EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 1333 / APRIL 2011

DISTRIBUTIONAL DYNAMICS UNDER SMOOTHLY STATE-DEPENDENT PRICING

by James Costain
and Anton Nakov

EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 1333 / APRIL 2011

DISTRIBUTIONAL DYNAMICS UNDER SMOOTHLY STATE-DEPENDENT PRICING[1]

by James Costain[2]
and Anton Nakov[3]

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

1 An earlier version of this paper circulated as "Dynamics of the price distribution in a general model of state-dependent pricing". The authors thank R. Bachmann, M. Dotsey, O. Licandro, B. Mackowiak, V. Midrigan, A. Reiff, M. Reiter, R. Wouters, and K. Sheedy for helpful comments, as well as seminar participants at IAS Vienna, EUI, ECB, CERGE-EI, ESSIM 2008, CEF 2008, REDg 2008, SNDE 2009, SED 2009, Banco de España "Workshop on Monetary Policy" (2009), and Banque de France "Understanding Price Dynamics" (2009). We especially thank V. Midrigan, E. Gagnon, and O. Kryvtsov for providing data. The views expressed in this paper are those of the authors and do not necessarily coincide with those of the Bank of Spain.

2 Corresponding author: Banco de España, Calle Alcalá 48, 28014 Madrid, Spain; phone number: +34-91-338-5732; e-mail: james.costain@bde.es

3 Banco de España and European Central Bank, Kaisestrasse 29, D-60311 Frankfurt am Main, Germany; e-mail: anton.nakov@ecb.europa.eu

Address
Kaiserstrasse 29
60311 Frankfurt am Main, Germany

Postal address
Postfach 16 03 19
60066 Frankfurt am Main, Germany

Telephone
+49 69 1344 0

Internet
http://www.ecb.europa.eu

Fax
+49 69 1344 6000

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

ISSN 1725-2806 (online)

CONTENTS

Abstract

Starting from the assumption that firms are more likely to adjust their prices when doing so is more valuable, this paper analyzes monetary policy shocks in a DSGE model with firm-level heterogeneity. The model is calibrated to retail price microdata, and inflation responses are decomposed into "intensive", "extensive", and "selection" margins. Money growth and Taylor rule shocks both have nontrivial real effects, because the low state dependence implied by the data rules out the strong selection effect associated with fixed menu costs. The response to firm-specific shocks is gradual, though inappropriate econometrics might make it appear immediate.

Keywords: Nominal rigidity, state-dependent pricing, menu costs, heterogeneity, Taylor rule

JEL classification: E31, E52, D81

Non-Technical Summary

Sticky prices are an important ingredient in modern dynamic general equilibrium models, including those used by central banks for policy analysis. But how best to model price stickiness, and to what extent stickiness of individual prices implies rigidity of the aggregate price level, remains as controversial as ever. Calvo (1983) proposed — only for simplicity — assuming that the probability of price adjustment remains constant over time. More realistically, the probability of price adjustment will be "state-dependent"(it will depend on the current situation of the firm), which is true, for example, if a firm must pay a constant "menu cost" each time it adjusts its price. This difference matters: shocks to monetary policy stimulate the real economy strongly in models based on Calvo's assumption, whereas Golosov and Lucas (2007) have shown that the real effects of monetary shocks are very small in a menu cost model.

This paper studies a general model of state-dependent pricing that nests the Calvo (1983) and menu cost models as two opposite limiting cases. Our setup rests on one fundamental assumption: firms are more likely to adjust their prices when doing so is more valuable. The parameters of the relationship between the probability of price adjustment and the value of price adjustment are chosen so that the distribution of price changes predicted by the model resembles the distribution observed in recent microeconomic data from the US retail sector. When we choose the parameters this way, we find that the degree of state dependence is quite low: that is, there is relatively little variation in the probability of price adjustment. Therefore, when we simulate the effects of monetary policy shocks we find substantial monetary non-neutrality, with real effects only slightly weaker than the Calvo model implies.

Allowing the adjustment probability to be state-dependent makes it harder to calculate macroeconomic dynamics, because it requires keeping track of the specific situations of all the firms in the economy; one cannot simply consider a "representative"firm. We employ a recent method for computing heterogeneous agent economies (Reiter, 2009) which is well-suited to the context of state-dependent pricing. The method describes the distribution of firms nonparametrically, calculating a histogram of firm characteristics in each time period.

Our computational method allows us to calculate how various cross-sectional statistics change over time. In particular, we decompose inflation into an "intensive margin" relating to the average desired price change, an "extensive margin" relating to the fraction of firms adjusting, and a "selection effect" relating to *which* firms adjust. Our decomposition vindicates the claim of Golosov and Lucas (challenged by Caballero and Engel, 2007), that the strength of the selection effect is the main reason why money shocks are so much weaker in the menu cost setup, as compared with other models of sticky prices.

Our calculations also address a number of issues not considered in previous papers on state-dependent pricing. We show that the shape and persistence of the macroeconomic response to money supply shocks depends mainly on the degree of state dependence, not on the degree of autocorrelation of the money process. Likewise, assuming a realistic degree of state dependence, monetary policy has strong real effects regardless of whether it is described by a money growth rule or by a Taylor rule. A final contribution of this paper is to calculate the response of prices to idiosyncratic as well as aggregate shocks.

1. Introduction

Sticky prices are an important ingredient in modern dynamic general equilibrium models, including those used by central banks for policy analysis. But how best to model price stickiness, and to what extent stickiness of individual prices implies rigidity of the aggregate price level, remains as controversial as ever. Calvo's (1983) assumption of a constant adjustment probability is popular for its analytical tractability, and implies that monetary shocks have large and persistent real effects. However, Golosov and Lucas (2007, henceforth GL07) have argued that microfounding price rigidity on a fixed "menu cost"and calibrating to microdata implies that monetary shocks are almost neutral.

This paper calibrates and simulates a general model of state-dependent pricing that nests the Calvo (1983) and fixed menu cost (FMC) models as two opposite limiting cases, with a continuum of smooth intermediate cases lying between them. As in Dotsey, King, and Wolman (1999) and Caballero and Engel (2007), the setup rests on one fundamental property: firms are more likely to adjust their prices when doing so is more valuable. Implementing this assumption requires the selection of a parameterized family of functions to describe the adjustment hazard; the exercise is disciplined by fitting the model to the size distribution of price changes found in recent US retail microdata (Klenow and Kryvtsov 2008; Midrigan 2008; Nakamura and Steinsson 2008).[1] One of the calibrated parameters controls the degree of state dependence; matching the smooth distribution of price changes seen in microdata requires rather low state dependence. Therefore, an impulse response analysis of the effects of monetary policy shocks reveals substantial monetary nonneutrality, with real effects only slightly weaker than the Calvo model implies.

The impulse response analysis considers a number of issues unaddressed by previous work on state-dependent pricing. GL07 restricted attention to *iid* money growth shocks; this paper also considers the autocorrelated case, and shows that the shape and persistence of responses is primarily determined by the degree of state dependence, not by the autocorrelation of the driving process. Moreover, this paper also studies monetary policy governed by a Taylor rule, as opposed to an exogenous money growth process, which reinforces the conclusion that a calibrated model of state-dependent pricing

[1] A companion paper, Costain and Nakov (2008), discusses the calibration in greater detail, documenting the steady-state model's fit to cross-sectional microdata on price adjustments, both for low and high trend inflation rates.

has nontrivial real effects. This paper also decomposes inflation into an "intensive margin" relating to the average desired price change, an "extensive margin" relating to the fraction of firms adjusting, and a "selection effect" relating to *which* firms adjust. Our decomposition vindicates the claim of GL07, which was challenged by Caballero and Engel (2007), that the selection effect is crucial for the behavior of the FMC model. A fourth contribution of this paper is to calculate the impulse responses of prices to idiosyncratic as well as aggregate shocks. The paper also implements a recent algorithm for computing heterogeneous agent economies which is well-suited to modeling state-dependent pricing but has not yet been applied in this context.

1.1. Relation to previous literature

Most previous work on state-dependent pricing has obtained solutions by strongly limiting the analysis, either focusing on partial equilibrium (e.g. Caballero and Engel, 1993, 2007; Klenow and Kryvtsov, 2008), or assuming firms face aggregate shocks only (e.g. Dotsey *et al.*, 1999), or making strong assumptions about the distribution of idiosyncratic shocks (e.g. Caplin and Spulber, 1987; Gertler and Leahy, 2005). But Klenow and Kryvtsov (2008) argue convincingly that firms are frequently hit by large idiosyncratic shocks. And while heterogeneity may average out in many macroeconomic contexts, this is not true in the debate over nominal rigidities, because firm-level shocks could greatly alter firms' incentives to adjust prices. GL07 were the first to confront these issues head on, by studying a menu cost model in general equilibrium with idiosyncratic productivity shocks. They obtained a striking near-neutrality result, but their model's fit to price data is questionable, as our Figure 1 shows. A histogram based on retail microdata shows a wide range of price adjustments, whereas their FMC model generates just two sharp spikes of price increases and decreases occurring near the (S,s) bounds.

Other micro facts have been addressed in more recent papers on state-dependent pricing. Eichenbaum, Jaimovich, and Rebelo (2008) and Kehoe and Midrigan (2010) modeled "temporary" price changes (sales), assuming that these adjustments are cheaper than other price changes. However, they ultimately conclude that the possibility of sales has little relevance for monetary transmission, which depends instead on the frequency of regular non-sale price changes. Guimaraes and Sheedy's (2010) model of sales as stochastic price discrimination has the same implication. Thus, since the model developed in this paper has no natural motive for sales, it will be compared to a dataset of "regular" price changes from which apparent sales have been removed. In another branch of the literature, Boivin, Giannoni, and Mihov (2009) and Mackowiak, Moench, and Wiederholt (2009) calculate that prices respond much more quickly to idiosyncratic than to aggregate shocks. However, the present paper performs a Monte Carlo exercise which shows that this finding should be treated with caution. Remarkably, even when the true response to an idiosyncratic shock fades in and out gradually, the estimation routine of Mackowiak *et al.* can erroneously conclude that idiosyncratic shocks have an immediate, permanent impact on prices.

While matching pricing data makes it essential to allow for firm-specific shocks, this complicates computation, because the distribution of prices and productivities across firms becomes a relevant state variable. This paper shows how to compute a dynamic general equilibrium with state-dependent pricing via the two-step algorithm of Reiter (2009), which calculates steady state equilibrium using backwards induction on a grid, and then linearizes the equations at every grid point to calculate the dynamics. This avoids some complications (and simplifying assumptions) required by other applicable methods. In contrast to GL07, there is no need to assume that aggregate output remains constant after a money shock. In contrast to Dotsey, King, and Wolman (2008), it more fully exploits the recursive structure of the model, tracking the price distribution without needing to know who adjusted when. In contrast to the method of Krusell and Smith (1998), used by Midrigan (2008), there is no need to find an adequate summary statistic for the distribution. In contrast to Den Haan (1997), there is no need to impose a specific distributional form. The nonlinear, nonparametric treatment of firm-level heterogeneity in Reiter's algorithm makes it straightforward to calculate the time path of cross-sectional statistics, like our inflation decomposition; the linearization of aggregate dynamics makes it just as easy to analyze a variety of monetary policy rules or shock processes as it would be in a standard, low-dimensional DSGE model.

Several other closely related papers have also remarked that an FMC model implies a counterfactual distribution of price adjustments, in which small changes never occur. They proposed some more complex pricing models to fix this problem, including sectoral heterogeneity in menu costs (Klenow and Kryvtsov, 2008), multiple products on the same "menu" combined with leptokurtic technology shocks (Midrigan, 2008), or a mix of flexible- and sticky-price firms plus a mix of two distributions of productivity shocks (Dotsey *et al.*, 2008). This paper proposes a simpler approach: we just assume the probability of price adjustment increases with the value of adjustment, and treat the hazard function as a primitive of the model. A family of hazard functions with just three parameters suffices to match the distribution of price changes at least as well as the aforementioned papers do. Our setup can be interpreted as a stochastic menu cost model, like Dotsey *et al.* (1999) or Caballero and Engel (1999); under this interpretation the hazard function corresponds to the *c.d.f.* of the menu cost. Alternatively, our setup can be seen as a model of near-rational behavior, like Akerlof and Yellen (1985), in which firms sometimes make mistakes if they are not very costly; in this case the hazard function corresponds to the value distribution of errors.[2] Under either interpretation, the key point is that the adjustment hazard increases smoothly with the value of adjusting, in contrast with the discontinuous jump in probability implied by the FMC model. An appropriate calibration of the smoothness of the hazard function yields a smooth histogram of price changes consistent with microdata; this smoothness is the same property that eliminates the strong selection effect found by GL07. Thus, none of the complications Dotsey *et al.* and Midrigan tack on to the FMC framework are crucial for their most important finding: a state-dependent pricing model consistent with observed price changes implies nontrivial real effects of monetary shocks, similar to those found under the Calvo framework.

