Dynamics of the Price Distribution in a General Model of State-Dependent Pricing

James Costain
Banco de España

Anton Nakov
Banco de España

November 2008

Abstract

This paper analyzes the effects of monetary shocks in a DSGE model that allows for a general form of smoothly state-dependent pricing by firms. As in Dotsey, King, and Wolman (1999) and Caballero and Engel (2007), our setup is based on one fundamental property: firms are more likely to adjust their prices when doing so is more valuable. The exogenous timing (Calvo 1983) and fixed menu cost (Golosov and Lucas 2007) models are nested as limiting cases of our setup.

Our model is calibrated to match the steady-state distribution of price adjustments in microdata; realism calls for firm-specific shocks. Computing a dynamic general equilibrium requires us to calculate how the distribution of prices and productivities evolves over time. We solve the model using the method of Reiter (2008), which is well-suited to this type of problem because it combines a fully nonlinear treatment of firm-level state variables with a linearization of the aggregate dynamics.

We compute impulse responses to *iid* and autocorrelated money growth shocks, and decompose the inflation impact into 'intensive margin', 'extensive margin' and 'selection' components. Under our most successful calibration, increased money growth causes a persistent rise in inflation and output. The real effects are substantially larger if money growth is autocorrelated. In contrast, if we instead impose a fixed menu cost specification, money growth shocks cause a sharp spike in inflation (via the selection component) so that the real effects are small and short-lived, especially if money growth is *iid*.

An increase in aggregate productivity raises consumption but causes labor to fall. Also, impulse responses differ depending on the distribution at the time the shock occurs. In particular, increased money growth has different effects starting from the steady state distribution than it does if all firms have recently received an economy-wide productivity shock.

Keywords: Price stickiness, state-dependent pricing, stochastic menu costs, generalized (S,s), heterogeneous agents, distributional dynamics

JEL Codes: E31, E52, D81

1 Introduction[1]

Sticky prices are an important ingredient in modern dynamic general equilibrium models, including those used by central banks for policy analysis. But how exactly to model price stickiness remains just as controversial as ever. The Calvo (1983) model's fixed probability of adjustment is popular for its analytical tractability, but lacks the theoretical appeal of a microfounded framework immune to the Lucas critique. In an influential article, Golosov and Lucas (2007) studied a model of price stickiness microfounded on the basis of fixed "menu costs". They calibrated their model to match certain moments of the distribution of price changes in US microdata, and found only small and transitory real effects of monetary shocks. The implication is that the larger real effects found under the Calvo setup are exaggerated and therefore misleading for policy purposes.

In this paper, we calibrate and simulate a general model of smoothly state-dependent pricing by firms that nests a variety of influential pricing models. As in Dotsey, King, and Wolman (1999) and Caballero and Engel

[1]The authors wish to thank Michael Reiter, Virgiliu Midrigan, Rudi Bachmann, Omar Licandro, and seminar participants at the Vienna Institute for Advanced Studies, the European University Institute, ESSIM 2008, CEF 2008, REDg Madrid 2008, and ASSET 2008 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 and do not necessarily coincide with those of the Bank of Spain or the Eurosystem.

(2007), our setup rests on one fundamental property: firms are more likely to adjust their prices when doing so is more valuable. The Calvo and fixed menu cost models are nested as two opposite limiting special cases of our framework. We calibrate our model to match the distribution of price adjustments found in recent US microdata (Klenow and Kryvstov 2008; Midrigan 2008; Nakamura and Steinsson 2007), and in doing so we estimate a parameter that controls the degree of state dependence of pricing behavior. Unlike Golosov and Lucas (2007), our calibration is consistent with the fact that small price changes are common in the data, alongside frequent large price adjustments (Midrigan 2008).[2]

Our impulse response analysis shows that increased money growth causes a persistent rise in both inflation and output, with real effects only slightly weaker than those in the Calvo model. We show how to decompose the impulse response of inflation into three parts: the intensive margin (driven by changes in the average desired price adjustment), the extensive margin (driven by changes in the fraction of firms adjusting), and the selection effect (driven by changes in which firms adjust). Under our baseline parameterization, about two-thirds of the inflation response to a money supply shock comes from the intensive margin, and most of the rest from the selection effect. Golosov and Lucas' (2007) finding that uncorrelated money growth shocks have only a small and transitory effect on output derives from the strong selection effect their model generates. But we find that their specification exaggerates the selection effect because it implies a degree of state dependence inconsistent with microdata. In particular, extreme state dependence is the same property that generates a sharply bimodal distribution of price adjustments, which is why our estimate prefers milder state dependence.

Whereas Golosov and Lucas restrict attention to *iid* money growth shocks, we also study the autocorrelated case. In all versions of the model, including the fixed menu cost specification, making money shocks autocorrelated substantially increases their real effects. However, the shape and persistence of the response is primarily determined by the degree of state dependence, not by the degree of autocorrelation of the driving process. Thus we find large differences in behavior between our calibrated model and a fixed menu cost specification, but little difference between the impact of autocorrelated and uncorrelated money shocks, except for a rescaling.

As many authors have emphasized, matching microdata on price adjustments makes it essential to allow for firm-specific shocks, and the presence of these shocks potentially changes the stickiness of prices. Idiosyncratic shocks have other quantitatively important implications too: in particular, we find they imply that changes in price dispersion have a first-order impact on productivity. But obviously, including firm-specific shocks complicates the analysis, because it implies a heterogeneous agent problem in which the entire distribution of prices and productivities across firms becomes a state variable. Methodologically, our main contribution is to show how to characterize the general equilibrium dynamics using the algorithm of Reiter (2008). This method is well suited for problems in which idiosyncratic shocks matter more to the individual decision maker than aggregate shocks do, because it is fully nonlinear in idiosyncratic factors even though it imposes linearity in aggregate factors. Moreover, it is easy to implement because each step in the calculation is a familiar numerical procedure. First, it involves calculating the steady state equilibrium, which means solving a backwards induction problem on a grid repeatedly until a fixed point for the aggregate price level is found. Second, the aggregate dynamics are solved linearly, which can be done with standard methods (*e.g.* Klein 2000; Sims 2001) in spite of the fact that this involves a very large system of equations representing values and densities at all points on the grid.

1.1 Related literature

Few prior studies on state-dependent pricing have attempted to calculate a dynamic general equilibrium with firm-specific shocks. Instead, much research on how state-dependent pricing aggregates has looked at partial equilibrium models, as in Caballero and Engel (1993, 2007) and Klenow and Kryvstov (2008). Some papers have demonstrated surprising aggregation properties implied by special idiosyncratic shock processes, including Caplin and Spulber (1987), Caplin and Leahy (1997), Gertler and Leahy (2005), and Damjanovic and Nolan (2005). An important step forward to a more standard general equilibrium framework for state-dependent pricing was taken by Dotsey, King, and Wolman (1999). But their solution method relied on reducing the dimensionality of the aggregate state by ignoring idiosyncratic shocks, so that all firms that adjust at a given point in time choose the same price. While heterogeneity may average out in many macroeconomic contexts, it is not so easily ignored in the debate over price stickiness, because firm-level shocks could be crucial for

[2] In the data, small and large adjustments coexist even within narrowly defined product categories.

firms' incentives to adjust prices. This makes it hard to draw unambiguous conclusions about the importance of state-dependent pricing from the Dotsey *et al.* (1999) setup.

Golosov and Lucas (2007) were the first to confront these issues head on, studying a menu cost model in general equilibrium with idiosyncratic productivity shocks, and they obtained a striking near-neutrality result. There is some debate between Golosov and Lucas (2007) and Midrigan (2008) on one hand, and Caballero and Engel (2007) on the other, as to whether this result is caused by selection effects; we provide a quantitative definition of the selection component which does, in fact, prove to be the main factor explaining monetary near-neutrality in Golosov and Lucas' setup. We also extend Golosov and Lucas' work by comparing *iid* and autocorrelated money growth shocks. Moreover, our solution method avoids making the assumption of constant consumption that their solution required. While this assumption was roughly valid for the case they considered, we show that it is violated under more general forms of state-dependent pricing and under autocorrelated money growth shocks.

Several recent papers, like our own, have remarked that Golosov and Lucas' model generates a counterfactual distribution of price adjustments in which small changes never occur. They have proposed some more complex pricing models to fix this problem, including sectoral heterogeneity in menu costs (Klenow and Kryvstov, 2008),[3] multiple products on the same "menu" combined with leptokurtic technology shocks (Midrigan, 2008), a mix of flexible- and sticky-price firms together with a mix of two distributions of productivity shocks (Dotsey, King, and Wolman, 2008). While the latter two models are quite successful at matching the distribution of price changes, we propose a much simpler approach: we just allow the probability of price adjustment to increase smoothly with the value of adjustment (in contrast with the discontinuous jump in probability that occurs in a fixed menu cost model). Implementing this approach requires us to impose some smooth family of functions to represent the adjustment probability, and estimate its parameters. There are just three free parameters in the family of functions we choose, but our success in reproducing the distribution of price changes is similar to that of the aforementioned papers.[4]

Like us, Dotsey, King, and Wolman (2008) and Midrigan (2008) calculate general equilibrium dynamics. Dotsey *et al.* (2008) build on their earlier (1999) model with identical firms by adding first two possible idiosyncratic states, then as many as five. They find that the non-monotonic impulse responses ("echo effects") observed in their original model disappear as additional microstates are included. Their computational approach is complicated by the need to keep track of how many firms changed prices at each prior point in time and in each possible idiosyncratic state. Midrigan (2008) instead computes general equilibrium dynamics by the method of Krusell and Smith (1998). This method has several disadvantages. First, it requires a guess about which moments of the distribution will best summarize shifts in the value function, which may not be obvious.[5] Second, one must verify the guess numerically by solving Krusell and Smith's fixed-point problem (mutual consistency between the value function and the law of motion). Third, the inaction region in fixed menu cost models implies that money supply shocks could have substantially different effects starting from different distributions that share the same mean, which calls into question the approximate aggregation property that underpins the Krusell and Smith method.

Reiter's (2008) method provides a more straightforward way of tackling the distributional dynamics of a state-dependent pricing model. In contrast to Dotsey *et al.* (2008), it more fully exploits the recursive structure of the model, simply keeping track of the distribution of prices and productivities, with no need to know who adjusted when. In contrast to Krusell and Smith (1998), there is no need to search for an adequate summary statistic for the distribution. In contrast to Den Haan (1997), there is no need to impose a specific functional form on the distribution. We hope to convince the reader that by combining a standard backward-induction problem with a standard linearization of the dynamics, Reiter's method allows us to characterize distributional dynamics in a way that is transparent to write down and straightforward to program, yet provides a thorough recursive description of general equilibrium. At the same time, we show that none of the complications Dotsey *et al.* and Midrigan tack on to the fixed menu cost framework are crucial for their most important finding. Simply smoothing out individual decisions suffices to reproduce the distribution of price changes, and leads to a calibration with substantial monetary non-neutrality. As in Dotsey *et al.* (2008) and Midrigan (2008), the

[3] Damjanovic and Nolan (2005) also study sectoral heterogeneity in menu costs, but they focus on explaining differences in frequency and timing of adjustment across sectors, whereas Klenow and Kryvstov (2008), like us, attempt to reproduce the distribution of price adjustments.

[4] Our model also behaves well in the face of large changes in the steady state inflation rate: see our companion paper, Costain and Nakov (2008).

[5] Midrigan's summary statistic is the cross-sectional mean of the product of the lagged price with current productivity.

effects of money growth shocks in our model resemble those of the Calvo framework much more than those found under fixed menu costs.

2 Partial equilibrium: the firm's problem

We begin by explaining how we model price stickiness. For this purpose, it suffices to study the partial equilibrium problem of a monopolistically competitive firm. Later we will show how our firm's problem fits into an otherwise-standard dynamic stochastic general equilibrium.

Klenow and Kryvstov (2008) and Golosov and Lucas (2007) have argued convincingly that firms often suffer large idiosyncratic shocks. Thus, if they are fully rational, fully informed, and capable of frictionless adjustment, firms will adjust their prices every time a new shock is realized. We instead assume 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. The adjustment benefit is calculated from the firm's Bellman equation: 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 is the adjustment benefit (or the loss from failing to adjust). The function $\lambda(L)$ that gives the adjustment probability in terms of the loss L from failing to adjust is taken as a primitive of the model. We will choose a specification for λ that makes it easy to nest different models by appropriately setting a few parameters.

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, as in Akerlof and Yellen (1985), "stickiness" can be seen as a minimal deviation from rational expectations behavior, in which firms sometimes fail to react to changed conditions if the cost of such errors is small. Perhaps failure to adjust occurs because information itself is "sticky" (as in Reis, 2006); or perhaps because managers face information processing constraints (as in Woodford 2008); rather than taking a stand on this, we just regard our assumption as an axiom to be imposed on near-rational, near-full-information behavior. Our framework stays close to 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, allowing us to deviate smoothly from the standard rational case to nest and compare other nearby forms of behavior.

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 an idiosyncratic productivity process A_{it}:

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

Firms are monopolistic competitors, facing the demand curve $Y_{it} = \vartheta_t P_{it}^{-\epsilon}$, where ϑ_t represents aggregate demand, and we assume they fulfill all demand at the price they set. They hire in competitive labor markets at wage rate W_t, so per-period profits are

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

We call the aggregate state of the economy Ω_t. There is no need to specify the structure of Ω_t yet, except to say that it is a Markov process which determines the aggregate endogenous variables: $\vartheta_t = \vartheta(\Omega_t)$, $W_t = W(\Omega_t)$. Idiosyncratic productivity A_{it} is driven by an unchanging Markov process, *iid* across firms and unrelated to Ω_t. Thus A_{it} is correlated with A_{it-1} but is uncorrelated with all other shock processes in the model. There is nothing essential about our assumption that variations in demand are related to aggregate conditions, with idiosyncratic shocks to productivity only; we focus on this case only for consistency with related papers (e.g. Golosov and Lucas, 2007, and Reis, 2006) and to keep notation simple. It could be interesting in the future to distinguish between shocks to demand and productivity both at the firm and at the aggregate level for greater quantitative realism, but for our current purpose of investigating the relevance of state dependence for price stickiness this seems unimportant.

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, $V(P_{it}, A_{it}, \Omega_t)$, which also depends on the aggregate state of the economy. When a firm adjusts, we assume it chooses the best price conditional on its current productivity shock and on the aggregate state. Therefore, the value function of an adjusting firm, after netting out any costs that may be required to make the adjustment, is just $V^*(A_{it}, \Omega_t) \equiv \max_{P'} V(P', A_{it}, \Omega_t)$. The value of adjusting to the optimal price, written in the same units as the value function, is then

$$D(P_{it}, A_{it}, \Omega_t) \equiv \max_{P'} V(P', A_{it}, \Omega_t) - V(P_{it}, A_{it}, \Omega_t) \tag{1}$$

Of course, we don't want the adjustment probability to differ when values are denominated in euros instead of pesetas. In order to take the function λ that maps the value of adjusting into the probability of adjusting as a primitive of the model, we must be sure to write it in the appropriate units. Under either interpretation 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 suppose the adjustment probability is related to the labor effort associated with obtaining new information and/or recomputing the optimal price.[6] 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 $\lambda\left(L\left(P_{it}, A_{it}, \Omega_t\right)\right)$, where $L\left(P_{it}, A_{it}, \Omega_t\right) = \frac{D(P_{it}, A_{it}, \Omega_t)}{W(\Omega_t)}$ and λ is a given weakly increasing function which we take as a primitive of the model.

For clarity, we will distinguish between the firm's beginning-of-period price, $\widetilde{P}_{it} \equiv P_{it-1}$, and the price at which it produces and sells at time t, P_{it}, which may or may not differ from $\widetilde{P}_{it}$. Adjustments occurs with state-dependent probability λ:

$$P_{it} = \begin{cases} P^*(A_{it}, \Omega_t) \equiv \arg\max_{P'} V(P', A_{it}, \Omega_t) & \text{with probability } \lambda\left(\frac{D(\widetilde{P}_{it}, A_{it}, \Omega_t)}{W(\Omega_t)}\right) \\ \widetilde{P}_{it} \equiv P_{i,t-1} & \text{with probability } 1 - \lambda\left(\frac{D(\widetilde{P}_{it}, A_{it}, \Omega_t)}{W(\Omega_t)}\right) \end{cases}$$

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

$$\lambda\left(L\right) \equiv \frac{\overline{\lambda}}{\overline{\lambda} + (1 - \overline{\lambda})\left(\frac{\alpha}{L}\right)^{\xi}} \tag{2}$$

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 $\xi > 1$. It has fatter tails than the normal *cdf*, which may help it match the fat tails of the observed adjustment distribution emphasized by Midrigan (2008).

Note that the parameter ξ can be interpreted as controlling the degree of state dependence. The value of adjustment, L, is the summary statistic relevant for determining whether or not a firm should adjust. In the limit $\xi = 0$, our model nests that of Calvo (1983), with $\lambda\left(L\right) = \overline{\lambda}$, so that literally speaking the adjustment probability is independent of the relevant state. At the opposite extreme, our setup nests a fixed menu cost model. Taking the limit as $\xi \to \infty$, $\lambda\left(L\right)$ becomes the indicator function $\mathbf{1}\left\{L \geq \alpha\right\}$, which has value 1 whenever $L \geq \alpha$ and is zero otherwise. This has very strong state dependence, in the sense that the adjustment probability jumps from 0 to 1 when the state L passes threshold α. Under all intermediate values of ξ, the probability increases smoothly as a function of the state L. In this sense, choosing ξ to match microdata means determining what degree of state dependence is most consistent with observed firm behavior.

We are now ready to write 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

[6] Studies of the managerial costs of price adjustment, like Zbaracki *et al.* (2005), naturally start by calculating all costs in units of labor time, even if these are then converted to dollar values.

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 $\left(P - \frac{W(\Omega)}{A}\right)\vartheta(\Omega)P^{-\epsilon}$. The firm anticipates adjusting or not adjusting in the next period depending on the benefits of adjusting at that time. Therefore, using primes to denote next period's values, the Bellman equation is:

$$V(P,A,\Omega) = \left(P - \frac{W(\Omega)}{A}\right)\vartheta(\Omega)P^{-\epsilon} +$$
$$+ E\left\{ Q\left(\Omega,\Omega'\right)\left[\left(1-\lambda\left(\tfrac{D(P,A',\Omega')}{W(\Omega')}\right)\right)V(P,A',\Omega') \ + \ \lambda\left(\tfrac{D(P,A',\Omega')}{W(\Omega')}\right)\max_{P'} V\left(P',A',\Omega'\right)\ \right]\middle| A,\Omega\right\}$$

where $Q\left(\Omega,\Omega'\right)$ is the firm's stochastic discount factor and 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 time of production. In the expectation on the right, P represents the price $\widetilde{P}_{i,t+1}$ 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.

