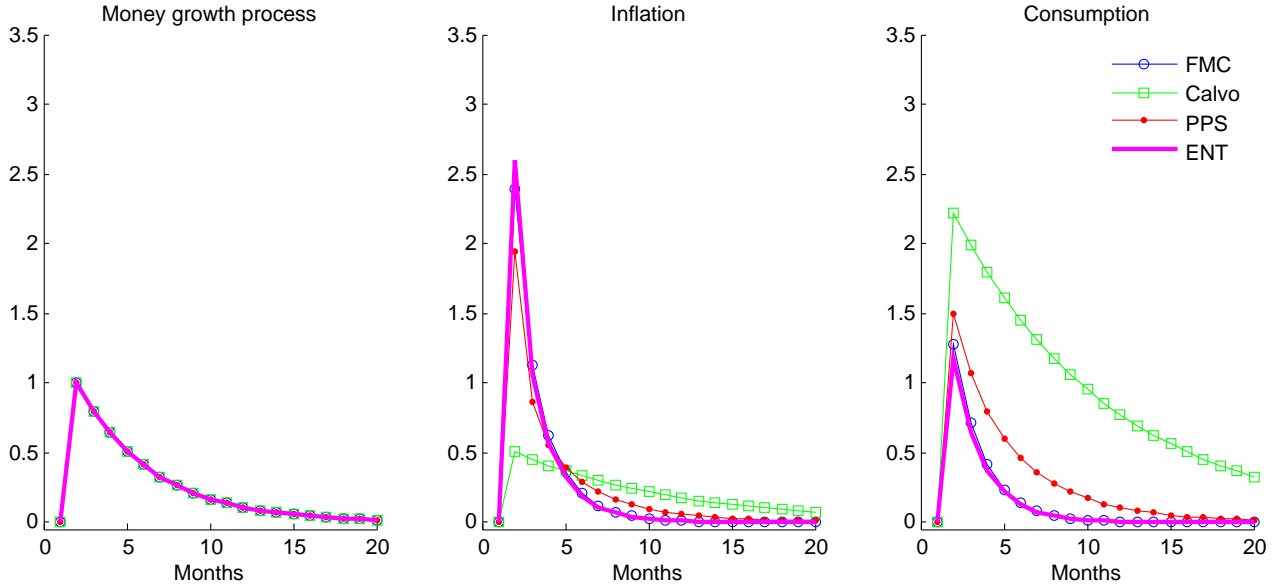


Fig. 9. Impulse responses to money growth shock: comparing models

Note: Simulated responses of inflation and consumption to a money growth shock, monthly autocorrelation 0.8.

Green squares: Calvo model; blue circles: fixed menu costs; red dots: PPS model; thick pink: ENT model.

added, but these wash out in the aggregate response. The PPS model behaves somewhat differently because its errors are so much bigger: indeed, a firm anticipates almost a 50% probability of *readjustment* immediately after any price change. This means that when it reacts to any shock, a firm expects to take several periods before it settles on a new price consistent with that shock. This slows down the convergence of the aggregate price too, and therefore implies a longer horizon over which a monetary shock has real effects. This explains why Table 2 shows real effects roughly twice as large in the PPS model as in the ENT model.

To separate the various components of the inflation response in a precise way, let us define the productivity-specific optimal price $p^*(a, \Xi_t) \equiv \arg \max_p v(p, a, \Xi_t)$, and also $x^*(\tilde{p}, a, \Xi_t) \equiv p^*(a, \Xi_t) - \tilde{p}$, the desired log price adjustment of a firm at the beginning of t with productivity a and real log price \tilde{p} . The actual log price adjustment of firm i can be decomposed as $x_{it} = x^*(\tilde{p}_{it}, a_{it}, \Xi_t) + \epsilon_{it}$, where ϵ_{it} is an error, in logs. We can then write the average desired adjustment at t as $\bar{x}_t^* \equiv \int \int x^*(\tilde{p}, a, \Xi_t) \tilde{\psi}_t(\tilde{p}, a) da d\tilde{p}$, and write the fraction of firms adjusting at t as $\bar{\lambda}_t \equiv \int \int \lambda(d(\tilde{p}, a, \Xi_t)) \tilde{\psi}_t(\tilde{p}, a) da d\tilde{p}$. Now, note that the average log error made by a firm with log productivity a that adjusts at the beginning of t is $\int (p - p^*(a, \Xi_t)) \pi(p|a, \Xi_t) dp$. Multiplying by the adjustment probability, and integrating over the population, the aggregate log error at t is

$$\bar{\epsilon}_t = \int \int \left(\int (p - p^*(a, \Xi_t)) \pi(p|a, \Xi_t) dp \right) \lambda(d(\tilde{p}, a, \Xi_t)) \tilde{\psi}_t(\tilde{p}, a) da d\tilde{p}. \quad (48)$$