2. Model

This discrete-time model embeds state-dependent pricing by firms in an otherwise-standard New Keynesian general equilibrium framework based on GL07. Besides the firms, there is a representative household and a monetary authority that either implements a Taylor rule or follows an exogenous growth process for nominal money balances.

The aggregate state of the economy at time t, which will be identified in Section 2.3., is called Ω_t. Whenever aggregate variables are subscripted by t, this is an abbreviation indicating dependence, in equilibrium, on aggregate conditions Ω_t. For example, consumption is denoted by $C_t \equiv C(\Omega_t)$.

[2]The two interpretations imply slightly different Bellman equations: in the first case, but not in the second, a flow of menu costs is subtracted out of the firm's flow of profits.

2.1. Household

The household's period utility function is $\frac{1}{1-\gamma}C_t^{1-\gamma} - \chi N_t + \nu \log(M_t/P_t)$, where C_t is consumption, N_t is labor supply, and M_t/P_t is real money balances. Utility is discounted by factor β per period. Consumption is a CES aggregate of differentiated products C_{it}, with elasticity of substitution ϵ:

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

The household's nominal period budget constraint is

$$\int_0^1 P_{it}C_{it}di + M_t + R_t^{-1}B_t = W_tN_t + M_{t-1} + T_t + B_{t-1} + U_t, \tag{2}$$

where $\int_0^1 P_{it}C_{it}di$ is total nominal consumption. B_t is nominal bond holdings, with interest rate $R_t - 1$; T_t is a lump sum transfer from the central bank, and U_t is a dividend payment from the firms.

Households choose $\{C_{it}, N_t, B_t, M_t\}_{t=0}^{\infty}$ to maximize expected discounted utility, subject to the budget constraint (2). Optimal consumption across the differentiated goods implies

$$C_{it} = (P_t/P_{it})^{\epsilon} C_t, \tag{3}$$

so nominal spending can be written as $P_tC_t = \int_0^1 P_{it}C_{it}di$ under the price index

$$P_t \equiv \left\{ \int_0^1 P_{it}{}^{1-\epsilon} di \right\}^{\frac{1}{1-\epsilon}}. \tag{4}$$

Defining inflation as $\Pi_{t+1} \equiv P_{t+1}/P_t$, the first-order conditions for labor supply, consumption, and money use can be written as:

$$\chi = C_t^{-\gamma} W_t/P_t, \tag{5}$$

$$R_t^{-1} = \beta E_t \left(\frac{C_{t+1}^{-\gamma}}{\Pi_{t+1} C_t^{-\gamma}} \right), \tag{6}$$

$$1 - \frac{v'(m_t)}{C_t^{-\gamma}} = \beta E_t \left(\frac{C_{t+1}^{-\gamma}}{\Pi_{t+1} C_t^{-\gamma}} \right). \tag{7}$$

2.2. Monopolistic firms

Each firm i produces output Y_{it} under a constant returns technology $Y_{it} = A_{it}N_{it}$, where A_{it} is an idiosyncratic productivity process, AR(1) in logs:

$$\log A_{it} = \rho \log A_{it-1} + \varepsilon_{it}^a, \tag{8}$$

and labor N_{it} is the only input. Firm i is a monopolistic competitor that sets a price P_{it}, facing the demand curve $Y_{it} = C_t P_t^{\epsilon} P_{it}^{-\epsilon}$, and must fulfill all demand at its chosen price. It hires in a competitive labor markets at wage rate W_t, generating profits

$$U_{it} = P_{it}Y_{it} - W_t N_{it} = \left(P_{it} - \frac{W_t}{A_{it}} \right) C_t P_t^{\epsilon} P_{it}^{-\epsilon} \equiv U(P_{it}, A_{it}, \Omega_t) \tag{9}$$

per period. Firms are owned by the household, so they discount nominal income between times t and $t+1$ at the rate $\beta \frac{P(\Omega_t)u'(C(\Omega_{t+1}))}{P(\Omega_{t+1})u'(C(\Omega_t))}$, consistent with the household's marginal rate of substitution.

Let $V(P_{it}, A_{it}, \Omega_t)$ denote the nominal value of a firm at time t that produces with productivity A_{it} and sells at nominal price P_{it}. Prices are sticky, so P_{it} may or may not be optimal. However, we assume that whenever a firm adjusts its price, it chooses the optimal price conditional on its current productivity, keeping in mind that it will sometimes be unable to adjust in the future. Hence, the value function of an adjusting firm, after netting out any costs that may be required to make the adjustment, is $V^*(A_{it}, \Omega_t) \equiv \max_P V(P, A_{it}, \Omega_t)$. For clarity, it helps to distinguish the firm's beginning-of-period price, $\widetilde{P}_{it} \equiv P_{it-1}$, from the end-of-period price at which it sells at time t, P_{it}, which may or may not be the same. The distributions of prices and productivities across firms at the beginning and end of t will be denoted $\widetilde{\Phi}_t(\widetilde{P}, A)$ and $\Phi_t(P, A)$, respectively.

The gain from adjusting at the beginning of t is:

$$D(\widetilde{P}_{it}, A_{it}, \Omega_t) \equiv \max_P V(P, A_{it}, \Omega_t) - V(\widetilde{P}_{it}, A_{it}, \Omega_t). \tag{10}$$

The main assumption of our framework is that the probability of price adjustment increases with the gain from adjustment. The weakly increasing function λ that governs this probability is taken as a primitive of the model. Invariance of this function requires that its argument, the gain from adjustment, be written in appropriate units. As was mentioned in the introduction, this setup can be interpreted as a stochastic menu cost model, or as a model of near-rational price decisions. In the case of stochastic menu costs, the labor effort of changing price tags or rewriting the menu is likely to be a large component of the cost; in the near-rational case, the adjustment probability should be related to the labor effort involved in obtaining new information or recomputing the optimal price. Therefore, under either interpretation, the most natural units for the argument of the λ function are units of labor time. Thus, the probability of adjustment will be defined as $\lambda\left(L\left(\widetilde{P}_{it}, A_{it}, \Omega_t\right)\right)$, where

$L\left(\widetilde{P}_{it}, A_{it}, \Omega_t\right) = \frac{D(\widetilde{P}_{it}, A_{it}, \Omega_t)}{W(\Omega_t)}$ expresses the gains from adjusting in time units by dividing by the wage. The functional form for λ will be specified in Sec. 2.2.1.

The value of selling at any given price equals current profits plus the expected value of future production, which may or may not occur at a new, adjusted price. Given the firm's idiosyncratic state variables (P, A) and the aggregate state Ω, and denoting next period's variables with primes, the Bellman equation under the near-rational interpretation of the model is

$$V(P, A, \Omega) = \left(P - \frac{W(\Omega)}{A}\right) C(\Omega) P(\Omega)^{\epsilon} P^{-\epsilon} + \tag{11}$$

$$\beta E \left\{ \frac{P(\Omega)C(\Omega')^{-\gamma}}{P(\Omega')C(\Omega)^{-\gamma}} \left[\left(1 - \lambda\left(\frac{D(P,A',\Omega')}{W(\Omega')}\right)\right) V(P, A', \Omega') \; + \; \lambda\left(\frac{D(P,A',\Omega')}{W(\Omega')}\right) \max_{P'} V\left(P', A', \Omega'\right) \right] \middle| A, \Omega \right\}.$$

Here the expectation refers to the distribution of A' and Ω' conditional on A and Ω. Note that on the left-hand side of the Bellman equation, and in the term that represents current profits, P refers to a given firm i's price P_{it} at the end of t, when transactions occur. In the expectation on the right, P represents the price $\widetilde{P}_{i,t+1}$ at the beginning of $t + 1$, which may (probability λ) or may not $(1 - \lambda)$ be adjusted prior to time $t + 1$ transactions to a new value P'.

The right-hand side of the Bellman equation can be simplified by using the notation from (9), and the rearrangement $(1 - \lambda) V + \lambda \max V = V + \lambda(\max V - V)$:

$$V(P, A, \Omega) \;\; = \;\; U(P, A, \Omega) + \beta E \left\{ \frac{P(\Omega)C(\Omega')^{-\gamma}}{P(\Omega')C(\Omega)^{-\gamma}} \left[V(P, A', \Omega') + G(P, A', \Omega')\right] \middle| A, \Omega \right\}, \tag{12}$$

where

$$G(P, A', \Omega') \equiv \lambda\left(\frac{D(P, A', \Omega')}{W(\Omega')}\right) D(P, A', \Omega'). \tag{13}$$

The terms inside the expectation in the Bellman equation represent the value V of continuing without adjustment, plus the flow of expected gains G due to adjustment. Since the firm sets the optimal price whenever it adjusts, the price process associated with (12) is

$$P_{it} = \begin{cases} P^*(A_{it}, \Omega_t) \equiv \arg\max_P V(P, A_{it}, \Omega_t) & \text{with probability } \lambda\left(\frac{D(\widetilde{P}_{it}, A_{it}, \Omega_t)}{W(\Omega_t)}\right) \\ \widetilde{P}_{it} \equiv P_{i,t-1} & \text{with probability } 1 - \lambda\left(\frac{D(\widetilde{P}_{it}, A_{it}, \Omega_t)}{W(\Omega_t)}\right) \end{cases}. \tag{14}$$

Equation (14) is written with time subscripts for additional clarity.

2.2.1. *Alternative sticky price frameworks*

Our assumptions require the function λ to be weakly increasing and to lie between zero and one. The paper focuses primarily on the following functional form:

$$\lambda(L) \equiv \frac{\overline{\lambda}}{\overline{\lambda} + (1-\overline{\lambda})\left(\frac{\alpha}{L}\right)^{\xi}} \tag{15}$$

with α and ξ positive, and $\overline{\lambda} \in [0,1]$. This function equals $\overline{\lambda}$ when $L = \alpha$, and is concave for $\xi \leq 1$ and S-shaped for $\xi > 1$ (see the second panel of Fig. 1). The parameter ξ effectively controls the degree of state dependence. In the limit $\xi = 0$, (15) nests Calvo (1983), with $\lambda(L) = \overline{\lambda}$, making the adjustment hazard literally independent of the relevant state variable, which is L. At the opposite extreme, as $\xi \to \infty$, $\lambda(L)$ becomes the indicator $\mathbf{1}\{L \geq \alpha\}$, which equals 1 whenever $L \geq \alpha$ and is zero otherwise. This implies very strong state dependence, in the sense that the adjustment probability jumps from 0 to 1 when the state L passes the threshold level α. For all intermediate values $0 < \xi < \infty$, the hazard increases smoothly with the state L. In this sense, choosing ξ to match microdata means finding the degree of state dependence most consistent with firms' observed pricing behavior.

TABLE 1 ABOUT HERE

The combination of Bellman equation (12) with (13) is based on a near-rational interpretation of our setup; for $0 < \xi < \infty$ this version of the model will be called "SSDP", for "smoothly state-dependent pricing". However, (12) nests several other models too, by appropriate choice of the gains function G and the hazard function λ, as Table 1 shows. Subtracting a flow of menu costs $E(\kappa|\kappa < L) \equiv \int_0^L \kappa\lambda(d\kappa)$ out of the gains G converts the SSDP model into a stochastic menu cost (SMC) model. The FMC model sets the adjustment probability to a step function, subtracting a constant menu cost α out of G; it is the limit of the SMC model as $\xi \to \infty$. The Calvo model is derived both from SSDP and from SMC as $\xi \to 0$.[3] An alternative hazard function derived from Woodford (2008) is also considered.

[3]In the limit of SMC as $\xi \to 0$, the menu cost is zero with probability $\overline{\lambda}$ and infinite with probability $1-\overline{\lambda}$, which is when firms do not adjust. The flow of menu costs paid is therefore zero.

