

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

Shaowen Luo and Daniel Villar*

Columbia University, November 20, 2015

*JOB MARKET PAPER

Abstract

We document the predictions of a broad class of existing price setting models on how various statistics of the price change distribution change with the rate of aggregate inflation. Notably, menu cost models uniformly feature the price change distribution becoming less dispersed and less skewed as inflation rises, while in the Calvo model both relations are positive. Using a novel data set, the micro data underlying the U.S. CPI from the late 1970's onwards, we evaluate these predictions using the large variation in inflation over this period. Price change dispersion does indeed fall with inflation, but skewness does not, meaning that none of the existing models can fit these patterns. We then present a model that does, in addition to matching the price change moments that existing models do. Our model features random menu costs, and we show that with a menu cost distribution that gives a significant probability to free price changes, and a high probability to very high menu costs, the model predicts a flat inflation-skewness relation. This menu cost distribution moves the model close to a Calvo model, and the model therefore exhibits a much higher degree of monetary non-neutrality than the Golosov and Lucas (2007) model, and higher even than in the subsequent menu cost models such as Midrigan (2011).

JEL classification codes: E31, E32, E47, E52

*Department of Economics, Columbia University, 420 w118th st., New York, NY 10027, sl3256@columbia.edu (Shaowen Luo), dv2173@columbia.edu (Daniel Villar). We would like to thank Emi Nakamura, Jón Steinsson, Ricardo Reis and Michael Woodford for their invaluable advice and support. We also thank Jennifer La'O, Martín Uribe, Stéphane Dupraz, Jorge Mejía-Licona, Savitar Sundaresan, Erick Sager, Timothy Erickson, and other seminar participants at Columbia and the Bureau of Labor Statistics for valuable feedback and suggestions. The data was made accessible to us by the BLS, and we thank Ted To and John Molino for their help as our BLS coordinators. All remaining errors are our own. Updates to this draft are available at <http://www.columbia.edu/~dv2173/docs/jmp.pdf>.

1 Introduction

The dynamics of price changes (when, how, and why firms change the prices of the goods and services that they sell) have been a major focus of the study of monetary economics for the past several decades. It is indeed well known that monetary variables have no influence on real economic activity (monetary neutrality) if all prices can be freely re-set at any point in time. This has drawn attention to the study of frictions in the price-setting process for a long time: Barro (1972) and Sheshinski and Weiss (1977) characterized the pricing behavior of a firm that faces a fixed price adjustment cost, while Calvo (1983) did so for a firm facing the random opportunity to change its price. What has also become well established is that the distinction between these two approaches in modeling price change dynamics matters greatly for monetary non-neutrality. While central banks have widely adopted Calvo-style staggered price setting into the models that they use to evaluate the effects of their policies, much of the literature has highlighted how this considerably over-states the effectiveness of monetary policy, compared to what it would be if prices are set based on adjustment (or menu) costs.

The literature has emphasized that monetary non-neutrality depends not only on how often prices change, but also crucially on which prices change. Caplin and Spulber (1987) and especially Golosov and Lucas (2007) demonstrated this by showing that if prices are sticky because of menu costs, money is close to neutral. These seminal studies showed that in the presence of menu costs, only relatively large price changes will justify the payment of the cost and occur at all, which makes the aggregate price level considerably more responsive to nominal shocks than in the Calvo model. This mechanism came to be known as the selection effect, and much research has been devoted to re-evaluating the results of Golosov and Lucas (2007), and the strength of the selection effect, in light of new empirical findings established with price micro data sets (most notably, Nakamura and Steinsson (2010) and Midrigan (2011)).

Understanding the selection effect, and to what extent it plays an important role, is necessary to determine the true degree of monetary non-neutrality, but this mechanism cannot be observed directly. It would be very difficult to observe whether the prices that change are those predicted by the selection effect, so its presence and strength must be inferred indirectly from observable price change statistics. The existing work in the field has done this primarily by bringing quantitative price setting models together with the price data that has become available in the past

decade. However, an important limitation with these studies is that they have, for the most part, only used unconditional moments of the price change distribution (such as the frequency or size of price changes, averaged over time) to discipline the models used. In this paper, we show that conditional moments, which have been seldom used, are extremely informative and yield new insights on the selection effect. In particular, we find that the selection effect makes very strong predictions about how the shape of the price change distribution should change with aggregate inflation. Using a new data set, the price data underlying the U.S. CPI from 1977 onwards, we show that these predictions are not supported empirically. Finally, using a flexible model in which the strength of the selection effect can be freely set, we show that the selection effect has to be much weaker than assumed by the existing menu cost models to match the empirical facts that we present. This, in turn, indicates that monetary non-neutrality is higher than these models predicted, and similar in magnitude to what is predicted by the Calvo model.

In menu cost models, the presence of a fixed adjustment cost induces a selection effect: only price changes that are large enough to justify the cost occur, leaving an inaction region of changes (centered at zero) that are too small to be justified. A positive monetary shock (raising nominal demand) will induce prices that were otherwise already strongly mis-aligned to change, meaning that average price changes would respond relatively strongly to such a shock. This implies, in turn, that the aggregate price level will be very responsive to monetary shocks, eliminating much of the effect of the monetary shock on real activity (money is close to neutral). We exploit the fact that this logic also has strong implications for how the distribution of price changes responds to such shocks: an inflationary shock will push more price changes out of the inaction region to the positive side, and into the inaction region from the negative side. There will therefore be more price changes concentrated on the positive side of the inaction region, leaving a price change distribution that is less dispersed and more asymmetric (negatively skewed). Indeed, all existing menu cost models, because of the selection effect created by the presence of an adjustment cost, imply a very strong negative correlation between inflation and both dispersion and skewness of price changes, and these are implications that can be empirically tested.

The literature on sticky prices has faced thus far been unable to test these types of predictions because the kind of price data that is necessary has only been available for periods of low and stable inflation. Although some studies (such as Alvarez et al. (2011a); and Gagnon (2009)) have used price data from countries that experienced

high inflation, they used this data to determine how the frequency of price change behaves at high inflation, without considering the higher moments of the price change distribution. For the U.S., the main source of price data in this line of work, the micro data underlying the Consumer Price Index, was, until recently, only available going back to 1988 (while other commonly used data sets go back even less far). However, we use the data set recently presented in Nakamura et al. (2015), which extends the C.P.I micro data 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 post-war U.S., 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.

We find that while the dispersion of price changes does go down considerably in high inflation periods, the skewness does not, contrary to the strong predictions of menu cost models. Since the counter-factual predictions are driven by the mechanism behind the selection effect, we modify the menu cost model in a way that weakens this mechanism: introducing random, heterogeneous menu costs that add randomness to whether the firm will have an opportunity to change its price. We find that if the probability that firms face a very high menu cost (such that it would almost never choose to change its price) is high, the model no longer predicts the negative inflation-skewness correlation, while still matching all the facts matched by previous models. In addition, such a model features a much higher level of monetary non-neutrality than any of the existing menu cost models: around six times higher than in a standard menu cost model, and 70% as high as in a Calvo model.

The rest of the paper is organized as follows. In what remains of the introduction we provide a more detailed overview of the work done in this branch of the literature. In section 2, 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 3 describes the data set that we use and evaluates the predictions of the different models based on the data. Section 4 presents the generalized menu cost model, comparing predictions to what is observed in the data and shows the degree of monetary non-neutrality exhibited by the different models. Finally, Section 5 provides some concluding remarks.

Literature Review

While a few empirical studies of price stickiness in certain industries have been around for some time (e.g. Cecchetti (1986); Carlton (1986); Kashyap (1995)), 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 facts established by Bils and Klenow and the subsequent empirical studies on price stickiness (most notably, Klenow and Kryvtsov (2008); and Nakamura and Steinsson (2008)) have enriched the discussion on monetary non-neutrality by providing the models that evaluate monetary non-neutrality with a standard by which to be measured.

Caplin and Spulber (1987) had used a very stylized model to show that if prices are sticky, state-dependent pricing implies that monetary shocks can still have little or no effect on economic activity. Golosov and Lucas (2007) then incorporated this mechanism into a quantitative menu cost model that was calibrated to match the new empirical facts of the sticky price literature, and they confirmed that under state-dependent pricing, monetary policy is close to neutral. The model matched the fraction of prices that change (frequency of price change) estimated by the empirical papers, but also the observation that when prices do change, the changes tend to be large. Since, under menu costs, firms will only change their prices when they really need to, and so will not bother incurring a menu cost for a small price change, this latter fact in particular lent credibility to the adoption of a menu cost as the foundation of price stickiness.

Since then, the literature has continued to combine quantitative, micro-founded, price setting models with empirical facts from micro price datasets, and in this way the non-neutrality debate has advanced. While the Golosov and Lucas model matched the frequency and average size of price changes, much subsequent work has modified the model to match other aspects of the distribution of price changes, generally finding that the degree of monetary non-neutrality predicted ends up being much larger than in the original model (for example, Nakamura and Steinsson (2010); Midrigan (2011); Alvarez et al. (2014)). In a slightly different style, Vavra (2013) showed that the frequency and dispersion of price changes are counter-cyclical in the U.S., and introduced counter-cyclical dispersion shocks to match this.

With the exception of Vavra (2013), however, the papers mentioned thus far match moments that are price change statistics averaged across time. Yet all the statistics that they consider can be computed period by period, as they pertain to a distribution of price changes, which is observed period by period. Obviously, focusing on averages across time abstracts from the time series variation in these statistics, which is observed to be quite significant in the data, and this misses out on potentially informative patterns. Our paper departs from most of the existing

literature by focusing on the variation of price change statistics over time to evaluate sticky price models. These models are aimed at understanding how the dynamic pricing behavior of firms aggregates up to the response of aggregate inflation to monetary shocks. A natural way to use the time series variation of price stickiness statistics is therefore to see how they co-move with inflation, both in models and empirically. However, as mentioned earlier in this section, most existing studies have faced the limitation of working with price data sets that only cover periods of low and stable inflation. It is in this way that our data set is novel, as it makes it possible to measure price stickiness statistics at high and low inflation.

Nevertheless, evaluating sticky price models with this kind of time series variation is not unprecedented. For example, Gagnon (2009) and Alvarez et al. (2011a) use price data from high inflation episodes in Mexico and Argentina, respectively, to show that the frequency of price change rises with inflation. This fact is consistent with menu cost models, but it goes against the core assumption of the Calvo pricing model, that firms face a constant probability of changing their prices over time. Our paper confirms this result, but documents more patterns based on other statistics that paint a more nuanced picture. While the relation between the frequency of price change and inflation provides strong evidence against the strict assumptions of the Calvo model, changes in the shape of the price change distribution (measured by its dispersion and skewness) are also informative to distinguish between the models.

Ultimately, we find that neither menu cost nor Calvo models are able to match all the patterns in the data that we present. In particular, the menu cost model makes very strong predictions about the shape of the price change distribution: the dispersion and the skewness fall sharply with inflation. In the data the dispersion of price changes does fall with inflation, but the skewness does not. We are not the first to find empirical failures of this model: Nakamura and Steinsson (2010) and Midrigan (2011) had 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 non-neutrality. 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 also inconsistent with the facts that we present. Finally, we also consider models of imperfect information in which firms adjust their prices infrequently (Alvarez et al. (2011b), Woodford (2009)), and find that these also fail to match the data, although each in different ways.

Based on these findings, and in search of a model that is consistent with our empirical results, we present a generalized menu cost model in which the size of the

menu cost (the cost paid to change one’s price) is random, and changes across firms and time. This generalizes a common theme in the approach taken by Nakamura and Steinsson (2010) and Midrigan (2011): to incorporate heterogeneity of menu costs, and in so doing making the firm’s decision of whether to change its price more exogenous to the firm. These models therefore include some of the features of the Calvo model, and can be thought of as hybrids between state- and time-dependent models. Our model builds on this by introducing a distribution of menu costs that gives it the flexibility to behave like a Calvo model, a menu cost model, and to cover the spectrum in between. By working with random menu costs, we follow the example of Dotsey et al. (1999), and our generalization of the Calvo-menu cost dichotomy is closely related to Caballero and Engel (1993)’s approach. They proposed modelling sticky prices with a continuous probability of price adjustment, as a function of the gap between the firm’s current and optimal price. Our random menu cost model maps naturally into their price adjustment hazard function approach. We adjust the distribution of menu costs in our generalized model to fit the new correlations that we report, and find that, especially to match the non-negative inflation-skewness correlation, the distribution of menu costs needs to feature a positive probability of price changes being free, and a high probability of menu costs being very high. These correlations allow us to restrict the menu cost distribution in a way that neither Dotsey et al. (1999) nor Caballero and Engel could, with important implications for monetary non-neutrality. Indeed, the real effects of monetary shocks in our model are considerably higher than in the Golosov and Lucas model, and higher even than in Midrigan (2011).

