
DOI: 10.1111/jmcb.12700

SHAOWEN LUO
DANIEL VILLAR

The Skewness of the Price Change Distribution: A New Touchstone for Sticky Price Models

We present a new way of empirically evaluating various sticky price models that are used to assess the degree of monetary nonneutrality. While menu cost models uniformly predict that price change skewness and dispersion fall with inflation, in the Calvo model, both rise. However, the U.S. Consumer Price Index (CPI) data from the late 1970s onward show that skewness does not fall with inflation, while dispersion does. We present a random menu cost model that, with a menu cost distribution that has a strong Calvo flavor, can match the empirical patterns. The model exhibits much more monetary nonneutrality than existing menu cost models.

JEL codes: E31, E32, E47, E52

Keywords: nominal rigidity, state-dependent pricing, random menu cost, monetary nonneutrality

THE DYNAMICS OF PRICE CHANGES—that is, when, how, and why firms change the prices of the goods and services that they sell—has been an area of active research in monetary economics over the past few decades. It is well known that monetary variables have no influence on real economic activity if all prices can be freely reset at any point in time. Therefore, much effort has been devoted toward incorporating frictions in price-setting models and using detailed price data to measure how sticky prices really are (Klenow and Kryvtsov 2008, Nakamura and Steinsson 2008, etc.). One important finding in this literature is that the degree of monetary

We would like to thank Emi Nakamura, Jón Steinsson, Ricardo Reis, and Michael Woodford for their invaluable advice and support. We also thank the editor, two anonymous referees, Jennifer La'O, Martín Uribe, Stéphane Dupraz, Jorge Mejía-Licona, Savitar Sundaresan, Erick Sager, Timothy Erickson. The data were made accessible to us by the Bureau of Labor Statistics under MOU 3594, and we thank Ted To and John Molino for their help as our BLS coordinators. All remaining errors are our own. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Board of Governors or the Federal Reserve System.

SHAOWEN LUO is an Assistant Professor in the Department of Economics at Virginia Tech (E-mail: sluo@vt.edu). DANIEL VILLAR is an Economist in the Division of Research and Statistics at the Federal Reserve Board of Governors (Email: daniel.villar@frb.gov).



Fig 1. Intuition for the Menu Cost Model.

NOTES: In the first three panels, the black curve represents the distribution of the desired price change. The dashed lines represent the Ss band. The gray shaded area represents the distribution of realized price changes.

neutrality depends not only on how often prices change, but also, crucially, on the extent to which the prices that change are selected based on their misalignment. This mechanism has come to be known as the selection effect (following Caplin and Spulber 1987, Golosov and Lucas 2007, etc.). Specifically, if the prices that change are those heavily misaligned from their optimal level (as would be the case if firms must pay price adjustment, or menu costs), money will have a much smaller real effect than if they were randomly selected (as is the case in Calvo-type models).

This paper evaluates the strength of the selection effect based on an empirical pattern that has not been previously considered—the correlation between inflation and price change skewness, and a new data set of prices covering high inflation periods—the Consumer Price Index (CPI) microdata going back to 1977. Moreover, we present a new random menu cost model of sticky prices (based on Dotsey, King, and Wolman 1999) to capture the patterns observed in the data and use the model to assess the degree of monetary nonneutrality.

The strength of the selection effect can be inferred through the distribution of realized price changes, although it is not directly observable in the data. We exploit the fact that the selection in price changes has a strong impact on the behavior of the distribution of realized price changes in response to aggregate shocks. In fixed menu cost models, the presence of a fixed adjustment cost induces a selection effect: the price adjustment occurs if and only if the profit gains from the adjustment outweigh the adjustment cost. This leads to an inaction region of price changes and rules out small price changes. Because adjusting firms change prices by large amounts, the real effect of monetary shocks is small. In addition, this selection effect implies that both dispersion and skewness of the *nonzero price change* distribution fall with inflation. As illustrated in Figure 1, an inflationary shock will push some price changes out of the inaction region to the positive side, and push some changes into the inaction region from the negative side. Thus, the realized price change distribution is less dispersed and more asymmetric with a negative skewness. The opposite is true for a

deflationary shock. These implications appear in a broad class of menu cost models and can be tested empirically.¹

In the data, we find that while the dispersion of nonzero price changes clearly falls in high inflation periods, the skewness does not. The latter is contrary to the predictions of menu cost models, and is therefore inconsistent with a very strong selection effect. Moreover, the negative inflation–dispersion correlation and the positive correlation between inflation and the frequency of price change contradict the predictions of the Calvo model. Overall, we find that no existing model can match all the empirical patterns that we observe.

We use the data set recently presented in Nakamura et al. (2018), that extends the CPI microdata back to 1977, to evaluate whether the dispersion and skewness of price changes do indeed fall with inflation. Since the newly recovered period includes the highest inflation episodes in the postwar United States, as well as the disinflation period initiated by the Federal Reserve under Paul Volcker, our data set is particularly well suited for the tests that we propose.² The extended data set thus overcomes an important limitation faced by the CPI data from 1988 onward (the main source of price data in the sticky price literature) covering only periods of low and stable inflation.³

To develop a model consistent with nonnegative inflation–skewness and negative inflation–dispersion correlations, we modify the menu cost model in a way that weakens the selection effect. We do this by introducing random menu costs that add randomness to whether the firm will have an opportunity to change its price following Dotsey, King, and Wolman (1999). The model therefore incorporates some Calvo features, and can be thought of as a hybrid between state- and time-dependent sticky price models. Random menu cost models have been proposed by previous studies. We follow the example of Dotsey, King, and Wolman (1999) and Dotsey and Wolman (2018) and adjust the distribution of menu costs to fit the new correlations that we report.⁴ We find that in order to capture the nonnegative inflation–skewness correlation, the probability of price changes being free must be nonzero and the probability of price changes being costly must be high. The fitted model features low state dependence and implies a high level of monetary nonneutrality, six times higher than that in a fixed menu cost model, and approximately 70% as large as that in a Calvo model.

Our paper contributes to a large literature devoted to studying the selection effect of price changes and monetary nonneutrality. Caplin and Spulber (1987) show

1. As discussed in Section 1, the inflation–dispersion relation is negative in the inflation region, and positive in the deflation region. In the data, we mostly observe positive inflation episodes. Because of this, and for convenience, we will mostly refer to the “negative inflation–dispersion correlation,” instead of “the negative inflation–dispersion correlation in the positive inflation region.”

2. Although some studies (such as Gagnon 2009, Alvarez et al. 2019) have used price data from countries that experienced high inflation, they study how the frequency of price change behaves at high inflation, without considering the higher moments of the price change distribution. Notably, Alvarez et al. (2019) look at the dispersion of price levels (within narrow product categories), but not of price changes.

3. Other commonly used data sets go back even less far, such as Dominick’s and the Nielsen Homescan Dataset.

4. Dotsey and Wolman (2018), in particular, use price microdata to estimate a random menu cost model in order to derive results on nonneutrality. However, they do not consider the informative correlations moments that we show in this paper.

that in a highly stylized menu cost model, money can be completely neutral even if prices change very infrequently. Golosov and Lucas (2007) then develop a quantitative standard menu cost model calibrated to match certain empirical price change facts. In their model, monetary shocks have very small effects, and these results seriously called into question whether monetary policy could influence the real economy to the degree shown by the Calvo model. However, more advanced menu cost models, such as Woodford (2009), Costain and Nakov (2011), Midrigan (2011), and Costain and Nakov (2018) later included Calvo features in the price change policy, which reduce the degree of selection. These models are able to match some important features of the data, and generated considerably higher levels of monetary nonneutrality than the standard menu cost models.

Empirical studies of price stickiness in certain industries have been around for some time (e.g., Carlton 1986, Cecchetti 1986, Kashyap 1995). However, it is only starting with Bils and Klenow (2004) that monetary economists have been able to start measuring statistics related to price stickiness for the economy as a whole. The work done by Bils and Klenow and the subsequent empirical studies on price stickiness (most notably Klenow and Kryvtsov 2008, Nakamura and Steinsson 2008, Klenow and Malin 2010) have enriched the discussion on monetary nonneutrality by providing empirical facts to evaluate the models. Since Golosov and Lucas (2007), the literature has continued to combine quantitative, microfounded, price setting models with empirical facts from microprice data sets, and in this way the nonneutrality debate has advanced (e.g., Nakamura and Steinsson 2010, Midrigan 2011, Alvarez, Bihan, and Lippi 2016). Nakamura and Steinsson (2010) and Midrigan (2011) have already pointed out problems with some of the predictions of the Golosov and Lucas model, and shown that changes to the model that corrected these problems overturned the result of low monetary nonneutrality. However, we show that even these modifications to the Golosov and Lucas model, though they reconcile the menu cost framework with the data in some ways, are inconsistent with the facts that we present. In particular, these studies have, for the most part, used unconditional moments of the price change distribution (such as the frequency or the size of price changes averaged over time) to discipline the models in question. In this paper, we show that conditional higher moments of prices changes are extremely informative. Finally, Vavra (2013), Alvarez et al. (2019), Alvarez, Bihan, and Lippi (2016), and Berger and Vavra (2018) also consider higher moments of the price change or price-level distribution and their implications for sticky price models.^{5,6} Connected to these studies, we focus on the behavior of the nonzero price change distribution, as opposed to the relative price

5. Alvarez, Bihan, and Lippi (2016) in particular show that in a broad class of models (covering many of the models that we consider) the kurtosis of price changes provides information on the degree of monetary nonneutrality. Using data from the CPI in France, they find that the degree of nonneutrality is between that in menu cost models and that in the Calvo model, broadly consistent with our final results. As we explain further on, we believe that our approach to inferring nonneutrality has practical empirical advantages.

6. Berger and Vavra (2018) also measure the time-series variation of skewness of the nonzero price change distribution. However, they focus on the relationship between moments of the price change distribution and the business cycle. Their analysis shows no notable cyclical variation in skewness.

distribution. Moreover, we emphasize that the relationship between price change higher moments and aggregate inflation is particularly informative about the selection mechanism.

The rest of the paper is organized as follows. In Section 1, we present the predictions of a large class of sticky price models, and explain why time- and state-dependent models give such different predictions. Section 2 describes the data set that we use and evaluates the predictions of different models based on the data. Section 3 presents the random menu cost model, comparing predictions to empirical findings, and shows the degree of monetary nonneutrality exhibited by the different models. Finally, Section 4 provides some concluding remarks.