2.3. *Monetary policy and aggregate consistency*

Two specifications for monetary policy are compared: a money growth rule and a Taylor rule. In both cases the systematic component of monetary policy is perturbed by an AR(1) process z,

$$z_t = \phi_z z_{t-1} + \epsilon_t^z, \tag{16}$$

where $0 \leq \phi_z < 1$ and $\epsilon_t^z \sim i.i.d.N(0, \sigma_z^2)$. Under the money growth rule, which is analyzed first to build intuition and for comparison with previous studies, z affects money supply growth:

$$M_t/M_{t-1} \equiv \mu_t = \mu^* \exp(z_t). \tag{17}$$

Alternatively, under a Taylor interest rate rule, which is a better approximation to actual monetary policy, the nominal interest rate follows

$$\frac{R_t}{R^*} = \exp(-z_t)\left(\left(\frac{P_t/P_{t-1}}{\Pi^*}\right)^{\phi_\pi}\left(\frac{C_t}{C^*}\right)^{\phi_c}\right)^{1-\phi_R}\left(\frac{R_{t-1}}{R^*}\right)^{\phi_R}, \tag{18}$$

where $\phi_c \geq 0$, $\phi_\pi > 1$, and $0 < \phi_R < 1$, so that when inflation Π_t exceeds its target Π^* or consumption C_t exceeds its target C^*, R_t tends to rise above its target $R^* \equiv \Pi^*/\beta$. For comparability between the two monetary regimes, the inflation target is set to $\Pi^* \equiv \mu^*$, and the rules are specified so that in both cases, a positive z represents an expansive shock.

Seigniorage revenues are paid to the household as a lump sum transfer T_t, and the government budget is balanced each period, so that $M_t = M_{t-1} + T_t$. Bond market clearing is simply $B_t = 0$. When supply equals demand for each good i, total labor supply and demand satisfy

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

Equation (19) also defines a measure of price dispersion, $\Delta_t \equiv P_t^\epsilon \int_0^1 P_{it}^{-\epsilon} A_{it}^{-1} di$, weighted to allow for heterogeneous productivity. As in Yun (2005), an increase in Δ_t decreases the goods produced per unit of labor, effectively acting like a negative aggregate shock.

At this point, all equilibrium conditions have been spelled out, so an appropriate aggregate state variable Ω_t can be identified. At time t, the lagged distribution of transaction prices $\Phi_{t-1}(P, A)$ is predetermined. Knowing Φ_{t-1}, the lagged price level can be substituted out of the Taylor rule, using $P_{t-1} = \left[\iint P^{1-\epsilon}\Phi_{t-1}(dP, dA)\right]^{1/(1-\epsilon)}$. It can thus be seen that $\Omega \equiv (z_t, R_{t-1}, \Phi_{t-1})$ suffices to define the aggregate state. Given this Ω_t, equations (4), (5), (6), (8), (9), (10), (12), (13), (14), (16), (18), and (19) together give enough conditions to determine the distributions $\widetilde{\Phi}_t$ and Φ_t, the price level P_t, the functions $V_t \equiv V(P, A, \Omega_t)$, U_t, D_t, and G_t, and the variables R_t, C_t, N_t, W_t, and z_{t+1}. Thus they determine the next state, $\Omega_{t+1} \equiv (z_{t+1}, R_t, \Phi_t)$.

Under a money growth rule, the time t state can instead be defined as $\Omega_t \equiv (z_t, M_{t-1}, \Phi_{t-1})$. Substituting (7) for (6) and (17) for (18), knowing $\Omega_t \equiv (z_t, M_{t-1}, \Phi_{t-1})$ suffices to determine $\widetilde{\Phi}_t$, Φ_t, P_t, V_t, U_t, D_t, G_t, C_t, N_t, W_t, z_{t+1}, and M_t. Thus the next state, $\Omega_{t+1} \equiv (z_{t+1}, M_t, \Phi_t(P, A))$, can be calculated.

3. Computation

The fact that this model's state variable includes the distribution Φ, an infinite-dimensional object, makes computing equilibrium a challenge. The popularity of the Calvo model reflects its implication that general equilibrium can be solved up to a first-order approximation by keeping track of the average price only. Unfortunately, this result typically fails to hold if pricing is state-dependent; instead, computation requires tracking the whole distribution Φ.

Equilibrium will be computed following the two-step algorithm of Reiter (2009), which is intended for contexts, like this model, with relatively large idiosyncratic shocks and also relatively small aggregate shocks. In the first step, the aggregate steady state of the model is computed on a finite grid, using backwards induction.[4] Second, the stochastic aggregate dynamics are computed by linearization, grid point by grid point. In other words, the Bellman equation is treated as a large system of expectational difference equations, instead of as a functional equation.

3.1. Detrending

Calculating a steady state requires detrending to make the economy stationary. Here it suffices to scale all nominal variables by the aggregate price level, defining the real wage and money supply $w_t = W_t/P_t$ and $m_t \equiv M_t/P_t$, and the real prices at the beginning and end of t, $\widetilde{p}_{it} \equiv \widetilde{P}_{it}/P_t$

[4] Actually, Reiter's algorithm permits calculation of the aggregate steady state using a variety of finite-element methods; we choose backwards induction on a grid since it is a familiar and transparent procedure.

and $p_{it} \equiv P_{it}/P_t$. The beginning-of-t and end-of-t distributions will be written as $\widetilde{\Psi}_t(\widetilde{p}_{it}, A_{it})$ and $\Psi_t(p_{it}, A_{it})$, respectively. At the end of t, when goods are sold, the real price level is one by definition:

$$1 = \left\{ \iint p_{t-1}^{1-\epsilon} \Psi_t(dp, dA) \right\}^{1/(1-\epsilon)}. \tag{20}$$

For this detrending to make sense, the nominal price level P_t must be irrelevant for real quantities, which must instead be functions of a real state variable Ξ_t that is independent of nominal prices and the nominal money supply. A time subscript on any aggregate variable must now denote dependence on the *real* state, implying for example $w_t = w(\Xi_t) = W(\Omega_t)/P(\Omega_t)$ and $C_t = C(\Xi_t) = C(\Omega_t)$. While the price level will cancel out, inflation will still appear in the model, and must be determined by real variables, satisfying $\Pi_t = \Pi(\Xi_{t-1}, \Xi_t) = P(\Omega_t)/P(\Omega_{t-1})$.

A similar property applies to the value function and profits, which must be homogeneous of degree one in prices. Thus, define real profits u and real value v as follows:

$$u\,(p, A, \Xi) \;=\; u\,(P/P(\Omega), A, \Xi) \equiv P(\Omega)^{-1} U(P, A, \Omega), \tag{21}$$

$$v\,(p, A, \Xi) \;=\; v\,(P/P(\Omega), A, \Xi) \equiv P(\Omega)^{-1} V(P, A, \Omega). \tag{22}$$

To verify homogeneity, divide through the nominal Bellman equation (12) by $P(\Omega)$ to obtain

$$v(p, A, \Xi) = u(p, A, \Xi) + \beta E \left\{ \frac{C(\Xi')^{-\gamma}}{C(\Xi)^{-\gamma}} \left[v\left(\frac{p}{\Pi(\Xi, \Xi')}, A', \Xi' \right) + g\left(\frac{p}{\Pi(\Xi, \Xi')}, A', \Xi' \right) \right] \middle| A, \Xi \right\}, \tag{23}$$

using the definitions

$$g\,(\widetilde{p}, A, \Xi) \;\equiv\; \lambda \left(\frac{d\,(\widetilde{p}, A, \Xi)}{w(\Xi)} \right) d\,(\widetilde{p}, A, \Xi), \tag{24}$$

$$d\,(\widetilde{p}, A, \Xi) \;\equiv\; \max_p v(p, A, \Xi) - v(\widetilde{p}, A, \Xi), \tag{25}$$

which satisfy $g\,(\widetilde{p}, A, \Xi) = G(P(\Omega)\widetilde{p}, A, \Omega)/P(\Omega)$ and $d\,(\widetilde{p}, A, \Xi) = D(P(\Omega)\widetilde{p}, A, \Omega)/P(\Omega)$.[5] This detrending implies that when a firm's nominal price remains unadjusted at time t, its real price is deflated

[5] In deriving (23) from (12), initially a term of the form $\frac{C(\Omega')^{-\gamma}}{P(\Omega')C(\Omega)^{-\gamma}} V\,(P, A', \Omega')$ appears on the right-hand side; using (22) this reduces to $\frac{C(\Xi')^{-\gamma}}{C(\Xi)^{-\gamma}} v \left(\frac{p}{\Pi(\Xi', \Xi)}, A', \Xi' \right)$. Reducing the G term in the same way yields (23).

by factor $\Pi_t = P_t/P_{t-1}$. Therefore the real price process is

$$p_{it} = \begin{cases} p^*(A_{it}, \Xi_t) \equiv \arg\max_p v(p, A_{it}, \Xi_t) & \text{with probability } \lambda\left(\frac{d\left(\Pi_t^{-1}p_{i,t-1}, A_{it}, \Xi_t\right)}{w(\Xi_t)}\right) \\ \Pi_t^{-1}p_{i,t-1} & \text{with probability } 1-\lambda\left(\frac{d\left(\Pi_t^{-1}p_{i,t-1}, A_{it}, \Xi_t\right)}{w(\Xi_t)}\right) \end{cases} . \tag{26}$$

To see that these definitions of real quantities suffice to detrend the model, define the real state as $\Xi_t \equiv (z_t, R_{t-1}, \Psi_{t-1})$. Knowing Ξ_t, in the case of a Taylor rule, equations (5), (6), (8), (16), (18), (19), (20), (21), (23), (24), (25), and (26), with substitutions of real for nominal variables where necessary, suffice to determine the distributions $\widetilde{\Psi}_t$ and Ψ_t, inflation Π_t, the functions u_t, v_t, g_t, and d_t, and the variables C_t, w_t, N_t, R_t, and z_{t+1}. For a money growth rule, the real state can be defined as $\Xi_t \equiv (z_t, m_{t-1}, \Psi_{t-1})$, and equation (18) is replaced by (7) and by

$$m_t = \mu^* \exp(z_t) m_{t-1}/\Pi_t, \tag{27}$$

which together determine R_t and m_t. Thus next period's state Ξ_{t+1} can be calculated if Ξ_t is known.

3.2. Discretization

The price process (26) takes a continuum of possible values, but to solve this model numerically the idiosyncratic states must be restricted to a finite-dimensional support. Hence, the continuous model will be approximated on a two-dimensional grid $\Gamma \equiv \Gamma^p \times \Gamma^a$, where $\Gamma^p \equiv \{p^1, p^2, ...p^{\#^p}\}$ and $\Gamma^a \equiv \{a^1, a^2, ...a^{\#^a}\}$ are logarithmically-spaced grids of possible values of of p_{it} and A_{it}. Thus the time-varying distributions will be treated as matrices $\widetilde{\mathbf{\Psi}}_t$ and $\mathbf{\Psi}_t$ of size $\#^p \times \#^a$, in which the row j, column k elements, called $\widetilde{\Psi}_t^{jk}$ and Ψ_t^{jk}, represent the fraction of firms in state (p^j, a^k) before and after price adjustments in period t, respectively. From here on, bold face is used to identify matrices, and superscripts are used to identify notation related to grids.