2 The Skewness of Price Change in Sticky Price Models

We begin by presenting the models that we will be evaluating, and describing the predictions that we will focus on testing. Our analysis will consider the models that have been used in the sticky price literature, including the Calvo model, the Golosov and Lucas menu cost model and the variants of it that have appeared since. First, we describe the set-up of the various models, both the common framework and the differences that set them apart, before explaining how we derive the predictions, and we finally summarize the predictions.

2.1 General Set-Up

All the sticky price models that we consider have certain features in common, that are also used in the sticky price literature in general. First, households maximize expected discounted utility of the following form:

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

All our analysis will focus on the firm's dynamic price setting, so the set up of the household problem matters for our purposes insofar as it determines the relationship between aggregate consumption and the real wage, which will be the firm's main cost. There is then a continuum of monopolistically competitive firms, indexed by z , producing a differentiated product, and aggregate consumption is given by a constant elasticity of substitution aggregator, meaning that each firm faces the standard demand function for its good:

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

, 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:

$$y_t(z) = A_t(z) L_t(z)$$

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 large 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). These shocks are typically modelled as first-order autoregressive processes with normal innovations, but Midrigan (2011) argues that such a process yields a distribution of price changes with tails that are too thin, relative to what is observed in the data. He therefore introduces Poisson shocks in the productivity process in the following way:

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

This set-up nests the standard AR(1) productivity, which can be obtained by

simply setting the probability of a shock occurring (p_ϵ) to 1. Since we will consider various models with AR(1) productivity, as well as Midrigan'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 modelling nominal output as a log random walk with drift:

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

This process stands in for monetary policy in these models: nominal output is determined exogenously, and firms' price responses to these shocks determine how inflation, and how real output respond. We will use the same parameter values for this process (to match the behavior of US aggregate activity) across the different models, and we define monetary non-neutrality 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 (2015) 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. Next, we describe the price setting problem faced by firms, which is the main dimension along which the different models vary.

2.2 Price-Setting

In the standard Golosov and Lucas (2007) menu cost model, firms must pay a fixed cost (in units of labor) whenever they change their price. The period profit function therefore takes the following form:

$$\Pi_t(z) = p_t(z)y_t(z) - W_t L_t(z) - \chi W_t I\{p_t(z) \neq p_{t-1}(z)\}$$

The menu cost (χ) can then be calibrated to match the frequency of price changes, while the standard deviation of the idiosyncratic shocks can be set to match the average size of price changes (we also set the probability of an idiosyncratic shock

occurring, p_ϵ , to 1 to make the process an AR(1), as in the original model). This is, in a way, the most “state-dependent” model, as under the fixed menu cost firms are fully in control of the decision of when to change the price for each good (subject to the constant menu cost). It is this feature that makes prices very responsive to aggregate demand shocks, and that famously yields very low monetary non-neutrality.

The first extension to the menu cost model that we consider is the Nakamura and Steinsson (2010) multi-sector menu cost model. The only change here is that firms are separated into sectors, with firms in different sectors facing different menu costs, and a different variance of idiosyncratic shocks. This reflects the fact, documented in the paper and in Nakamura and Steinsson (2008), that the frequency of price change varies considerably across sectors, as does the average size of price changes. Golosov and Lucas (2007) calibrated their model to match the average frequency of price changes across sectors, and Nakamura and Steinsson show that calibrating sector by sector makes a major difference for the degree of monetary non-neutrality in the models, as the multi-sector model predicts much higher non-neutrality than the standard model.

Midrigan (2011) modified the standard menu cost model in two ways: 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. In the standard model, since a firm always has to pay a fixed cost to change its price, there will be a threshold for the size of the price change, such that changes below a certain size are not profitable and do not occur. Midrigan (2011) models multi-product firms that can change the prices of all their products for the payment of the menu cost. Because of this, a firm might choose to pay the menu cost to change the price of a particularly mis-aligned product price, and then also take the opportunity to change the price of another product by a small amount. 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 therefore 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. This adds, as an additional parameter to calibrate, the probability of drawing a zero menu cost (free price change): p_z . With the additional parameters in this model, we target the fraction of price changes that are small, as in Midrigan (2011).¹

¹Midrigan (2011) 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

We also consider a Calvo model, which has the set-up described above, except that firms, instead of facing a menu cost, have a fixed probability every period of receiving the opportunity to freely change their price (otherwise, they do not get to change price). 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. The variance needs to be higher than in menu cost models, because menu costs induce the selection effect that naturally leads to large price changes to be more likely.

Finally, we also include two models involving imperfect information: the Alvarez et al. (2011b) 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. 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 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 mis-alignment. This is the kind of model that we work with in section 4, and we view the rational inattention framework as a potential micro-foundation for this.

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. Finally, Midrigan (2011) also emphasized the failure of the Golosov and Lucas model to match the kurtosis of the price change distribution, and the introduction of Poisson idiosyncratic shocks helps to get the kurtosis in the model closer to what it is in the data. However, it turns out to be very difficult to match (it seems to be very high in the data), and Midrigan (2011) does not achieve it completely. We therefore do not match the kurtosis either.

2.3 Solution and Simulation

We solve each of the models mentioned above by value function iteration, mostly with the parameter values used by the original authors, which were set for the models to match various features of the micro price data. One difficulty in solving these models is that in all of them the price level (P_t) is an aggregate endogenous variable whose evolution depends on the behavior of all firms. This means that, in principle, every firm's relevant state should include the state of every other firm, which makes for an infinitely large state space. As done elsewhere in the literature, we use an approach analogous to Krusell and Smith (1998) to solve the model assuming a relationship between the price level and a small number of variables, and to then verify that the resulting solution is consistent with the assumed relationship. In the appendix, we provide more details about the procedure, as well as the calibration of the different models. The parameters of the process for nominal aggregate demand, described above, are calibrated to match the average growth and volatility of U.S. nominal GDP, and the same values are used for all the models.

The first aim of our paper is to document what these different models imply for the price change distribution at different inflation rates. Our approach is to simulate each model, for 1,000 periods (months) and 40,000 firms. From these, we obtain a simulated series for aggregate inflation (determined by the endogenous response of prices to the nominal aggregate demand shocks) and a distribution of price changes for each period. Since the models are calibrated to match the frequency of price change that is observed empirically, the vast majority of prices do not change every period. Our analysis is therefore based on the distribution of price changes, conditional on a non-zero price change, and this applies for the rest of the paper, including in our empirical work. We compute various moments of each period's price change distribution, giving us a time series for each moment, and compute correlations between inflation and each of the moments, and this is how we determine how the price change distribution changes with inflation.

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. We present the correlation between frequency and inflation in the models, but also consider other correlations with other moments: the standard deviation of price changes, and the skewness of price changes. As we will show, the menu cost models have very strong and clear implications for these correlations that are markedly different from those of the Calvo model. Furthermore, as seen in Midrigan (2011), the shape of

the price distribution can be very informative about the importance or presence of the mechanisms that weaken the role of monetary shocks, and it is therefore to be expected that the way in which the shape of this distribution changes (as described by the dispersion and skewness) with inflation would also be informative about these mechanisms.

We present a summary of these theoretical results in Table 1, indicating whether each correlation is positive (+), close to zero (0), or negative (-) in the different models:

Model	Table 1: Correlation of Inflation and		
	Frequency	Std. Deviation	Skewness
Calvo	0	+	+
Golosov and Lucas	+	-	-
Nakamura and Steinsson	+	-	-
Midrigan	+	-	-
Alvarez et al.	+	0	-
Woodford	+	0	+

In order to further illustrate these results, we present scatter plots between inflation and the different moments from the simulations (in which one point represents one period in the model simulations). Figure 1 shows the correlations for the frequency of price change, while Figures 2 and 3 do so for the dispersion and skewness of price changes, respectively. These bring out the fact that in the menu cost models, the relationships between inflation and dispersion and skewness are very clear and strong (especially in the Golosov and Lucas model for the dispersion). In contrast, the same relations in the Calvo and imperfect information models are not so strong.

