

Working Papers

RESEARCH DEPARTMENT

WP 19-09

January 2019

<https://doi.org/10.21799/frbp.wp.2019.09>

Investigating Nonneutrality in a State-Dependent Pricing Model with Firm-Level Productivity Shocks

Michael Dotsey

Federal Reserve Bank of Philadelphia

Alexander L. Wolman

Federal Reserve Bank of Richmond

ISSN: 1962-5361

Disclaimer: This Philadelphia Fed working paper represents preliminary research that is being circulated for discussion purposes. The views expressed in these papers are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors or omissions are the responsibility of the authors. Philadelphia Fed working papers are free to download at: <https://philadelphiafed.org/research-and-data/publications/working-papers>.

Investigating Nonneutrality in a State-Dependent Pricing Model with Firm-Level Productivity Shocks*

Michael Dotsey

Alexander L. Wolman

Federal Reserve Bank of Philadelphia Federal Reserve Bank of Richmond[†]

January 2019

Abstract

In recent years there has been an abundance of empirical work examining price setting behavior at the micro level. First generation models with price setting rigidities were generally at odds with much of the micro price data. A second generation of models, with fixed costs of price adjustment and idiosyncratic shocks, have attempted to rectify this shortcoming. Using a model that matches a large set of microeconomic facts we find significant nonneutrality. We decompose the nonneutrality and find that state-dependence plays an important part in the responses of output and inflation to a monetary shock. We also examine how aggregating firm behavior can generate flat hazards. Lastly we find that the steady state statistic developed by Alvarez, Le Bihan, and Lippi (2016) is an imperfect guide to characterizing nonneutrality in our model.

*We are deeply indebted to Bob King, without whose contributions this paper would not exist. We are grateful to Andrew Foerster and Michael Johnston for outstanding research assistance, Thomas Laubach and John Leahy for valuable comments on earlier drafts, and seminar and conference participants at the Federal Reserve Board and the Federal Reserve Banks of New York, Kansas City and Atlanta, Duke University, Goethe University in Frankfurt, the Banque de France, the 2009 SED Conference, the Reserve Bank of New Zealand, Oxford University, the University of Glasgow, and St. Andrews University for their comments and questions. We also wish to thank William Goffe for providing his code for the simulated annealing algorithm used in this paper.

[†]**Disclaimer:** This Philadelphia Fed working paper represents preliminary research that is being circulated for discussion purposes. The views expressed in these papers are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia, the Federal Reserve Bank of Richmond, or the Federal Reserve System. Any errors or omissions are the responsibility of the authors. Philadelphia Fed working papers are free to download at <https://www.philadelphiafed.org/research-and-data/working-papers>.

1 Introduction

In recent years there has been an abundance of empirical work examining price setting behavior at the micro level. In order to match the common features that have been uncovered, economists have developed a second generation of sticky-price models that include both fixed costs of price adjustment and idiosyncratic shocks. Those models have then served as laboratories for studying monetary nonneutrality.¹ Our paper contributes to this literature by constructing a model that matches an extensive set of micro facts and then studying the implications for monetary nonneutrality. In particular, our model matches the following seven facts: (1) there are both large positive and negative price changes that are intermixed with many small price changes; (2) average price changes are an order of magnitude larger than needed to keep up with inflation; (3) many prices are set infrequently, with changes occurring less frequently than once a year, while some prices appear to be completely flexible; (4) aggregate hazards are relatively flat; (5) the idiosyncratic part of price changes is more volatile and less persistent than that due to aggregate factors; (6), the size of a price change is not related to the time since the last price change, and (7) price changes are positively autocorrelated. These are among the facts that lead Klenow and Malin (2011) to conclude that idiosyncratic forces are dominant in accounting for actual pricing behavior and are largely consistent with a state-dependent approach to price setting. Thus, our paper is most closely linked to recent work by Midrigan (2011), Alvarez, Le Bihan, and Lippi (2016), and Karadi and Reiff (2018). Our contribution is that we match the features of the micro price data listed above while working in a relatively standard DSGE framework, albeit one in which there is a finite-state firm-level productivity process.² Because we use perturbation methods, our modeling of state-dependent pricing may lend itself to estimation via Bayesian econometric methods as computational speeds improve.

Many papers have matched subsets of these seven facts, but to our knowledge none have simultaneously matched all of them. Midrigan (2011), Alvarez, Le Bihan, and Lippi (2016), Nakamura and Steinsson (2008), and Karadi and Reiff (2018) all have theoretical environments consistent with attributes (1) and (2), and Midrigan like us also matches the actual histogram of price changes. Nakamura and Steinsson (2010) have multiple sectors with varying degrees of price flexibility and so have a model that is consistent with fact (3) as well. Klenow and Willis (2016) matches the autocorrelation. Notably, we find two model parameterizations consistent with these seven facts, and both imply a significant degree of

¹For example, see Golosov and Lucas (2007), Burstein and Hellwig (2007), Midrigan (2011), Gertler and Leahy (2008), Nakamura and Steinsson (2008,2010), Alvarez and Lippi (2014), Alvarez, Le Bihan, and Lippi (2016), Karadi and Reiff (2018), Vavra (2014), and Luo and Villar (2018) to name a few.

²The model is a direct descendant of the one in Dotsey, King and Wolman (1999).

nonneutrality.

Caplin and Spulber (1987) first showed that when price adjustment is costly, the extent of nonneutrality can be highly sensitive to the details of firm-level heterogeneity. A series of more recent papers have sought to characterize the relationship between firm-level heterogeneity and nonneutrality. Midrigan (2011) was a notable contribution, emphasizing that the presence of both small and large price changes in the micro data works against the selection effect. In both Caplin and Spulber (1987) and Golosov and Lucas (2007) a strong selection effect works against nonneutrality. Alvarez, Le Bihan and Lippi (2016) proved that in a relatively wide class of models, a sufficient statistic for the extent of nonneutrality is the ratio of the kurtosis of the size-distribution of price changes and the average number of price changes per year, both measured in steady state. Our model does not satisfy the conditions for the theorem in that paper, and we find that the theorem does not provide a good approximation to our model's behavior. Karadi and Reiff (2018) also explore the relationship between heterogeneity and nonneutrality using a stochastic process that also departs from Alvarez, Le Bihan and Lippi's assumptions, but nests the processes in Golosov and Lucas (2007) and Midrigan (2011). They are able to generate significant nonneutrality, as are Costain and Nakov (2011a, 2011b) who use scanner-level data to calibrate their model. Finally, Luo and Villar (2018), explore the relationship between nonneutrality and the skewness and dispersion of price-changes in a theoretical framework that is akin to ours.

Capturing the rich distribution of price changes requires us to embed significant heterogeneity across firms. As is standard in the literature, we assume that firms face idiosyncratic shocks to their productivity or marginal cost. However, unlike much of that literature, here productivity follows a finite-state Markov process. This approach, while less parsimonious than some others, has a high payoff in terms of allowing the model to match the many facts listed above. And, it is easily incorporated into our DSGE framework. We also incorporate a fraction of firms that set prices flexibly, which helps the model account for the many small price changes in the data. As in Dotsey, King and Wolman (1999), firms face idiosyncratic variation in fixed costs of price adjustment. We calibrate the model to match the distribution of price changes as well as the average duration of prices documented in Klenow and Kryvtsov (2008) and the positive autocorrelation of price changes discussed in Klenow and Willis (2016). It matches these facts with firms incurring relatively small menu costs. An outgrowth of this calibration is that the model is consistent with the other pricing facts listed above. We are also able to identify features that are essential to finding flat aggregate hazard functions. Notably, we find that the micro pricing features are consistent with both flat hazards and a significant amount of nonneutrality, with the degree of nonneutrality roughly that found in the data and similar to that produced by a Calvo model (see Costain and

Nakov (2011a, 2011b). As mentioned, we find more than one specification of our model can match the steady-state distribution of price changes and the duration of prices as estimated by Klenow and Kryvtsov (2008). The two parameterizations illustrate that in matching the steady-state distribution, there is a trade-off between the fraction of sticky-price firms and the menu costs incurred by those firms.

We begin in Section 2 with a description of the model, paying particular attention to the details of the individual firm’s pricing decisions. In Section 3, we present our steady-state results, which match the facts about price changes listed above. These facts can be matched with either a relatively small or large fractions of flexible-price firms. Section 4 discusses the dynamics in response to monetary shocks. We do this by following up on work by Klenow and Kryvtsov (2008), Caballero and Engel (2007) and Costain and Nakov (2011a), who present various decompositions of the dynamics of the price level in response to a monetary shock. Here we find that the state-dependent elements of pricing are indeed important for the model’s behavior. In Section 5, we examine the model’s ability to generate flat or downward-sloping aggregate hazards even though the state-specific hazards are upward sloping. In Section 6 we report how our model relates to the striking result in Alvarez, Le Bihan and Lippi (2016). Section 7 concludes.

2 The Model

The basic elements of the model are drawn from the state-dependent pricing framework of Dotsey, King and Wolman (1999), or DKW for short, with the addition of stochastic variation in productivity at the firm level. We refer to the different productivity levels as microstates. Firms are heterogeneous with respect to productivity realizations but share a common stochastic process.³ As in the DKW framework, some firms are heterogeneous with respect to the realization of fixed costs, but we also allow for a simple form of heterogeneity in the distribution of fixed costs: a fraction of firms are flexible-price firms, facing zero adjustment costs. After analyzing the stochastic adjustment framework, we describe the standard features of the rest of the model.

³The framework can easily incorporate firm-specific demand disturbances. However, we focus on productivity shocks because they are the main focus of the related literature and because studies such as Boivin, Giannoni, and Mihov (2009) indicate that much of the idiosyncratic variation can be attributed to supply-type shocks.

2.1 A stochastic adjustment framework

To study the effect of serially correlated firm-level shocks along with heterogeneity in adjustment costs, we extend the DKW framework while retaining much of its tractability.

Heterogeneous productivity: The relative productivity level, e_t , of a firm is its microstate. We assume that there are K different levels of the microstate that may occur, \underline{e}_k , $k = 1, 2, \dots, K$, so that a firm of type k at date t has total factor productivity

$$a_{kt} = a_t \underline{e}_k \tag{1}$$

where a_t is a common productivity shock. We assume that the relative productivity levels are ordered so that $\underline{e}_1 < \underline{e}_2 < \dots < \underline{e}_K$.

Transitions between microstates are governed by a state transition matrix, Q , where

$$q_{kf} = \text{prob}(e_{t+1} = \underline{e}_f | e_t = \underline{e}_k). \tag{2}$$

We also use the notation $q(e'|e)$ to denote the conditional probability of state e' occurring next period if the current microstate is e . We assume that these transitions are independent across firms and that there is a continuum of firms, so that the law of large numbers applies. The stationary distribution of firms across microstates is then given by a vector Φ such that

$$\Phi = Q^T * \Phi.$$

That is, the k^{th} element of Φ (denoted ϕ_k) gives the fraction of firms in the k^{th} microstate (these firms have productivity level \underline{e}_k).⁴

Heterogeneous adjustment costs: A fraction f of firms face zero adjustment costs and thus have perfectly flexible prices. In each period, the remaining fraction $1 - f$ of firms draw from a nondegenerate distribution $G()$ of fixed costs. Because the fraction and the identity of flexible-price firms are fixed over time, the only interaction between the two types arises from general equilibrium considerations. We focus on the details of optimal pricing and adjustment for the firms that do face adjustment costs and then briefly discuss the flexible-price firms.

⁴The stationary probability vector can be calculated as the eigenvector associated with the unit eigenvalue of the transpose of Q . See, for example, Kemeny and Snell (1976).

2.1.1 Optimal pricing and adjustment

Consider a firm that faces demand $d(p, s)y(s)$ for its output if it is charging relative price p in aggregate state s . Following DKW, we assume that this firm faces a fixed labor cost of adjustment ξ , which is i.i.d. and is drawn from a continuous distribution with support $[0, \mathcal{B}]$.

It is convenient to describe the optimal adjustment using three value functions: \underline{v} , the value of the firm if it does not adjust; v^o , its value if it has a current adjustment cost realization of 0; and v , its maximized value given its actual adjustment cost ξ .

In terms of these value functions, the market value of a firm is governed by

$$v(p, e, s, \xi) = \max \{ \underline{v}(p, e, s), [v^o(e, s) - \lambda(s)w(s)\xi] \} \quad (3)$$

if it has a relative price of p , is in microstate e , aggregate state s , and draws a stochastic adjustment cost of ξ . In (3), $\lambda(s)$ is the marginal value of state contingent cash flow and $w(s)$ is the real wage in state s .⁵ The market value, v , is then the maximum of two components, one being the value conditional on adjustment ($[v^o(e, s) - \lambda(s)w(s)\xi]$) and the other the value of nonadjustment ($\underline{v}(p, e, s)$).