1. THE SKEWNESS AND DISPERSION OF PRICE CHANGE IN STICKY PRICE MODELS

In this section, we explain and illustrate how the comovement between inflation and the higher moments of the price change distribution provides information on the strength of the selection effect, and therefore on the degree of monetary nonneutrality. First, we derive analytical expressions for the variance and skewness of price changes in a simple menu cost model, and show that both are decreasing in inflation. Second, we provide an intuitive explanation for this result based on the mechanics of menu cost models. Then, we present simulations based on a broader set of menu cost models to show that the results hold more generally.

1.1 Analytical Results

In this subsection, we derive an analytical expression for the inflation–dispersion as well as inflation–skewness correlations in a simple menu cost model, based on a version of the model in Alvarez et al. (2019) (itself similar to the model of Golosov and Lucas 2007).

In a continuous time set-up, monopolistically competitive firms face idiosyncratic shocks z , an aggregate price level that grows at a constant inflation rate π , and a menu cost whenever they adjust prices. With constant elasticity of demand (θ) and constant returns to scale, the instantaneous profit of a firm is given by the function,

$$F(p, z) = e^{-\theta p}(e^p - e^z), \quad (1)$$

where p is the log of its price relative to the aggregate price. The idiosyncratic real marginal cost shock z follows a diffusion process

$$dz = \sigma dW - z dN, \quad (2)$$

where dW is a standard Brownian motion, and N is the counter of a Poisson process with a constant arrival rate per unit of time, $\dot{\rho}$. This can be interpreted as the rate at

which firms die and are replaced with a new firm. This allows the model to maintain a stationary distribution of relative prices while the process for marginal costs contains a trend.

Firms choose a sequence of price adjustment periods, and a corresponding sequence of price change quantities, to maximize the discounted value of future profits net of menu costs. The menu cost depends on the realization of the marginal cost, and is proportional to the profit function evaluated at the static profit-maximizing relative price $p^*(z)$:

$$\chi(z) = cF(p^*(z), z), \quad (3)$$

where c is a constant.

The continuous time nature of the model, along with the assumptions about the shock process, imply that a firm will change its price exactly at the point in time in which its relative price reaches an upper or lower threshold. This means that price changes will only take a single positive or a single negative value, so that the distribution of price changes will have two mass points. We exploit the simplicity of this distribution to derive expressions for the dispersion and skewness.

Alvarez et al. (2019) show that the optimal policy is for the firm to follow a threshold rule summarized by a vector of constants $X \equiv [\underline{x}, \hat{x}, \bar{x}]$, with $\underline{x} < \hat{x} < \bar{x}$. The price is left unchanged as long as $p \in (\underline{x} + z, \bar{x} + z)$, and is adjusted to be $\hat{x} + z$ when it reaches either boundary of the inaction region. When the relative price reaches the lower bound of the inaction region, the size of the price change is,

$$\Delta^+ \equiv (\hat{x} + z) - (\underline{x} + z) = \hat{x} - \underline{x}. \quad (4)$$

When the relative price reaches the upper bound, the size of the price change is,

$$\Delta^- \equiv (\bar{x} + z) - (\hat{x} + z) = \bar{x} - \hat{x}. \quad (5)$$

Regardless of the realization of z , positive price changes only take one value, while negative price changes also take one value that does not depend on z . Since there are no aggregate shocks that affect the growth rate of the price level, which is effectively being tracked, the size of price change does not vary with the realization of z , and only depends on whether the upper or lower bound of the inaction region has been reached. The continuous nature of time and of the shock process is important for this result.⁷

We focus on the realized *price change distribution excluding zero price changes*. As proved in Online Appendix A, the variance and skewness of price changes distribution only depend on the size of price increases (Δ^+) and decreases (Δ^-), the

7. In discrete time models (which will be the focus of the analysis in the next section), a firm also changes its price once a threshold is reached. However, because shocks arrive at discrete intervals, the firm will generally change its price once its relative price has been pushed beyond the inaction region (the relative price will almost never be exactly at the boundary of the inaction region), so that different firms will choose different price changes. Under continuous time, in contrast, firms wait until their relative price is exactly equal to the optimal threshold.

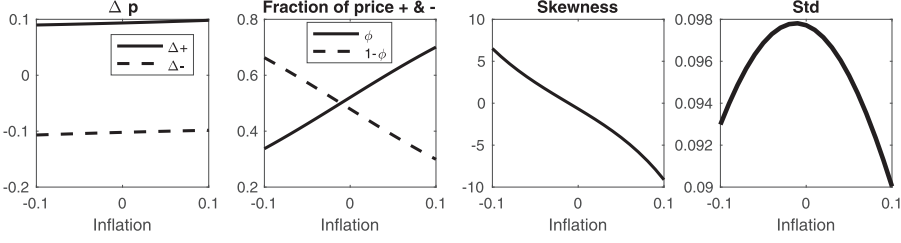


Fig 2. Price Change Moments of the Continuous Time Model.

NOTES: X and ϕ are solved following Propositions 4 and 5 in Alvarez et al. (2019) by setting $\dot{\rho} = 0.1$, $\theta = 3$, $c = 0.02$, $\sigma = 0.15$, and firm's discount parameter equals 0.06.

frequency of price increases (λ^+) and decreases (λ^-), with the following functional form:

$$\text{standard deviation: } E|\Delta p - \gamma| = \sqrt{\phi(1-\phi)}(\Delta^+ + \Delta^-), \quad (6)$$

$$\text{skewness: } \frac{E(\Delta p - \gamma)^3}{[E(\Delta p - \gamma)^2]^{\frac{3}{2}}} = \frac{1 - 2\phi}{\phi^{\frac{1}{2}}(1-\phi)^{\frac{1}{2}}}, \quad (7)$$

where $\gamma \equiv E(\Delta p | \Delta p \neq 0) = (\lambda^+ \Delta^+ - \lambda^- \Delta^-) / (\lambda^+ + \lambda^-)$ denotes the mean of the price change distribution, $\phi \equiv \lambda^+ / (\lambda^+ + \lambda^-)$ denotes the fraction of price changes that are increases.⁸

For low to intermediate values of inflation, the size of price changes (Δ^+ and Δ^-) do not change appreciably, as shown in Alvarez et al. (2019). However, the fraction of positive price changes, ϕ , rises with inflation. It is straightforward to prove that the skewness of the price change distribution decreases with ϕ , given that $\partial \text{skewness} / \partial \phi < 0$. In turn, this means that the relationship between price change skewness and inflation is negative. Moreover, the dispersion of the price change distribution falls with inflation when $\phi \geq 0.5$, and rises with inflation when $\phi < 0.5$. Figure 2 illustrates the size of price changes, the fraction of price increases and decreases, along with the second and the third moments of the price change distribution, conditional on inflation.

In the following subsections, we will show that the same results hold in more general menu cost models. Before proceeding, it is worth pointing out an important difference between our focus here and the focus in Alvarez et al. (2019). First, Alvarez et al. (2019) study the relations between inflation and dispersion of the *price level*

8. The kurtosis of the price change distribution in this set-up is closely related to the skewness of the distribution, given that $\text{kurtosis} = E(\Delta p - \gamma)^4 / [E(\Delta p - \gamma)^2]^2 = \text{skewness}^2 + \text{skewness} \cdot \phi(1+\phi) / [\phi^{0.5}(1-\phi)^{0.5}]$. This result supports our point that the kurtosis measure as presented by Alvarez, Bihan, and Lippi (2016) is related to the skewness measure as presented in this paper. However, the skewness measure has an empirical advantage, which is discussed in Section 2.

distribution, while we analyze the relation between inflation and different moments of the *price change* distribution, including frequency, dispersion, and skewness. The higher moments of the price change distribution display clear patterns that are different from those of the distribution of relative prices. We also believe that the distribution of price changes is particularly worth analyzing because it can be easily constructed and observed with price microdata. On the other hand, the distribution of relative prices requires settling on an aggregate price index with which to normalize individual nominal prices, which makes it difficult to deal with product heterogeneity, as pointed out by Nakamura et al. (2018).

1.2 Intuition for the Menu Cost Model

In the previous subsection, we exploited the fact that the distribution of price changes had only two mass points in that particular version of the continuous time model. As we will illustrate, the negative relationships between inflation and dispersion as well as inflation and skewness hold in discrete time models, which will be the focus of our analysis from now on. Indeed, the same patterns can be seen in menu cost models for which price changes take a range of values, more in line with the distributions seen in the data. We now present the intuition for the relationship between inflation and higher moments of the price change distribution.

Idiosyncratic and aggregate nominal shocks move the price level of a firm away from its optimal level and yield a distribution of desired price changes. The presence of a menu cost means that only desired price changes above a certain size will actually occur. Thus, the realized price change distribution is the distribution of desired price changes with a band containing zero removed, as illustrated in Figure 1.

Nominal aggregate shocks that change the desired price of every firm by the same amount change the position and shape of the realized nonzero price change distribution. For example, the price change distribution is symmetric in an economy with zero inflation as presented in the middle panel of Figure 1. A positive aggregate shock, however, moves the desired price change distribution to the right and leads to higher inflation (illustrated in the right panel of Figure 1). The same shock also shifts the realized price change distribution to the right in a way that forms a left tail due to the presence of the inaction region. As inflation rises, this left tail is left further and further to the left of the average price change, while price changes are more concentrated on the right side of the inaction region. Consequently, the distribution becomes more asymmetric; the skewness of the distribution is more negative. The opposite is true after a deflationary shock (as illustrated in the left panel of Figure 1). Thus, inflation is negatively correlated with skewness of the price change distribution.⁹ Nonetheless, these responses do not occur in a Calvo model: in such a model, every desired price

9. Notably, the relationship between skewness and inflation is nonmonotonic during extreme inflation scenarios: as inflation approaches infinity, the skewness increases and approaches zero, as the selection effect plays little role when the desired price change distribution shifts far to the right and almost all price changes are positive (illustrated in Figure 10 in Online Appendix A). However, we do not observe this kind of hyperinflation in our sample.

change has a fixed probability of being realized, so aggregate shocks have a minimal influence on the shape of the realized price change distribution.