Figure 1: Price Change Frequency & Inflation

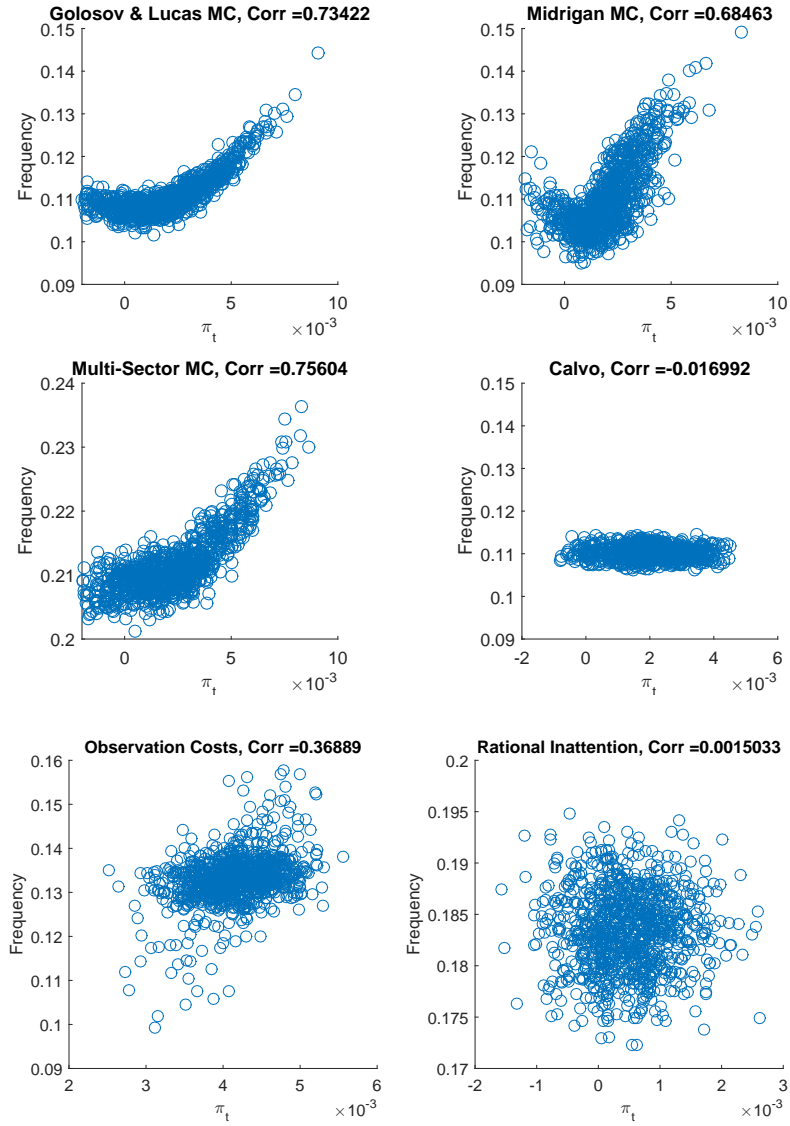


Figure 2: Price Change Dispersion & Inflation

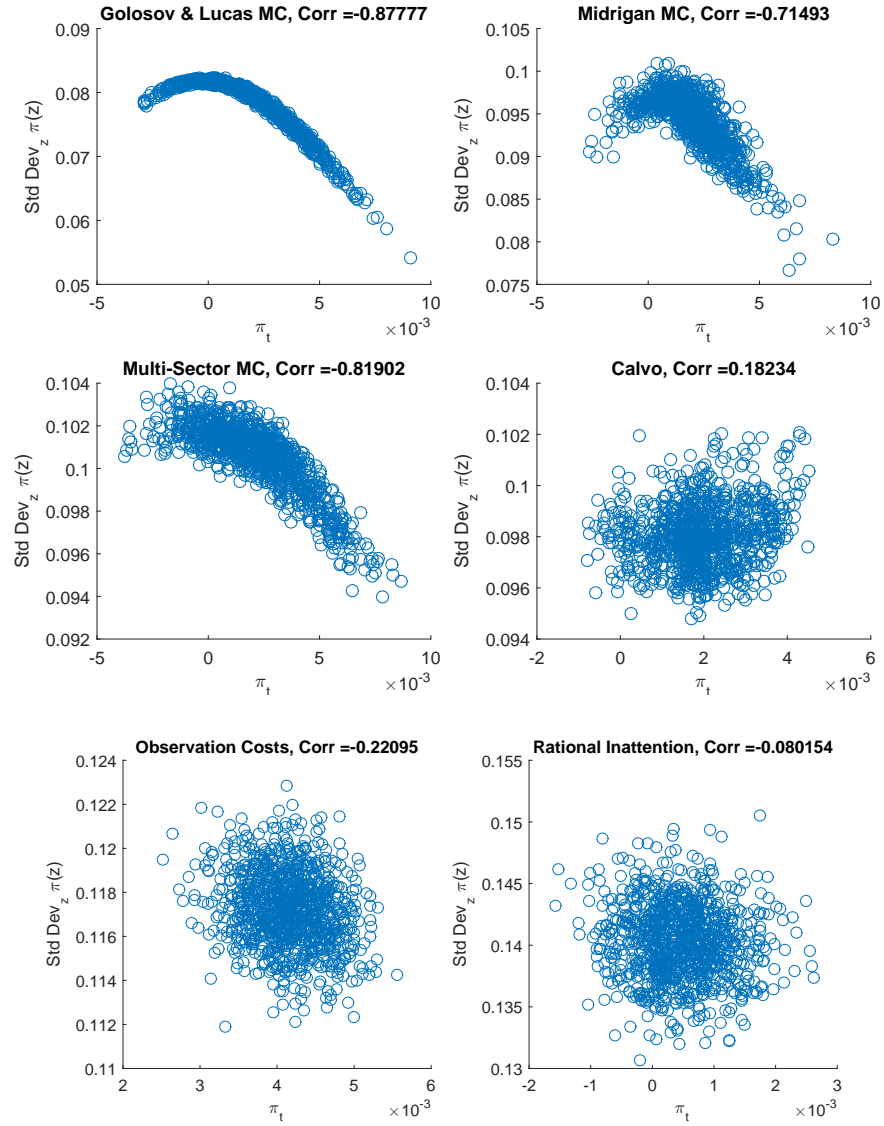
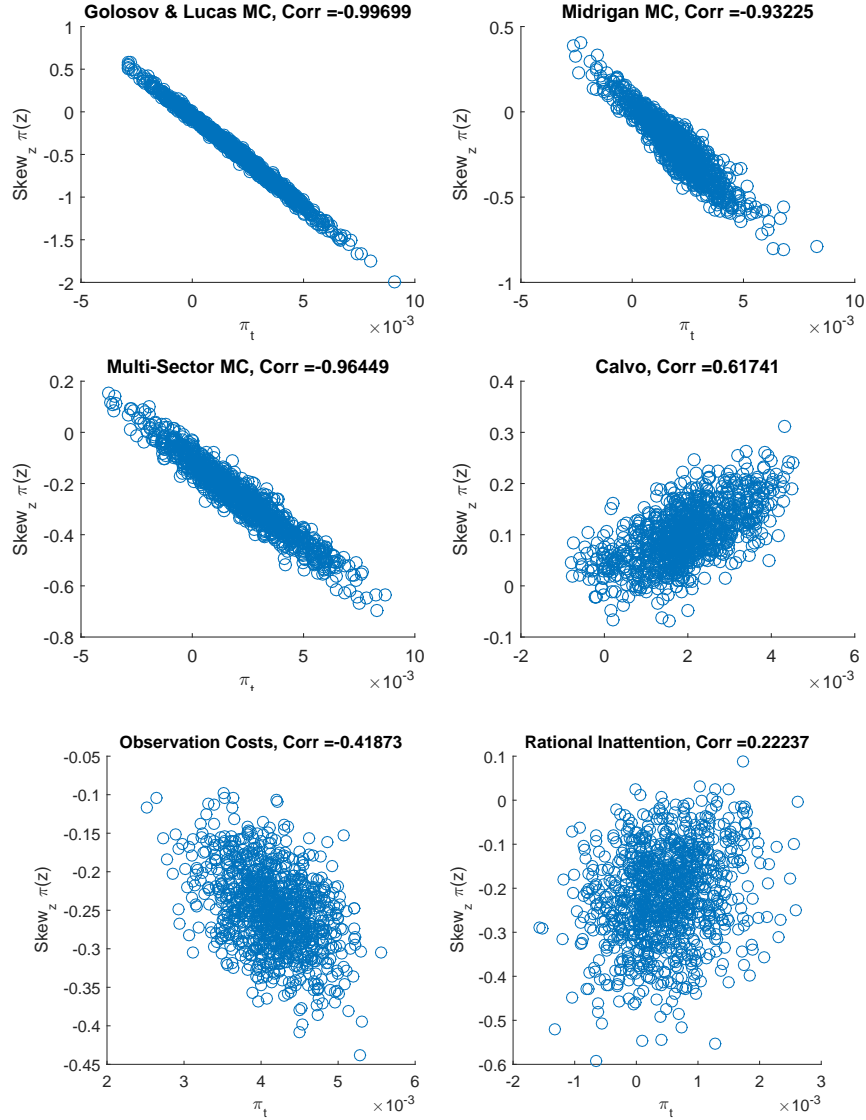


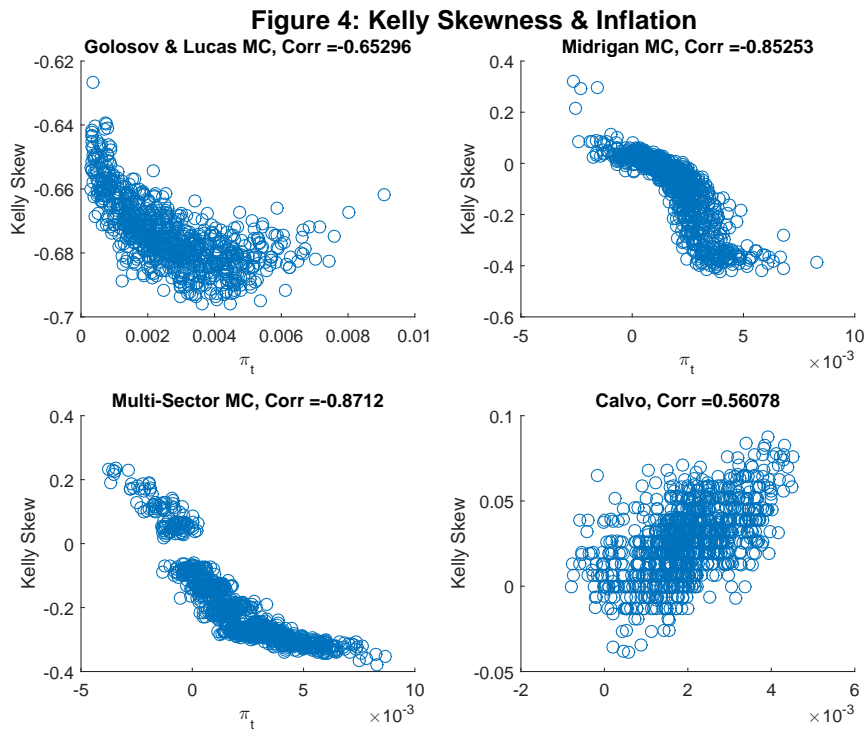
Figure 3: Price Change Skewness & Inflation



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 inter-quartile range (for dispersion) and Kelly’s coefficient of skewness²(as opposed to “moment skewness”, which is what

²These statistics are defined as follows, with Q_i representing the i^{th} percentile. Inter-quartile

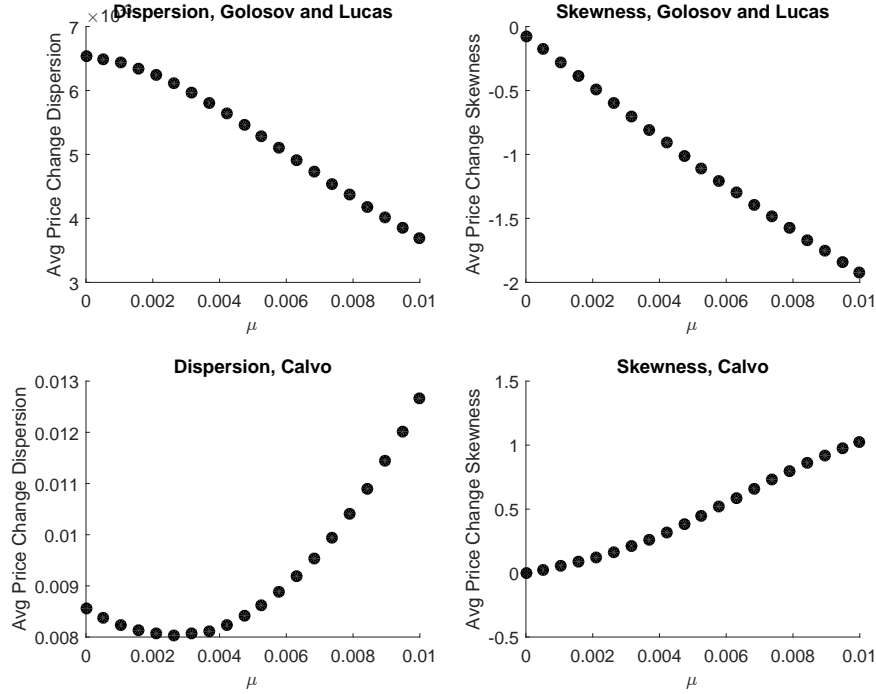
we have been estimating so far). Since these statistics are quantile-based, they are well-defined for any distribution, and they are also less sensitive to outliers. The correlations are similar for all the models (inter-quartile range compared with standard deviation, and moment skewness with Kelly Skewness). Figure 4 below shows scatter plots of Kelly Skewness in the different models.



Another concern could be that these simulations all assume that the value of steady-state inflation is held constant throughout the simulated time period. This could be problematic in terms of testing the predictions on data, as the U.S. clearly went from a moderate to a low inflation regime over our sample period. To address this, we also conduct the following exercise: we solve each model for different values of the trend inflation parameter (μ), and for each solution compute the average dispersion and skewness of price change (either from the stationary distribution of price changes, or averaging over simulated time periods; they are almost the same). In Figure 6, we plot the results.

range = $Q_{75} - Q_{25}$. Kelly Skewness = $\frac{(Q_{90} - Q_{50}) - (Q_{50} - Q_{10})}{Q_{90} - Q_{10}}$. Kelly skewness essentially measures the degree of asymmetry in a distribution, comparing the size of the right and left tails.

Figure 5: Steady-State Correlations




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 (we are only plotting results for the Golosov and Lucas model, but the same pattern holds for the other menu cost models). Here too, the Calvo model predicts weak positive relations for both moments. This 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 (Golosov and Lucas, Nakamura and Steinsson, Midrigan, and observation costs). The scatter plots show that the values taken by moments we report do vary across the models (for example, in the Golosov and Lucas model the skewness of price changes takes a wider range of values than in the other models), but the fact that the sign and strength of the correlations across the models are similar is notable. Indeed, the Nakamura and Steinsson and Midrigan menu cost models were developed as extensions of the Golosov and 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 to the baseline menu cost model, they all have the same implications along the dimensions that we are considering. Next, we discuss the intuition behind these theoretical results.

2.4 Intuition for the Menu Cost Model

Menu cost models are often also known as “Ss” models, due to the fact that they tend to feature an inaction region for price changes (the edges of which can be labelled with “S” and “s”), and this makes it easier to understand the theoretical correlations between inflation and the moments of the price change distribution that we find in this section. Price change dynamics in the menu cost model can be thought of in the following way: both idiosyncratic and aggregate nominal shocks give a distribution of desired price changes (the price change a firm would choose if it changed its price, or in the absence of price change frictions). The presence of a menu cost means that only desired prices above a certain size (positive and negative) will actually occur, as only those will yield a benefit to the firm big enough to compensate for the menu cost. The realized price change distribution in this model is therefore the underlying distribution with a band containing 0 removed, as illustrated in Figure 6 below.



The presence of idiosyncratic shocks yields variation in firms’ desired prices, and nominal aggregate shocks move the position (average) of the underlying distribution. For example, a positive aggregate shocks moves the distribution to the right, which also leads to realized prices being higher on average, resulting in higher inflation (the reverse is true for negative aggregate shocks). Such shocks also result in a higher fraction of price changes being positive, which are separated from the negative ones by the inaction region. This reduces the dispersion of price changes because a bigger fraction of them are on one side of the inaction region, and therefore 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, and by the same logic, higher than when inflation is negative (when more price changes are decreases), which is what we see in the dispersion plots for the menu cost model: dispersion decreasing with inflation in the positive region, and increasing in the negative region, with the maximum attained at zero inflation.

The logic for why the skewness falls with inflation is related. The skewness, as a statistic, measures the asymmetry of a distribution, or the relative sizes of the

right and left tails. As a positive aggregate shock raises the average desired price change, and the average realized price change, some negative price changes (to the left of the inaction region) remain and form the left tail. This makes the skewness negative: the resulting distribution has a left tail (price decreases relatively distant from the average price change, which is positive), without a corresponding right tail (as price increases are to the right of the inaction region and relatively close to each other). Furthermore, for the range of values that inflation takes in our simulations (which corresponds roughly to the historical range for inflation since the late 1970's), there is always a significant proportion of negative price changes. This means that as inflation rises (due to larger positive aggregate shocks), these negative price changes form a left tail in the price change distribution that is further and further (to the left) of the average of the price change distribution, leading to a skewness that is more negative.³ What this also implies is that the relationship between skewness and inflation is monotonic, decreasing for positive and negative values of inflation.