Defining the real flow of profits as $z(p, e, s)$, the nonadjustment value \underline{v} – the value of continuing with the current relative price p for at least one more period – obeys the Bellman equation,

$$\underline{v}(p, e, s) = [\lambda(s)z(p, e, s) + \beta E v(p', e', s', \xi') | (p, e, s)] \quad (4)$$

with $p' = p/\pi(s')$ and $\pi(s')$ denoting the gross inflation rate in the next period. Therefore, the value \underline{v} depends on the current relative price p , marginal utility $\lambda(s)$, the flow of real profits $z(p, e, s)$, and the discounted expected future value $\beta E v(p', e', s', \xi')$.⁶

The “costly adjustment value” is given by the value of the firm if it is free to adjust, $v^o(e, s)$, less the cost of adjustment, which depends on the macro state through $\lambda(s)w(s)$ and also on the realization of the random adjustment cost ξ . In turn, the “free adjustment value” v^o obeys

$$v^o(e, s) = \max_{p^*} \{ \lambda(s)z(p^*, e, s) + \beta E v(p', e', s', \xi') | p^*, e, s \} \quad (5)$$

with $p' = p^*/\pi(s')$.

⁵The aggregate state, s , will be determined as part of the general equilibrium but can be left unspecified at present. Note that the values are in marginal utility units, which can be converted into commodity units by dividing through by marginal utility.

⁶In forming conditional expectations it is not necessary to condition on ξ because it is i.i.d.

As in other generalized partial adjustment models, the firm adjusts if

$$[v^o(e, s) - \lambda(s)w(s)\xi] > \underline{v}(p, e, s).$$

Accordingly, there is a threshold value of the adjustment cost, such that

$$\underline{\xi}(p, e, s) = \frac{v^o(e, s) - \underline{v}(p, e, s)}{\lambda(s)w(s)}. \quad (6)$$

Firms facing a lower adjustment cost adjust. Those with a higher cost charge the same nominal price as they did last period, implying that their relative price moves in the opposite direction of the price level.

The fraction of (non-flexible-price) firms that adjusts, then, is

$$\alpha(p, e, s) = G(\underline{\xi}(p, e, s)) \quad (7)$$

where G is the cumulative distribution of adjustment costs. The adjustment decision depends on the state of the economy, but it also depends on the distribution of adjustment costs in two ways. First as highlighted by (7), the adjustment cost distribution governs the fraction of adjusting firms given the threshold. Second, stochastic adjustment costs provide the firm with an incentive to wait for a low adjustment cost realization: the adjustment cost distribution affects this incentive and thus the adjustment threshold.

We study an environment with positive trend inflation. Because adjustment costs are bounded above and inflation continually erodes a firm's relative price if it does not adjust, there will exist a maximum number of periods over which a firm may choose not to adjust its price. Because historical prices depend on the state h that a firm was in when it last adjusted, as well as the current state k , this maximal number, J_{hk} , will depend on both h and k in steady state. Thus, the stationary distribution of firms in a steady state equilibrium takes on a finite but potentially large number of elements.

2.1.2 Flexible-price firms

For the fraction ω^f of firms with zero adjustment costs, things are much simpler. These flexible-price firms face the same productivity process as other firms and simply charge the price in state k that maximizes profits *in state k*. Thus, they set a price, p_k^f , that is a constant markup over their marginal cost, which depends on both their microstate and the aggregate state of the economy.

2.1.3 Dynamics and accounting

We need to track the distribution of firms over the set of historical prices and the evolving levels of micro productivity. The core mechanics are as follows. We start with a joint distribution of relative prices and productivity that prevailed last period. This distribution is then influenced by the effects of microproductivity transitions (Q), the adjustment decisions of firms ($\alpha(p, e, s)$); and the effects of inflation on relative prices. The net effect is to produce a new distribution of relative prices prevailing in the economy. Table 1 summarizes some of the key notation and equations.

Initial conditions and sticky prices. Let $p_{j-1,h,t-1}$ be the previous period's relative price of a firm that last changed its price at date $t-1-(j-1) = t-j$ when it was in microstate h . Let $\omega_{j-1,h,l,t-1}$ be the fraction of (sticky-price) firms in this situation, that charged this price and also had productivity level l . Ranging over all the admissible j, h, l , this information gives the joint distribution of productivity and prices at date $t-1$.

If such a firm chooses not to adjust, its relative price evolves according to

$$p_{j,h,t} = p_{j-1,h,t-1} / \pi_t \tag{8}$$

where π_t is the current inflation rate (π_t is shorthand for $\pi(s_t)$ from above). That is, one effect on the date t distribution of relative prices comes from inflation.

Endogenous fractions: Two micro-disturbances hit a firm in our model, so that its ultimate decisions are conditioned on its productivity (e) and its adjustment cost (ξ). For the purpose of accounting and ease of presentation, it is convenient to specify that the productivity shock occurs first and the adjustment cost shock second, but doing so has no substantive implications.

As above, let $\omega_{j-1,h,l,t-1}$ be the fraction of (sticky-price) firms that charged the price $p_{j-1,h,t-1}$ when they were in microstate l last period. Let $q_{l,k}$ denote the probability of transitioning from microstate l to microstate k . As a result of stochastic productivity transitions there will be a fraction

$$\theta_{j,h,k,t} = \sum_l q_{l,k} \omega_{j-1,h,l,t-1}$$

of sticky-price firms that start period t with a j -period old nominal price set in microstate h and have a microstate k in the current period.

However, not all of these firms will continue to charge the nominal price that they set in the past. To be concrete, consider firms with a j -period old price set in microstate h that are now in state k . Of these firms, the adjustment rate is $\alpha_{j,h,k,t}$. Then, the fraction of

sticky-price firms choosing to continue charging the nominal price set j periods ago will be

$$\omega_{j,h,k,t} = (1 - \alpha_{j,h,k,t})\theta_{j,h,k,t} \quad (9)$$

Taking all of these features into account, we can see that transitions are governed by two mechanisms: the exogenous stochastic transitions of the microstates ($q_{l,k}$) and the endogenous adjustment decisions of firms ($\alpha_{j,h,k,t}$). As discussed, the adjustment decision depends on the firm's relative price, its microstate and the macroeconomic states.

Given that firms currently in microstate k adjust from a variety of historical states, it follows that the fraction of adjusting firms is given by

$$\omega_k^f + \omega_{0,k,k,t} = \omega^f \phi_k + \sum_j \sum_h \alpha_{j,h,k,t} \theta_{j,h,k,t}. \quad (10)$$

We use the redundant notation $\omega_{0,k,k,t}$ to denote the fraction of sticky-price firms that adjust in microstate k so that this is compatible with (9), and we use the notation ω_k^f to denote the fraction of firms that have flexible prices and are in state k . Since the distribution of microstates is assumed to be stationary, there is a constraint on the fractions,

$$\phi_k = \omega^f \phi_k + \left(\omega_{0,k,k,t} + \sum_h \sum_j \omega_{j,h,k,t} \right) (1 - \omega^f).$$

Equivalently, a fraction ϕ_k of both flexible- and sticky-price firms are in state k at the end of the period.

2.1.4 State variables suggested by the accounting

As suggested by the discussion above, there are two groups of natural endogenous state variables in the model. One is the vector of past relative prices $p_{j-1,h,t-1}$ for $h = 1, 2, \dots, K$ and $j = 1, 2, \dots, J_h$. The other is the fraction of sticky-price firms that enter the period with a particular past microstate (l) and a relative price that was set j periods ago in microstate h ,

$$\omega_{j-1,h,l,t-1}. \quad (11)$$

Thus, the addition of microstates raises the dimension of the minimum state space from roughly $2 * J$ to roughly $J * K + J * K^2$, where J is the maximum number of periods of nonadjustment and K is the number of microstates. However, this is only an approximation because the maximum number of periods can differ across microstates: J_{kh} is the endpoint suitable for firms currently in microstate k that last adjusted in historical microstate h .

2.1.5 The adjustment process

Recall, that for each price lag $(j - 1)$, microstate last period (l) and historical state (h) , a fraction $\omega_{j-1,h,l,t-1}$ of firms enters the period. Then, the microstate transition process leads to a fraction of sticky-price firms $\theta_{j,h,k,t} = \sum_l q_{l,k} \omega_{j-1,h,l,t-1}$ having a price lag j , a current microstate k ; and a historical state (h) . Of these firms, a fraction $\alpha_{j,h,k,t}$ chooses to adjust while a fraction $\eta_{j,h,k,t} = 1 - \alpha_{j,h,k,t}$ chooses not to adjust, leaving $\omega_{j,h,k,t} = \eta_{j,h,k,t} \theta_{j,h,k,t}$ charging relative price $p_{j,h,t}$ and experiencing microstate k . One thing that is important to stress at this stage is that we allow for zero adjustment or for complete adjustment in various situations (particular j, k, h entries).

2.2 The DSGE model

We now embed this generalized partial adjustment apparatus into a particular DSGE model, which is designed to be simple on all dimensions other than pricing.

2.2.1 The household

As is conventional, there are two parts of the specification of household behavior, aggregates and individual goods. We assume that there are many identical households that maximize

$$\begin{aligned} \max_{c_t, n_t} E_0 \left\{ \sum_t \beta^t \left[\frac{1}{1-\sigma} c_t^{1-\sigma} - \frac{\chi}{1+\gamma} n_t^{1+\gamma} \right] \right\} \\ \text{subject to: } c_t \leq w_t n_t + (1 - \omega^f) \sum_j \sum_h \sum_k \omega_{j,h,k,t} z_{j,h,k,t} + \omega^f \sum_k \phi_k z_{k,t}^f, \end{aligned}$$

where c_t and n_t are consumption and labor effort, respectively, $z_{j,h,k,t}$ is the profits remitted to the household by a type (j, h, k) firm, and $z_{k,t}^f$ is the profits of a flexible price firm in state k . The consumption aggregate, c , is given by the standard Dixit-Stiglitz demand aggregator. Thus, $c = \left(\int_0^1 (y(i))^{\frac{\epsilon-1}{\epsilon}} di \right)^{\frac{\epsilon}{\epsilon-1}}$. There is an economy-wide factor market for the sole input, labor, which earns a wage, w . The first-order condition determining labor supply is

$$\lambda_t w_t = \chi n_t^\gamma, \tag{12}$$

and, hence, γ^{-1} is the Frisch labor supply elasticity. The first-order condition determining consumption is

$$c_t^{-\sigma} = \lambda_t \tag{13}$$

where λ_t is the multiplier on the household's budget constraint, which serves also to value the firms.

2.2.2 Firms

Two aspects of the firms' specification warrant discussion. First, we adopt a simple production structure, but we think of it as standing in for some of the elements in the "flexible supply side" model of Dotsey and King (2006). Thus, production is linear in labor, $y(i) = a(i)n(i)$, where $y(i)$ is the output of an individual firm, $a(i)$ is the level of its technology, and $n(i)$ is hours worked at a particular firm. Hence, real marginal cost, ψ_t , is given by $\psi_t = w_t/(a_t e_k)$ for a firm that is in microstate k at date t .

$$0 = \lambda(s)z_p(p^*, e, s) + \beta\eta E[\underline{v}_p(p', e', s')], \quad (14)$$

with $p' = p^*/\pi(s')$ and the nonadjustment probability being $\eta(p, e, s) = 1 - \alpha(p', e', s')$.