Another implication is that positive aggregate shocks reduce the dispersion of price changes. When a bigger fraction of the realized price changes are on one side of the inaction region, price changes are relatively close to each other. It is when the share of price changes on either side of the inaction region is equal that the dispersion is highest. Thus, the dispersion decreases with inflation in the positive region, and increases in the negative region, with the maximum attained at zero inflation. The intuition for this relationship has been applied by Vavra (2013) to explain why, in standard menu cost models, the frequency and the dispersion of price changes move in opposite directions in response to aggregate shocks. What we show here is that the same logic leads to observable inflation–dispersion and inflation–skewness relationships.¹⁰

What makes these correlations interesting is that they have to do with the central mechanism of the menu cost model: the selection effect. Indeed, monetary nonneutrality is low in menu cost models (relative to the Calvo model) because the menu cost ensures that only large price changes occur, which makes the average price response to aggregate shocks (and therefore aggregate flexibility) relatively high. The same mechanism of the selection effect leads the price change distribution to become less dispersed and more asymmetric (lower skewness) in response to inflationary shocks. This makes the correlations that we emphasize particularly informative about the strength of the selection effect.

Naturally, the selection effect has received much attention in recent research on sticky prices, as it makes a crucial difference to the degree of monetary nonneutrality. However, the fundamental difficulty in empirically evaluating the strength of the selection effect is that it involves the desired price change of firms. Since most firms' prices do not change in any given month, the desired price change is unobserved in most cases.¹¹ This makes it impossible to directly test whether the prices that change are those that are most misaligned, in line with the selection effect. Instead, one must make an inference based on the implications made by models for realized (and therefore observable) price changes. In this paper, we are presenting and implementing a new way of testing for the strength of the selection effect: the presence of selection in menu cost models implies the negative skewness and dispersion correlations.

1.3 Existing Models

In this subsection, we consider the empirical implications of the selection effect in a set of existing sticky price models that can be separated into four categories:

10. Relatedly, simpler statistics, such as the frequency of price increases, or the fraction of price changes that are increases, are not helpful to distinguish between models. Indeed, while it is true that a positive aggregate shock leads to more price increases in a menu cost model, the same is true in a Calvo model. What is different about the dispersion and skewness is that because they capture the shape of the distribution, they are affected by the presence of the inaction region in menu cost models.

11. Eichenbaum, Jaimovich, and Rebelo (2011) study the selection effect using a grocery store data set containing input price changes. However, the exact desired price change is still missing, and their sample is restricted to grocery products.

(i) Calvo, (ii) menu cost, (iii) observation cost, and (iv) rational inattention. We choose six models in those categories to evaluate, namely, the standard Calvo (1983) (augmented with idiosyncratic shocks), Golosov and Lucas (2007), Nakamura and Steinsson (2010), Midrigan (2011), Alvarez, Lippi, and Paciello (2011), and Woodford (2009).

The menu cost models that we consider have a common basic structure: firms produce a differentiated output with labor and a production technology subject to idiosyncratic shocks. In addition, they face constraints on changing their nominal price. Different models introduce different constraints, and in some cases different processes for the idiosyncratic shocks. All models, however, include aggregate nominal demand shocks. By shifting marginal costs, the aggregate shocks shift the desired price of all firms. However, since the constraints to changing prices are different across models, the response of prices (both the average price change and the distribution of price changes) will also be different across models. Below we provide a formal set-up of the models.

First, households choose consumption (C_t) and labor supply (L_t) to maximize expected discounted utility of the following form:

$$E_t \sum_{\tau=t}^{\infty} \beta^{\tau-t} [\log C_{\tau+t} - \omega L_{\tau+t}]. \quad (8)$$

There is a continuum of monopolistically competitive firms, indexed by z , producing a differentiated product. Aggregate consumption is given by a constant elasticity of substitution aggregator, meaning that each firm faces the standard demand function for its good:

$$c_t(z) = \left(\frac{p_t(z)}{P_t} \right)^{-\theta} C_t, \quad (9)$$

where θ is the elasticity of demand, and P_t is the CES price aggregator. Firms produce output based on a linear production function, with labor as the only input:

$$y_t(z) = A_t(z)L_t(z). \quad (10)$$

Productivity is subject to idiosyncratic shocks, which have been an important feature of sticky price models since Golosov and Lucas (2007). Large idiosyncratic shocks make it possible for such models to match the significant heterogeneity and high average size of price changes observed in the data, which was documented notably by Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008). Following Midrigan (2011) and Vavra (2013), we assume that idiosyncratic shocks arrive infrequently with a Poisson probability p_ϵ , and model the process in the following way:

$$\log A_t(z) = \begin{cases} \rho \log A_{t-1}(z) + \epsilon_t, & \text{with probability } p_\epsilon \\ \log A_{t-1}(z), & \text{with probability } 1 - p_\epsilon \end{cases}, \quad \epsilon_t \stackrel{iid}{\sim} N(0, \sigma_\epsilon^2). \quad (11)$$

As Midrigan (2011) has noted, this Poisson set-up allows the model to imply a distribution of price changes with fatter tails than the standard AR(1) productivity (used by Golosov and Lucas 2007, Nakamura and Steinsson 2010, for example), which is closer to what is seen in the data. It also nests the AR(1) set up when the probability of a shock occurring (p_ϵ) is 1. Since we will consider various models with AR(1) productivity, as well as Midrigan (2011)'s model with Poisson shocks, we maintain this set-up, and cover the different models by adjusting the relevant parameters.

In order to generate aggregate fluctuations, the sticky price models that we look at incorporate a stochastic process for nominal aggregate demand. Again, we stick to what is most often used in the literature by modeling nominal output as a log random walk with drift,

$$\log S_t = \mu + \log S_{t-1} + \eta_t, \quad \eta_t \stackrel{iid}{\sim} N(0, \sigma_\eta^2), \quad (12)$$

where $S_t \equiv P_t C_t$. This process stands in for monetary policy in these models: nominal output is determined exogenously, and therefore firms' price responses to these shocks determine how inflation and real output respond. We will use the same parameter values for this process (to match the behavior of U.S. aggregate activity) across the different models. Also, we define monetary nonneutrality as the variation in aggregate real consumption induced by the nominal shocks. This has become the main way of introducing monetary variables in the menu cost literature because it lends itself much more easily to the global solution methods that are used for such models than explicitly incorporating systematic monetary policy. Although Blanco (2016) developed a menu cost model with a Taylor-type policy rule, we do not attempt this for the models in this section. Our goal is to show how the price change distribution changes with inflation under different sticky price models, and the aggregate demand process that we use enables us to do this.

The general price-setting constraint takes the form of a (potentially time- and firm-varying) cost in terms of units of labor that must be paid for a firm to change its nominal price. Specifically, the period profit function therefore takes the form:

$$\Pi_t(z) = p_t(z)y_t(z) - W_t L_t(z) - \chi_t(z)W_t I\{p_t(z) \neq p_{t-1}(z)\}. \quad (13)$$

In the standard Golosov and Lucas (2007) menu cost model, the cost χ is fixed for all firms and periods, and can be calibrated to match the frequency of price changes observed in the data. The idiosyncratic shock process is normal AR(1), so $p_\epsilon = 1$. The standard deviation of shocks is calibrated to match the average size of price changes. This is, in a way, the most "state-dependent" model, as firms are fully in control of the decision of when to change the price for each good (subject to the constant menu cost).

The first extension to the menu cost model that we consider is the Nakamura and Steinsson (2010) multisector menu cost model, in which firms are separated into sectors. Firms in different sectors face a different menu cost and variance of idiosyncratic shocks. Second, we also analyze the model in Midrigan (2011), who introduced other

modifications to the standard menu cost model: first by changing the idiosyncratic shock process so that it would feature fat tails (which we described above), and giving firms a motive to make small price changes.¹² In his model, multiproduct firms can change the prices of all their products by paying the menu cost once. This enables the model to match the considerable fraction of small price changes that are observed in the data, but it also makes the model much more difficult to solve. We follow Vavra (2013) in simplifying the Midrigan model by assuming that, instead of producing multiple products, firms each period are randomly given the possibility of changing their price for free (with a low probability), or by paying a menu cost. The random menu cost structure yields similar results for monetary nonneutrality as introducing multiproduct firms. This is also a variation of the CalvoPlus model presented by Nakamura and Steinsson, and adds the probability of drawing a zero menu cost (free price change, p_z) as an additional parameter to calibrate. With the additional parameters in this model, we target the fraction of price changes that are small, as in Midrigan.¹³

We also consider a Calvo model, which has the set-up described above, except that every period firms have a fixed probability p_z of receiving the opportunity to freely change their price and a probability $1 - p_z$ that the price will remain unchanged. This is equivalent to the simplified Midrigan model that we describe, but with the high menu cost set to infinity, and the probability of a free price change set to equal the average frequency of price change in the data. This model includes idiosyncratic shocks to obtain a distribution of price changes, and we also set the variance of these shocks to match the average size of price changes.

Finally, we also include two models involving imperfect information: the Alvarez, Lippi, and Paciello (2011) model of observation and menu costs, and the rational inattention model of Woodford (2009). In the former, firms must pay a fixed cost to observe the relevant state (or conduct a “price review”), and a menu cost to change their price. Facing such costs, firms conducting a price review choose the date of the next review, and a price plan until that date. Because the Alvarez, Lippi, and Paciello model includes a menu cost, it features a high degree of selection. Woodford (2009) considers the same type of price-setting problem, but within the rational inattention framework proposed by Sims (2003): firms face a cost based on how much information they process, and therefore choose to receive limited information based on which they choose when to review prices. In this model, the cost of processing information is a crucial parameter, and both the Calvo model and standard menu cost model are nested as extreme cases of the information cost in this set-up (infinite and zero, respectively). Furthermore, intermediate values of the information cost result in what is

12. In Midrigan’s model, firms can also carry out temporary price changes, or sales, by setting regular prices and posted prices that can be different from each other. However, this feature of the model does not have a major effect on monetary nonneutrality, and we abstract from temporary price changes in our analysis.

13. Midrigan defines a small price change as a price change that is less than half, in absolute value, of the average size of price change. Due to the variation in the average size of price changes over time and across sectors, we prefer to use an absolute measure, and focus instead on the fraction of price changes that are smaller than 1% in absolute value.

described as a “generalized Ss model”: while a simple Ss model involves a threshold rule for price adjustment, a generalized Ss model features a probability of price adjustment as a function of the degree of price misalignment. This is the kind of model that we work with in Section 3, and we view the rational inattention framework as a potential microfoundation for this.