It is important to emphasize that these correlations have to do with the central mechanism of the menu cost model: the selection effect. When firms face a fixed cost to changing their price, only relatively large price changes will occur, leading to the presence of the inaction region. As the average of the underlying distribution rises (moved by aggregate shocks), there is a large response of inflation because there is a large share of price increases that are marginal: without the shock they would not occur, but are pushed outside the inaction region (and many marginal price decreases do not occur with the shock), leads to a relatively large rise in inflation, muting the real effect of the aggregate shock. This is the logic for why state-dependent models are known to imply low levels of monetary non-neutrality.

However, what we show is that this same mechanism leads to predictions that are in principle observable: the presence of an inaction region means that positive aggregate shocks should lead to not only more price increases, but to a distribution with price changes more concentrated on the right, leading to a declining dispersion and skewness. This does not occur in a Calvo model: in such a model every desired price change has a fixed probability of being realized, so as the desired price changes rise, the shape of the realized price change distribution does not change in a meaningful way.

³This also means that if the aggregate shock were so high that all price changes were positive (to the right of the inaction region), the relationship would break down, as price decreases would no longer be separated from price increases. However, this would also mean that all prices would change, and that inflation would be extremely high. This kind of situation, or anything resembling it, never occurs in the period we are considering.

The intuition for this theoretical result is easiest to explain in the case of the “standard” Golosov and Lucas model, or in general any menu cost model with a single fixed menu cost. The other menu cost models that we consider feature a richer structure of menu costs that led to some very different empirical predictions. However, we have shown that these models also imply negative correlations for the dispersion and skewness of price changes, and the intuition for this is the same as for the standard model. In the multi-sector menu cost model, different sectors face different menu costs, and this 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.

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 no 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 (as those will be the only ones for which a firm will be willing to pay the positive menu cost, when it faces the positive menu cost). It is this difference between small and large price changes that makes the same mechanism present in this model and drives the correlations, even though small price changes do occur (as they do in the data, but do not in the Golosov and Lucas model).

We have shown that menu cost models, under the assumptions commonly made in the literature, make clear, consistent predictions about how the shape of the price change distribution changes with inflation, and that these do not change much based on the type of menu cost model in question, and that the predictions are strikingly different from those of the Calvo model. Furthermore, these are predictions that can be tested with the price data available to us, which enables us to evaluate this broad class of sticky price models. In the following section, we do this by presenting the empirical counterpart to the correlations presented in Table 1, and we discuss how each of the models falls short of matching the data.

3 Empirical Evidence from High Inflation Periods

In the previous section, we documented the predictions made by various sticky price models on the behavior of price changes at different inflation rates. In this section,

we describe the data set that we will use to test these predictions, and report that while the inflation-dispersion correlation is consistent with the empirical evidence, the inflation-skewness correlation is not.

3.1 Previous Empirical Work

The micro data that underlies the U.S. Consumer Price Index (CPI), gathered by the Bureau of Labor Statistics, is one of the most widely used data sets in the literature on monetary price-setting models. Bils and Klenow (2004) were the first to use this data set to provide estimates for the frequency of price change. Since then, other studies have documented additional features of the price change distribution using this data set (e.g. Nakamura and Steinsson (2008); Klenow and Kryvstov (2008)). The availability of a large, representative data set that makes it possible to observe the price changes of very specific products has lead monetary economists to develop models that match the behavior of price changes as closely as possible.

The data set that has been used in this line of work covers the period 1988 to the present, as 1988 marked a major revision of the structure of the CPI. However, a limitation of the data used thus far is that throughout this period, aggregate inflation has been relatively low and stable, especially compared to the years before. Since 1988, the maximum twelve month change in the headline CPI has been 6.2% (4.6% for the Core PCE), and the average has been 2.8% (2.2% for the Core PCE). Partly because of this, most research on sticky price models up until now has focused on matching moments of the price change distribution that are averaged over time (the main exception being Vavra (2013), who uses the CPI micro data to investigate the cyclicity of price change moments). But as we showed in the previous section, the models imply that these moments would change over time, and in a way that is closely related to aggregate inflation, with implications that differ strongly across models.

Motivated by this, a few studies have used data from other countries that experienced episodes of high inflation, such as Argentina (Alvarez et al. (2011a)) and Mexico (Gagnon (2009)). These studies also used the micro data that underlies the CPI's of these countries, and reported how various price change statistics change as inflation goes from low, to moderate, to high. They find that the frequency of price change is fairly constant, and not very responsive to inflation, at low levels of inflation (below 10% annual). Once inflation rises even higher, however, the frequency of price change begins to rise sharply with inflation. In addition, they show that

a standard menu cost model matches this relationship very well. What this shows is that, first, the Calvo (1983) assumption of a constant frequency, while possibly approximately valid for low inflation, becomes problematic when inflation rises beyond a certain level. Second, the evidence presented in these papers is shown to be consistent with standard menu cost models, suggesting that they better explain the behavior of price changes when inflation is high. However, Alvarez et al. (2011a) and Gagnon (2009) do not look at the higher moments of the price change distribution that we emphasized in the previous section, which is what we do in this paper.

3.2 Data Set and Construction of Statistics

The CPI Research Database collected and maintained by the U.S. BLS contains about 80,000 monthly prices collected from around the U.S, classified into about 300 categories called Entry Level Items (ELI's). As mentioned before, the data going back to 1988 has been available for a little over a decade. The data going back to 1977 has 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. (2015), and that paper also describes in detail just how the data set was made available. As explained in the BLS handbook of Methods (BLS, 2013), there were several changes made to how the BLS samples prices and computes the CPI. While there are many variables present in the post-1988 data set that are not available for the older period, we are able to study the price change distribution in a way that is consistent throughout the whole period, and with the theoretical framework that we are testing. First, we have access to the variables that identify specific products, and that reveal when a substitution has occurred (when a new version of a product has replaced the old one). Second, the data set contains information on when any given price is a temporary sale, or imputed (or 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.

The empirical literature on price setting has emphasized the importance of identifying “pure” or regular price changes, as opposed to price changes coming from temporary sales or substitutions. The reason is that sales and substitutions have features that make them different in terms of their relevance for the study of the role of monetary policy and aggregate shocks. Indeed, when a product goes on sale, its price will change, but it is not clear that this happens in response to any changes

in aggregate conditions. What’s more, products on sale tend to revert back quickly to their pre-sale price. This distinction was pointed out notably by Nakamura and Steinsson (2008), and Anderson et al. (2015) document the ways in which sale prices behave differently from regular prices. In a similar way, the distinction between regular price changes and substitutions is made because a price change coming from a product substitution could reflect the changes in product characteristics or in quality that could be behind the substitution. Although it is possible in some cases to estimate the contribution of quality or characteristic changes to a substitution price change (and the BLS does for certain products), we prefer to use the product identifiers to focus on price changes involving identical products.

In order to test the predictions that we presented in the previous section, we use the data set to construct distributions of price changes for each month, and a few observations on how these are constructed are in order. First, since the vast majority of prices do not change in any given month, these distributions only include non-zero price changes (which corresponds to what we look at in the theoretical results). Second, because estimates of higher moments are very sensitive to outliers, we follow other empirical work in excluding price changes whose absolute value is above a certain value (e.g. Klenow and Kryvstov (2008); Alvarez et al. (2014)), (our threshold is one log point). Third, Eichenbaum et al. (2013) have shed light on problems with the methods of reporting and collecting prices in some of the product categories of data sets such as the CPI. They show that this leads to erroneous small price changes appearing in the data, price changes that come from the price collection methods, and that do not reflect actual price changes. This is particularly important for us, as estimates of dispersion and skewness will be sensitive to the relative amounts of small and large price changes. We deal with this by constructing statistics that exclude very small price changes ($<1\%$ in absolute value) in the ELI’s that Eichenbaum et al. flagged as problematic as a robustness check.

Finally, it has been pointed out by Nakamura and Steinsson (2008 and 2010) that there is significant heterogeneity of price change statistics across sectors. To report the average overall frequency of price change, they estimate the frequency for each ELI, and then take a weighted average of each frequency (with the expenditure weights that go into the CPI). The same method is used by many of the other cited empirical studies. For the frequency of price change, we use the same method, considering both the weighted median and mean frequency⁴. For the dispersion and

⁴Nakamura 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

skewness, we follow a similar approach: we first estimate each moment by sector-month. However, as ELI’s are fairly narrow categories, most of them have a handful of price change observations in any given month, fewer than would be necessary to estimate higher moments with any precision. We therefore do not use ELI’s as our definition of sectors, but instead separate products into 13 “major groups”, which are listed in the appendix. While this sectoral classification is fairly broad, it 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., and then for each month take weighted averages of the statistics.

This approach then leaves us with monthly series of the different moments of the price change distribution. We believe that our approach, following the empirical price setting literature, gives us the most valid estimates to compare with those from model simulations. Indeed, the models that we are testing involve “pure” price changes, and abstract from temporary sales and product substitutions, which is why we try as much as possible to include only regular price changes in our empirical estimates. Perhaps more importantly, the models do not allow for differences across sectors. Such differences, such as sector-specific shocks, have the potential to strongly affect the shape of the overall price change distribution (when all price changes across sectors are pooled together), in turn affecting the higher moments of the distribution. Because of this, we might see the moments of the “pooled” distribution of price changes vary over time due to such sector-specific shocks, which would be unrelated to the mechanisms that are behind the predictions of the models that we described in the previous section. For this reason, we attempt to “control for” these kinds of effects by computing statistics sector by sector.

3.3 Results

The goal of our empirical work is to determine whether the theoretical patterns documented above are borne out by the data. As in the theoretical section, we focus on the correlations between aggregate inflation and price change dispersion, and between inflation and price change skewness. The price change moments are calculated as described above, and our preferred measure for aggregate inflation is

a few ELI’s having very high frequencies. They argue that the median is a better measure of the average frequency in the sense that a single-sector menu cost model calibrated to match the median frequency is a much better approximation of a multi-sector model, of the kind described in Section 2. In this way, the median frequency is a statistic that better describes the degree of price stickiness (as it relates to monetary non-neutrality). This is also why we calibrate all the single sector models to match the median frequency.

monthly core PCE inflation. We prefer to use core inflation because the sharp changes in headline inflation tend to be driven by changes in the global market prices of food and commodities, which would not be well described by the price-setting models that we are working with. However, we will also compute correlations with headline inflation as a robustness check (as well as using estimates of the moments excluding price changes from food and energy categories). Finally, to control for seasonality in the inflation and moment series, we calculate the correlations after removing month dummies from the series, and after applying a moving average smoother to them.

The price data is monthly, and inflation series are monthly, so we can compute the correlations at a monthly frequency. However, the drawback of using monthly series is that each period's moment estimates are based on relatively few observations, making them less precise (this is especially important for higher moments such as the dispersion or skewness). The alternative is to group price change observations by quarters or years (but still separating them by sector) and to estimate the moments from these samples, which gives us more precise estimates (as they are based on distributions with more observations), but only quarterly or annual moment series. Since quarterly and annual inflation averages also have the advantage of containing less noise than monthly inflation series, we consider monthly, quarterly, and annual correlations. We present the results in two ways: first, with raw correlations and scatter plots, as with the models, to give a simple illustration of the signs and strength of the relationships in the data, and a qualitative comparison with the models. which we correlate with inflation series of the corresponding frequency. Secondly, we estimate these relationships with regressions. This allows us to more formally test for significance, and to control for other variables that might conceivably affect the price change distribution.

3.3.1 Correlations

Our sample period for the price data is 1977-2014, and the early, high inflation, part of the period is particularly important. We want to answer whether the dispersion and skewness of price changes move inversely with aggregate inflation, as predicted by most sticky price models, and in order to do this it is very helpful to see how the statistics change when inflation was high. However, we first verify that the frequency of price change rises with inflation, as found by Gagnon (2009) and Alvarez et al. (2011a). Table 2 reports correlations between the frequency of price change and inflation, and Figure 6 is the empirical counterpart to Figure 1 from the simulations (scatter plots of the average frequency and inflation for the months in the sample,

for both the weighted mean and median frequency).






Table 2: Corr(Frequency, Inflation)

	Weighted Median					
	Monthly		Quarterly		Annual	
	1977-2014	1985-2014	1977-2014	1985-2014	1977-2014	1985-2014
Raw	0.575	0.399	0.671	0.536	0.764	0.618
Smoothed	0.769	0.552	0.785	0.628	-	
Weighted Mean						
Raw	0.311	-0.019	0.314	-0.216	0.374	-0.243
Smoothed	0.371	-0.337	0.36	-0.295	-	

The table and figure confirms that there is a positive association between the frequency and inflation, although this is considerably clearer for the median than the mean frequency. As argued in the previous studies that had looked into this relation, this provides strong evidence against the Calvo assumption of time-dependent price setting. Figures 7 and 8 illustrate the other correlations that are presented in Table 2: those involving quarterly and annual averages of inflation and the frequency, and here the same pattern holds.



