PRICE ADJUSTMENTS IN A GENERAL MODEL OF STATE-DEPENDENT PRICING

2008

James Costain and Antón Nákov

Documentos de Trabajo
N.º 0824

BANCO DE ESPAÑA
Eurosistema

PRICE ADJUSTMENTS IN A GENERAL MODEL OF STATE-DEPENDENT PRICING

PRICE ADJUSTMENTS IN A GENERAL MODEL OF STATE-DEPENDENT PRICING(*)

James Costain and Antón Nákov

BANCO DE ESPAÑA

(*) We wish to thank Michael Reiter, Virgiliu Midrigan, and seminar participants at the Vienna Institute for Advanced Studies, the Bank of Spain, the 2008 SNDE meetings and the REDg-DGEM workshop, for helpful comments. We especially thank Virgiliu Midrigan, Etienne Gagnon, and Oleksiy Kryvstov for providing their data. The views expressed in this paper are those of the authors.

Documentos de Trabajo. N.º 0824
2008

The Working Paper Series seeks to disseminate original research in economics and finance. All papers have been anonymously refereed. By publishing these papers, the Banco de España aims to contribute to economic analysis and, in particular, to knowledge of the Spanish economy and its international environment.

The opinions and analyses in the Working Paper Series are the responsibility of the authors and, therefore, do not necessarily coincide with those of the Banco de España or the Eurosystem.

The Banco de España disseminates its main reports and most of its publications via the INTERNET at the following website: http://www.bde.es.

ISSN: 0213-2710 (print)
ISSN: 1579-8666 (on line)
Depósito legal:
Unidad de Publicaciones, Banco de España

Abstract

In this paper, we show that a simple model of smoothly state-dependent pricing generates a distribution of price adjustments similar to that observed in microeconomic data, both for low and high inflation. Our setup is based on one fundamental assumption: price adjustment is more likely when it is more valuable. The constant probability model (Calvo 1983) and the fixed and stochastic menu cost models (Golosov and Lucas 2007; Dotsey, King and Wolman 1999) are nested as special cases of our framework.

All parameterizations of our model can be ranked according to a measure of state dependence. The fixed menu cost model has the highest possible degree of state dependence; the parameterization which best fits US microdata has low state dependence. The fixed menu cost model is inconsistent with the evidence both because it never generates small price adjustments, and because it implies a large fall in the standard deviation of price adjustments as trend inflation increases. Even though the state dependence of our preferred parameterization is almost as low as that of the Calvo model, it is well-behaved when we change the steady state inflation rate, matching the data at least as well as Golosov and Lucas' model.

Keywords: Price stickiness, state-dependent pricing, stochastic menu costs, generalized (S,s), bounded rationality

JEL Codes: E31, D81

1 Introduction

Sticky prices are an essential component of contemporary dynamic general equilibrium models, including those used by central banks for policy analysis. Nonetheless, the way in which price stickiness is modeled remains as controversial as ever. The Calvo (1983) formulation of a constant adjustment probability has become popular due to its analytical tractability, but lacks the theoretical appeal of a microfounded model immune to the Lucas critique. It has been criticized in an influential article by Golosov and Lucas (2007), who propose a model of price stickiness based on a fixed "menu cost" of adjusting prices. They calibrate their model to match certain features of the empirical distribution of price changes and find that it predicts much weaker and less persistent real effects of nominal shocks than the Calvo model.

Given this controversy, it is fortunate that a variety of micro-level evidence on pricing behavior has recently become available, in papers by Klenow and Kryvstov (2008), Midrigan (2008), Nakamura and Steinsson (2008), and Gagnon (2007), among others. These studies provide cross-sectional evidence on the distribution of price adjustments in a variety of retail contexts, differing in the range of sellers and the underlying inflation rate in the countries and time periods considered. Hopefully these new data will permit further progress in selecting a useful, microfounded model of price rigidity.

In this paper, we study a general model of state-dependent pricing which nests a number of well-known sticky price frameworks as special cases. We make one main assumption: a firm is more likely to adjust its price when adjustment is more valuable. The time-dependent model of Calvo (1983) is a limiting case of our model in which the probability of adjustment is constant; the state-dependent menu cost model of Golosov and Lucas (2007) is a limiting case of our model in which the probability of adjustment is a step function taking values zero and one only. We investigate which version of our model works best by estimating its parameters and comparing the pricing behavior it generates with the cross-sectional distribution of price adjustments observed in micro data. Our estimation results allow us to address a number of questions: to what extent is price adjustment state dependent rather than time dependent? How does the distribution of price adjustments change when the underlying inflation rate changes? What is the welfare loss associated with price stickiness?

Golosov and Lucas' study of these issues stressed the fact that firm-level data show price increases and decreases much larger than those that could be justified by aggregate inflation alone, which suggests that most price adjustments are driven by idiosyncratic rather than by aggregate factors. They showed that incorporating idiosyncratic shocks into their model helps it reproduce some price adjustment statistics, such as the average absolute size of price changes. Furthermore, they stress that the combination of fixed menu costs with large idiosyncratic shocks in their model implies a strong selection effect— meaning that firms which need to make a larger price adjustment are more likely to adjust— which makes aggregate prices very sensitive (and real quantities insensitive) to nominal shocks.

We share Golosov and Lucas' assessment that idiosyncratic shocks are the main driving force behind firm-level price adjustments. However, a closer look at the microeconomic evidence reveals a number of ways in which their model's predictions are inconsistent with observed price setting behavior. The histogram of price adjustments generated by their model shows a large spike of price increases, a large spike of price decreases, and nothing in between; the actual distribution of price adjustments is much smoother (Kashyap 1995, Midrigan 2008). Another, related, prediction of their model is that the standard deviation of price adjustments decreases sharply as the inflation rate rises, whereas cross-sectional data from periods of low and high inflation like those of Gagnon (2007) show a small, nonmonotonic change in the standard deviation of price adjustments with increased inflation. Both of these features of the data call for a model in which prices adjust more smoothly than fixed menu costs imply.

Summarizing our main findings, an appropriately calibrated model in which (a) the probability of price adjustment increases smoothly with the value of adjustment, and (b) firms are subject to idiosyncratic shocks, does a good job of matching the cross-sectional distribution of price changes. In contrast, the fixed menu cost and Calvo alternatives (the two limiting cases of our general model) work poorly because they both fail to reproduce the frequency of small price changes (see Figure 1). However, along several dimensions, our estimated model is much closer to the Calvo specification than it is to the fixed menu cost model. In particular, we propose a measure of state dependence, based on the cross-sectional variance of the adjustment probability. The parameters that best fit our model imply very low state dependence, almost as low as in the Calvo case. Nonetheless, our estimated model is reasonably consistent with price adjustment data even at high rates of aggregate inflation under which the Calvo model breaks down.

After a brief literature review, in section 2 we study a simple model of the partial equilibrium pricing decision of an imperfectly competitive firm, showing how it nests a variety of adjustment models. We then incorporate our setup into a dynamic stochastic general equilibrium framework in section 3 and explain how to solve for the steady state distribution of prices. In section 4 we estimate the parameters that best match the distribution of price changes in the AC Nielsen dataset documented by Midrigan (2008), under two specifications of the adjustment probability function, one of which is taken from Woodford (2008). We compare the distribution of price changes implied by our model to the distributions observed under Calvo pricing and under the fixed menu cost model, both for low and high steady state inflation rates. In section 5 we briefly consider a generalization of our model that allows firms to set a default inflation rate. Section 6 concludes.

1.1 Related literature

A new empirical literature on pricing began with Klenow and Kryvtsov (2008), who analyzed a large dataset of prices from the US Bureau of Labor Statistics. They showed that while price changes are large on average (around 10% in absolute value), about 40% of price changes are less than 5% in absolute value. Similar results are found by Midrigan (2008), based on scanner data from selected retail stores, and by Nakamura and Steinsson (2008) using a more detailed dataset than Klenow and Kryvtsov (2008).

The high frequency of tiny price changes found in these studies is clearly at odds with the implications of the standard one-sector fixed menu cost model of Golosov and Lucas (2007). A possible explanation is that some sectors might have higher fixed costs than others. But Klenow and Kryvtsov (2008) find that even allowing for 67 sectors, each with a different fixed menu cost, the model still fails to match the distribution of price changes, displaying a big hole in the middle of the distribution.[1] Moreover, they show that large and small price changes coexist within narrowly defined product categories too. That is, even in a sector-by-sector calibration, there is a tradeoff between picking a high menu cost to match the large average (absolute) size of price changes and a low menu cost to reproduce the small price changes in that sector.

Midrigan (2008) points out that this puzzle might disappear in a multi-product setting where firms face economies of scale when adjusting prices: once a firm pays the menu cost to correct some large price misalignment, they can also change other (less misaligned) prices on the same menu costlessly. He also emphasizes the necessity of leptokurtic technology shocks in order to reproduce the excess kurtosis of the distribution of price changes. But while Midrigan's multi-product model seems like a plausible explanation of some small price adjustments, it may take the existence of fixed menu costs too literally. We feel that on grounds of realism it is essential to impose some smoothness on individual behavior. Our calibrated simulations show that simply requiring the adjustment probability to vary smoothly with the value of adjustment suffices to construct a model that fits microdata well, including the presence of frequent small price changes.

In other words, our simulations show that a pricing model like the stochastic menu cost setup of Dotsey, King, and Wolman (1999) fits well if it incorporates large idiosyncratic shocks (Dotsey *et al.* ignored idiosyncratic shocks to simplify the calculation of aggregate dynamics). While Caballero and Engel (1993, 1999, 2006, 2007) have long advocated a similar mechanism, calling it the "generalized (S,s) approach", most of the new papers on price stickiness have instead opted to assume fixed menu costs (including Golosov and Lucas 2007, Midrigan 2008, Nakamura and Steinsson 2008, and Gagnon 2007). Also, while Caballero and Engel (1999) begin by writing their adjustment probability in terms of values, they mostly focus on a "semi-structural" setup in which the adjustment probability is a function of the distance between the current policy and the optimal policy. We hope to convince the reader that both the optimization problem and the distributional dynamics can be stated in an elegant and numerically tractable way (see our matrix formulation in Sec. 3.5) if we treat the adjustment probability as a function of the value of adjustment, which is anyway more correct.

Another recent paper with implications for the cross-sectional distribution of price adjustments is that of Gertler and Leahy (2006), who construct a state-dependent pricing model that yields the same tractable Phillips curve as that of the Calvo model. However, they make some strong assumptions to simplify the aggregation of (S,s) policies, including a uniform distribution of idiosyncratic shocks (as in Danziger, 1999), implying a counterfactual distribution of price changes. Woodford (2008) studies pricing decisions when information processing is costly, and derives the implied functional form for the probability of adjustment, but does not calibrate his model to microeconomic data. Dorich (2007) calibrates time-dependent and state-dependent pricing models to

[1] We reproduce Klenow and Kryvstov's data and results in the last panel of Figure 1.

measure the implied welfare losses due to stickiness. Caballero and Engel (2006) also calibrate their model to the Nakamura and Steinsson data, and calculate the Calvo parameter that best matches the data.

In contrast to these recent related papers, we study how the various versions of our calibrated model (from a low inflation environment) perform as the steady state inflation rate increases. In particular, we find that in the presence of fixed menu costs, the histogram of price adjustments collapses from bimodality to unimodality as inflation increases, leading to a large decrease in the standard deviation of price changes. This contrasts with Gagnon's (2007) finding in Mexican data that the standard deviation of price changes remains roughly constant as the inflation rate rises. In addition, our paper offers a measure of the degree of state dependence, according to which the Calvo model has the least state dependence, while the Golosov-Lucas model has the most. We show that the data favor a parameterization that is quite close to the Calvo model in its degree of state dependence.

2 Sticky prices in partial equilibrium

Before we describe the full general equilibrium structure of our economy, it is helpful to study the partial equilibrium pricing decision of a monopolistic producer, to see how our model nests a variety of popular pricing frameworks.

Like Golosov and Lucas (2007), we assume price changes are driven primarily by idiosyncratic shocks. If firms are entirely rational, fully informed, and capable of frictionless adjustment, they will adjust their prices every time a new idiosyncratic or aggregate shock is realized. We instead assume that prices are "sticky", in a well-defined sense: the probability of adjusting is less than one, but is greater when the benefit from adjusting is greater. What we mean by "the benefit from adjusting" becomes clear as soon as we write down the Bellman equations that describe the firm's decision. There is a value associated with optimally choosing a new price today (while bearing in mind that prices will not always be adjusted in the future); likewise there is a value associated with leaving the current price unchanged today (likewise bearing in mind that prices will not always be adjusted in the future). The difference between these two values is the benefit from adjusting (or the loss from failing to adjust). The smoothly increasing function $\lambda(L)$ that gives the adjustment probability as a function of the loss L from failing to adjust is taken as a primitive of the model.

There are at least two ways of interpreting this framework. It could be seen as a model of stochastic menu costs, as in Dotsey, King, and Wolman (1999) or Caballero and Engel (1999). If rational, fully-informed firms draw an *iid* adjustment cost x every period, with cumulative distribution function $\lambda(x)$, then they will adjust their behavior whenever the adjustment cost x is less than or equal to the loss L from failing to adjust. Therefore, their probability of adjustment is $\lambda(L)$ when the loss from nonadjustment is L.

But perhaps this is an unnecessarily literal interpretation of the model. Alternatively, following Akerlof and Yellen (1985), "stickiness" can be modeled as a minimal deviation from rational expectations behavior, in which agents occasionally fail to make precisely optimal choices if the cost of such errors is small. Under full rationality, full information, and zero adjustment costs, economic agents adjust to a new optimal setting of their control variables in every period; here instead we assume that they sometimes fail to make this adjustment. Perhaps failure to adjust occurs because information itself is "sticky" (as in Reis, 2006), or perhaps because calculation requires effort; rather than taking a stand on this, we simply regard our assumption as an axiom that should be imposed on near-rational, near-full-information behavior.

For analytical tractability, many authors have implemented near rationality as Akerlof and Yellen did, making most agents fully rational while assuming that a few others follow a rule of thumb. Here, instead, we implement it by assuming all agents are close to full rationality, but not quite there. Our framework permits us to deviate smoothly from full rationality, both because we can choose a λ function that is close to one for most L, and more importantly because large mistakes are less likely than trivial ones. To us, this version of near rationality seems less arbitrary since there is no need to specify a rule of thumb (though we have to choose a functional form for λ), and also more realistic, especially in terms of its compatibility with microdata. The main cost of our assumption is that the model requires a numerical analysis.

2.1 The monopolistic competitor's decision

Suppose then, following Golosov and Lucas (2007), that each firm i produces output Y_{it} under a constant returns technology, with labor N_{it} as the only input, and faces idiosyncratic productivity shocks A_{it}:

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

We assume firms are monopolistic competitors, facing the demand curve $Y_{it} = \vartheta P_{it}^{-\epsilon}$, where ϑ represents aggregate demand, and that they must fulfill all demand at the price they set. They hire in competitive labor markets at wage rate W, so the period profit function is

$$\Pi_{it} = P_{it} Y_{it} - W N_{it} = \left(P_{it} - \frac{W}{A_{it}} \right) Y_{it} = \left(P_{it} - \frac{W}{A_{it}} \right) \vartheta P_{it}^{-\epsilon}$$

Since we focus in this paper on the cross-sectional distribution of price adjustments, we assume the aggregate state of the economy is constant. That is, there are no aggregate shocks, and the distribution of productivity shocks and prices across firms has converged to a steady state.[2] This is why we treat aggregate demand ϑ and the wage rate W as constants. We assume the idiosyncratic productivity shocks A_{it} are given by a time-invariant Markov process, *iid* across firms. Thus A_{it} is correlated with A_{it-1} but is uncorrelated with all other firms' shock processes. The assumption that productivity shocks are the only shocks affecting the firm (other than the random arrival of adjustment opportunities) is inessential for our methodology; we ignore more general cases only to keep notation simple.

To implement our assumption that adjustment is more likely when it is more valuable, we must define the values of adjustment and nonadjustment. If a firm fails to adjust (so that $P_{it} = P_{it-1}$), then its current profits and its future prospects will both depend on its productivity A_{it} and on its price P_{it}. Therefore these both enter as state variables in the value function of a nonadjusting firm, $V(P_{it}, A_{it})$. When a firm adjusts, we assume it chooses the best price conditional on its current productivity shock (of course, taking into account the fact that it may not adjust in all future periods). Therefore, the value function of an adjusting firm, after netting out any costs that may be required to make the adjustment, is just $V^*(A_{it}) \equiv \max_P V(P, A_{it})$. The value of adjusting to the optimal price, written in the same units as the value function, is then

$$D(P_{it}, A_{it}) \equiv \max_P V(P, A_{it}) - V(P_{it}, A_{it})$$

Of course, we don't want the real probability of adjustment to differ when values are denominated in euros instead of pesetas. In order to treat the function λ as a primitive of the model, we need to write it in the appropriate units. Under either of the proposed interpretations of the model, the most natural units are those of labor time. Under the stochastic menu cost interpretation, the labor effort of changing price tags or rewriting the menu is likely to be a large component of the cost. Under the bounded rationality interpretation, even though we don't explicitly model the computation process, we take the probability of adjustment to be related to the labor effort associated with obtaining new information and/or recomputing the optimal price. Therefore, the function λ should depend on the loss from failing to adjust, converted into units of labor time by dividing by the wage rate. That is, the probability of adjustment is $\lambda\left(L\left(P_{it}, A_{it}\right)\right)$, where $L\left(P_{it}, A_{it}\right) = D(P_{it}, A_{it})/W$ and λ is a given increasing function.