The marginal value for a nonadjusting firm is

$$\underline{v}_p(p, e, s) = \lambda(s)z_p(p, e, s) + \beta E[\eta(p', e', s')\frac{1}{\pi(s')}\nu_p(p', e', s')], \quad (15)$$

with $p' = p/\pi(s')$.⁷

For the fraction, ω^f , of firms with flexible prices, the profit maximization problem yields

⁷Maximizing the "free adjustment value" (5) implies a first-order condition,

$$0 = \lambda(\varsigma)z_p(p^*, v, \varsigma) + \beta E[\frac{1}{\pi(\varsigma')}v_p(\frac{p^*}{\pi(\varsigma')}, v', \varsigma', \xi')].$$

The value function v takes the form

$$v(p, v, \varsigma, \xi) = \left\{ \begin{array}{ll} \underline{v}(p, v, \varsigma) & \text{if } \xi \geq \bar{\xi}(p, v, \varsigma) \\ [v^o(v, \varsigma) - \lambda(\varsigma)w(\varsigma)\xi] & \text{if } \xi \leq \bar{\xi}(p, v, \varsigma) \end{array} \right\}$$

so that

$$v_p(p, v, \varsigma, \xi) = \left\{ \begin{array}{ll} \underline{v}_p(p, v, \varsigma) & \text{if } \xi \geq \bar{\xi}(p, v, \varsigma) \\ 0 & \text{if } \xi \leq \bar{\xi}(p, v, \varsigma) \end{array} \right\}.$$

Since \underline{v}_p does not depend on ξ , we can express the FOC as in the text. A similar line of reasoning leads to condition (15). We use this first-order approach as an element of producing candidate steady-state equilibria. We then confirm that the candidate is indeed an equilibrium, by checking that adjusting firms' behavior satisfies (5). John and Wolman (2008) discuss the possibility that solutions to the first order conditions may not satisfy (5).

the static first-order condition:

$$z_p^f(p^{*,f}, e, s) = 0.$$

2.2.3 Money and equilibrium

To close the model, it is necessary to specify money supply and demand and to detail the conditions of macroeconomic equilibrium. We impose the money demand relationship $M_t/P_t = c_t$. Ultimately, the level of nominal aggregate demand is governed by this relationship along with the central bank's supply of money. The model is closed by assuming that nominal money supply growth follows an autoregressive process,

$$\Delta \log(M_t) = \rho \Delta \log(M_{t-1}) + x_{mt},$$

where x_{mt} is a mean zero random variable.

There are three conditions for equilibrium. First, labor supply is equal to labor demand, which is a linear aggregate across all the production input requirements of firms plus labor for price adjustment:

$$n_t = (1 - \omega^f) \sum_j \sum_h \sum_k \omega_{j,h,k,t} n_{j,h,k,t} + \omega^f \sum_k \phi_k n_{k,t}^f + (1 - \omega^f) \sum_j \sum_h \sum_k \omega_{j,h,k,t} E \xi_{j,h,k,t}.$$

In this expression, $n_{j,h,k,t}$ is the labor used in production in period t by a sticky-price firm in state k that last adjusted its price in period $t - j$ when it was in state h ; $n_{k,t}^f$ is the labor used in production in period t by a flexible-price firm in state k ; and $E \xi_{j,h,k,t}$ is expected price adjustment costs for a firm in period t in state k that last adjusted its price in period $t - j$ when it was in state h :

$$E \xi_{j,h,k,t} = E (\xi | \xi < G^{-1}(\alpha_{j,h,k,t})).$$

The second equilibrium condition is that consumption must equal output, and the third is that money demand must equal money supply.

2.3 Calibration of macroeconomic parameters

Typically, we will be calibrating at a monthly frequency. Our benchmark settings for the macroeconomic parameters are as follows: $\beta = .97^{1/n}$ where n is the number of periods in a year. The steady-state annualized inflation is 2.5%. The coefficient of relative risk aversion, $\sigma = .25$, labor supply elasticity γ^{-1} is 20, and the demand elasticity ε is 5. The potential for any type of endogenous propagation in this type of model is closely related to the elasticity

of marginal cost with respect to output. Thus, low labor supply elasticities or coefficients of relative risk aversion greater than one severely compromise persistence in this simple and stark setting.⁸ Dotsey and King (2005, 2006) explore many features of more sophisticated models that are capable of generating low marginal cost responses to output. We view our parameter settings as stand-ins for the more realistic persistence-generating mechanisms that are present in larger models. For robustness, we will also look at how setting $\sigma = 1$ or 2 affects some of our results.

3 Benchmark Parameterizations and Properties of our Steady-State Price Distribution

We choose the parameters for our benchmark cases so that the steady-state distribution of price changes from the model is consistent with the facts reported by Klenow and Kryvtsov (2008). The critical parameters are those governing the distributions of fixed costs and firm-level productivity.

3.1 Choosing parameters

Klenow and Kryvtsov (2008), referred to as KK below, report on the distribution of price changes from individual price data underlying the U.S. CPI from 1988 to 2004. Figure 1, uses KK’s data on standardized non-zero price changes and displays a seven bin histogram (the figure is analogous to their Figure 3) where the set of regular price changes is meant to exclude sales.⁹ Standardized price changes are the preferred measure for comparing the data to our one-good economy as the measure controls for different behavior of various categories of goods (see Alvarez, Le Bihan and Lippi (2016) for further discussion). As Midrigan (2011) and Guimaraes and Sheedy (2011) argue, the existence of sales has little relevance for the nonneutrality of money, implying that for our purposes the emphasis in terms of matching the micro-price data should be placed on matching regular price changes. Other relevant price change statistics reported by Klenow and Kryvtsov (2008) are listed

⁸Combining the household’s first-order conditions and approximating the wage by marginal cost and labor and consumption by output yields

$$d \ln mc / d \ln y = \gamma + \sigma,$$

which indicates that values of $\sigma > 1$ preclude low elasticities of marginal cost with respect to output.

We should note that our labor supply elasticity while large is less than the infinite elasticity that is commonly used for example in Midrigan (2011), Costain and Nakov (2011a), or Karadi and Reiff (2018).

⁹All prices and changes in prices are expressed in logs and log deviations.

in the first column of Table 2. In choosing parameters, we target the histogram in Figure 1, the median adjustment probability of 0.139 from Table 2, and the autocorrelation of price changes (0.32) reported in Klenow and Willis (2016). Note that while the histogram represents a direct measurement of price changes, in calculating adjustment probabilities KK impose a Calvo pricing structure on the data. For each category of goods and services they assume that there is a fixed adjustment probability. They then use the price data to compute maximum likelihood estimates of the adjustment probabilities for each category. The median adjustment probability across categories is 0.139.

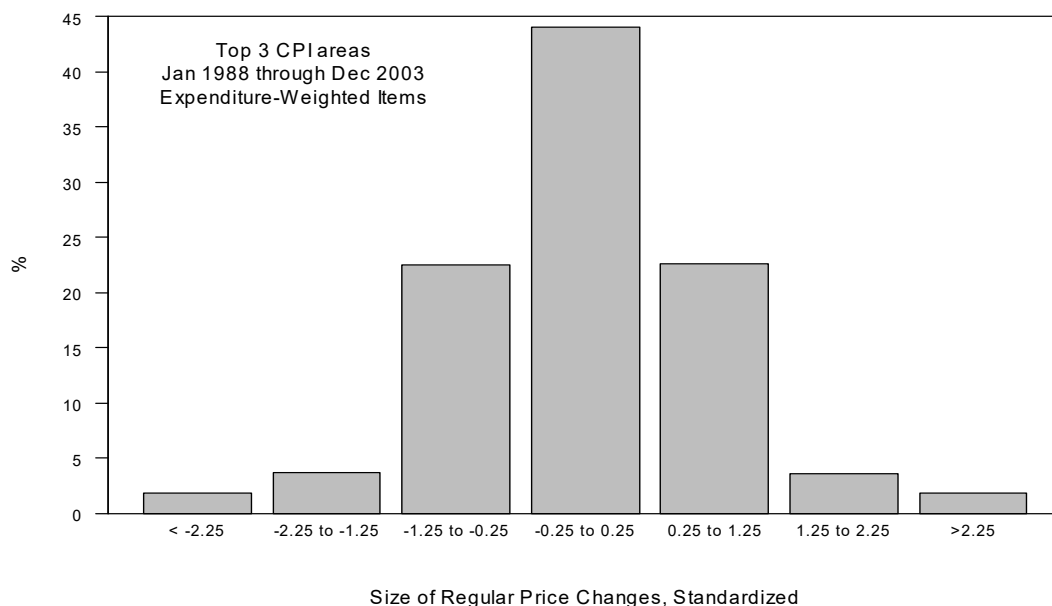


Figure 1. Price Change Distribution from Klenow and Kryvtsov

To choose parameters, we minimize a weighted sum of the squared deviations of the sample statistics (histogram elements, median duration, *and autocorrelation*) from the corresponding properties of our model's steady-state distribution. The parameters we allow to vary are (1) the elements of the productivity transition matrix Q , (2) the productivity realizations e_k , (3) the fraction of firms that are flexible-price firms ω^f , and (4) the upper bound of the fixed-cost distribution $G()$. We fix the number of productivity realizations at seven, restricting the interior elements of that distribution to be equally spaced, but allowing the minimum and maximum productivities to be chosen freely.¹⁰

¹⁰To find parameters such that the model's steady-state distribution matches the data, we use a simulated annealing algorithm developed by William Goffe (1994,1992).

Regarding the cost of adjusting prices, as in DKW we employ a flexible distribution based on the tangent function. The distribution is governed by three parameters: the upper bound on the fixed cost (B) and the curvature parameters (cc and dd) :

$$G(\xi) = \frac{1}{cc} \left\{ \tan \left(\frac{\xi - k_2}{k_1} \right) + dd \cdot \pi \right\}, \quad (16)$$

with

$$k_1 = \frac{B}{[\arctan(cc - dd \cdot \pi) + \arctan(dd \cdot \pi)]} \quad (17)$$

$$k_2 = \arctan(dd \cdot \pi) \cdot k_1. \quad (18)$$

This function can produce a wide range of distributions depending on the values of cc and dd . For example, it can approximate a single fixed cost of adjustment, adjustment probabilities that are relatively flat over a wide range, or the nearly quadratic adjustment such as has been used by Caballero and Engel (2007). In calibrating the distribution to fit the relevant statistics, we allow cc and dd as well as B , the upper bound of the support, to be a free parameters. As indicated above, we also allow for a fraction, ω^f , of firms to adjust their prices flexibly. The presence of these firms is important for generating small price changes in the distribution.¹¹

3.2 Effect of parameters on pricing statistics

Each parameter – or set of parameters – plays an important role in determining the steady-state distribution of price changes. Here we will provide a basic discussion of each parameter’s role, proceeding as though all other parameters were fixed. In matching moments from the data we will implicitly be varying all parameters simultaneously, but this discussion should provide some intuition.

The upper bound on fixed costs has an obvious effect on the length of time firms hold their prices fixed: higher B means greater duration of prices, and therefore larger average price changes for the sticky-price firms. Increasing the fraction of flexible-price firms ω^f leads to a larger number of small price changes and a higher overall frequency of price adjustment. Greater dispersion of idiosyncratic productivity shocks tends to raise the dispersion of relative prices and thus to spread out the distribution of price changes. This reasoning is exact for

¹¹Midrigan (2011) and Costain and Nakov (2011b) emphasize the importance of small price changes for the model’s ability to generate nonneutralities. Midrigan explores two methods, both different from ours, for generating small price changes: complementarity in price adjustment and the probability of drawing a zero menu cost. He finds that the degree of nonneutrality is not significantly affected by the way in which small price changes are introduced.

flexible-price firms, but also occurs for a given persistence of prices in the sticky price sector, because optimal prices will vary by more across states.

The productivity transition matrix also affects price behavior. For a given median adjustment probability, greater persistence in the productivity process tends to increase the dispersion of prices. If productivity is entirely transitory (i.i.d.) then expected future productivity is invariant to current productivity. In this case, with prices adjusting infrequently an adjusting firm's optimal price will not vary much with its current state. In contrast, a highly persistent productivity process makes expected future productivity vary closely with current productivity, and for a given degree of price stickiness optimal prices will also be sensitive to current productivity. This reasoning does not apply exactly, because the degree of price stickiness varies with the productivity process. Greater persistence in the productivity process will tend to increase the degree of price stickiness; greater persistence means a lower probability of a productivity change and thus a lower probability of change in the firm's optimal price. Note that some of the reasoning in this paragraph abstracts from the fact that nonzero steady-state inflation interacts with the productivity process in affecting the incentives for price adjustment.

3.3 Benchmark parameter values

We have found two sets of parameters that enable the model's steady state to closely match the chosen moments.¹² The next subsection will provide more details on the closeness of that match and will discuss in detail the properties of the steady-state equilibrium for both cases. Here we simply provide information on the parameter values.

One case will be called the low-flexibility benchmark because it has a relatively low proportion of flexible-price firms. The other will be called the high-flexibility benchmark because it has a higher fraction of flexible-price firms. The low-flexibility benchmark has the smallest fraction of flexible-price firms that allowed us to match the desired moments. The high-flexibility benchmark has roughly twice as many flexible-price firms and matched the data moments equally well. In each case, the estimated idiosyncratic shock process required a large dispersion of productivity for the extreme states in order to match the fat tails of the empirical price change distribution. However, the two cases required very different upper bounds on the cost distribution as well as very different transition matrices. Importantly, both cases yield hazards that are close to that of a truncated Calvo model.