Next, we look at the results for the moments that our discussion has focused on: the dispersion and skewness of price changes. Tables 3 and 4 report the correlations for the dispersion and skewness respectively. 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; “Skewness” in the tables and graphs refers to moment skewness), the correlation is either strongly positive (over the whole sample period) or close to zero (post-1984). This can also be seen in Figures 9 and 10⁵, which are scatter plots illustrating the correlations (with each period corresponding to a month).

⁵In the following figures, the left panel uses the statistics estimated using all available observations, while the right panel uses the estimates that exclude price changes below 1% in absolute value in ELI’s deemed problematic by Eichenbaum et al. (2013) (EJRS).





Table 3: Corr(IQR, Inflation)

	All Observations					
	Monthly		Quarterly		Annual	
	1977-2014	1985-2014	1977-2014	1985-2014	1977-2014	1985-2014
Raw	-0.602	-0.446	-0.716	-0.665	-0.776	-0.751
Smoothed	-0.675	-0.706	-0.719	-0.742	-	
EJRS						
Raw	-0.666	-0.434	-0.711	-0.689	-0.775	-0.779
Smoothed	-0.792	-0.701	-0.709	-0.769	-	

Table 4: Corr(Skewness, Inflation)

	All Observations					
	Monthly		Quarterly		Annual	
	1977-2014	1985-2014	1977-2014	1985-2014	1977-2014	1985-2014
Raw	0.265	0.084	0.345	0.067	0.473	0.122
Smoothed	0.506	0.136	0.474	0.133	-	-

EJRS						
Raw	0.272	0.068	0.327	0.053	0.447	0.102
Smoothed	0.462	0.144	0.452	0.105	-	-

Table 5: Corr(Kelly Skewness, Inflation)

	All Observations					
	Monthly		Quarterly		Annual	
	1977-2014	1985-2014	1977-2014	1985-2014	1977-2014	1985-2014
Raw	0.584	0.069	0.674	-0.106	0.744	-0.165
Smoothed	0.696	-0.067	0.697	-0.199	-	-

Figures 11-14 show these correlations with the quarterly and annual measures (including the Kelly Skewness correlation using annual data), illustrating how the same patterns hold.




Figure 13: Skewness & Inflation, Quarterly




Figure 14: IQR & Inflation, Annual






Figure 15: Skewness & Inflation, Annual





To summarize, we find first that the dispersion of price changes falls sharply with inflation throughout the sample period. Second, the skewness, while varying over time, does change with inflation in a systematic way for low levels of inflation. However, there does seem to be a positive relationship when inflation is high. We see this from the different correlations for the different sample periods (which roughly correspond to the high and low inflation periods). Finally, all these patterns hold true regardless of whether we exclude potentially spurious small price changes or apply seasonal adjustment and smoothing to the data series. Next, we formalize this analysis with linear regressions.

3.3.2 Regressions

Although the correlations and scatter plots provide a general picture of what the data shows on the relationships in question, we turn to regressions to determine whether these correlations are statistically significant, and to consider different control variables. However, the question of interest for us is not merely whether they are statistically significant from zero, but also whether they are significantly different from what the models predict. To do this, we estimate regressions of the frequency, dispersion (inter-quartile 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:

$$y_t = \alpha + \beta\pi_t + \gamma Controls_t + e_t$$

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. Since expected inflation could affect firms' price setting decisions separately from present realized nominal shocks, we include expected inflation (measured by the University of Michigan Survey of Consumers) as a control. 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 (no controls, Fed dummies only, or Fed dummies with expected inflation) and the different periods. Tables 6 to 9 show the estimates for β from these different specifications, with the standard errors below them. All standard errors are calculated according to Newey and West (1987), and allow for serial correlation in the residuals.

Table 6
Coefficients for Frequency Regressions

Specification	Weighted Median		Weighted Mean	
	1977-2014	1985-2014	1977-2014	1985-2014
All	0.708*** (0.071)	0.777*** (0.224)	0.164 (0.203)	0.018 (0.196)
Fed Dummies	0.728*** (0.095)	0.810*** (0.208)	0.686*** (0.104)	0.339** (0.167)
Inflation Only	0.771*** (0.237)	0.587** (0.252)	0.438*** (0.108)	-0.087 (0.236)

Notes. *** Significant at 1% level (** at 5% level; * at 10% level.)

This table reports the correlation coefficients from regressions of the weighted average (median and mean) frequency of price changes on aggregate CPI inflation.

The regressions are run using quarterly series, where quarterly inflation is defined the mean of the 12-month log changes in the CPI for the three months in every quarter. The different cells indicate different specifications, which change with respect to the sample period used and inclusion exclusion of small price changes (columns), and what controls are used. Standard errors that are consistent for heteroskedasticity and auto-correlation of the residuals (Newey-West) are reported. The same observations apply to the other regression tables, which report coefficients on inflation in regressions with other dependent variables.

Table 7
Coefficients for IQR Regressions

Specification	All Observations		EJRS	
	1977-2014	1985-2014	1977-2014	1985-2014
Inflation Only	-0.296*** (0.042)	-0.428*** (0.070)	-0.327*** (0.046)	-0.491*** (0.082)
Fed Dummies	-0.186*** (0.038)	-0.414*** (0.077)	-0.204*** (0.044)	-0.476*** (0.089)
Fed & Expected Infl	-0.257*** (0.089)	-0.222** (0.086)	-0.261*** (0.095)	-0.224*** (0.092)

Table 8
Coefficients for Skewness Regressions

Specification	All Observations		EJRS	
	1977-2014	1985-2014	1977-2014	1985-2014
Inflation Only	3.936*** (0.827)	1.732 (1.641)	3.501*** (0.828)	1.108 (1.534)
Fed Dummies	4.309*** (1.012)	1.541 (1.857)	3.928*** (0.966)	1.130 (1.705)
Fed & Expected Infl	2.665 (2.788)	3.634 (3.279)	1.947 (2.538)	2.963 (2.985)

Table 9
Coefficients for Kelly Skewness Regressions

Specification	All Observations	
	1977-2014	1985-2014
Inflation Only	2.499*** (0.354)	0.320 (0.454)
Fed Dummies	2.439*** (0.363)	0.710* (0.423)
Fed & Expected Infl	1.658 (0.948)	0.942 (0.595)

These results confirm what the correlations showed: the frequency of price change rises with inflation (although for the mean frequency this is not so clear), the relationship between dispersion and inflation is negative and statistically significant in all specifications and sample periods. The skewness correlation, however, is significantly positive for the whole sample, but not significantly different from zero when the early, high-inflation period is excluded (and this applies for both measures of skewness). This indicates (as we can also see from the scatter plots), that this relation is close to flat for low inflation periods, but clearly positive for high inflation periods. The fact that the skewness of price change is higher on average in high inflation periods is important, because it also goes against the menu cost models' predictions at high values of steady-state inflation, as we showed in Figure 6. It is also notable that the skewness coefficients change considerably when expected inflation is included as a regressor. Since expected inflation is very highly correlated with realized inflation,

the estimates are much less precise (as shown by the high standard errors), so this is not surprising. However, this makes little difference to the comparisons with the coefficients predicted by the models, which we turn to with Table 10.

Table 10
Coefficients on Inflation for Price Change Moment

Model	Frequency	IQR	Skewness	Kelly Skewness
Golosov & Lucas	0.139	-0.937	-17.7	-0.40
Multisector Menu Cost	0.143	-0.218	-5.39	-4.33
Midrigan	0.348	-0.896	-9.84	-6.53
Observation and Menu Costs	.268	-0.071	-4.32	
Calvo	-0.003	0.040	2.93	1.00
Rational Inattention	0.001	0.007	3.03	

The table presents the coefficients on inflation from regressions of the same type, but run on simulated data from the different models. The first four models (menu cost models) have negative coefficients for the inter-quartile range, although for all but the multi-sector model, they are outside the 95% confidence intervals of the coefficients that we estimate. 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 for the skewness coefficients under all specifications, and the same is true for Kelly skewness⁶. We summarize our findings in Table 11 below, which “updates” Table 1 by adding the signs of the correlations in the data to those predicted by the models.

⁶The one exception is the coefficient for the Golosov and Lucas 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 (bi-modal) in this model, leading to this weak relationship. The model’s Kelly Skewness coefficient is still rejected in all the other specifications, however.

Table 11: Correlation of Inflation and
Frequency Std. Deviation Skewness

Model	Frequency	Std. Deviation	Skewness
Calvo	0	0	+
Golosov and Lucas	+	-	-
Multi Sector Menu Cost	+	-	-
Midrigan	+	-	-
Observation and Menu Costs	+	0	-
Rational Inattention	+	0	+
Data	+	-	+

For each model, the signs that match the data are colored in blue, while those that do not are red. We do this to highlight the fact that in the broad class of state-dependent price setting models that we consider, none match the data in all the dimensions that we have presented. In particular, this highlights the usefulness of the inflation-skewness correlation as a statistic to test the existing menu cost models. As we have already argued, these models make a counterfactual prediction with this statistic because of the state-dependence that underlies them. It is also worth noting that the Calvo and rational inattention models have the same predictions, and therefore disagree with the data in the same ways. In the next section, we consider a menu cost model that weakens state-dependence and can be reconciled with the empirical correlations that we find.

4 A Generalized Menu Cost Model

We present a model that fits into the general framework of Section 2: the demand system and technology faced by the firm is the same, but we generalize the price setting problem in the following way: the menu cost faced by each firm every period is random. Formally, the period profit function of the firm takes on this form:

$$\Pi_t(z) = p_t(z)y_t(z) - W_t L_t(z) - \chi_t(z)W_t I\{p_t(z) \neq p_{t-1}(z)\}, \chi_t(z) \stackrel{iid}{\sim} G(\chi)$$

The difference with the Golosov and Lucas model is that now the menu cost can vary over time and across firms, the difference with the Midrigan model is that the distribution of menu costs is generalized, and as opposed to the Nakamura and Steinsson model, the menu cost for any given firm here varies over time. ⁷The

⁷This set-up can replicate the Golosov and Lucas model, if the menu cost distribution is degenerate.

assumption of random menu costs is similar to that made by Dotsey et al. (1999), but we present it within the framework we have been using until now⁸.

4.1 Background on Random Menu Costs

We choose to modify the model in this way for several reasons. The first is that it gives us a very general model that nests the menu cost models we have considered, as described above. In addition, this approach has a close relation to another, even more general approach already pursued by Caballero and Engel in a series of papers (1993, 2006a, 2006b). They propose thinking about price adjustment through the price adjustment hazard function⁹, which is the probability of a price change occurring as a function of the deviation of the current price from its optimal value (p^*):

$$H(x) = P(\Delta p | p^* - p = x)$$

Any of the models we have considered will imply a price adjustment hazard function, so the hazard function can be a helpful object to summarize the important features of each model. This function also describes the degree of the selection effect, as it indicates to what extent prices have a higher chance of adjusting the more misaligned they are. Caballero and Engel (2006a) show that, within this framework, aggregate price flexibility (or the inverse of monetary non-neutrality) can be expressed as the sum of two components:

$$\int H(x)f(x)dx + \int xH'(x)f(x)dx$$

where $f(x)$ is the probability density of the desired price gap, x . The first term in this sum is the frequency of price change, and the second is what Caballero and

erate, and the Midrigan model, if the distribution is discrete with two support points (one being zero, the other being positive). The Calvo model is replicated when the higher support point is infinite. Since the Nakamura and Steinsson model involves different firms facing different menu costs that are fixed over time, it is not encompassed by our set-up.

⁸The key differences with Dotsey et al. (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 micro data to place restrictions on the menu cost distribution, which is what the present exercise is about.

⁹It must be noted that this is distinct from the hazard function of price change estimated by Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008), which is the function $\lambda(t)$ that gives the probability that a price will change after t periods, given that it has already stayed constant for t periods. This function gets at timing of price changes, and at whether prices become more likely to change the longer they have stayed constant. While this can also help to distinguish between models, we do not estimate the timing hazard function in this paper.

Engel refer to the extensive margin, which incorporates the selection effect. To illustrate how this works, the hazard function corresponding to the Calvo model is a constant (the average frequency of price change), so that the first term is equal to that constant, and the second is zero. This shows that in general, and as long as the hazard function is increasing in the absolute value of the price gap (which is true for all models considered), the Calvo model gives a lower bound on the degree of flexibility. In our random menu cost model, a particular menu cost distribution will imply a particular hazard function, and will therefore determine aggregate flexibility (and monetary non-neutrality) as shown by the expression above. In this way, there is a very tight relation between these approaches¹⁰.