Conditional on adjustment, we have assumed that the firm sets the optimal price, $P^*(A_{it}) \equiv \arg\max_P V(P, A_{it})$. For clarity, we will distinguish between the firm's beginning-of-period price, $\widetilde{P}_{it} = P_{it-1}$, and the price at which it produces and sells at time t, P_{it}, which may or may not differ from $\widetilde{P}_{it}$. The adjustments are determined by the function λ:

$$P_{it} = \begin{cases} P^*(A_{it}) & \text{with probability } \lambda(D(\widetilde{P}_{it}, A_{it})/W) \\ \widetilde{P}_{it} = P_{i,t-1} & \text{with probability } 1 - \lambda(D(\widetilde{P}_{it}, A_{it})/W). \end{cases}$$

Function λ must satisfy $\lambda' \geq 0$. In particular, we will consider the class

$$\lambda\left(L\right) \equiv \frac{\bar{\lambda}}{\bar{\lambda} + \left(1 - \bar{\lambda}\right)\left(\frac{\alpha}{L}\right)^{\xi}} \tag{1}$$

with α and ξ positive and $\bar{\lambda} \in [0, 1]$. This function equals $\bar{\lambda}$ when $L = \alpha$, and is concave for $\xi \leq 1$ and S-shaped for $\xi > 1$. We choose this class because it has fatter tails than the normal *cdf*, which may help it match the

[2] Studying the transitional dynamics of this economy, or its response to aggregate shocks, requires calculating the dynamics of the distribution of firm-level productivity and prices. This problem is interesting both for monetary policy and from the perspective of numerical methodology. We address the dynamics in a companion paper, Costain and Nakov (2008b).

fat tails of the observed price adjustment distribution emphasized by Midrigan (2008).[3] Importantly for the purposes of our estimation, with $\xi = 0$ the function becomes constant at $\bar{\lambda}$, while with $\xi \to \infty$ it approaches a step function taking values 0 (when $L < \alpha$) and 1 (when $L > \alpha$). In this sense the Calvo and fixed menu costs specifications are nested as (extreme) special cases of this more general hazard function. Also, as α goes to zero or $\bar{\lambda}$ goes to one, $\lambda(L)$ equals one for all L, so the case of fully flexible prices is nested too.

We can now state the Bellman equation that defines the value of producing at any given price. It differs somewhat depending on whether we impose the stochastic menu cost interpretation of our model or the bounded rationality interpretation; we begin with the latter because it is slightly simpler. Given the firm's price P and its productivity shock A, current profits are $\left(P - \frac{W}{A}\right)\vartheta P^{-\epsilon}$. The firm anticipates adjusting or not adjusting in the next period depending on the benefits of adjusting at that time. Therefore the Bellman equation is:

$$V(P,A) = \left(P - \frac{W}{A}\right)\vartheta P^{-\epsilon} + R^{-1}E\left\{ (1-\lambda(L(P,A')))\,V(P,A') + \lambda(L(P,A'))\max_{P'} V(P',A') \Big| A \right\}$$

where R^{-1} is the firm's discount factor and the expectation refers to the distribution of A' conditional on A. Note that on the left-hand side of this equation, and in the current profits term, P refers to a given firm i's price P_{it} at the time of production. In the expectation on the right, P represents the price $\widetilde{P}_{i,t+1} = P_{it}$ at the beginning of period $t+1$, which may (probability λ) or may not (probability $1-\lambda$) be adjusted prior to time $t+1$ production.

We can simplify substantially by noticing that the value on the right-hand side of the equation is just the value of continuing without adjustment, plus the expected gains due to adjustment:

$$V(P,A) = \left(P - \frac{W}{A}\right)\vartheta P^{-\epsilon} + R^{-1}E\left\{ V(P,A') + \lambda(L(P,A'))\,D(P,A') \middle| A\right\} \tag{2}$$

or equivalently

$$V(P,A) = \left(P - \frac{W}{A}\right)\vartheta P^{-\epsilon} + R^{-1}E\left\{ V(P,A') + G(P,A') \middle| A\right\} \tag{3}$$

where $G(P,A') \equiv \lambda(L(P,A'))\,D(P,A')$ represents the expected gains due to adjustment. A natural solution method for this equation is backwards induction on a two-dimensional grid of possible prices and productivity shocks. However, we postpone discretizing the state space until Sec. 3.3, when we detrend the model to express all prices in real terms.

2.2 Alternative sticky price frameworks

1. **Calvo pricing** (Calvo 1983): Suppose prices adjust each period with constant exogenous probability $\bar{\lambda}$. This corresponds to a special case of our hazard function (1) with $\xi = 0$. The Bellman equation is the limiting case of (2) in which $\lambda(L)$ reduces to the constant function:

$$\lambda(L) \equiv \bar{\lambda} \tag{4}$$

2. **Fixed menu costs** (Golosov and Lucas 2007; Mankiw, 1985): Suppose the cost of adjusting prices in any given period is κ units of labor (where κ is an exogenous constant called the "menu cost"). This corresponds to a special case of our hazard function (1) with $\xi \to \infty$ and $\alpha = \kappa$. The Bellman equation is the same as (3), with expected gains function equal to

$$G(P,A') = \mathbf{1}\{L(P,A') \geq \kappa\}\,(D(P,A') - \kappa W) \tag{5}$$

That is, the adjustment probability λ is given by the indicator function $\lambda(L) = \mathbf{1}\{L \geq \kappa\}$, which takes value one when $L \geq \kappa$ and zero otherwise. The gain function G also subtracts off the term $\lambda(L)\,\kappa W$, that is, the firm anticipates paying the nominal menu cost κW with probability $\lambda(L)$.

[3] We started with the specification $\lambda(L) = L/(\alpha + L)$, which is (1) with $\overline{\lambda} = 0.5$ and $\xi = 1$, but found we needed more degrees of freedom to match the data. Upon reading Woodford (2008), we realized that by including the three parameters α, ξ, and $\overline{\lambda}$ we could explicitly nest the Calvo and fixed menu cost setups into our general model.

3. **Stochastic menu costs** (Dotsey, King, and Wolman, 1999): Suppose the cost of adjusting prices in any given period is κ units of labor, where κ is an *i.i.d.* random variable with *cdf.* $\lambda(\kappa)$. Then the Bellman equation is the same as (3), setting the expected gains function to

$$G(P, A') = \lambda(L(P, A'))[D(P, A') - WE(\kappa | L(P, A') > \kappa)] \tag{6}$$

That is, with probability $\lambda(L)$, the firm realizes nominal gains D, and it pays the (conditional) expected menu cost $WE(\kappa | L > \kappa)$.

Thus, the only difference between the "bounded rationality" interpretation of our model and the stochastic menu cost interpretation is that in the latter case, the conditional expected menu cost $WE(\kappa|L > \kappa)$ is netted out of the flow of gains G. Note also that the Calvo model can be written as a stochastic menu cost model too, by assuming the menu cost κ takes value 0 with probability $\bar{\lambda}$ and value ∞ with probability $1 - \bar{\lambda}$; in this case the flow of expected menu costs is zero.

It is also useful to consider how our framework compares to some alternative price adjustment mechanisms:

1. **Generalized (S,s) model** (Caballero and Engel 1999): while their setup is initially defined in terms of stochastic menu costs, Caballero and Engel go on to write λ as a function of the distance between the current price and the optimal price. This "semi-structural" formulation (as they call it) is an approximation to the behavior of a stochastic menu cost model.

2. **Information processing** (Woodford 2008): Woodford shows that under certain conditions, optimal information processing implies adjustment according to the following probability function:

$$\lambda(L) \equiv \frac{\bar{\lambda}}{\bar{\lambda} + \left(1 - \bar{\lambda}\right) \exp\left(\xi\left(\alpha - L\right)\right)} \tag{7}$$

where α and ξ are positive constants and $\bar{\lambda} \in [0, 1]$. This suggests another possible functional form for use in our model.

3. **Quadratic adjustment costs** (Rotemberg 1982): this model is substantially different since all agents make a partial adjustment instead of some agents making a complete adjustment. As for the aggregate implications, quadratic adjustment costs imply that the average deviation from the optimal price converges linearly to zero over time. In our model, convergence is faster for large deviations, but slower for small deviations.

2.3 Measuring state dependence

As we have emphasized, the Calvo model and the fixed menu cost model are both limiting cases of our general model, corresponding to different versions of the function λ. The Calvo model is purely time dependent because λ does not vary with the benefit of adjusting; the Golosov and Lucas model is state dependent because adjustment occurs if and only if the benefit from adjusting exceeds a given threshold.

Since we have ignored adjustment mechanisms that vary explicitly with time (as in Taylor, 1979), the Calvo model is in fact the only version of our model that is entirely time dependent; all other cases have some degree of state dependence. Interestingly, if we take the variance of the adjustment probability as a measure of state dependence, then the Golosov and Lucas model places an upper bound on the degree of state dependence, and all other variants of our model can be ranked between the Calvo and Golosov-Lucas extremes. This motivates the following definition.

Definition. Let $\bar{\lambda} = E(\lambda_i)$ be the mean adjustment probability across firms i in the population; let $\sigma_\lambda^2 = E\left(\left(\lambda_i - \bar{\lambda}\right)^2\right)$ be the corresponding variance. Then the **degree of state dependence** S is given by

$$S = \frac{\sigma_\lambda^2}{\bar{\lambda}\left(1 - \bar{\lambda}\right)}$$

The Calvo model and the Golosov-Lucas model provide lower and upper bounds on the degree of state dependence, as stated in the following proposition.

Proposition. The degree of state dependence satisfies $S \in [0,1]$, with $S = 0$ in the Calvo model, and $S = 1$ in the fixed menu cost model.

Proof. Since σ_{λ}^2 is a variance, it is nonnegative; it equals zero in the Calvo model. Let $f(\lambda_i) \in [0,1]$ be the cross-sectional frequency of the adjustment probabilities $\lambda_i \in [0,1]$ across firms i in the population. Then (writing expectations with sums, under the assumption that there are a finite number of values of λ_i in the population)

$$\begin{aligned}
\sigma_{\lambda}^2 &= E\left((\lambda_i)^2\right) - \bar{\lambda}^2 = \sum_{\lambda_i} f(\lambda_i)\lambda_i^2 - \bar{\lambda}^2 = \sum_{\lambda_i} f(\lambda_i)\lambda_i\left(\lambda_i - \bar{\lambda}\right) \\
&= \sum_{\lambda_i < \bar{\lambda}} f(\lambda_i)\lambda_i\left(\lambda_i - \bar{\lambda}\right) + \sum_{\lambda_i > \bar{\lambda}} f(\lambda_i)\lambda_i\left(\lambda_i - \bar{\lambda}\right) \leq \sum_{\lambda_i > \bar{\lambda}} f(\lambda_i)\lambda_i\left(\lambda_i - \bar{\lambda}\right) \\
&\leq \left(1 - \bar{\lambda}\right)\sum_{\lambda_i} f(\lambda_i)\lambda_i = \left(1 - \bar{\lambda}\right)\bar{\lambda}
\end{aligned}$$

Thus we have an upper bound for the cross-sectional variance of λ_i. In the case of the fixed menu cost model, λ_i only takes values 0 or 1, and $\bar{\lambda} = f(1)$. In this case, $\sigma_{\lambda}^2 = f(0) * 0 * (0 - \bar{\lambda}) + f(1) * 1 * (1 - \bar{\lambda}) = \left(1 - \bar{\lambda}\right)\bar{\lambda}$. Therefore $S = 1$ in the fixed menu cost model. **QED.**

3 General equilibrium

We next embed this partial equilibrium decision framework in a dynamic New Keynesian general equilibrium like that in Golosov and Lucas (2007). In addition to the firms, there is a representative household and a monetary authority that chooses the money supply.

3.1 Households

The household's period utility function is

$$u(C_t) - x(N_t) + v\left(M_t/P_t\right)$$

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

$$\int_0^1 P_{it}C_{it}di + M_t + R_t^{-1}B_t = W_tN_t + M_{t-1} + T_t + B_{t-1} + \Pi_t$$

where $\int_0^1 P_{it}C_{it}di$ is total nominal spending on the differentiated goods. B_t is nominal bond holdings, with interest rate $R_t - 1$; T_t represents lump sum transfers received from the monetary authority, and Π_t represents dividend payments received from the firms.

Assuming households choose $\{C_{it}, N_t, B_t, M_t\}_{t=0}^{\infty}$ so as to maximize expected discounted utility subject to the budget constraint, we obtain the following necessary conditions. 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:

$$P_t \equiv \left\{\int_0^1 P_{it}^{1-\epsilon}di\right\}^{\frac{1}{1-\epsilon}}$$

Given this consumption allocation, we can also write nominal spending as $P_tC_t = \int_0^1 P_{it}C_{it}di$. Optimal labor supply and money holdings imply the first-order conditions

$$\begin{aligned} x'(N_t) &= u'(C_t)W_t/P_t \\ \nu'\left(\frac{M_t}{P_t}\right) &= u'(C_t)(1-R_t^{-1}) \end{aligned}$$

and the Euler equation is

$$R_t^{-1} = \beta E_t\left(\frac{P_t u'(C_{t+1})}{P_{t+1}u'(C_t)}\right)$$

3.2 Monetary policy and aggregate consistency

We assume the growth rate of the money supply follows an autoregressive process:

$$M_t = \mu_t M_{t-1}$$

where $\mu_t - \mu = \phi\left(\mu_{t-1}-\mu\right) + \varepsilon_t^{\mu}$, $\mu = E\left[\mu_t\right] \geq 1$, and $\varepsilon_t^{\mu} \sim i.i.d.N(0,\sigma_{\mu}^2)$ is a money growth shock. Seigniorage revenues are paid to the household as a lump sum transfer. Therefore aggregate consistency in the money markets requires $M_t = M_{t-1} + T_t$, while aggregate consistency in the bond market is simply $B_t = 0$.

Market clearing for good i implies the following demand and supply relations for firm i:

$$Y_{it} = A_{it}N_{it} = C_{it} = P_t^{\epsilon}C_tP_{it}^{-\epsilon}$$

This is consistent with our description of the firm's problem if we set $\vartheta_t = P_t^{\epsilon}C_t$. We can then also calculate total labor demand:

$$N_t = \int_0^1 \frac{C_{it}}{A_{it}}di = P_t^{\epsilon}C_t\int_0^1 P_{it}^{-\epsilon}A_{it}^{-1}di$$

3.3 Discrete detrended steady state

Since our focus in this paper is the distribution of prices and price adjustments across firms, we restrict our simulations to steady-state cross-sectional implications. To obtain a steady state, we have to detrend with respect to the money supply, which is nonstationary. So next we rewrite all nominal variables in real terms by dividing by the money supply. We define the real aggregate price level as $p_t = P_t/M_t$ and the real wage as $w_t = W_t/M_t$; we write the real production price of firm i at time t as $p_{it} = P_{it}/M_t$, and its real price at the beginning of t as $\widetilde{p}_{it} = \widetilde{P}_{it}/M_t$.

Furthermore, for numerical tractability we assume that real prices p_{it} and $\widetilde{p}_{it}$ can only lie in a discrete grid $\Gamma^p \equiv \left\{p^1, p^2, ...p^{\#p}\right\}$ (we use superscripts to indicate grid elements, and $\#p$ is the number of points in the grid). It is convenient to set Γ^p so that it is evenly spaced in logarithms: $\log(p^i/p^{i-1}) = \Delta_p > 0$ for all integers $i \in \{2, 3, ..., \#p\}$. Likewise, we assume the firm's productivity always takes one of a finite number of values: $A_{it} \in \Gamma^a \equiv \left\{a^1, a^2, ..., a^{\#a}\right\}$ for all i and t. Therefore, if the aggregate economy is in a steady state, any given firm's real state is summarized by a point in a finite, two-dimensional grid: $(p_{it}, A_{it}) \in \Gamma^p \times \Gamma^a$. Let the fraction of firms in state $(p^j, a^k) \in \Gamma^p \times \Gamma^a$ at the time of period t production be $\psi_t\left(p^j, a^k\right)$. Similarly, using tildes again to represent variables prior to adjustment, let the fraction of firms in state $(p^j, a^k) \in \Gamma^p \times \Gamma^a$ at the beginning of t be $\widetilde{\psi}_t\left(p^j, a^k\right)$. In other words, distribution $\widetilde{\psi}_t$ involves the beginning-of-period prices $\widetilde{p}_{it}$, and distribution ψ_t involves the prices p_{it} at the time of production.

Convergence to a steady state means that the distribution of prices and productivity across firms converges to two alternating ergodic distributions: $\widetilde{\psi}_t\left(p, A\right) \rightarrow \widetilde{\psi}\left(p, A\right)$ at the beginning of the period, and $\psi_t\left(p, A\right) \rightarrow \psi\left(p, A\right)$ at the time of production, where we indicate steady states by suppressing time subscripts. Likewise, aggregate variables converge to steady states: $p_t \rightarrow p$, $w_t \rightarrow w$, $C_t \rightarrow C$, $N_t \rightarrow N$, and $R_t \rightarrow R$. In this steady state, the Euler equation becomes

$$R^{-1} = \beta/\mu \tag{8}$$

The household's first-order conditions converge to

$$x'(N) = u'(C)\,w/p \tag{9}$$
$$\nu'(1/p) = (1-\beta/\mu)\,u'(C) \tag{10}$$

In steady state, the demand for a product produced by a firm with price p_{it} is $C_{it} = p^{\epsilon}Cp_{it}^{-\epsilon}$. Thus we can express labor market clearing in terms of the distribution of prices and productivity at the time of production, as follows:

$$N = p^{\epsilon}C\sum_{j=1}^{\#p}\sum_{k=1}^{\#a}\left(p^{j}\right)^{-\epsilon}\left(a^{k}\right)^{-1}\psi\left(p^{j},a^{k}\right) \tag{11}$$

Aggregate consumption must satisfy

$$C^{\frac{\epsilon-1}{\epsilon}} = \sum_{j=1}^{\#p}\sum_{k=1}^{\#a}\left[p^{\epsilon}C\left(p^{j}\right)^{-\epsilon}\right]^{\frac{\epsilon-1}{\epsilon}}\psi\left(p^{j},a^{k}\right) \tag{12}$$

and the price level must satisfy

$$p^{1-\epsilon} = \sum_{j=1}^{\#p}\sum_{k=1}^{\#a}\left(p^{j}\right)^{1-\epsilon}\psi\left(p^{j},a^{k}\right) \tag{13}$$

It is straightforward to verify that the last two equations are equivalent, so only one of them will be needed for calculating equilibrium.