¹²While we can not prove that there is not another comparable steady state, we did search extensively over various fractions of flexible price firms and found only two steady states that matched the histogram with great precision.

Distribution of Fixed Costs




Figure 2. Distribution of Fixed Costs

For the low-flexibility benchmark, the fraction of firms with flexible prices is $\omega^f = 0.053$. The maximal fixed cost of changing prices is $B = 0.033$, which corresponds to 16.5 percent of labor hours. This does not imply that high costs of adjustment are actually incurred as only .03 percent of time is used in changing prices in the steady state. The productivity levels are given by $e_k \in \{0.93, 0.96, 0.98, 1.0, 1.02, 1.03, 1.06\}$. The dispersion in productivity appears at first glance to be somewhat large. Note, however, that while there is no direct mapping between the products in our model and plants, our productivity process displays considerably less dispersion than has been found in the literature on productivity of plants. For example, Foster, Haltiwanger, and Krizan (2006) find interquartile differences in productivity of .57 for the U.S. retail trade sector.¹³

We also allow the two other parameters of our cost of adjustment distribution, cc and dd , to vary. In the low-flexibility case they equal 41.9 and 2.8, respectively, and in the high-flex case they are 50.4 and 2.3. These parameters give rise to the distributions depicted in Figure 2 and look very much like a truncated constant hazard model and thus very Calvoesque. Along with the productivity transition process, the distribution of fixed adjustment costs helps determine the resulting distribution of prices.

¹³See also Syverson (2004) and Foster, Haltiwanger, and Syverson (2008).

For the high-flexibility benchmark, the fraction of firms with flexible prices is naturally higher: $\omega^f = 0.091$ and the productivity levels are given by $e_k \in \{0.92, 0.96, 0.98, 1.0, 1.02, 1.04, 1.12\}$. Given that both cases match the distribution of price changes and the median price duration, sticky-price firms need to face higher costs of price adjustment: the upper bound on the cost distribution is somewhat larger, $B = 0.075$. While the productivity levels are quite similar to the low-flexibility case, the productivity transition matrix differs greatly. Figure 3 plots the paths of expected productivity, one-period ahead, conditional on each productivity state. In the low-flexibility benchmark, productivity is relatively persistent, whereas in the high-flexibility benchmark productivity is relatively transitory.

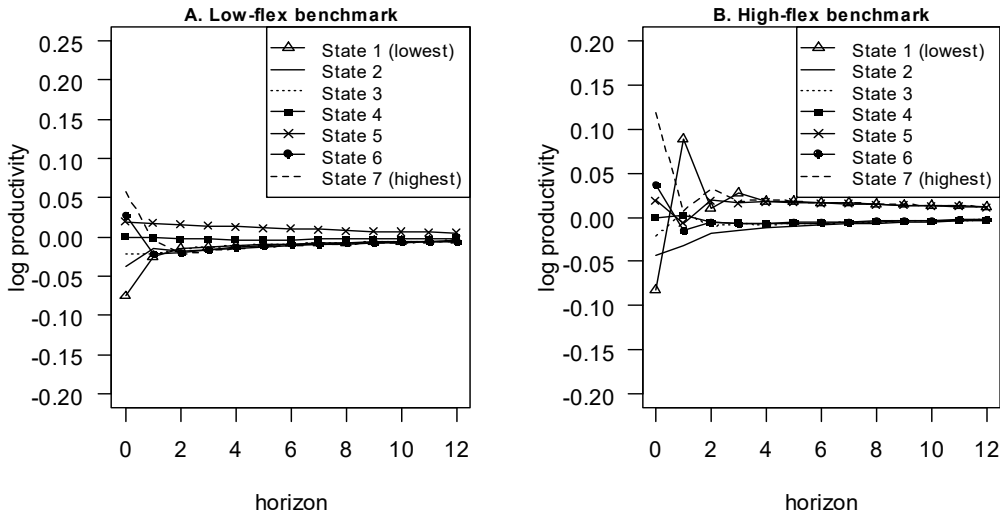


Figure 3. Expected productivity state by state

3.4 Steady-state properties

In both of the benchmark cases, the hypothetical Calvo adjustment probability calculated from the model's steady-state equilibrium matches KK's estimate exactly. KK calculate price duration (λ_{kk}) as the inverse of the Calvo parameter, so our benchmark cases imply identical price durations to KK's median sector.¹⁴ Panel A of Figures 4 and 5 shows that the histograms of price adjustment match the data quite closely. Across these two important dimensions of observable data then, the high- and low-flexibility benchmarks are nearly identical. Of course, we know that in other ways the two parameterizations are not identical.

¹⁴

$$\lambda_{KK} = -\ln \left((1 - \omega^f) \cdot \left(1 - \sum_h \omega_{0,k} \right) \right)$$

In this subsection, we examine the implications of the two benchmark cases for other aspects of price adjustment behavior. In the next section we turn to the implications for monetary nonneutrality.

In addition to the Calvo parameter and implied median duration, Table 2 lists several other statistics reported by KK describing the behavior of prices. For each of those statistics, the low-flex and high-flex benchmarks are quite close to each other. Compared to the data, the models produce an average price duration that is somewhat low (7.9 months vs. 8.6 months in the data), and they produce a fraction of unchanged prices that is somewhat high (0.87 vs. 0.73 in the data). Overall however, for both benchmark cases the steady-state price adjustment behavior implied by the model is quite close to that measured in the data.

In terms of matching the data on price changes in the CPI, our model does so in more detail than most of the existing literature. Golosov and Lucas (2007) match three moments: the fraction of prices that don't change, the mean of positive price changes, and the standard deviation of positive price changes. In restricting their attention to these three moments, their estimated model drastically underestimates the fraction of small price changes and as a result generates no nonneutralities. Nakamura and Steinsson (2010) match the mean frequency and absolute size of price changes across sectors and investigate the role of sectoral heterogeneity in the dynamic responses of their economy to monetary shocks. Like us, Midrigan (2011), using both grocery store data and the CPI, matches a histogram of price changes. Costain and Nakov (2011a,b), using a flexible adjustment cost structure somewhat different from ours, are able to match the distribution of price changes in grocery store scanner data. Unlike in our model, the shape of the state-contingent hazards is insensitive to value function losses associated with not changing prices. Thus, their approach yields a more Calvoesque model with the distribution of price changes being approximately invariant to inflation. Karadi and Reiff (2018) match the frequency and average size of absolute price changes in the Hungarian processed food sector, but they do not closely match of the histogram of these price changes.

| Table 2 | | | |
|---|-------------|-------------------------------|--------------------------------|
| | data | low-flex benchmark | high-flex benchmark |
| fraction flexible | | 0.052 | 0.091 |
| hypothetical Calvo parameter | 0.139 | 0.135 | 0.15 |
| average duration | 8.6 | 7.93 | 7.12 |
| Pr(same) | 0.73 | 0.87 | 0.86 |
| Pr(small)* | 0.30 | 0.36 | 0.44 |
| mean abs(dlnp) | 0.09 | 0.02398 | 0.049 |
| autocorrelation | .32 | .30 | .30 |

* Pr(small) is the fraction of price changes less than 5%.

Having established that both the benchmark model and the model with more flexible-price firms are broadly consistent with the microeconomic data on price adjustment, we now describe additional features of the model's stationary distribution of prices and price adjustment. This discussion will refer to Figures 4 and 5. We have already discussed the close correspondence between the histogram of nonzero price changes in Panel A and the corresponding histogram in Klenow and Kryvtsov (2008). Panel B of Figures 4 and 5 plots the distribution of prices by age. In the low-flex benchmark, 13.5% of prices are newly set in any period; 46% of prices are less than six months old; and 74% of prices are no more than twelve months old. The mean age of prices is 7.7 months. In contrast, there is a somewhat larger fraction of long-lived prices in the high-flex case; the fraction of newly set prices differs little, at 15%, but 38% of prices are less than six months old; and only 61% of prices are no more than twelve months old. The mean age is 10.3 months. The data described in Nakamura and Steinsson (2008) indicates that around 5.1% of prices can reasonably be classified as flexible and we have been unable to find any mention of the price age distribution in the literature. Based on the fraction of flexible prices in the data, one might argue that the low-flex benchmark is the more empirically relevant. However, it is of theoretical interest that there is more than one set of parameters that allow for an accurate depiction of many of the statistics that characterize the price change data. This result is

related to the work of Karadi and Reiff (2018), who find that both their mixed process and their Poisson process generate similar theoretical price change distributions.

Panel C of both figures displays the aggregate hazard function with respect to time, by which we mean the probability of price adjustment conditional on the time since the last adjustment. These figures share two notable features. The first is the relative flatness over the 40 months plotted in the figure. The second is the sharp decline in the hazard from one to two months. In Section 6 we discuss in detail how the properties of the idiosyncratic productivity process determine the shape of the aggregate hazard function. At this juncture we will just make two relatively simple points. First, the sharp decline is due to all the flex-price firms changing prices and leaving only sticky-price firms over the remainder of the hazard distribution. This is the reason that there is a greater initial drop for the model with a relatively high fraction of flexible-price firms. Second, the relatively flat hazard results from an interaction between positive average inflation and the nature of the productivity process. To see this, refer to Panel D of Figures 4 and 5, which plots adjustment probabilities (hazards) as a function of relative price (price charged divided by price level) for each level of productivity. If a nominal price is not adjusted, the corresponding relative price decreases because of positive inflation. If productivity is unchanged, such a move to the left in Panel D corresponds to a higher adjustment probability – a rising hazard with respect to time. One sees that the price distribution of newly set prices is relatively concentrated and that the adjustment probabilities do not change much over most of the relative price values due to the approximately constant hazards associated with the estimated adjustment cost distribution. It takes a fairly sizeable erosion of the relative price before the hazards begin to increase and this accounts for the relatively flat hazards displayed in panel C. Also, because of the small dispersion between the prices set at the lowest and highest productivity levels, prices never rise much above 1.0 and the hazard is low when one moves between productivity states after a recent price change.

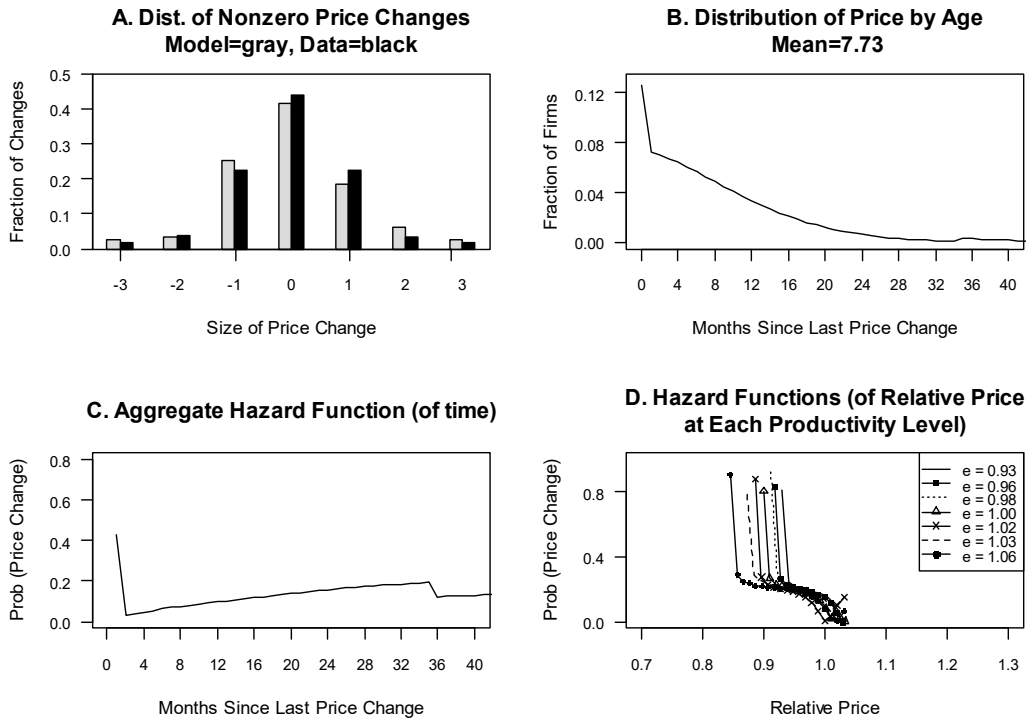


Figure 4. Steady state of low-flex benchmark.