Now inflation can be written as

$$i_t = i(\Xi_t, \Xi_{t-1}) = \int \int \left(\int x^*(\tilde{p}, a, \Xi_t) \pi(p|a, \Xi_t) dp \right) \lambda(d(\tilde{p}, a, \Xi_t)) \tilde{\psi}_t(\tilde{p}, a) da d\tilde{p} + \bar{\epsilon}_t. \quad (49)$$

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

$$\Delta i_t = i_t - i = \bar{\lambda} \Delta \bar{x}_t^* + \bar{x}^* \Delta \bar{\lambda}_t + \Delta \int \int x^*(\tilde{p}, a, \Xi_t) [\lambda(d(\tilde{p}, a, \Xi_t)) - \bar{\lambda}_t] \tilde{\psi}_t(\tilde{p}, a) da d\tilde{p} + \Delta \bar{\epsilon}_t, \quad (50)$$

where terms without time subscripts represent steady states, and Δ represents a change relative to steady state.¹⁴

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,

$$\mathcal{S}_t \equiv \Delta \int \int x^*(\tilde{p}, a, \Xi_t) [\lambda(d(\tilde{p}, a, \Xi_t)) - \bar{\lambda}_t] \tilde{\psi}_t(\tilde{p}, a) da d\tilde{p} \quad (51)$$

represents the part of inflation due to redistributing adjustment opportunities across firms desiring adjustments of different sizes (or signs), while fixing the total number adjusting. Finally, there is a fourth term, $\Delta \bar{\epsilon}_t$, which is the change in the aggregate log error.

The inflation impulse responses associated with all the specifications considered are decomposed in Figure 10. By construction, all changes in inflation in the Calvo model are 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 shifts the distribution in a way that leaves many firms with an especially low real price, so that these firms then prefer to adjust (likewise, some firms that would otherwise have had an especially high real price are discouraged from adjusting). These shifts are the selection effect. The larger price increase in these specifications explains the small consumption impact shown in Figure 9. However, the selection effect begins somewhat lower but is especially persistent in the PPS case. This is because the large errors in the PPS model move many firms out to the tails of the distribution; this shift is gradually corrected as firms eventually hit satisfactory prices. This slows down the equilibration of the price level in the PPS specification, so consumption reacts more strongly than it does in ENT and FMC, though still much less than it does in the Calvo model.

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.

¹⁴See Costain and Nakov (2011B) for further discussion of this decomposition.

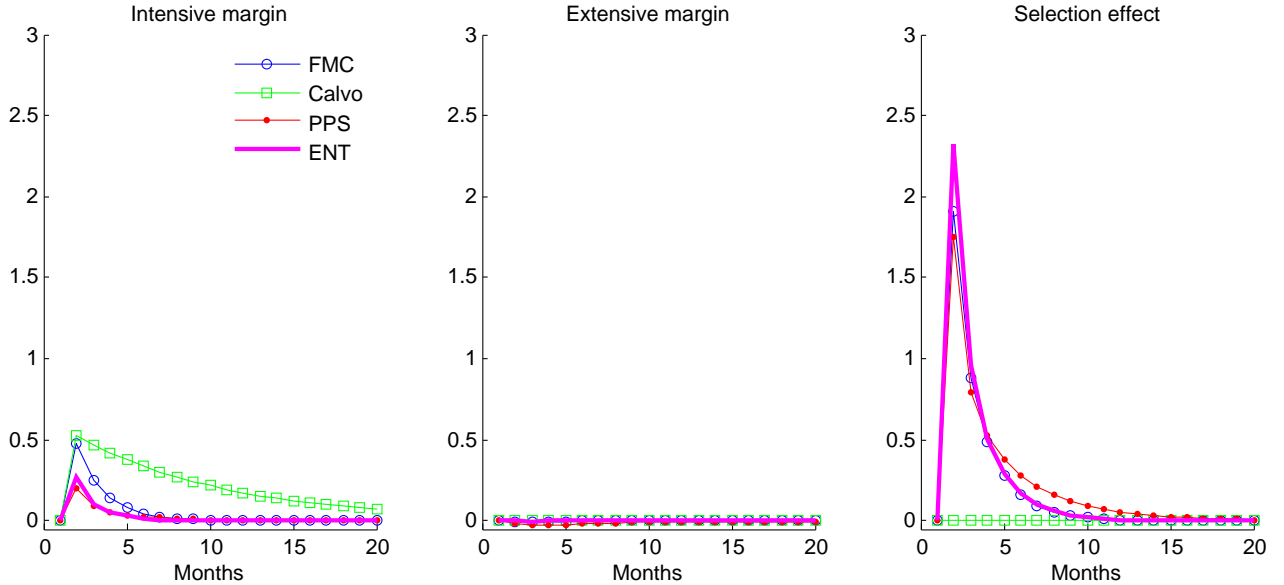


Fig. 10. Inflation decomposition after money shock: comparing models

Note: Intensive, extensive, and selection margins of inflation response to a money growth shock, monthly auto-correlation 0.8.