3.4 Firm behavior in a discrete detrended steady state

Previously we wrote the value function in nominal terms. We next restate the firm's problem in steady state general equilibrium, detrending in the same way we detrended the rest of the model, dividing nominal quantities by the nominal money supply. Thus, the real value function v must satisfy:

$$V(P_{it},A_{it}) = M_t v\left(\frac{P_{it}}{M_t},A_{it}\right) = M_t v\left(p_{it},A_{it}\right)$$

Likewise, the detrended versions of the difference function and the expected gains function that appear in the Bellman equation must satisfy

$$D(P_{it},A_{it}) = M_t d\left(\frac{P_{it}}{M_t},A_{it}\right) = M_t d\left(p_{it},A_{it}\right)$$
$$G(P_{it},A_{it}) = M_t g\left(\frac{P_{it}}{M_t},A_{it}\right) = M_t g\left(p_{it},A_{it}\right)$$

Writing the Bellman equation in detrended notation also requires us to treat time t and time $t+1$ quantities in a consistent way. Since the money stock is multiplied by the factor μ from one period to the next, calculating the real price at the beginning of $t+1$ requires division by μ:

$$\widetilde{p}_{it+1} \equiv \frac{\widetilde{P}_{it+1}}{M_{it+1}} = \frac{P_{it}}{M_{it+1}} = \frac{M_{it}}{M_{it+1}}\frac{P_{it}}{M_{it}} = p_{it}/\mu$$

Then for any real state $(p_{it},A_{it}) = (p^{j},a^{k}) \in \Gamma^{p}\times\Gamma^{a}$, the Bellman equation (3) that nests all versions of our model can be written as

$$v(p^{j},a^{k}) = \left(p^{j}-\frac{w}{a^{k}}\right)C\left(\frac{p^{j}}{p}\right)^{-\epsilon} + \beta E\left\{v(\mu^{-1}p^{j},A') + g(\mu^{-1}p^{j},A')|a^{k}\right\}$$

Here we have rewritten real aggregate demand as $\vartheta/M = Cp^{\epsilon}$. As before, prices on the left-hand side refer to the time of production, while prices inside the expectation on the right-hand side are those at the beginning of the next period, which is why they are deflated by μ. The expectation is taken with respect to next period's productivity shock (written as A'), conditional on this period's productivity, $A_{it} = a^k$.

Note that our notation implicitly assumes that $\mu^{-1}p^j$ is a grid point in Γ^p. This is not a problem, because we are free to choose the steady state inflation rate and the distance between grid points. By choosing money growth so that $\log \mu$ is an integer multiple of the step size in the grid, we ensure that $\mu^{-1}p^j$ will be a grid point. That is, if $\log \mu = \#\mu\Delta_p$, where $\#\mu$ is a positive integer, then $\mu^{-1}p^j = p^{j-\#\mu} \in \Gamma^p$. Therefore, to simplify notation, from here on we assume $\log \mu$ is an integer multiple of Δ_p.[4]

3.5 Firm behavior: matrix formulation

Given that the real dynamics are constrained to a discrete grid, our Bellman equations can be written in an even more compact notation which will also help us write down the distributional dynamics explicitly and concisely. The key is to notice that we can track the dynamics of the exogenous state (productivity) separately from the dynamics of the endogenous state (prices). The discrete Markov process that governs productivity can be summarized by a matrix $\mathbf{S}$ with row m, column k element

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

where $a^m, a^k \in \Gamma^a$. Similarly, the real period t price p_{it} is deflated to $\widetilde{p}_{it+1} = \mu^{-1}p_{it}$ at the beginning of period $t+1$. This adjustment can be summarized by a matrix $\mathbf{R}$ in which the row l, column j element is

$$R^{lj} = prob(\widetilde{p}_{i,t+1} = p^l | p_{it} = p^j)$$

Since we have assumed that trend money growth equals $\#\mu$ steps in the price grid (where $\#\mu$ is an integer), column j of matrix $\mathbf{R}$ must have a one in row $j - \#\mu$, and zeros elsewhere (assuming $j > \#\mu$). For columns $j \leq \#\mu$, deflating by factor μ would leave prices outside of the grid Γ^p. Therefore, we instead assume that any prices which fall off the grid are automatically rounded up to the minimum price p^1 (that is, columns $j \leq \#\mu$ of matrix $\mathbf{R}$ have a one in the first row and zeros elsewhere.) This assumption is made for numerical purposes only, and has negligible impact on the equilibrium as long as we choose a sufficiently wide price grid Γ^p.

For conformability with this matrix notation, it makes sense to write the value function as a matrix too, corresponding to the values of all prices and productivities in our discrete grid. Let $\mathbf{V}$ be the matrix with the value of state $\left(p^j, a^k\right) \in \Gamma^p \times \Gamma^a$ in row j, column k:

$$v^{jk} = v\left(p^j, a^k\right)$$

Similarly, let $\mathbf{G}$ be the matrix with the row j, column k element

$$g^{jk} = g\left(p^j, a^k\right)$$

for all prices and productivities in our grid. Finally, let $\mathbf{U}$ be the $\#p \times \#a$ matrix of current payoffs, with elements $u^{jk} = \left(p^j - \frac{w}{a^k}\right) C \left(\frac{p^j}{p}\right)^{-\epsilon}$ for all $\left(p^j, a^k\right) \in \Gamma^p \times \Gamma^a$. Then it is easily verified that Bellman equation (3) simplifies to

$$\mathbf{V} = \mathbf{U} + \beta\mathbf{R}' * (\mathbf{V} + \mathbf{G}) * \mathbf{S} \tag{14}$$

where $*$ denotes matrix multiplication.[5]

To write the distributional dynamics in matrix notation, let $\mathbf{\Psi}$ be the $\#p \times \#a$ matrix in which the row j, column k element

$$\psi^{jk} = \psi\left(p^j, a^k\right)$$

represents the fraction of firms in state $\left(p^j, a^k\right) \in \Gamma^p \times \Gamma^a$ at the time of production. Likewise, let $\widetilde{\mathbf{\Psi}}$ represent the beginning-of-period distribution, so that the row j, column k element is

$$\widetilde{\psi}^{jk} = \widetilde{\psi}\left(p^j, a^k\right)$$

[4] We emphasize that this assumption is for notational simplicity only. It is not required for our calculations, and we do not always impose it on our calculations.

[5] To clarify, we emphasize that (14) contains an expectational term even though no E_t operator is seen. The expectation over future idiosyncratic productivity $A_{i,t+1}$ conditional on A_{it} is captured by multiplying by the matrix $\mathbf{S}$.

Now, since rows relate to prices, and columns relate to productivities, the price dynamics can be represented by left multiplication, and the productivity dynamics by right multiplication. Therefore, the relation between the time t distribution and the time $t+1$ beginning-of-period distribution is simply

$$\widetilde{\Psi}_{t+1} = \mathbf{R} * \Psi_t * \mathbf{S}' \tag{15}$$

Likewise, the relation between the beginning-of-t distribution and the distribution at the time of production in period t can also be summarized in matrix notation. Let $\mathbf{E}$ be a $\#p \times \#p$ matrix of ones. Let $\lambda(\mathbf{D}/w)$ be a $\#p \times \#a$ matrix with element (j,k) equal to $\lambda(d\left(p^j, a^k\right)/w)$, where $\left(p^j, a^k\right) \in \Gamma^p \times \Gamma^a$, representing the probability of adjustment from any beginning-of-period state. Also, for any $a^k \in \Gamma^a$, define $l(k)$ so that $p^{l(k)} = \arg\max_{p \in \Gamma^p} v\left(p, a^k\right)$. Then let $\mathbf{P}$ be the $\#^p \times \#^a$ matrix with a one in row $l(k)$ of each column k, and zero elsewhere. Then the production-time distribution Ψ_t is calculated from the beginning-of-period distribution $\widetilde{\Psi}_t$ as follows:

$$\Psi_t = (\mathbf{E} - \lambda(\mathbf{D}/w)) .* \widetilde{\Psi}_t + \mathbf{P} .* (\mathbf{E} * (\lambda(\mathbf{D}/w) .* \widetilde{\Psi}_t)) \tag{16}$$

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

3.6 Computing steady state equilibrium

Steady state equilibrium is described by the first-order conditions (9) and (10), the aggregate consistency conditions (11) and (13), and the matrix formulation of the firm-level dynamics, (14), (15), and (16). These equations suffice to determine the four scalars p, C, w, and N, and the three matrices $\mathbf{V}$, $\mathbf{\Psi}$, and $\widetilde{\mathbf{\Psi}}$.

This system of equations is easily converted into a two-dimensional fixed point problem in p and N. But our simulations will focus on the even simpler case where the disutility of labor is linear, $x(N) = \chi N$. In this case, the equilibrium calculation reduces to a one-dimensional fixed point problem:

1. Guess p.
2. Use (10) to calculate C, then (9) to calculate w; then construct $\mathbf{U}$.
3. Find the fixed point $\mathbf{V}$ of (14), then construct $\mathbf{P}$ from the optimal policy associated with $\mathbf{V}$.
4. Find the fixed points $\mathbf{\Psi}$ and $\widetilde{\mathbf{\Psi}}$ associated with (15) and (16).
5. Use (13) to determine a new p, and return to step 2.

Note that in this calculation, $\mathbf{R}$ and $\mathbf{S}$ are known *ex ante*, and $\mathbf{G}$ and $\lambda(\mathbf{D}/w)$ are just abbreviations denoting transformations of the value function $\mathbf{V}$.

4 Results

4.1 Parameterization

We can now compute the steady state of our model, parameterizing it to fit cross-sectional microeconomic data on price changes. In particular we seek to match the distribution of price changes in the AC Nielsen dataset of household product purchases documented by Midrigan (2008). Therefore we will simulate our model at monthly frequency, and we report all parameters at monthly frequency unless otherwise specified. Furthermore, we set the steady state growth rate of money to 0%, consistent with the zero average price change in the AC Nielsen dataset. Also, since Midrigan removes price changes attributable to temporary "sales", our simulation results should be interpreted as a model of "regular" price changes unrelated to sales. This is also consistent with the frequency and methodology of Klenow and Kryvstov (2008) and Nakamura and Steinsson (2008).

We take our utility parameterization from Golosov and Lucas (2007), setting 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 $\nu = 1$. But ν is just a normalization for our purposes, because changing it affects the wage and the aggregate price p, but does not alter the distribution of relative prices or price adjustments.

Given these parameters, we can calibrate the idiosyncratic productivity shock process and the adjustment process to match the distribution of regular price changes. 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. The free parameters here are ρ and σ_ε^2, the variance of the innovation ε_t^a. Since our numerical method requires us to treat A as a discrete variable, we use Tauchen's method to approximate this AR(1) on a discrete grid Γ^A. As for the adjustment process (1), it has three free parameters, $\bar{\lambda}$, α and ξ. We estimate all five parameters by minimizing a criterion consisting of two terms: (1) the distance between the average frequency of adjustment in the data and that in the model, and (2) the distance between the distribution of log price changes in the data and that in the model. We search for the optimal parameter vector with a global constrained minimization routine initialized from a grid of alternative starting values.

For the frequency of adjustment, we choose to match the median monthly frequency of regular price changes of 10% reported by Nakamura and Steinsson (2008). This is lower than the 20% mean frequency of regular price changes reported by Midrigan (2008), both because Nakamura and Steinsson control more carefully for sales, and because the median frequency is less driven by outliers than the mean. We prefer Nakamura and Steinsson's number for its apparent robustness. The parameter $\bar{\lambda}$ is closely linked to the frequency of adjustment, so our estimate of $\bar{\lambda}$ is always close to 0.1. For the distribution of price adjustments, we use a histogram of 25 equally spaced bins representing log changes from -0.5 to 0.5 in the AC Nielsen data, and we measure the distance between the histograms in the data and the model as Euclidean distance. We then divide the distance by the number of bins (25) so that the frequency term and the distribution term in our minimization criterion are comparably weighted.

4.2 Estimating the degree of state dependence

Table 1 and Figure 1 summarize our main estimation results. The table reports the best estimate for each specification, along with a number of statistics describing the implied distribution of price changes. Clearly many of these measures are correlated, but we report them all because different empirical studies have focused on different statistics. The last four columns reproduce the corresponding statistics reported by Midrigan (2008) for his AC Nielsen and Dominick's datasets, as well as those from Nakamura and Steinsson (2008) and Klenow and Kryvstov (2008).

The columns marked SDSP (for "state-dependent sticky prices") are based on estimating the two parameters of the productivity process and the three parameters of the baseline adjustment process (1). We run the estimate both on a 25-by-25 grid and on a 101-by-101 grid; there is little difference in the results.[6] The estimated model matches our target adjustment frequency of 10% per month almost exactly, and also does a good job of hitting the moments of the distribution of price adjustments. The mean absolute price change is 10% in the model, and 10.5% in the data; the median is slightly lower in both cases because the distribution has fat tails. In the model, the standard deviation of the distribution of price changes is 11.9%, and the kurtosis is 2.7; in the data they are 13.2% and 3.5. Half of all price adjustments in the data are increases, and we obtain very nearly the same figure in the model. The fit of the distribution is also illustrated by the fourth panel of Figure 1, which shows the histogram of nonzero log price adjustments for the AC Nielsen data (shaded), together with the corresponding histogram from our model. The fit is especially good in the middle range, though the tails are somewhat fatter in the data than in the model, as the difference in kurtosis indicates.

Our baseline SDSP estimate is further illustrated in Figure 2, which shows equilibrium objects like the value function $\mathbf{V}$, the value of adjustment $\mathbf{D}$, and the adjustment probability $\boldsymbol{\lambda}$, graphed as functions of inverse productivity A_{it}^{-1} and price p_{it}.[7] The fourth and sixth panels show the beginning-of-period distribution $\widetilde{\boldsymbol{\Psi}}$ and the distribution at the time of production, $\boldsymbol{\Psi}$; the sharp ridge seen in distribution $\boldsymbol{\Psi}$ represents the mass

[6] In the coarser estimate, Γ^A is a grid of 25 points covering plus or minus 2.5 standard deviations of log A. The price grid Γ^P extends 10% past the prices that would be chosen at the highest and lowest values of A if prices were fully flexible. In the finer estimate, Γ^A covers ± 5 standard deviations, so the distance between grid points (0.015) is half that in the coarser estimate (0.031).

[7] The axes of the graph are shifted so that a firm at the origin has the mean productivity level and has the price that a flexible-price firm would choose at that productivity.

of firms that have just changed their prices. The eighth graph compares the price policy function under our estimate to that of a firm with perfectly flexible prices (the dotted line). Note that a sticky-price firm prices "conservatively": if its productivity strays far away from the mean, it will adjust its price by less than a flexible-price firm, anticipating mean-reversion in productivity.

Table 1 also shows estimates of the fixed menu cost model, $\xi = \infty$ (marked MC), and the Calvo model, $\xi = 0$.[8] The productivity process is allowed to differ in each estimate. These models do a good job of replicating the frequency of price adjustments, but fare less well at matching the absolute size of price changes, measured as the mean or the median of the absolute (log) price change. In particular, the absolute price changes predicted by the menu cost model are too large on average, while in the Calvo model they are too small.

Moreover, both models fail dramatically in one respect. In the AC Nielsen data, 25% of the absolute price changes are smaller than 5%, whereas the Calvo model implies twice this many small changes, and in the fixed menu cost model there are no changes of this size whatsoever. This can also be seen from the model-generated histograms in the top panels of Figure 1. In the menu cost model (top left panel), the lack of small price adjustments is seen as a big hole in the middle of the histogram. That is, generating the large price changes found in the data requires such a high fixed menu cost that small price changes become prohibitively expensive. On the other hand, the Calvo model (top right panel) generates too many small price changes, implying a unimodal distribution of adjustments, in contrast with the mildly bimodal distribution in the data.

While these inconsistencies may at first sound like irrelevant details, they are potentially important in terms of welfare and policy analysis. In particular, given the observed frequency of adjustment, the fact that the fixed menu cost model underpredicts the incidence of small price adjustments means that it overpredicts the occurrence of (more valuable) large ones. Figure 3 illustrates this, showing the distribution of losses L (at the time of production), together with the adjustment probability function $\lambda(L)$. We see that the distribution is truncated on the right in the menu cost specification, eliminating all the most costly failures to adjust. Therefore the MC model is likely to exaggerate aggregate flexibility and also to understate firms' lost value due to price stickiness.

The difference in losses can also be seen at the bottom of Table 1. The table shows both the median and average loss L from failure to adjust, as a percentage of the median value V of a firm, in each of the calibrations. The median losses are small, ranging from a 0.042% loss in the fixed menu cost model to a loss of 0.114% in the Calvo case. The median loss under SDSP (0.074% of median firm value) lies in between these two. But the loss distribution is highly skewed in all cases except the fixed menu cost specification, where the long right tail of uneliminated losses is truncated. Therefore mean losses are always several times larger than median losses, except under fixed menu costs. Thus while median losses differ by less than a factor of three, average losses range from 0.086% of median firm value for MC, to 0.271% for SDSP, and 0.393% for Calvo: average losses differ by a factor of 4.6.