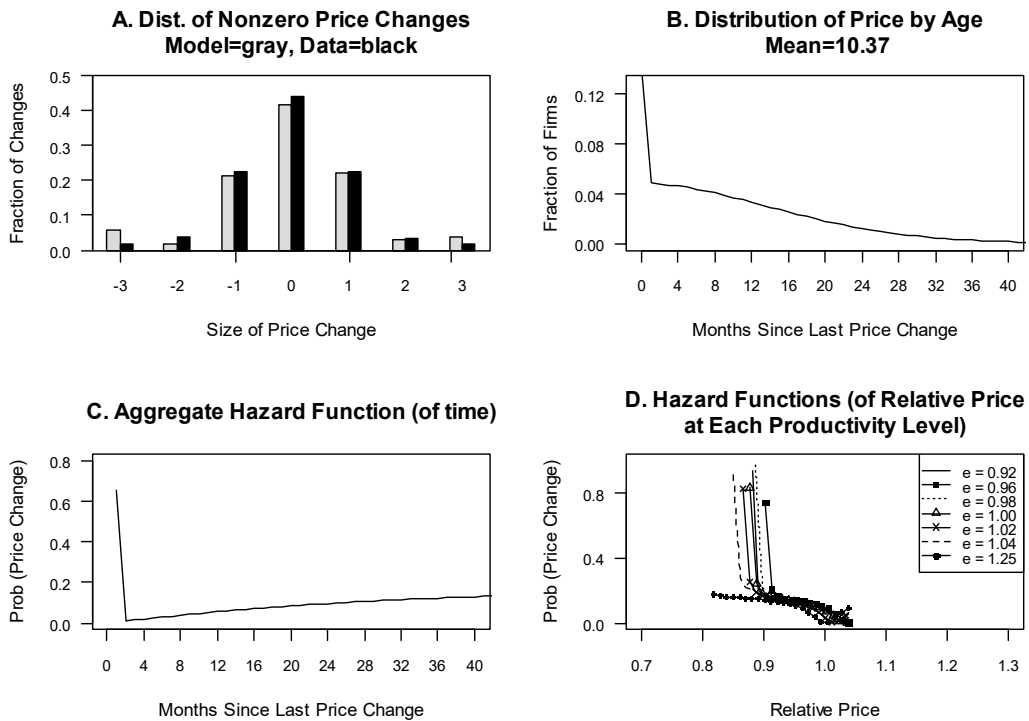


Figure 5. Steady state of high-flex benchmark.

While we have been stressing the commonalities across the two benchmark cases, there are important differences across the two cases in the parameters governing both fixed costs and productivity transitions. The low-flexibility benchmark has low fixed costs for the sticky-price firms and fewer firms (5.2%) with flexible prices, whereas the high-flexibility benchmark has higher fixed costs for the sticky-price firms and more firms (9.1%) with flexible prices. In addition, the productivity process is more persistent in the low-flexibility benchmark. Those differences offset each other and there is little difference across the two figures. Although firms face the potential for higher fixed costs in the high-flexibility benchmark, in the steady-state equilibrium the resources used for price adjustment increase only modestly: 0.030% of labor is used in final goods production in the low-flexibility benchmark and 0.06% in the high-flexibility benchmark.

In Panel D, we observe that in the high-flexibility benchmark the hazard functions are also flat over a wide region, which follows from a similarly calibrated adjustment cost distribution. Another feature of Panel D in the two figures is that the vector of steady-state relative prices is highly concentrated and non-monotonic for both cases. In the high-flexibility benchmark reset prices are given by: $\{1.015, 1.039, 1.033, 1.030, 1.014, 1.028, 1.001\}$ as opposed to $\{1.032, 1.030, 1.027, 1.016, 1.001, 1.023, 1.017\}$ in the low-flexibility benchmark. With a relatively transitory productivity process optimal prices for the sticky-price firms do not vary much across productivity states. The same is true for the low-flexibility case.

Figure 6 provides a different perspective on each model's stationary distribution. Each of the panels has price relative to reset price on the horizontal axis. Firms are grouped into 13 bins, representing different intervals of price relative to reset price before the adjustment decision is made. The top and bottom panels represent the low- and high-flex benchmarks, respectively. The left-hand panels display the distribution of relative prices prior to the adjustment decision, indicating how big an adjustment will be made if it occurs. These panels also provide some information about the distribution of the value of adjustment – firms whose price is close to their optimal price have a relatively small value of adjusting. The right-hand panels display the probability of adjustment.¹⁵ Viewing the four panels of Figure 6 together, we see that the distribution of relative prices before adjustment is concentrated in a region where the hazard function is fairly flat. Following Caballero and Engel (2007), these pictures suggest that the response to an aggregate shock will likely not differ very much from what would occur in a model with fixed adjustment probabilities. Related to Midrigan (2011), there is a relatively small mass of firms that will want to change

¹⁵It is computed by taking a weighted average of the individual hazard functions from Figures 3.D and 4.D, where the weights come from the distributions in the top panels. Also, before averaging, the individual hazards are normalized by their respective relative reset prices.

their price in the presence of a monetary shock, and we should, therefore, expect the economy to exhibit a small change in the aggregate price level in response to such a shock. That is, monetary shocks will be nonneutral.

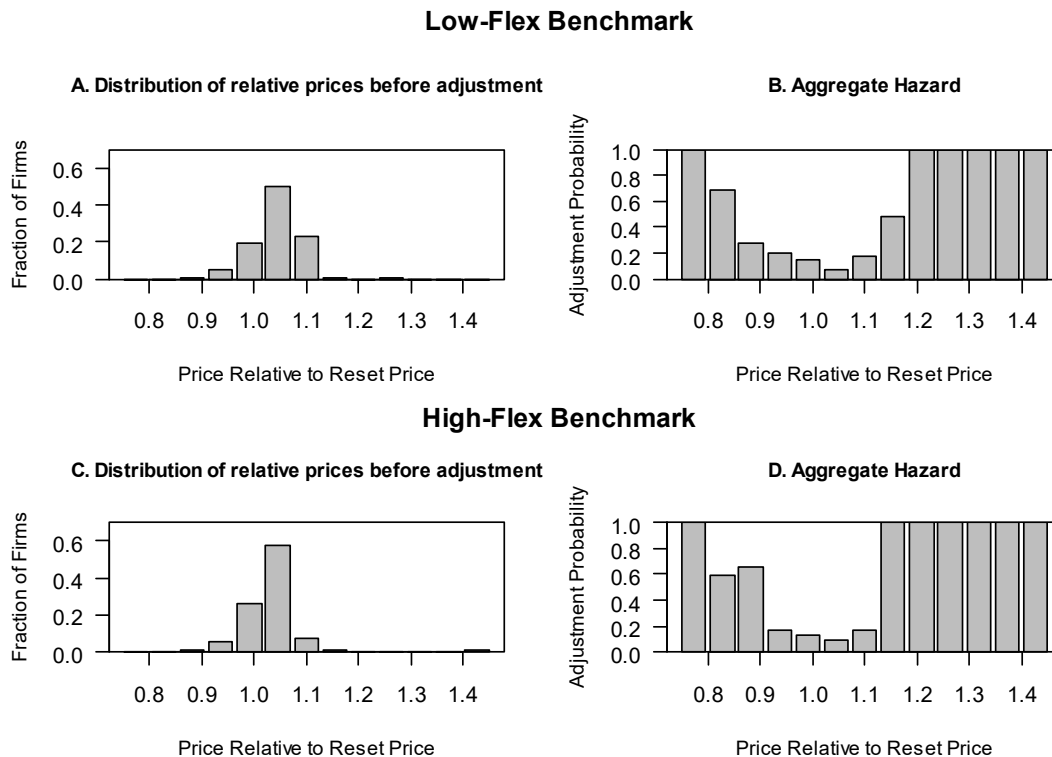


Figure 6. Relative prices and adjustment probabilities.

Another attribute of price adjustment displayed by our two benchmark steady states is the relationship between the size of price changes and the length of time since a price was last adjusted. One might expect this relationship to be positive: the longer a firm waits to adjust its price, the further its price has drifted from optimal. However, in the data there is basically no relationship between these two elements. In our model as well as in the data, the simplistic reasoning does not hold because variation in price duration results from variation in productivity realizations as well as variation in fixed costs. Figure 7 displays our model's implications for the relationship between size and duration of price changes. The low-flex benchmark is graphed in panel A and the high-flex benchmark is graphed in panel B. The time since the last price change is on the horizontal axis and the average size of price changes made by firms that last changed their price is on the vertical axis. Figure 7 shows that the size of price changes is nearly invariant to the time elapsed since the last price change.

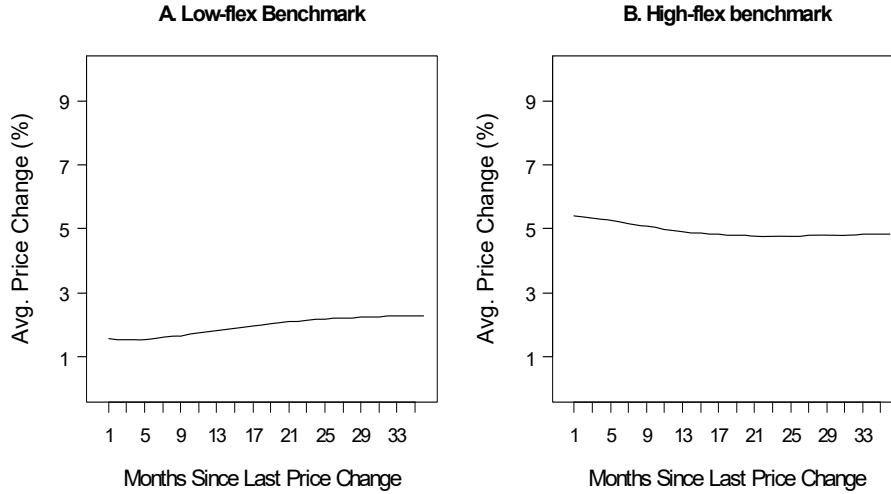


Figure 7. Duration of price and size of price change.

4 Dynamics

Next we examine and compare the local linear dynamics generated by our two benchmark specifications. First, we examine each model's response to a random walk shock. Because the shock has no dynamics, this case is interesting for studying the model's internal dynamics. Subsequently we analyze the case where the monetary shock is persistent. One aspect of state-dependent pricing models that has received much attention is the decomposition of price changes into components associated with fixed and varying adjustment probabilities. Because such a decomposition will appear in our discussion of the model's dynamics, that decomposition needs to be made clear before proceeding to the dynamics.

4.1 Decomposing price changes

In order to better understand the role that state dependence has on the model's dynamics, we decompose price level changes along the lines of Klenow and Kryvtsov (2008), Caballero and Engel (2007) and Costain and Nakov (2011a). Starting from the price index equation:

$$P_t^{1-\varepsilon} = \sum_{h=1}^S \sum_{k=1}^S \theta_{J,h,k,t} P_{0,k,t}^{1-\varepsilon} + \sum_{j=1}^{J-1} \sum_{h=1}^S \sum_{k=1}^S \alpha_{j,h,k,t} \theta_{j,h,k,t} P_{0,k,t}^{1-\varepsilon} + \sum_{j=1}^{J-1} \sum_{h=1}^S \sum_{k=1}^S (1-\alpha_{j,h,k,t}) \theta_{j,h,k,t} P_{0,h,t-j}^{1-\varepsilon},$$

we decompose the log-linearized detrended price index (\hat{P}_t) into four components: (i) a time-dependent component \mathfrak{T}_t , which reflects changes in reset prices ($\hat{p}_{0,k,t} + \hat{P}_t$) holding constant

the distribution of firms and their adjustment probabilities:¹⁶

$$\begin{aligned} \mathfrak{T}_t = & \sum_{j=1}^J \sum_{h=1}^S \sum_{k=1}^S \alpha_{jhk} \theta_{jhk} p_{0,k}^{1-\varepsilon} \cdot \left(\hat{p}_{0,k,t} + \hat{P}_t \right) + \\ & \sum_{j=1}^{J-1} \sum_{h=1}^S \sum_{k=1}^S (1 - \alpha_{jhk}) \theta_{jhk} p_{j,h,k}^{1-\varepsilon} \cdot \left(\hat{p}_{0,h,t-j} + \hat{P}_{t-j} \right); \end{aligned}$$

(ii) an extensive-margin component \mathcal{E}_t , which reflects changes in the average adjustment probability applied to a representative firm:

$$\mathcal{E}_t = (1 - \varepsilon)^{-1} \bar{p} \cdot \bar{\alpha}_t,$$

where \bar{p} can be thought of as representing the average desired price change (though note that \bar{p} is negative if the desired price change is positive):

$$\bar{p} \equiv \sum_{j=1}^J \sum_{h=1}^S \sum_{k=1}^S \theta_{jhk} \left(p_{0,k}^{1-\varepsilon} - p_{j,h,k}^{1-\varepsilon} \right),$$

and $\bar{\alpha}_t$ is the mean deviation from steady state of adjustment probabilities:

$$\bar{\alpha}_t \equiv \sum_{j=1}^{J-1} \sum_{h=1}^S \sum_{k=1}^S \theta_{jhk} \cdot d\alpha_{j,h,k,t};$$

(iii) a selection-effect component \mathcal{S}_t , which reflects the fact that adjustment decisions may be related to the magnitude of desired price changes:

$$\mathcal{S}_t = \left(\frac{1}{1 - \varepsilon} \right) \left(\sum_{j=1}^{J-1} \sum_{h=1}^S \sum_{k=1}^S \theta_{jhk} \left(p_{0,k}^{1-\varepsilon} - p_{j,h,k}^{1-\varepsilon} \right) \cdot d\alpha_{j,h,k,t} - \bar{p} \cdot \bar{\alpha}_t \right);$$

and finally, (iv) a shifting distribution effect \mathcal{D}_t , which quantifies the contribution to the price level of changes in the distribution of firms, holding fixed the prices they charge and their adjustment behavior:¹⁷

¹⁶In these expressions, \hat{x}_t means percent deviation from steady state of x , and dx_t means level deviation from steady state of x .