Green squares: Calvo model; blue circles: fixed menu costs; red dots: PPS model; thick pink: ENT model.

4.4 Effects of error-prone pricing

Finally, in Tables 1-2 and Figure 11 we can also see how our model's behavior varies with κ , the parameter that controls the amount of noise in the choice process. In Table 1, cutting noise in half raises the frequency of price adjustment from 10% to 13.3% per month (ENT model), or from 10% to 12.3% (PPS). At the same time, the average price change becomes smaller, for two reasons. 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. Of course, reducing errors also reduces the size of price changes directly, since the errors are an independent component of the price adjustment process.

In Figure 11, we see that lower noise increases the initial spike in inflation, and thus decreases the real effect of the nominal shock (see also Table 2). This is true both in the ENT and PPS specifications, though the baseline calibration of ENT is already so close to full flexibility that a further decrease in noise has only a minimal effect. In other words, as rationality increases, both the ENT and PPS models eventually converge to a fully flexible setup in which money is entirely neutral.

Differences in noise also help explain the main contrasts between the ENT and PPS specifications. In terms of microeconomic implications, Table 1 and Figure 3 show that ENT implies a more concentrated distribution, with smaller adjustments, than PPS. While PPS and ENT both

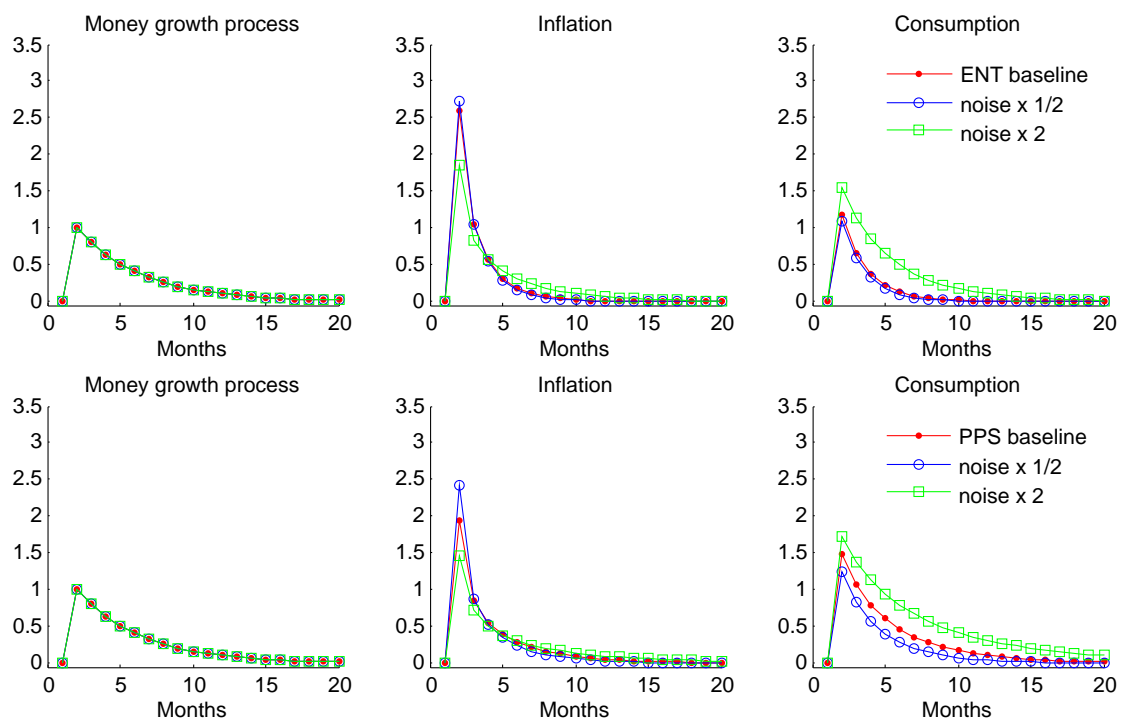


Fig. 11. Impulse responses to money growth shock: effects of noise parameter

Note: Simulated responses of inflation and consumption to a money growth shock, monthly autocorrelation 0.8.