To sum up, our preferred calibration of the SDSP model does a good job of matching the cross-sectional distribution of price adjustments, and it works much better than either of the extreme special cases. Both the menu cost model ($\xi = \infty$) and the Calvo model ($\xi = 0$) have particular difficulty capturing the frequency of the smallest price adjustments, while the simulated distribution under our preferred value of ξ (the middle left panel of Figure 1) very closely resembles the data of Midrigan (2008). For comparison, we also show, in the bottom left panel of Figure 1, the distribution produced by Midrigan's (2008) model of multiproduct firms with fixed menu costs. The bottom right panel graphs the data of Klenow and Kryvstov (2008) together with the distribution from their fixed menu cost model with heterogeneity across sectors. Evidently allowing the probability of adjustment to vary smoothly with the value of adjustment is more useful for matching data on price changes than either of these sophisticated generalizations of the fixed menu cost model.

While neither of the extreme cases of the model fits well, many findings point to the conclusion that the Calvo specification gets closer to the data than the fixed menu cost specification does. The data favor an adjustment function that allows sizeable fractions of small and large price changes to coexist. This requires a value of ξ substantially below one, meaning that the probability of adjustment rises rapidly at zero and is fairly flat thereafter, resembling the Calvo model over most of the range of observed adjustments. This is also reflected in our measure of state-dependence: the estimated SDSP model scores much closer (0.025) to the Calvo model (0), than to the fixed menu cost model (1). Given these results, it is reasonable to expect that

[8] When $\xi = 0$, the "menu cost" parameter α is unidentified, and when $\xi = \infty$, the "Calvo" parameter $\bar{\lambda}$ is unidentified. Therefore we estimate these two specifications efficiently, which is to estimate the model (4) or the model (5) directly, instead of estimating (1) subject to a parameter restriction.

under our preferred calibration, monetary shocks will cause large real effects on output and labor, like those in the Calvo specification. We study impulse responses to monetary shocks in a companion paper, Costain and Nakov (2008).

4.3 Fixing the technology process

The estimates in Table 1 give each model its best shot at matching the histogram of actual price changes, because in addition to estimating the adjustment parameters, we also reestimate the productivity process for each specification. In particular, the standard deviation of the productivity process is 0.165 in our Calvo estimate, as opposed to 0.148 in the baseline SDSP estimate. Therefore, to isolate the effect of the different adjustment mechanisms themselves, we also report estimates of each adjustment specification conditional on the same productivity process (the one that was obtained for the SDSP baseline).

The results are shown in Table 2. The main thing to notice is that the statistics in Table 2 are very similar to those reported in Table 1. Thus the contrasting findings we emphasized earlier derive mainly from the different price adjustment mechanisms in the menu cost, Calvo, and SDSP models, and owe little to differences in the estimated productivity process.

4.4 Estimating Woodford's specification

Up to now, we have only considered our baseline adjustment probability function (1). Woodford (2008) instead advocates the following specification, on the basis of a model of information processing:

$$\lambda(L) \equiv \frac{\bar{\lambda}}{\bar{\lambda} + \left(1-\bar{\lambda}\right)\exp\left(\xi(\alpha - L)\right)}. \tag{17}$$

This nests the Calvo model ($\xi \to 0$) and the fixed menu cost model ($\xi \to \infty$) (with menu cost $\kappa = \alpha$) in the same way our baseline specification does. The differences between the specifications involve their implications when the loss L from failure to adjust is very large or very small. Woodford's model implies that the probability of adjustment converges to one exponentially as $L \to \infty$, whereas in our baseline specification it converges polynomially. Since the latter convergence is slower, in principle it may allow larger losses to persist in our model than in Woodford's. At the same time, Woodford's model allows for a strictly positive probability of "adjustment", $\lambda(0) = \left(1 + \frac{1-\bar{\lambda}}{\bar{\lambda}} e^{\xi a}\right)^{-1}$, even in the limit of no loss. In contrast, the probability of adjustment at zero loss is always zero in our baseline specification: $\lim_{L\to 0} \lambda(L) = 0$ for all strictly positive ξ. This means that convergence to the Calvo model (as $\xi \to 0$) is pointwise in our SDSP specification, whereas it is uniform in the model of Woodford.

Estimates of Woodford's specification are shown in the columns marked 'Woodf' in Table 1 (where we reestimate the productivity process too) and Table 2 (where the productivity process is assumed the same as in the benchmark SDSP estimate). Overall, Woodford's model has success in matching the data similar to that of our SDSP model. In particular, Woodford's model does a somewhat better job of capturing the tails of the distribution of price changes (and it gets the kurtosis statistic right), whereas our model works better in the middle of the histogram (especially for the lower frequency of very small price changes).

Looking at the adjustment probability function in Figure 3 helps clarify what is going on in the estimation of Woodford's model. His model implies that $\lambda(L)$ has a positive limit at $L = 0$, and the data appear not to favor an adjustment probability near one anywhere in the observed range of adjustments. Therefore, the estimated function $\lambda(L)$ is remarkably flat over its whole observed range. As a consequence, its measure of state dependence is 0.008, even closer to a Calvo model than our SDSP estimate (state dependence 0.025). That is, even though $\lambda(L)$ approaches one exponentially in Woodford's specification, this is irrelevant over the observed range of losses. This makes large losses somewhat more frequent, so the mean and median losses due to price stickiness are slightly larger in Woodford's case than in SDSP, almost as large as in the Calvo model.[9]

[9] In fact, in Table 1 the losses in the Woodford specification are even larger than those in the Calvo specification. But this is due to larger estimated shocks. As Table 2 shows, if we fix the productivity process, then the biggest losses occur under the Calvo model.

4.5 Effects of steady state inflation

While our model is quite successful in reproducing the observed distribution of price changes, a more interesting and more demanding way of testing it is to change the steady state inflation rate, as in Golosov and Lucas (2007). Useful evidence for evaluating the effects of a change in inflation is reported by Gagnon (2007), who analyzes the distribution of price adjustments in Mexico at times of substantially differing inflation rates (annual rates of 4.6%, 29%, and 63%). Gagnon's data on the distribution of price adjustments at each of these inflation rates are shown in Figure 5.

Table 3 and Figure 9 show how the various versions of our model behave when the inflation rate changes from its baseline rate of 0% to a high rate of 63% per annum. In the MC, SDSP, and Woodford models, all the calculations are based on the estimated parameters from Table 1, changing only the aggregate rate of money growth. The upper panel of Figure 9 shows how the frequency of price adjustment rises with the inflation rate. Going from 5% to 63% inflation, the adjustment frequency increases from 10.3% to 14.3% (a 33% log difference) in our smoothly state-dependent model and from 10.6% to 18.2% (a 54% log difference) in the menu cost model. For comparison, in Gagnon's data (see the first panel of Figure 9) an increase in inflation from 5% to 63% raises the monthly frequency of price changes from 25% to 41% (a 49% log difference), lying between the predicted changes from the MC model and the SDSP model.[10]

The Calvo specification is omitted from Figure 9 because it is not even possible to compute equilibrium at all these inflation rates under our baseline parameterization. The problem, of course, is that the Calvo model treats the adjustment probability as a "deep" parameter, unaffected by trend inflation. But with the baseline calibration $\overline{\lambda} = 0.1$, average losses from failure to adjust rise quickly, to 0.99% of median firm value at 4.6% annual inflation, 3.9% of firm value at 10% inflation, and 10.1% of firm value at 15% inflation. Equilibrium fails to exist at 29% and 63% inflation, because the value of the firm becomes negative if we keep $\overline{\lambda}$ fixed at these inflation rates. For illustrative purposes only, Table 3 reports a calculation of the Calvo model with a 63% inflation rate in which we have arbitrarily changed the adjustment probability to $\overline{\lambda} = 0.25$. Even with this adjustment rate, the Calvo model implies huge losses from inaction, averaging 49% of median firm value.

While the fixed menu cost model does a good job of modeling how the adjustment frequency varies with inflation, it is inconsistent with another implication of increased inflation. In Figure 5, and in the second panel of Fig. 9, we see that the standard deviation of price adjustments is roughly unchanged in Gagnon's data as the inflation rate rises from 5% to 63%.[11] In the fixed menu cost model, on the other hand, this change in the inflation rate causes the strongly bimodal distribution of price changes (like the top left panel of Fig. 1) to collapse to an almost unimodal distribution as (see the top left panel of Figure 6). As the distribution collapses to a single peak, there is a large decrease in the standard deviation of price adjustments, from 18.7% to 12.4% (a 41% log difference), in contrast with the Mexican data.

The SDSP model, unfortunately, suffers from the opposite problem: the standard deviation of price changes increases from 12.3% to 19.3% under the same increase in inflation. This problem is even more severe in our estimate of the Woodford model, where the standard deviation of price changes increases from 14.4% to 28%. The fact that the SDSP model and especially the Woodford model err in the opposite direction from the MC model suggests that our preferred estimate may be somewhat understating the true degree of state dependence in the data (while the MC estimate overstates it). Therefore a careful approach to parameterizing a model for policy analysis may require using comparable data from periods with more than one inflation rate in order to correctly calibrate the degree of state dependence. Nonetheless, even though our estimate may be somewhat too close to the Calvo model, the contrast between its performance and that of the Calvo model is remarkable. It is inappropriate in principle to analyze a change in steady state inflation while keeping $\overline{\lambda}$ fixed, and in practice we see that doing so gives absurd results long before we reach a hyperinflationary environment. The SDSP model, on the other hand, yields reasonable results over the whole range of inflation rates observed recently in Mexico, with an overall quantitative performance no worse than that of Golosov and Lucas (2007).

[10] There is a level shift of the frequency of adjustment in Gagnon's data compared to our models, because his data are from a different country and because he does not control for sales. Therefore we focus on the change in frequency caused by inflation, instead of comparing the levels.

[11] In fact the data are mildly nonmonotonic: the standard deviation of price adjustments at first decreases from 12% to 10% and then rises again to 11%.

5 Sticky recursive plans

As we have just seen, our SDSP estimate from a low inflation environment performs reasonably well in predicting pricing behavior under higher inflation rates. But under higher inflation rates, the losses from inaction increase. At 63% inflation, the average loss is 3.2% of median firm value, giving firms a fairly strong incentive to look for a more adequate class of policies. In particular, if we take seriously the bounded rationality interpretation of our model instead of the stochastic menu cost interpretation, then what is costly for the firm is not changing prices *per se*, but making new plans. Therefore, it might be better for the firm to make a simple plan that can take inflation into account, such as choosing a pair (P_i, π_i) consisting of a price level P_i, together with a default inflation rate π_i that it automatically applies to determine next period's price whenever it is unable to fully reoptimize:

$$(P_{it+1}, \pi_{it+1}) = ((1+\pi_{it})P_{it}, \pi_{it})$$

We call this a *recursive plan* since next period's plan (P_{it+1}, π_{it+1}) is determined by this period's plan (P_{it}, π_{it}) without the need for any new information or new decision-making. If the reason for nominal stickiness is the cost of planning, rather than the cost of changing prices, then we can imagine that firms follow a given plan as long as this is relatively close to optimal behavior. More precisely, as in our previous model of sticky prices, it is natural to assume that a firm's probability of changing plans is a smoothly increasing function of the value of changing plans. This leads to a Bellman equation similar to the ones we have seen already, except that the sticky policy is now a pair (P_{it}, π_{it}) instead of a single number P_{it}.

As before, we study an aggregate steady state. We detrend nominal variables by dividing by the current money stock, so that the price at the time of production in period t is $p_{it} \equiv P_{it}/M_{it}$. In real terms, the relation between today's plan and the plan at the beginning of the next period (indicated by tildes) is:

$$(\widetilde{p}_{it+1}, \widetilde{\pi}_{it+1}) = \left(\frac{\widetilde{P}_{it+1}}{M_{it+1}}, \widetilde{\pi}_{it+1}\right) = \left(\frac{(1+\pi_{it})P_{it}}{M_{it+1}}, \pi_{it}\right) = \left(\frac{(1+\pi_{it})M_{it}}{M_{it+1}}\frac{P_{it}}{M_{it}}, \pi_{it}\right) = \left(\mu^{-1}(1+\pi_{it})p_{it}, \pi_{it}\right)$$

Therefore, a firm's real value of producing under its current plan (p, π), given productivity A, is described by the following Bellman equation:

$$v(p,\pi,A) \;=\; \left(p - \frac{w}{A}\right)\vartheta p^{-\epsilon} \;+ R^{-1}E\left\{v(\mu^{-1}(1+\pi)p, \pi, A') + g\left(\mu^{-1}(1+\pi)p, \pi, A'\right)|A\right\} \tag{18}$$

where primes indicate next period's values. In the boundedly rational interpretation of our model, the gains function is

$$g\left(\mu^{-1}(1+\pi)p, \pi, A'\right) \equiv \lambda\left(\frac{d(\mu^{-1}(1+\pi)p, \pi, A')}{w}\right) d\left(\mu^{-1}(1+\pi)p, \pi, A'\right)$$

where

$$d(\mu^{-1}(1+\pi)p, \pi, A') \equiv \max_{p',\pi'} v(p', \pi', A'_i) - v(\mu^{-1}(1+\pi)p, \pi, A')$$

The alternative interpretations of the model impose other functions g analogous to those we saw earlier.

Again, for numerical tractability, we restrict dynamics to a finite grid in real terms. This time, the grid is three-dimensional: $(p_{it}, \pi_{it}, A_{it}) \in \Gamma \equiv \Gamma^p \times \Gamma^\pi \times \Gamma^a$, where Γ^p, Γ^π, and Γ^a are grids of size $\#p$, $\#\pi$, and $\#a$ describing the possible values of p_{it}, π_{it}, and A_{it}, respectively. Again, we write the distributions Ψ and $\widetilde{\Psi}$ and the value function $\mathbf{V}$ as matrices with each row representing a possible policy (p, π) and each column representing a possible shock A, so that the matrices have size $\#p\#\pi \times \#a$. The general equilibrium Bellman equation and distributional dynamics are then formally identical to the sticky price model we studied earlier:

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

$$\Psi_t = (\mathbf{E} - \lambda(\mathbf{D}/w)). * \widetilde{\Psi}_t + \mathbf{P}. * (\mathbf{E} * (\lambda(\mathbf{D}/w). * \widetilde{\Psi}_t))$$

$$\widetilde{\Psi}_{t+1} = \mathbf{R} * \Psi_t * \mathbf{S}'$$

As before, matrix $\mathbf{S}$ describes the exogenous dynamics, so its definition is unchanged. Matrix $\mathbf{R}$, of size $\#p\#\pi \times \#p\#\pi$, now combines three steps that relate time t production prices p_{it} and time $t+1$ beginning-of-period prices $\widetilde{p}_{i,t+1}$: detrending with respect to the monetary growth rate μ, inflating by the idiosyncratic default inflation rate π_{it}, and rounding off stochastically to the nearest grid points in Γ^p. To be precise, consider a firm with policy (p^l, π^m) at the time of production in period t. Then, the only policies that may be in place at this firm with nonzero probability at the beginning of $t+1$ are given by

$$prob(\widetilde{p}_{i,t+1} = p^n, \pi_{i,t+1} = \pi^m | p_{it} = p^l, \pi_{it} = \pi^m) =$$
$$\begin{cases} 1 \text{ if } \mu^{-1}p^l \leq p^1 = p^n \\ \left(\frac{\mu^{-1}(1+\pi^m)p^l - p^{n-1}}{p^n - p^{n-1}}\right) \text{ if } p^1 < p^n = \min\{p \in \Gamma^p : p \geq \mu^{-1}(1+\pi^m)p^l\} \\ \left(\frac{p^{n+1} - \mu^{-1}(1+\pi^m)p^l}{p^{n+1} - p^n}\right) \text{ if } p^1 \leq p^n = \max\{p \in \Gamma^p : p < \mu^{-1}(1+\pi^m)p^l\} \end{cases}$$

These transition probabilities describe the matrix $\mathbf{R}$ for this case (note that $\mathbf{R}$ is huge, but contains mostly zeros, so in the computation we treat it as a sparse matrix). Finally, matrix $\mathbf{P}$ identifies the optimal policy. Column k of matrix $\mathbf{P}$ has a one in the row that represents the optimal policy:

$$(p^{*k}, \pi^{*k}) \equiv \arg \max_{p \in \Gamma^p, \pi \in \Gamma^\pi} v(p, \pi, a^k) \tag{19}$$

and zeros elsewhere.

5.1 Results of the recursive plans model

We simulate the sticky recursive plans model using the productivity process and adjustment probability function we estimated for our sticky price model SDSP. We augment the grid of productivities and prices with an evenly-spaced grid of 25 possible default inflation rates around the steady-state inflation rate (from -12 to +12 steps of 0.002 from the steady-state inflation rate). Figure 10 shows that firms optimally choose to exploit both the price margin and the inflation margin of the planning space. Given our AR(1) technology specification, it is optimal for a firm that experiences a very high productivity shock to set a low price, but also to set a relatively high default inflation rate π. Thus the price tends to trend back to the mean over time, tracking the expected dynamics of the productivity shock.

We simulate the model both for low and high inflation rates. Some sample paths are shown (in real and nominal terms) in Fig. 11, and statistics on price adjustments and the cost of errors are presented in Table 4. Interestingly, at low inflation (4.6% annually), choosing a recursive plan instead of choosing a price only has little effect on profits. But for high inflation (63% annually), the recursive plan specification implies a large increase in value relative to sticky prices: the average loss is 3.2% of median firm value under sticky prices, but only 0.65% under sticky plans. That is, average and median losses rise much more slowly with the inflation rate under recursive plans, giving firms a strong incentive to move to a policy like that of the recursive plans model if the inflation rate rises.

Nevertheless, this model does a poor job of matching the distribution of price adjustments. While a small fraction of firms may optimally choose $\pi_{it} = 0$, if trend inflation is positive the vast majority of firms at any point in time set a nonzero default inflation rate. Therefore, the frequency of price adjustment in our simulations is over 90% at a 4.6% trend inflation, and is virtually 100% per period at 29% or 63% inflation. Also, the distribution of price adjustments is tightly clustered around the monthly inflation rate (see Figure 12). Thus, 93% of the monthly price changes are less than 5% in absolute value in the low inflation simulation, and even at 63% annual inflation, more than 65% of price changes are less than 5% in absolute value.

