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

Shaowen Luo†

Department of Economics

Virginia Tech

Daniel Villar‡

Federal Reserve Board of Governors

August 29, 2017

Abstract

We present a new way of empirically evaluating various sticky price models that are used to assess the degree of monetary non-neutrality. While menu cost models uniformly predict that price change skewness and dispersion fall with inflation, in the Calvo model both rise. However, CPI price data from the late 1970's onwards shows that skewness does not fall with inflation, while dispersion does. We develop a random menu cost model that, with a menu cost distribution that has a strong Calvo flavor, can match the empirical patterns. The model therefore exhibits much more monetary non-neutrality than existing menu cost models.

JEL classification codes: E31, E32, E47, E52

1 Introduction

The dynamics of price changes — that is, when, how, and why firms change the prices of the goods and services that they sell — has been an area of active research in monetary economics over the past few decades. It is well known that in sticky price models monetary variables have no

*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. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Board of Governors or the Federal Reserve System.

†sluo@vt.edu, 3016 Pamplin Hall, 880 West Campus Drive, Blacksburg, VA 24060

‡daniel.villar@frb.gov, 20th & Constitution Ave. NW, Washington, DC 20551

influence on real economic activity (monetary neutrality) if all prices can be freely re-set to their profit-maximizing value at any point in time. Therefore, much effort has been devoted towards incorporating frictions in price-setting models and using detailed price data to measure how sticky prices really are (Klenow and Kryvtsov, 2008; Nakamura and Steinsson, 2008, etc.). One important finding in this literature is that the degree of monetary neutrality depends not only on how often prices change, but also, crucially, on the extent to which the prices that change are selected based on their misalignment (Caplin and Spulber, 1987; Golosov and Lucas, 2007, etc.). Specifically, if the prices that change are those heavily misaligned from their optimal level (as would be the case if firms must pay price adjustment, or menu costs), money will have a much smaller real effect than if they were randomly selected, as is the case in Calvo-type models.

This paper evaluates the strength of the selection effect based on an empirical pattern that has not been previously considered — the correlation between inflation and price change skewness, and a new data set of prices covering high inflation periods — the Consumer Price Index (C.P.I.) micro data going back to 1977. Moreover, we present a new random menu cost model of sticky prices to capture the pattern and use the model to assess the degree of monetary non-neutrality.

We exploit the fact that the selection in price changes (which cannot be directly observed) has a strong impact on the behavior of the distribution of realized price changes (which is observed in our data) in response to shocks (such as money supply shocks or aggregate demand shocks). In menu cost models, the presence of a fixed adjustment cost induces a selection effect: the price of a good or service changes if and only if the price change generates a profit greater than the menu cost. This leads to an inaction region of price changes that are too small to be justified. Relatively large price changes occur and respond to monetary shocks, which means that the average price response is large and monetary shocks have a small real effect. This selection effect also implies that an inflationary shock will push some price changes out of the inaction region to the positive side, and into the inaction region from the negative side. Thus, a positive shock will lead to 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). The opposite is true for a deflationary shock. This phenomenon is illustrated in Figure 1. Indeed, we show that a broad class of menu cost models, because of the selection effect created by the presence of an adjustment cost, implies a strongly negative correlation between inflation and both dispersion and skewness of the *non-zero price change* distribution in the positive inflation region. Importantly, these are implications that can be empirically tested.

Figure 1: Intuition for the Menu Cost Model

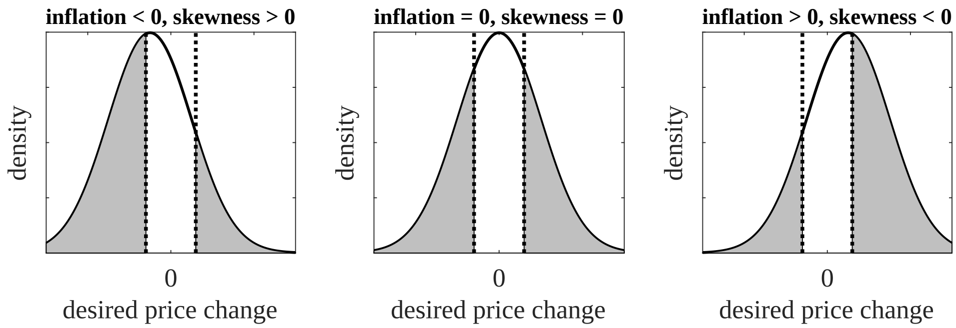

Note: In the first three panels, the black curve represents the distribution of the desired price change. The dashed lines represent the Ss band. The grey shaded area represents the distribution of realized price changes.