Top row: ENT specification. Bottom row: PPS specification

Red line with dots indicates baseline noise parameter: $\kappa = 0.0050$ (ENT); $\kappa = 0.0428$ (PPS).

Green line with squares indicate higher noise: $\kappa = 0.0100$ (ENT); $\kappa = 0.0856$ (PPS).

Blue line with circles indicate lower noise: $\kappa = 0.0025$ (ENT); $\kappa = 0.0214$ (PPS).

imply that small and large price adjustments coexist, the histogram in the ENT specification is more sharply bimodal than the data, whereas for PPS it is more weakly bimodal and more spread out. Figures 4-6 show that PPS and ENT have qualitatively similar implications for several other micro facts, though the adjustment hazard is much more strongly decreasing in PPS than it is in ENT. In terms of impulse responses, Figures 9-11 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 $\kappa = 0.0428$ 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 $\kappa \mathcal{D}(\pi||u)$. Therefore matching the same 10% adjustment frequency requires a substantially lower degree of noise; the required calibration for the ENT specification is $\kappa = 0.0050$. This very low level of noise explains why the mean loss reported in the first row of Table 1 is lower for ENT than it is for PPS, in spite of the fact that choosing the right price entails a decision cost in ENT but not in PPS.

Since the ENT model combines a substantially lower noise parameter with a nonzero decision cost upon adjustment, it behaves more like the FMC model than PPS does. This is especially true in terms of impulse responses, where the real effects of money shocks are very small and almost identical under the FMC and ENT specifications. However, the presence of errors helps the ENT model fit microdata better than a fixed menu cost model does. Three of our findings especially favor ENT relative to FMC. First, ENT generates a more realistic distribution of price changes than FMC does, with a smoother histogram that includes some very small adjustments. Second, it implies that the adjustment hazard is approximately independent of the time since last adjustment. Third, it implies that the standard deviation of price adjustments only varies weakly with the aggregate inflation rate, while this standard deviation collapses quickly in the FMC model, as the fraction of price increases goes rapidly to 100%. The ENT model also raises the variability of prices relative to the variability of costs, though in our calibration it does not succeed in raising the ratio above one. Of the various empirical issues we have explored, the only point where FMC outperforms ENT is in matching the size of the price adjustment, as a function of the time since last adjustment (see Fig. 5). On this point, the FMC model matches the data almost perfectly.

5 Conclusions

We have analyzed the pricing behavior of near-rational firms which can freely adjust their prices at any time in response to idiosyncratic and aggregate shocks, but sometimes make mistakes. We model error-prone behavior by assuming that more accurate decisions are more costly. In equilibrium, firms devote some time to improving the precision of their price choices, but not so much that errors are completely eliminated. We focus on the special case where the cost of precision is measured by relative entropy; in this case, prices are distributed according to a multinomial logit. Prices are endogenously sticky, because when a firm leaves its price unchanged it avoids the decision costs associated with choosing a new price, and it avoids the risk of accidentally setting a price worse than the current one.

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. The presence of errors implies that the adjustment hazard, and the average size of the adjustment, are only weakly related to the time since last adjustment. As in the data, we find that extreme prices are more likely to be young, and that prices are more volatile than costs. However, our model of price stickiness does not seem to offer an explanation of monetary nonneutrality. Since our setup guarantees that firms that are making sufficiently large errors will choose to adjust, it generates a “selection effect” in response to nominal shocks that largely eliminates 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 fits microdata better than the Calvo model does, but implies a much more flexible aggregate price level, because of a strong selection effect. Thus, resolving the “puzzle” of small price adjustments need not, by itself, alter Golosov and Lucas’ conclusion that money shocks have small real effects. At the same time, our model also performs better at matching microdata than the fixed menu cost model does. In other words, taking seriously the idea that intermittent price adjustment is caused by decision costs, rather than menu costs *per se*, has observable implications which seem consistent with

microeconomic evidence; but it does not necessarily change the macroeconomic implications.