An obvious problem with this version of the model is that it is written as if there were an unambiguous definition of one "period", which we interpret in our simulation as one month. Since prices are updated once per period in the model, even when the plan is unchanged, the model overpredicts the monthly probability of price adjustment. Perhaps a more flexible interpretation of the model would be more reasonable, since it is not clear how to define one period in the data. But for now we conclude that the sticky price model does a better job of describing microdata on price adjustments than the sticky plans model, even at high inflation rates that give firms a nontrivial incentive to adopt more sophisticated policies.

6 Conclusions

In this paper, we have estimated a model of state-dependent pricing that nests both the exogenous timing (Calvo 1983) and fixed menu cost (Golosov and Lucas 2007) models as limiting cases. Our model is based on two assumptions which we consider undeniably realistic: we assume the probability of adjustment is a smooth function of the value of adjustment, and we assume firms face idiosyncratic shocks. The model fits steady state microdata better than that of Golosov and Lucas, and also provides a simple alternative to extensions of the fixed menu cost framework like Midrigan's (2008) multiproduct setting with leptokurtic shocks and Klenow and Kryvstov's (2008) model with sectoral heterogeneity.

By fitting our model to the microdata reported in Midrigan's paper, we find that the probability of adjustment is nonzero even for small errors, but that it converges very slowly towards one. That is, the behavior of our estimated model is much closer to that of a Calvo model than it is to that of a fixed menu cost model. In particular, we calculate that the degree of state dependence, measured in terms of the variance of the adjustment probability, is close to the Calvo case. While our measure of state dependence (based on cross-sectional variation) differs from that of Klenow and Kryvstov (who use a measure based on time series variation), we provide further support for their claim that state dependence is low in US data.

While our findings constitute a partial defense of the Calvo model, this should not be taken as a green light to apply the Calvo model in all policy analysis contexts. We think the low state dependence in our estimated model helps explain why nominal shocks are observed to have such large real effects. But since the Calvo adjustment probability is not a "deep" parameter, there are many policy questions which cannot even be addressed using the Calvo framework. We have seen that the Calvo model fails completely when we consider changes in steady state inflation as large as those observed in recent Mexican data.

Therefore, while the Calvo model will surely continue to serve as a workhorse in many studies, we believe further exploration of monetary policy under state-dependent pricing is important and is likely to prove fruitful. This will require macroeconomists to face the time-varying distributions that a general equilibrium model of state-dependent pricing implies, perhaps via analytical approximations or perhaps via numerical methods designed to compute equilibrium distributional dynamics. In a companion paper (Costain and Nakov 2008), we argue that calculating distributional dynamics is more tractable than is commonly believed, and we study the effects of monetary policy shocks in simulations calibrated on the basis of this paper's estimates.

References

[1] AKERLOF, G., and J. YELLEN (1985). "A near-rational model of the business cycle with wage and price inertia", *Quarterly Journal of Economics,* 100 (Supplement), pp. 823-838.

[2] ÁLVAREZ, L. (2007). *What do micro price data tell us on the validity of the New Keynesian Phillips curve?,* Banco de España Working Paper 0729.

[3] BASU, S. (2005). "Comment on: Implications of state-dependent pricing for dynamic macroeconomic modelling", *Journal of Monetary Economics,* 52 (1), pp. 243-247.

[4] BEAN, CH. (1993). "Comment on: Microeconomic rigidities and aggregate price dynamics", *European Economic Review,* 37 (4), pp. 712-714.

[5] BURSTEIN, A. (2006). "Inflation and output dynamics with state-dependent pricing decisions", *Journal of Monetary Economics,* 53, pp. 1235-1257.

[6] CABALLERO, R., and E. ENGEL (1993). "Microeconomic rigidities and aggregate price dynamics", *European Economic Review,* 37 (4), pp. 697-711.

[7] – (1999). "Explaining investment dynamics in U.S. manufacturing: a generalized (S,s) approach", *Econometrica,* 67 (4), pp. 783-862.

[8] – (2006). *Price stickiness in Ss models: basic properties,* manuscript, Massachusetts Institute of Technology.

[9] – (2007). "Price stickiness in Ss models: new interpretations of old results", *Journal of Monetary Economics,* 54, pp. 100-121.

[10] CALVO, G. (1983). "Staggered prices in a utility-maximizing framework", *Journal of Monetary Economics,* 12, pp. 383-98.

[11] CAPLIN, A., and J. LEAHY (1991). "State-dependent pricing and the dynamics of money and output", *Quarterly Journal of Economics,* 106 (3), pp. 683-708.

[12] CARVALHO, C. (2006). "Heterogeneity in price stickiness and the real effects of monetary shocks", *Berkeley Electronic Press Frontiers of Macroeconomics,* 2 (1), Article 1.

[13] COSTAIN, J., and A. NAKOV (2008). *Distributional dynamics in a general model of state-dependent pricing,* manuscript, Banco de España.

[14] DORICH, J. (2007). *The welfare losses of price rigidities,* manuscript, Univ. Pompeu Fabra.

[15] DOTSEY, M., R. KING and A. WOLMAN (1999). "State-dependent pricing and the general equilibrium dynamics of money and output", *Quarterly Journal of Economics,* 114 (2), pp. 655-690.

[16] DOTSEY, M., and R. KING (2005). "Implications of state-dependent pricing for dynamic macroeconomic modelling", *Journal of Monetary Economics,* 52 (1), pp. 213-242.

[17] GAGNON, E. (2007). *Price setting under low and high inflation: evidence from Mexico,* manuscript, Federal Reserve Board.

[18] GERTLER, M., and J. LEAHY (2008). "A Phillips curve with an Ss foundation", *Journal of Political Economy,* 116 (3), pp. 533-572.

[19] GOLOSOV, M., and R. E. LUCAS Jr. (2007). "Menu costs and Phillips curves", *Journal of Political Economy,* 115 (2), pp. 171-199.

[20] KASHYAP, A. (1995). "Sticky prices: new evidence from retail catalogs", *Quarterly Journal of Economics,* 110 (1), pp. 45-74.

[21] KLENOW, P., and O. KRYVSTOV (2008). "State-dependent or time-dependent pricing: does it matter for recent US inflation?", *Quarterly Journal of Economics,* 123, pp. 863-904.

[22] MANKIW, N. G. (1985). "Small menu costs and large business cycles: a macroeconomic model of monopoly", *Quarterly Journal of Economics,* 100 (2), pp. 529-537.

[23] MIDRIGAN, V. (2008). *Menu Costs, Multi-Product Firms and Aggregate Fluctuations,* manuscript, New York University

[24] NAKAMURA, E., and J. STEINSSON (2008). "Five facts about prices: a reevaluation of menu cost models", *Quarterly Journal of Economics,* forthcoming.

[25] ROTEMBERG, J. (1982). "Monopolistic price adjustment and aggregate output", *Review of Economic Studies,* 49, pp. 517-531.

[26] SHEEDY, K. (2007a). *Intrinsic inflation persistence,* manuscript, London School of Economics.

[27] – (2007b). *Inflation persistence when price stickiness differs between industries,* manuscript, London School of Economics.

[28] SIMS, CH. (2003). "Implications of rational inattention", *Journal of Monetary Economics,* 50, pp. 665-690.

[29] TAYLOR, J. (1979). "Staggered wage setting in a macro model", *American Economic Review,* 69 (2), pp. 108-113.

[30] – (1993). "Comment on: Microeconomic rigidities and aggregate price dynamics", *European Economic Review,* 37 (4), pp. 714-717.

[31] WOODFORD, M. (2008). *Information-constrained state-dependent pricing,* manuscript, Columbia University.

Table 1. Baseline Estimates, Simulated Moments and Evidence (zero trend inflation)

SDSP (101x101): $\left(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \alpha, \xi\right) = (0.0049, 0.8808, 0.1091, 0.0310, 0.2900)$
SDSP (25x25): $\left(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \alpha, \xi\right) = (0.0049, 0.8812, 0.1089, 0.0311, 0.2937)$
Woodford: $\left(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \alpha, \xi\right) = (0.0085, 0.8596, 0.0946, 0.0609, 1.3341)$
Calvo: $\left(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \xi\right) = (0.0072, 0.8576, 0.10, 0)$
Menu cost: $\left(\sigma_\varepsilon^2, \rho, \alpha, \xi\right) = (0.0059, 0.8469, 0.0631, \infty)$

Model	MC	Calvo	Wdfd	SDSP		Evidence			
	$\xi = \infty$	$\xi = 0$		101x101	25x25	MAC	MD	NS	KK
Frequency of price changes	10	10	10	10	10	20.5	19.2	10	13.9
Mean absolute price change	18.3	6.4	10.1	10.1	10	10.5	7.7		11.3
Median absolute price change	17.9	5.0	8.0	8.7	8.7	8.0	5.3	8.5	9.7
Std of price changes	18.8	8.2	13.2	12.1	11.9	13.2	10.4		
Kurtosis of price changes	1.3	3.4	3.7	2.8	2.7	3.5	5.4		
Percent of price increases	54.8	50.1	50.2	50.5	50.3	50	65.5	66	56
% of price changes $\leq$5% in abs value	0	49.7	35.4	25.2	25.4	25	47		44
Flow of menu cost as % of revenues	1.46								0.50
Mean abs distance from optimal price	5.4	5.9	7.7	6.4	6.4				
Median abs dist. from optimal price	4.7	4.4	5.7	4.8	4.8				
Mean loss as % of median firm value	0.086	0.393	0.504	0.271	0.262				
Median loss as % of median value	0.042	0.114	0.134	0.073	0.074				
Std of loss as % of median firm value	0.098	0.717	0.912	0.576	0.506				
State dependence metric $[0, 1]$	1	0	0.008	0.025	0.025				

Note: All statistics refer to regular consumer price changes excluding sales, and are stated in percent.
The last four columns reproduce the statistics reported by Midrigan (2008) for AC Nielsen (MAC) and Dominick's (MD), Nakamura and Steinsson (2008) (NS), and Klenow and Kryvtsov (2008) (KK)

Table 2. Estimates with a Fixed Productivity Process (zero trend inflation)

SDSP (101x101): $\left(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \alpha, \xi\right) = (0.0049, 0.8808, 0.1091, 0.0310, 0.2900)$
SDSP (25x25): $\left(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \alpha, \xi\right) = (0.0049, 0.8812, 0.1089, 0.0311, 0.2937)$
Woodford: $\left(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \alpha, \xi\right) = (0.0049, 0.8808, 0.0887, 1.8 \times 10^{-6}, 1.7489)$
Calvo: $\left(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \xi\right) = (0.0049, 0.8808, 0.10, 0)$
Menu cost: $\left(\sigma_\varepsilon^2, \rho, \alpha, \xi\right) = (0.0049, 0.8808, 0.0591, \infty)$

Model	MC	Calvo	Wdfd	SDSP		Evidence			
	$\xi = \infty$	$\xi = 0$		101x101	25x25	MAC	MD	NS	KK
Frequency of price changes	10	10	10	10	10	20.5	19.2	10	13.9
Mean absolute price change	17.2	6.0	8.3	10.1	10	10.5	7.7		11.3
Median absolute price change	16.7	4.8	6.5	8.7	8.7	8.0	5.3	8.5	9.7
Std of price changes	17.6	7.7	10.7	12.1	11.9	13.2	10.4		
Kurtosis of price changes	1.2	3.5	3.5	2.8	2.7	3.5	5.4		
Percent of price increases	54.3	50.7	50.1	50.5	50.3	50	65.5	66	56
% of price changes $\leq$5% in abs value	0	52.2	41.4	25.2	25.4	25	47		44
Flow of menu cost as % of revenues	1.38								0.50
Mean abs distance from optimal price	5.2	5.5	6.5	6.4	6.4				
Median abs dist. from optimal price	4.5	4.1	4.8	4.8	4.8				
Mean loss as % of median firm value	0.078	0.349	0.317	0.271	0.262				
Median loss as % of median value	0.043	0.103	0.096	0.073	0.074				
Std of loss as % of median firm value	0.089	0.646	0.544	0.576	0.506				
State dependence metric $[0, 1]$	1	0	0.005	0.025	0.025				

Note: All statistics refer to regular consumer price changes excluding sales, and are stated in percent.
The last four columns reproduce the statistics reported by Midrigan (2008) for AC Nielsen (MAC) and Dominick's (MD), Nakamura and Steinsson (2008) (NS), and Klenow and Kryvtsov (2008) (KK)

Figure 1. Distribution of monthly non-zero price changes*
(zero trend inflation)

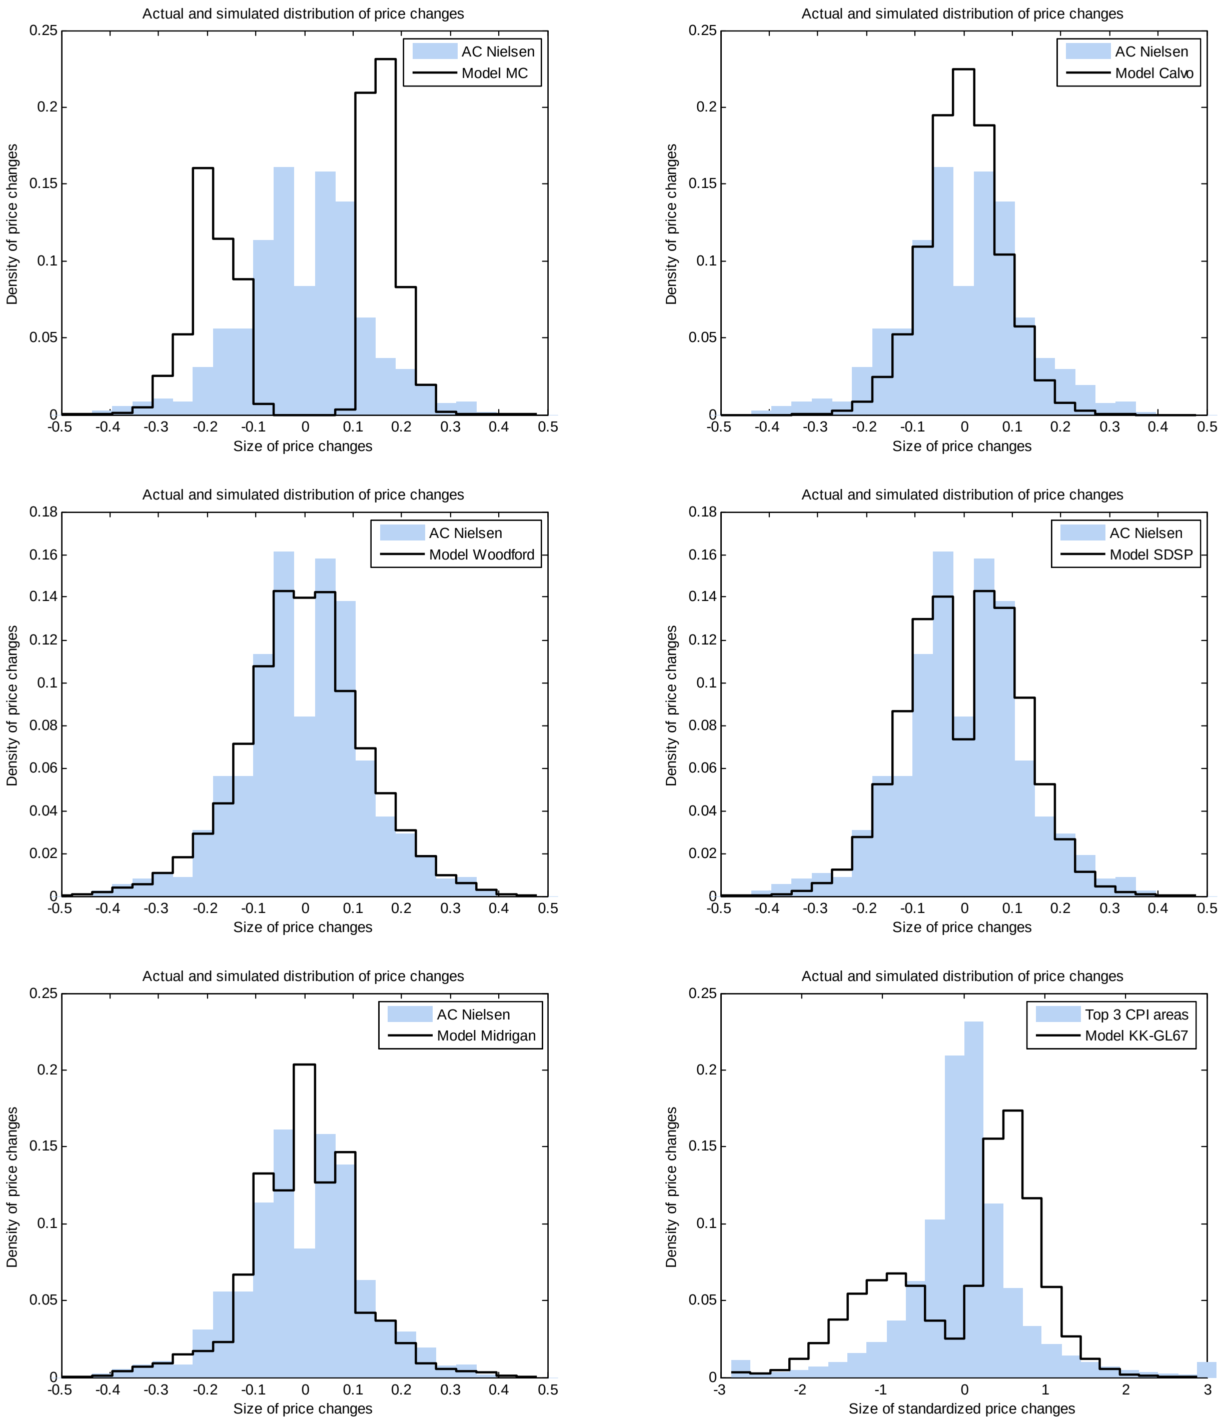

* Horizontal axis: log price change. Vertical axis: frequency.
Shaded area: actual data. First five panels: AC Nielsen data provided by Midrigan (2008).
Last panel: BLS data from top 3 CPI areas provided by Klenow and Kryvtsov (2008).
Solid line: model simulations. First row: fixed menu cost (left) and Calvo (right);
Second row: Woodford (left) and SDSP (right);
Third row: Midrigan (left), and Klenow and Kryvtsov (right).