A more structural approach to price stickiness that is also related to ours is Woodford (2009)'s model of rational inattention. He shows that by varying the cost of processing information, price setting under rational inattention in the style of Sims (2003) can also nest, as extreme cases, the single menu cost model (free information) and the Calvo model (infinitely costly information), as well as the spectrum in between, which he also describes with the adjustment hazard function implied by different information costs. In addition, there will also be a mapping between the value for the cost of information in that framework, and a specific distribution of menu costs in ours. While we do not derive this mapping, we do believe that the rational inattention framework (or another type of informational constraint) could provide a micro-foundation for the distribution of menu costs that we assume to be general at first, and then adjust to allow the model to fit the empirical facts that we have presented¹¹.

¹⁰It would naturally also be interesting to directly estimate this hazard function. In this paper we continue to work in the menu cost framework to maintain the structure of a General Equilibrium model and obtain a quantitative response to the question of monetary non-neutrality. However, Caballero and Engel (2006b) attempt to do this, although the price change moments that they had access to were limited in the informational value they provided for this. By the same logic we have put forth in this paper, the higher moment price change correlations will be very informative to estimate the hazard function, and we do this in a separate paper (Luo and Villar, 2015).

¹¹As 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 it is plausible that the 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 explicit modelling of the informational constraints or consumer considerations that underly it to future research.

4.2 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. One extreme case of this is of course perfect price flexibility, or firms being free to change their prices every period without facing any kind of cost for doing so (although this is inconsistent with the fact that most prices do not change on any given month). But right after this comes a menu cost environment such as the one in Golosov and Lucas: firms are still able to choose when to change their prices, but are subject to a fixed cost (that is small in typical calibrations, to match the frequency of price change in the data). Adding randomness to the menu cost makes the price change decision more exogenous to the firm, as an additional dimension of the problem (how much changing the price will cost) is now outside the firm's control (with the extreme being the Calvo model, where the opportunity to change price is completely exogenous). The Midrigan model (both in Midrigan (2011), and the simplification of it that we present) goes in this direction, and as a result the degree of monetary non-neutrality in that model 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 parametrize the distribution of menu costs in a way that enables us to do this.

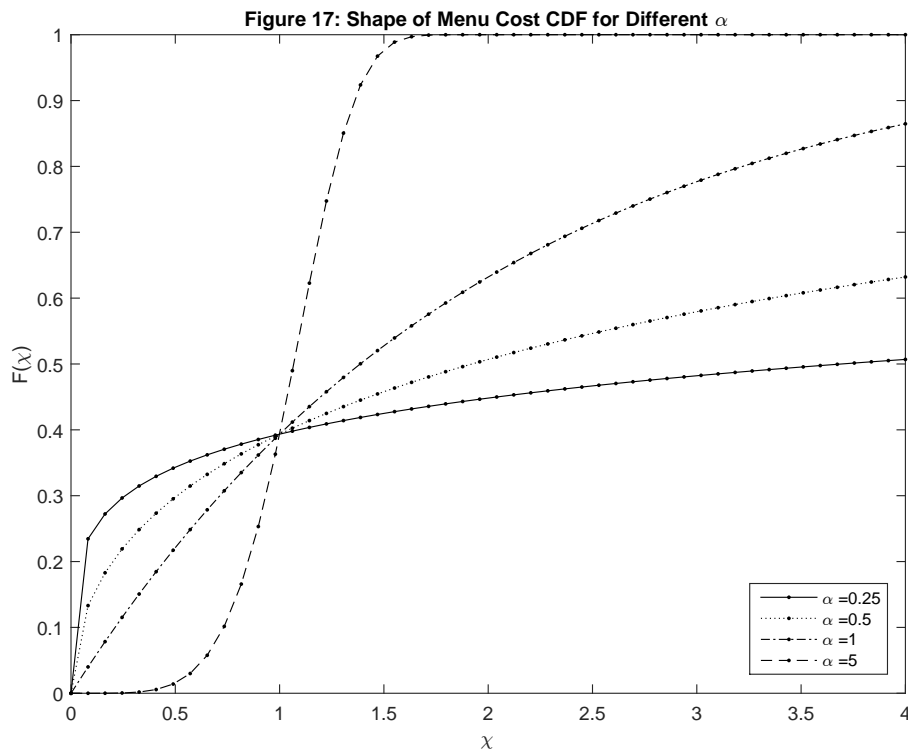
There are two important features that the menu cost distribution will need in order to achieve this: a positive probability of the menu cost being zero (of a free price change), which eliminates the “Ss” band or inaction region in the price setting problem, as some firms, facing a free price change, will choose to change their prices even if it is by a small amount. However, the Midrigan model already includes this, and as we have shown it also predicts a counterfactual inflation-skewness correlation. The other feature is that there must also 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 faced with these menu costs. This is important, because in the existing 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 in our generalized model 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. We assume that menu costs are iid 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:

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

In our version of the Midrigan model, the menu cost was either zero or a fixed positive value. The difference here is that instead of the positive value being fixed, it is drawn from a non-degenerate distribution. This distribution is a transformation of the exponential distribution (it is the same when $\alpha = 1$), and shares the important feature that the random variable is always positive. The difference is that α governs the curvature of the distribution function, which roughly corresponds to the fatness of the tails. Figure 16 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 values is relatively high, but the probability of very high values is also

quite high. When α is high, these extreme probabilities are low, and as α rises, the density concentrates on one value, approximating the case of a unique menu cost.

4.3 Calibration and Results

Our set-up has introduced 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 the shocks (p_ϵ), and the probability of a free price change (p_z) which was used in the Midrigan model. We set these parameters so that the model can match the empirical facts that we have discussed so far, which we divide into two categories:

1. From existing models: although these have not been the focus of our discussion, all the existing models match the average monthly frequency of price change and the average size of price change. Our model therefore matches the median of these statistics measured in our data. In addition, our empirical work has confirmed that, as previous studies had shown, the correlation between inflation and the frequency of price change is positive, so our model also matches this fact.
2. New moments: like the existing menu cost models, and consistent with the data, our model will imply a strongly negative correlation between inflation and the dispersion of price changes. The novelty will be that the implied correlation between inflation and the skewness of price changes will be non-negative, as in the data.

Table 12 presents the parameter values that we choose to match these moments, and Table 13 shows the moments attained by the model, compared to their empirical values.


Table 12

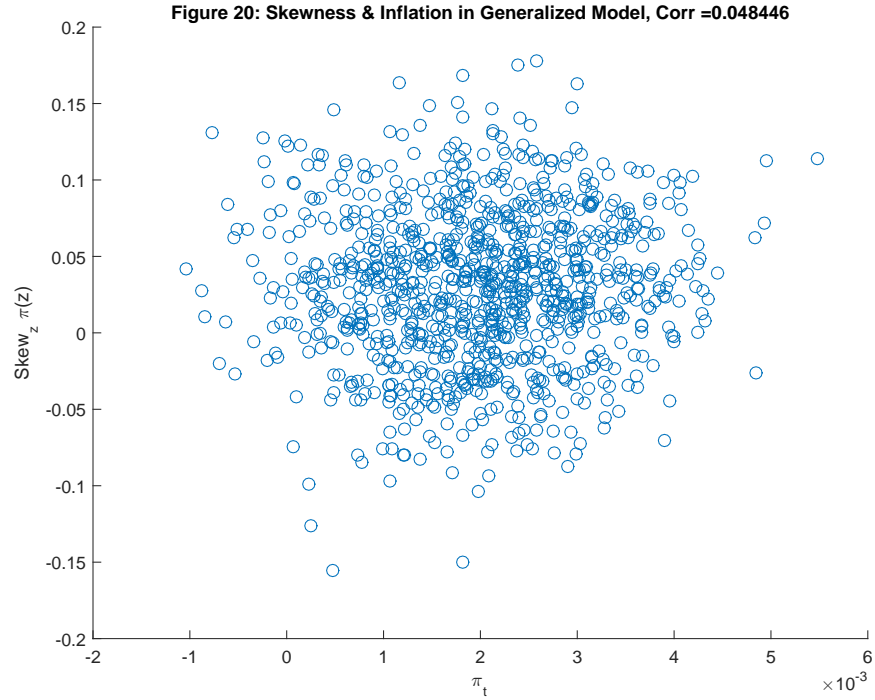

Parameter	Description	Value
λ	Inv. average menu cost	0.1925
α	Fatness of tails of MC	0.27
p_z	P(zero MC)	0.056
p_ϵ	P(idio. shock)	0.345
σ_ϵ	Size of idio. shocks	0.101

Table 13

Moment	Model	Data
Avg. Frequency	11.3%	11.3%
Avg. Size	8.0%	8.0%
Corr(IQR, π)	-0.59	-0.70
Corr(Skew, π)	0.05	0.39
Corr(Freq, π)	0.58	0.63

The first two moments are matched almost exactly. For the empirical value of the correlations, we present the results for the quarterly correlations involving the raw data, including all time periods, and excluding suspicious small price changes (for dispersion and skewness), and the weighted median for the frequency. The model matches the dispersion and frequency correlations quite closely. However, the skewness correlation in the model is close to zero, while it is clearly positive in the data for the whole sample. Before explaining this in more detail, we illustrate these correlations with scatter plots for the generalized model under the calibration above in Figures 17-19.






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%)¹², the correlation also ap-

¹²Inflation 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

appears to be close to zero in the data. We carry out the same “long-run” analysis as in Figure 5 above: 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 rises, and the correlation between period-by-period price change skewness and inflation (the same correlations we have been focusing on so far) also rises. 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), we should expect to see the skewness rising with inflation. In addition, this also makes our model stand out even more from the existing ones, as the other menu cost models feature a declining average price change skewness as steady-state inflation rises (and a period-by-period skewness correlation that is always negative). Figure 20 below shows this clearly by plotting the steady-state skewness correlations for the Midrigan model and our heterogeneous menu cost model separately.



This pattern highlights how the steady-state (or trend) inflation plays an important role behind our model’s non-negative skewness correlation. Indeed, positive trend inflation leads firms to expect positive future inflation when considering whether to re-set their prices. This will lead them to be less likely to cut their prices, even when facing an idiosyncratic (or aggregate) shock that would reduce their current desired price. This asymmetry in firms’ willingness to cut prices also means that the left tail of the price change distribution will be less responsive to aggregate

are the same. This is a direct result of the differences in monetary non-neutrality in the models, as higher non-neutrality means that the same nominal shocks have a greater effect on real consumption (and induce greater real variation), leading to less variation in inflation. This is shown below.

shocks, weakening the mechanism that led to the negative skewness correlation in the existing models.

What these results and figures make clear is that the generalized menu cost model that we presented, in making menu costs random in a way that weakens the selection effect, matches the important empirical facts that have been the focus of previous work on sticky prices as well as the existing models, and overturns the counterfactual prediction of these models that we have emphasized. We now show what this means for the degree of monetary non-neutrality.

4.4 Monetary Non-Neutrality

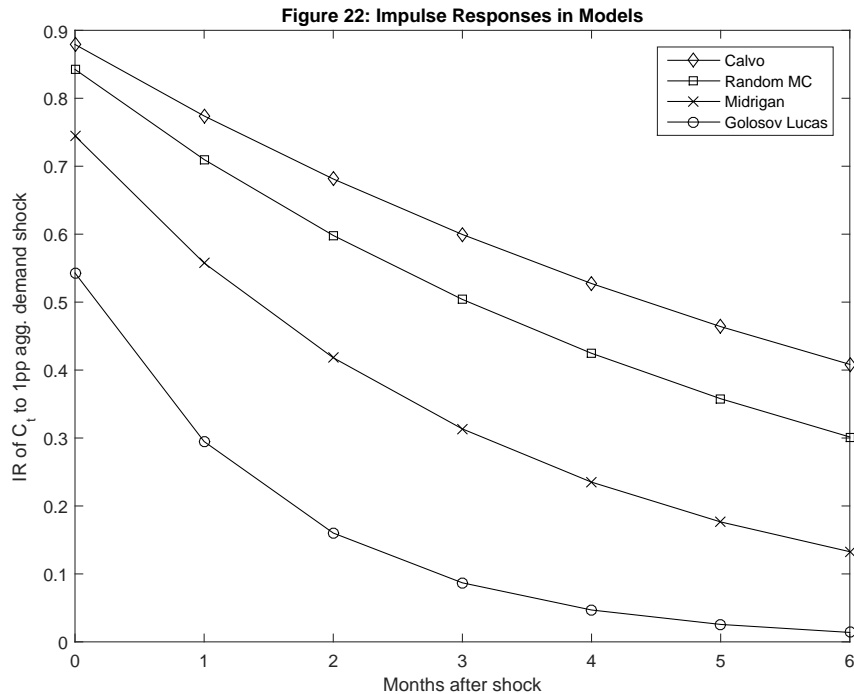
Monetary non-neutrality in these models is defined as the variation in real consumption induced by the nominal aggregate demand shocks, which are the only aggregate shocks, and we compare this statistic for the Calvo model, the Golosov and Lucas and Midrigan menu cost models, and our generalized menu cost model. As we have explained, making the menu cost distribution 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 this model would imply a greater degree of monetary non-neutrality. Table 14 below provides a quantitative illustration of this.