¹⁷Costain and Nakov (2011a) decompose the behavior of inflation in their model into intensive margin (\mathcal{I}_t), extensive margin and selection components. Our time-dependent, \mathfrak{T}_t , component is closely related to their intensive margin, but not identical. Our \mathfrak{T}_t is simply the effect on the price level of changes in reset prices, holding fixed all distributions and adjustment probabilities. In contrast, their \mathcal{I}_t “leaves out” some of that effect if changes in reset prices are correlated with (steady-state) adjustment probabilities, putting it instead into the selection effect.

In the long run, because adjustment probabilities and the distribution of firms return to steady state, the change in the price level must be entirely accounted for by the time-dependent component. Initially, however, there can be substantial divergence between those paths. Caballero and Engel (2007) refer to the extensive margin effect, which corresponds to our selection effect, as a critical barometer for determining the importance of state dependence.

4.2 Low-flexibility Benchmark: Random Walk Money Shock

Figure 8 displays the response to a permanent 1% increase in M in period zero, for the low-flex parameterization. In response to the shock, the price level *barely* rises on impact (Panel A, dashed line). The half life is about 10 months, which is larger than the average duration in steady state, and that of output is also about 10 months (Panel B, solid line). Output and inflation only gradually return to steady state and the behavior exhibits a modest oscillation. It takes roughly 40 months for the effects of the shock to fully dissipate. While the overall response is around half as big as that found in the empirical literature, it is much larger than the effect found by Golosov and Lucas (2007) and somewhat larger than obtained in Costain and Nakov (2011b). The frequency of price adjustment responds very little to the monetary shock (Panel D), which is suggestive of the findings in Boivin, Giannoni and Mihov (2009) and Mackowiak, Moench and Wiederholt (2009) that idiosyncratic or sectoral shocks are more important than aggregate shocks in determining firms' price adjustment behavior. However, the frequency of price adjustment does play an important role here, exhibiting oscillatory behavior that translates into the dynamic responses of output and inflation.

Panel C presents the price level decomposition described above. The state-dependent part of price changing is indeed important: the changes in the distribution of price changers are almost as important as the time dependent element over the first 18 months. However, the selection effect is not playing an important role, which is not surprising given the bunching of the reset price distribution, implying that the firms who are adjusting do not have a significantly different distribution of desired price changes than the unconditional distribution. The extensive margin contributes almost nothing to the behavior of the price level. As discussed in the derivation of our decomposition, in the long run all price changes are due to the time-dependent component.

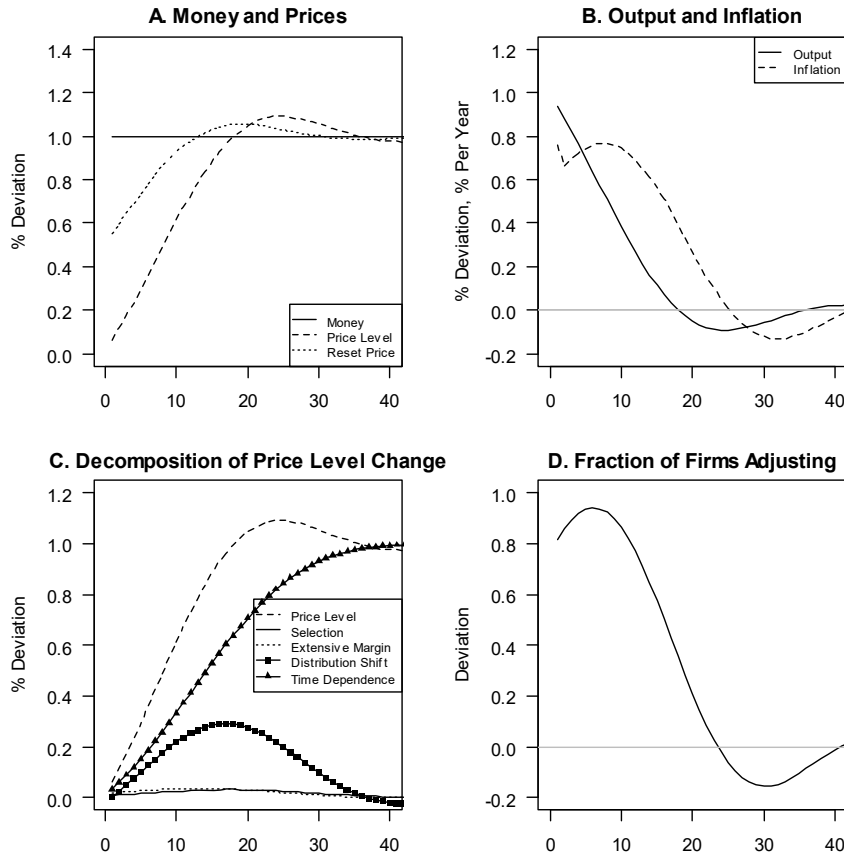


Figure 8. IRF to random-walk money shock, low-flex benchmark.

4.3 High-flexibility Benchmark: Random Walk Money Shock

Figure 9 displays the same set of impulse response functions for the high-flexibility benchmark. There is little difference in the two cases although the change in the fraction of firms adjusting is roughly half that in the low-flexibility case. This result is probably due to the greater fraction of firms that are freely adjusting their price in the steady state. There is somewhat greater nonneutrality as the dynamics of inflation and output are a bit more drawn out. The half life for output is closer to 12 months and for inflation around 25 months. Inflation is again hump shaped, reaching its peak deviation from steady state after approximately eighteen months. As in the low-flexibility case, after 40 months the distribution has still not returned to steady state. And again, the shift in the distribution of price changers is an important contributor to the model economy's dynamic response.

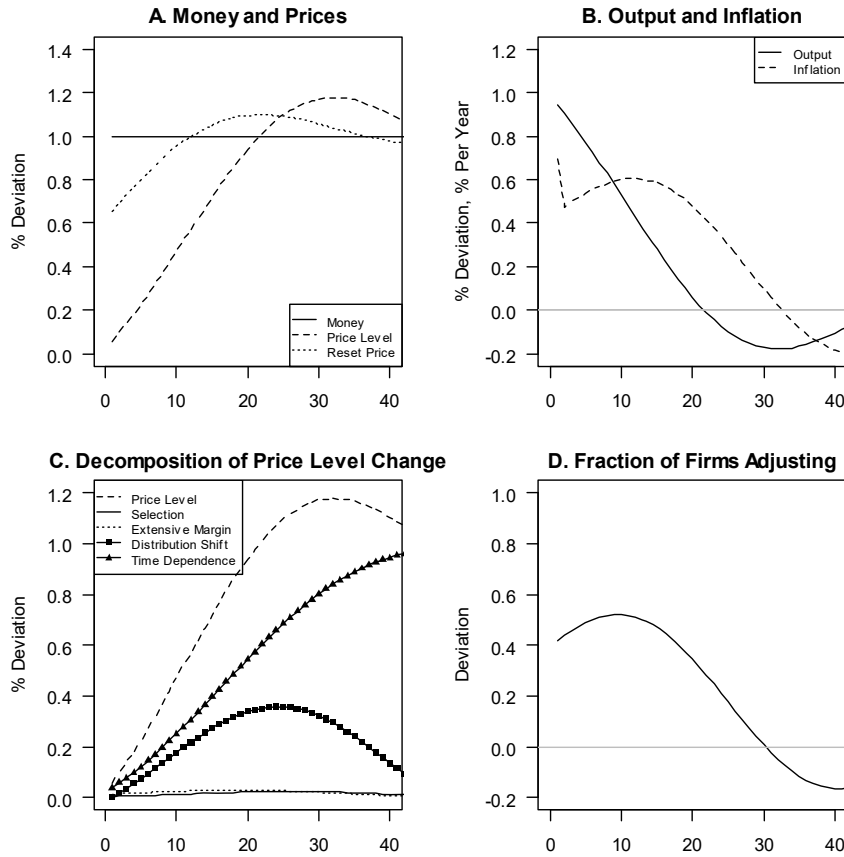


Figure 9. IRF to random-walk money shock, high-flex benchmark.

4.4 Low-flexibility benchmark: Persistent Money Growth Shock

Figure 10 displays impulse response functions to a persistent money growth shock (autocorrelation of .50 at a quarterly frequency) for the low-flex benchmark. There are substantially greater effects on output and inflation in terms of magnitude, but surprisingly little effect in terms of persistence. However, the dynamic responses of both inflation and output are now decidedly hump shaped with the behavior of output leading that of inflation. As in the random walk case, the price level does not respond much on impact. It takes roughly 3 years for both variables to return to steady state. The fraction of firms adjusting is much larger than in the random walk case due to the persistent nature of the shock, which results in the price level rising by nearly 5 times as much as it does in the random walk case. This magnified effect on the price level induces more firms to adjust prices in response to the shock as their current price is now further away from where it will optimally be reset. As in the previous case, the price level decomposition indicates that state-dependent aspects are important for price setting, but the effects are almost entirely manifested through distributional shifts in

price changers. The hump shaped response is also obtained by Karadi and Reiff (2018), but their preferred specification generates a much smaller degree of nonneutrality than we find here.

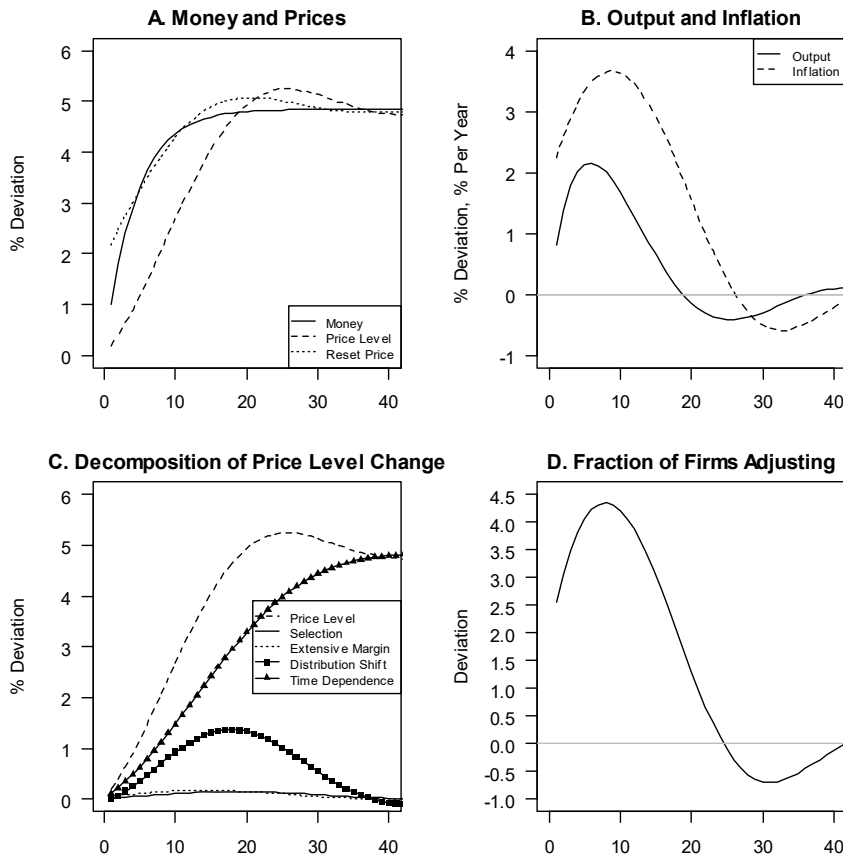


Figure 10. IRF to persistent money shock, low-flex benchmark.