In the data, we find that while the dispersion of price changes clearly decreases in high inflation periods, the skewness does not. The latter is contrary to the predictions of menu cost models, and is therefore inconsistent with a very strong selection effect. Moreover, the negative inflation-dispersion correlation (and the fact that the frequency of price change rises with inflation in our data) both contradict what one would see in the Calvo model. Overall, we find that no existing model can match all the empirical patterns that we present. We use the data set recently presented in Nakamura et al. (2017), that 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.[1] The dataset overcomes an important limitation faced by the main source of price data in the sticky price literature, the micro data underlying the C.P.I., which was, until recently, only available going back to 1988 (while other commonly used data sets go back even less far), covering only periods of low and stable inflation.

To develop a model consistent with non-negative inflation-skewness and negative inflation-

[1] Although some studies (such as Alvarez et al., 2016a; 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. Notably, Alvarez et al. (2016a) look at the dispersion of prices (within narrow product categories), but not of price changes.

dispersion correlations, we modify the menu cost model in a way that weakens the selection effect. We do this by introducing random, heterogeneous menu costs that add randomness to whether the firm will have an opportunity to change its price. The model therefore incorporates some of the features of the Calvo model, and can be thought of as a hybrid between state- and time-dependent models. By working with random menu costs, we follow the example of Dotsey et al. (1999) and adjust the distribution of menu costs to fit the new correlations that we report. We find that in order to capture the non-negative inflation-skewness correlation, the probability of price changes being free must be nonzero and the probability of price changes being costly must be high. These correlations allow us to restrict the menu cost distribution in a way that Dotsey et al. (1999) could not, and allow us to obtain important implications for monetary non-neutrality. Indeed, our model features a much higher level of monetary non-neutrality than any of the existing menu cost models: it is approximately six times higher than in a standard menu cost model, and approximately 70% as high as in a Calvo model.

Understanding the selection effect is necessary to determine the extent of monetary non-neutrality due to price rigidity, and our paper contributes to a large literature devoted to studying this mechanism. Caplin and Spulber (1987) showed that in a highly stylized menu cost model, money can be completely neutral even if prices change very infrequently. Golosov and Lucas (2007) then developed a quantitative standard menu cost model calibrated to match certain empirical price change facts. In their model, monetary shocks have very small effects, and these results seriously called into question whether monetary policy could influence the real economy to the degree shown by the Calvo model. However, more advanced menu cost models, such as Midrigan (2011) and Woodford (2009), later included Calvo features in the price change policy, which reduce the degree of selection. These models are able to match some important features of the data, and generated considerably higher levels of monetary non-neutrality than the standard menu cost models.

While the sticky price literature has made clear the importance of the selection effect, assessing its strength is challenging given that it is a mechanism that 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. Existing work in this field has done this primarily by bringing quantitative price setting models together with the price data that has become available in the past decade (Midrigan, 2011; Nakamura and Steinsson, 2010, for example). These studies have, for the most part, used unconditional moments of the price change distribution (such as the frequency or size of price changes, averaged over time) to discipline the models in question. In this paper, we show that conditional higher moments of prices changes are extremely informative 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. This provides an opportunity to infer the degree of monetary non-neutrality.

Empirical studies of price stickiness in certain industries have been around for some time (e.g. Carlton, 1986; Cecchetti, 1986; Kashyap, 1995). However, it is only starting with Bils and Klenow (2004) that monetary economists have been able to start measuring statistics related to price stickiness for the economy as a whole. The work done by Bils and Klenow and the subsequent empirical studies on price stickiness (most notably Klenow and Kryvtsov, 2008; Nakamura and Steinsson, 2008) have enriched the discussion on monetary non-neutrality by providing empirical facts to evaluate the models that are used to study non-neutrality . Since Golosov and Lucas (2007), 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 (for example Alvarez et al., 2016b; Midrigan, 2011; Nakamura and Steinsson, 2010). 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. In addition, as mentioned before, our analysis is based on an extended version of the CPI micro data set that allows us to evaluate models with data from high inflation periods. Finally, Vavra (2013), Alvarez et al. (2016a) and Alvarez et al. (2016b) also consider higher moments of the price change or price level distribution and the implications for sticky price models.[2] We emphasize that the relationship between price change higher moments and aggregate inflation is particularly informative about the