Table 14: Monetary Non-Neutrality

Model	$\text{Var}(C_t) * 10^4$
Golosov and Lucas	0.05778
Midrigan	0.17588
Generalized Menu Cost	0.33617
Calvo	0.47696

As Golosov and Lucas (2007) had famously shown, their model features a trivial amount of monetary non-neutrality compared to the Calvo model. Between the menu cost models, the major difference is between the baseline (Golosov and Lucas) and the others. Allowing for small price changes, as the Midrigan model does, leads to a very large increase in monetary non-neutrality, and this was emphasized by Midrigan (2011). However, our generalized model goes further by giving firms a large probability of effectively not being able to change their price, and yields an even higher level of non-neutrality. The Calvo model still has a higher degree of monetary non-neutrality, but our model gets significantly closer than the others. To

further illustrate the differences between the models, in Figure 21 we plot the impulse response of real aggregate consumption to a one percentage point increase in nominal aggregate demand in the same four models:



Not only is the effect on real activity greater on impact in our random menu cost model, but the effect is also considerably more persistent than in the existing menu cost models. It is clear that while in the Golosov and Lucas model the real effect of a nominal shock is small and transient, it is not so in our model, which has used the inflation-skewness correlation to evaluate the strength of the selection effect. Finally, the Calvo model gives a closer approximation to monetary non-neutrality than the other menu cost models.

5 Conclusion

The literature on sticky prices has made extensive use of price micro data to discipline models of price setting, and in this way the models have conformed more and more to important aspects of the dynamics of price changes. This line of work has notably enabled the study of monetary non-neutrality to be more grounded in data. However, an important limitation of the work done so far is that it has mostly used data for low inflation environments. Since the models in question are designed to study how

prices respond to aggregate shocks, it is helpful to be able to observe the behavior of price changes under large aggregate shocks and high inflation.

Our paper contributes to this by using price data from the U.S. going back to the late 1970's to compare how the price change distribution changes with inflation, to the predictions of a wide range of sticky price models. 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 site, follows very intuitively from the selection effect that is central to menu cost models. 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 non-neutrality that is considerably higher than what is predicted by the Golosov and Lucas model, and higher still than the Midrigan model.

The distinction between menu cost and Calvo models, or between state- and time-dependent pricing models has taken an important place in this literature. Much work has been done to show how these two ways of modelling pricing stickiness yield such different implications on monetary non-neutrality, and to determine which models are best at matching empirical facts. Our paper contributes to this line of work by introducing statistics not previously considered that are very useful to discriminate between the different models. In addition, we follow Woodford (2009) in presenting the distinction between time- and state-dependent models as a continuum, or spectrum. Woodford (2009) 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, we 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. For now, along with Nakamura and Steinsson (2010) and Midrigan (2011), we show that the assumption made by Golosov and Lucas (2007) of firms facing a single, constant menu cost is starkly at odds with many aspects of the price data, and that monetary policy can be expected to have substantial and persistent effects on real economic activity.

References

- Alvarez, Fernando, Martin Gonzalez-Rozada, Andy Neumeyer, and Martin Bereja (2011a), "From Hyperinflation to Stable Prices: Argentina's Evidence on Menu Cost Models", Working Paper, University of Chicago
- Alvarez, Fernando, Francesco Lippi, and Luigi Paciello (2011b), "Optimal Price Setting With Observation and Menu Costs", *Quarterly Journal of Economics*, 126, 1909-1960
- Alvarez, Fernando, Herve Le Bihan, and Francesco Lippi (2014), "Small and Large Price Changes and the Propagation of Monetary Shocks", Working Paper
- Anderson, Eric T., and Duncan I. Simester (2010), "Price Stickiness and Customer Antagonism," *Quarterly Journal of Economics*, 125, 729-765
- Anderson, Eric T., Benjamin A. Malin, Emi Nakamura, Duncan I. Simester, and Jón Steinsson (2015), "Informational Rigidities and the Stickiness of Temporary Sales", Working Paper
- Barro, Robert J. (1972), "A Theory of Monopolistic Price Adjustment", *The Review of Economic Studies*, 39(1), 17-26
- Bils, Mark, and Peter J. Klenow (2004), "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112, 947-985
- Blanco, Julio A. (2015), "Optimal Inflation Target in an Economy with Menu Costs and an Occasionally Binding Zero Lower Bound", Working Paper
- Caballero, Ricardo J., and Eduardo Engel (1993), "Microeconomic Adjustment Hazards and Aggregate Dynamics," *Quarterly Journal of Economics*, 108 (2), 359-383
- Caballero, Ricardo J., and Eduardo Engel (2006a), "Price Stickiness in Ss Models: New Interpretations of Old Results," *Journal of Monetary Economics*, 54(1), 100-121
- Caballero, Ricardo J., and Eduardo Engel (2006b), "Price Stickiness in Ss Models: Basic Properties", Working Paper
- Calvo, Guillermo A. (1983), "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12, 383-398

- Caplin, Andrew, and Daniel Spulber (1987), "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics*, 102, 703–725
- Carlton, Dennis W. (1986), "The Rigidity of Prices," *American Economic Review*, 76, 637–658
- Cecchetti, Stephen G., (1986) "The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines," *Journal of Econometrics*, 31, 255–274
- 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(2), 655-690
- Eichenbaum, Martin, Nir Jaimovich, Sergio Rebelo, and Josephine Smith (2013), "How Frequent Are Small Price Changes", Working Paper
- Gagnon, Etienne (2009), "Price Setting during Low and High Inflation: Evidence from Mexico", *Quarterly Journal of Economics*, 124 (4), 1221-1263
- Golosov, Mikhail, and Robert E. Lucas (2007), "Menu Costs and Phillips Curves," *Journal of Political Economy*, 115, 171–199
- Kashyap, Anil K. (1995), "Sticky Prices: New Evidence from Retail Catalogs," *Quarterly Journal of Economics*, 110, 245–274
- 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 (3), 863-904
- Krusell, Per, and Anthony Smith (1998), "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy*, 106, 867-896
- Luo, Shaowen, and Daniel Villar (2015), "The State-Dependent Price Adjustment Hazard Function: Evidence from High Inflation Environments", Working Paper forthcoming
- Midrigan, Virgiliu (2011), "Menu Costs, Multi-Product Firms and Aggregate Fluctuations", *Econometrica*, 79(4), 1139-1180
- Nakamura, Emi, and Jón Steinsson (2008), "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 123, 1415–1464


- Nakamura, Emi, and Jón Steinsson (2010), “Monetary Non-Neutrality in a Multisector Menu Cost Model,” *Quarterly Journal of Economics*, 154 (4), 961-1013
- Nakamura, Emi, Jón Steinsson, Patrick Sun, and Daniel Villar (2015), “The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation”, Working Paper
- Newey, Whitney K., and Kenneth D. West (1987), “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix”, *Econometrica*, 55(3), 703-708
- Sheshinski, Eytan and Yoram Weiss (1977), “Inflation and Costs of Price Adjustment”, *Review of Economic Studies*, 44, 287-303
- Sims, Christopher A. (2003), “Implications of Rational Inattention”, *Journal of Monetary Economics*, 50, 665-690
- Tauchen, George (1986), “Finite State Markov-Chain Approximation to Univariate and Vector Autoregressions”, *Economic Letters*, 20, 177-181
- Vavra, Joseph (2014), "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation", *Quarterly Journal of Economics*, 129(1)
- Woodford, Michael (2009), “Information-Constrained State-Dependent Pricing”, *Journal of Monetary Economics*, 56, 100-124
- 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, 514-533

Appendix

A Data Set and Statistics

As mentioned in the main text, the data set we use for our empirical analysis is the micro data underlying the U.S. CPI for the period 1977-2014, with the previously unavailable period being 1977-1986. One of us (Daniel Villar) worked intensively in the process of re-constructing this data set from the micro film made available by the Bureau of Labor Statistics. This process is described in detail in Nakamura et

al. (2015), and it leaves us with a large data set that tracks the prices of individual, narrowly-defined products in a monthly or bi-monthly frequency. We then combine this data set with the existing CPI data (1987 onwards), and that forms the data set for our analysis. Figure A1 below shows the size of our sample month by month. We plot both the number of non-missing available prices each month, as well as the number of price change observations available. The distinction is important, because we are always interested in price *change* statistics. The number of price observations is greater than the number of price change observations because for the price change to be observed in a particular month, we need both the current month’s price, and last month’s price. So when a product has a missing price for some month, the price change will be missing for that month and the following month.



The BLS makes a considerable effort to ensure that the prices of individual products are tracked, so that the price changes cannot be attributable to changes in any product characteristics. This conforms with our goals very well, as we are also only interested in price changes of identical products. An individual product could be, for example, a two quart bottle of Diet Coke in a particular supermarket location in New York City, or a specific futon model in a particular furniture store in Los Angeles. The BLS also identifies whenever a product substitution occurs, or when a new “version” of a particular product is introduced. Since a change of version indicates that some characteristic of the product has changed, we treat a new version as an entirely new product, and only compute price changes by comparing price changes within identical versions. We compute price changes as the difference of the log price, or:

$$\Delta p_{it} = \log\left(\frac{P_{it}}{P_{it-1}}\right)$$

As discussed in Section 3, we exclude observations for which there is any indication that the price was not actually observed but imputed, and for which the product was on sale. There are therefore missing observations in the price spells that we use. To compute the price change for any given month, we compare the price for that month to the previous month's price, when it is available. When the previous month's price is not available, we compare the current price to the price from two months before. Without this, we would have to drop a significant amount of data, as many prices are only sampled every two months. Since price changes are relatively infrequent, we believe that it is overwhelmingly likely that if a price changed between any two months, it only changed once, which means that we are observing the true price change, whether it occurred in the current or previous month. This is then not extremely important, as for much of our analysis we combine the price changes by quarter or year.

With the price change observations, we then form distributions of these price changes, keeping only the non-zero changes, for each period (either month, quarter, or year). For the diperson and skewness statistics, we first separate observations into categories that we label major groups. There are thirteen of these, and table A1 below provides a list, along with the share of expenditure weight that they represent.

Table A1

Major Group	Weight (%)
Processed Food	8.2
Unprocessed Food	5.9
House Furnishings	5.0
Apparel	6.5
Transportation	8.3
Medical Care	1.7
Recreation	3.6
Edu. Supplies	0.5
Miscellaneous	3.2
Services	38.5
Utilities	5.3
Gasoline	5.1
Travel Services	5.5

Services represent the lion's share of the weight. We then compute the dispersion and skewness statistics from each major group, and for each time period we then take

an expenditure-weighted average of the statistics, which represents the value of the statistics that we will use. If, for example, $Skew_{kt}$ is the skewness of the distribution of price changes in major group k and period t , then the value of skewness that we use in our analysis, $Skew_t$, is given by:

$$Skew_t = \sum_k w_k Skew_{kt}$$

We follow the same method for the dispersion, and thus obtain time series for the skewness and dispersion of price changes. This also applies for the frequency, but there we calculate the frequency first by ELI, which is a much narrower category. That is because the frequency is merely an average of the dummy variable indicating whether a price has changed or not, and it is calculated based on the number of price change observations (zero or non-zero), while the other moments are only calculated based on the non-zero changes (which gives fewer observations). This means that the frequency can be estimated with reasonable precision by ELI. Finally, the expenditure weights that we use are those from the 1998 revision of the CPI, which are the latest ones available. Different weights were used for 1977-1987 and 1988-1997, but we keep the weights constant throughout the sample so that changes in the weights do not induce changes in the statistics that we estimate.

B Computational Procedure and Calibration

We solve the sticky price models in this paper by value function iteration, following the method described in Nakamura and Steinsson (2010). The main difficulty with this method applied to this type of problem is that an important variable entering the firm's profit function is the aggregate price level. Since its future evolution depends on each firm's price, every firm's current state is, in principle, a state variable for all firms, making the problem intractable. To get around this, we follow the example of Krusell and Smith (1998) and approximate the law of motion of the price level with a finite number of moments, as in Nakamura and Steinsson (2010). In particular, we impose that firms perceive future inflation to depend only on future nominal aggregate demand (S_t , which is exogenous), and the current price level:

$$\pi_t \equiv \log\left(\frac{P_t}{P_{t-1}}\right) = \Gamma\left(\frac{S_t}{P_{t-1}}\right)$$

Under this assumption, the state space can be reduced to three dimensions: the firm's idiosyncratic productivity (exogenous), the firm's relative price (choice vari-

able), and real aggregate demand (C_t , which determines the real wage in equilibrium). The latter is endogenously determined, but the probability distribution of its future value is known fully with the law of motion of nominal aggregate demand, and the assumed law of motion of inflation.