4.5 High-flexibility benchmark: Persistent Money Growth Shock

In response to a persistent shock, the more flexible specification displays roughly the same degree of nonneutrality (Figure 11). The output and inflation responses show a bit more persistence and there is less adjustment in the fraction of firms changing prices. However, there is more similarity between the two cases than there are differences and the differences in steady state behavior do not translate to any marked degree into the dynamics. Thus, there is no direct translation between nonneutrality in the dynamics and the detailed statistics that inform us about stickiness in the steady state – as originally shown in stark terms by Caplin and Spulber (1987). Although Klenow and Kryvtov’s histogram and the overall average duration of prices are derived from a large amount of data, they are insufficient to pin down the parameters of our model.

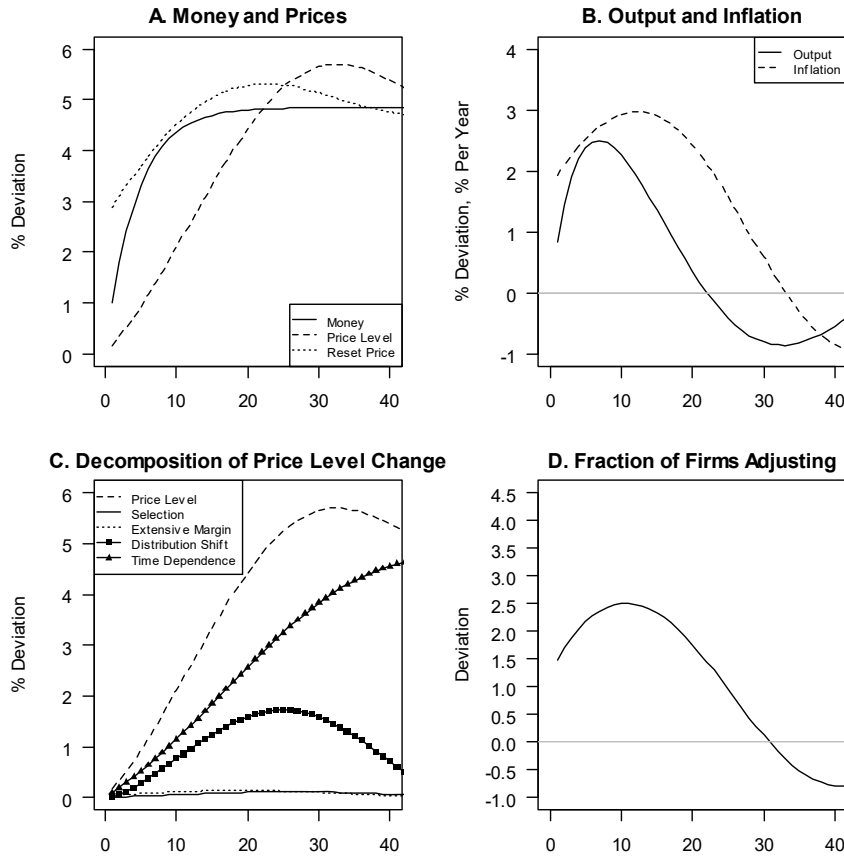


Figure 11. IRF to persistent money shock, high-flex benchmark.

5 Aggregation and Upward- and Downward- Sloping Hazard Functions

A notable feature of both our benchmark case is the wide range over which the aggregate hazard is relatively flat or even downward-sloping, even though conditional on remaining in the same state all hazards are upward sloping. This feature of our model is consistent with hazard functions estimated on micro-data underlying the Japanese CPI (see Ikeda and Nishioka (2007)). Along any pricing spell, the hazards for most goods and services in Japan are increasing and none are decreasing. Appropriately aggregating these various hazards produces a downward sloping hazard for goods and a relatively flat hazard for services. In order for our model to replicate the broad features of hazards in the U.S. and Japanese data, there must be many firms with a low hazard, and the fraction of firms with a low hazard must be increasing as we consider older and older prices (i.e., as we move out along

the hazard function). The reasoning is analogous to, but more involved than that which explains a downward-sloping aggregate hazard in a model with two types of firms with different, constant adjustment probabilities as in Calvo. Aggregating across these two types of firms gives a downward-sloping hazard even though the hazard rates are constant for each type. As the age of a price increases, the fraction of firms that have not adjusted is increasingly dominated by those with the lower hazard – implying a downward-sloping aggregate hazard rate. In the limit, the aggregate hazard approaches that of the low hazard type firms. In this example, the proportion of firms with a low hazard is increasing in age over a significant range and aggregation implies a downward-sloping hazard.

To show how this can happen we modify the parameters so that the hazards are nearly quadratic, which is more in line with the Japanese data. One can see this by examining panels A and C of Figure 12, where we graph the state-specific hazards (Panels A and C) and the weights placed on firms with those particular hazards (Panels B and D) as functions of relative prices in a model, which for simplicity has only two microstates. First, examine Figure 12.A, which gives the hazard rates for the various types of firms when shocks are persistent. The age of a price is represented by the number of symbols one must count to the right in order to reach the firm’s reset price (at which point the hazard is zero). For example, the downward-sloping line with circles depicts the upward-sloping hazard (upward in age, that is) of low-productivity firms that remain low productivity firms, and the fourth circle to the left of the reset price is the relative price of a firm that reset its price while in the low productivity state four periods ago and is still in the low productivity state at that time. The upward sloping squares indicate the downward sloping (age) hazards of firms that were low productivity but switch to being high productivity. This hazard is downward sloping because over time inflation erodes the relative price and makes it closer to the optimal price of resetting high productivity firms. If there are enough of these types of firms, then the aggregate hazard will be downward sloping.

Thus, one needs to know the evolution of the fraction of firms over time. This evolution is given in Figure 12.B, where again the time elapsed since prices were reset is depicted by the distance from the reset price. The fraction of firms that transit from the low- to high-productivity state $\theta_{j,1,2}$ is increasing over time, eventually becoming more than half the firms. Because these are the firms that experience a falling hazard rate and their share is rising, the aggregate hazard is downward sloping.

When the idiosyncratic productivity shocks are i.i.d., the downward-sloping portion of the hazard function for the low to high transiting firms is smaller (see Figure 12.C and 12.D). The distribution of optimal prices is much narrower, so the fraction of firms that are on the downward-sloping portion never gets large enough to offset the upward-sloping hazards

faced by most firms. Thus, a combination of relatively large persistence and dispersion in productivity shocks is required in order to generate a downward-sloping aggregate hazard.

The effect that price dispersion and persistence have on the shape of the aggregate hazard is shown in three dimensions in Figure 13. The left-hand panel displays the effect of changing persistence for a dispersion of 0.156 and the right-hand panel displays the effects of dispersion for a shock process in which there is an 80% probability of remaining in the same idiosyncratic state. Look first at the right-most slice of the right-hand panel; it represents the aggregate hazard function for a model with zero dispersion of the idiosyncratic shocks – that is, a model without idiosyncratic shocks. The hazard function is everywhere upward sloping, as are the hazards displayed in DKW for a model without idiosyncratic shocks. With zero dispersion all firms have the same optimal price. Constant nonzero inflation means that *age* moves all firms uniformly away from their optimal price, raising the adjustment probability for all firms. As the dispersion of idiosyncratic shocks increases (moving left along the axis labeled “dispersion”), the successive slices become flatter, and for dispersion around 0.7, the hazards begin to have downward-sloping portions. It is in this region that there is a large fraction of firms transiting from low to high productivity and then letting their relative price depreciate as described above. For very high degrees of dispersion, the hazards again become upward sloping. In this region reset prices vary greatly across productivity levels, so changes in productivity inevitably involve changes in price.

The left-hand plot in Figure 13 illustrates how the aggregate hazard function varies with the idiosyncratic shock’s persistence. Without persistence, the right-most slice of the figure, the hazard is everywhere upward sloping. In this case, reset prices are relatively insensitive to productivity because firms’ future productivity levels are independent of productivity at the time they set their prices. As persistence increases, the reset prices spread apart, creating the condition described above under which hazards may slope downward. For very high degrees of persistence, transitions across states are rare, so that although they can lead to decreasing hazards for certain productivity trajectories, there are too few firms with such trajectories for the aggregate hazard to slope downward.

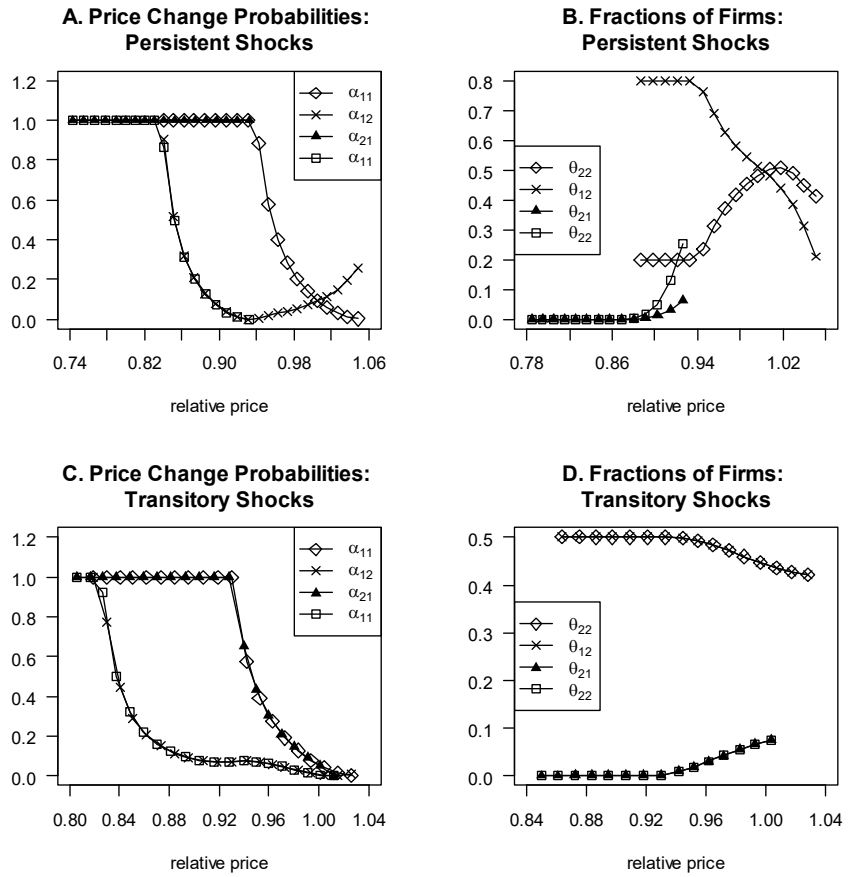


Figure 12. Explaining the aggregate hazard.

Hazards as Functions of Persistence and Dispersion

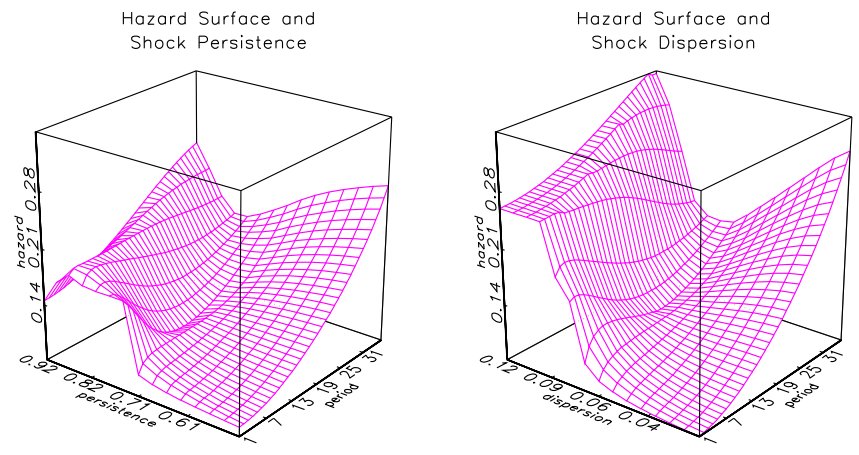


Figure 13. Hazard surfaces.

6 A steady-state sufficient statistic for the degree of nonneutrality?

The approach taken in this paper can be described as brute-force: we calibrate a richly parameterized model so that it matches a large number of facts about the distribution of individual price changes, then use numerical methods to learn about the extent of nonneutrality the model exhibits in response to a monetary policy shock. In contrast, Alvarez, Le Bihan and Lippi (2016) – referred to below as ALL – lay out a parsimonious model and derive a simple analytical relationship between the steady-state distribution of price changes and the degree of nonneutrality in response to a shock. In this section we investigate whether their result holds approximately in our model.