Figure 2: Value function, distributions, and related objects (SDSP)*
(zero trend inflation)

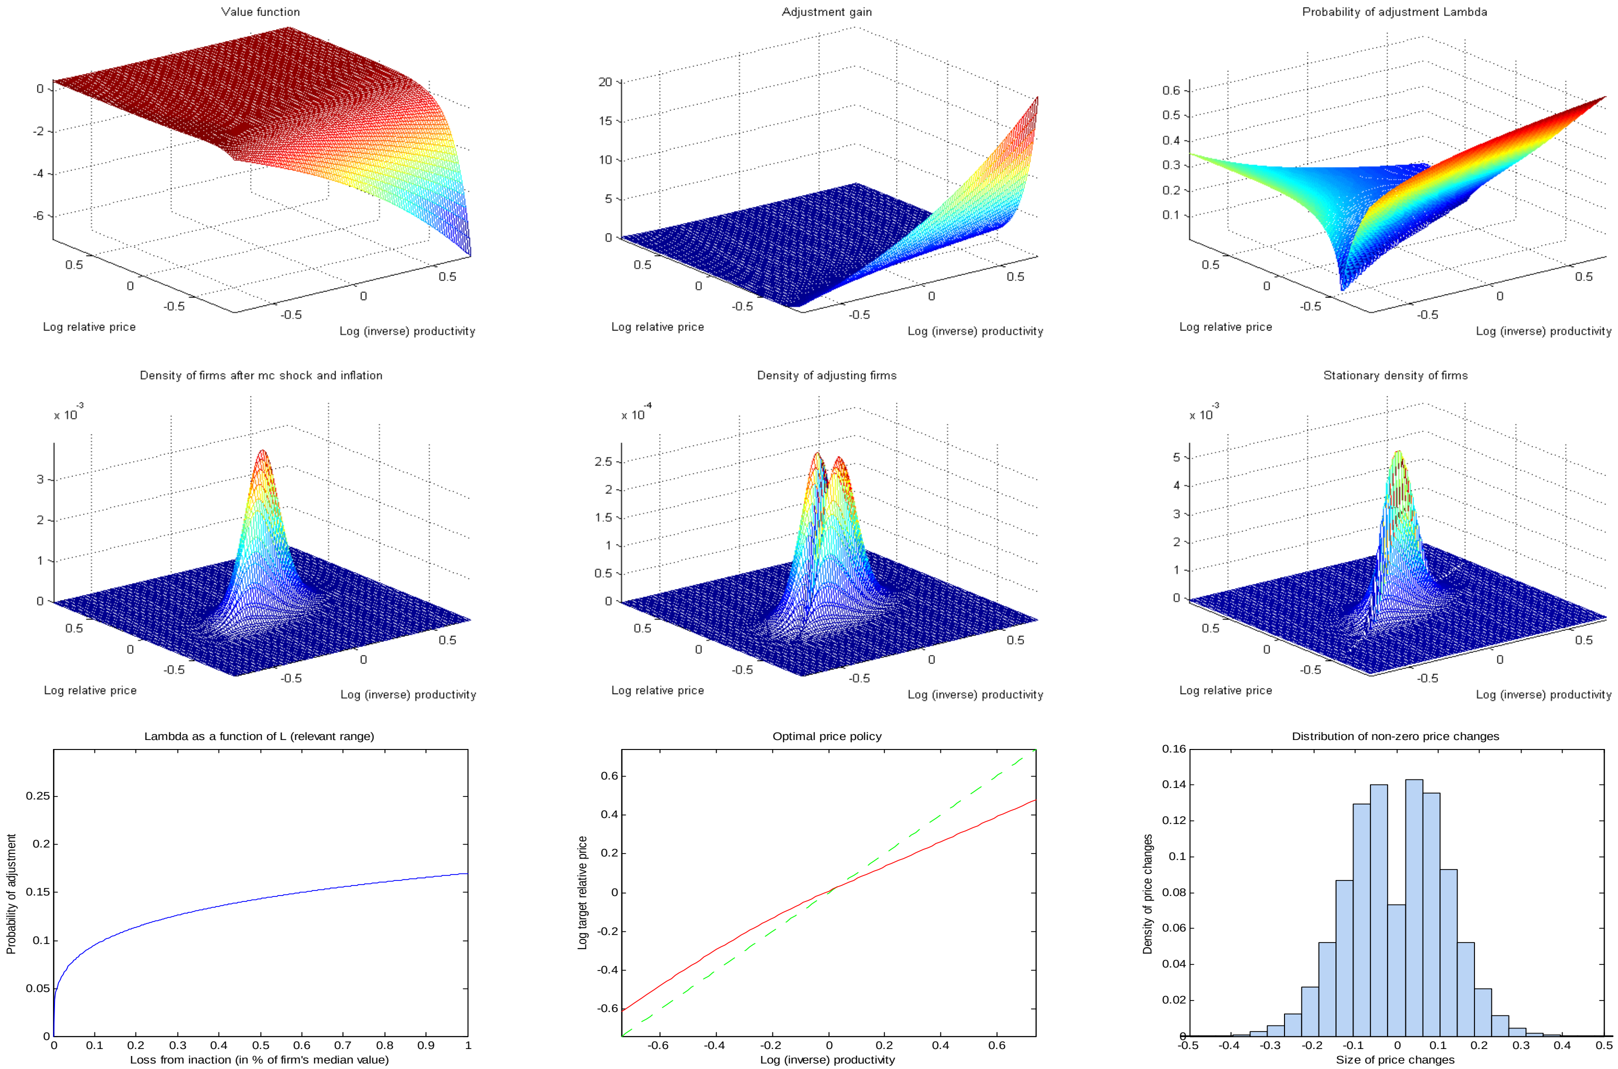

*Simulations from SDSP model. First line: value V, adjustment gain, and adjustment probability lambda, as functions of real price and productivity shock. Second line: beginning of period distribution, adjustment distribution, and distribution at time of production, as functions of real price and productivity shock. Third line: adjustment probability as function of the loss from inaction, policy function, and distribution of monthly non-zero price changes.

Figure 3. Distribution of losses from failure to adjust and adjustment function*
(zero trend inflation)

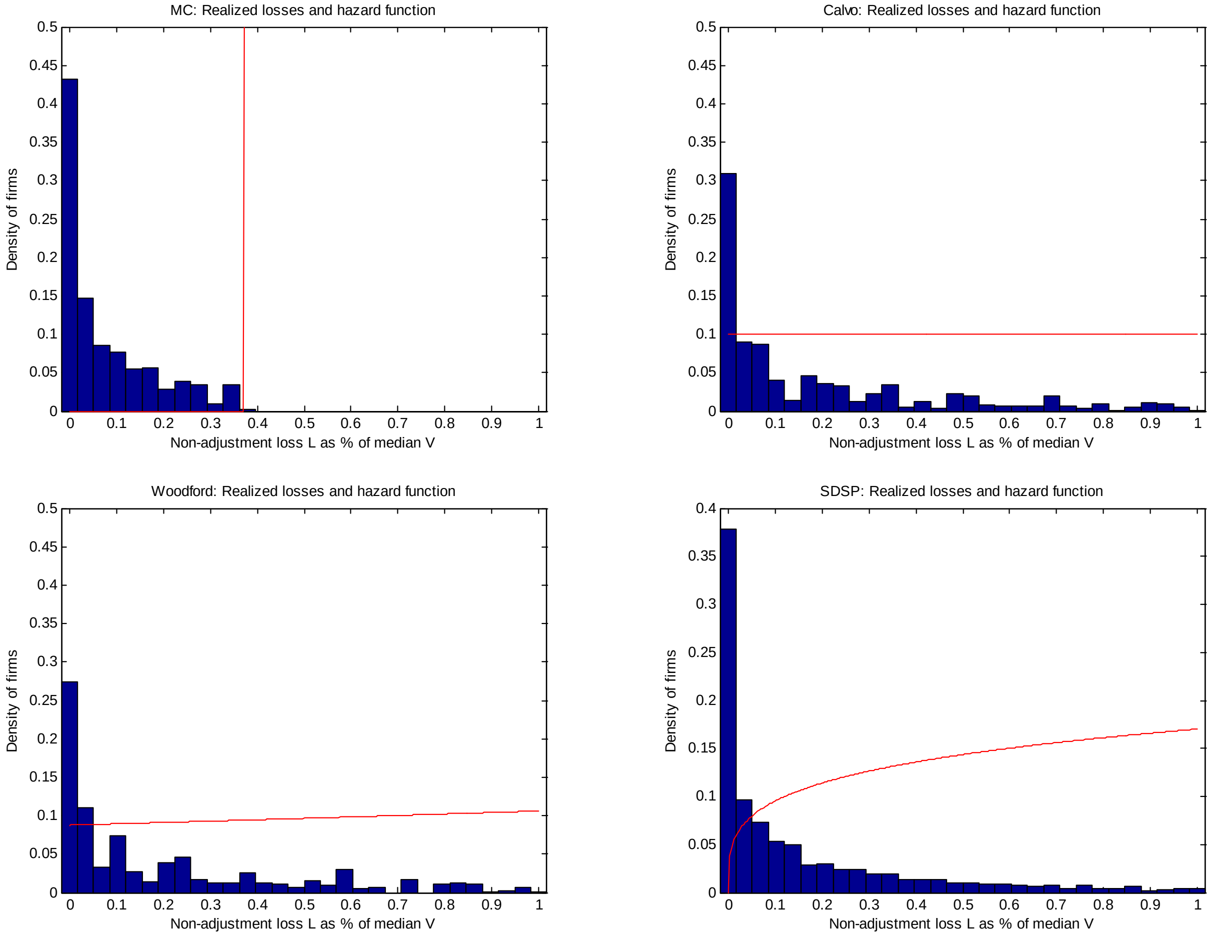

*Horizontal axis: average loss L from failure to adjust, as percentage of median firm value V. Vertical axis: frequency.

Figure 4: Policy functions: firm's price as a function of productivity shock*
(zero trend inflation)

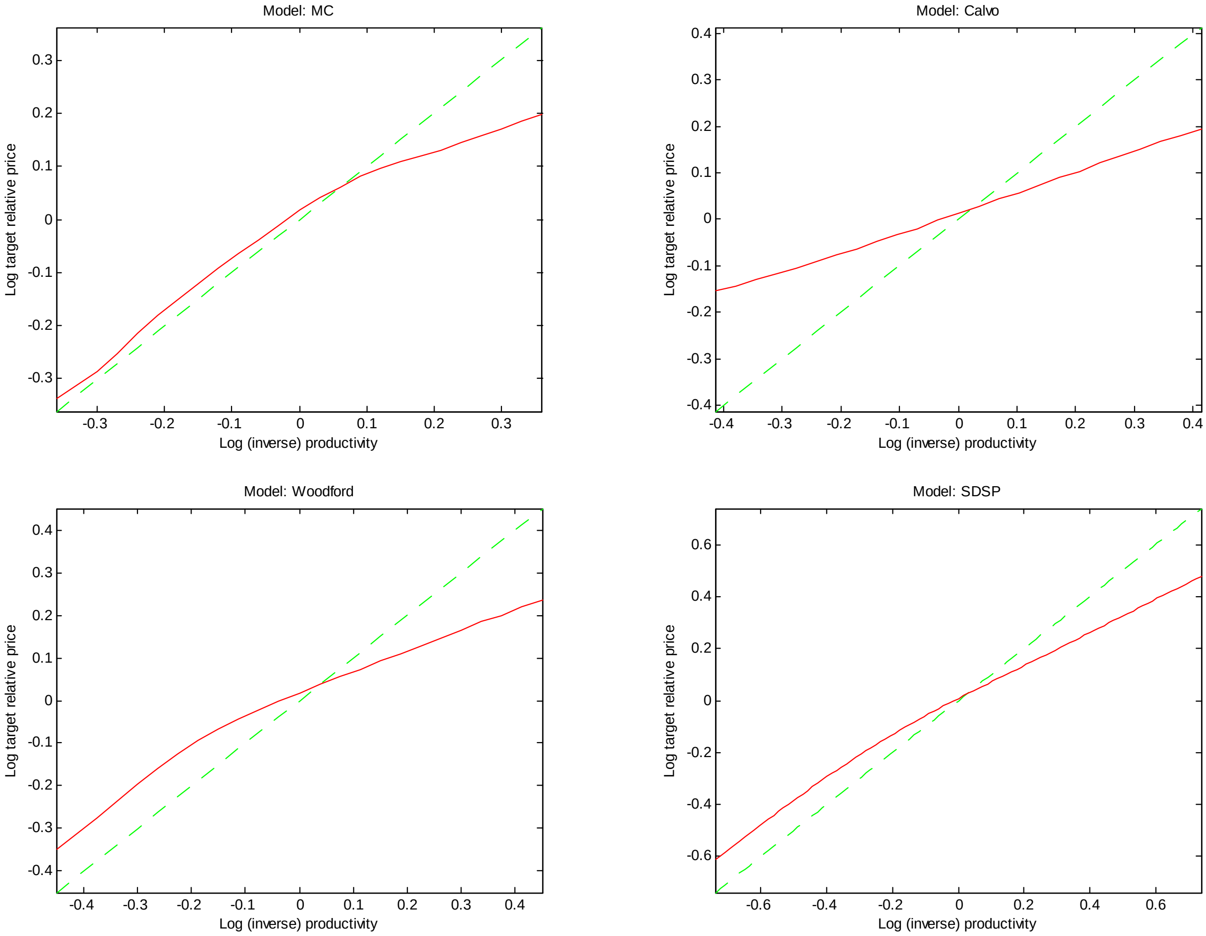

*Dashed line shows policy function under flexible prices (log price as function of log inverse productivity)
Solid line shows policy function under sticky prices.

Figure 5: Distribution of monthly non-zero price changes in Mexico (Gagnon 2007)*
(annualized inflation rates 4.6%, 28.9%, 63.2%)

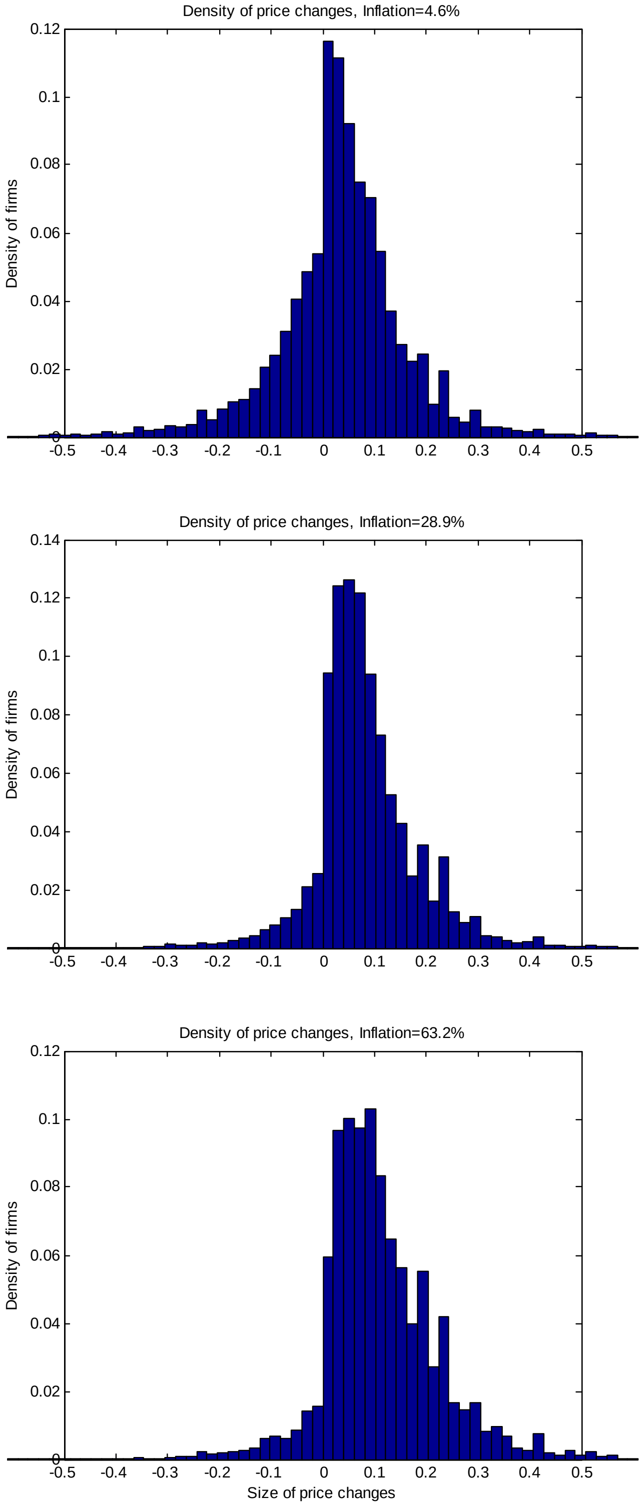

*Monthly log price changes of nonregulated goods, including some sale prices,
from three months with low, medium, and high inflation rates. Data provided by Etienne Gagnon (2007)
Horizontal axis: percent price change. Vertical axis: frequency.