[2] Alvarez et al. (2016b) in particular shows that in a broad class of models (covering many of the models that we consider) the kurtosis of price changes provides information on the degree of monetary non-neutrality. Using data from the CPI in France, they find that the degree of non-neutrality is between that in menu cost models and that in the Calvo model, broadly consistent with our final results. As we explain further on, we believe that our approach to inferring non-neutrality has practical empirical advantages.

selection mechanism.

The rest of the paper is organized as follows. 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.

2 The Skewness of Price Change in Sticky Price Models

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

2.1 Analytical Results

The main theoretical result of our paper is that in menu cost models, the dispersion and skewness of the distribution of price changes both fall as inflation rises. To obtain an analytical expression for the price change moments in the menu cost model, we derive a version of Alvarez et al. (2016a) (itself similar to the model of Golosov and Lucas (2007)) in this subsection.

The model is a continuous time model in which monopolistically competitive firms face idiosyncratic shocks, an aggregate price level that grows at a constant rate and a fixed cost (a menu cost) whenever they change their nominal price. The continuous time nature of the model, along with certain assumptions about the shock process, yield that a firm will change its price exactly at the point in time in which its relative price reaches an upper or lower threshold. This means that price changes will only take a single positive or a single negative value, so that the distribution of price changes will have two mass points. We will exploit the simplicity of this distribution to derive expressions for the dispersion and skewness, but first formally lay out the model.

With constant elasticity of demand (η) and constant returns to scale, the instantaneous profit of the firm is given by the function,

$$F(p,z) = e^{-\eta p}(e^p - e^z)$$

where p is the log of its price relative to the aggregate price (which grows at the inflation rate π), and z is the idiosyncratic shock and the log of the marginal cost. z follows a diffusion process

$$dz = \sigma dW - z dN,$$

where dW is a standard Brownian motion, and N is the counter of a Poisson process with a constant arrival rate per unit of time, ρ. This can be interpreted as the rate at which firms die and are replaced with a new firm (for which $z = 0$). This allows the model to maintain a stationary distribution of relative prices while the process for marginal costs contains a trend.

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

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

Alvarez et al. (2016a) show that the optimal policy is for the firm to follow a threshold rule summarized by a vector of constants $X \equiv [\underline{x}, \hat{x}, \overline{x}]$, with $\underline{x} < \hat{x} < \overline{x}$. The price is left unchanged as long as $p \in (\underline{x} + z, \overline{x} + z)$, and is adjusted to be set to $\hat{x} + z$ when it reaches either boundary of the inaction region.

When the relative price reaches the lower bound of the inaction region, the size of the price change is given by:

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

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

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

Regardless of the realization of z, positive price changes will only take one value, while negative price changes will also take one value that does not depend on z. This should not be particularly surprising. Indeed, given the assumptions made about demand and technology, a firm in this model is choosing to re-set its relative price, always to the same value, once the relative price reaches a certain value. Since there are not aggregate shocks that affect the growth rate of the price level, which is effectively being tracked, the percentage by which the nominal price has to change does not vary with the realization of z, and only depends on whether the upper or lower bound of the inaction region has been reached. The continuous nature of time and of the shock process is also important for this result.[3]

We focus on the realized price change distribution *excluding zero price changes*. As we show in the appendix, the variance and skewness of price changes distribution only depend on the size of price increases (Δ^+) and decreases (Δ^-), the frequency of price increases (λ^+) and decreases (λ^-), with the following functional form:

$$\begin{aligned} \text{standard deviation:} \qquad & E|\Delta p - \mu| = \sqrt{\beta(1-\beta)}(\Delta^+ + \Delta^-), \\ \text{skewness:} \qquad & \frac{E(\Delta p - \mu)^3}{[E(\Delta p - \mu)^2]^{\frac{3}{2}}} = \frac{1-2\beta}{\beta^{\frac{1}{2}}(1-\beta)^{\frac{1}{2}}}, \end{aligned}$$

where $\mu \equiv E(\Delta p | \Delta p \neq 0) = \frac{\lambda^+\Delta^+ - \lambda^-\Delta^-}{\lambda^+ + \lambda^-}$ denotes the mean of the price change distribution, $\beta \equiv \frac{\lambda^+}{\lambda^+ + \lambda^-}$ denotes the fraction of price changes that are increases.[4]