Making the rearrangement $(1-\lambda)V + \lambda \max V = V + \lambda(\max V - V)$ allows us to simplify the Bellman equation substantially on the right-hand side. We notice that the terms inside the expectation represent the value of continuing without adjustment, plus the flow of gains due to adjustment. The Bellman equation becomes:

Bellman equation in partial equilibrium, with aggregate shocks:

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

where

$$G(P,A',\Omega') \equiv \lambda\left(\frac{D(P,A',\Omega')}{W(\Omega')}\right)D(P,A',\Omega') \tag{4}$$

represents the expected gains due to adjustment.

This model represents a computational challenge, because the wage, the aggregate demand factor, the stochastic discount factor, and therefore also the value function all depend on the aggregate state Ω. In general equilibrium, at any time t, there will be many firms i facing different idiosyncratic shocks A_{it} and stuck at different prices P_{it}. The state of the economy will therefore include the entire distribution of prices and productivities. The reason for the popularity of the Calvo model is that even though firms have many different prices, up to a first-order approximation only the average price matters for equilibrium. Unfortunately, this property does not hold in general, and in the current context, we need to treat all equilibrium quantities explicitly as functions of the distribution of prices and productivity across the economy. To calculate equilibrium, we therefore need an algorithm that takes account of the distributional dynamics.

We attack this problem by implementing Reiter's (2008) solution method for dynamic general equilibrium models with heterogeneous agents and aggregate shocks. The first step in Reiter's algorithm is to calculate the steady state general equilibrium that obtains in the absence of aggregate shocks. Idiosyncratic shocks are still active, but are assumed to have converged to their ergodic distribution, so an aggregate steady state means that Ω, W, and ϑ are all constant.[7] We indicate the steady state by dropping Ω as an argument of the value function and other equilibrium objects, so the Bellman equation can be written as:

Bellman equation in partial equilibrium steady state:

$$V(P,A) = \left(P - \frac{W}{A}\right)\vartheta P^{-\epsilon} + R^{-1}E\left\{V(P,A') + G\left(P,A'\right)|A\right\} \tag{5}$$

Here R^{-1} is the steady state of the stochastic discount factor Q, and

$$G(P,A') \equiv \lambda\left(\frac{D(P,A')}{W}\right)D(P,A'), \qquad D(P,A') \equiv \max_{P'} V(P',A') - V(P,A') \tag{6}$$

This steady state Bellman equation is a standard dynamic programming problem, except for the timing of the max operator. A natural solution method is backwards induction on a two-dimensional grid $\Gamma \equiv \Gamma^P \times \Gamma^A$,

[7] More precisely, we allow these variables to have a nominal trend, and search for a steady state in real terms. Detrending is discussed in Sec. 3.5.

where Γ^P is a finite grid of possible values of P_i, and Γ^A is a grid of possible values of A_i. However, before we define notation that confines the dynamics to a grid, it is useful to describe the general equilibrium and detrend the model with respect to money growth, leaving all quantities in real terms.

2.2 Alternative sticky price frameworks

As we have stressed, our model can nest a number of alternative pricing frameworks, either by changing the parameters or the functional form of the adjustment probability λ, or by redefining the gains function G. All the following cases are then nested in a Bellman equation of form (3).

1. **Calvo pricing**: Suppose prices adjust each period with probability $\bar{\lambda}$, where $\bar{\lambda}$ is an exogenous constant. Then the Bellman equation is the same as (3), if we set $\lambda(D/W) \equiv \bar{\lambda}$. This is the special case of (2) in which $\xi = 0$.

2. **Fixed menu costs**: Suppose it costs κ units of labor to adjust prices in any given period, where κ is an exogenous constant called the "menu cost". Then the Bellman equation is given by (3), with $G = \lambda(D/W)D$ replaced by $G = \mathbf{1}\{D \geq \kappa W\}(D - \kappa W)$, where $\mathbf{1}\{D \geq \kappa W\}$ is an indicator function taking value one when $D \geq \kappa W$ and zero otherwise. In this case the function λ has form (2), with $\alpha = \kappa$ and $\xi = \infty$.

3. **Stochastic menu costs**: Suppose it costs κ units of labor to adjust prices in any given period, where κ is an *i.i.d.* random variable with *c.d.f.* $\lambda(\kappa)$. Then the Bellman equation is given by (3), with $G = \lambda(D/W)D$ replaced by $G = \lambda(D/W)[D - WE(\kappa|D > \kappa W)]$.

4. **Information-constrained pricing**: Woodford (2008) proposes a model in which managers decide on when to review a price based on imprecise awareness of current market conditions. His model implies the following adjustment probability function:

$$\lambda(D/W) \equiv \frac{\overline{\lambda}}{\overline{\lambda} + (1 - \overline{\lambda})\exp\left(\xi(D/W - \alpha)\right)} \tag{7}$$

where α is a fixed cost of purchasing information, and ξ^{-1} represents the marginal cost of information.

3 General equilibrium

We next embed this partial equilibrium decision framework in an otherwise standard New Keynesian general equilibrium, following the setup of 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

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

where u and v are increasing, concave functions, and x is an increasing, convex function. Utility is discounted by factor β per period. Consumption C_t is a Spence-Dixit-Stiglitz aggregate of differentiated products:

$$C_t = \left\{\int_0^1 C_{it}^{\frac{\epsilon-1}{\epsilon}} di\right\}^{\frac{\epsilon}{\epsilon-1}} \tag{8}$$

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

$$\int_0^1 P_{it}C_{it}di + M_t + R_t^{-1}B_t = W_tN_t + M_{t-1} + T_t + B_{t-1} + \Pi_t \tag{9}$$

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.

Households choose $\{C_{it}, N_t, B_t, M_t\}_{t=0}^{\infty}$ to maximize expected discounted utility, subject to the budget constraint (9). Optimal allocation of consumption across differentiated goods implies

$$C_{it} = (P_t/P_{it})^{\epsilon} C_t \tag{10}$$

where P_t is the following price index:

$$P_t \equiv \left\{ \int_0^1 {P_{it}}^{1-\epsilon} di \right\}^{\frac{1}{1-\epsilon}} \tag{11}$$

This means we can rewrite nominal spending as $P_tC_t = \int_0^1 P_{it}C_{it}di$.[8] Optimal labor supply and money holdings imply the first-order conditions

$$x'(N_t) = u'(C_t)W_t/P_t \tag{12}$$

$$v'\left(\frac{M_t}{P_t}\right) = u'(C_t)(1 - R_t^{-1}) \tag{13}$$

and the Euler equation is

$$R_t^{-1} = \beta E_t \left(\frac{P_t u'(C_{t+1})}{P_{t+1}u'(C_t)} \right) \tag{14}$$

3.2 Monetary policy

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

$$M_t = \mu_t M_{t-1} \tag{15}$$

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

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

Here $0 \leq \phi_z < 1$ and $\epsilon_t^z \sim i.i.d.N(0, \sigma_z^2)$ is a money growth shock. Thus the money supply trends upward by approximately factor $\mu \geq 1$ per period on average.

Seigniorage revenues are paid to the household as a lump sum transfer, and the government budget is balanced each period. Therefore the government's budget constraint is

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

3.3 Aggregate consistency

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

$$Y_{it} = A_{it}N_{it} = C_{it} = P_t^{\epsilon} C_t P_{it}^{-\epsilon} \tag{17}$$

Also, total labor supply must equal total labor demand:

$$N_t = \int_0^1 \frac{C_{it}}{A_{it}} di = P_t^{\epsilon} C_t \int_0^1 P_{it}^{-\epsilon} A_{it}^{-1} di \equiv \Delta_t C_t \tag{18}$$

Labor market clearing condition (18) also defines a weighted measure of price dispersion, $\Delta_t \equiv P_t^{\epsilon} \int_0^1 P_{it}^{-\epsilon} A_{it}^{-1} di$, which generalizes the dispersion measure in Yun (2005) to allow for heterogeneous productivity. As in Yun's paper, an increase in Δ_t decreases the consumption goods produced per unit of labor, effectively acting like a negative shock to aggregate productivity.[9]

[8] One of the preceding equations is superfluous: (10) plus (8) implies (11), and likewise (10) plus (11) implies (8).

[9] Dorich (2007) also introduces a heterogeneity-adjusted dispersion measure that acts like an aggregate productivity shock.

At this point, we have spelled out all equilibrium conditions, so we are ready to consider how to define the aggregate state variable Ω_t. To do so, it helps to distinguish idiosyncratic states (prices and productivities) before price adjustment from those after adjustment. Thus we again use the notation $\widetilde{P}_{it}$ to refer to firm i's price at the beginning of period t, prior to adjustment, and P_{it} to indicate the price at which it actually produces. Likewise, we indicate the distribution of beginning-of-period prices and productivities as $\widetilde{\Phi}_t(\widetilde{P}_{it}, A_{it})$, writing the distribution of prices and productivities at the time of production as $\Phi_t(P_{it}, A_{it})$.

The aggregate state of the economy at time t depends, among other things, on the money supply M_t. Since the growth rate of money is AR(1) over time, the latest deviation in growth rates, z_t, is a state variable too. For any given firm i, the individual state variables that are relevant for time t decisions are the beginning-of-period price and productivity $(\widetilde{P}_{it}, A_{it})$. Thus one possible definition of the time t aggregate state is $\Omega_t \equiv (M_t, z_t, \widetilde{\Phi}_t)$: M_t and z_t are a sufficient statistic for the Markov process driving the money supply, and $\widetilde{\Phi}_t$ is the beginning-of-t distribution of idiosyncratic states.

However, this is not the only possible definition of the aggregate state. Nominal prices are unadjusted between the time of production in $t-1$ and the beginning of t, so the distributions Φ_{t-1} and $\widetilde{\Phi}_t$ differ only due to the exogenous Markov process that drives idiosyncratic productivity. In other words, Φ_{t-1} is known if and only if $\widetilde{\Phi}_t$ is known. Therefore an alternative representation of the time t aggregate state is

$$\Omega_t \equiv (M_t, z_t, \Phi_{t-1}) \tag{19}$$

This is the representation of the aggregate state that we will actually use, because it turns out to be algebraically convenient. As we will see in Section 5.1, representation (19) allows us to define the full dynamic equilibrium equation system in a form that substitutes out many variables and equations.

3.4 The firm's problem in general equilibrium

The setup of sections 3.1-3.3 holds regardless of how firms set prices. That is, regardless of the price-setting mechanism, C_t, N_t, P_t, W_t, R_t, C_{it}, P_{it}, and M_t must obey equations (8) - (18). In particular, to make the firm's problem (3) consistent with the goods market clearing conditions (17), the aggregate demand shift factor must be

$$\vartheta(\Omega) = C(\Omega)P(\Omega)^{\epsilon} \tag{20}$$

Also, we assume that the representative household owns the firms, so the stochastic discount factor in the firm's problem must be consistent with the household's Euler equation (14). Therefore the appropriate stochastic discount factor is

$$Q(\Omega, \Omega') = \beta \frac{P(\Omega)u'(C(\Omega'))}{P(\Omega')u'(C(\Omega))} \tag{21}$$

To write the firm's problem in general equilibrium, we simply plug (20) and (21) into the firm's problem (3). Showing time subscripts for transparency, the value of producing with price P_{it} and productivity A_{it} is
Bellman equation in general equilibrium:

$$\begin{aligned} V(P_{it}, A_{it}, \Omega_t) = &\left(P_{it} - \frac{W(\Omega_t)}{A_{it}}\right) C(\Omega_t)P(\Omega_t)^{\epsilon} P_{it}^{-\epsilon} + \\ &+ \beta E_t \left\{ \frac{P(\Omega_t)u'(C(\Omega_{t+1}))}{P(\Omega_{t+1})u'(C(\Omega_t))} \left[V(P_{it}, A_{i,t+1}, \Omega_{t+1}) + G(P_{it}, A_{i,t+1}, \Omega_{t+1}) \right] \middle| A_{it}, \Omega_t \right\} \end{aligned} \tag{22}$$

where $G(P_{it}, A_{i,t+1}, \Omega_{t+1})$ has the form described in (4) or one of the forms associated with the alternative sticky price frameworks mentioned in Section 2.2.

3.5 Detrending

So far we have written the value function and all prices in nominal terms, but it is natural to assume that we can rewrite the model in real terms. Thus, suppose we deflate all prices by the nominal money stock, defining $p_t \equiv P_t/M_t$, $p_{it} \equiv P_{it}/M_t$, and $w_t \equiv W_t/M_t$. Given the nominal distribution $\Phi_t(P_i, A_i)$ and the money stock M_t, let us denote by $\Psi_t(p_i, A_i)$ the distribution over real production prices $p_{it} \equiv P_{it}/M_t$. Likewise, let $\widetilde{\Psi}_t(\widetilde{p}_i, A_i)$ be the distribution of real beginning-of-period prices $\widetilde{p}_{it} \equiv \widetilde{P}_{it}/M_t$, in analogy to the beginning-of-period distribution of nominal prices $\widetilde{\Phi}_t(\widetilde{P}_i, A_i)$. If the model can be rewritten in real terms, then the level of the money supply, M_t, must be irrelevant for determining real quantities. Therefore, to describe the real equilibrium, it suffices to condition on the real state variable $\Xi_t \equiv (z_t, \Psi_{t-1})$, instead of the full nominal state $\Omega_t \equiv (M_t, z_t, \Phi_{t-1})$.[10] The "real" value function v should likewise be the nominal value function, divided by the current money stock, and should be written as a function of real variables. That is,

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

Deflating in this way, the Bellman equation can be rewritten as follows (see the appendix for details).
Detrended Bellman equation, general equilibrium:

$$\begin{aligned} v(p_{it}, A_{it}, \Xi_t) \;=\; & \left(p_{it} - \frac{w(\Xi_t)}{A_{it}}\right)\left(\frac{p_{it}}{p(\Xi_t)}\right)^{-\epsilon} C(\Xi_t) + \\ & \beta E_t \left\{ \frac{p(\Xi_t)u'(C(\Xi_{t+1}))}{p(\Xi_{t+1})u'(C(\Xi_t))} \left[v\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1}\right) + g\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1}\right)\right] \middle| A_{it}, \Xi_t \right\} \end{aligned}$$

where

$$g\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1}\right) \equiv \lambda\left(\frac{d\left(p_{it}/\mu_{t+1}, A_{i,t+1}, \Xi_{t+1}\right)}{w(\Xi_{t+1})}\right) d\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1}\right)$$

$$d\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1}\right) \equiv \max_{p'} v(p', A_{i,t+1}, \Xi_{t+1}) - v\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1}\right)$$

Let $p^*(A_{i,t+1}, \Xi_{t+1})$ denote the optimal choice in the maximization problem above. Taking into account the fact that the firm starts period $t+1$ with the eroded price $\widetilde{p}_{i,t+1} \equiv p_{it}/\mu_{t+1}$, the real price process is

$$p_{i,t+1} = \begin{cases} p^*(A_{i,t+1}, \Xi_{t+1}) & \text{with probability } \lambda\left(\frac{d\left(p_{it}/\mu_{t+1}, A_{i,t+1}, \Xi_{t+1}\right)}{w(\Xi_{t+1})}\right) \\ \frac{p_{it}}{\mu_{t+1}} & \text{with probability } 1 - \lambda\left(\frac{d\left(p_{it}/\mu_{t+1}, A_{i,t+1}, \Xi_{t+1}\right)}{w(\Xi_{t+1})}\right) \end{cases}.$$

In other words, when the firm's nominal price remains unadjusted at time $t+1$, its real price is deflated by factor μ_{t+1}.

4 Computing general equilibrium: steady state

4.1 Discrete numerical model

For computational purposes, we next approximate the economy by assuming that individual states, in real terms, always lie on a finite grid.[11] This allows us to solve the firm's problem numerically by backwards induction. This procedure is entirely standard, but we will spell out the details, both to see how it nests into the general equilibrium and in order to clarify our calculations later when we study the effects of aggregate shocks.

Thus, consider the two-dimensional grid $\Gamma \equiv \Gamma^p \times \Gamma^a$, where $\Gamma^p \equiv \{p^1, p^2, ...p^{\#^p}\}$ is a logarithmically-spaced grid of possible values of p_i, and $\Gamma^a \equiv \{a^1, a^2, ...a^{\#^a}\}$ is a logarithmically-spaced grid of possible values of A_i. It is now natural to treat the distributions as matrices $\widetilde{\mathbf{\Psi}}$ and $\mathbf{\Psi}$ of size $\#^p \times \#^a$, in which the row j, column k elements, called $\widetilde{\Psi}^{jk}$ and Ψ^{jk}, represent the fraction of firms in state (p^j, a^k) at the beginning of the period and at the time of production, respectively. From here on, we frequently use bold face to identify matrices, and superscripts to identify notation related to grids.

[10] The deflation factor for real prices between periods $t-1$ and t is known if z_t is known. Therefore if z_t is known, knowing the real distribution at the time of production in period $t-1$, Ψ_{t-1}, is equivalent to knowing the real distribution at the beginning of t, $\widetilde{\Psi}_t$. So the real aggregate state at time t can be defined as $(z_t, \widetilde{\Psi}_t)$ or (z_t, Ψ_{t-1}), but the latter turns out to be algebraically convenient, as we will see in Section 5.1.

[11] Reiter (2008) calls his computational method "projection and perturbation". In this subsection we are implementing the "projection" step, by projecting our infinite-dimensional economy onto an approximately equivalent finite-dimensional economy. The "perturbation" step is the linearization of the aggregate dynamics, as discussed in subsection 5.1.