Table 3: Simulations with High Inflation (63% annual)

SDSP (101x101): $\left(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \alpha, \xi\right) = (0.0049, 0.8808, 0.1091, 0.0310, 0.2900)$
SDSP (25x25): $\left(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \alpha, \xi\right) = (0.0049, 0.8812, 0.1089, 0.0311, 0.2937)$
Woodford: $\left(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \alpha, \xi\right) = (0.0085, 0.8596, 0.0946, 0.0609, 1.3341)$
Calvo: $\left(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \xi\right) = (0.0072, 0.8576, 0.25, 0)$
Menu cost: $\left(\sigma_\varepsilon^2, \rho, \alpha, \xi\right) = (0.0059, 0.8469, 0.0631, \infty)$

Model	MC	Calvo	Wdfd	SDSP	
	$\xi = \infty$	$\xi = 0$		101x101	25x25
Frequency of price changes	18.2	25	12.8	14.3	14.2
Mean absolute price change	24.9	16.3	34.9	29.1	29.6
Median absolute price change	24.7	12.1	31.7	26.2	26.4
Std of price changes	12.4	14.5	28	19.4	20.3
Kurtosis of price changes	10.0	8.7	2.4	3.3	3.3
Percent of price increases	93.8	99	87.2	95.4	93.9
% of price changes $\leq$5% in abs value	0.002	18.3	10	5.5	6.4
Flow of menu cost as % of revenues	2.7				
Mean abs distance from optimal price	7.6	12.6	20.6	16.2	16.5
Median abs dist. from optimal price	6.6	7.9	15.6	11.7	11.1
Mean loss as % of median firm value	0.099	49.1	3.25	3.2	3.3
Median loss as % of median value	0.049	0.207	0.673	0.316	0.299
Std of loss as % of median firm value	0.116	1e5	5.82	13.5	14.1
State dependence metric $[0, 1]$	1	0	0.068	0.051	0.053

*All statistics are stated in percentage points. Structural parameters are the same as in Table 1, except for Calvo, where Calvo parameter is raised to 0.25 to ensure equilibrium existence.

Figure 6: Distribution of monthly nonzero price changes*
(trend inflation = 63% annual)

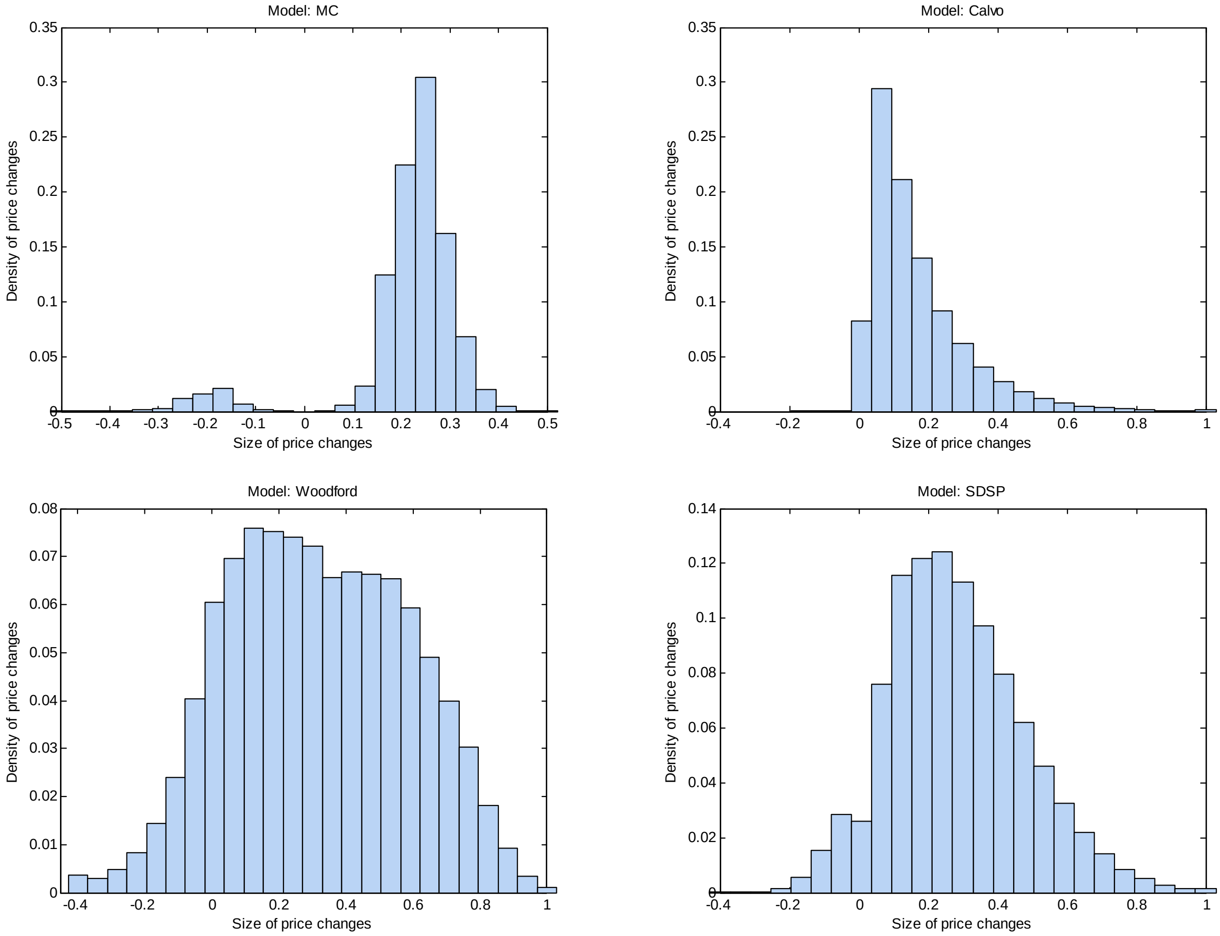

*Horizontal axis: log price change. Vertical axis: frequency.

Note: in Calvo specification, adjustment parameter is raised to 25% monthly.

Figure 7: Distribution of losses from failure to adjust
(trend inflation = 63% annual)

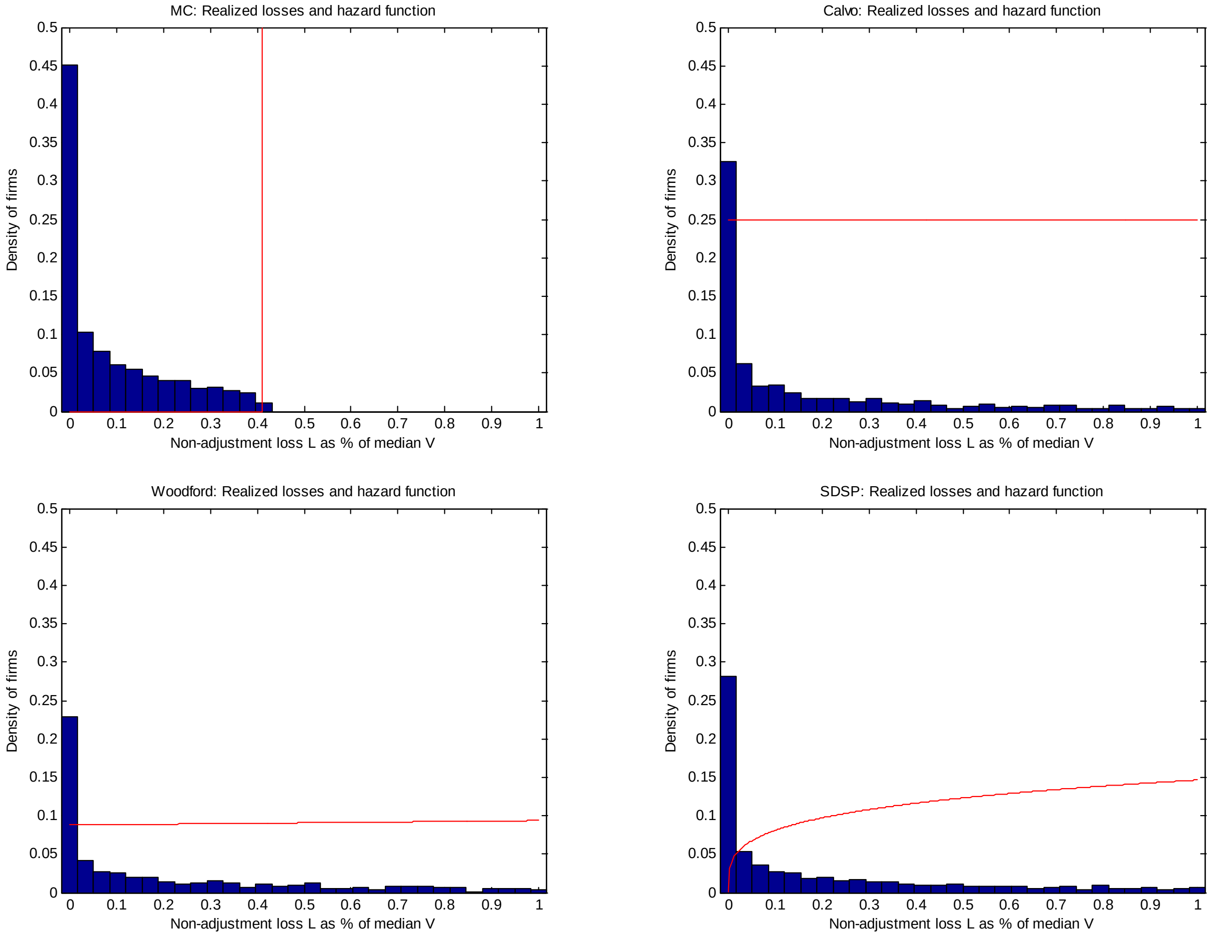

*Horizontal axis: average loss L from failure to adjust, as percentage of median firm value V. Vertical axis: frequency.

Figure 8: Policy functions: firm's price as a function of cost shock
(trend inflation = 63% annual)

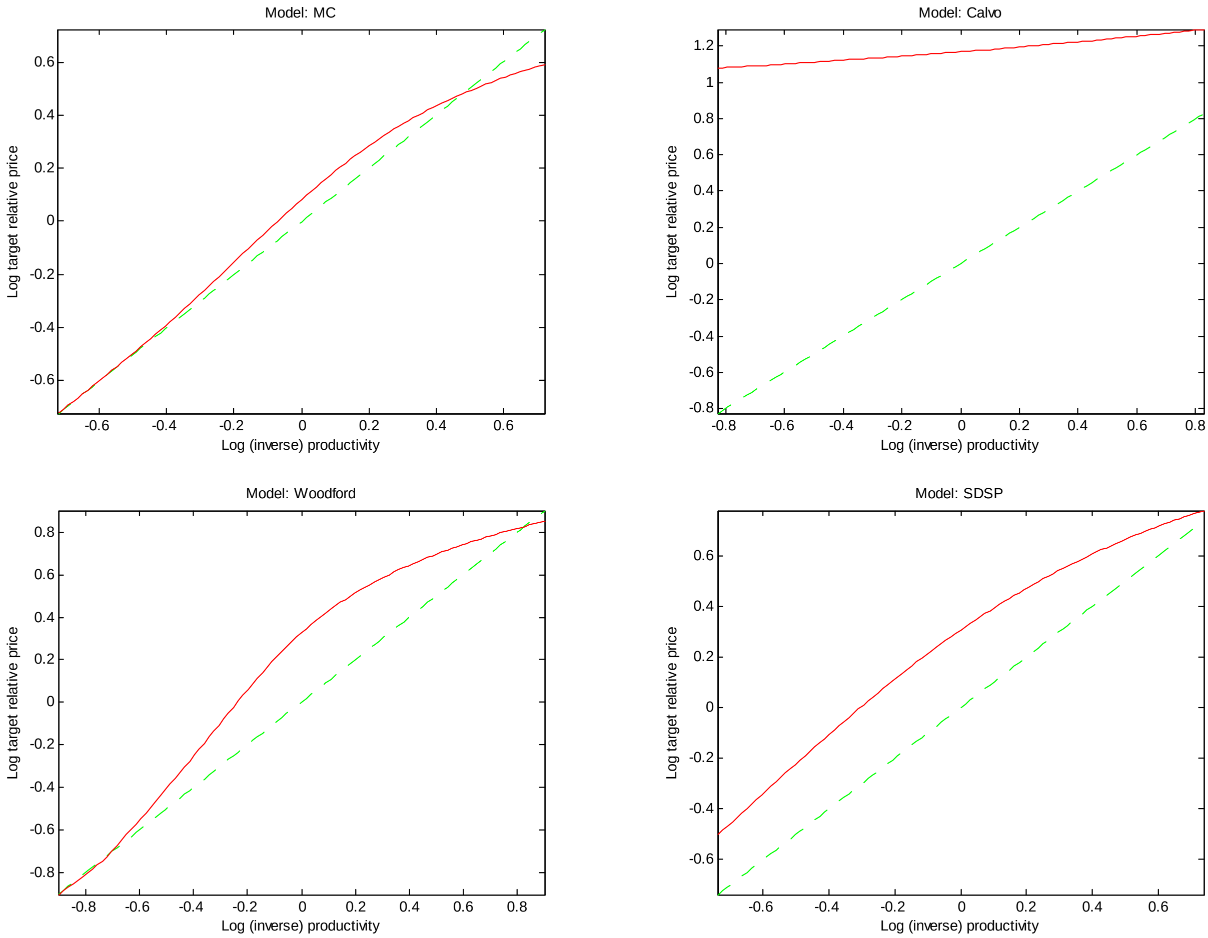

*Dashed line shows policy function under flexible prices (log price as function of log cost). Solid line shows policy function under sticky prices.

Note: in Calvo specification, adjustment parameter is raised to 25% monthly.

Figure 9: Effect of trend inflation on frequency and size of adjustments
(Gagnon low, medium and high)

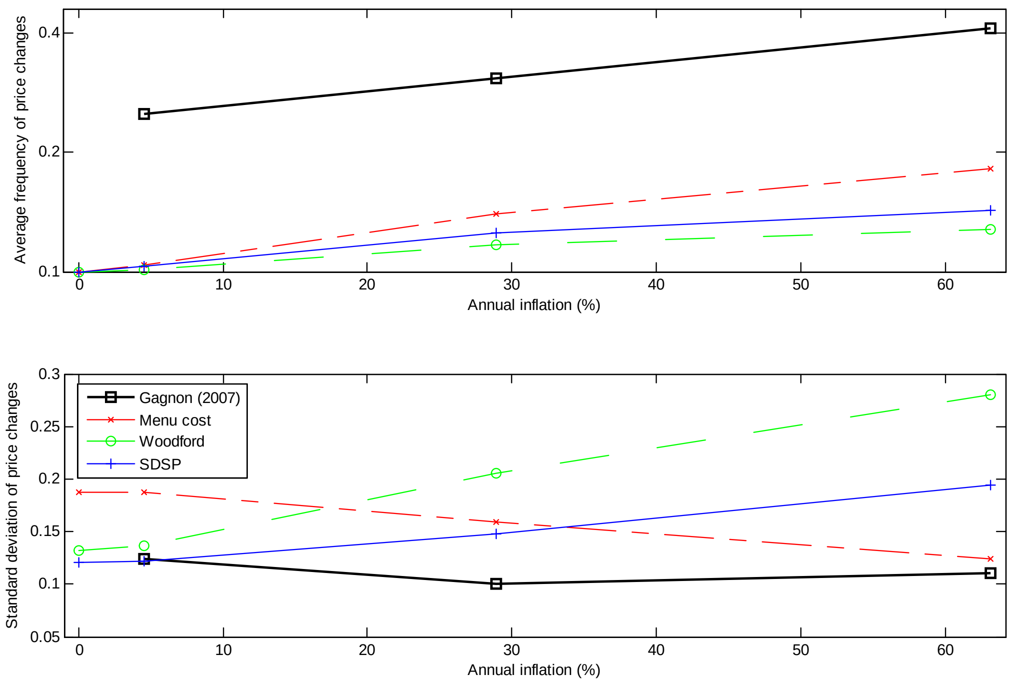

*Bold black line: Mexican data provided by Etienne Gagnon. Other lines: model simulations. Horizontal axis: annualized inflation rate. Top graph shows frequency of nonzero price changes (on a log scale); bottom graph shows standard deviation of log price change.

Table 4. Recursive (p, π) plans (low. medium, and high inflation)

SDSP: $\left(\sigma_{\varepsilon}^{2}, \rho, \bar{\lambda}, \alpha, \xi\right) = (0.0049, 0.8812, 0.1089, 0.0311, 0.2937)$

	Annual inflation rate		
	$\pi = 4.6\%$	$\pi = 28.9\%$	$\pi = 63.1\%$
Frequency of price changes	90.3	100	100
Mean absolute price change	1.58	3.01	5.71
Median absolute price change	0.369	2.12	4.08
Std of price changes	3.97	3.8	5.5
Kurtosis of price changes	3330	2980	2150
Percent of price increases	92.3	93.1	93.7
% of price changes $\leq$5% in abs value	93.4	85.9	65.5
Mean loss as % of median firm value	0.321	0.285	0.654
Median loss as % of median value	0.131	0.123	0.252
Std of loss as % of median firm value	0.553	0.449	0.909

Note: All statistics refer to regular consumer price changes and are stated in percent.