Similarly, the value function is written as a $\#^p \times \#^a$ matrix $\mathbf{V}_t$ of values $v_t^{jk} \equiv v(p^j, a^k, \Xi_t)$ associated with the prices and productivities $\left(p^j, a^k\right) \in \Gamma$. The time subscript indicates the fact that the value function shifts due to changes in the aggregate state Ξ_t. When necessary, the value is evaluated using splines at points $p \notin \Gamma^p$ off the price grid. In particular, the policy function

$$p_t^*(A) \equiv \arg\max_p v(p, A, \Xi_t) \tag{28}$$

is defined without requiring that it be chosen from the grid Γ^p, because our solution method will require policies to vary continuously with their arguments. The policies at the productivity grid points $a^k \in \Gamma^a$ are written as a row vector $\mathbf{p}_t^* \equiv \left\{p_t^{*1}...p_t^{*\#^a}\right\} \equiv \left\{p_t^*(a^1)...p_t^*(a^{\#^a})\right\}$. Various other equilibrium functions are also treated as $\#^p \times \#^a$ matrices. The adjustment values $\mathbf{D}_t$, the probabilities $\mathbf{\Lambda}_t$, and the expected gains $\mathbf{G}_t$ have (j,k) elements given by[6]

$$d_t^{jk} \equiv \max_p v_t(p, a^k) - v_t^{jk}, \tag{29}$$

$$\lambda_t^{jk} \equiv \lambda\left(d_t^{jk}/w_t\right), \tag{30}$$

$$g_t^{jk} \equiv \lambda_t^{jk} d_t^{jk}. \tag{31}$$

Given this discrete representation, the distributional dynamics can be written in a more explicit way. First, to keep productivity A on the grid Γ^a, it is assumed to follow a Markov chain defined by a matrix $\mathbf{S}$ of size $\#^a \times \#^a$. The row m, column k element of $\mathbf{S}$ represents the probability

$$S^{mk} = prob(A_{it} = a^m | A_{i,t-1} = a^k). \tag{32}$$

Also, beginning-of-t real prices must be adjusted for inflation. Ignoring grids, the time $t-1$ price $p_{i,t-1}$ would be deflated to $\widetilde{p}_{it} \equiv p_{i,t-1}/\Pi_t$ at the beginning of t. Prices are forced to remain on the grid by a $\#^p \times \#^p$ Markov matrix $\mathbf{R}_t$ in which the row m, column l element represents

$$R_t^{ml} \equiv prob(\widetilde{p}_{it} = p^m | p_{i,t-1} = p^l, \Pi_t = \Pi(\Xi_t, \Xi_{t-1})). \tag{33}$$

When the deflated price $p_{i,t-1}/\Pi_t$ falls between two grid points, $\mathbf{R}_t$ rounds it up or down stochastically without changing its mean. Also, if $p_{i,t-1}/\Pi_t$ drifts up or down past the largest or smallest grid points, then $\mathbf{R}_t$ rounds it down or up to keep prices on the grid. Thus the transition probabilities are

$$R_t^{ml} = \begin{cases} 1 & \text{if } p^l/\Pi_t \leq p^1 = p^m \\ \frac{p^l/\Pi_t - p^{m-1}}{p^m - p^{m-1}} & \text{if } p^1 < p^m = \min\{p \in \Gamma^p : p \geq p^l/\Pi_t\} < p^{\#^p} \\ \frac{p^{m+1} - p^l/\Pi_t}{p^{m+1} - p^n} & \text{if } p^1 \leq p^m = \max\{p \in \Gamma^p : p < p^l/\Pi_t\} < p^{\#^p} \\ 1 & \text{if } p^l/\Pi_t > p^{\#^p} = p^m \\ 0 & \text{otherwise} \end{cases}. \tag{34}$$

[6] The max in (29), like the arg max in (28), ignores the grid Γ^p so that d_t^{jk} varies continuously in response to any change in the value function.

Combining the adjustments of prices and productivities, the beginning-of-t distribution $\widetilde{\mathbf{\Psi}}_t$ can be calculated from the lagged distribution $\mathbf{\Psi}_{t-1}$ as follows:

$$\widetilde{\mathbf{\Psi}}_t = \mathbf{R}_t * \mathbf{\Psi}_{t-1} * \mathbf{S}', \tag{35}$$

where the operator $*$ represents matrix multiplication. Two facts explain the simplicity of this equation. First, the exogenous shocks to A_{it} are independent of the inflation adjustment linking $\widetilde{p}_{it}$ with $p_{i,t-1}$. Second, productivity is arranged from left to right in the matrix $\mathbf{\Psi}_{t-1}$, so productivity transitions are represented by right multiplication, while prices are arranged vertically, so price transitions are represented by left multiplication.

Next, a firm with beginning-of-t state $(\widetilde{p}_{it}, A_{it}) = (p^j, a^k) \in \Gamma$ will adjust its price to $p_{it} = p_t^{*k}$ with probability λ_t^{jk}, and otherwise leave it unchanged. If adjustment occurs, prices are kept on the grid by rounding p_t^{*k} up or down stochastically to the nearest grid points, without changing the mean. For concise notation, let $\mathbf{E}_{pp}$ and $\mathbf{E}_{pa}$ be matrices of ones of size $\#^p \times \#^p$ and $\#^p \times \#^a$, respectively. Let Γ^p be wide enough so that $p^1 < p_t^{*k} < p^{\#^p}$ for all $k \in \{1, 2, ...\#^a\}$. For each k, define $l_t(k)$ so that $p^{l_t(k)} = \min\{p \in \Gamma^p : p \geq p_t^{*k}\}$. Then the following $\#^p \times \#^a$ matrix governs the stochastic rounding:

$$\mathbf{P}_t \equiv \begin{cases} \frac{p^{l_t(k)} - p_t^{*k}}{p^{l_t(k)} - p^{l_t(k)-1}} & \text{in column } k, \text{ row } l_t(k) - 1 \\ \frac{p_t^{*k} - p^{l_t(k)-1}}{p^{l_t(k)} - p^{l_t(k)-1}} & \text{in column } k, \text{ row } l_t(k) \\ 0 & \text{elsewhere} \end{cases}. \tag{36}$$

The end-of-t distribution $\mathbf{\Psi}_t$ can then be calculated from $\widetilde{\mathbf{\Psi}}_t$ as follows:

$$\mathbf{\Psi}_t = (\mathbf{E}_{pa} - \mathbf{\Lambda}_t) \;.*\; \widetilde{\mathbf{\Psi}}_t + \mathbf{P}_t \;.*\; (\mathbf{E}_{pp} * (\mathbf{\Lambda}_t \;.*\; \widetilde{\mathbf{\Psi}}_t)). \tag{37}$$

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

The same transition matrices show up when the Bellman equation is written in matrix form. Let $\mathbf{U}_t$ be the $\#^p \times \#^a$ matrix of current profits, with elements

$$u_t^{jk} = u(p^j, a^k, \Xi_t) = \left(p^j - \frac{w_t}{a^k}\right) C_t p_j^{-\epsilon} \tag{38}$$

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

$$\mathbf{V}_t = \mathbf{U}_t + \beta E_t \left\{ \frac{C_{t+1}^{-\gamma}}{C_t^{-\gamma}} \mathbf{R}'_{t+1} * (\mathbf{V}_{t+1} + \mathbf{G}_{t+1}) * \mathbf{S} \right\}, \tag{39}$$

where $\mathbf{G}_{t+1} = \mathbf{\Lambda}_{t+1} \; .* \, \mathbf{D}_{t+1}$ was defined by (31).

Several comments may help clarify this Bellman equation. Note that the expectation E_t refers only to the effects of the time $t+1$ aggregate shock z_{t+1}, because multiplying by $\mathbf{R}'_{t+1}$ and $\mathbf{S}$ fully describes the expectation over the idiosyncratic state $(p^j, a^k) \in \Gamma$. $\mathbf{S}$ has no time subscript, since the Markov productivity process is not subject to aggregate shocks, whereas the inflation adjustment represented by $\mathbf{R}'_{t+1}$ varies with the policy shock. Also, while the distributional dynamics iterate forward in time, with transitions governed by $\mathbf{R}$ and $\mathbf{S}'$, the Bellman equation iterates backwards, so its transitions are described by $\mathbf{R}'$ and $\mathbf{S}$.

3.3. Computation: steady state

In an aggregate steady state, monetary policy shocks z are zero, and transaction prices converge to an ergodic distribution $\mathbf{\Psi}$, so the aggregate state of the economy is constant: $\Xi_t = (z_t, R_{t-1}, \mathbf{\Psi}_{t-1}) = (0, R, \mathbf{\Psi}) \equiv \Xi$ under the Taylor rule, or $\Xi_t = (z_t, m_{t-1}, \mathbf{\Psi}_{t-1}) = (0, m, \mathbf{\Psi}) \equiv \Xi$ under a money growth rule. The steady state of any aggregate equilibrium object is indicated by dropping the subscript t.

The steady state calculation nests the firm's backwards induction problem inside a loop that determines the steady-state real wage w. Note first that inflation and the interest rate must satisfy $\Pi = \mu^* = \beta R$; hence the matrix $\mathbf{R}$ is known. Then, guessing w, the first-order condition (5) determines C, making it possible to calculate all elements u^{jk} of $\mathbf{U}$ from (38). Thus backwards induction on the grid Γ can be used to solve the Bellman equation

$$\mathbf{V} = \mathbf{U} + \beta \mathbf{R}' * (\mathbf{V} + \mathbf{G}) * \mathbf{S}. \tag{40}$$

Solving (40) involves finding the matrices $\mathbf{V}$, $\mathbf{D}$, $\mathbf{\Lambda}$, and $\mathbf{G}$, so the matrix $\mathbf{P}$ can also be calculated from (36). Thus (35) and (37) can be used to find the distributions $\widetilde{\mathbf{\Psi}}$ and $\mathbf{\Psi}$, and finally (4) can be used to check the guessed value of w. In the discretized notation, equation (4) becomes

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

If (41) holds at the ergodic distribution $\mathbf{\Psi}_t = \mathbf{\Psi}$, then a steady-state equilibrium has been found.

3.4. Computation: dynamics

While the Bellman equation (39) and distributional dynamics (37) can be interpreted as functional equations, under the discrete approximation of Sec. 3.2 they can alternatively be seen as two long lists of expectational difference equations that describe the values and probabilities at all grid points. Thus Reiter (2009) proposes linearizing these equations around their steady state, calculated in Sec. 3.3. To do so, it is first convenient to reduce the number of variables by eliminating simple intratemporal relationships. Assuming a money growth rule, the model can be described by the following vector of endogenous variables:

$$\overrightarrow{X}_t \equiv \left(vec\left(\mathbf{V}_t\right)', \quad C_t, \quad \Pi_t, \quad vec\left(\mathbf{\Psi}_{t-1}\right)', \quad m_{t-1} \right)' \tag{42}$$

Vector $\overrightarrow{X}_t$, together with the shock process z_t, consists of $2\#^p\#^a + 4$ variables determined by the following system of $2\#^p\#^a + 4$ equations: (39), (7), (41), (37), (27), and (16). Under a Taylor rule, m_{t-1} is replaced by R_{t-1}, and (7) and (27) are replaced by (6) and (18). Thus the expectational difference equations governing dynamic equilibrium constitute a first-order system of the form

$$E_t \mathcal{F}\left(\overrightarrow{X}_{t+1}, \overrightarrow{X}_t, z_{t+1}, z_t\right) = 0, \tag{43}$$

where E_t is an expectation conditional on z_t and all previous shocks.[7] Next, system $\mathcal{F}$ can be linearized numerically to construct the Jacobian matrices $\mathcal{A} \equiv D_{\overrightarrow{X}_{t+1}}\mathcal{F}$, $\mathcal{B} \equiv D_{\overrightarrow{X}_t}\mathcal{F}$, $\mathcal{C} \equiv D_{z_{t+1}}\mathcal{F}$, and $\mathcal{D} \equiv D_{z_t}\mathcal{F}$. This yields the following first-order linear expectational difference equation system:

$$E_t \mathcal{A}\Delta\overrightarrow{X}_{t+1} + \mathcal{B}\Delta\overrightarrow{X}_t + E_t\mathcal{C}z_{t+1} + \mathcal{D}z_t = 0, \tag{44}$$

where Δ represents a deviation from steady state. This system has the form considered by Klein (2000), so it will be solved using his QZ decomposition method, though other linear rational expectations solvers would be applicable as well.

The virtue of Reiter's method is that it combines linearity and nonlinearity in a way appropriate for the context of price adjustment, where idiosyncratic shocks are larger and more economically important

[7]Note that given $(\overrightarrow{X}_{t+1}, \overrightarrow{X}_t, z_{t+1}, z_t)$, all other variables appearing in (39), (7), (41), (37), (27), and (16) can be substituted out using intratemporal equations. Given Π_t and Π_{t+1}, $\mathbf{R}_t$ and $\mathbf{R}_{t+1}$ are known; thus $\widetilde{\mathbf{\Psi}}_t = \mathbf{R}_t * \mathbf{\Psi}_{t-1} * \mathbf{S}'$ can be calculated too. The wage can be calculated from (5), so $\mathbf{U}_t$ can be constructed. Finally, given $\mathbf{V}_t$ and $\mathbf{V}_{t+1}$ it is possible to construct $\mathbf{P}_t$, $\mathbf{D}_t$, and $\mathbf{D}_{t+1}$, and thus $\mathbf{\Lambda}_t$ and $\mathbf{G}_{t+1}$. Therefore the arguments of $\mathcal{F}$ are indeed sufficient to evaluate the system (43).

for individual firms than aggregate shocks. To deal with large idiosyncratic shocks, it treats functions of idiosyncratic states nonlinearly (calculating them on a grid). But in linearizing each equation at each grid point, it recognizes that aggregate changes (monetary shocks z, or shifts of the distribution $\mathbf{\Psi}$) are unlikely to affect individual value functions in a strongly nonlinear way. On the other hand, it makes no assumption of approximate aggregation like that of Krusell and Smith (1998).