As mentioned in the introduction, the studies that have examined price change statistics in high-inflation environments have mostly focused on whether the frequency of price change rises with inflation, as the menu cost model predicts. Motivated by the logic explained above about the implications of the selection effect for the shape of the price change distribution in menu cost models, we also consider the dispersion and skewness of price changes. We do this in two different ways: by analyzing short-run fluctuations in inflation and changes in the value of steady-state inflation. Notably, the kind of analysis that we can carry out with Alvarez, Lippi, and Paciello (2011) and Woodford (2009) is more restricted than the perfect information models. We provide details on the simulation procedure for these two models in Online Appendix B.

First, to analyze short-run fluctuations, we solve each model with a fixed value for the parameters of the nominal aggregate demand process (μ and σ_η), and simulate each for a large number of firms and periods. From the simulated price series, we then compute the various price change moments for each period (obtaining a time series for each moment), and look at the relationship with the time series for inflation endogenously derived. Second, we perform steady-state analysis that is more in line with what is done by other papers (such as Alvarez et al. 2019) and analyze the long-run relationship between inflation and price change moments. Because much of the variation in inflation throughout our sample period is generally understood to reflect regime changes caused by systematic changes to the conduct of monetary policy, it is important to consider whether the correlations in question are the same when it is steady-state inflation that changes. For this analysis, we solve each model with different values for the steady-state inflation parameter (μ , keeping all other parameters fixed), and for each solution compute the values of the price change moments from the model’s stationary distribution. We find that the relationships between inflation and price change moments are qualitatively the same based on both short-run and long-run inflation analyses. The parameter values used to solve and simulate the models are listed and discussed in Online Appendix B.

To illustrate results with short-run inflation fluctuations, we present scatter plots between inflation and the different moments of the price change distribution from the simulations in Figure 3.¹⁴ We simulate each model for 1,000 periods and 50,000 firms. In each panel, the x-axis is the value of monthly inflation (so, e.g., a value

14. The Alvarez, Lippi, and Paciello (2011) model contains no aggregate shocks. Therefore, the “short-run” analysis of this model is excluded. Strictly speaking, the Woodford (2009) model cannot be solved with aggregate nominal disturbances. Nonetheless, we take a simplified approach following Section 5 of Woodford (2009). We simulate the model with the dynamics of aggregate nominal expenditure being i.i.d. and mean zero to conduct the “short-run” analysis (refer to Online Appendix B for detail). The “long-run” analysis of this model is excluded.

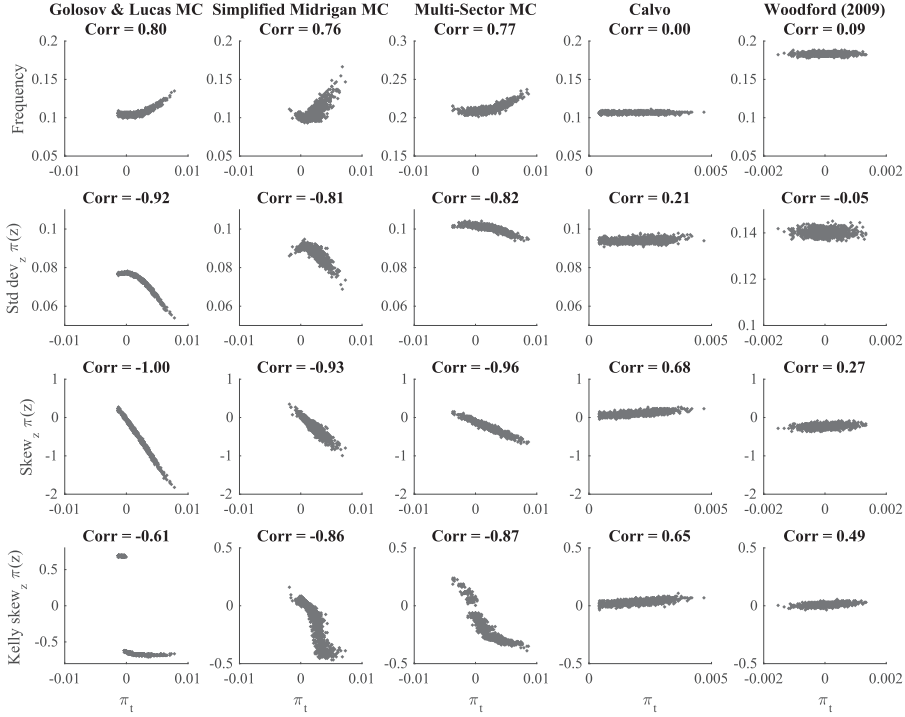


Fig 3. Simulated Moments and Inflation from Different Models.

of 0.01 represents 12% annualized inflation), while each point represents a time period in the simulation. These plots bring out the fact that in the menu cost models, the inflation–dispersion relationship and the inflation–skewness relationship are very clear and strong (especially in the Golosov and Lucas 2007, model for the dispersion): the skewness of price change falls very sharply with inflation in menu cost models, as does the dispersion for positive values of inflation. Moreover, the inflation–dispersion relationship is nonmonotonic, as explained in Section 1.2. In contrast, the same relations in the Calvo and imperfect information models are not so strong. However, the Calvo and rational inattention models feature weakly positive relationships for inflation–skewness as well as inflation–dispersion. That is because price changes are not selected in the Calvo model, so the mechanism described earlier is entirely absent.

The intuition for the correlations is easiest to explain in the case of the “standard” Golosov and Lucas model, as in Section 1.2, yet it also applies to the other menu cost models. The multisector menu cost model can be thought of as sectors facing different inaction regions, with each sector behaving as described for the standard menu cost model. Therefore, the aggregate price change distribution behaves similarly to how each sector’s distribution does. Our simplified version of the Midrigan model involves firms randomly facing either a positive or zero menu cost. This weakens the selection

effect, because there is now a positive probability that a firm will change its price even if it will be a small change, so that price changes are not entirely “selected” based on how out of line the original price is. However, the selection effect is still present to a certain extent, because it is only relatively large price changes that will happen with certainty. The tails of the price change distribution will therefore be very sensitive to the aggregate shocks that drive inflation in the model, leading to the same relationships for price change dispersion and skewness as in the Golosov and Lucas (2007) model.

Although the relationships come out very clearly in these simulations, it could be a concern that the higher moments that we are estimating might not be well defined in the distributions that we are working with. In addition, estimates of higher moments are very sensitive to outliers, which would be of concern particularly when we estimate from the data. That is why we also consider alternative measures for the dispersion and skewness of price change: the interquartile range (for dispersion) and Kelly’s coefficient of skewness (as opposed to “moment skewness,” which is what we have been estimating so far).¹⁵ Since these statistics are quantile based, they are well defined for any distribution, and they are also less sensitive to outliers. The correlations of inflation and higher moments are similar as the previous predictions. The last row of Figure 3 shows scatter plots of Kelly Skewness in the different models.¹⁶

In Figure 4, we plot the results for the long-run analysis, in which we vary the value of steady-state inflation. For each model solution, we can construct a stationary distribution of price changes, from which we can then compute the stationary value for the different price change moments, and these are the values plotted in the figure. What the scatter plots show is that, as in the “short-run” analysis, the dispersion and skewness of price changes fall with trend inflation in the menu cost model; meanwhile, the Calvo model predicts weak positive relations for both moments with respect to the steady-state inflation. These findings will be important when comparing the skewness of price change between the low- and high-inflation periods in the data.

To conclude our theoretical analysis, we emphasize that the correlations that we consider all have the same sign in the four menu cost models, namely, Golosov and Lucas (2007), Nakamura and Steinsson (2010), Midrigan (2011), and Alvarez, Lippi, and Paciello (2011). The scatter plots show that the values taken by moments we report do vary across the models, but it is notable that the sign and strength of the correlations across the models are similar. Indeed, the Nakamura and Steinsson and Midrigan menu cost models were developed as extensions of the Golosov and

15. These statistics are defined as follows, with Q_i representing the i th percentile. Interquartile range = $Q_{75} - Q_{25}$. Kelly skewness = $[(Q_{90} - Q_{50}) - (Q_{50} - Q_{10})]/(Q_{90} - Q_{10})$. Kelly skewness essentially measures the degree of asymmetry in a distribution, comparing the size of the right and left tails.

16. There is a discontinuous jump in the Kelly skewness values for the Golosov and Lucas (2007) model because the median price change (which is used to compute Kelly skewness) jumps discretely from the left to the right band of the inaction region. The jump also corresponds to a value of approximately 0 inflation, as that is consistent with an equal share of price increases and decreases. However, within the positive (or negative) inflation periods, the relationship between inflation and Kelly skewness is still negative.

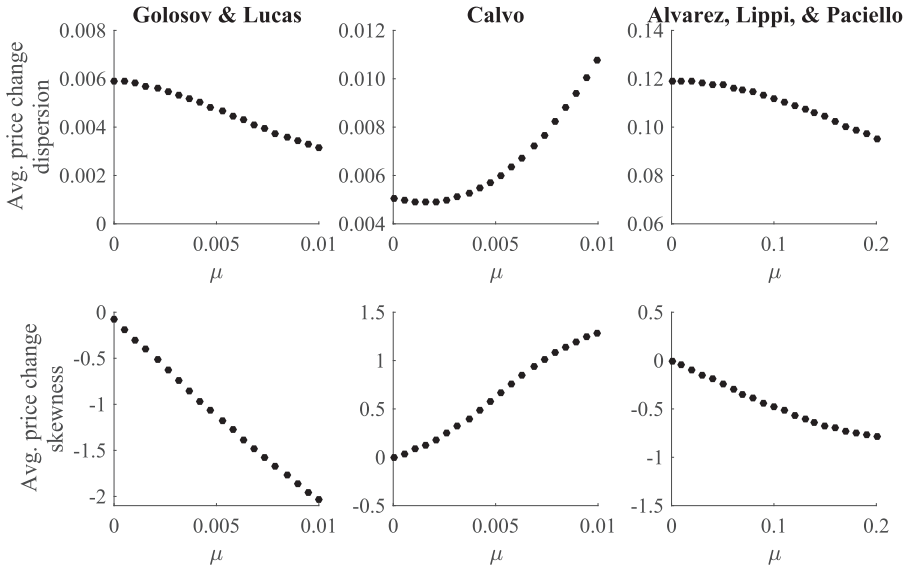


Fig 4. Simulated Long-Run Statistics from Different Models.