Likewise, we can write the value function as a $\#^p \times \#^a$ matrix $\mathbf{V}$ of values $v^{jk} \equiv v(p^j, a^k)$ associated with the prices and productivities $\left(p^j, a^k\right) \in \Gamma$. We can construct splines to evaluate the value function at points $p \notin \Gamma^p$ off the price grid, when necessary. In particular, we define the policy function

$$p^*(A) \equiv \arg\max_p v(p, A) \tag{23}$$

without requiring that it be chosen from the grid Γ^p, because our solution method requires policies to vary smoothly with aggregate conditions.[12] We will write the policies at the productivity grid points $a^k \in \Gamma^a$ as a row vector $\mathbf{p}^* \equiv \left\{p^{*1}...p^{*\#^a}\right\} \equiv \left\{p^*(a^1)...p^*(a^{\#^a})\right\}$. We also define several other $\#^p \times \#^a$ matrices: the adjustment values $\mathbf{D}$, the probabilities $\mathbf{\Lambda}$, and the expected gains $\mathbf{G}$, with (j,k) elements given by[13]

$$d^{jk} \equiv \max_p v(p, a^k) - v^{jk} \tag{24}$$

$$\lambda^{jk} \equiv \lambda\left(d^{jk}/w\right) \tag{25}$$

$$g^{jk} \equiv \lambda^{jk} d^{jk} \tag{26}$$

We can now write the discrete Bellman equation and the discrete distributional dynamics in a precise way. The dynamics involve three transitions. First, suppose a firm has beginning-of-t price $\widetilde{p}_{it} = p^j \in \Gamma^p$ and productivity $A_{it} = a^k \in \Gamma^a$. This firm will adjust its production price to $p_{it} = p^{*k}$ with probability λ^{jk}, or will leave it unchanged ($p_{it} = \widetilde{p}_{it} = p^j$) with probability $1 - \lambda^{jk}$. If adjustment occurs, we maintain our grid-based approximation by rounding p^{*k} up or down stochastically to the nearest grid points. To be precise, suppose p^{*k} lies between grid points $l-1$ and l, that is, $p^{l-1} < p^{*k} \leq p^l$. Then we round p^{*k} up to p^l with probability $(p^{*k} - p^{l-1})/(p^l - p^{l-1})$, and down to p^{l-1} with probability $(p^l - p^{*k})/(p^l - p^{l-1})$. This transition can be summarized in matrix notation. Let $\mathbf{E}_{IJ}$ be an $I \times J$ matrix of ones. Assume Γ^p is chosen wide enough so that $p^1 < p^{*k} < p^{\#^p}$ for all $k \in \{1, 2, ...\#^a\}$. Then for each k, define $l(k)$ so that $p^{l(k)} = \min\{p \in \Gamma^p : p \geq p^{*k}\}$. This allows us to define a $\#^p \times \#^a$ matrix $\mathbf{P}$ that rounds the policy function stochastically up or down to the nearest grid points:

$$\mathbf{P} \equiv \begin{cases} \frac{p^{l(k)} - p^{*k}}{p^{l(k)} - p^{l(k)-1}} & \text{in column } k, \text{ row } l(k)-1 \\ \frac{p^{*k} - p^{l(k)-1}}{p^{l(k)} - p^{l(k)-1}} & \text{in column } k, \text{ row } l(k) \\ 0 & \text{elsewhere} \end{cases} \tag{27}$$

Then we can calculate distribution $\mathbf{\Psi}_t$ from $\widetilde{\mathbf{\Psi}}_t$ as follows:

$$\mathbf{\Psi}_t = \left(\mathbf{E}_{\#^p \#^a} - \mathbf{\Lambda}\right) .* \widetilde{\mathbf{\Psi}}_t + \mathbf{P} .* \left(\mathbf{E}_{\#^p \#^p} * \left(\mathbf{\Lambda} .* \widetilde{\mathbf{\Psi}}_t\right)\right) \tag{28}$$

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

The second step in the distributional dynamics is to adjust real prices to take account of steady state money growth. Ignoring grids, the time t price p_{it} is deflated to $\widetilde{p}_{i,t+1} \equiv p_{it}/\mu$ at the beginning of $t+1$. To keep prices on the grid, we define a $\#^p \times \#^p$ Markov matrix $\mathbf{R}$ in which the row m, column l element is

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

When $\widetilde{p}_{i,t+1}$ falls between two grid points, matrix $\mathbf{R}$ must round up or down stochastically. When $\widetilde{p}_{i,t+1}$ moves down past the least element of the grid, matrix $\mathbf{R}$ must round up to keep prices on the grid.[14] Therefore we construct $\mathbf{R}$ according to

$$R^{ml} = prob(\widetilde{p}_{i,t+1} = p^m | p_{it} = p^l) = \begin{cases} 1 & \text{if } \mu^{-1} p^l \leq p^1 = p^m \\ \frac{\mu^{-1} p^l - p^{m-1}}{p^m - p^{m-1}} & \text{if } p^1 < p^m = \min\{p \in \Gamma^p : p \geq \mu^{-1} p^l\} \\ \frac{p^{m+1} - \mu^{-1} p^l}{p^{m+1} - p^m} & \text{if } p^1 \leq p^m = \max\{p \in \Gamma^p : p < \mu^{-1} p^l\} \\ 0 & \text{otherwise} \end{cases} \tag{29}$$

[12] Ensuring differentiability of all equilibrium objects is discussed in Appendix B.

[13] The max in (24), like the arg max in (23), ignores the grid Γ^p so that d^{jk} varies smoothly in response to any shift in the value function.

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

The third and final step in the distributional dynamics is to take into account the Markov matrix $\mathbf{S}$ that governs the idiosyncratic productivity shocks A_i. The row m, column k element of $\mathbf{S}$ is the exogenous probability

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

Combining the second and third steps, we can calculate the beginning-of-period distribution $\widetilde{\mathbf{\Psi}}_{t+1}$ at $t+1$ as a function of the time t distribution of production prices $\mathbf{\Psi}_t$:

$$\widetilde{\mathbf{\Psi}}_{t+1} = \mathbf{R} * \mathbf{\Psi}_t * \mathbf{S}' \tag{30}$$

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

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

$$u^{jk} = \left(p^j - \frac{w}{a^k}\right) C \left(\frac{p_j}{p}\right)^{-\epsilon} \tag{31}$$

for $\left(p^j, a^k\right) \in \Gamma$. Then the Bellman equation is
Steady state general equilibrium Bellman equation, matrix version:

$$\mathbf{V} = \mathbf{U} + \beta \mathbf{R}' * (\mathbf{V} + \mathbf{G}) * \mathbf{S} \tag{32}$$

Since the Bellman equation iterates backwards in time, it involves probability transitions represented by $\mathbf{R}'$ and $\mathbf{S}$, whereas the distributional dynamics iterate forward in time and therefore contain $\mathbf{R}$ and $\mathbf{S}'$.

Finally, in addition to these matrix equations, there are four scalar equations relating to first-order conditions and aggregate consistency conditions:

$$u'(C) = \frac{px'(N)}{w} \tag{33}$$

$$1 - \frac{v'(1/p)}{u'(C)} = \beta/\mu \tag{34}$$

$$N = \sum_{j=1}^{\#^p} \sum_{k=1}^{\#^a} \Psi^{jk} \left(\frac{p^j}{p}\right)^{-\epsilon} \frac{C}{a^k} \tag{35}$$

$$p^{1-\epsilon} = \sum_{j=1}^{\#^p} \sum_{k=1}^{\#^a} \Psi^{jk} \left(p^j\right)^{1-\epsilon} \tag{36}$$

Equations (23)-(36) fully describe the steady state general equilibrium. The unknowns are the matrices $\mathbf{V}$, $\mathbf{D}$, $\mathbf{\Lambda}$, $\mathbf{G}$, $\mathbf{P}$, $\mathbf{R}$, $\mathbf{U}$, $\mathbf{\Psi}$, and $\widetilde{\mathbf{\Psi}}$; the vector $\mathbf{p}^*$; and the scalars w, p, N, and C.

While this seems like a huge system of equations, it is easy to solve because it reduces to a small scalar fixed-point problem. We solve it under linear labor disutility, $x(N) = \chi N$. In this case, guessing p permits us to calculate w and C analytically from equations (33) and (34). We can then construct matrix $\mathbf{U}$ from (31), so we are ready to solve the Bellman equation (32) to find $\mathbf{V}$ and $\mathbf{P}$. We then find the steady state price distributions $\mathbf{\Psi}$ and $\widetilde{\mathbf{\Psi}}$ from (28) and (30). Knowing distribution $\mathbf{\Psi}$, we can calculate the price level p from (36). Thus, finding a fixed point in p allows us to construct the entire steady state equilibrium.[15]

[15] If labor disutility is nonlinear, the problem is only slightly harder: it reduces to a two-dimensional fixed-point problem in p and N.

4.2 Results: steady state

This steady state model can be calibrated by comparing its predictions to cross-sectional data on price changes, like those reported in Klenow and Kryvstov (2005), Midrigan (2008), and Nakamura and Steinsson (2007). In a companion paper, Costain and Nakov (2008), we report detailed results from a variety of specifications, and compare the model's behavior under low and high steady-state inflation rates. Here we simply briefly discuss our preferred estimate from that paper, which minimizes an equally-weighted sum of two terms: the absolute difference between the mean adjustment frequency in the model and that in the data, plus the distance between the histogram of price changes in the model and that in the data. We will simulate our model at monthly frequency, for consistency with the results reported in the empirical literature. Also, since these papers all attempt to remove price changes attributable to temporary "sales", our simulation results should be interpreted as a model of "regular" price changes unrelated to sales.

We take our utility parameterization from Golosov and Lucas (2007). Hence, we set the discount factor to $\beta = 1.04^{-1}$ per year; 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 utility of real money holdings is logarithmic, $v(m) = \nu \log(m)$, with $\nu = 1$. The elasticity of substitution in the consumption aggregator is $\epsilon = 7$. The steady state growth rate of money is set to zero, consistent with the zero average price change found in the AC Nielsen dataset used to estimate the model.

Given these utility parameters, we estimate the idiosyncratic productivity shock process and the adjustment process to match data on the distribution of regular price changes. We assume productivity is AR(1) in logs:[16]

$$\log A_{it} = \rho \log A_{it-1} + \varepsilon^a_{it}$$

where ε^a_{it} is a mean-zero, normal, *iid* shock. There are two free parameters of the productivity process: ρ and σ^2_ε, the variance of ε^a_{it}, while the adjustment process has three free parameters, $\overline{\lambda}$, α, and ξ, in the function class we have imposed,

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

Table 1 reports our preferred estimate from Costain and Nakov (2008), (labelled SDSP, for 'state-dependent sticky prices'), together with evidence from four empirical studies. It also reports a Calvo (1983) version of the model, a fixed menu cost version (as in Golosov and Lucas 2007, labelled 'MC'), and a version based on Woodford's (2008) adjustment function (7). Each of them is estimated to best fit the size distribution of price adjustments in the AC Nielsen data, and to match Nakamura and Steinsson's (2007) measure of the median frequency of price adjustments (which is lower, but presumably more robust, than measures based on means). All versions of the model match the target adjustment frequency of 10% per month almost exactly. Our preferred estimate 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 11.8%, 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.

Figure 1 graphs a variety of objects that characterize the stationary equilibrium under the SDSP specification.[17] In the first plot we see the value function, as a function of prices and marginal cost (one over productivity); the lowest value occurs when the highest marginal cost is paired with the lowest price. The fourth and sixth plots show the distributions $\widetilde{\mathbf{\Psi}}$ at the beginning of the period and $\mathbf{\Psi}$ at the time of production. The production distribution $\mathbf{\Psi}$ looks rather like a sail-backed dinosaur: the "sail" represents the mass of firms that have adjusted to the optimal price conditional on current productivity. At the beginning of the next period, this mass gets spread out by the productivity shock process, resulting in the smooth distribution $\widetilde{\mathbf{\Psi}}$ seen in the fourth graph. Graphing the policy function in the eighth plot shows that the firm sets prices closer to the mean than would be the case under flexible prices, in anticipation of mean reversion of the technology process.

[16] Our numerical method requires us to treat A as a discrete variable, so we use Tauchen's method (Mertens, 2006) to approximate this AR(1) process on the discrete grid Γ^a. We use a grid of 201 points representing five standard deviations of A. The price grid Γ^p also contains 201 logarithmically-spaced points, which results in price steps of 0.74%.

[17] As far as we can tell, steady state equilibrium is unique (both under our baseline specification, and for the Calvo and fixed menu cost cases). Large changes in the initial guess for p and/or the terminal value from which backward induction begins do not change the fixed point to which our steady-state solution algorithm converges.

The last graph shows the distribution of nominal price adjustments, which is mildly bimodal around zero, and resembles quite closely the AC Nielsen data of supermarket price changes reported in Midrigan (2008) and reproduced here in shaded bars. 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.

The seventh panel of Figure 1 shows the probability of price adjustment, as a function of the value of adjustment. The probability of adjustment initially rises very quickly, but it then quickly levels off, remaining below 20% when the loss from failing to adjust is 1% of the value of the firm (which is the largest value seen in the graph). On one hand, this λ function points to a low degree of state dependence: the probability of adjustment in any given month is still well below 0.5 even if several percentage points of the value of the firm are at stake. On the other hand, few firms actually come to suffer large losses, since the range shown in the seventh panel of Fig. 1 includes almost all the variation actually observed in the simulation; adjustment normally occurs before leaving this range. Thus the median loss is just 0.07% of median firm value, and the mean loss (which is larger since the loss distribution is highly skewed) is 0.27% of median firm value.

Finally, it is helpful to consider the computational implications of the relatively large but infrequent price adjustments seen in the data. On the 201 by 201 grid with a median absolute price change of around 9%, a typical price movement by firms in our baseline SDSP simulation is a jump of about 12 steps in the price grid Γ^p. Clearly then, at any point in time most firms lie several steps away from their optimal prices; the table shows that the typical deviation from the optimal price ranges from 4.4% (in terms of the median) to 7.7% (on average), depending on the model. This suggests that constraining price adjustment to a finite grid is relatively unimportant both for price dynamics and for welfare analysis. We confirm this fact in Table 1 by recomputing the model (under the SDSP calibration) on a much coarser grid, with only 25 possible productivities (spanning ± 2.5 standard deviations instead of ± 5 standard deviations) and only 31 possible prices. Thus, in the coarser grid, each price step represents a 2.5% price change, instead of the 0.7% in the previous calculation.

This dramatic coarsening of the grid has only minor consequences for the performance of the model. The statistic that changes most is the fraction of small price changes, which decreases from 25.2% to 24.8%. The other statistics are barely altered, including the welfare losses caused by price stickiness. Thus, computing the dynamics on a finite grid seems unimportant for the results, even when the grid is quite coarse. This is very helpful for our purposes, because it suggests that the more numerically challenging problem of characterizing the distributional dynamics can also be studied on a coarse grid.

5 Computing general equilibrium: dynamics

To characterize our model's general equilibrium dynamics in the presence of both idiosyncratic and aggregate shocks, we implement the algorithm of Reiter (2008). Reiter's method recognizes that the large system of nonlinear equations we solved to calculate the general equilibrium steady state can also be interpreted as a system of nonlinear first-order autonomous difference equations describing the dynamics of a grid-based approximation to general equilibrium away from steady state. In the absence of strong strategic complementarities or an inappropriate Taylor rule that might give rise to indeterminacy, such an equation system can be solved by perfectly standard linear simulation techniques. We will solve for the saddle-path stable solution of our linearized model using the QZ decomposition, following Klein (2000).

The crucial thing to notice about Reiter's method is that it combines linearity and nonlinearity in a way appropriate for the model at hand. In our model, idiosyncratic shocks are likely to be larger and more economically important for individual firms' decisions than aggregate shocks. This is true in many macroeconomic contexts (e.g. precautionary saving) and in particular Klenow and Kryvstov (2008), Golosov and Lucas (2007), and Midrigan (2008) argue that firms' pricing decisions appear to be driven primarily by idiosyncratic shocks. Therefore, to deal with large idiosyncratic shocks, we treat functions of idiosyncratic states in a fully nonlinear way, by calculating them on a grid. As we emphasized above, this grid-based solution can be regarded as a large system of nonlinear equations, with equations specific to each of the grid points. By linearizing each of these equations with respect to the aggregate dynamics, we recognize that aggregate changes are unlikely to affect individual value functions in a strongly nonlinear way. That is, we are implicitly assuming that both money supply shocks μ and changes in the distributions $\mathbf{\Psi}$ and $\widetilde{\mathbf{\Psi}}$ have sufficiently smooth effects on individual values that a linear treatment of these effects is sufficient. On the other hand, we need *not* start from any assumption of approximate aggregation like that required for the method of Krusell and Smith (1998).

Thus, we will write the general equilibrium dynamics as a system of difference equations. For parsimonious notation in this context, we indicate dependence on the aggregate state by time subscripts, instead of by writing endogenous variables as functions of Ξ_t. We will see that the difference equation system is a straightforward generalization of the steady state equations from the previous section. First, the time t money growth process is $\mu_t = \mu \exp(z_t)$, where

$$z_t = \phi_z z_{t-1} + \epsilon_t^z \tag{37}$$

where ϵ_t^z is an *iid* normal shock with mean zero and standard deviation σ_z.

Second, the firms' Bellman equation can be written as a $\#^p \times \#^a$ matrix system of equations for each $(p^j, a^k) \in \Gamma$. Let $\mathbf{U}_t$ be the matrix of current profits, so that the (j,k) element of $\mathbf{U}_t$ is

$$u_t^{jk} \equiv \left(p^j - \frac{w_t}{a^k}\right) C_t \left(\frac{p^j}{p_t}\right)^{-\epsilon} \equiv \left(p^j - \frac{w(\Xi_t)}{a^k}\right) C(\Xi_t) \left(\frac{p^j}{p(\Xi_t)}\right)^{-\epsilon} \tag{38}$$

Write the value function as a matrix $\mathbf{V}_t$, with (j,k) element equal to $v_t^{jk} \equiv v_t(p^j, a^k) \equiv v(p^j, a^k, \Xi_t)$ for $(p^j, a^k) \in \Gamma$. We can write the Bellman equation as

Dynamic general equilibrium Bellman equation, matrix version:

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

All quantities in the Bellman equation are analogous to corresponding quantities in the steady state equilibrium. The matrix $\mathbf{G}_{t+1}$ is defined by

$$\mathbf{G}_{t+1} \equiv \mathbf{\Lambda}_{t+1} \; .* \, \mathbf{D}_{t+1} \tag{40}$$

where the (l,m) element of $\mathbf{D}_{t+1}$ is

$$d_{t+1}^{lm} \equiv d_{t+1}(p^l, a^m) \equiv \max_{p'} v_{t+1}(p', a^m) - v_{t+1}(p^l, a^m) \tag{41}$$

and $\mathbf{\Lambda}_{t+1}$ is the matrix with (l,m) element $\lambda_{t+1}^{lm} \equiv \lambda\left(d_{t+1}^{lm}/w_{t+1}\right)$.