Figure 10. (p, π) policy functions: price and inflation as function of cost shock

SDSP, $\pi = 4.6\%$

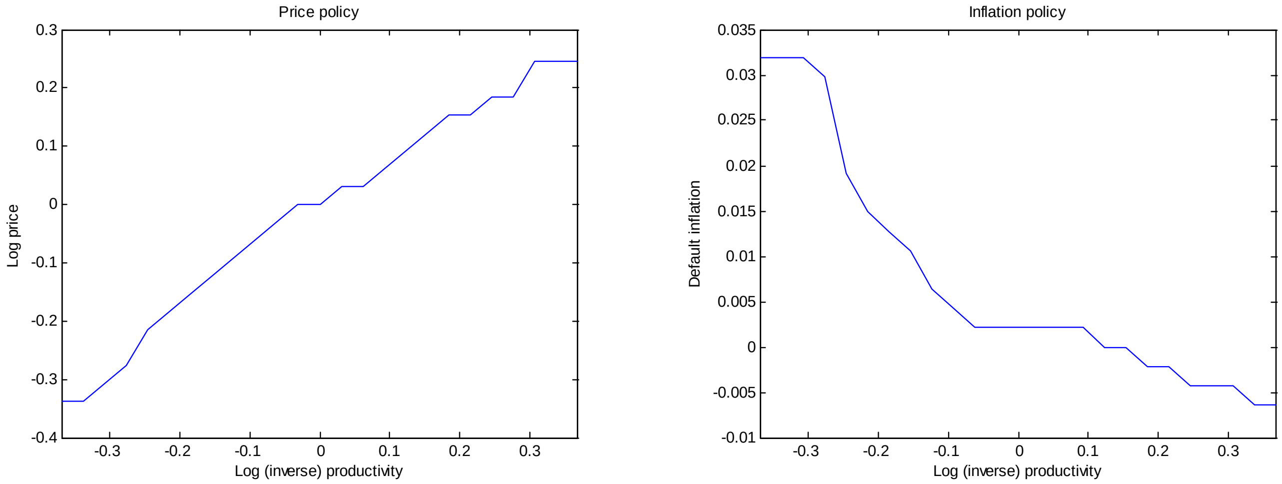

SDSP, $\pi = 28.9\%$

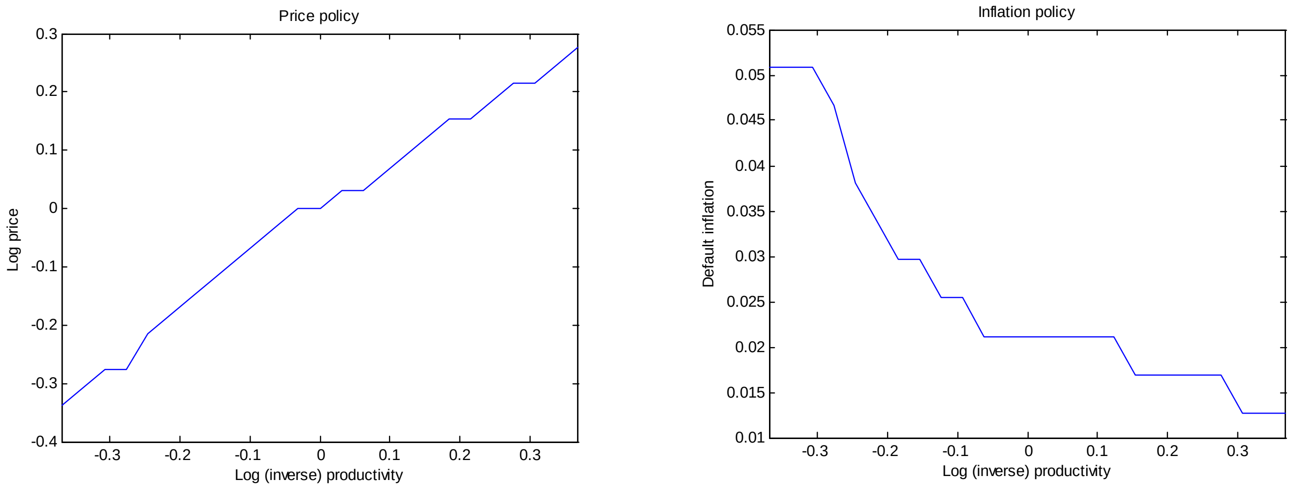

SDSP, $\pi = 63.1\%$

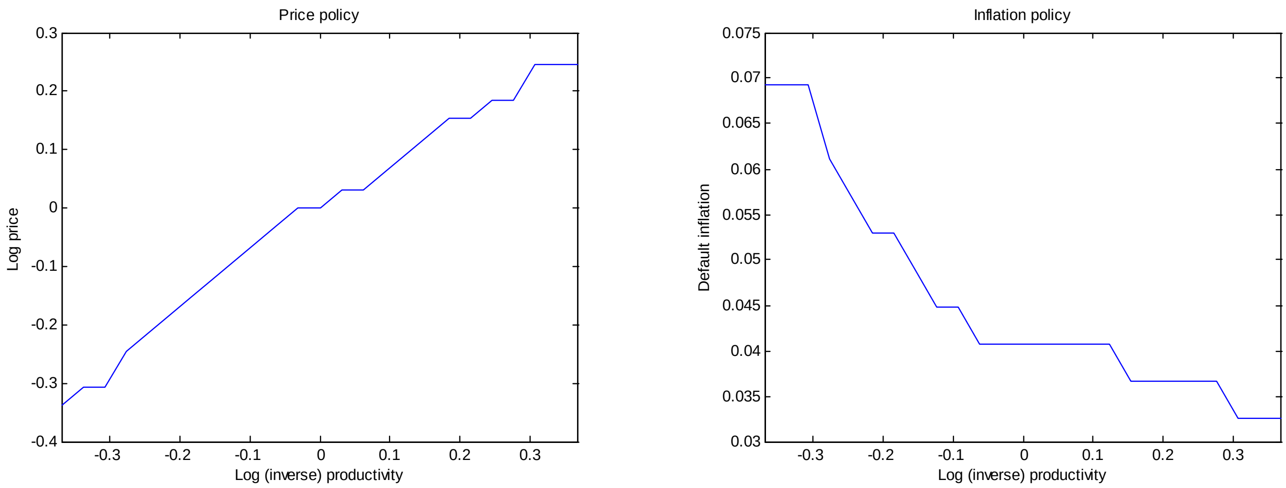

Figure 11. Simulation: price paths under recursive (p, π) policies

Ten price histories, $\pi = 4.6\%$ Ten price histories, $\pi = 63.1\%$

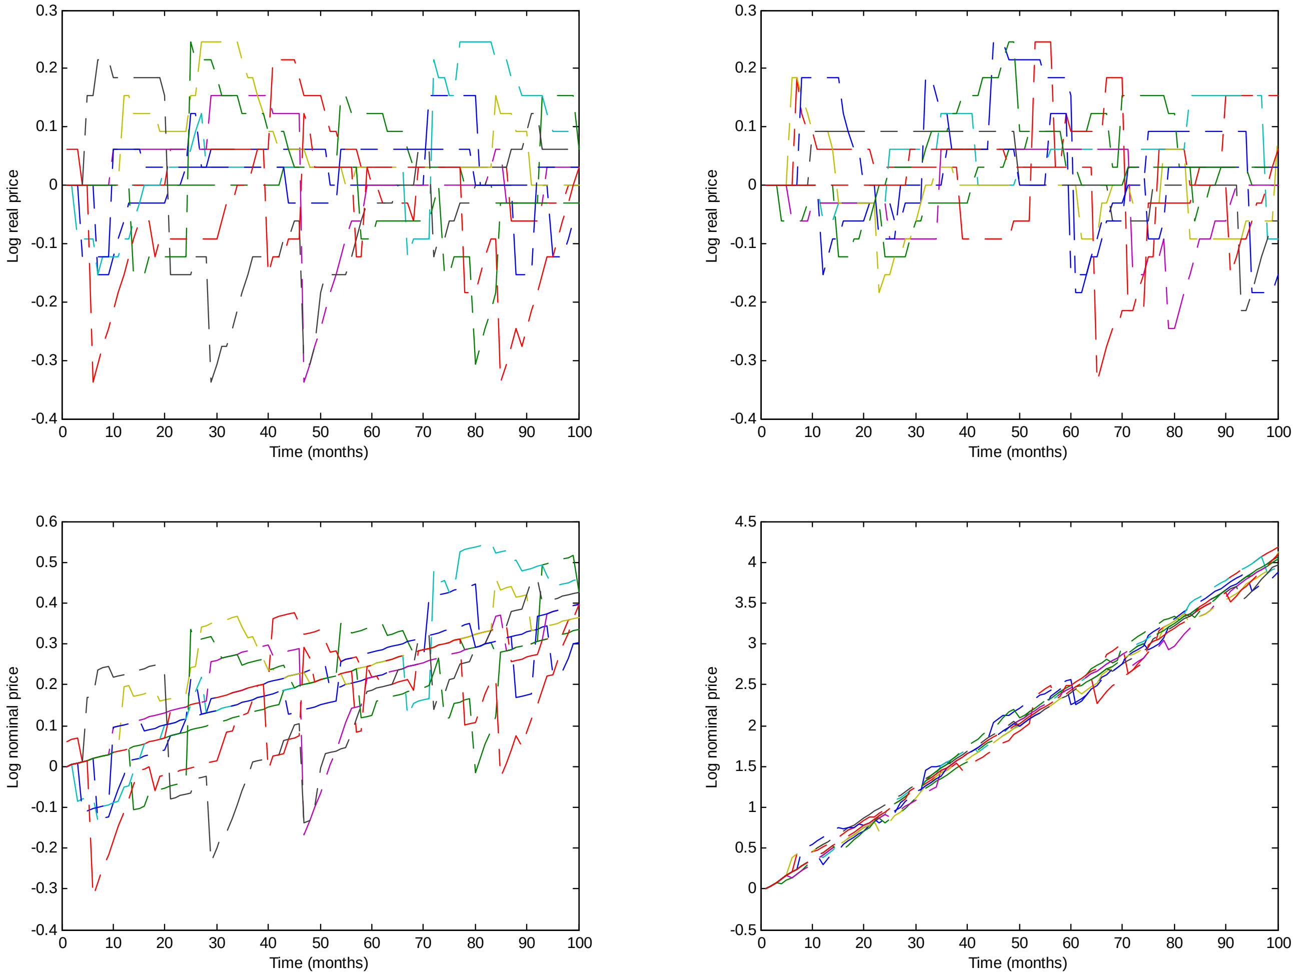

Figure 12 Distribution of monthly nonzero price changes: recursive (p, π) policy with SDSP parameters

$\pi = 4.6\%$

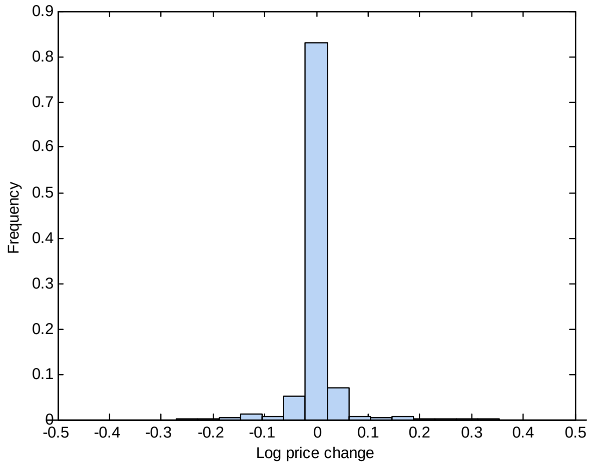

$\pi = 28.9\%$

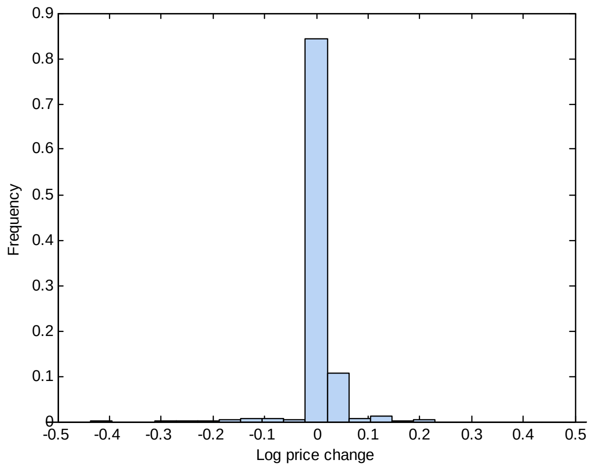

$\pi = 63.1\%$

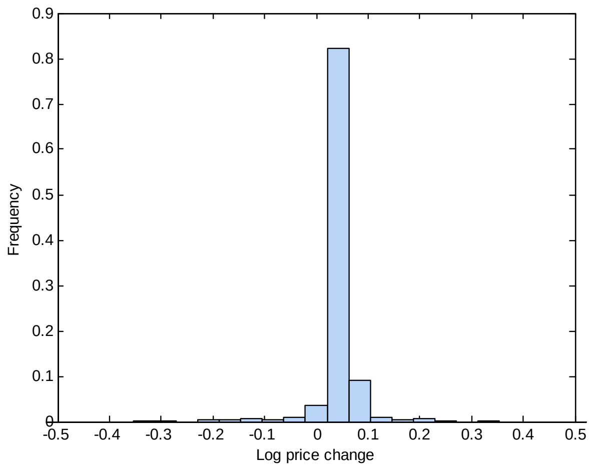

BANCO DE ESPAÑA PUBLICATIONS

WORKING PAPERS[1]

0701 PRAVEEN KUJAL AND JUAN RUIZ: Cost effectiveness of R&D and strategic trade policy.

0702 MARÍA J. NIETO AND LARRY D. WALL: Preconditions for a successful implementation of supervisors' prompt corrective action: Is there a case for a banking standard in the EU?

0703 PHILIP VERMEULEN, DANIEL DIAS, MAARTEN DOSSCHE, ERWAN GAUTIER, IGNACIO HERNANDO, ROBERTO SABBATINI AND HARALD STAHL: Price setting in the euro area: Some stylised facts from individual producer price data.

0704 ROBERTO BLANCO AND FERNANDO RESTOY: Have real interest rates really fallen that much in Spain?

0705 OLYMPIA BOVER AND JUAN F. JIMENO: House prices and employment reallocation: International evidence.

0706 ENRIQUE ALBEROLA AND JOSÉ M.ª SERENA: Global financial integration, monetary policy and reserve accumulation. Assessing the limits in emerging economies.

0707 ÁNGEL LEÓN, JAVIER MENCÍA AND ENRIQUE SENTANA: Parametric properties of semi-nonparametric distributions, with applications to option valuation.

0708 ENRIQUE ALBEROLA AND DANIEL NAVIA: Equilibrium exchange rates in the new EU members: external imbalances vs. real convergence.

0709 GABRIEL JIMÉNEZ AND JAVIER MENCÍA: Modelling the distribution of credit losses with observable and latent factors.

0710 JAVIER ANDRÉS, RAFAEL DOMÉNECH AND ANTONIO FATÁS: The stabilizing role of government size.

0711 ALFREDO MARTÍN-OLIVER, VICENTE SALAS-FUMÁS AND JESÚS SAURINA: Measurement of capital stock and input services of Spanish banks.

0712 JESÚS SAURINA AND CARLOS TRUCHARTE: An assessment of Basel II procyclicality in mortgage portfolios.

0713 JOSÉ MANUEL CAMPA AND IGNACIO HERNANDO: The reaction by industry insiders to M&As in the European financial industry.

0714 MARIO IZQUIERDO, JUAN F. JIMENO AND JUAN A. ROJAS: On the aggregate effects of immigration in Spain.

0715 FABIO CANOVA AND LUCA SALA: Back to square one: identification issues in DSGE models.

0716 FERNANDO NIETO: The determinants of household credit in Spain.

0717 EVA ORTEGA, PABLO BURRIEL, JOSÉ LUIS FERNÁNDEZ, EVA FERRAZ AND SAMUEL HURTADO: Update of the quarterly model of the Bank of Spain. (The Spanish original of this publication has the same number.)

0718 JAVIER ANDRÉS AND FERNANDO RESTOY: Macroeconomic modelling in EMU: how relevant is the change in regime?

0719 FABIO CANOVA, DAVID LÓPEZ-SALIDO AND CLAUDIO MICHELACCI: The labor market effects of technology shocks.

0720 JUAN M. RUIZ AND JOSEP M. VILARRUBIA: The wise use of dummies in gravity models: Export potentials in the Euromed region.

0721 CLAUDIA CANALS, XAVIER GABAIX, JOSEP M. VILARRUBIA AND DAVID WEINSTEIN: Trade patterns, trade balances and idiosyncratic shocks.

0722 MARTÍN VALLCORBA AND JAVIER DELGADO: Determinantes de la morosidad bancaria en una economía dolarizada. El caso uruguayo.

0723 ANTÓN NÁKOV AND ANDREA PESCATORI: Inflation-output gap trade-off with a dominant oil supplier.

0724 JUAN AYUSO, JUAN F. JIMENO AND ERNESTO VILLANUEVA: The effects of the introduction of tax incentives on retirement savings.

0725 DONATO MASCIANDARO, MARÍA J. NIETO AND HENRIETTE PRAST: Financial governance of banking supervision.

0726 LUIS GUTIÉRREZ DE ROZAS: Testing for competition in the Spanish banking industry: The Panzar-Rosse approach revisited.

0727 LUCÍA CUADRO SÁEZ, MARCEL FRATZSCHER AND CHRISTIAN THIMANN: The transmission of emerging market shocks to global equity markets.

0728 AGUSTÍN MARAVALL AND ANA DEL RÍO: Temporal aggregation, systematic sampling, and the Hodrick-Prescott filter.

0729 LUIS J. ÁLVAREZ: What do micro price data tell us on the validity of the New Keynesian Phillips Curve?

0730 ALFREDO MARTÍN-OLIVER AND VICENTE SALAS-FUMÁS: How do intangible assets create economic value? An application to banks.

0731 REBECA JIMÉNEZ-RODRÍGUEZ: The industrial impact of oil price shocks: Evidence from the industries of six OECD countries.

1. Previously published Working Papers are listed in the Banco de España publications catalogue.

BANCO DE ESPAÑA
Eurosistema

Unidad de Publicaciones
Alcalá, 522; 28027 Madrid
Telephone +34 91 338 6363. Fax +34 91 338 6488
e-mail: publicaciones@bde.es
www.bde.es