Lucas model to make it match new empirical facts, and the changes made considerably weakened the selection effect that reduces the importance of monetary shocks. However, what we find here is that, despite the important changes made, they all have the same implications along the dimensions that we are considering.

2. EMPIRICAL EVIDENCE

We next present the data set and the empirical results that are used to test the model predictions of Section 1.

2.1 Data Set and Construction of Statistics

Along with much of the sticky price literature, we make use of the microdata that underlies the U.S. CPI. The CPI Research Database collected and maintained by the U.S. Bureau of Labor Statistics (BLS) contains about 80,000 monthly prices collected from around the U.S., classified into about 300 categories called Entry Level Items (ELIs). As mentioned before, the data going back to 1988 has been available for a little over a decade. The data going back to 1977 have recently become available, and this is the novel part of the data set that we use extensively. This new data set has thus far only been used by Nakamura et al. (2018), and that paper also describes in detail just how the data set was reconstructed. We have access to the variables that identify specific

products, and that reveal when a substitution has occurred (i.e., when a new version of a product has replaced the old one). In addition, the data set contains information on when any given price is a temporary sale, or an imputation (not properly collected). Because of this, we are confident that we are observing the price changes of identical products and services, with the price being actually observed, and all of this with the same standards throughout the sample period.

In order to test the predictions that we presented in the previous section, we construct distributions of price changes for each month, from which the different moments of interest can be estimated period by period. We calculate the log price change for all the goods and services in our sample, and then construct the distributions subject to a few restrictions. We keep only nonzero price changes to compute the dispersion and skewness (while the frequency measures the fraction of nonzero price changes), and exclude temporary sales, substitutions, along with price changes that are implausibly large in absolute value. We provide further details on these restrictions in Online Appendix C.

To control for the heterogeneity of price change statistics across sectors, we follow Nakamura and Steinsson (2008, 2010) to report the average overall frequency of price change by taking a weighted average of the ELI-level frequencies (using the expenditure weights that go into the CPI). We consider both the aggregate weighted median and mean frequency.¹⁷ For the dispersion and skewness, we follow a similar approach: estimate each moment by sector-month and then take a weighted average. However, we do not categorize products by ELI's but categorize them into 13 "major groups," based on the classification of Nakamura and Steinsson (which is listed in Online Appendix C and is also used by Vavra 2013). The reason is that ELIs are fairly narrow categories with number of price change observations insufficient to estimate higher moments. The major group classification allows us to separate goods and services into similar categories, while leaving enough observations in each sector-month to obtain good estimates of the dispersion and skewness.^{18,19}

Measuring price change statistics using sectoral average can also control for the effects of unobserved sectoral shocks. When all price changes across sectors are pooled together, sector-specific shocks have the potential to strongly affect the shape of the

17. Nakamura and Steinsson (2008) highlight the difference between the mean and the median, arising from the fact that the distribution of frequencies by ELI is very skewed to the right, with a few ELI's having very high frequencies. They argue that the median is a better measure of the average frequency, as a single-sector menu cost model calibrated to match the median frequency better approximates a multisector model. In this way, the median frequency is a statistic that better describes the degree of price stickiness. Therefore, we calibrate all the single sector models to match the median frequency.

18. We present the average sample size by major group in the Online Appendix C. The weighted median of monthly nonzero price change observation per major group is slightly under 1,600, averaged across all months. There is also significant heterogeneity of sample size across major groups.

19. While our analysis focuses on the per period weighted average of price change skewness and dispersion, another way of evaluating the empirical relationship with inflation would be to use panel regressions at the ELI or major group level. We do not follow this approach for the following two reasons. First, sectoral measures of inflation are more subject to sampling errors than the measure of aggregate inflation. Second, types of pricing constraints vary across sectors, that could imply different signs for inflation-skewness and inflation-dispersion correlations. In this paper, we focus on the average moments to analyze the aggregate economy.

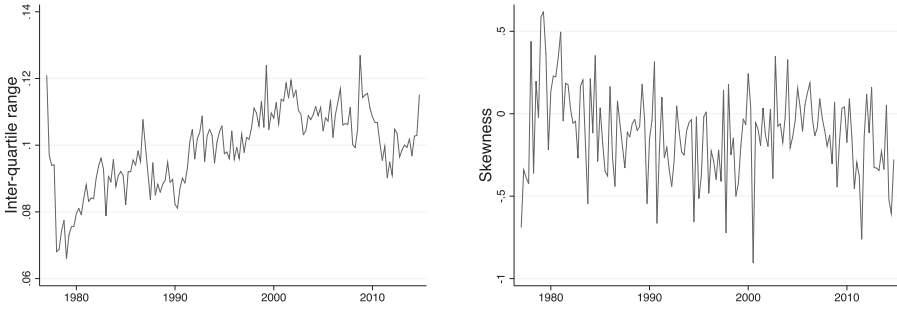


Fig 5. IQR and Skewness of Price Change Distribution, Quarterly.

overall price change distribution and its moments variation over time. To test the model predictions that we focus on, we attempt to control for these types of effects by computing statistics sector by sector.

Our preferred measure for aggregate inflation is monthly core Personal Consumption Expenditures (PCE) inflation. Sharp changes in headline inflation tend to be driven by the global market prices of food and commodities, which would not be well described by the price-setting models that we are working with, making core inflation preferable for us. However, we also compute correlations of price change moments with headline (CPI) inflation as well as estimates the moments excluding price changes from food and energy categories as robustness checks in Online Appendix D. Finally, to control for seasonality in moment series, we calculate the correlations after removing the seasonal factor using monthly dummies, and after applying a moving average smoother to them in Online Appendix D.

Given that the price data and inflation series are monthly, we can compute the correlations at a monthly frequency. However, the drawback of using monthly series is that the small number of observations makes the estimation of monthly price change moments less precise (especially for higher moments such as the dispersion or skewness). The alternative is to group price change observations and to estimate the moments by quarters or years, which gives us more precise estimates. Quarterly and annual inflation averages also have the advantage of containing less noise than monthly inflation series, so we will focus on presenting results using quarterly series. Figure 5 plots the quarterly time series that we construct for the interquartile range and skewness of price changes. Monthly and annual results are presented in Online Appendix D.

2.2 Results

The goal of our empirical work is to determine whether the theoretical patterns documented above are borne out by the data. As in Section 1, we focus on the correlations between aggregate inflation and price change dispersion, and between

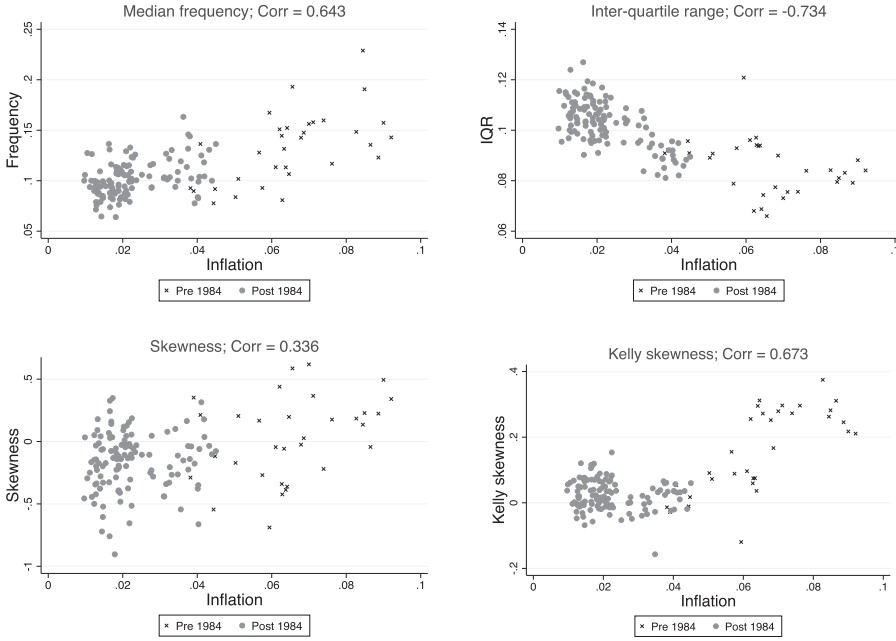


Fig 6. Moments of Price Change and Inflation, Quarterly.

inflation and price change skewness. We present the correlation results in two ways, scatter plots of raw correlations and regressions with additional controls.

First, we present scatter plots using the quarterly moment and inflation series in Figure 6 (as the empirical counterpart to the simulation scatter plots of Figure 3). Correlation values are reported in Table 10 in Online Appendix D. We first verify that the frequency of price change rises with inflation, as found by Gagnon (2009) and Alvarez et al. (2019). As argued in the previous studies, this provides strong evidence against the Calvo assumption of time-dependent price setting.

Next, we investigate the inflation–dispersion and inflation–skewness relationships.²⁰ Our main results is that while there does seem to be a clear negative relationship between inflation and dispersion, there is no such relation between inflation and skewness. Indeed, for both measures of skewness (moment skewness and Kelly skewness), the correlation is either strongly positive (over the whole sample period) or close to zero (post-1984), as presented in Figure 6. Skewness, while varying over time, does not change with inflation in a systematic way for low levels of inflation (although there does seem to be a positive relationship when inflation is high). Inspection of the series of Figure 5 already illustrates the patterns. The left panel shows

20. The results reported here focus on the inter-quartile range as the measure of dispersion. Although it is somewhat less volatile than the standard deviation, all the same patterns also hold with the standard deviation.

that the interquartile range is considerably lower in the high inflation period of the late 1970s and early 1980s. In the right panel, one can see that there was no similar drop in skewness in the high inflation period. If anything, skewness was higher then.