The expectation E_t in the Bellman equation refers only to the effects of the time $t+1$ money shock μ_{t+1}, because the shocks and dynamics of the idiosyncratic state $(p^j, a^k) \in \Gamma$ are completely described by the matrices $\mathbf{R}'_{t+1}$ and $\mathbf{S}$. Note that $\mathbf{S}$ has no time subscript, and is exactly the same matrix described in the previous section. The Markov matrix $\mathbf{R}_{t+1}$ differs from the steady state matrix $\mathbf{R}$ only because in the fully dynamic equilibrium we must detrend by the realized money shock μ_{t+1} instead of trend money growth μ. The row n, column l element of $\mathbf{R}_{t+1}$, which we will call R_{t+1}^{nl}, is

$$R_{t+1}^{nl} = prob(\widetilde{p}_{i,t+1} = p^n | p_{it} = p^l, \mu_{t+1}) = \begin{cases} 1 & \text{if } p^l/\mu_{t+1} \leq p^1 = p^n \\ \frac{p^l/\mu_{t+1} - p^{n-1}}{p^n - p^{n-1}} & \text{if } p^1 < p^n = \min\{p \in \Gamma^p : p \geq p^l/\mu_{t+1}\} \\ \frac{p^{n+1} - p^l/\mu_{t+1}}{p^{n+1} - p^n} & \text{if } p^1 \leq p^n = \max\{p \in \Gamma^p : p < p^l/\mu_{t+1}\} \\ 1 & \text{if } p^l/\mu_{t+1} > p^{\#^p} = p^n \\ 0 & \text{otherwise} \end{cases}$$

As for the distributional dynamics, the two steps are analogous to the steady state case:

$$\mathbf{\Psi}_t = (\mathbf{E}_{\#^p\#^a} - \mathbf{\Lambda}_t) \; .* \, \widetilde{\mathbf{\Psi}}_t + \mathbf{P}_t \; .* \, (\mathbf{E}_{\#^p\#^p} * (\mathbf{\Lambda}_t \; .* \, \widetilde{\mathbf{\Psi}}_t)) \tag{42}$$

$$\widetilde{\mathbf{\Psi}}_{t+1} = \mathbf{R}_{t+1} * \mathbf{\Psi}_t * \mathbf{S}' \tag{43}$$

Matrix $\mathbf{P}_t$ is constructed from the policy function

$$p_t^{*k} \equiv p_t^*(a^k) \equiv \arg\max_p v(p, a^k, \Xi_t) \tag{44}$$

in the same way as in the steady state.[18] If $p^{l(k)}$ is the first price grid point greater than or equal to p_t^{*k}, then $\mathbf{P}_t$ takes value $\left(\frac{p_t^{*k} - p^{l(k)-1}}{p^{l(k)} - p^{l(k)-1}}\right)$ in row $l(k)$, column k; and value $\left(\frac{p^{l(k)} - p_t^{*k}}{p^{l(k)} - p^{l(k)-1}}\right)$ in row $l(k)-1$, column k; and is zero elsewhere.

[18] As in the steady state calculation, the max in (41) and (44) is not restricted to the grid Γ^p, allowing d_t^{jk} and p_t^{*k} to vary smoothly with changes in aggregate conditions Ξ_t, which is necessary for our linearized solution. In the limiting case $\xi = \infty$, we must also take care to ensure that λ_t^{jk} varies smoothly. Both issues are discussed in Appendix B.

Finally, the remaining equations that must be satisfied by the dynamic general equilibrium are

$$\begin{aligned} x'(N_t) &= \frac{w_t}{p_t} u'(C_t) && (45) \\ 1 - \frac{v'(1/p_t)}{u'(C_t)} &= \beta E_t \left(\frac{p_t u'(C_{t+1})}{\mu_{t+1} p_{t+1} u'(C_t)} \right) && (46) \\ N_t &= \sum_{j=1}^{\#^p} \sum_{k=1}^{\#^a} \Psi_t^{jk} \left(\frac{p^j}{p_t} \right)^{-\epsilon} \frac{C_t}{a^k} && (47) \\ p_t^{1-\epsilon} &= \sum_{j=1}^{\#^p} \sum_{k=1}^{\#^a} \Psi_t^{jk} \left(p^j \right)^{1-\epsilon} && (48) \end{aligned}$$

5.1 Linearization

We are now ready to calculate the general equilibrium dynamics by linearization. To do so, we eliminate as many variables from the equation system as we can. For additional simplicity, we assume linear labor disutility, $x(N) = \chi N$. Thus the first-order condition for labor reduces to $\chi p_t = w_t u'(C_t)$, so we don't actually need to solve for N_t in order to calculate the rest of the equilibrium.[19]

Under the linear labor disutility assumption, we can summarize the whole general equilibrium system in terms of the exogenous shock process z_t, the endogenous 'jump' variables $\mathbf{V}_t$, C_t, and p_t, and the lagged distribution of idiosyncratic states $\mathbf{\Psi}_{t-1}$, which is the endogenous component of the time t aggregate state. The full system of equations reduces to

$$\mathbf{V}_t = \mathbf{U}_t + \beta E_t \left\{ \frac{p_t u'(C_{t+1})}{p_{t+1} u'(C_t)} \mathbf{R}'_{t+1} * (\mathbf{V}_{t+1} + \mathbf{G}_{t+1}) * \mathbf{S} \right\} \tag{49}$$

$$1 - \frac{v'(1/p_t)}{u'(C_t)} = \beta E_t \left(\mu_{t+1} \frac{p_t u'(C_{t+1})}{p_{t+1} u'(C_t)} \right) \tag{50}$$

$$p_t^{1-\epsilon} = \sum_{j=1}^{\#^p} \sum_{k=1}^{\#^a} \Psi_t^{jk} \left(p^j \right)^{1-\epsilon} \tag{51}$$

$$\mathbf{\Psi}_t = (\mathbf{E}_{\#^p\#^a} - \mathbf{\Lambda}_t) \ .* \widetilde{\Psi}_t + \mathbf{P}_t \ .* (\mathbf{E}_{\#^p\#^p} * (\mathbf{\Lambda}_t \ .* \widetilde{\mathbf{\Psi}}_t)) \tag{52}$$

$$z_{t+1} = \phi_z z_t + \epsilon^z_{t+1} \tag{53}$$

Counting element by element the matrix equations (49) and (52) each contain $\#^p\#^a$ scalar equations, and therefore the whole system contains $2\#^p\#^a + 3$ nonlinear difference equations. If we now collapse all the endogenous variables into a single vector

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

then the whole set of expectational difference equations (49)-(53) governing the dynamic equilibrium becomes a first-order system of the following form:

$$E_t \mathcal{F} \left(\overrightarrow{X}_{t+1}, \overrightarrow{X}_t, z_{t+1}, z_t \right) = 0 \tag{54}$$

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

Together, the endogenous vector $\overrightarrow{X}_t$ and the shock process z_t amount to a list of $2\#^p\#^a + 3$ variables. To see that these are in fact the only variables we need, because all others can be substituted out. Given z_t and z_{t+1}we can construct μ_t, and μ_{t+1}, and thus $\mathbf{R}_t$ and $\mathbf{R}_{t+1}$. Given $\mathbf{R}_t$, we can construct $\widetilde{\mathbf{\Psi}}_t = \mathbf{R}_t * \mathbf{\Psi}_{t-1} * \mathbf{S}'$ from $\mathbf{\Psi}_{t-1}$. Under linear labor disutility, we can calculate $w_t = \chi p_t / u'(C_t)$, which gives us all the information

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

needed to construct $\mathbf{U}_t$. Finally, given $\mathbf{V}_t$ and $\mathbf{V}_{t+1}$ we can construct $\mathbf{P}_t$, $\mathbf{D}_t$, and $\mathbf{D}_{t+1}$, and thus $\mathbf{\Lambda}_t$ and $\mathbf{G}_{t+1}$. Therefore the variables in $\overrightarrow{X}_t$ and z_t are indeed sufficient to evaluate the system (49)-(53).

Finally, if we linearize system $\mathcal{F}$ numerically with respect to all its arguments 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 $\mathcal{D} \equiv D_{z_t}\mathcal{F}$, then we obtain the following first-order linear expectational difference equation system:

$$E_t\mathcal{A}\Delta\overrightarrow{X}_{t+1} + \mathcal{B}\Delta\overrightarrow{X}_t + E_t\mathcal{C}z_{t+1} + \mathcal{D}z_t = 0 \tag{55}$$

where Δ represents a deviation from steady state. This system has the form considered by Klein (2000), so we solve our model using his QZ decomposition method.[20]

6 Results: dynamics

6.1 Effects of money growth shocks

We now study the impulse response functions implied by money supply shocks in several versions of our model.[21] Figures 3 and 5 show the response to a 1% increase in money supply growth in our SDSP calibration, and compare it with the response in the Calvo and fixed menu cost cases. All three versions are simulated under the same parameters, and the same aggregate and idiosyncratic shock processes, assuming zero trend inflation. Only the parameterization of the adjustment probability function λ differs across specifications. Each impulse response is calculated starting from the steady state distribution of prices and productivities associated with the corresponding specification. Figure 3 shows impulse responses under the assumption that money supply growth is *iid*, while Figure 5 assumes money growth has monthly autocorrelation of 0.8 (0.51 at quarterly frequency).[22]

As in any New Keynesian model, the impulse responses show that an increase in money growth stimulates consumption. Since not all prices adjust instantaneously, increased money growth raises households' real money balances, thus increasing consumption demand. However, as Golosov and Lucas (2007) emphasized, the average price level adjusts rapidly in the fixed menu cost specification, so there is a large, short-lived spike in the inflation rate. Therefore changes in real variables are small and not very persistent, approaching the monetary neutrality associated with full price flexibility. On the other hand, the response of our SDSP model (lines with round dots) mostly lies between the responses seen in the Calvo (lines with squares) and fixed menu cost (lines with crosses) specifications, and it is typically much closer to the Calvo case. In particular, since λ is much less than one for almost all firms in our SDSP model, prices rise much more gradually in our preferred model than they would under fixed menu costs, leading to a large and persistent increase in output resembling the response in the Calvo specification.

Impulse responses are also plotted for a number of other macroeconomic variables. It can be shown analytically that nominal interest rates depend on expected future money growth only (insert (13) into (46) to rewrite the Euler equation in terms of R_t, R_{t+1}, and μ_{t+1} only). Deviations of the nominal interest rate from steady state are proportional to those of money growth, with factor of proportionality $\frac{(R-1)\phi_z}{\mu R(R-\phi_z)}$.[23] Thus when money growth is uncorrelated ($\phi_z = 0$), nominal interest rates are constant, as we see in the third panel of Fig. 3. As for real interest rates, the fact that consumption rises above its long term level in response to an increase in money growth means that real interest rates must fall. Also, the real wage must rise, because labor must increase in order to produce the additional consumption goods.

We also plot the response of price dispersion, defined as

$$\Delta_t = \sum_{j,k} \frac{\Psi_t^{jk}}{a^k}\left(\frac{p^j}{p_t}\right)^{-\epsilon}$$

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

[21] Our results can be reproduced by running our MATLAB programs, which are available for download at http://www.econ.upf.edu/~nakov/dyn_programs.zip.

[22] For numerical tractability we compute equilibrium on the coarse grid of 25 productivities and 31 price levels analyzed in Table 1 and Figure 2, which yields a distribution of price changes similar to that on a much finer grid.

[23] Therefore, as is common in sticky-price models driven by money supply shocks, the 'anticipated inflation effect' is stronger than the 'liquidity effect', so increased money growth (if autocorrelated) causes the nominal interest rate to rise. The issue of the 'liquidity effect' in sticky price models is discussed in Gali (2003).

In our setup, part of the reason firms set different prices is that they face different productivities. But additional price dispersion, caused by failure to adjust when necessary, implies inefficient variation in demand across goods that acts like a negative aggregate productivity shock: $N_t = \Delta_t C_t$. In a representative agent model near a zero-inflation steady state, the dispersion wedge Δ_t is negligible because it is roughly proportional to the cross-sectional variance of prices, a quantity which is of second order in the inflation rate.[24] But the cross-sectional variance of prices is not second order in the inflation rate when idiosyncratic shocks are present. Thus variations in Δ_t are often quantitatively important, especially since $\epsilon = 7$ strongly magnifies variations in the ratio p^j/p_t. In the third row of Fig. 3 we see that increased money growth throws firms' prices further out of line with fundamentals, increasing dispersion, especially in the SDSP and Calvo specifications where prices are less flexible. Therefore increasing consumption requires a proportionally larger increase in labor in these specifications.

Next, comparing Figures 3 and 5, we note that while the shape of the inflation and output responses differs substantially across models, it is qualitatively similar under *iid* money growth and autocorrelated money growth. Unsurprisingly, inflation spikes immediately under fixed menu costs and *iid* money growth. More interestingly, it does the same under fixed menu costs and autocorrelated money growth, because the average price increase rises by much more than 1% (inflation immediately jumps by 2% when money growth jumps by 1%, because firms anticipate that money growth will remain positive for some time). On the other hand, there is a smaller but more persistent rise in inflation in our SDSP specification and in the Calvo specification. 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 price stickiness. Thus the big difference between the inflation and output responses in Figures 3 and 5 is one of size, not of shape: the overall response is larger when money growth is more autocorrelated.

Table 2 reports additional calculations regarding the extent of monetary nonneutrality in various parameterizations of our model. As in Section VI of Golosov and Lucas (2007), we address this issue by asking the following question: if money supply shocks were the only source of inflation variation, how much output variation would they cause? That is, for each specification, we pick the variance of the money supply shock to match 100% of observed US inflation volatility, and then calculate the implied variability of output. Interestingly, in the SDSP case, money shocks would explain almost all US output fluctuation (96%, if money growth is assumed *iid*; or 91%, if autocorrelated). Under the Calvo specification, implied output fluctuations would be about 40% larger. With menu costs, the figure is much lower: money supply shocks would only explain 26% or 29% of output fluctuations.[25] In addition, we calculate a "Phillips curve" coefficient by regressing output growth on money growth; the coefficient is at least twice as large in the SDSP specification as it is in the fixed menu cost specification (and is much larger if money growth is autocorrelated). None of these calculations should be taken as conclusive, since both inflation and output are affected by many other shocks besides money shocks; but they do all demonstrate that a model calibrated to match microdata on price adjustments yields much greater monetary nonneutrality than a model with fixed menu costs does.

6.2 Inflation decompositions

To further understand how the real effects of money shocks differ across models, we next decompose changes in the inflation rate into three main 'effects' mentioned in many recent papers. A variety of decompositions have been proposed, which differ both in details and in substance, but all start from the observation that inflation is an average of log nominal price changes. In our framework, all nominal price changes occur at the beginning of the period, starting from the beginning-of-period distribution $\widetilde{\mathbf{\Psi}}_t$. The nominal inflation rate at time t is

$$\pi_t = \sum_{j=1}^{\#^p} \sum_{k=1}^{\#^a} x_t^{jk} \lambda_t^{jk} \widetilde{\Psi}_t^{jk} \tag{56}$$

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 at the beginning of period t after shocks have been revealed.

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

[25] Even this figure is larger than Golosov and Lucas found, because of calibration differences, especially the fact that we use Nakamura and Steinsson's (2007) measure of median price adjustment frequency. But the important point is that a calibration chosen to match the distribution of price adjustments greatly increases monetary nonneutrality, compared with a calibration that assumes fixed menu costs.

Klenow and Kryvstov (2008) point out that (56) can be rewritten as the product of the average log price adjustment $\overline{x}_t$ times the frequency of price adjustment $\overline{\lambda}_t$:

$$\pi_t = \overline{x}_t \overline{\lambda}_t, \qquad \overline{x}_t \equiv \frac{\sum_{j,k} x_t^{jk} \lambda_t^{jk} \widetilde{\Psi}_t^{jk}}{\sum_{j,k} \lambda_t^{jk} \widetilde{\Psi}_t^{jk}}, \qquad \overline{\lambda}_t \equiv \sum_{j,k} \lambda_t^{jk} \widetilde{\Psi}_t^{jk} \tag{57}$$

This leads to the following inflation decomposition:

$$\Delta\pi_t = \overline{\lambda}\Delta\overline{x}_t + \overline{x}\Delta\overline{\lambda}_t + h.o.t. \tag{58}$$

where, as in Section 4, variables without time subscripts represent steady states, and Δ represents a deviation from steady state.[26] Klenow and Kryvstov'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 of inflation attributable to changes in the frequency of price adjustment.

A weakness of Klenow and Kryvstov's decomposition is that an increase in the average log price adjustment $\overline{x}_t$ may be caused by a rise in all firms' desired price adjustments, or by a reallocation of adjustment opportunities from firms desiring small or negative price changes to others desiring large positive price changes. That is, $\mathcal{I}_t^{KK}$ mixes the effect of changes in desired adjustments (the only relevant issue in time-dependent pricing models like that of Calvo) with the "selection effect" emphasized by Golosov and Lucas in their paper on state-dependent pricing. Therefore, we prefer a decomposition that breaks inflation into three terms: an intensive margin that captures changes in the average *desired* log price change, an extensive margin that captures changes in *how many* firms adjust, and a selection effect that captures changes in *who* adjusts.

All three effects can be defined clearly if we start by rewriting (56) as

$$\pi_t = \overline{x}_t^* \overline{\lambda}_t + \sum_{j,k} x_t^{jk} \left(\lambda_t^{jk} - \overline{\lambda}_t\right) \widetilde{\Psi}_t^{jk}, \qquad \overline{x}_t^* \equiv \sum_{j,k} x_t^{jk} \widetilde{\Psi}_t^{jk} \tag{59}$$

Note that in (59), $\overline{x}_t^*$ is the average *desired* log price change, whereas in (57), $\overline{x}_t$ is the average log price change *among those who adjust*. Thus (59) says that inflation equals the mean preferred 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 whenever some sizes of price adjustments x_t^{jk} are more or less likely than the mean probability of adjustment $\overline{\lambda}_t$, or (equivalently) when firms with different probabilities of adjustment λ_t^{jk} tend to prefer adjustments that differ from the mean preferred adjustment $\overline{x}_t^*$.

Equation (59) leads us to the following inflation decomposition:

$$\Delta\pi_t = \overline{\lambda}\Delta\overline{x}_t^* + \overline{x}^*\Delta\overline{\lambda}_t + \Delta\sum_{j,k} x_t^{jk} \left(\lambda_t^{jk} - \overline{\lambda}_t\right) \widetilde{\Psi}_t^{jk} + h.o.t. \tag{60}$$

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). Obviously $\mathcal{I}$ is the only nonzero term in the Calvo model, where $\lambda_t^{jk} = \overline{\lambda}$ for all j, k, and t. Our extensive margin effect, $\mathcal{E}_t \equiv \overline{x}^*\Delta\overline{\lambda}_t$, is the effect of changing the fraction of firms that adjust, if we select the new adjusters (or new nonadjusters) 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 adjustments x_t^{jk}, while fixing the overall fraction that adjust. The selection term is zero in the Calvo model, and also in a state-dependent model if we happen to start from a distribution with no heterogeneity ($\widetilde{\Psi}_t^{jk} = 1$ for some particular j and k).[27]

Caballero and Engel (2007) propose an alternative decomposition, which is also based on differencing (56):

$$\Delta\pi_t = \sum_{j,k} \Delta x_t^{jk} \lambda^{jk} \widetilde{\Psi}^{jk} + \sum_{j,k} x^{jk} \Delta\lambda_t^{jk} \widetilde{\Psi}^{jk} + \sum_{j,k} x^{jk} \lambda^{jk} \Delta\widetilde{\Psi}_t^{jk} + h.o.t. \tag{61}$$

[26] Actually, Klenow and Kryvstov (2008) propose a time series variance decomposition, whereas (57) is a decomposition of each period's inflation realization. But the logic of (57) is the same as that in their paper.

[27] To see this, note that $\overline{\lambda}_t = \lambda_t^{jk}$ for the (j,k) pair such that $\widetilde{\Psi}_t^{jk} = 1$, and that $\widetilde{\Psi}_t^{jk} = 0$ for all other j and k.