4. Results

4.1. Parameterization

Our calibration seeks price adjustment and productivity processes consistent with microdata on price changes, like those in Klenow and Kryvtsov (2008), Nakamura and Steinsson (2008), and Midrigan (2008). Since utility parameters are not the main focus, these are set to the values used by GL07. The discount factor is set to $\beta = 1.04^{-1}$ per year; the coefficient of relative risk aversion of consumption is set at $\gamma = 2$. The coefficients on labor disutility and the utility of money are $\chi = 6$ and $\nu = 1$, respectively, and the elasticity of substitution in the consumption aggregator is $\epsilon = 7$.

The main price data that will serve as an empirical benchmark are the AC Nielsen data reported by Midrigan (2008). Therefore, the model will be simulated at monthly frequency, with a zero steady state money growth rate, consistent with the zero average price change found in the monthly AC Nielsen dataset. Midrigan reports the data after removing price changes attributable to temporary "sales", so our simulation results should be interpreted as a model of "regular" price changes unrelated to sales. Conditional on these specification choices, the parameters of the adjustment process ($\overline{\lambda}$, α, and ξ) and of the productivity process (ρ and σ_ε^2) are chosen to minimize a distance criterion between the data and the model's steady state.[8] The criterion sums two terms, scaled for comparability: one relating to the frequency of adjustment, and the other relating to the histogram of nonzero price adjustments.

TABLE 2 ABOUT HERE

Table 2 summarizes the steady-state behavior of the model under the estimated parameters, to-

[8]The productivity process (8) is approximated on the grid Γ^a using Tauchen's method; we thank Elmar Mertens for making his software available. The productivity grid has 25 points, and the price grid Γ^p has 31 points. Both grids are logarithmically spaced; steps in Γ^p represent 4% changes. Results are not sensitive to the use of this coarse grid, since the average absolute price adjustment is much larger (around 10%).

gether with evidence from four empirical studies. The baseline specification, in which $\overline{\lambda}$, α, and ξ are all estimated, is labelled SSDP. The table also reports Calvo ($\overline{\lambda}$ estimated, $\xi \equiv 0$, and α undefined) and FMC specifications (α estimated, $\xi \equiv \infty$, and $\overline{\lambda}$ undefined), as well as a version based on Woodford's (2008) adjustment function and an SMC specification. All versions of the model match the target adjustment frequency of 10% per month.[9] But the extreme cases of the model (Calvo and FMC) are much less successful in fitting the size distribution of price adjustments than are the smooth intermediate cases; the Calvo model understates the average size and standard deviation of price adjustments, whereas the FMC model overstates both.

The trouble with the FMC model, as Fig. 1 shows, is that it only produces price changes lying just outside the (S,s) bands, whereas the adjustments observed in the data are very diverse.[10] Thus the FMC model that best fits the data produces adjustments that are too large on average; no adjustments in the model are less than 5%, whereas one quarter of all adjustments are below the 5% threshold in the AC Nielsen data. The Calvo model errs in the opposite direction, with too many small price adjustments, though its fit statistics are better than those of the FMC model. In contrast, all three specifications with a smoothly increasing adjustment hazard (SSDP, SMC, and Woodford) match the data well, since they permit large and small price adjustments to coexist. In fact, there is so little difference between these models that only SSDP will be discussed from here on.[11]

FIGURE 1 ABOUT HERE

4.2. Effects of monetary policy shocks

Since all specifications are calibrated to the same observed adjustment frequency, the fact that only large, valuable price changes occur in the FMC model, whereas some changes in the SSDP and Calvo frameworks are trivial, has important implications for monetary transmission. Fig. 2 compares responses to several types of monetary shocks across these three adjustment specifications. All simula-

[9] We fit the model to Nakamura and Steinsson's (2008) measure of the *median* frequency of price adjustments which is lower, but presumably more robust, than measures based on means.

[10] Klenow and Kryvtsov (2008) document that large and small price changes coexist even within narrow product categories, and that the FMC model performs poorly even when menu costs are allowed to differ across sectors.

[11] Our companion paper, Costain and Nakov (2008), shows that the SSDP model performs somewhat better than Woodford's specification at high (*e.g.* 70% annually) inflation rates. But at low inflation rates, the responses to monetary shocks (available from the authors) are indistinguishable across the SSDP, SMC, and Woodford specifications.

tions assume the same utility parameters, and zero baseline inflation, and are calculated starting from the steady state distribution associated with the corresponding specification. The first two rows show impulse responses to one percentage point money growth shocks, comparing the *i.i.d.* case with that of monthly autocorrelation $\phi_z = 0.8$. The third row shows the responses to a 25 basis point interest rate shock under a Taylor rule.

FIGURE 2 ABOUT HERE

In all three models, an increase in money growth stimulates consumption. The fact that some prices fail to adjust immediately means expected inflation rises, decreasing the *ex ante* real interest rate; it also means households' real money balances increase; both these effects raise consumption demand. However, as GL07 emphasized, the average price level adjusts rapidly in the FMC specification (lines with circles), with a large, short-lived spike in inflation. This makes changes in real variables small and transitory, approaching the monetary neutrality associated with full price flexibility. At the opposite extreme, prices adjust gradually in the Calvo specification (lines with squares), leading to a large, persistent increase in output. The response of the SSDP model (lines with dots) mostly lies between the other two, but is generally much closer to that of the Calvo model.

Comparing the first and second rows of Fig. 2 shows that while the shape of the inflation and output responses differs substantially across models, it is qualitatively similar under *iid* and autocorrelated money growth processes. Inflation spikes immediately in the FMC model with autocorrelated money, because the average price increase rises by much more than 1%, as firms anticipate that money growth will remain positive for some time. The rise in inflation is smaller but more persistent in the SSDP and Calvo cases. Note that the persistence of inflation does not differ noticeably depending on the autocorrelation of money growth, but instead appears to be determined primarily by the degree of state dependence. Thus the big difference between the impulse responses in the first and second rows is one of size, not of shape: the overall response is larger with autocorrelated money.

The third row of Fig. 2 shows responses under a Taylor rule, assuming that the underlying shock z is *i.i.d.*, and that the rule has inflation and output coefficients $\phi_\pi = 2$ and $\phi_c = 0.5$, and smoothing coefficient $\phi_R = 0.9$. While money growth shocks are small, permanent, changes to the *level* of the nominal money supply, Taylor rule shocks involve large but mostly transitory changes in the level of nominal money. Nonetheless, the two types of monetary policy shocks have similar real effects, and moreover, the finding that a micro-calibrated model of state-dependent pricing implies substantial monetary nonneutrality is strengthened in several ways by considering a Taylor rule. First, under the Taylor rule, the SSDP and Calvo impulse responses of inflation and consumption are even closer together than they were under the money growth rule. In fact, for consumption, both SSDP *and* FMC imply virtually the same effect on impact as that occurring in the Calvo model, though the effect is less persistent in the FMC case.

Recall, though, that the Taylor rule responses in Fig. 2 suppose an initial drop in the nominal interest rate of 25 basis points. Since the interest rate is endogenous, the required underlying shock ϵ^z varies across models, and it is particularly large in the FMC case. Therefore, it is useful to consider additional ways of comparing the real effects of monetary shocks across models. Thus, instead of comparing shocks with the same initial interest rate effect, Table 3 compares monetary policies with the same implied inflation variability. As in Sec. VI of GL07, the calculation asks the following question: if monetary policy shocks were the only source of observed US inflation volatility, how much output variation would they cause? Under the SSDP specification, money growth shocks alone would explain 65% of observed output fluctuations; the figure rises to 116% under the Calvo specification, and falls to 15% in the FMC case.[12] Assuming a Taylor rule, the differences across models are even stronger, and the monetary nonneutrality associated with the SSDP and Calvo specifications is even larger. Taylor rule shocks alone would explain 110% of US output fluctuation under the SSDP specification, rising to 306% in the Calvo case. The table also reports a "Phillips curve" coefficient, calculated by regressing log output on inflation, instrumented by the exogenous monetary policy shock. The conclusions are similar: the SSDP model implies large real effects of monetary shocks, closer to the Calvo specification than to the FMC specification, and the differences across the three models increase under a Taylor rule, compared with a money growth rule.

Next, Fig. 3 plots the response of price dispersion, Δ_t, defined in (19). In our model, one reason prices vary is that firms face different productivities. But additional price dispersion, caused by failure to adjust when necessary, implies inefficient variation in demand across goods that acts as a decrease

[12] The table considers autocorrelated money growth shocks. The results for *i.i.d.* money growth are very similar, since correlation mostly changes the scale of the impulse responses, rather than their shape.

in aggregate productivity: $C_t = N_t/\Delta_t$. In a representative agent model near a zero-inflation steady state, Δ_t is negligible because it is roughly proportional to the cross-sectional variance of prices, a quantity of second order in the inflation rate.[13] But cross-sectional price variance is not second order when large idiosyncratic shocks are present, so the dispersion wedge Δ_t may be quantitatively important, especially since $\epsilon = 7$ magnifies variations in the ratio P_{it}/P_t. The first row of Fig. 3 shows that for SSDP and Calvo, increased money growth throws firms' prices further out of line with fundamentals, increasing dispersion; raising consumption therefore requires a larger increase in labor in these specifications. In contrast, the FMC case shows little change in Δ_t, because all firms with severe price misalignments do in fact adjust. Interestingly, since the Taylor rule leans against inflationary shocks, there is much less variation in the price level for the SSDP and Calvo cases in our Taylor rule simulation than there is under autocorrelated money growth. The result is that in all three specifications, variation in Δ_t is negligible after a shock to the Taylor rule.

FIGURE 3 ABOUT HERE

4.3. Inflation decompositions

To a first-order approximation, inflation can be calculated as an average of log nominal price changes. Using our grid-based notation, and starting from the beginning-of-period distribution $\widetilde{\boldsymbol{\Psi}}_t$,

$$\pi_t = \log \Pi_t = \sum_{j=1}^{\#^p} \sum_{k=1}^{\#^a} x_t^{jk} \lambda_t^{jk} \widetilde{\Psi}_t^{jk}, \tag{45}$$

where $x_t^{jk} \equiv \log\left(\frac{p_t^*(a^k)}{p^j}\right)$ is the *desired* log price adjustment of a firm with price p^j and productivity a^k. Formula (45) can be decomposed in several ways to investigate the sources of monetary nonneutrality in the model. Klenow and Kryvtsov (2008) rewrite (45) as the product of the average log price adjustment $\overline{x}_t$ times the frequency of price adjustment $\overline{\lambda}_t$:

$$\pi_t = \overline{x}_t \overline{\lambda}_t, \qquad \overline{x}_t \equiv \frac{\sum_{j,k} x_t^{jk} \lambda_t^{jk} \widetilde{\Psi}_t^{jk}}{\sum_{j,k} \lambda_t^{jk} \widetilde{\Psi}_t^{jk}}, \qquad \overline{\lambda}_t \equiv \sum_{j,k} \lambda_t^{jk} \widetilde{\Psi}_t^{jk}. \tag{46}$$

Dropping higher-order terms, this implies the following inflation decomposition:

$$\Delta\pi_t = \overline{\lambda} \Delta \overline{x}_t + \overline{x} \Delta \overline{\lambda}_t, \tag{47}$$

[13] See for example Galí (2008), p. 46 and Appendix 3.3.

where variables without time subscripts represent steady states, and Δ represents a deviation from steady state.[14] Klenow and Kryvtsov's "intensive margin", $\mathcal{I}_t^{KK} \equiv \overline{\lambda}\Delta\overline{x}_t$, is the part of inflation attributable to changes in the average price adjustment; their "extensive margin", $\mathcal{E}_t^{KK} \equiv \overline{x}\Delta\overline{\lambda}_t$, is the part due to changes in the frequency of adjustment.