We now turn to regressions to determine whether these correlations are statistically significant, and to consider different control variables. The question of interest about the coefficients on inflation is not merely whether they are statistically significantly different from zero, but also whether they are significantly different from what the models predict. To do this, we estimate regressions of the frequency, dispersion (interquartile range), and skewness (both moment and Kelly skewness) of the price change distribution on inflation, with different specifications allowing for different sets of controls and sample periods. As before, we run the regressions both on the whole sample period and on only after 1984. This allows us to see if the relationship looks different between the low- and high-inflation periods. The regressions all take the following form:

$$y_t = \alpha + \beta\pi_t + \gamma Controls_t + e_t, \quad (14)$$

where y_t denotes the different price change moments (frequency, dispersion, and skewness). Controls are included to address the fact that many important changes occurred in the U.S. monetary environment over our sample period, which could conceivably have a direct effect on the price change distribution. For example, expected inflation could affect firms' price setting decisions separately from present realized nominal shocks, so we include expected inflation (measured by the University of Michigan Survey of Consumers) as a control. Because the survey of inflation expectations asks about expectations of headline inflation specifically, we use headline CPI inflation in regression analyses here. Additional regression analyses presented in Online Appendix D use PCE inflation as robustness checks. We also include dummy variables for the different Federal Reserve chair's times in office, to control for differences in the conduct of monetary policy. The different specifications cover different combinations of controls (expected inflation only, Fed dummies only, or Fed dummies with expected inflation) and the different periods. Results are reported in Table 1.

These regression results support what the three correlations showed in the scatter plots. (i) The frequency of price change rises with inflation. (ii) The relationship between dispersion and inflation is negative and statistically significant in all specifications and sample periods. (iii) The skewness correlation is significantly positive for the whole sample, but not significantly different from zero when the early high-inflation period is excluded. In particular, the finding about the positive inflation-skewness correlation in high inflation periods is important, as it goes against the menu cost models' predictions at high values of steady-state inflation, as shown in Figure 4.

Table 2 presents the regression estimates of price change moments on inflation using simulated data from the different models. The last row presents the coefficients using CPI data, which replicates the third column of Table 1. The first three menu cost models have negative coefficients for the interquartile range, although for all but the

TABLE 1
COEFFICIENTS ON INFLATION FOR PRICE CHANGE MOMENTS—USING CPI DATA

	1977–2014			1985–2014		
	Both controls	Fed dummies	Inflation only	Both controls	Fed dummies	Inflation only
Frequency	0.335* (0.156)	0.659** (0.117)	0.685** (0.100)	0.276 (0.157)	0.463** (0.134)	0.444* (0.154)
IQR	−0.237** (0.064)	−0.199** (0.038)	−0.308** (0.037)	−0.140* (0.059)	−0.304** (0.060)	−0.301* (0.064)
Skewness	1.366 (2.070)	3.324** (1.032)	3.610** (0.963)	1.869 (2.272)	0.820 (1.385)	1.212 (1.314)
Kelly skewness	0.924 (0.661)	2.289** (0.420)	2.527** (0.419)	0.776 (0.424)	0.589 (0.327)	0.526 (0.399)

NOTES: Significant ** at 1% level (* at 5% level). The regressions are run using quarterly series, where quarterly inflation is defined as the mean of the 12-month log changes in the CPI for the 3 months in every quarter. The different cells indicate different specifications, which change with respect to the sample period used and what controls are used. Standard errors that are corrected for heteroskedasticity and autocorrelation of the residuals (Newey–West) are reported.

TABLE 2
COEFFICIENTS ON INFLATION FOR PRICE CHANGE MOMENT—USING SIMULATED DATA

Model	Frequency	IQR	Skewness	Kelly skewness
Golosov and Lucas (2007)	0.14	−0.94	−17.70	−0.40
Nakamura and Steinsson (2010)	0.14	−0.22	−5.39	−4.33
Simplified Midrigan	0.35	−0.90	−9.84	−6.53
Calvo	0.00	0.04	2.93	1.00
Woodford (2009)	0.02	0.03	2.87	1.00
BLS CPI data	0.68	−0.31	3.61	2.53

multisector model, they are outside the 95% confidence intervals of the coefficients that we estimate using the CPI data. However, the disagreement with the data is much starker with the skewness coefficients. These are all very far outside the confidence intervals that we estimate using the CPI data under all specifications, and the same is true for Kelly skewness.²¹

Finally, all these patterns hold true regardless of whether we exclude potentially spurious small price changes (as defined by Eichenbaum et al. 2014) or apply seasonal adjustment and smoothing to the data series (refer to Online Appendix D for detail). Similar patterns of higher order moments of price changes have been found in Klenow and Malin (2010) using price micro data for the three largest cities in the U.S. CPI from 1988 to 2009. They find a negative variance–inflation correlation (−0.29 for

21. The one exception is the coefficient for the Golosov and Lucas (2007) model, which is much smaller in magnitude than in the other menu cost models, and is marginally accepted in the specification that restricts the sample to the post-1984 period and uses only Fed chair controls. It appears that the value of the Kelly skewness is extremely sensitive to the unusual shape of the price change distribution (bimodal) in this model, leading to this weak relationship. The model's Kelly skewness coefficient is still rejected in all the other specifications, however.

regular prices and -0.16 for posted prices) as well as a positive skewness–inflation correlation (0.15 for regular prices and 0.17 for posted prices).

As our empirical results have shown, while the dispersion of price changes falls as inflation rises, there is no evidence for such a relationship with the skewness of price change. This goes against the predictions of a broad class of menu cost models and calls into question the importance of the selection effect that these models emphasize. We will show in the next section how weakening the selection effect can reconcile menu cost models with the data, while also raising the implied degree of monetary nonneutrality.

Before proceeding, we compare our approach to that of Alvarez, Bihan, and Lippi (2016), who show that in a broad class of sticky price models nonneutrality can be measured by the kurtosis of price changes. Although the theoretical result of Alvarez, Bihan, and Lippi is noteworthy, a significant advantage of our approach lies in its practical implementation of the empirical estimation. The key statistic in Alvarez, Bihan, and Lippi is the kurtosis of price changes. However, the result on the degree of monetary nonneutrality will be only as reliable as the estimate of price change kurtosis is precise. Obtaining precise estimates for the kurtosis is difficult for a few reasons, not least of which is the simple fact that moment estimates are less precise in finite samples the higher is the order of the moment. Perhaps more important in practice, however, is the fact that estimates of kurtosis are potentially very sensitive to just how they are computed. For example, Alvarez, Bihan, and Lippi point out that in order to compute the kurtosis from a combined sample of potentially heterogeneous distributions, it is necessary to normalize each price change observation in order to obtain an unbiased estimate. As we show in Luo and Villar (2017a), the kurtosis of price change estimated from the U.S. CPI data varies considerably depending on just how this normalization is done, and over time. Naturally, estimates of the skewness are also subject to such concerns (although to a slightly lower degree because of the lower order of the moment). However, the advantage of our approach is that it emphasizes correctly estimating the sign of the relationship between skewness and inflation, and not the precise value of skewness. As we have shown, the various specifications that we use lead to the same results in what matters for nonneutrality and for evaluating the models, which is that skewness does not fall with inflation. In this sense, our approach yields more consistent results when implemented in practice.

3. A RANDOM MENU COST MODEL

In this section, we present a random menu cost model (based on Dotsey, King, and Wolman 1999) that matches the empirical patterns we observe.

3.1 *Background on Random Menu Costs*

The model has a similar set-up as the menu cost models presented in Section 1 with the same demand system and technology faced by firms. However, we generalize the

price setting problem by making the menu cost random. Formally, the period profit function of the firm is,

$$\Pi_t(z) = p_t(z)y_t(z) - W_tL_t(z) - \chi_t(z)W_tI\{p_t(z) \neq p_{t-1}(z)\}, \quad \chi_t(z) \stackrel{iid}{\sim} G(\chi). \quad (15)$$

The difference with the Golosov and Lucas (2007) model is that here the menu cost can vary over time and across firms; the difference with the Midrigan (2011) model is that the distribution of menu costs is generalized, and as opposed to the Nakamura and Steinsson (2010) model, the menu cost for any given firm here varies over time.²²

The assumption of random menu costs we propose is similar to that made by Dotsey, King, and Wolman (1999), but we present it within the framework of the models in Section 1.²³ Also, the random menu cost model is connected to the smoothly state-dependent pricing (SSDP) model proposed by Costain and Nakov (2011). The probability of price adjustment depends on the expected gain of adjustment in the SSDP model, which is observationally equivalent to a random menu cost model. In future work, Costain and Nakov (2018) show that a similar process for price adjustment can be derived from a set up in which firms face costly decision making. Costain and Nakov (2011) also find that the price adjustment probability function must be such that the degree of selection is weak in order to match the data (notably the presence of small price changes). In Online Appendix E, we provide a more detailed discussion about the difference between the SSDP model and the random menu cost model presented in this paper.

Our approach also has a close relation to a hazard function approach, which is more general and has been pursued in a series of papers by Caballero and Engel (1993, 2006, 2007). They propose considering price adjustment through the price adjustment hazard function of the current price deviation from its optimal value (denoted by p^*),

$$H(x) = P(\Delta p \neq 0 | p^* - p = x). \quad (16)$$

Any of the models we have considered or a random menu cost model with a particular menu cost distribution would imply a particular price adjustment hazard function, as expressed above. The hazard function therefore determines aggregate price flexibility and monetary nonneutrality. In a separate paper (Luo and Villar 2017b), we show that the same data and empirical patterns can be used to estimate the price adjustment hazard function.

A more structural approach to price stickiness (that is also related to ours) is Woodford (2009)'s model of rational inattention, which we have described in Section 1.

22. This set-up can replicate the Golosov and Lucas (2007) model if the menu cost distribution is degenerate, the simplified Midrigan model if the distribution is discrete with two support points (one being zero, the other being positive), and the Calvo model if the higher support point is infinite. Since the Nakamura and Steinsson (2010) model involves different firms facing different menu costs that are fixed over time, it is not encompassed by our set-up.

23. The key differences with Dotsey, King, and Wolman (1999) are that their model does not include idiosyncratic shocks, that it does include capital as an input to production, and that they did not have a way of using information from price microdata to place restrictions on the menu cost distribution, which is what the present exercise is about.

He presents various adjustment hazard functions implied by different information costs. A menu cost model with inattention as a source of randomized discrete adjustment is observationally equivalent to a random-menu-cost model (see Woodford 2008, 2009). Hence, we believe that a decision-theoretic justification for this random menu cost model can be derived based on the rational inattention framework. As Woodford (2009) also points out, the direct empirical evidence on the actual costs of price adjustment put forth by Zbaracki et al. (2004) indicates that the most important part of those costs are related to the process of gathering the necessary information for a price review. In addition, Anderson and Simester (2010) give evidence on how price changes can antagonize consumers, which introduces costs to changing prices. To the extent that the menu costs in the menu cost framework represent these costs, we believe that menu costs are random to some extent, and vary across firms and time. This lends plausibility to our random menu costs assumption, although we leave the explicitly modeling of the informational constraints or consumer considerations that underly it to future research.