While this paper has emphasized the potential of price errors to explain a number of facts from microdata, we should stress that the *timing* of adjustment is perfectly rational here: a firm adjusts its price if and only if the expected value of doing so exceeds the expected value of leaving it unchanged. It would also be natural to consider a model with mistakes in both the *size* and the *timing* of price adjustment. We are studying an extension of this type in ongoing work (Costain and Nakov 2013). 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, Fernando; Francesco Lippi; and Luigi Paciello (2011), “Optimal price setting with observation and menu costs.” *Quarterly Journal of Economics* 126 (4), pp. 1909-1960.
- [3] Álvarez, Fernando; Martín González-Rozada; Andy Neumeyer; and Martín Beraja (2011), “From hyperinflation to stable prices: Argentina’s evidence on menu cost models.” Manuscript, Univ. of Chicago.
- [4] Anderson, Simon; André de Palma, and Jacques-François Thisse (1992), *Discrete Choice Theory of Product Differentiation*, MIT Press.
- [5] 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.
- [6] Baron, Richard, Jacques Durieu, Hans Haller, and Philippe Solal (2002), “Control costs and potential functions for spatial games.” *International Journal of Game Theory* 31, pp. 541-561.
- [7] Barro, Robert (1972), “A theory of monopolistic price adjustment.” *Review of Economic Studies* 39 (1), pp. 17-26.
- [8] Caballero, Ricardo (1992), “A fallacy of composition.” *American Economic Review* 82 (5), pp. 1279-1292.
- [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] 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.
- [12] Costain, James, and Anton Nakov (2011B), “Distributional dynamics under smoothly state-dependent pricing.” *Journal of Monetary Economics* 58 (6-8), pp. 646-665.
- [13] Costain, James, and Anton Nakov (2013), “Logit price dynamics.” Banco de España Working Paper 1301.

- [14] Danziger, Leif (1999), "A dynamic economy with costly price adjustments." *American Economic Review* 89 (4), pp. 878-901.
- [15] 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.
- [16] Dotsey, Michael; Robert King; and Alexander Wolman (2009), "Inflation and real activity with firm-level productivity shocks." Manuscript, Boston University.
- [17] Eichenbaum, M.; N. Jaimovich, and S. Rebelo (2011), "Reference prices and nominal rigidities." *American Economic Review* 101 (1), pp. 234-262.
- [18] Gagnon, Etienne (2009), "Price setting under low and high inflation: evidence from Mexico." *Quarterly Journal of Economics* 124, pp. 1221-1263.
- [19] Gertler, Mark, and John Leahy (2008), "A Phillips curve with an Ss foundation." *Journal of Political Economy* 116 (3), pp. 533-72.
- [20] Golosov, Mikhail, and Robert E. Lucas, Jr. (2007), "Menu costs and Phillips curves." *Journal of Political Economy* 115 (2), pp. 171-99.
- [21] 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.
- [22] Klenow, Peter, and Benjamin Malin (2009), "Microeconomic evidence on price setting." Forthcoming, *Handbook of Monetary Economics*, edited by Benjamin Friedman and Michael Woodford.
- [23] 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.
- [24] Marsili, Matteo (1999), "On the multinomial logit model." *Physica A* 269, pp. 9-15.
- [25] Matejka, Filip (2010), "Rationally inattentive seller: sales and discrete pricing." CERGE-EI Working Paper wp408.
- [26] Matejka, Filip, and Alisdair McKay (2013), "Rational inattention to discrete choices: a new foundation for the multinomial logit model." Manuscript, CERGE-EI and Boston University.
- [27] Mattsson, Lars-Goran, and Jorgen Weibull (2002), "Probabilistic choice and procedurally bounded rationality." *Games and Economic Behavior* 41, pp. 61-78.
- [28] McKelvey, Richard, and Thomas Palfrey (1995), "Quantal response equilibrium for normal form games." *Games and Economic Behavior* 10, pp. 6-38.
- [29] McKelvey, Richard, and Thomas Palfrey (1998), "Quantal response equilibrium for extensive form games." *Experimental Economics* 1, pp. 9-41.
- [30] Midrigan, Virgiliu (2011), "Menu Costs, Multiproduct Firms and Aggregate Fluctuations." *Econometrica* 79 (4), pp. 1139-1180.

- [31] 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.
- [32] Reiter, Michael (2009), “Solving heterogeneous-agent models by projection and perturbation.” *Journal of Economic Dynamics and Control* 33 (3), pp. 649-65.
- [33] Sheshinski, Eytan, and Yoram Weiss (1977), “Inflation and costs of price adjustment.” *Review of Economic Studies* 44 (2), pp. 287-303.
- [34] Sims, Christopher (2003), “Implications of rational inattention.” *Journal of Monetary Economics* 50, pp. 665-90.
- [35] Stahl, Dale (1990), “Entropy control costs and entropic equilibrium.” *International Journal of Game Theory* 19, pp. 129-138.
- [36] Woodford, Michael (2009), “Information-constrained state-dependent pricing.” *Journal of Monetary Economics* 56, pp. S100-S124.
- [37] 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.