Unfortunately, this decomposition does not reveal whether a rise in the average log price adjustment $\overline{x}_t$ is caused by a rise in all firms' desired adjustments, or by a reallocation of adjustment opportunities from firms desiring small or negative price changes to others wanting large price increases. That is, $\mathcal{I}_t^{KK}$ confounds changes in *desired* adjustments (the only relevant changes in the Calvo model) with the "selection effect" emphasized by GL07. To distinguish between these last two effects, inflation can instead be broken into three terms: an intensive margin capturing changes in the average *desired* log price change, an extensive margin capturing changes in *how many* firms adjust, and a selection effect capturing changes in *who* adjusts. These three effects are distinguished by rewriting (45) as

$$\pi_t = \overline{x}_t^* \overline{\lambda}_t + \sum_{j,k} x_t^{jk} \left(\lambda_t^{jk} - \overline{\lambda}_t \right) \widetilde{\Psi}_t^{jk}, \qquad \overline{x}_t^* \equiv \sum_{j,k} x_t^{jk} \widetilde{\Psi}_t^{jk}. \tag{48}$$

Note that in (48), $\overline{x}_t^*$ is the average *desired* log price change, whereas in (46), $\overline{x}_t$ is the average log price change *among those who adjust*. Thus (48) says that inflation equals the mean desired adjustment times the adjustment frequency plus a selection term $\sum_{j,k} x_t^{jk} \left(\lambda_t^{jk} - \overline{\lambda}_t \right) \widetilde{\Psi}_t^{jk} = \sum_{j,k} \lambda_t^{jk} \left(x_t^{jk} - \overline{x}_t^* \right) \widetilde{\Psi}_t^{jk}$ that can be nonzero if some changes x_t^{jk} are more or less likely than the mean adjustment probability $\overline{\lambda}_t$, or (equivalently) if firms with different probabilities of adjustment λ_t^{jk} tend to prefer adjustments that differ from the mean desired change $\overline{x}_t^*$.

Equation (48) leads to the following inflation decomposition:

$$\Delta\pi_t = \overline{\lambda}\Delta\overline{x}_t^* + \overline{x}^*\Delta\overline{\lambda}_t + \Delta \sum_{j,k} x_t^{jk} \left(\lambda_t^{jk} - \overline{\lambda}_t \right) \widetilde{\Psi}_t^{jk}. \tag{49}$$

Our intensive margin effect, $\mathcal{I}_t \equiv \overline{\lambda}\Delta\overline{x}_t^*$, is the effect of changing all firms' desired adjustment by the same amount (or more generally, changing the mean preferred adjustment in a way that is uncorrelated with the adjustment probability). $\mathcal{I}_t$ is the only nonzero term in the Calvo model, where $\lambda_t^{jk} = \overline{\lambda}$ for all j, k, t. Our extensive margin effect, $\mathcal{E}_t \equiv \overline{x}^*\Delta\overline{\lambda}_t$, is the effect of changing the fraction of

[14] Actually, Klenow and Kryvtsov (2008) propose a time series variance decomposition, whereas (46) is a decomposition of each period's inflation realization. But the logic of (46) is the same as that in their paper.

firms that adjust, assuming the new adjusters are selected randomly. Our selection effect, $\mathcal{S}_t \equiv \Delta \sum_{j,k} x_t^{jk} \left(\lambda_t^{jk} - \overline{\lambda}_t \right) \widetilde{\Psi}_t^{jk}$, is the effect of redistributing adjustment opportunities across firms with different desired changes x_t^{jk}, while fixing the overall fraction that adjust.

An alternative decomposition, proposed by Caballero and Engel (2007), also differences (45):

$$\Delta \pi_t = \sum_{j,k} \Delta x_t^{jk} \lambda^{jk} \widetilde{\Psi}^{jk} + \sum_{j,k} x^{jk} \Delta \lambda_t^{jk} \widetilde{\Psi}^{jk} + \sum_{j,k} x^{jk} \lambda^{jk} \Delta \widetilde{\Psi}_t^{jk} \tag{50}$$

They further simplify this to

$$\Delta \pi_t = \overline{\lambda} \Delta \mu_t + \sum_{j,k} x^{jk} \Delta \lambda_t^{jk} \widetilde{\Psi}^{jk} \tag{51}$$

under the assumption that all desired price adjustments change by $\Delta x_t^{jk} = \Delta \mu_t$ when money growth increases by $\Delta \mu_t$, and by taking an ergodic average so that the last term drops out.[15] Their first term, $\mathcal{I}_t^{CE} \equiv \Delta \mu_t \overline{\lambda}$, is the same as our intensive margin $\mathcal{I}_t$, if their assumption that all desired price adjustments change by $\Delta \mu_t$ is correct. But therefore, their "extensive margin" term $\mathcal{E}_t^{CE} \equiv \sum_{j,k} x^{jk} \Delta \lambda_t^{jk} \widetilde{\Psi}^{jk}$, confounds the question of *how many* firms adjust (our extensive margin $\mathcal{E}_t$) with the question of *who* adjusts (our selection effect $\mathcal{S}_t$), which is the mechanism stressed by GL07.

FIGURE 4 ABOUT HERE

The importance of identifying the selection effect separately becomes clear in Fig. 4, which illustrates our decomposition of the inflation impulse response to monetary shocks. The three components of inflation, $\mathcal{I}_t$, $\mathcal{E}_t$, and $\mathcal{S}_t$, are shown to the same scale for better comparison. The graphs demonstrate clearly (in contrast to Caballero and Engel's claim) that the short, sharp rise in inflation observed in the FMC specification results from the selection effect. This is true both under Taylor rule shocks, where inflation spikes to 1.5% on impact, of which 1.25% is the selection component, and under (autocorrelated) money growth shocks, where inflation spikes to 2.8%, with 2.25% due to selection. In contrast, inflation in the Calvo model is caused by the intensive margin only; in SSDP there is a nontrivial selection effect but it still only accounts for around one-third of the inflation response.

[15] Our equation (49) is intended to decompose each period's inflation realization, so it allows for shifts in the current distribution $\widetilde{\Psi}_t^{jk}$. Caballero and Engel instead propose a decomposition (see their eq. 17) of the *average* impact of a monetary shock. Therefore they evaluate their decomposition at the ergodic distribution (the time average over all cross-sectional distributions, called $f_A(x)$ in their paper). Since this is a fixed starting point of their calculation, they do not need to include a $\Delta f_A(x)$ term.

On the other hand, the extensive margin $\mathcal{E}_t \equiv \overline{x}^* \Delta \overline{\lambda}_t$ plays a negligible role in the inflation response. This makes sense, because the simulation assumes a steady state with zero inflation, so steady state price adjustments are responses to idiosyncratic shocks only, and the average desired adjustment $\overline{x}^*$ is essentially zero. Therefore $\mathcal{E}_t$ is negligible even though the adjustment frequency $\overline{\lambda}_t$ itself does vary.[16] The extensive margin only becomes important when there is high trend inflation, so that the average desired adjustment $\overline{x}^*$ is large and positive.

As for the intensive margin, its initial effect after a money growth shock is similar across all adjustment specifications, but it is more persistent in the Calvo and SSDP cases than in the FMC case. The scale of the intensive margin depends on the autocorrelation of money growth: the mean desired price change rises roughly one-for-one after an *i.i.d.* money growth shock (not shown), and rises by roughly five percentage points when money growth has autocorrelation $\phi_z = 0.8$ (first row of Fig. 4). Thus, in the autocorrelated case, the intensive margin is initially $\mathcal{I}_1 \equiv \overline{\lambda} \Delta \overline{x}_1^* \approx 0.5\%$. In other words, firms wish to "frontload" price adjustment by approximately the same amount in all three specifications; but many of these changes occur immediately in the FMC case (showing up as a redistribution of adjustment opportunities, *i.e.*, a selection effect), whereas they are realized gradually in the other specifications. Under a Taylor rule, the intuition is similar, bearing in mind that Fig. 4 is scaled to give an initial decline of 25 basis points in the nominal interest rate. This requires a larger underlying shock z in the FMC specification than in the other cases; thus the effect on the intensive margin is larger (but less persistent) for FMC than it is for Calvo and SSDP.

4.4. Effects of idiosyncratic shocks

Two recent empirical papers have compared how prices respond to idiosyncratic, as well as aggregate shocks (Boivin, Giannoni, and Mihov, 2007; Mackowiak, Moench, and Wiederholt, 2009). Fig. 5 shows our model's implications for idiosyncratic shocks. Specifically, it shows the *expected* response to a one standard deviation idiosyncratic productivity decrease, in the Calvo, FMC, and SSDP specifications. A productivity decrease causes a gradual price increase over time, followed by a decline back to the mean price level as the autoregressive productivity process (8) reverts. Any individual response, of

[16]The fact that the steady state has exactly zero inflation is not crucial here; $\mathcal{E}_t$ is quantitatively insignificant compared to the other inflation components at any typical OECD inflation rate.

course, is a large discrete price adjustment; but since the probability of adjustment in any given month is much less than one, the average response is slow, reaching its peak after eight months in the SSDP model.[17] The speed of response is similar to that in the Calvo specification, in contrast with the FMC case, where the maximum impact occurs after only three months. On the other hand, the maximum response is much larger in the SSDP case than in the Calvo model; this is one dimension along which SSDP resembles the FMC model. In other words, even though firms in the SSDP model suffer Calvo-like adjustment lags, the fact that the probability of adjustment increases with the value of adjustment protects them from the risk of exceptionally bad price misalignments. They are therefore more willing to react to idiosyncratic shocks than Calvo firms are.

FIGURE 5 ABOUT HERE

These results might raise doubts about the SSDP model's consistency with empirical evidence, since Boivin *et al.* (2007) and Mackowiak *et al.* (2009) claim that the response to idiosyncratic shocks is much faster than that to aggregate shocks. To see whether our findings contradict these previous papers, Fig. 6 reports the results of running the estimation routine of Mackowiak *et al.* on panel data produced by the SSDP model. The simulated data cover the prices of 79 firms over 245 months, which is the same structure of observations as in Mackowiak *et al.*, except that their observations correspond to sectors, whereas ours correspond to individual firms. The results are remarkably similar to those in Figs. 1-2 of Mackowiak *et al.* (2009). In particular, the estimated response to an idiosyncratic shock is immediate and essentially permanent, as those authors found. In contrast, the estimated response to aggregate shocks appears more sluggish (and is statistically indistinguishable from zero in our case).

FIGURE 6 ABOUT HERE

What causes the estimation routine to characterize the response to an idiosyncratic productivity shock in this way, in stark contrast with the true response, shown in Fig. 5? The problem is that the true idiosyncratic shocks in microdata are unknown to an econometrician, so Mackowiak *et al.* identify them as *residual price increases* not explained by aggregate shocks. In the SSDP model, individual prices typically respond with a lag to the true productivity shock. But in the estimation routine, the moment of the shock *corresponds by assumption* to the moment of the price increase, so the response

[17]The responses shown are averaged both with respect to the steady state distribution of prices and productivities, and with respect to the realization of the adjustment process λ.

is estimated to be immediate.

Mean reversion of the idiosyncratic component occurs by discrete individual price jumps in the model, whereas Mackowiak *et al.* assume that this component (Bv in their eq. 1) decays smoothly. Hence their estimation procedure interprets price adjustments that revert to the mean as a sequence of new idiosyncratic shocks that happen to go in the opposite direction (which is why the initial shock is interpreted as permanent). Thus, results from this procedure (or others that identify idiosyncratic shocks as a residual, *e.g.* Boivin *et al.*) should be treated with caution. Our Monte Carlo exercise shows that, at least in some cases, it may exaggerate the speed of response to idiosyncratic shocks, which might suggest stronger state dependence than the data actually warrant.

5. Conclusions

This paper has computed the impact of monetary policy shocks in a quantitative macroeconomic model of state-dependent pricing. It has calibrated the model for consistency with microeconomic data on firms' pricing behavior, estimating how the probability of price adjustment depends on the value of adjustment. Given the estimated adjustment function, the paper has characterized the dynamics of the distribution of prices and productivities in general equilibrium.