They further simplify this to

$$\Delta \pi_t = \overline{\lambda} \Delta \mu_t + \sum_{j,k} x^{jk} \Delta \lambda_t^{jk} \widetilde{\Psi}^{jk} \tag{62}$$

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.[28] Note that their first term, $\mathcal{I}^{CE} \equiv \Delta \mu_t \overline{\lambda}$, is the same as our intensive margin $\mathcal{I}$, as long as their assumption that all desired price adjustments change by $\Delta \mu_t$ is correct. But therefore, their "extensive margin" term $\mathcal{E}^{CE} \equiv \sum_{j,k} x^{jk} \Delta \lambda_t^{jk} \widetilde{\Psi}^{jk}$, combines the issue of *how many* firms adjust (our extensive margin $\mathcal{E}_t$) with the issue of *who* adjusts (our selection effect $\mathcal{S}_t$), which we think it is clearer to consider separately.

The second rows of Figures 3 and 5 illustrate our decomposition of the inflation impulse response. (The inflation decomposition at the time of the shock, $t = 1$, is also presented in Table 3.) The panels representing inflation and its three components $\mathcal{I}_t$, $\mathcal{E}_t$, and $\mathcal{S}_t$ are shown to the same scale for better comparison. The graphs unambiguously illustrate that the short, sharp rise in inflation observed in the fixed menu cost specification results from the selection effect. This is true both under *iid* money shocks, where inflation spikes to 0.451% on impact, of which 0.353% is the selection component, and under autocorrelated shocks, where inflation spikes to 1.94%, with 1.48% due to selection. In contrast, inflation in the Calvo model is caused by the intensive margin only; in SDSP there is a nontrivial selection effect but it still only accounts for around one-third of the impulse response of inflation.

We also see that the extensive margin, $\mathcal{E}_t \equiv \overline{x}^* \Delta \overline{\lambda}_t$, plays a negligible role in the inflation impulse response. This makes sense, because we are considering a steady state with zero inflation, so steady state price adjustments are responses to idiosyncratic shocks only, and the average desired adjustment $\overline{x}^*$ must be close to zero. Therefore $\mathcal{E}_t$ is tiny even though the adjustment frequency $\overline{\lambda}_t$ does vary, rising from 10% to 12.5% on impact in the MC model, and from 10% to 10.3% on impact in SDSP.[29] The extensive margin only becomes important when there is high trend inflation: then the average desired adjustment $\overline{x}^*$ is large and positive, because firms that have not adjusted recently need to make substantial price increases. Variations in $\overline{\lambda}_t$ then account for a large part of the inflationary impact of money growth, as we see in Table 3 and Figure 6, which analyze money supply shocks in an equilibrium with 63% inflation *per annum* (the highest rate observed in the Mexican data of Gagnon, 2007).[30] Money supply shocks imply a bigger spike in inflation when trend inflation is high, bringing the model even closer to monetary neutrality in this case than it was in Figure 3, and much of this difference is due to the extensive margin.

Turning to the intensive margin, on impact the effect is the same in all three models. This reflects the fact that an uncorrelated 1% increase in money growth raises the mean desired price change by approximately 1 percentage point in all three models, implying an initial intensive margin effect $\mathcal{I}_1 \equiv \overline{\lambda} \Delta \overline{x}_1^*$ of 0.1 percentage points (since we calibrated all models to $\overline{\lambda} = 0.1$). The importance of the intensive margin fades quickly in the MC model, since those desiring the largest price changes do in fact adjust immediately, but remains important in the specifications with lower state dependence. Likewise, the initial intensive margin effect is larger, but the same across models, under autocorrelated money growth: the mean desired price change rises roughly five percentage points in all three specifications, so that $\mathcal{I}_1 \equiv \overline{\lambda} \Delta \overline{x}_1^* \approx 0.5\%$. That is, firms want to "frontload" prices in response to autocorrelated money shocks by roughly the same amount in all three specifications; the difference is that in the MC case many of these changes take place immediately, whereas they are only gradually realized under the other specifications.

We can also understand the three effects by examining Figure 4, which illustrates the distribution of price adjustments, before and after an increase in money growth. As we have emphasized, the distribution of adjustments in the MC model is strongly bimodal. The main effect of an increase in money growth is to make large price decreases less likely, and large price increases more likely, causing a large increase in inflation overall.

[28] Our equation (60) is intended to decompose each period's inflation realization, and therefore it must take into account shifts in the current distribution $\widetilde{\Psi}_t^{jk}$. Caballero and Engel instead propose a decomposition (see their equation 17) of the *average* impact of a money supply shock. Therefore they evaluate their decomposition at the ergodic distribution (the time average over all cross-sectional distributions, which is 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.

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

[30] All parameters are the same as in our earlier simulations, except the steady state money growth rate. We suppress the Calvo specification in this case, because no Calvo equilibrium exists if we fix parameters (including the adjustment rate $\overline{\lambda} = 0.1$) and raise the inflation rate to 63%.

This redistribution in adjustment probabilities is, by our definition, a selection effect. In contrast, in the Calvo case, the whole distribution of adjustments shifts right with no change in shape, which by our definition is an intensive margin change. In the SDSP case, the change mostly resembles a rightward shift of the distribution, but we can also see some redistribution of adjustment probability from the left mode of the distribution to the right mode, so both the intensive margin and the selection effect are active in this case. As for the extensive margin, little difference in the overall adjustment frequency $\overline{\lambda}_t$ is visible in the graphs, even in the MC case.

In summary, the sharp state dependence of the fixed menu cost model is the key to understanding both its implications for the distribution of price adjustments and its implications for monetary neutrality. We say that the state dependence is strong under fixed menu costs because it implies that λ is a step function: at the threshold, a tiny increase in the value of adjustment suffices to increase the adjustment probability from 0 to 1. This behavior implies that the distribution of price changes consists of two spikes: there are no small changes, and firms cannot drift outside the adjustment thresholds before change takes place. Hence, in steady state, those firms whose behavior might be affected by a money shock are all near the two adjustment thresholds. Therefore, the main effect of an increase in money growth is to decrease λ from 1 to 0 for some firms desiring a price decrease, while increasing λ from 0 to 1 for some firms desiring a price increase. This selection effect implies a strong change in inflation, and monetary near-neutrality, as Golosov and Lucas argued. But it depends on an extreme degree of state dependence which our estimates reject. This may help explain why money shocks do, in aggregate time series, appear to have important real effects.

6.3 The role of the distribution

Besides calculating the impact of money growth shocks, we can also calculate transitional dynamics, which shed light on a number of issues. Figure 7 illustrates transitional dynamics of the SDSP calibration of the model, starting from a variety of different initial conditions. Figure 8 shows how the impulse response to a money supply shock changes when starting from different initial conditions. Both figures also illustrate nonlinear aspects of the dynamics that our hybrid linear/nonlinear solution method can capture.

Figure 7.1 shows the impact of a large monetary shock, but the calculations are carried out in a different way from those behind Figure 3. Instead of starting at the steady state distribution and feeding in a money shock, we simply shift the distribution of real prices two grid points to the left. That is, we start from an initial distribution $\mathbf{\Psi}_0$ such that $\Psi_0^{jk} = \Psi^{j+2,k}$: the lagged distribution from period 0 is the steady state distribution, shifted left by two points in the price grid. Each step in the price grid is a difference of 2.5%, so this amounts to a 5% real price decrease, which is also equivalent to an uncorrelated increase in money supply growth by 5%. By calculating the effects of the shock in this way, we take nonlinear changes in the impulse response into account, since our computational method allows full nonlinearity between one grid point and the next.[31] Some of the impulse responses are proportional to those shown before; for example, the response of inflation in Figure 7.1 (a 0.77% rise on impact) is roughly five times larger than that shown for the SDSP calibration in the second panel of Figure 3 (a 0.15% rise). But the fraction of firms adjusting (not shown) increases in a more-than-linear way, increasing eight times as much in this example compared with the example of Figure 3, which makes sense since the value of adjustment increases nonlinearly in the distance from the optimal price. Since this especially affects the prices that are furthest out of line, price dispersion *falls* substantially, whereas it increased in Figure 3. This permits labor to decrease even as consumption rises.

In Figure 7.2, we see how our model can also shed light on the effects of technology shocks. The exercise is similar to that in Fig. 7.1, but instead of shifting the distribution in the price direction, we shift the distribution by two grid points in the productivity direction: $\Psi_0^{jk} = \Psi^{j,k+2}$. Thus the transition dynamics in Fig. 7.2 represent the effect of a persistent, but not permanent, 6.1% increase in productivity.[32] The effects confirm that the implications of productivity shocks known from macroeconomic models based on Calvo pricing (e.g. Galí 1999) also occur under state-dependent pricing. Higher productivity drives inflation down, and permits households to increase consumption while decreasing labor. Note that since prices take time to adjust,

[31] When calculating impulse responses by feeding a change in z_0 into (55), doubling the size of the shock exactly doubles the size of all responses, since is (55) linear. But since each grid point is governed by a different equation, there is no linear relationship between the dynamic coefficients at one grid point and another. Therefore, the impulse response derived by shifting the distribution across grid points incorporates non-linear effects.

[32] This productivity shock is somewhat nonstandard, since it is the result of an unexpected correlated increase in all firms' productivity, instead of an aggregate shock *per se*. Nonetheless we see that its effects are very similar to those of an aggregate technology shock.

consumption takes time to reach its peak level; since consumption is initially growing and thereafter decreasing, the real interest rate first rises above its steady state level and then falls back below it. An additional effect in our state-dependent model is that those prices pushed furthest out of line by the productivity shock are most likely to adjust (downwards). Therefore part of the decrease in inflation is a selection effect, and price dispersion falls as the most extreme prices adjust more than usual, permitting labor to fall further than it otherwise would.

Figure 7.3 shows one more example of transitional dynamics, in which we assume that all firms start at their optimal sticky prices. That is, the initial (lagged) price distribution is $\Psi_0^{jk} = \Psi^{l(k),k}$, where $l(k)$ is the grid point associated with the steady state optimal price: $p^{l(k)} = p^*(a^k)$. This might be seen as the effect of "introducing the euro": it is as if all firms have just been forced to change prices, taking into account that their prices will be sticky in the future. Conditional on this one-time change, the fraction of firms adjusting is thereafter below steady state (not shown), since they start out at their preferred price. We also see that the price dispersion measure starts out substantially below its steady state value. This acts as a positive productivity shock: firms are on average closer to their efficient prices, permitting more consumption with less labor input.

Finally, Figure 8 shows an example of how different initial conditions alter the effects of a monetary shock, which is another non-linear phenomenon picked up by our solution method. For the SDSP calibration, it shows the effect of an uncorrelated 1% increase in money supply at time 1, either starting from the steady state distribution (as seen previously in Fig. 3), or occurring simultaneously with an increase in aggregate productivity (that is, starting from $\Psi_0^{jk} = \Psi^{j,k+2}$, as in Fig. 7.2). In both cases, we aim to show only the effect of the money supply shock itself, so to graph the blue circled curve in Fig. 8 we first compute a path starting from a technology shock plus a money shock; then we compute a path starting from a technology shock only (that is, the path shown in Fig. 7.2), and then we take the difference between the two. Most of the impulse responses are essentially unchanged; in particular, the response of real consumption to the money shock is not altered by the starting distribution. However, note that the money supply shock, by offsetting the incentive to decrease prices that results from the technology shock, decreases price dispersion in this case, requiring less increase in labor to finance the same amount of consumption. Since labor is the most elastic choice variable here, it absorbs all the difference.

7 Conclusions

In this paper, we have computed the impact of money growth shocks in a quantitative macroeconomic model of state-dependent pricing. We have calibrated the model for consistency with microeconomic data on firms' pricing behavior. In particular, we have estimated how firms' probability of price adjustment depends on the value of adjustment, adopting a flexible specification that nests the Calvo specification as one extreme, and the fixed menu cost specification as the other extreme. Given our estimate of this adjustment function, we have then characterized the dynamics of the distribution of prices and productivities in general equilibrium.

In our calibrated model, we find that shocks to money growth have large, persistent effects on real variables, only slightly weaker than the effects found in the Calvo model. Prices rise gradually in response to increased money growth, leading to a persistent stimulative effect on consumption and labor. Real interest rates fall; real wages and real money holdings rise; the nominal interest rate is constant if money growth is *iid*, and rises if money growth is autocorrelated. We also find that the main factor determining how monetary shocks propagate to the rest of the economy is the degree of state dependence. That is, the autocorrelation of money shocks has little effect on the shape of the impulse responses of most variables and little effect on their persistence. Instead, increasing the autocorrelation of money growth shocks simply makes their real effects larger.

We show how the impulse response of inflation can be decomposed into an intensive margin effect relating to the average desired price adjustment, an extensive margin effect relating to the number of firms adjusting, and a selection effect relating to changes in the relative frequencies of small and large or negative and positive adjustments. In our calibrated model, starting from a low baseline inflation rate, about two-thirds of the effect of a money growth shock comes through the intensive margin, and most of the rest through the selection effect. The extensive margin only matters when starting from a high baseline inflation rate.

As Golosov and Lucas (2007) argued, in a model of fixed menu costs the real effects of money supply shocks are greatly decreased, because prices jump strongly on impact. They rightly attributed this to a selection effect. However, such strong selection effects only arise if a small change in the value of adjustment can cause a large jump in the probability of adjustment. Our estimate of the function governing the probability of adjustment rejects such extreme state dependence, favoring a specification that behaves more like the Calvo model. As state dependence increases towards the fixed menu cost specification, the distribution of price adjustments becomes more and more strongly bimodal. Weaker state dependence yields a more bell-shaped distribution of adjustments, which is more consistent with microdata. The weaker state dependence also implies that money growth shocks have nontrivial real effects, as VAR evidence suggests.

8 Appendix A: Detrending

Suppose the model can be rewritten in real terms by deflating all prices by the nominal money stock, defining $p_t \equiv P_t/M_t$, $p_{it} \equiv P_{it}/M_t$, and $w_t \equiv W_t/M_t$. Given the nominal distribution $\Phi_t(P_i, A_i)$ and the money stock M_t, we denote by $\Psi_t(p_i, A_i)$ the distribution over real production prices $p_{it} \equiv P_{it}/M_t$. Likewise, let $\widetilde{\Psi}_t(p_i, A_i)$ be the distribution of real beginning-of-period prices $\widetilde{p}_{it} \equiv \widetilde{P}_{it}/M_t$, in analogy to the beginning-of-period distribution of nominal prices $\widetilde{\Phi}_t(\widetilde{P}_i, A_i)$. If it is true that the model can be rewritten in real terms, then it is not necessary to condition equilibrium behavior on M_t; conditioning on $\Xi_t \equiv (z_t, \Psi_{t-1})$ suffices.[33]

Therefore, if there exists a real equilibrium, the real aggregate functions can be written in terms of Ξ_t only, and must satisfy $C_t = C(\Xi_t) = C(\Omega_t)$, $N_t = N(\Xi_t) = N(\Omega_t)$, $p_t = p(\Xi_t) = P(\Omega_t)/M_t$ and $w_t = w(\Xi_t) = W(\Omega_t)/M_t$. Deflating from one period to the next will depend on the growth rate of money supply from one period to the next (but not on the level of the money supply). Thus the stochastic discount factor will be

$$q(\Xi_t, \Xi_{t+1}) = \beta \frac{M_t p(\Xi_t) u'(C(\Xi_{t+1}))}{M_{t+1} p(\Xi_{t+1}) u'(C(\Xi_t))} = \beta \frac{p(\Xi_t) u'(C(\Xi_{t+1}))}{\mu e^{z_{t+1}} p(\Xi_{t+1}) u'(C(\Xi_t))}$$

The "real" value function v should likewise be the nominal value function, divided by the current money stock, and should be written as a function of real prices. Therefore we have

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

If a firm's nominal price at time t is P_{it}, then the value of maintaining this price fixed at time $t+1$ can be written in nominal or real terms as

$$\begin{aligned} V(P_{it}, A_{i,t+1}, \Omega_{t+1}) &= M_{t+1} v\left(\frac{P_{it}}{M_{t+1}}, A_{i,t+1}, \Xi_{t+1}\right) \\ &= M_{t+1} v\left(\frac{M_t p_{it}}{M_{t+1}}, A_{i,t+1}, \Xi_{t+1}\right) = M_{t+1} v\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1}\right) \end{aligned}$$

Likewise, if for any time t nominal price P_{it} we have the definitions

$$D(P_{it}, A_{i,t+1}, \Omega_{t+1}) \equiv \max_{P'} V(P', A_{i,t+1}, \Omega_{t+1}) - V(P_{it}, A_{i,t+1}, \Omega_{t+1})$$

$$G\left(P_{it}, A_{i,t+1}, \Omega_{t+1}\right) \equiv \lambda\left(\frac{D\left(P_{it}, A_{i,t+1}, \Omega_{t+1}\right)}{W(\Omega_{t+1})}\right) D\left(P_{it}, A_{i,t+1}, \Omega_{t+1}\right)$$

then we can define

$$D(P_{it}, A_{i,t+1}, \Omega_{t+1}) \equiv M_{t+1} d\left(\frac{P_{it}}{M_{t+1}}, A_{i,t+1}, \Xi_{t+1}\right) = M_{t+1} d\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1}\right)$$

$$G(P_{it}, A_{i,t+1}, \Omega_{t+1}) \equiv M_{t+1} g\left(\frac{P_{it}}{M_{t+1}}, A_{i,t+1}, \Xi_{t+1}\right) = M_{t+1} g\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1}\right)$$

Using this deflated notation, we can rewrite the Bellman equation as

[33] If money growth is uncorrelated, then z_t has no effect on the distribution of μ_{t+1}, so the aggregate state can be summarized by $\widetilde{\Phi}_t$ only. But when money growth is correlated we must keep track of $\Xi_t \equiv (\widetilde{\Phi}_t, z_t)$ instead.

$$M_t v(p_{it}, A_{it}, \Xi_t) = M_t \left(p_{it} - \frac{w(\Xi_t)}{A_{it}} \right) \left(\frac{p_{it}}{p(\Xi_t)} \right)^{-\epsilon} C(\Xi_t) +$$

$$\beta E_t \left\{ \frac{M_t p(\Xi_t) u'(C(\Xi_{t+1}))}{M_{t+1} p(\Xi_{t+1}) u'(C(\Xi_t))} \; M_{t+1} \left[v\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1} \right) + g\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1} \right) \right] \middle| A_{it}, \Xi_t \right\}$$

Note that M_t cancels from both sides of the equation, and M_{t+1} cancels inside the expectation. Therefore we obtain

$$v(p_{it}, A_{it}, \Xi_t) = \left(p_{it} - \frac{w(\Xi_t)}{A_{it}} \right) \left(\frac{p_{it}}{p(\Xi_t)} \right)^{-\epsilon} C(\Xi_t) +$$

$$\beta E_t \left\{ \frac{p(\Xi_t) u'(C(\Xi_{t+1}))}{p(\Xi_{t+1}) u'(C(\Xi_t))} \left[v\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1} \right) + g\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1} \right) \right] \middle| A_{it}, \Xi_t \right\}$$

where

$$g\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1} \right) \equiv \lambda \left(\frac{d\left(p_{it}/\mu_{t+1}, A_{i,t+1}, \Xi_{t+1}\right)}{w(\Xi_{t+1})} \right) d\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1} \right)$$

$$d\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1} \right) \equiv \max_{p'} v(p', A_{i,t+1}, \Xi_{t+1}) - v\left(\frac{p_{it}}{\mu_{t+1}}, A_{i,t+1}, \Xi_{t+1} \right)$$

Let $p^*(A_{i,t+1}, \Xi_{t+1})$ denote the optimal choice in the maximization problem above. Taking into account the fact that the firm starts period $t+1$ with the eroded price $\widetilde{p}_{i,t+1} \equiv p_{it}/\mu_{t+1}$, the price process is

$$p_{i,t+1} = \begin{cases} p(A_{i,t+1}, \Xi_{t+1}) & \text{with } prob = \lambda \left(\frac{d\left(p_{it}/\mu_{t+1}, A_{i,t+1}, \Xi_{t+1}\right)}{w(\Xi_{t+1})} \right) \\ p_{it}/\mu_{t+1} & \text{with } prob = 1 - \lambda \left(\frac{d\left(p_{it}/\mu_{t+1}, A_{i,t+1}, \Xi_{t+1}\right)}{w(\Xi_{t+1})} \right). \end{cases}$$

In other words, when the firm's nominal price is not adjusted at time $t+1$, its real price is deflated by factor $\mu_{t+1} = \mu e^{z_{t+1}}$.