The firm's problem can therefore be written recursively with the following Bellman equation:

$$V(A_t(z), \frac{p_{t-1}(z)}{P_t}, \frac{S_t}{P_t}) = \max_{p_t(z)} \left\{ \Pi_t^R(z) + E_t \left[D_{t,t+1}^R V(A_{t+1}(z), \frac{p_t(z)}{P_{t+1}}, \frac{S_{t+1}}{P_{t+1}}) \right] \right\}$$

, where $V(\cdot)$ is firm z 's value function, $\Pi_t^R(z)$ ¹³ is firm z 's real profits at time t , and $D_{t,t+1}^R$ is the real stochastic discount factor between time t and $t+1$. Our procedure to solve the model then closely follows Nakamura and Steinsson (2010): First, we discretize the state variables and propose a guess for the function $\Gamma(\frac{S_t}{P_{t-1}})$ on the grid. Then, we solve for the firm's policy function, F ¹⁴, by value function iteration, using the proposed $\Gamma(\cdot)$ function, the stochastic processes for the exogenous variables (applied using the Tauchen (1986) method), and the menu cost structure of the firm's problem. We then check whether F and Γ are consistent, by computing the price level (and inflation) implied by F for each value on the $\frac{S_t}{P_{t-1}}$ grid and comparing it to the value given by Γ . If they are consistent, we stop and use F to simulate the models. If they are not consistent, we update Γ and go back to the value function iteration step and continue. To determine whether they are consistent, we compare the inflation values, grid point by grid point, and consider that they are consistent when the difference is smaller the difference in values between grid points.

The method described above applies to all the menu cost models (including the Calvo model). However, the imperfect information models are markedly different in several ways, and therefore require different methods. We solve these models using the same methods and parameter values used in the original papers (Alvarez et al. (2011b) for the observation costs model; Woodford (2009) for the rational inattention model), and use the policy functions to simulate the models.

As mentioned in Section 2, the existing menu cost models and the Calvo model are calibrated to match the median frequency of price change and the median average size of price change in the data. The way we compute these moments is by first

¹³It can be shown that the profit function under CES preferences and linear production using only labor can be written as $\Pi^R(A, \tilde{p}, C) = C\tilde{p}^{-\theta}[\tilde{p} - \frac{\omega C}{A}]$

¹⁴Because the value of the menu cost in our general model is stochastic, the policy function is also a function of the menu cost. However, because we assume that the menu costs are iid over time, they are not a state variable.

calculating the frequency of monthly price changes and the mean absolute value of price change by ELI-year. We then compute the median across the ELI frequencies for each year (to obtain an annual series for the median frequency) and to then take the mean across years. The average frequency that we obtain is 11.3%, and the average size of price change is 8.0%. For the Midrigan model (as well as our random menu cost model), we also target the fraction of price changes that are small (less than 1% in absolute value). We compute this as with the frequency and average size: evaluate fractions by ELI-year, and take weighted medians across ELI's. WE find a value of 12%. Table A2 below shows the model-implied moments for the Golosov and Lucas, Midrigan, and Calvo models, as well as our random menu cost model, and compares them to their empirical values:

Table A2

Model	Average Frequency (%)	Average Size (%)	Fraction Small (%)
Golosov and Lucas	11.1	8.0	0
Midrigan	11.0	8.0	12.4
Calvo	11.0	7.9	17.3
Random MC	11.3	8.0	12.2
Data	11.3	8.0	12.0

All the models match the frequency and size moments almost exactly, and the Midrigan and random menu cost models match the fraction of small changes very closely. The Calvo and Golosov and Lucas models over- and undershoot the empirical value, respectively, as they do not target it. Table A3 below shows the parameter values that we choose for these models.


Table A3

Parameter	Golosov and Lucas	Value
χ	Menu cost (as share of steady state revenue)	0.019
σ_ϵ	Std. dev. of idiosyncratic tech. shocks	0.042
Midrigan		
χ^{High}	Menu cost (when positive)	0.034
σ_ϵ	Std. dev. of idiosyncratic tech. shocks	0.076
p_z	Probability of free price change	0.037
p_ϵ	Probability of receiving idio. shock	0.153
Calvo		
α	Probability of price change	0.111
σ_ϵ	Std. dev. of idiosyncratic tech. shocks	0.197

For the multi-sector model, we use the same parameter values as in Nakamura and Steinsson (2010), which make the model match the average frequency and size of price change for each of 14 sectors.

C Additional Empirical Results

In Section 2, we presented results on the empirical result between inflation and various price change moments, using both scatter plots and regressions. For the scatter plots, the measure of inflation that we used was Core PCE inflation, which excludes food and energy prices that tend to be quite volatile (and that could be influenced by sectoral shocks that we do not consider in the models). In addition, since the PCE index is chained, it tends to yield a lower value for inflation than the CPI. However, for the regressions, we used CPI inflation because we include expected inflation as a control, and the survey of inflation expectations asks about expectations of CPI inflation specifically. We therefore used CPI inflation to make the two variables more comparable. In Figure A2 below, we plot the twelve month log change for both indexes. They both co-move very strongly, although the peak is much higher for the CPI.



In this section we show that our results do not depend on which inflation measure we use, so we present scatter plots with CPI inflation, and regression results with Core PCE inflation as the regressor. The only difference that this makes is that in the regressions, the absolute value of the coefficients on inflation are slightly larger, because core PCE inflation does not attain as high a value, so the estimated slope of the moments on inflation is smaller. We also present results using series filtered by a moving average smoother and seasonally adjusted by removing quarterly dummies. Again, the the same results hold, but they come out a bit more clearly. For all of these results, we focus on using the quarterly inflation and moment series, although the same results would hold with the monthly and annual series.

Figures A3-A6 below present scatter plots of the smoothed moment and inflation series.




Figure A4: IQR & Inflation, Quarterly Smoothed




Figure A5: Skewness & Inflation Smoothed, Quarterly


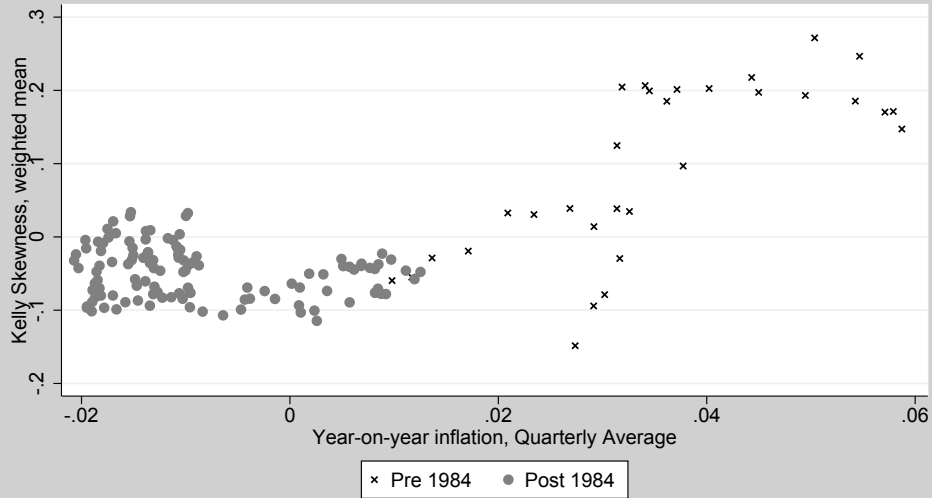


Figure A6: Kelly Skewness & Inflation Smoothed, Quarterly; Corr = 0.734



Figures A7-A10 are scatter plots using CPI inflation.

Figure A7: Frequency of Price Change & CPI Inflation, Quarterly




Figure A8: IQR & CPI Inflation, Quarterly






Figure A9: Skewness & CPI Inflation, Quarterly





The patterns in these scatter plots are the same as in the ones presented in Section 3. We further confirm these results with the regression tables below.

Table A4 (Core Inflation as regressor)
Coefficients for Frequency Regressions

Specification	Weighted Median		Weighted Mean	
	1977-2014	1985-2014	1977-2014	1985-2014
All	0.906*** (0.271)	1.362*** (0.313)	-0.046 (0.244)	-0.231 (0.305)
Fed Dummies	1.248*** (0.220)	1.503*** (0.214)	0.978*** (0.223)	0.281** (0.258)
Inflation Only	0.877*** (0.122)	1.083*** (0.253)	0.374** (0.173)	-0.580** (0.296)

Table A5 (Smoothed and Seas. Adj. Series)
Coefficients for Frequency Regressions

Specification	Weighted Median		Weighted Mean	
	1977-2014	1985-2014	1977-2014	1985-2014
Fed & Expected Infl	0.711*** (0.125)	0.796*** (0.210)	0.462 (0.138)	0.326* (0.189)
Fed Dummies	0.778 *** (0.075)	0.889*** (0.207)	0.723*** (0.109)	0.284* (0.163)
Inflation Only	0.716*** (0.062)	0.824*** (0.223)	0.437*** (0.105)	-0.178 (0.240)

Table A6 (Core Inflation as regressor)
Coefficients for IQR Regressions

Specification	All Observations		EJRS	
	1977-2014	1985-2014	1977-2014	1985-2014
Inflation Only	-0.412*** (0.060)	-0.676*** (0.081)	-0.461*** (0.068)	-0.803*** (-0.086)
Fed Dummies	-0.354*** (0.082)	-0.686*** (0.095)	-0.401*** (0.095)	-0.824*** (0.099)
Fed & Expected Infl	-0.366*** (0.127)	-0.485** (0.117)	-0.429*** (0.142)	-0.594*** (0.128)

Table A7 (Smoothed and Seas. Adj. Series)
Coefficients for IQR Regressions

Specification	All Observations		EJRS	
	1977-2014	1985-2014	1977-2014	1985-2014
Inflation Only	-0.301*** (0.043)	-0.493*** (0.073)	-0.330*** (0.047)	-0.561*** (0.086)
Fed Dummies	-0.241*** (0.048)	-0.495*** (0.084)	-0.249*** (0.054)	-0.556*** (0.097)
Fed & Expected Infl	-0.164** (0.069)	-0.377** (0.073)	-0.178** (0.075)	-0.431 *** (0.083)

Table A8 (Core Inflation as regressor)
Coefficients for Skewness Regressions

Specification	All Observations		EJRS	
	1977-2014	1985-2014	1977-2014	1985-2014
Inflation Only	4.537*** (1.306)	2.131 (2.062)	4.315*** (1.285)	1.658 (1.895)
Fed Dummies	7.546*** (1.686)	3.716 (2.270)	6.997*** (1.572)	3.396 (2.087)
Fed & Expected Infl	4.683 (2.870)	6.224* (3.316)	4.039* (2.657)	5.991 (3.136)

Table A9 (Smoothed and Seas. Adj. Series)
Coefficients for Skewness Regressions

Specification	All Observations		EJRS	
	1977-2014	1985-2014	1977-2014	1985-2014
Inflation Only	3.656*** (0.776)	1.208 (1.222)	3.263*** (0.776)	0.699 (1.148)
Fed Dummies	3.683*** (0.689)	0.925 (1.349)	3.404*** (0.680)	0.688 (1.245)
Fed & Expected Infl	0.969 (1.206)	0.453 (1.504)	0.785 (1.182)	0.152 (1.367)

Table A10 (Core Inflation as regressor)
Coefficients for Kelly Skewness Regressions

Specification	All Observations	
	1977-2014	1985-2014
Inflation Only	2.973*** (0.537)	-0.603 (0.512)
Fed Dummies	4.035*** (0.713)	0.504 (0.606)
Fed & Expected Infl	2.066** (1.047)	0.136* (0.721)

Table A11 (Smoothed and Seas. Adj. Series)
Coefficients for Kelly Skewness Regressions

Specification	All Observations	
	1977-2014	1985-2014
Inflation Only	2.465*** (0.342)	-0.088 (0.394)
Fed Dummies	2.479*** (0.329)	0.282 (0.435)
Fed & Expected Infl	1.636** (0.731)	0.204 (-0.430)

What these tables show is that while the size of the coefficients varies somewhat across specifications, the results presented in section 2 still hold: the frequency of price change rises with inflation, the dispersion falls, and the skewness does not fall with inflation (the relationship is positive but not significant in the low inflation period, and positive and mostly significant in the whole sample).