The calibrated model implies that prices rise gradually in response to monetary stimulus, causing a large, persistent rise in consumption and labor. Looking across specifications, the main factor determining how monetary shocks propagate through the economy is the degree of state dependence. That is, increasing the autocorrelation of money growth shocks simply makes their effects larger, without any notable change in the shape or persistence of the implied impulse responses. In contrast, decreasing the degree of state dependence from the extreme of fixed menu costs (FMC) to the opposite extreme of the Calvo (1983) model strongly damps the initial inflation spike caused by a money growth shock and increases its effect on real variables. The parameterization most consistent with microdata (labelled "SSDP" throughout the paper) is fairly close to the Calvo model in terms of its quantitative effects. The conclusions are similar if the monetary authority follows a Taylor rule instead of a money growth rule, except that the difference across adjustment specifications becomes even stronger, and the monetary nonneutrality of the SSDP specification is increased.

This paper also decomposes the impulse response of inflation into an intensive margin effect relating to the average desired price change, an extensive margin effect relating to the number of firms adjusting, and a selection effect relating to the relative frequencies of small and large or negative and positive adjustments. Under the preferred (SSDP) calibration, about two-thirds of the effect of a monetary shock comes through the intensive margin, and most of the rest through the selection effect. The extensive margin is negligible unless the economy starts from a high baseline inflation rate. Under the FMC specification, a monetary shock instead causes a quick increase in inflation, driven by the selection effect, which eliminates most of its effects on real variables.

Since the selection effect represents changes in the adjustment probability across firms, its strength depends directly on the degree of state dependence. We say that the state dependence is strong in a model of fixed menu costs because they make λ a step function: at the threshold, a tiny increase in the value of adjustment suffices to raise the adjustment probability from 0 to 1. Therefore the distribution of price changes consists of two spikes: there are no small changes, and firms change their prices as soon as they pass the adjustment thresholds. Hence, in steady state, those firms that might react to monetary policy are all near the two adjustment thresholds; a monetary stimulus decreases λ from 1 to 0 for some firms desiring a price decrease, while increasing λ from 0 to 1 for others preferring an increase, making the inflation response quick and intense. That is, the same property which makes money nearly neutral in the FMC model is the one which makes that model inconsistent with price microdata. A model in which adjustment depends more smoothly on the value of adjusting fits microdata better and yields larger real effects of monetary policy.

References

Akerlof, G., and Yellen, J., 1985. A near-rational model of the business cycle with wage and price inertia. Quarterly Journal of Economics 100 (Supplement), 823-838.

Boivin, J., Giannoni, M., Mihov, I., 2009. Sticky prices and monetary policy: Evidence from disaggregated US data. American Economic Review 99, 350-384.

Caballero, R., Engel, E., 1993. Microeconomic rigidities and aggregate price dynamics. European Economic Review 37, 697-711.

Caballero, R., Engel, E., 1999. Explaining investment dynamics in U.S. manufacturing: a generalized (S,s) approach. Econometrica 67, 741-82.

Caballero, R., Engel, E., 2007. Price stickiness in Ss models: new interpretations of old results. Journal of Monetary Economics 54, 100-121.

Calvo, G., 1983. Staggered prices in a utility-maximizing framework. Journal of Monetary Economics 12, 383-398.

Caplin, A., Spulber, D., 1987. Menu costs and the neutrality of money. Quarterly Journal of Economics 102, 703-726.

Costain, J., Nakov, A., 2008. Price adjustments in a general model of state-dependent pricing. Banco de España Working Paper #0824. Forthcoming, Journal of Money, Credit and Banking.

Den Haan, W., 1997. Solving dynamic models with aggregate shocks and heterogeneous agents. Macroeconomic Dynamics 1, 355-386.

Dotsey, M., King, R., Wolman, A., 1999. State-dependent pricing and the general equilibrium dynamics of money and output. Quarterly Journal of Economics 114, 655-690.

Dotsey, M., King, R., Wolman, A., 2008. Inflation and real activity with firm-level productivity shocks. Manuscript, Boston University.

Eichenbaum, M., Jaimovich, N., Rebelo, S., 2008. Reference prices and nominal rigidities. NBER Working Paper #13829.

Gagnon, E., 2009. Price setting during low and high inflation: evidence from Mexico. Quarterly Journal of Economics 124, 1221-1263.

Galí, J., 2008. Monetary Policy, Inflation, and the Business Cycle: an Introduction to the New Keynesian Framework. Princeton University Press, Princeton, New Jersey.

Gertler, M., Leahy, J., 2008. A Phillips curve with an (S,s) foundation. Journal of Political Economy 116, 533-572.

Golosov, M., Lucas, R., 2007. Menu costs and Phillips curves. Journal of Political Economy 115, 171-199.

Guimaraes, B., Sheedy, K., 2010. Sales and monetary policy. Forthcoming, American Economic Review.

Kehoe, P., Midrigan, V., 2010. Prices are sticky after all. NBER Working Paper 16364.

Klein, P., 2000. Using the generalized Schur form to solve a multivariate linear rational expectations model. Journal of Economic Dynamics and Control 24, 1405-1423.

Klenow, P., Kryvtsov, O., 2008. State-dependent or time-dependent pricing: does it matter for recent US inflation? Quarterly Journal of Economics 123, 863-894.

Krusell, P., Smith, A., 1998. Income and wealth heterogeneity in the macroeconomy. Journal of Political Economy 106, 245-272.

Mackowiak, B., Moench, E., Wiederholt, M., 2009. Sectoral price data and models of price setting. Journal of Monetary Economics 56, S78-S99.

Midrigan, V., 2008. Menu costs, multi-product firms and aggregate fluctuations, Manuscript, New York University.

Nakamura, E., Steinsson, J., 2008. Five facts about prices: a reevaluation of menu cost models. Quarterly Journal of Economics, 123, 1415-1464.

Reiter, M., 2009. Solving heterogeneous-agent models by projection and perturbation. Journal of Economic Dynamics and Control 33, 649-655.

Woodford, M., 2009. Information-constrained state-dependent pricing. Journal of Monetary Economics 56, S100-S124.

Yun, T., 2005. Optimal monetary policy with relative price distortions. American Economic Review 95, 89-109.

Tables and Figures

Table 1: Adjustment specifications

Specification	Adjustment probability $\lambda(L)$	Mean gains, in units of time: $G(P, A, \Omega)/W(\Omega)$
Calvo	$\bar{\lambda}$	$\bar{\lambda}L(P, A, \Omega)$
Fixed MC	$\mathbf{1}\{L \geq \alpha\}$	$\lambda(L(P, A, \Omega))[L(P, A, \Omega) - \alpha]$
Woodford	$\bar{\lambda}/[\bar{\lambda} + (1 - \bar{\lambda})\exp(\xi(\alpha - L))]$	$\lambda(L(P, A, \Omega))L(P, A, \Omega)$
Stoch. MC	$\bar{\lambda}/[\bar{\lambda} + (1 - \bar{\lambda})(\alpha/L)^{\xi}]$	$\lambda(L(P, A, \Omega))[L(P, A, \Omega) - E(\kappa \vert \kappa < \lambda(L(P, A, \Omega)))]$
SSDP	$\bar{\lambda}/[\bar{\lambda} + (1 - \bar{\lambda})(\alpha/L)^{\xi}]$	$\lambda(L(P, A, \Omega))L(P, A, \Omega)$

Note: $\lambda(L)$ is the probability of price adjustment; L is the real loss from failure to adjust, as a function of firm's price P and productivity A, and aggregate conditions Ω. G represents mean nominal gains from adjustment; dividing by the nominal wage W converts gains to real terms. $\bar{\lambda}$, α and ξ are parameters to be estimated.

Table 2. Steady-state simulated moments for alternative estimated models and evidence

	Productivity parameters See eq. (8) for definitions	Adjustment parameters See Table 1 for definitions
Calvo	$(\sigma_\varepsilon, \rho) = (0.0850, 0.8540)$	$\bar{\lambda} = 0.10$
Fixed MC	$(\sigma_\varepsilon, \rho) = (0.0771, 0.8280)$	$\alpha = 0.0665$
Woodford	$(\sigma_\varepsilon, \rho) = (0.0924, 0.8575)$	$(\bar{\lambda}, \alpha, \xi) = (0.0945, 0.0611, 1.3335)$
Stoch. MC	$(\sigma_\varepsilon, \rho) = (0.0676, 0.9003)$	$(\bar{\lambda}, \alpha, \xi) = (0.1100, 0.0373, 0.2351)$
SSDP	$(\sigma_\varepsilon, \rho) = (0.0677, 0.9002)$	$(\bar{\lambda}, \alpha, \xi) = (0.1101, 0.0372, 0.2346)$

Moments	Model					Evidence			
	Calvo	FMC	Wdfd	SMC	SSDP	MAC	MD	NS	KK
Frequency of price changes	10.0	10.0	10.0	10.0	10.0	20.5	19.2	10	13.9
Mean absolute price change	6.4	17.9	10.3	10.0	10.1	10.5	7.7		11.3
Std of price changes	8.2	18.4	13.6	12.2	12.2	13.2	10.4		
Kurtosis of price changes	3.5	1.3	4.0	2.9	2.9	3.5	5.4		
% price changes $\leq$5% in abs value	47.9	0.0	37.0	26.3	26.3	25	47		44
Mean loss in % of frictionless profit	36.8	10.6	37.4	25.6	25.6				
Mean loss in % of frictionless revenue	5.2	1.5	5.3	3.6	3.6				
Fit: Kolmogorov-Smirnov statistic	0.111	0.356	0.038	0.024	0.025				
Fit: Euclidean distance	0.159	0.409	0.072	0.060	0.056				

Note: Price statistics refer to non-sale consumer price changes and are stated in percent. The last four columns report statistics from Midrigan (2008) for AC Nielsen (MAC) and Dominick's (MD), Nakamura and Steinsson (2008) (NS), and Klenow and Kryvtsov (2008) (KK). To calibrate the productivity parameters ρ and σ_ε^2, together with the adjustment parameters $\bar{\lambda}$, α and ξ, we minimize a distance criterion with two terms, (1) the difference between the median frequency of price changes in the model (fr) and in the data, and (2) the distance between the histogram of log price changes in the model ($histM$) and the data ($histD$): $\min(25 \|fr - 0.10\| + \|histM - histD\|)$.

Table 3. Variance decomposition and Phillips curves of alternative models

	Data	SSDP	Calvo	FMC
Std of quarterly inflation (×100)	0.246	0.246	0.246	0.246
% explained by nominal shock		100	100	100
Money growth rule (see eq. 16-17)				
Std of money growth shock (×100)		0.174	0.224	0.111
Std of detrended output (×100)	0.909	0.586	1.053	0.121
% explained by money growth shock		64.5	115.9	13.3
Slope coeff. of the Phillips curve		0.598	1.069	0.134
Standard error		0.004	0.039	0.005
Taylor rule (see eq. 18)				
Std of Taylor rule shock (×100)		0.393	0.918	0.129
Std of detrended output (×100)	0.909	0.995	2.741	0.134
% explained by Taylor rule shock		109.6	301.6	14.7
Slope coeff. of the Phillips curve		1.055	2.785	0.126
Standard error		0.093	0.290	0.006

Note: for each monetary regime (Taylor or money growth rule) and each pricing model, the nominal shock is scaled to account for 100% of the standard deviation of inflation. The volatility of output in the data is measured as the standard deviation of HP-filtered quarterly log real GDP. The "slope coefficients" are the estimates of β_2 in a 2SLS regression of (log) consumption on inflation, instrumented by the exogenous nominal shock. The first stage regression is $\pi_t^q = \alpha_1 + \alpha_2 \mu_t^q + \epsilon_t$, and the second stage is $c_t^q = \beta_1 + \beta_2 (4\hat{\pi}_t^q) + \varepsilon_t$, where $\hat{\pi}_t^q$ is the prediction for inflation from the first-stage and the superscript q denotes conversion to quarterly frequency.