Appendix A: Computation

Outline of algorithm

Computing this model is challenging due to heterogeneity: at any time t , firms will face different productivity shocks A_{it} and will be stuck at different prices P_{it} . The Calvo model is popular because, up to a first-order approximation, only the average price matters for equilibrium. But this property does not hold in most models; here we must treat all equilibrium quantities as functions of the time-varying distribution of prices and productivity across firms.

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, the algorithm calculates the steady-state general equilibrium in the absence of aggregate shocks. Idiosyncratic shocks are still active, but are assumed to have converged to their ergodic distribution, so the real aggregate state of the economy is a constant, Ξ . The algorithm solves for a discretized approximation to this steady state; here we restrict real log prices p_{it} and log productivities a_{it} to 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 both logarithmically spaced. We can then view the steady state value function as a matrix \mathbf{V} of size $\#^p \times \#^a$, comprising the values $v^{jk} \equiv v(p^j, a^k, \Xi)$ associated with prices and productivities $(p^j, a^k) \in \Gamma$.¹⁵ Likewise, the price distribution can be viewed as a $\#^p \times \#^a$ matrix $\mathbf{\Psi}$ in which the row j , column k element Ψ^{jk} represents the fraction of firms in state (p^j, a^k) at the end of any given period. Under this discrete representation, we can calculate steady state general equilibrium by guessing the wage w , 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. That is, the value function is represented by a $\#^p \times \#^a$ matrix \mathbf{V}_t with row j ,

¹⁵In this appendix, bold face indicates matrices, and superscripts represent indices of matrices or grids.

column k element $v_t^{j,k} \equiv v(p^j, a^k, \Xi_t)$, thus summarizing the time t values at all grid points $(p^j, a^k) \in \Gamma^p$. Then, instead of viewing the Bellman equation as a functional equation that defines $v(p, a, \Xi)$ for all possible idiosyncratic and aggregate states p , a , and Ξ , we simply regard it as an expectational relation between the matrices \mathbf{V}_t and \mathbf{V}_{t+1} . This amounts to a (large!) system of $\#^p \#^a$ first-order expectational difference equations that determine the dynamics of the $\#^p \#^a$ variables $v_t^{j,k}$. We linearize these equations numerically (together with the $\#^p \#^a$ equations that describe the evolution of the mass of firms at each grid point, and a few other scalar equations). We then solve for the saddle-path stable solution of the linearized model using the QZ decomposition, following Klein (2000).

This method combines linearity and nonlinearity in a way appropriate for models of price setting, where idiosyncratic shocks tend to have a bigger effect on firms' decisions than aggregate shocks do (*e.g.* Klenow and Kryvtsov, 2008; Golosov and Lucas, 2007; Midrigan, 2011). When we linearize the model's aggregate dynamics, we recognize that changes in the aggregate shock z_t or in the distribution Ψ_t are unlikely to have a strongly nonlinear effect on the value function. Note that this smoothness does not require any "approximate aggregation" property, in contrast with the Krusell and Smith (1998) method; nor do we need to impose any particular functional form on the distribution Ψ . However, to allow for the importance of firm-specific shocks, Reiter's method treats variation along idiosyncratic dimensions in a fully nonlinear way: the value at each grid point is determined by a distinct equation, which could in principle be entirely different from the equations governing the value at neighboring points.

The discretized model

In the discretized model, the value function \mathbf{V}_t is a matrix of size $\#^p \times \#^a$ with elements $v_t^{j,k} \equiv v(p^j, a^k, \Xi_t)$ for $(p^j, a^k) \in \Gamma$. The expected value of setting a new price is a column vector $\tilde{\mathbf{v}}_t$ of length $\#^a$, with k th element

$$\tilde{v}_t^k \equiv \kappa w(\Xi_t) \log \left(\frac{1}{\#^p} \sum_{j=1}^{\#^p} \exp \left(\frac{v_t^{j,k}}{\kappa w(\Xi_t)} \right) \right). \quad (52)$$

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

$$d_t^{j,k} \equiv \tilde{v}_t^k - v_t^{j,k}, \quad (53)$$

$$\lambda_t^{j,k} \equiv \mathbf{1} \left(d_t^{j,k} \geq 0 \right), \quad (54)$$

$$g_t^{j,k} \equiv \lambda_t^{j,k} d_t^{j,k}. \quad (55)$$

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

$$\pi_t^{j,k} = \pi(p^j | a^k, \Xi_t) \equiv \frac{\exp \left(v_t^{j,k} / (\kappa w(\Xi_t)) \right)}{\#^p \exp \left(\tilde{v}_t^k / (\kappa w(\Xi_t)) \right)},$$

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 .