ALL prove that under certain conditions, the steady-state distribution of price changes provides a sufficient statistic for the cumulative output effect of a monetary shock:

$$\mathcal{M} = \frac{\delta}{6\varepsilon} \frac{Kur(\Delta p_i)}{N(\Delta p_i)}, \quad (19)$$

where \mathcal{M} is the cumulative output response to an unexpected permanent increase in the money supply; δ is the size of the money supply shock; $1/\varepsilon - 1$ is the labor supply elasticity; $Kur(\Delta p_i)$ is the kurtosis of the size-distribution of price changes, and $N(\Delta p_i)$ is the average number of price changes per year. While we do not provide a complete formal description of their environment here, key elements behind the result are that (1) each firm’s log productivity follows an independent Brownian motion, and (2) costless price adjustment opportunities arise in the form of a Poisson process; otherwise price adjustment requires that the firm hire a fixed quantity of labor. In contrast, firm-level productivity in our model follows a finite-state stationary process, and the fixed cost of price adjustment is drawn each period (in discrete time) from a continuous distribution. That is, we have a discrete productivity distribution and a continuous distribution of price adjustment costs, whereas ALL have a smooth productivity distribution and a discrete distribution of price adjustment costs. Because of these and other differences between our model and that of ALL, we can be fairly certain that (19) will not hold exactly in our model. However, as ALL state that the main advantage of their result is its robustness, it is natural to investigate whether that robustness extends (approximately) to our model, which we have shown can be calibrated to match a large number of facts about the distribution of price changes.

The usefulness of ALL’s result lies in comparing different model specifications, whether they be fundamentally different models or different calibrations of the same model. Therefore, to determine whether that result holds in our model we need to look across different

values of the parameter vector. One way to do that would be to start with one of our benchmarks and then perturb some parameters. However, because the benchmarks involve parameter settings that are not particularly transparent (they are chosen not for simplicity but to match facts about the distribution of price changes), we take a different approach. We choose a relatively simple parameterization of our model, albeit one that does not deviate fundamentally from our benchmarks, perturb that parameterization in one dimension, and then look at how the right- and left-hand sides of (19) change.

Specifically, we hold all parameters other than those governing fixed costs of adjustment and the idiosyncratic productivity process fixed at their benchmark values. We choose a distribution for the fixed costs of price adjustment that has an approximately quadratic CDF ($cc = 4$ and $dd = 0$ in (16)), with the maximum adjustment cost $B = 0.01$. The fraction of flexible-price firms is 0.025. We use Tauchen’s (1986) approach to approximating a continuous autoregressive process with a finite-state Markov process. As in the benchmark, the number of states is seven. The standard deviation of the innovation to productivity is 0.017, and the approximation covers two standard deviations of the process for the *level* of productivity. We then vary the autocorrelation of the productivity process over a range from 0.94 to 0.99 – recall that in ALL the productivity process is a random walk. For each value of the autocorrelation, we recompute the steady state, record the right-hand side of (19), and then compute the dynamics and record \mathcal{M} . Table 3 shows the right hand side of (19), and \mathcal{M} for five different degrees of persistence of the productivity process. The sufficient statistic from ALL does not carry over to our model. In fact, as the ratio of kurtosis to number of price changes nearly triples, the cumulative output response ratio actually declines ,by almost 30 percent.

| Table 2 | |
|---|---------------|
| Ratio $\frac{Kur(\Delta p_i)}{N(\Delta p_i)}$ | \mathcal{M} |
| 2.08 | 6.36 |
| 2.32 | 6.09 |
| 4.11 | 5.62 |
| 4.97 | 5.28 |
| 5.93 | 4.96 |

While we have not undertaken a systematic analysis of what elements of our model cause the result in ALL to break down, based on their robustness exercises it seems likely that

the culprit is the dynamics of the firm-level productivity process. In our analysis, firm-level productivity is persistent but stationary, whereas in ALL it follows a random walk. Note that while the level of productivity is stationary for each firm in our model, this does not imply that a firm’s optimal nominal price is stationary, as it would in ALL. All the parameterizations we study have positive average inflation, so that the price level is nonstationary. In ALL average inflation is zero, and individual firms’ prices drift because of the nonstationary productivity process.

7 Summary and Conclusions

Within the confines of a standard DSGE model, we append a state-dependent pricing structure with idiosyncratic productivity variation. Employing small fixed costs of price adjustment, two parameterizations of the model are capable of matching many of the facts that have recently been uncovered concerning firms’ pricing behavior. In particular, not only do they match the distribution of price changes, they also match the moderate degree of stickiness, measured by average price duration. Further, the dispersion in productivity across firms needed by the models to account for the size and dispersion of price changes in the data does not seem overly large when looked at through the lens of the plant productivity literature. Also, the aggregate hazard functions generated by the benchmark models are rather flat, which is consistent with the data. This result occurs despite the fact that conditional on productivity, all hazards are upward sloping, a feature that appears to be consistent with micro-hazard data from the Japanese CPI and U.S. scanner data. We are able to trace out the way that aggregation works in our model and show that flat hazards are a feature of the dispersion and persistence of idiosyncratic productivity shocks. Additionally, our models are consistent with the autocorrelation of price changes and with the lack of correlation between size of a price change and time since the last price change.

Despite the relatively high degree of steady-state price flexibility in our model, there is significant nonneutrality of monetary disturbances. There is also substantial persistence in response to shocks. From our pricing decompositions, it is clear that state-dependence plays an important role in the way monetary shocks propagate in the model. Our analysis also reinforces Cabellero and Engel’s (2007) suggestion that steady-state stickiness and dynamic persistence are not as closely related in state-dependent models as in time-dependent models. As well, we confirm the analysis in Karadi and Reiff (2018) who indicate that “without direct information on the menu costs or the distribution of idiosyncratic shocks, the unconditional cross-sectional price change distribution does not identify the extent of money nonneutrality.” Finally, we investigate the interesting theoretical work of Alvarez, Le Bihan,

and Lippi (2016) and find that the relationship between nonneutrality and the kurtosis of the steady-state distribution of price changes is not what their theory predicts. Thus, the extent of nonneutrality in heterogeneous firm models depends on a host of interconnected features among which is the presence of small price changes. Importantly, our model is solved by perturbation methods holding out the hope that in the not-too-distant future state-dependent pricing models can be estimated by Bayesian means and serve as a useful tool for predicting the effects of changes in monetary policy.

References

- [1] Alvarez, Fernando E. and Francesco Lippi (2014). “Price setting with menu cost for multiproduct firms,” *Econometrica* 82, 89-135.
- [2] Alvarez, Fernando E., Herve Le Bihan, and Francesco Lippi (2016). “The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach,” *American Economic Review* 106, 2817-2851.
- [3] Boivin, Jean, Marc P. Giannoni, and Ilian Mihov (2009). “Sticky Prices and Monetary Policy: Evidence from Disaggregated U.S. Data,” *American Economic Review*, 99, 350-384.
- [4] Burstein, Ariel and Christian Hellwig (2007). “Prices and Market Shares in a Menu Cost Model,” manuscript, September 2007.
- [5] Caballero, Ricardo J. and Eduardo M. R. A. Engel (2007). “Price stickiness in Ss models: New interpretations of old results,” *Journal of Monetary Economics* 54, 100-121.
- [6] Calvo Guillermo (1983). “Staggered Prices in a Utility Maximizing Framework,” *Journal of Monetary Economics* 12, 383-398.
- [7] Caplin, Andrew S. and Daniel F. Spulber (1987). “Menu costs and the neutrality of money,” *Quarterly Journal of Economics* 102, 703-726.
- [8] Carvalho, Carlos (2006). “Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks,” *The B.E. Journal of Macroeconomics*. 2.
- [9] Costain, James and Anton Nakov (2011a). “Distributional Dynamics under Smoothly State-Dependent Pricing,” *Journal of Monetary Economics*, 58, 646-665.

- [10] Costain, James and Anton Nakov (2011b). "Price Adjustments in a General Model of State-Dependent Pricing," *Journal of Money, Credit and Banking* 43, 385-406.
- [11] Dixit, A.K., and Joseph E. Stiglitz (1976). "Monopolistic Competition and Optimum Product Diversity," *American Economic Review* 67, 297-308.
- [12] Dotsey, M., King, R.G., Wolman, A. L. (1999). "State-dependent pricing and the general equilibrium dynamics of money and output." *Quarterly Journal of Economics* 114, 655–90.
- [13] Dotsey, Michael, and Robert G. King (2005). "Implications of State-Dependent Pricing for Dynamic Macroeconomic Models," *Journal of Monetary Economics* 52, 213-242.
- [14] Dotsey, Michael, and Robert G. King (2006) "Pricing, Production, and Persistence," *Journal of European Economic Association* 4, 893-928.
- [15] Foster, Lucia, John Haltiwanger, and C. J. Krizan (2006). "Market Selection, Reallocation, and Restructuring in the U.S. Retail Trade Sector in the 1990s," *Review of Economics and Statistics* 88, 748-758.
- [16] Foster, Lucia, John Haltiwanger, and Chad Syverson (2008). "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review* 98, 394-425.
- [17] Gertler, Mark and John Leahy (2008). "A Phillips Curve with an Ss Foundation," *Journal of Political Economy* 116, 533-572.
- [18] Goffe, William L., Gary D. Ferrier, and John Rogers (1992). "Simulated Annealing: An Initial Application in Econometrics," *Computer Science in Economics & Management*, Society for Computational Economics. 5, 133-46.
- [19] Goffe, William L., Gary D. Ferrier, and John Rogers (1994) "Global Optimization of Statistical Functions with Simulated Annealing," *Journal of Econometrics* 60, 65-99.
- [20] Golosov, M. and R.E. Lucas (2007). "Menu costs and Phillips curves ". *Journal of Political Economy* 115, 171-199.
- [21] Guimaraes, Bernardo and Kevin Sheedy (2011). "Sales and Monetary Policy," *American Economic Review* 101. 844-76.
- [22] Ikeda, Daisuke, and Shinichi Nishioka (2007). "Price Setting Behavior and Hazard Functions: Evidence from Japanese CPI Micro Data," manuscript.

- [23] John, A. Andrew and Alexander L. Wolman (2008). “Steady-State Equilibrium with State-Dependent Pricing”, *Journal of Monetary Economics* 55, 383-405.
- [24] Karadi, Peter and Adam Reiff (2018). “Menu Costs, Aggregate Fluctuations, and Large Shocks,” *American Economic Journal: Macroeconomics*. Forthcoming.
- [25] Kemeny, J. G. and J. L. Snell (1976). *Finite Markov Chains*. New York: Springer-Verlag.
- [26] Klenow, Peter J. and Oleksiy Kryvtsov (2008). “State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?” *Quarterly Journal of Economics* 123, 863-904.
- [27] Klenow, Peter J. and Benjamin A. Malin (2011). “Microeconomic Evidence on Price-Setting,” In *Handbook of Monetary Economics*, Volume 3A, Elsevier B.V., 231-284.
- [28] Klenow, Peter J and Jonathan L. Willis (2016). “Real Rigidities and Nominal Price Changes,” *Economica* 83, 443-472.
- [29] Luo, Shaowen and Daniel Villar (2018). “The Skewness of the Price Change Distribution: A New Touchstone for Sticky Price Models.” Manuscript.
- [30] Mackowiak, Bartosz, Emanuel Moench, and Mirko Wiederholt (2009). “Sectoral Price Data and Models of Price Setting,” *Journal of Monetary Economics* 56, S78-99.
- [31] Midrigan, Virgiliu (2011). “Menu Costs, Multi-Product Firms, and Aggregate Fluctuations,” *Econometrica* 79, 1139-1180.
- [32] Nakamura, Emi and Jon Steinsson (2008). “Five Facts About Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics* 123, 1415-1464.
- [33] Nakamura, Emi and Jon Steinsson (2010). “Monetary nonneutrality in a Multisector Menu Cost Model,” *Quarterly Journal of Economics* 125, 961-1013.
- [34] Syverson, Chad (2004). “Product Substitutability and Productivity Dispersion,” *Review of Economics and Statistics* 86, 534-550.
- [35] Tauchen, George (1986). “Finite state markov-chain approximations to univariate and vector autoregressions.” *Economics Letters* 20, 177-181
- [36] Vavra, Joseph (2014). “Inflation Dynamics and Time-varying Volatility: New Evidence and an Ss Interpretation.” *Quarterly Journal of Economics* 129, 215-258.