3.2 *The Distribution of Menu Costs*

Introducing random menu costs allows us to determine the extent of state dependence present in the model, or to what extent firms choose when to change their prices. An extreme case is perfect price flexibility (with the menu cost always equal to zero), while a slightly less extreme case is a fixed menu cost model, such as the one in Golosov and Lucas (2007). Adding randomness to the menu cost makes the price change decision more exogenous to the firm, as this additional dimension of the problem is outside the firm's control. The Midrigan model (both in Midrigan 2011, and the simplification of it that we present) goes in this direction. As a result, the degree of monetary nonneutrality in Midrigan is much higher. We interpret our results so far as indicating that a model would need even more exogeneity (but less than the Calvo model) to match the empirical facts that we have presented. Therefore, we parameterize the distribution of menu costs in a way that enables us to set the degree of exogeneity.

The distribution of menu costs would need two important features to generate a nonnegative inflation-skewness correlation: a nonzero probability of free price adjustment and a large probability of high adjustment cost. First, a positive probability of free price adjustment eliminates the inaction region in the price setting problem, as the firms with zero menu cost will choose to change their prices even if it is by a small amount. The Midrigan (2011) model already includes this, but still predicts a counterfactual inflation-skewness correlation. That is why another feature is required: there must be a positive and considerable probability that the menu cost will be very high, so high that firms will not choose to change their price when this high-level cost is realized. Indeed, in the existing menu cost models, the skewness of price changes falls with inflation because a positive aggregate shock induces more firms that face a positive menu cost to pay it, effectively pushing them over a threshold, leading to an important shift in the shape of the distribution. Having a positive probability of very

high menu costs means that fewer firms will be pushed over this threshold, weakening this effect. It is also helpful to note that the Calvo model contains both of these features in the extreme, as it gives a positive probability of a free price change, and in all other cases the menu cost is infinite. Because of this, we say that the menu cost distribution will incorporate a strong “Calvo feature”, without going all the way to the Calvo extreme.

In order to achieve this, we present a relatively flexible distribution for menu costs. Menu costs are i.i.d. across time and firms, so that every period each firm draws a menu cost χ from a mixed distribution. First, with a certain probability, the menu cost is zero, and otherwise it is drawn from a continuous distribution, that is,

$$\chi = \begin{cases} 0, & \text{with probability } p_z \\ \tilde{\chi}, & \text{with probability } 1 - p_z \end{cases}, \text{ where } F(k) = P(\tilde{\chi} \leq k) = 1 - e^{-\lambda k^\alpha}. \quad (17)$$

This distribution is a transformation of the exponential distribution (and is an exponential distribution when $\alpha = 1$), so that the random variable is always positive. α governs the curvature of the distribution function, which roughly corresponds to the fatness of the tail. Figure 17 in Online Appendix E shows how the shape of the cumulative distribution function changes with α . For our purposes, what is important is that for low values of α , the probability of very low menu costs is relatively high, but the probability of very high menu costs is also quite high. When α is high, these extreme probabilities are low, and as α rises, the density is concentrated on one value, approximating the case of a unique menu cost.

3.3 Calibration and Results

Our set-up has introduced two new parameters, relative to the models we have been considering, the inverse of the average menu cost (λ), and the curvature of the menu cost distribution (α). The other parameters important for the firm’s price setting problem are the variance of the idiosyncratic shocks (σ_ϵ^2), the arrival probability of those shocks (p_ϵ), and the probability of a free price change (p_z), which was used in the simplified Midrigan model. We set these parameters so that the model can match the empirical facts that we have discussed so far. The solution and simulation method follow the discussion of Online Appendix B.

First, our model will match the unconditional price change moments matched by existing models. These include the average monthly frequency of price change and the average size of price change. These have not been the focus of our discussions so far, but in order to compare the degree of monetary nonneutrality implied by different models, it is necessary to target the same values for these moments. Our model therefore matches the (expenditure-weighted) median of these statistics measured in our data.

Second, and in line with the focus of our paper, we will target the signs of the correlations between inflation and the different price change moments. As previous studies had shown (and we confirmed), the correlation between inflation and the frequency

TABLE 3
PARAMETER VALUES

Parameter	Description	Value
λ	Inv. average menu cost	0.177
α	Fatness of tail of MC	0.27
p_z	P(zero MC)	0.056
p_ϵ	P(idio. shock)	0.345
σ_ϵ	Size of idio. shocks	0.0967

TABLE 4
SIMULATION RESULTS

Moment	Model	Data
Avg. frequency	10.7%	10.7%
Avg. size	7.6%	7.5%
Corr(IQR, π)	-0.67	-0.68
Corr(skew, π)	0.19	0.33
Corr(freq, π)	0.67	0.68

of price change is positive, so our model also matches this fact. In addition, our model implies a strongly negative inflation–dispersion correlation (as menu cost models do). The novelty is that the implied inflation–skewness correlation is nonnegative, as in the data.²⁴

Table 3 presents the parameter values that we choose to match these moments, and Table 4 shows the moments attained by the model, compared to their empirical values. The first two moments are matched almost exactly. The model matches the inflation–dispersion and inflation–frequency correlations quite closely. In addition, we match a nonnegative inflation–skewness correlation, as opposed to a clear positive correlation as in the data for the whole sample. Figure 7 illustrates the scatter plots of these correlations.

While the skewness correlation in this model is lower than in the data, for the range of inflation that occurs in the simulations (0–6%), the correlation also appears

24. In this exercise, we do not target the relationship between the average size of price changes and inflation, which is an important part of the analysis in Nakamura et al. (2018). That paper highlights that the average size of price changes was no higher during high inflation periods, using the same data. As they emphasize, this is consistent with the menu cost model, but not with the Calvo model. Nonetheless, in Luo and Villar (2017b) we estimate a hazard function of price adjustment, and find that a hazard function with both Calvo and menu cost features can match the signs of all the correlations, including the average size of price changes being constant with respect to inflation. Given that the hazard function approach and the random menu cost approach are closely related, we believe that the random menu cost model presented in this paper can capture the zero inflation–size correlation if that moment is targeted. However, without targeting a zero inflation–size correlation, the random menu cost model estimated in this section predicts that the average size of price changes rises somewhat with inflation, but considerably less so than the Calvo model.

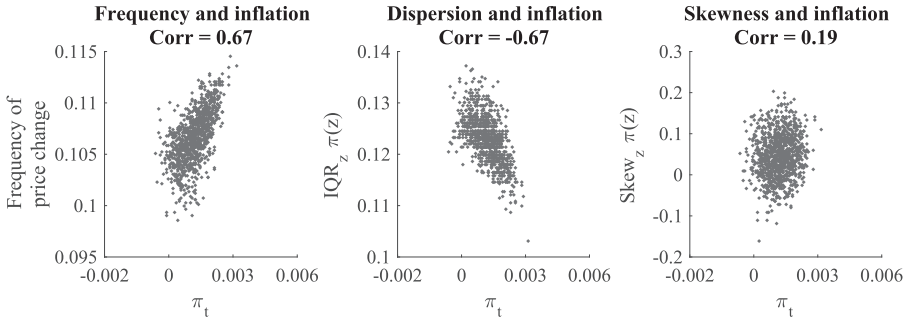


Fig 7. Scatter Plots, Random MC Model.

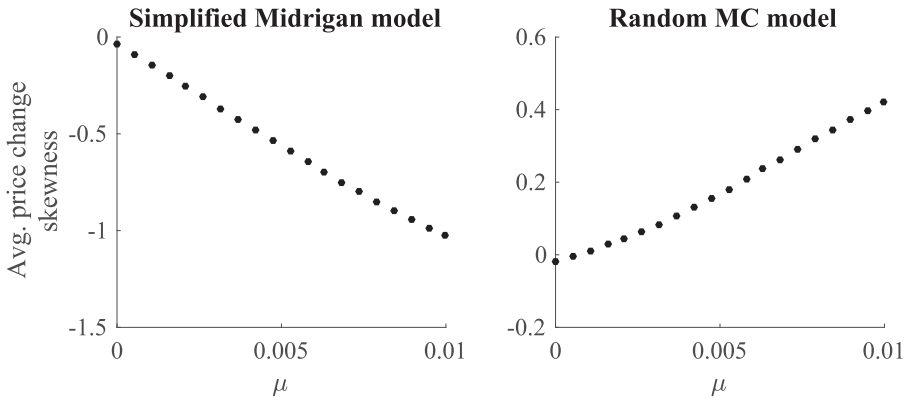


Fig 8. Steady-State Skewness Correlation.

to be close to zero in the data.²⁵ We carry out the same “long-run” analysis as in Figure 4: solving the model for different values of trend inflation. We find that for higher steady-state inflation, the average level of skewness in the price change distribution is larger. This result makes our model even more consistent with the data, as it shows that when steady-state inflation is higher (as it surely was in the early high-inflation part of our sample), skewness rises with inflation. This also makes our model stand out even more from the existing ones, as the other menu cost models feature a negative inflation–skewness correlation both in the short-run and in the long-run analysis. Figure 8 shows this clearly by plotting the steady-state skewness correlations

25. Inflation is less volatile and moves within a narrower range in our generalized model than in the other menu cost models, even though the parameters of the nominal aggregate demand process are the same. This is a direct result of the differences in monetary nonneutrality in the models, as higher nonneutrality means that the same nominal shocks have a greater effect on real consumption (and induce greater real variation), leading to less variation in inflation.

TABLE 5
MONETARY NONNEUTRALITY

Model	Var (C_t) * 10^4
Golosov and Lucas (2007)	0.058
Simplified Midrigan	0.170
Random menu cost	0.344
Calvo	0.517

for the simplified Midrigan model (as an example) and our heterogeneous menu cost model separately.

To check the robustness of the model implications to the size of the simulated sample, we implemented two variations of the random menu cost model simulations presented here. First, we simulate a random menu cost model with the number of firms close to the number observations in the average CPI sector (which is smaller than the total number of observations that we use in our empirical analysis). Second, we simulate a multi-sector random menu cost model by calibrating the model sector-by-sector and computing price change moments using a weighted sectoral average (as in our empirical approach). In both extensions, the patterns of the correlations are essentially the same as with the baseline approach. We present the details in Online Appendix E.