In this discrete representation, the productivity process (14) can be written in terms of a $\#^m \times \#^a$ matrix \mathbf{S} , where the (m, k) element represents the following transition probability:

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

It is helpful to introduce analogous Markovian notation for the price process. Let \mathbf{R}_t be a $\#^p \times \#^p$ Markov matrix in which the row m , column l element represents the probability that firm i 's beginning-of-period log real price $\tilde{p}_{i,t}$ equals $p^m \in \Gamma^p$ if its log real price at the end of the previous period was $p^l \in \Gamma^p$:

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

Generically, the deflated log price $p_{i,t-1} - i(\Xi_{t-1}, \Xi_t)$ will fall between two grid points; then the matrix \mathbf{R}_t must round up or down stochastically. Hence, for any $i_t \equiv i(\Xi_{t-1}, \Xi_t)$, we construct \mathbf{R}_t according to¹⁶

$$R_t^{ml} = \text{prob}(\tilde{p}_{it} = p^m | p_{i,t-1} = p^l, i_t) = \begin{cases} \frac{i_t^{-1} p^l - p^{m-1}}{p^m - p^{m-1}} & \text{if } p^m = \min\{p \in \Gamma^p : p \geq i_t^{-1} p^l\} \\ \frac{p^{m+1} - i_t^{-1} p^l}{p^{m+1} - p^m} & \text{if } p^m = \max\{p \in \Gamma^p : p < i_t^{-1} p^l\} \\ 0 & \text{otherwise} \end{cases} \quad (56)$$

Given this notation, we can now write the distributional dynamics in a more compact form. Equations (40) and (42) become

$$\tilde{\Psi}_t = \mathbf{R}_t * \Psi_{t-1} * \mathbf{S}', \quad (57)$$

where $*$ represents ordinary matrix multiplication. Note that exogenous shocks are represented from left to right in the matrix Ψ_t , so that their transitions can be described by right multiplication, while policies are represented vertically, so that transitions related to policies can be described by left multiplication. Next, 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. Equations (40) and (42) can then be rewritten as

$$\Psi_t = (\mathbf{E}_{pa} - \mathbf{\Lambda}_t) .* \tilde{\Psi}_t + \mathbf{\Pi}_t .* (\mathbf{E}_{pp} * (\mathbf{\Lambda}_t .* \tilde{\Psi}_t)), \quad (58)$$

where (as in MATLAB) the operator $.*$ represents element-by-element 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

$$u_t^{jk} \equiv \left(\exp(p^j) - \frac{w(\Xi_t)}{\exp(a^k)} \right) \frac{C(\Xi_t)}{\exp(\epsilon p^j)}, \quad (59)$$

for any $(p^j, a^k) \in \Gamma$. Then the Bellman equation can be written as

$$\mathbf{V}_t = \mathbf{U}_t + \beta E_t \left\{ \frac{u'(C(\Xi_{t+1}))}{u'(C(\Xi_t))} [\mathbf{R}'_{t+1} * (\mathbf{V}_{t+1} + \mathbf{G}_{t+1}) * \mathbf{S}] \right\}. \quad (60)$$

The expectation E_t in (60) refers only to the effects of the time $t+1$ aggregate shock z_{t+1} , because the 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 (60) iterates backwards in time, its transitions are governed by \mathbf{R}' and \mathbf{S} , whereas (57) iterates forward in time and therefore involves \mathbf{R} and \mathbf{S}' .

Next, we discuss how we apply the two steps of the Reiter (2009) method to this discrete model.

¹⁶Note that in this model, price adjustment occurs with probability one if the current price is sufficiently far from its desired level. Therefore, we can choose Γ^p sufficiently wide so that that prices never drift beyond the endpoints of the grid. Otherwise, the definition of \mathbf{R}_t would have to allow for the cases $i_t^{-1} p^l < p^1$ or $i_t^{-1} p^l > p^{\#^p}$.

Step 1: steady state

In the aggregate steady state, the aggregate shocks are zero, and the (end-of-period) distribution is some constant Ψ , so the state of the economy is constant: $\Xi_t \equiv (z_t, \Psi_{t-1}) = (0, \Psi) \equiv \Xi$. We abbreviate the names of steady state objects by suppressing time subscripts and the function argument Ξ ; for example, the steady state value function is \mathbf{V} , with elements $v^{jk} \equiv v(p^j, a^k, \Xi)$.