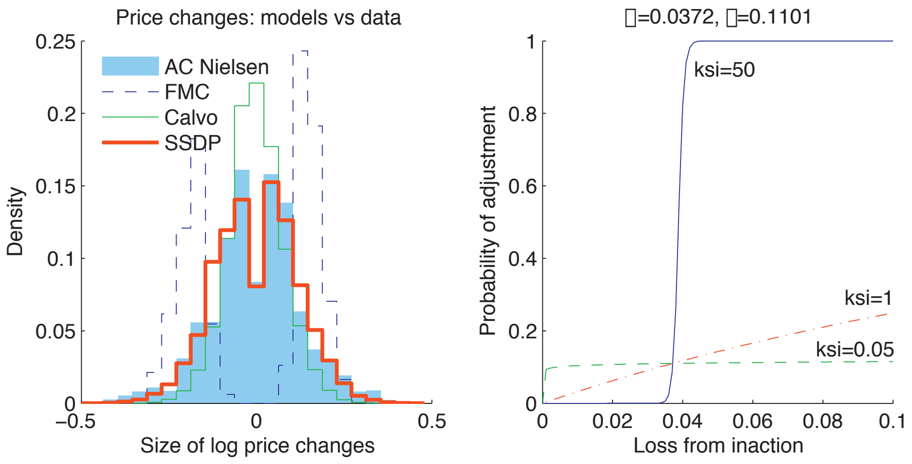

Fig. 1. Price change distributions and adjustment function

Note: (left panel) size distribution of price changes: data vs. models; (right panel) Adjustment function for alternative values of ξ

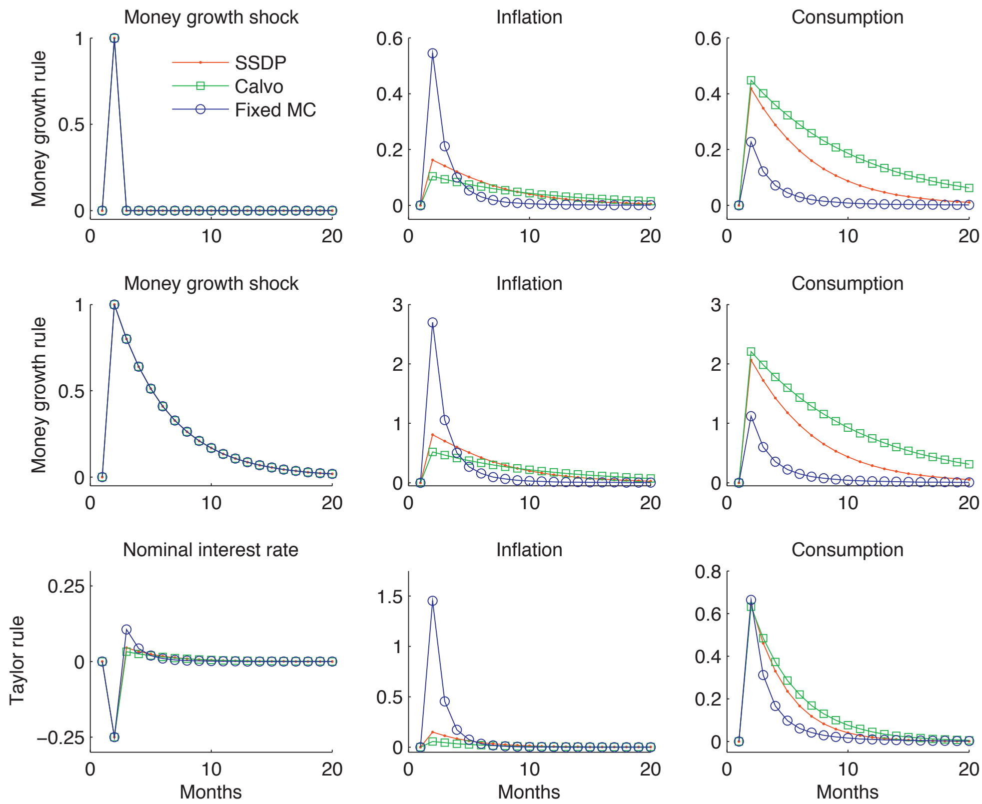

Fig. 2. The real effects of nominal shocks across models

Note: (top row) responses of inflation and consumption to an iid money growth shock; (middle row) responses to a correlated money growth shock; (bottom row) responses to a Taylor rule shock. Inflation responses are in percentage points; consumption responses are in percent deviation from steady-state. Lines with dots - benchmark SSDP model; lines with squares - Calvo; lines with circles - fixed menu cost

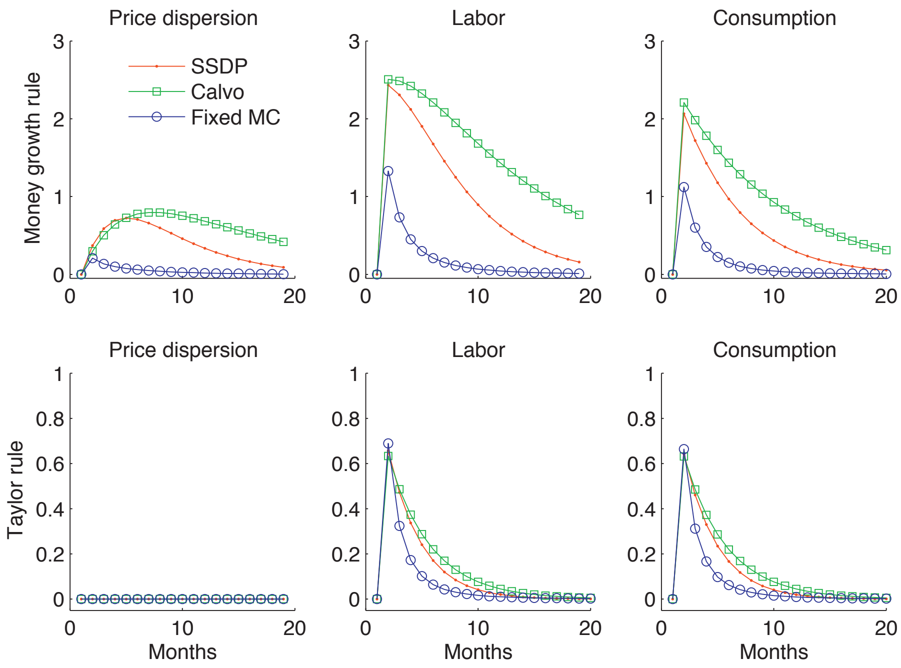

Fig. 3. Price dispersion across models

Note: (top row): responses to a correlated money growth shock; (bottom row): responses to a Taylor rule shock. The responses are in percent deviation from steady-state. Lines with dots - benchmark SSDP model; lines with squares - Calvo; lines with circles - fixed menu cost

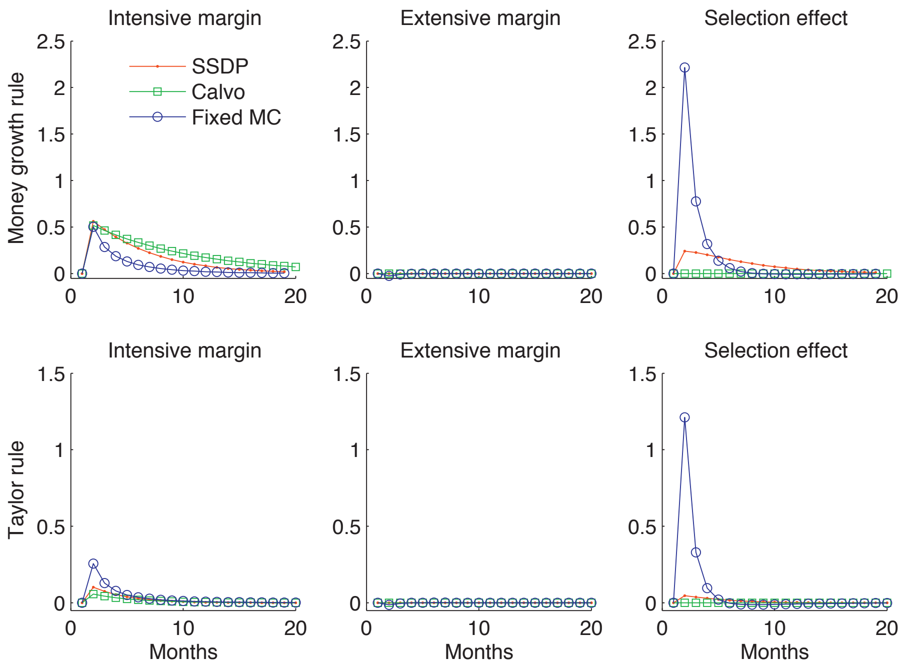

Fig. 4. Inflation decomposition across models

Note: decomposition of the inflation response into an intensive margin, extensive margin, and selection effect (see eq.54) (top row): responses to a correlated money growth shock. (bottom row): responses to a Taylor rule shock. The responses are in percentage points and sum up to the total inflation response shown in figure 2. Lines with dots - benchmark SSDP model; lines with squares - Calvo; lines with circles - fixed menu cost

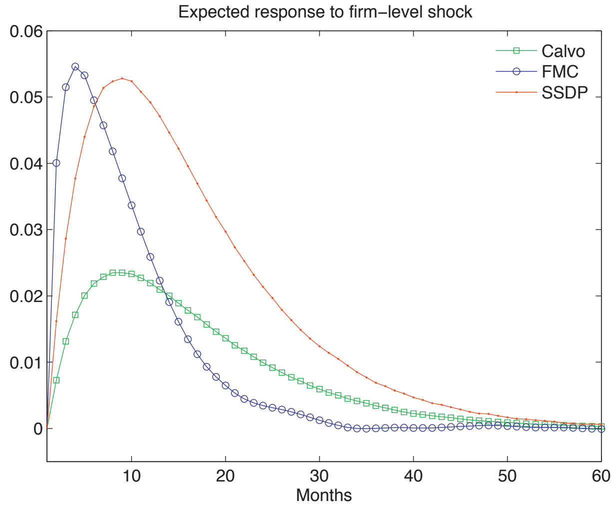

Fig. 5. Theoretical mean response to an idiosyncratic productivity shock across models

Lines with dots - benchmark SSDP model; lines with squares - Calvo; lines with circles - fixed menu cost

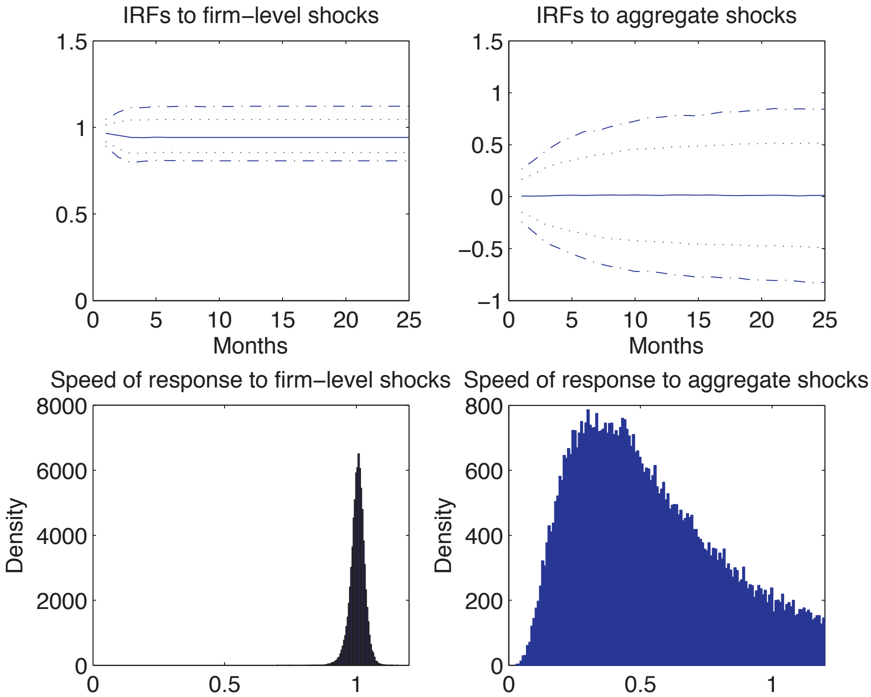

Fig. 6. Estimated price responses and speed of response from model-generated data

Note: (top, left): price responses to a firm-specific shock estimated on SSDP model-generated data; (top, right): price responses to an aggregate shock estimated on SSDP model-generated data; (bottom, left): speed of price response to a firm-specific shock; (bottom, right): speed of responses to an aggregate shock. The estimation is done by applying the procedure of Mackowiak, Moench, and Wiederholt to a panel of price series generated by the SSDP model with both idiosyncratic and aggregate shocks present.