In summary, these results make clear that the random menu cost model weakens the selection effect, matches the important empirical facts that have been the focus of previous work on sticky prices, and overturns the counterfactual prediction of these models that we have emphasized. We next show what this means for the degree of monetary nonneutrality.

3.4 Monetary Nonneutrality

Monetary nonneutrality in these models is defined as the variation in real consumption induced by the nominal aggregate demand shocks.²⁶ We compare this statistic for the Calvo model, the Golosov and Lucas (2007), the simplified Midrigan menu cost models, and our random menu cost model. As we have explained, making the menu costs random in the way that we have proposed weakens the selection effect that is at work in menu cost models, so it is to be expected that our model would imply a greater degree of monetary nonneutrality. Table 5 provides a quantitative illustration of this.

As Golosov and Lucas (2007) had famously shown, their model features a trivial amount of monetary nonneutrality compared to the Calvo model. Between the menu

26. Since the nominal aggregate demand shocks are the only aggregate shocks in the model, any variation in real consumption is due to those shocks and reflects to what extent they have real effects. In addition, if money were fully neutral, the nominal shocks would have an effect only on the aggregate price level, and thus no effect on real consumption. This measure for monetary nonneutrality is used by Golosov and Lucas (2007), Nakamura and Steinsson (2010), and Midrigan (2011), among others.

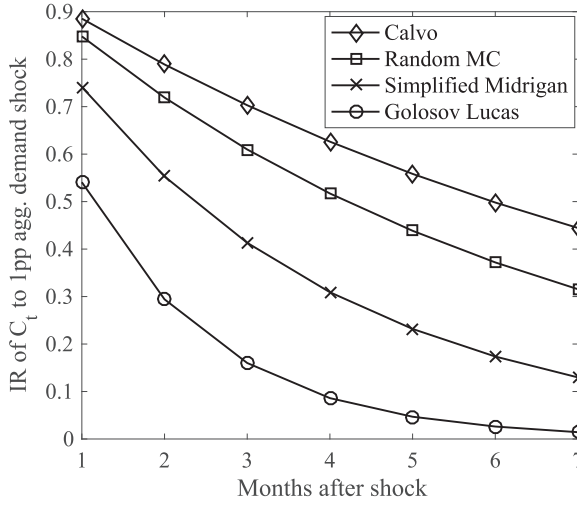


Fig 9. Impulse Responses of Consumption.

cost models, the major difference is between the baseline Golosov and Lucas and the others. Allowing for small price changes, as the Midrigan (2011) model does, leads to a very large increase in monetary nonneutrality. However, our generalized model goes further, and yields an even higher level of nonneutrality. The Calvo model still has a higher degree of monetary nonneutrality, but our model gets significantly closer to Calvo than the others. To further illustrate the differences between the models, in Figure 9 we plot the impulse response of real aggregate consumption to a 1 percentage point increase in nominal aggregate demand in the same four models. The effect on real activity is not only large, but also quite persistent in our model, and much more so than in the menu cost models. In this sense, our model is also much closer to the Calvo model.

4. CONCLUSION

The literature on sticky prices has paid considerable attention to the role of selection in price setting in determining the size of the real effects of monetary policy. Our paper has contributed to the debate on the importance of the selection effect by using new historical data from moderate to high inflation environments in the United States, and by focusing on statistics that have previously not been considered. Our main finding is that the menu cost models that have been most used in the literature fail to match the positive relationship between inflation and the skewness of price changes in the data, because they uniformly predict a sharp negative relationship. In addition, we argue that this relationship, although not obvious at first sight, follows

very intuitively from the selection effect that is central to menu cost models, and that makes these models imply relatively low monetary nonneutrality. We also show how a model with random menu costs can overcome this problem when the distribution of menu costs features a significant probability of very high and very low menu costs, making it resemble a Calvo model and weakening the selection effect. Finally, this model predicts a degree of monetary nonneutrality that is considerably higher than what is predicted by the Golosov and Lucas (2007) model, and higher still than the Midrigan (2011) model.

In the context of the debate between time-dependent and state-dependent pricing models, we follow Woodford (2009) in presenting the distinction between time- and state-dependent models as a continuum, or spectrum. Woodford shows how different values for the firm's cost of processing information leads to a different point on this spectrum. In contrast, our approach is agnostic as to what ultimately underlies the randomness of menu costs that allows our model to span the time- versus state-dependent spectrum. Instead, our contribution is to determine what point on the spectrum is most consistent with the data. Future research could combine these two approaches to gain a better understanding into the nature and importance of the informational constraints that underly price rigidity.

LITERATURE CITED

- Alvarez, Fernando, Herve Le Bihan, and Francesco Lippi. (2016) "The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistics Approach." *American Economic Review*, 106, 2817–51.
- Alvarez, Fernando, Martin Gonzalez-Rozada, Andy Neumeyer, and Martin Bereja. (2019) "From Hyperinflation to Stable Prices: Argentina's Evidence on Menu Cost Models." *Quarterly Journal of Economics*, 134, 451–505.
- Alvarez, Fernando, Francesco Lippi, and Luigi Paciello. (2011) "Optimal Price Setting with Observation and Menu Costs." *Quarterly Journal of Economics*, 126, 1909–60.
- Anderson, Eric T., and Duncan I. Simester. (2010) "Price Stickiness and Customer Antagonism." *Quarterly Journal of Economics*, 125, 729–65.
- Berger, David, and Joseph Vavra. (2018) "Dynamics of the U.S. Price Distribution." *European Economic Review*, 103, 60–82.
- Bils, Mark, and Peter J. Klenow. (2004) "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy*, 112, 947–85.
- Blanco, Julio A. (2016) "Optimal Inflation Target in an Economy with Menu Costs and Zero Lower Bound." Manuscript, University of Michigan.
- Caballero, Ricardo J., and Eduardo Engel. (1993) "Microeconomic Adjustment Hazards and Aggregate Dynamics." *Quarterly Journal of Economics*, 108, 359–83.
- Caballero, Ricardo J., and Eduardo Engel. (2006) "Price Stickiness in Ss Models: Basic Properties." Unpublished Manuscript, Massachusetts Institute of Technology.
- Caballero, Ricardo J., and Eduardo Engel. (2007) "Price Stickiness in Ss Models: New Interpretations of Old Results." *Journal of Monetary Economics*, 54, 100–21.

- Calvo, Guillermo A. (1983) "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics*, 12, 383–98.
- Caplin, Andrew, and Daniel Spulber. (1987) "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics*, 102, 703–25.
- Carlton, Dennis W. (1986) "The Rigidity of Prices." *American Economic Review*, 76, 637–58.
- Cecchetti, Stephen G. (1986) "The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines." *Journal of Econometrics*, 31, 255–74.
- Costain, James, and Anton Nakov. (2011) "Price Adjustments in a General Model of State-Dependent Pricing." *Journal of Money, Credit, and Banking*, 43, 385–406.
- Costain, James, and Anton Nakov. (2018) "Logit Price Dynamics." *Journal of Money, Credit, and Banking*, 51, 43–78.
- Dotsey, Michael, Robert G. King, and Alexander L. Wolman. (1999) "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output." *Quarterly Journal of Economics*, 114, 655–90.
- Dotsey, Michael, and Alexander L. Wolman. (2018) "Inflation and Real Activity with Firm-Level Productivity Shocks." Manuscript, Federal Reserve Bank of Philadelphia.
- Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo. (2011) "Reference Prices, Costs, and Nominal Rigidities." *American Economic Review*, 101, 234–62.
- Eichenbaum, Martin, Nir Jaimovich, Sergio Rebelo, and Josephine Smith. (2014) "How Frequent Are Small Price Changes." *American Economic Journal: Macroeconomics*, 6, 137–55.
- Gagnon, Etienne. (2009) "Price Setting during Low and High Inflation: Evidence from Mexico." *Quarterly Journal of Economics*, 124, 1221–63.
- Golosov, Mikhail, and Robert E. Lucas. (2007) "Menu Costs and Phillips Curves." *Journal of Political Economy*, 115, 171–99.
- Kashyap, Anil K. (1995) "Sticky Prices: New Evidence from Retail Catalogs." *Quarterly Journal of Economics*, 110, 245–74.
- 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.
- Klenow, Peter J., and Benjamin A. Malin. (2010) "Microeconomic Evidence on Price-Setting." In *Handbook of Monetary Economics*, edited by Benjamin M. Friedman and Michael Woodford, volume 3, pp. 231–84. Amsterdam: Elsevier.
- Luo, Shaowen, and Daniel Villar. (2017a). "Evidence on Monetary Non-Neutrality in U.S. Price Data." Unpublished Manuscript, Federal Reserve Board.
- Luo, Shaowen, and Daniel Villar. (2017b). "The State-Dependent Price Adjustment Hazard Function: Evidence from High Inflation Environments." Manuscript, Federal Reserve Board.
- Midrigan, Virgiliu. (2011) "Menu Costs, Multi-Product Firms and Aggregate Fluctuations." *Econometrica*, 79, 1139–80.
- Nakamura, Emi, and Jón Steinsson. (2008) "Five Facts about Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics*, 123, 1415–64.
- Nakamura, Emi, and Jón Steinsson. (2010) "Monetary Non-Neutrality in a Multisector Menu Cost Model." *Quarterly Journal of Economics*, 125, 961–1013.
- Nakamura, Emi, Jón Steinsson, Patrick Sun, and Daniel Villar. (2018) "The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation." *Quarterly Journal of Economics*, 133, 1933–80.

- Sims, Christopher A. (2003) "Implications of Rational Inattention." *Journal of Monetary Economics*, 50, 665–90.
- Vavra, Joseph. (2013) "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation." *Quarterly Journal of Economics*, 129, 215–58.
- Woodford, Michael. (2008) "Inattention as a Source of Randomized Discrete Adjustment." Manuscript, Columbia University.
- Woodford, Michael. (2009) "Information-Constrained State-Dependent Pricing." *Journal of Monetary Economics*, 56, 100–24.
- Zbaracki, Mark, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen. (2004) "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets." *Review of Economics and Statistics*, 86, 514–33.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.