9 Appendix B. Differentiability of the discretized equation system

Computing aggregate dynamics by linearization requires that we define the system $\mathcal{F}$ so that it varies smoothly with respect to all its arguments. As long as $\xi < \infty$, this only requires us to be careful about how we specify the maximization problem that appears in (23), (24), (41), and (44). If we were to calculate v_t^{jk} by choosing a price p on the grid Γ^p, then p_t^{*k} could vary discontinuously, by jumping from one grid point to another in response to some small change in aggregate conditions Ξ_t. Therefore we instead assume p may be chosen at points off the grid Γ^p. This requires us to use splines to interpolate the value function $v_t(p, a^k)$ for points $p \notin \Gamma^p$. Then, to map prices back onto the grid, we stochastically round p_t^{*k} up or down to the nearest grid points as described in (27). For $\xi < \infty$, this suffices to ensure differentiability of all equations in system $\mathcal{F}$.

For the fixed menu cost case $\xi = \infty$, one additional issue arises. In this case, the adjustment probability is a step function, $\lambda(L) = \mathbf{1}\{L \geq \alpha\}$, so it jumps discontinuously from 0 to 1 depending on the value of adjustment L. Therefore, if we interpret λ_t^{jk} as the probability of adjustment in state (p^j, a^k, Ξ_t), it will only take values 0 or 1 and may jump discontinuously in response to changes in aggregate conditions Ξ_t. Instead, in this case, we interpret λ_t^{jk} as the probability of adjustment in state (a^k, Ξ_t) when the price lies in the interval $I^j \equiv \left(\frac{p^{j-1}+p^j}{2}, \frac{p^j+p^{j+1}}{2} \right)$. To do so, we linearly interpolate the gain from adjustment $d\left(p, a^k, \Xi_t\right)$ at points $p \notin \Gamma^p$ using the values d_t^{jk} at the grid points $p^j \in \Gamma^p$. Using this interpolation, we define λ_t^{jk} as the fraction of interval I^j on which $d\left(p, a^k, \Xi_t\right) \geq \alpha$. In other words, λ_t^{jk} represents the fraction of interval I^j on which adjustment occurs. Thus λ_t^{jk} may take values 0 or 1, but may also lie strictly between 0 and 1, and will vary continuously as Ξ_t changes. This also ensures differentiability almost everywhere: that is, λ_t^{jk} will fail to be differentiable only if $d\left(p, a^k, \Xi_t\right) = \alpha$ exactly at $p = \frac{p^{j-1}+p^j}{2}$ or $p = \frac{p^j+p^{j+1}}{2}$, but generically this does not occur.

There is also a second, simpler way of ensuring differentiability for the fixed menu cost case: evaluate it by considering a large, finite ξ instead of the limiting case $\xi = \infty$. We have verified that this alternative gives quantitatively similar results.

References

[1] Akerlof, George, and Janet Yellen (1985), "A near-rational model of the business cycle with wage and price inertia." *Quarterly Journal of Economics* 100 (Supplement), pp. 823-38.

[2] Álvarez, Luis (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, Susanto (2005), "Comment on: Implications of state-dependent pricing for dynamic macroeconomic modeling." *Journal of Monetary Economics* 52 (1), pp. 243-7.

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

[5] Burstein, Ariel (2006), "Inflation and output dynamics with state-dependent pricing decisions." *Journal of Monetary Economics* 53, pp. 1235-57.

[6] Caballero, Ricardo, and Eduardo Engel (1993), "Microeconomic rigidities and aggregate price dynamics." *European Economic Review* 37 (4), pp. 697-711.

[7] Caballero, Ricardo, and Eduardo Engel (1999), "Explaining investment dynamics in U.S. manufacturing: a generalized (S,s) approach." *Econometrica* 67 (4), pp. 741-82.

[8] Caballero, Ricardo, and Eduardo Engel (2006), "Price stickiness in Ss models: basic properties." Manuscript, Massachusetts Institute of Technology.

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

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

[11] Caplin, Andrew, and John Leahy (1991), "State-dependent pricing and the dynamics of money and output." *Quarterly Journal of Economics* 106 (3), pp. 683-708.

[12] Caplin, Andrew, and John Leahy (1997), "Aggregation and optimization with state-dependent pricing." *Econometrica* 65, pp. 601-25.

[13] Caplin, Andrew, and Daniel Spulber (1987), "Menu costs and the neutrality of money." *Quarterly Journal of Economics* 102 (4), pp. 703-26.

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

[15] Costain, James, and Anton Nakov (2008), "Price adjustments in a general model of state-dependent pricing." Banco de España Working Paper #0824.

[16] Damjanovic, Vladislav, and Charles Nolan (2005), "Aggregate dynamics with heterogeneous agents and state-dependent pricing." Univ. of St. Andrews CDMA Working Paper 05/05.

[17] Den Haan, Wouter (1997), "Solving dynamic models with aggregate shocks and heterogeneous agents." Macroeconomic Dynamics, 1, 355–386.

[18] Dorich, José (2008), "The welfare losses of price rigidity." Manuscript, Univ. Pompeu Fabra.

[19] Dotsey, Michael; Robert King; and Alexander Wolman (1999), "State-dependent pricing and the general equilibrium dynamics of money and output." *Quarterly Journal of Economics* 114 (2), pp. 655-90.

[20] Dotsey, Michael, and Robert King (2005), "Implications of state-dependent pricing for dynamic macroeconomic modeling." *Journal of Monetary Economics* 52 (1), pp. 213-42.

[21] Dotsey, Michael; Robert King; and Alexander Wolman (2008), "Inflation and real activity with firm-level productivity shocks." Manuscript, Boston University.

[22] Gagnon, Etienne (2007), "Price setting under low and high inflation: evidence from Mexico." Manuscript, Federal Reserve Board.

[23] Galí, Jordi (1999), "Technology, employment, and the business cycle: do technology shocks explain aggregate fluctuations?" *American Economic Review* 89 (1), pp. 249-71.

[24] Galí, Jordi (2003), "New Perspectives on Monetary Policy, Inflation, and the Business Cycle", in Advances in Economic Theory, edited by: M. Dewatripont, L. Hansen, and S. Turnovsky, vol. III, 151-197, Cambridge University Press.

[25] Galí, Jordi (2008), *Monetary Policy, Inflation, and the Business Cycle: an Introduction to the New Keynesian Framework.* Princeton University Press.

[26] Gertler, Mark, and John Leahy (2006), "A Phillips curve with an (S,s) foundation." NBER Working Paper #11971.

[27] Golosov, Mikhail, and Robert E. Lucas, Jr. (2007), "Menu costs and Phillips curves." *Journal of Political Economy* 115 (2), pp. 171-99.

[28] Kashyap, Anil (1995), "Sticky prices: new evidence from retail catalogs." *Quarterly Journal of Economics* 110 (1), pp. 245-74.

[29] Klein, Paul (2000), "Using the generalized Schur form to multivariate linear rational expectations model." *Journal of Economic Dynamics and Control* 24, pp. 1405-23.

[30] Klenow, Peter, and Oleksiy Kryvstov (2008), "State-dependent or time-dependent pricing: does it matter for recent US inflation?" *Quarterly Journal of Economics* 123, pp. 863-94.

[31] Krusell, Per, and Anthony Smith (1998), "Income and wealth heterogeneity in the macroeconomy." *Journal of Political Economy* 106 (5), pp. 245-72.

[32] Mankiw, N. Gregory (1985), "Small menu costs and large business cycles: a macroeconomic model of monopoly." *Quarterly Journal of Economics* 100 (2), pp. 529-37.

[33] Midrigan, Virgiliu (2008), Menu Costs, Multi-Product Firms and Aggregate Fluctuations, Manuscript, New York University

[34] Nakamura, Emi, and Jón Steinsson (2007), "Five facts about prices: a reevaluation of menu cost models." *Quarterly Journal of Economics*, forthcoming.

[35] Reis, Ricardo (2006), "Inattentive producers." *Review of Economic Studies* 73, pp. 793-821.

[36] Reiter, Michael (2008), "Solving heterogeneous-agent models by projection and perturbation." *Journal of Economic Dynamics and Control*, forthcoming. doi:10.1016/j.jedc.2008.08.010

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

[38] Sheedy, Kevin (2007A), "Intrinsic inflation persistence." Manuscript, London School of Economics.

[39] Sheedy, Kevin (2007B), "Inflation persistence when price stickiness differs between industries." Manuscript, London School of Economics.

[40] Sims, Christopher (2001), “Solving linear rational expectations models.” *Computational Economics* 20 (1-2), pp. 1-20.

[41] Taylor, John (1993), “Comment on: Microeconomic rigidities and aggregate price dynamics.” *European Economic Review* 37 (4), pp. 714-7.

[42] Woodford, Michael (2008), “Information-constrained state-dependent pricing.” Manuscript, Columbia Univ.

[43] Yun, Tack (2005), "Optimal monetary policy with relative price distortions." *American Economic Review* 95 (1), pp. 89-109.

[44] Zbaracki, Mark; Mark Ritson; Daniel Levy; Shantanu Dutta; and Mark Bergen (2005), "Managerial and customer costs of price adjustment: direct evidence from industrial markets." *Review of Economics and Statistics* 86 (2), pp. 514-33.

Figure 1. Steady-state value function, distributions, and related objects (SDSP)

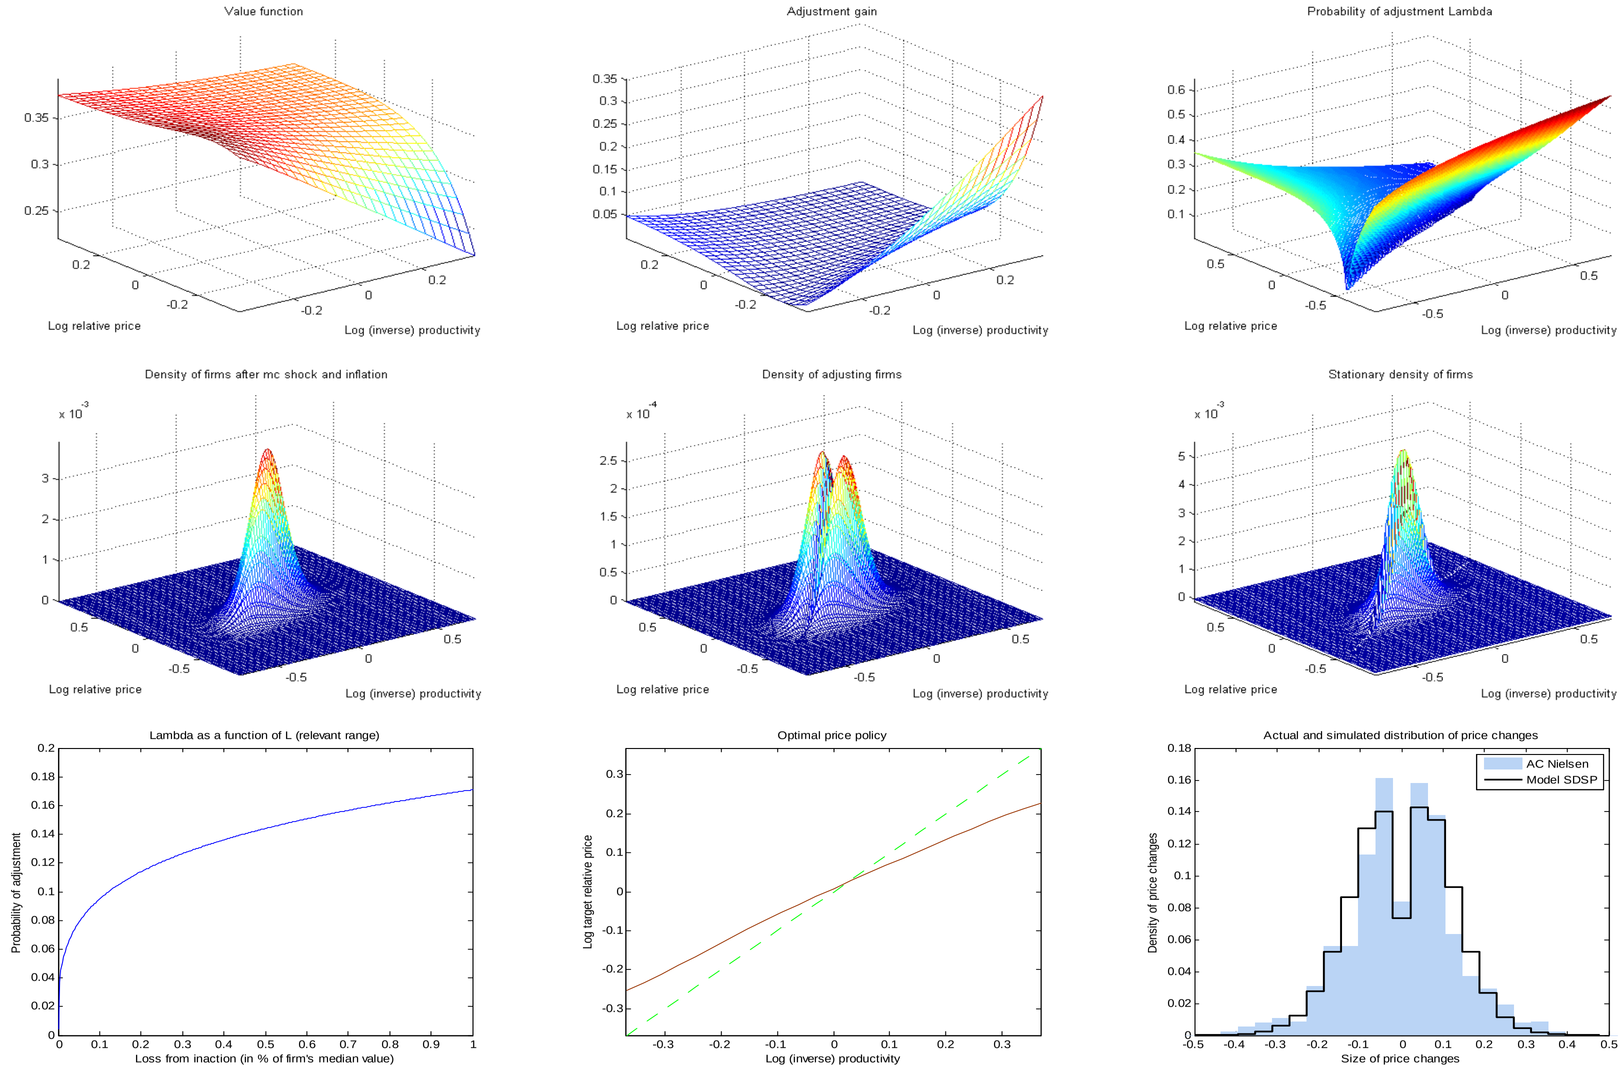

Simulations from SDSP model. First line: value V, adjustment gain, and adjustment probability lambda, as functions of real price and productivity shock. Second line: beginning of period distribution, adjustment distribution, and distribution at time of production, as functions of real price and productivity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

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

SDSP (201x201): $(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \alpha, \xi) = (0.0049, 0.8808, 0.1091, 0.0310, 0.2900)$
SDSP (31x25): $(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \alpha, \xi) = (0.0049, 0.8812, 0.1089, 0.0311, 0.2937)$
Woodford: $(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \alpha, \xi) = (0.0085, 0.8596, 0.0946, 0.0609, 1.3341)$
Calvo: $(\sigma_\varepsilon^2, \rho, \bar{\lambda}, \xi) = (0.0072, 0.8576, 0.10, 0)$
Menu cost: $(\sigma_\varepsilon^2, \rho, \alpha, \xi) = (0.0059, 0.8469, 0.0631, \infty)$

Model	MC	Calvo	Wdfd	SDSP			Evidence		
	$\xi = \infty$	$\xi = 0$		201x201	31x25	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	8.6	8.0	5.3	8.5	9.7
Std of price changes	18.8	8.2	13.2	12	11.8	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	50.6	50.4	50	65.5	66	56
% of price changes $\leq$5% in abs value	0	49.7	35.4	25.3	24.8	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	6.3				
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	0.269	0.260				
Median loss as % of median value	0.042	0.114	0.134	0.073	0.075				
Std of loss as % of median firm value	0.098	0.717	0.912	0.573	0.504				

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)