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

$$\mu = \exp(i).$$

Thus the steady-state transition matrix \mathbf{R} is known, since it depends only on steady state inflation i . Moreover, the Euler equation reduces to

$$\exp(i) = \beta R.$$

As in Section 3.4, we focus on the case of linear labor disutility, $x(N) = \chi N$. We can then calculate steady-state 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. 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 calculate $C = (w/\chi)^{1/\gamma}$. We can then construct the matrix \mathbf{U} with elements

$$u^{jk} \equiv \left(\exp(p^j) - \frac{w}{\exp(a^k)} \right) \frac{C}{\exp(\epsilon p^j)}. \quad (61)$$

We then find the fixed point of the value \mathbf{V} (simultaneously with $\tilde{\mathbf{v}}$, \mathbf{D} , $\mathbf{\Lambda}$, and \mathbf{G}):

$$\mathbf{V} = \mathbf{U} + \beta \mathbf{R}' * (\mathbf{V} + \mathbf{G}) * \mathbf{S}. \quad (62)$$

This also allows us to calculate the matrix of logit probabilities $\mathbf{\Pi}$, with elements

$$\pi^{jk} = \frac{\exp(v^{jk}/(\kappa w))}{\sum_{n=1}^{\#p} \exp(v^{jn}/(\kappa w))}. \quad (63)$$

We can then find the steady state distribution as the fixed point of these two equations:

$$\Psi = (\mathbf{E}_{pa} - \mathbf{\Lambda}) * \tilde{\Psi} + \mathbf{\Pi} * (\mathbf{E}_{pp} * (\mathbf{\Lambda} * \tilde{\Psi})), \quad (64)$$

$$\tilde{\Psi} = \mathbf{R} * \Psi * \mathbf{S}'. \quad (65)$$

Finally, we check whether

$$1 = \sum_{j=1}^{\#p} \sum_{k=1}^{\#a} \Psi^{jk} (p^j)^{1-\epsilon} \equiv p(w). \quad (66)$$

If $p(w) = 1$, then an equilibrium value of w has been found.

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. We can summarize the general equilibrium in terms of the exogenous shock process z_t ; the lagged distribution of idiosyncratic states Ψ_{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 i_t . The equations reduce to

$$z_{t+1} = \phi_z z_t + \epsilon_{t+1}^z, \quad (67)$$

$$\frac{\mu \exp(z_t)}{\exp(i_t)} = \frac{m_t}{m_{t-1}}, \quad (68)$$

$$\Psi_t = (\mathbf{E}_{pa} - \mathbf{\Lambda}_t) \cdot * \tilde{\Psi}_t + \mathbf{\Pi}_t \cdot * (\mathbf{E}_{pp} * (\mathbf{\Lambda}_t \cdot * \tilde{\Psi}_t)), \quad (69)$$

$$\mathbf{V}_t = \mathbf{U}_t + \beta E_t \left\{ \frac{u'(C_{t+1})}{u'(C_t)} [\mathbf{R}'_{t+1} * (\mathbf{V}_{t+1} + \mathbf{G}_{t+1}) * \mathbf{S}] \right\}, \quad (70)$$

$$1 = \sum_{j=1}^{\#p} \sum_{k=1}^{\#a} \Psi_t^{jk} (p^j)^{1-\epsilon}. \quad (71)$$

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

$$\vec{X}_t \equiv (\text{vec}(\Psi_{t-1})', \text{vec}(\mathbf{V}_t)', C_t, m_{t-1}, i_t)',$$

then the whole set of expectational difference equations (67)-(71) governing the dynamic equilibrium becomes a first-order system of the following form:

$$E_t \mathcal{F}(\vec{X}_{t+1}, \vec{X}_t, z_{t+1}, z_t) = 0, \quad (72)$$

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 i_t and i_{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 $u^{jk} \equiv \left(\exp(p^j) - \frac{w_t}{\exp(a^k)} \right) \frac{C_t}{\exp(ep^j)}$. Finally, given \mathbf{V}_t and \mathbf{V}_{t+1} , we can construct $\tilde{\mathbf{v}}_t$, $\mathbf{\Pi}_t$, \mathbf{D}_t , and \mathbf{D}_{t+1} , and thus $\mathbf{\Lambda}_t$ and \mathbf{G}_{t+1} . Therefore the variables in \vec{X}_t and z_t are indeed sufficient to evaluate the system (67)-(71). In all, we have $2\#p\#a + 3$ equations in $2\#p\#a + 3$ unknowns.

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:

$$E_t \mathcal{A} \Delta \vec{X}_{t+1} + \mathcal{B} \Delta \vec{X}_t + E_t \mathcal{C} z_{t+1} + \mathcal{D} z_t = 0, \quad (73)$$

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.