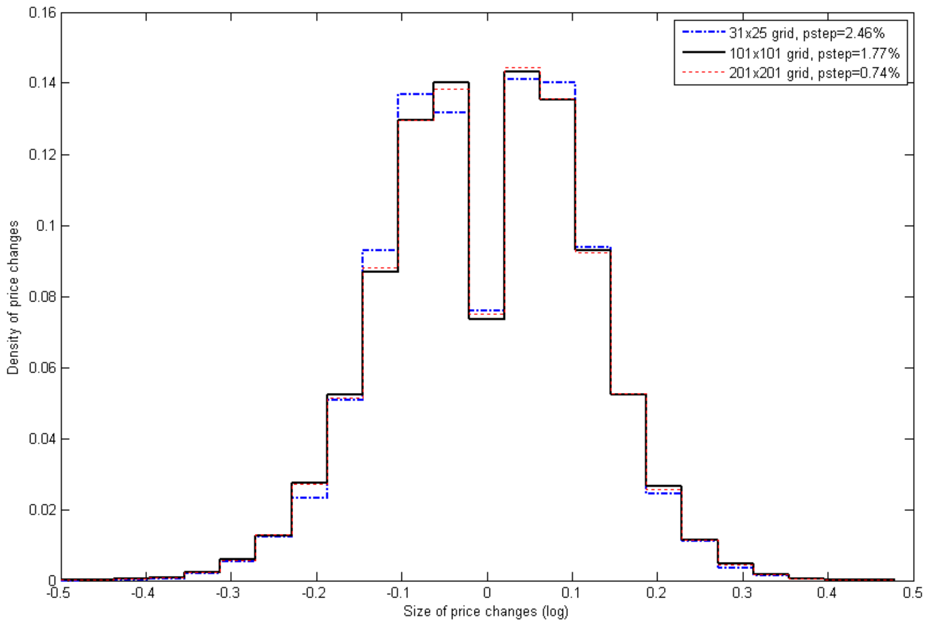

Figure 2. Steady-state distribution of nonzero price changes on different grids

Table 2. Variance decomposition and Phillips curves of alternative models

Uncorrelated money shock, $\phi_z = 0$	Data	SDSP	Calvo	Menu cost
Std of money shock (x100)		0.81	1.05	0.52
Std of quarterly inflation (x100)	0.25	0.25	0.25	0.25
Share explained by μ shock alone		100%	100%	100%
Std of quarterly output growth (x100)	0.51	0.49	0.72	0.13
Share explained by μ shock alone		96%	142%	26%
Slope coeff. of the Phillips curve		0.46	0.51	0.23
Standard error		0.02	0.03	0.00
R^2		0.33	0.20	0.71
Correlated money shock, $\phi_z = 0.8$				
Std of money shock (x100)		0.16	0.21	0.11
Std of quarterly inflation (x100)	0.25	0.25	0.25	0.25
Share explained by μ shock alone		100%	100%	100%
Std of quarterly output growth (x100)	0.51	0.47	0.67	0.15
Share explained by μ shock alone		91%	131%	29%
Slope coeff. of the Phillips curve		2.20	2.88	0.82
Standard error		0.00	0.03	0.01
R^2		0.99	0.88	0.89

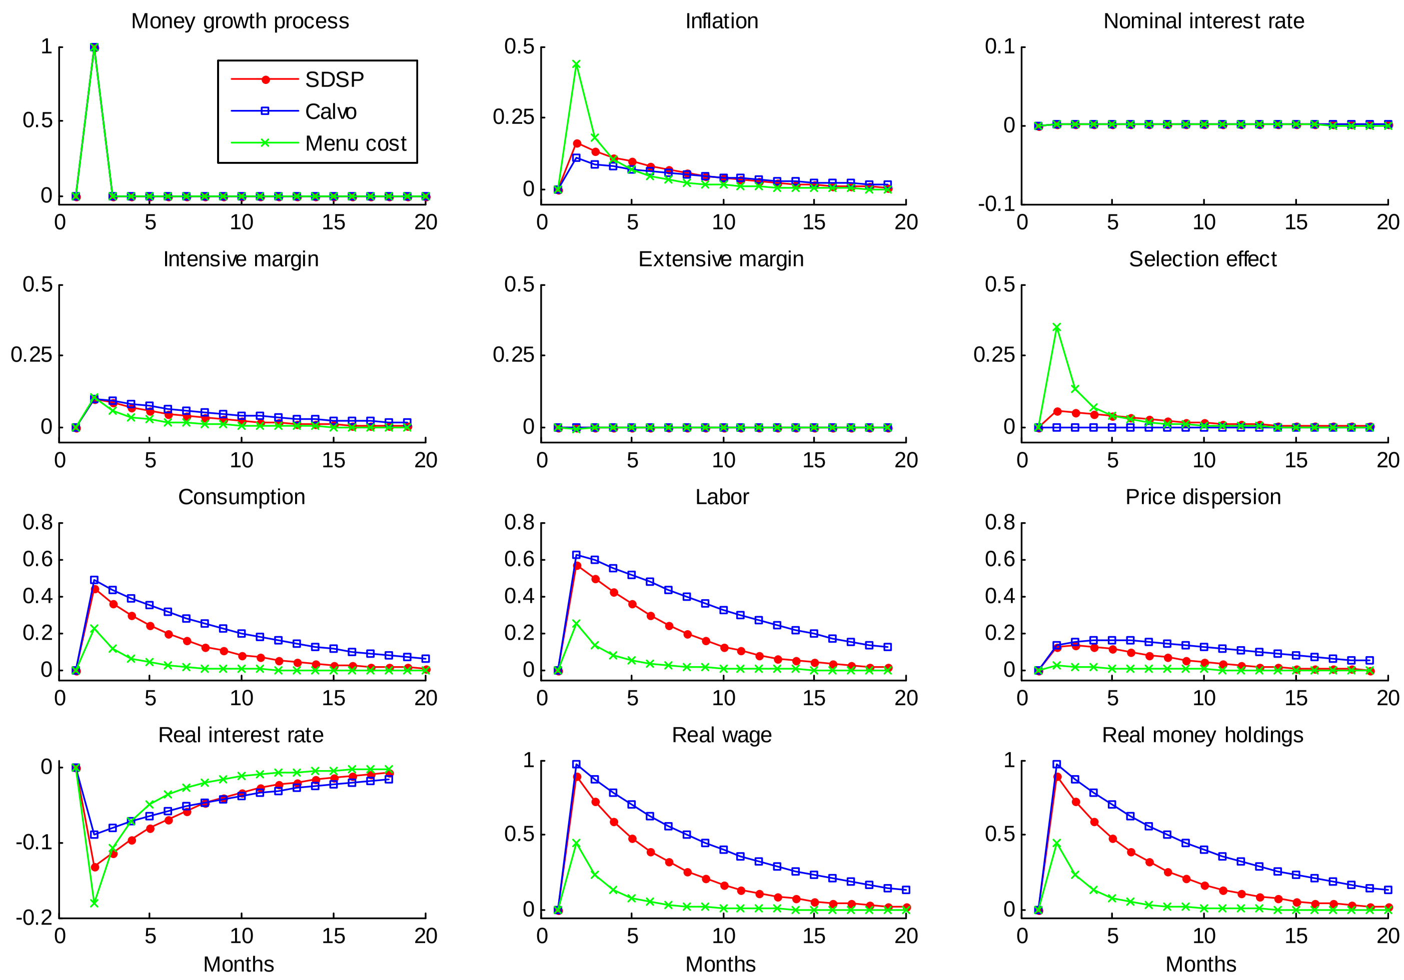

Figure 3. Impulse-response functions: uncorrelated money shock.

Panels 1-6 and 10: percentage points difference from steady state. Rest of panels: percent deviation from steady state

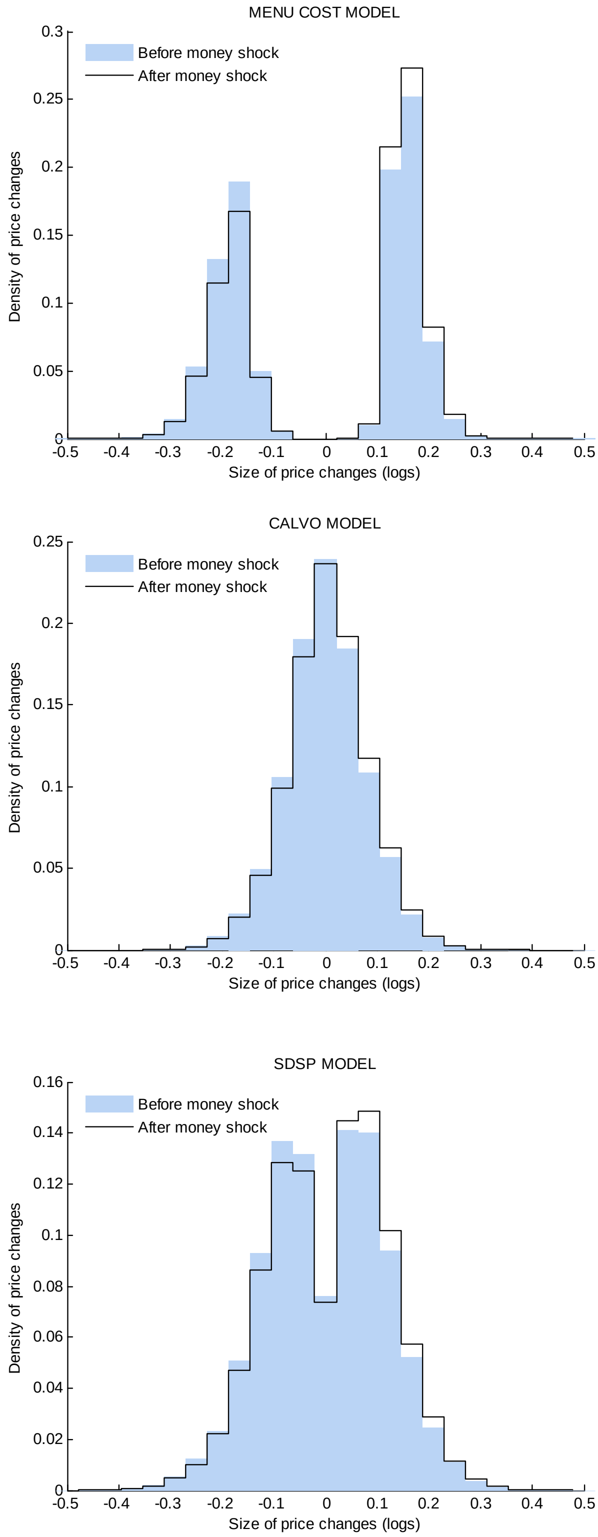

Figure 4. Distribution of price changes in menu cost, Calvo, and SDSP after 50 bp money growth increase.

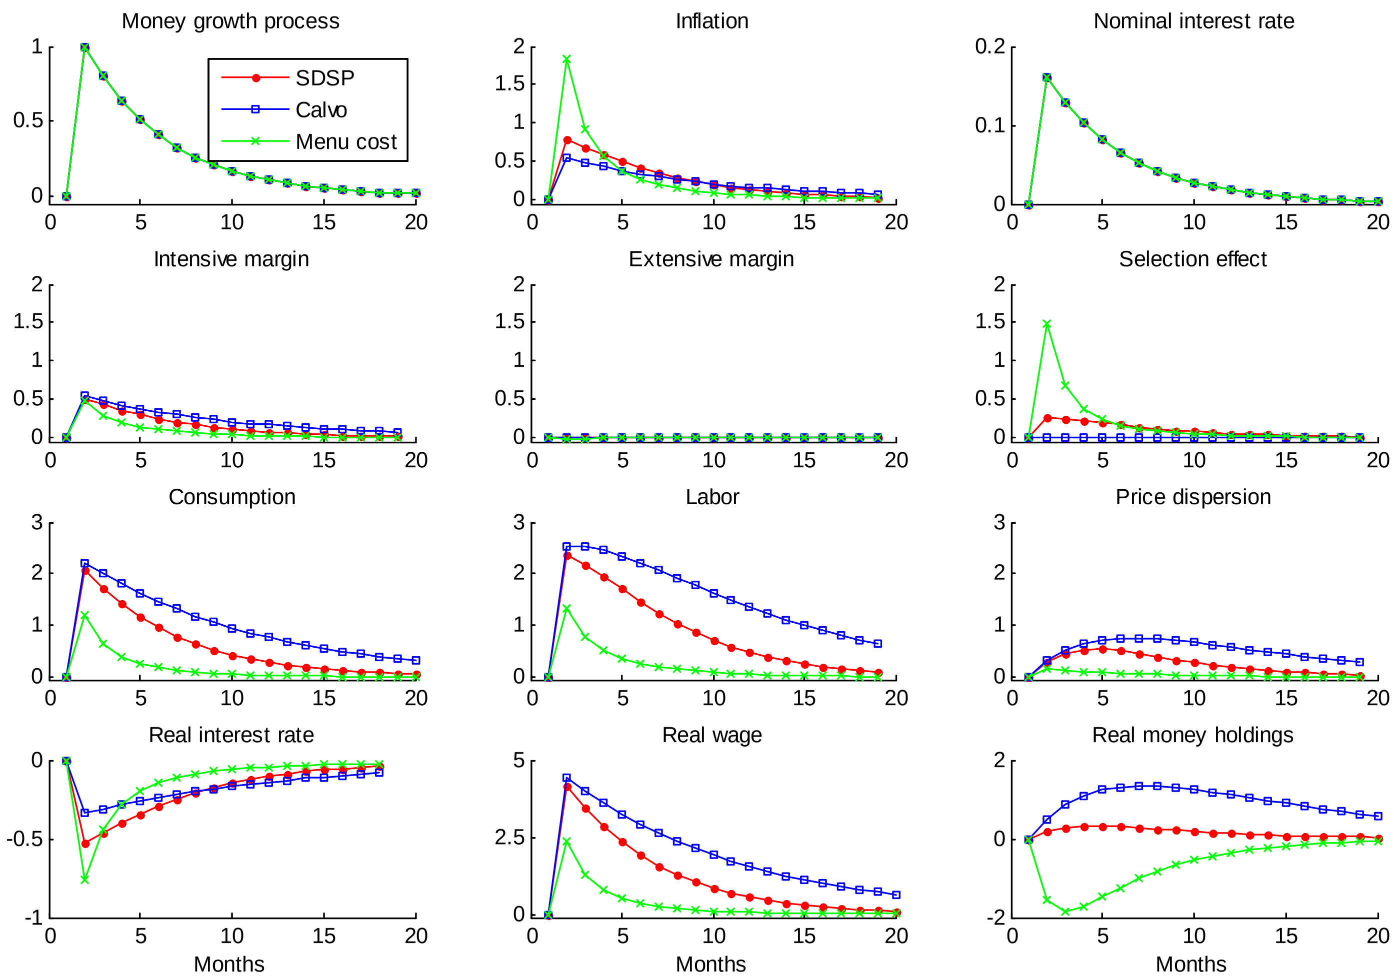

Figure 5. Impulse-response functions: autocorrelated money shock.

Panels 1-6 and 10: percentage points difference from steady state. Rest of panels: percent deviation from steady state

Table 3: Decomposing the initial impact of money shocks

	$\partial\pi_1/\partial\mu_1$	$\mathcal{I}_1$	$\mathcal{E}_1$	$\mathcal{S}_1$
Zero inflation				
Uncorrelated shocks				
Calvo	0.101	0.101	0	0
SDSP	0.157	0.100	-0.000	0.057
Menu cost	0.451	0.104	-0.006	0.353
Correlated shocks				
Calvo	0.548	0.548	0	0
SDSP	0.788	0.516	-0.000	0.272
Menu cost	1.940	0.480	-0.018	1.480
63% annual inflation				
Uncorrelated shocks				
SDSP	0.363	0.243	0.116	0.003
Menu cost	0.763	0.187	0.218	0.358
Correlated shocks				
SDSP	1.770	1.220	0.617	-0.068
Menu cost	3.050	0.862	0.823	1.370

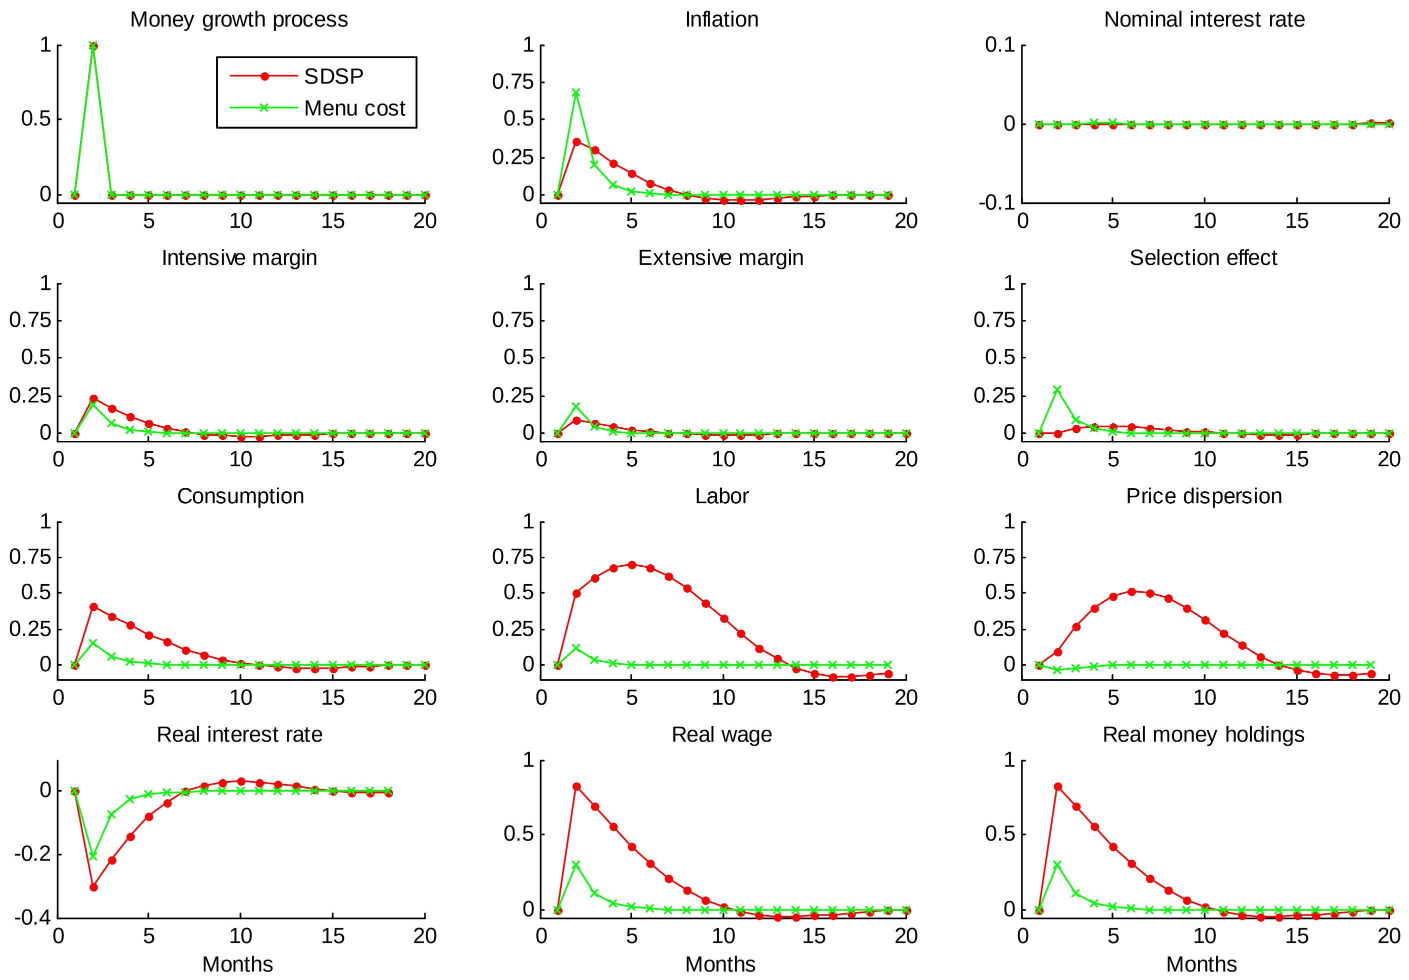

Figure 6. Impulse-response functions: 63% annual inflation.

Panels 1-6 and 10: percentage points difference from steady state. Rest of panels: percent deviation from steady state

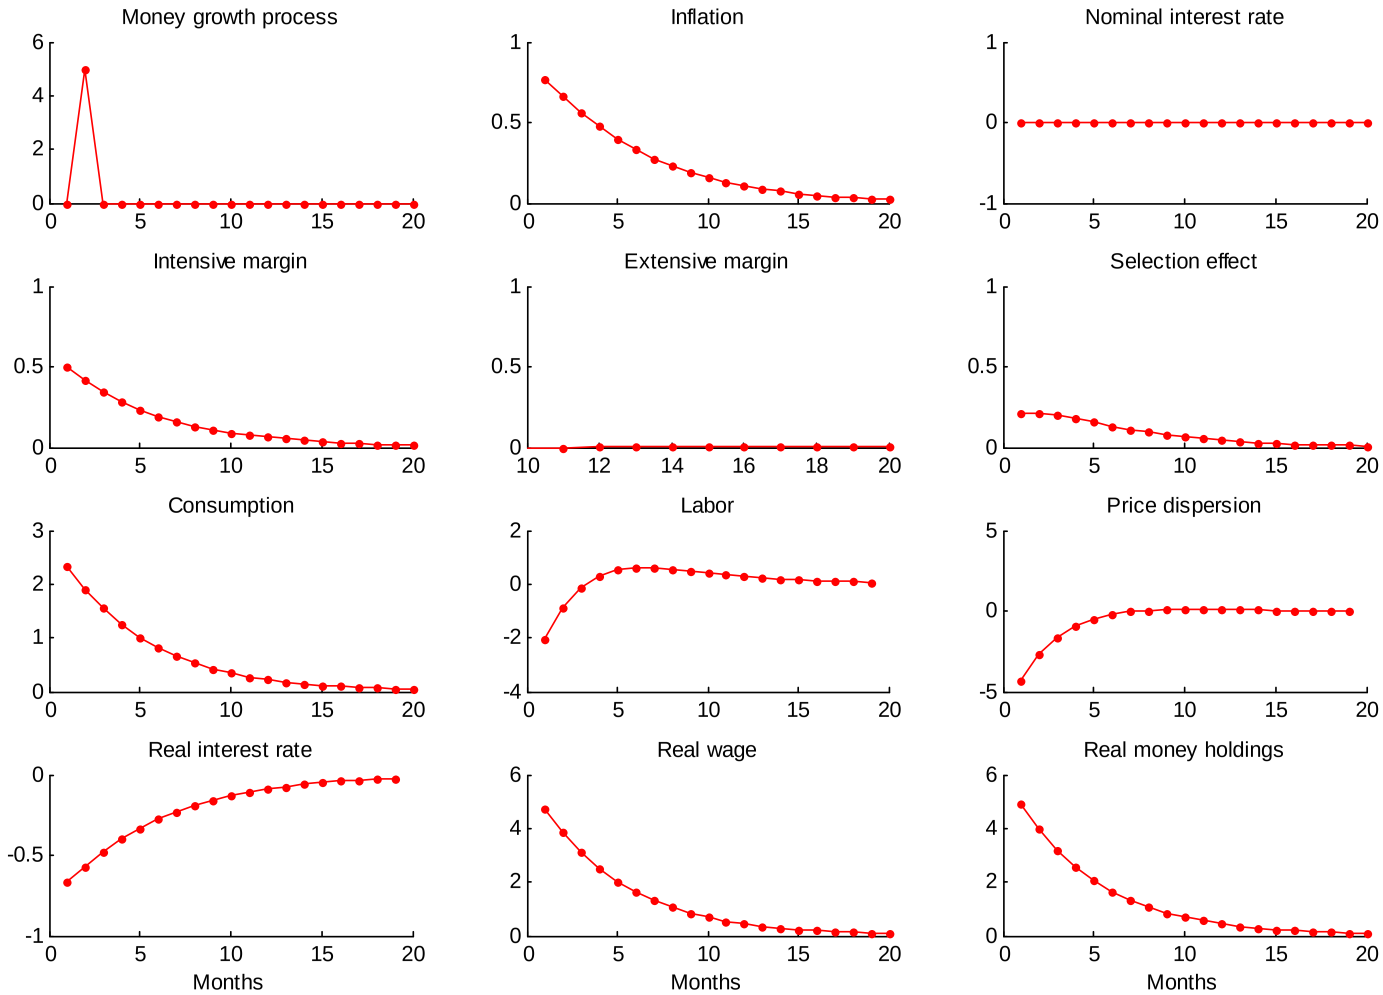

Figure 7.1 Transitional dynamics from shifted distribution: 5% increase in money.

Panels 1-6 and 10: percentage points deviation from steady state. Rest of panels: percent deviation from steady state.

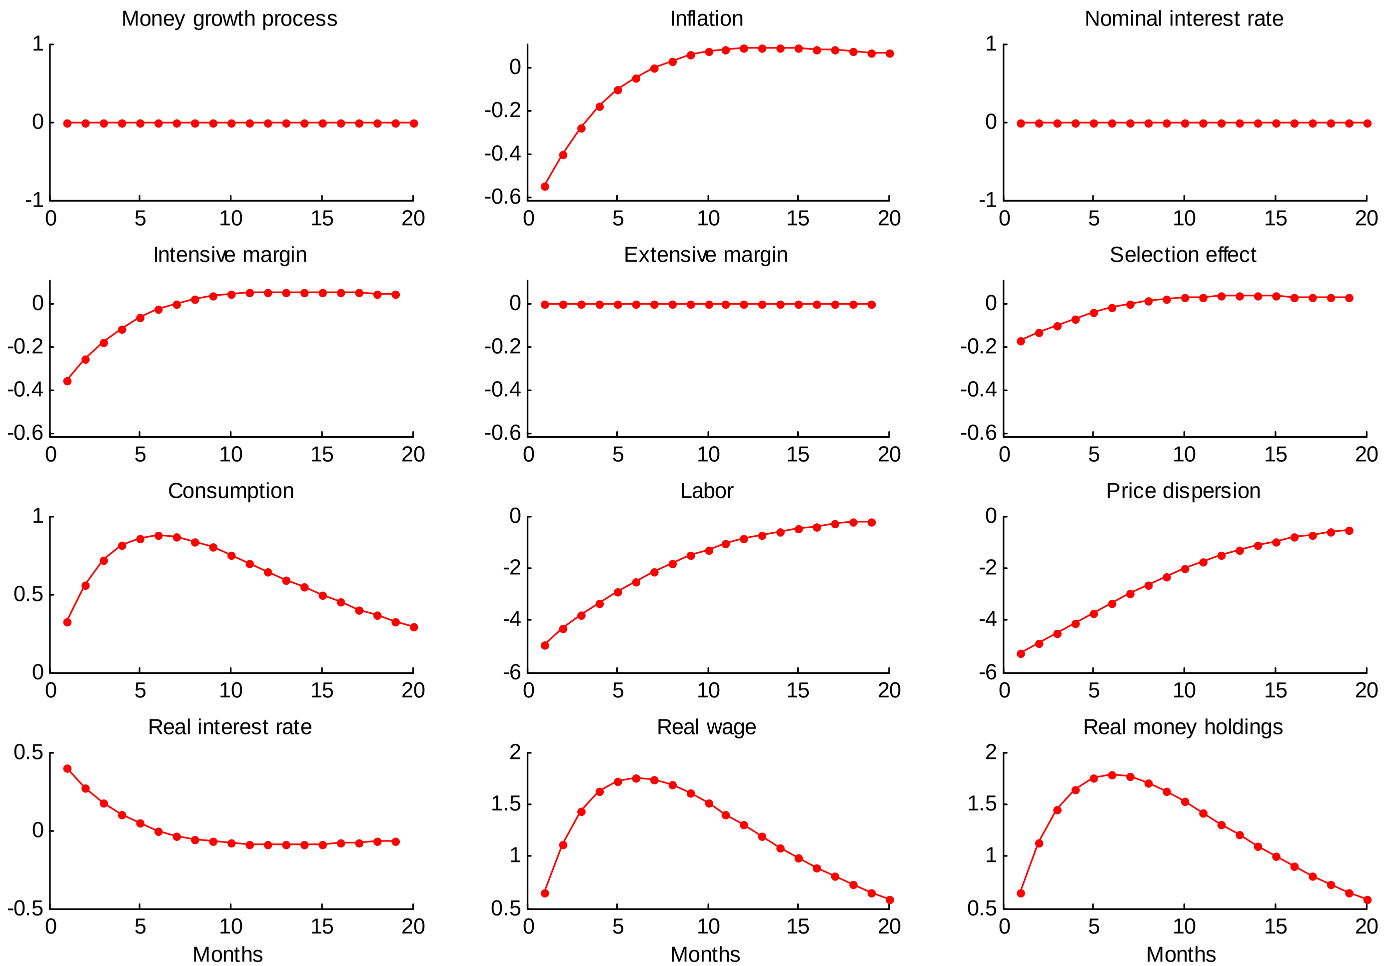

Figure 7.2 Transitional dynamics from shifted distribution: 6% increase in aggregate productivity.
Panels 1-6 and 10: percentage points deviation from steady state. Rest of panels: percent deviation from steady state.

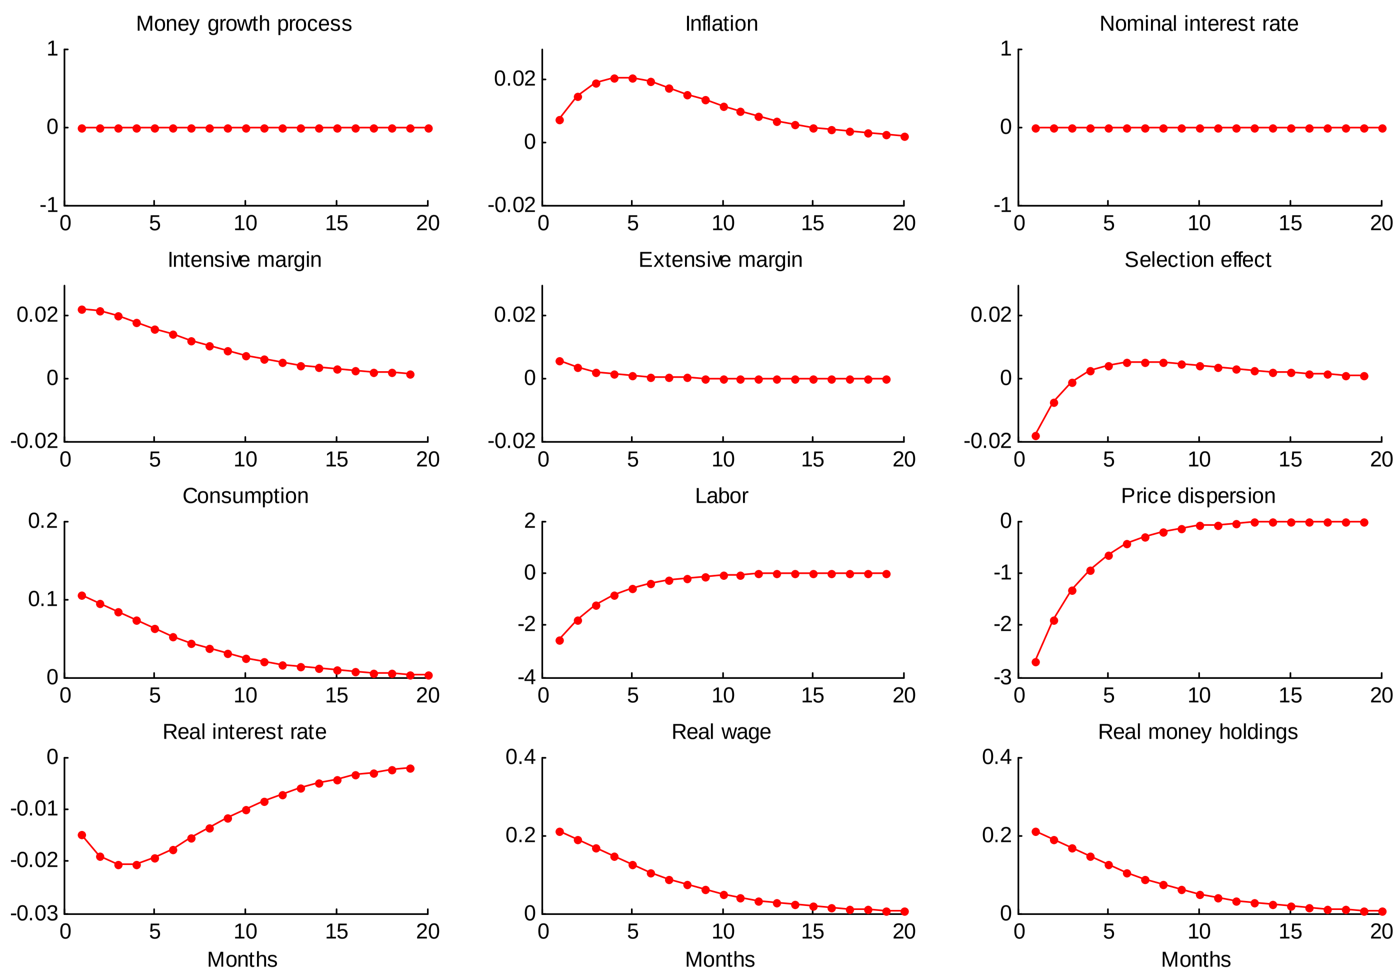

Figure 7.3 Transitional dynamics from shifted distribution: all firms from optimal sticky price.
Panels 1-6 and 10: percentage points deviation from steady state. Rest of panels: percent deviation from steady state.

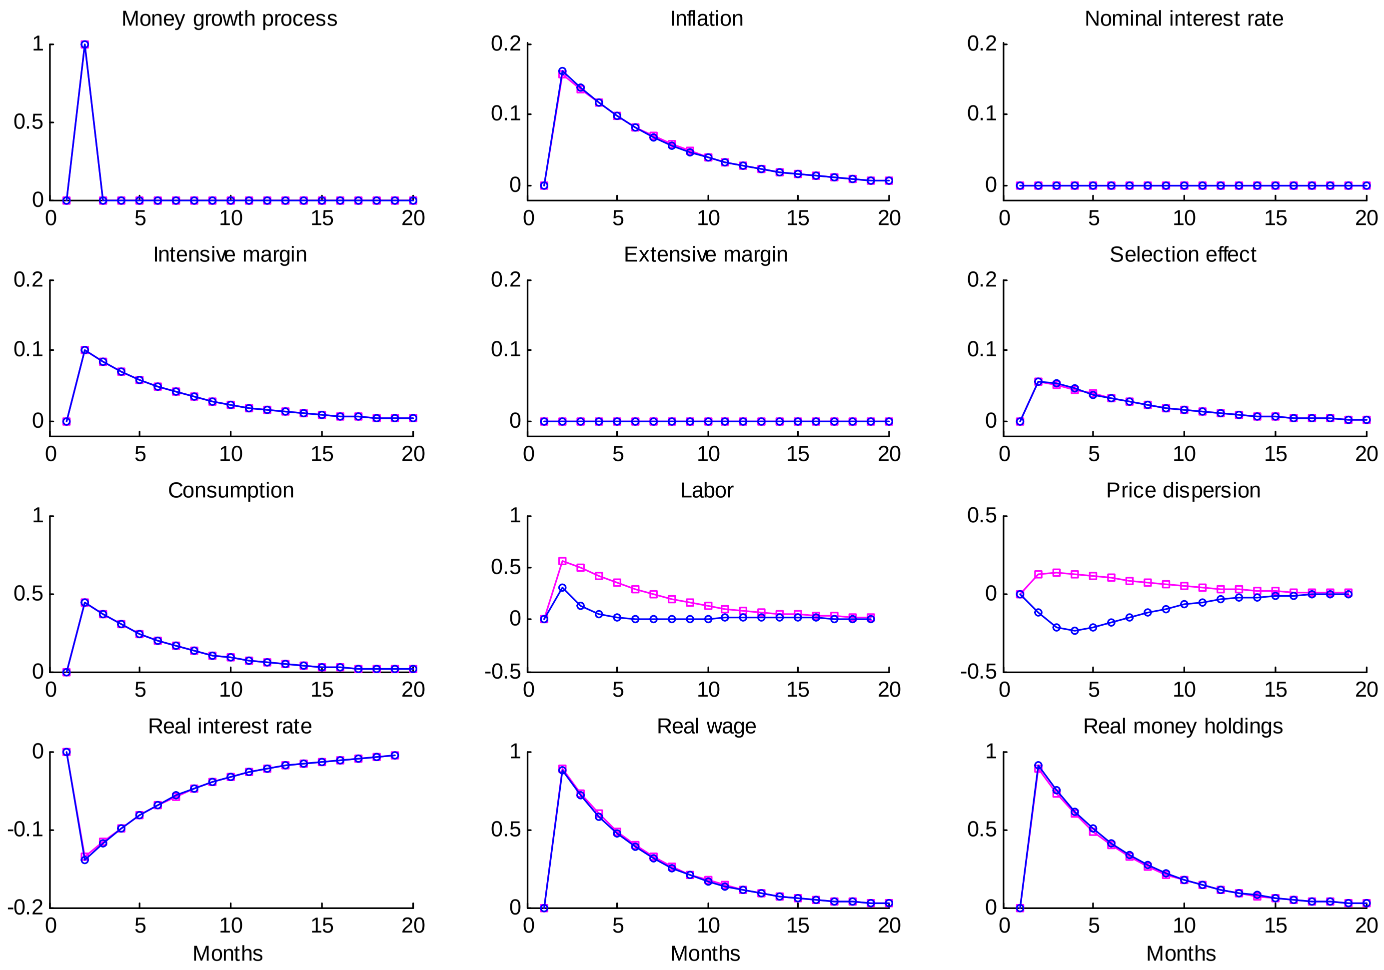

Figure 8. Impulse-responses from different initial distributions.
Lines with red squares: baseline IRFs from fig 3. Lines with blue circles: IRFs ocurring after technology shock.
Panels 1-6 and 10: percentage points deviation from steady state. Rest of panels: percent deviation from steady state.