For low to intermediate values of inflation, the size of price changes (Δ^+ and Δ^-) do not change appreciably, as shown in Alvarez et al. (2016a). What does change is the fraction of price changes that are increases, β, which rises with inflation. It is straightforward to prove that the skewness of the price change distribution decreases with β, given that $\frac{\partial \text{skewness}}{\partial \beta} < 0$. In turn, this means that the relationship between price change skewness and inflation is negative, as β rises

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

[4] The kurtosis of the price change distribution in this set-up is closed related to the skewness of the distribution, given that kurtosis $= \frac{E(\Delta p - \mu)^4}{(E(\Delta p - \mu)^2)^2} = \text{skewness}^2 + \text{skewness} \cdot \frac{\beta(1+\beta)}{\beta^{0.5}(1-\beta^{0.5}}$. This result supports our point that the kurtosis measure as presented by Alvarez et al. (2016b) is related to the skewness measure as presented in this paper. However, the skewness measure has an empirical advantage, which is discussed in section 3.

with inflation. Moreover, the dispersion of the price change distribution falls with inflation when $\beta \geq 0.5$, and rises with inflation when $\beta < 0.5$. Figure 2 illustrates the size of price changes, the fraction of price increases and decreases, along with the second and the third moment of the price change distribution.

Figure 2: Price change moments of the continuous time model

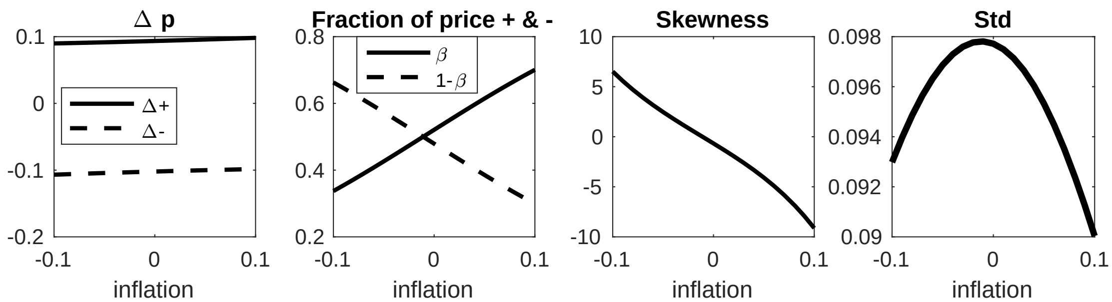

Note: X and β are solved following propositions 4 and 5 in Alvarez et al. (2016a) by setting $\rho = 0.1, \eta = 3, c = 0.02, \sigma = 0.15$, and firms discount parameter equals 0.06.

In what remains of the section, we will show that the same results hold in more general menu cost models, first by providing an intuitive explanation for the results, and then by presenting results from simulations. Before proceeding, it is worth pointing out the difference between our focus here and the focus in Alvarez et al. (2016a). First, Alvarez et al. (2016a) study the relations between inflation and different moments of the *price level* distribution, while we analyze the relation between inflation and different moments of the *price change* distribution. The moments that they focus on are the frequency of price change, average price changes, and the dispersion of relative prices. They do not, however, analyze the dispersion or skewness of the price change distribution.[5] We show here that the higher moments of the price change distribution display clear patterns that are different from those of the distribution of relative prices, and that provide information on the selection effect at work in menu cost models. We also believe that the moments of the price change distribution in sticky price models are particularly worth analyzing because the distribution of price changes can be easily constructed and observed with price micro data, while the distribution of relative prices requires settling on an aggregate price index with which to nor-

[5]The relation between the frequency of price change and inflation is useful to distinguish between the Calvo and menu cost models, and the dispersion of relative prices is closely related to the welfare costs due to inflation, so these objects are clearly worthy of attention.

malize individual nominal prices (which makes it difficult to deal with product heterogeneity) (As pointed out by Nakamura et al. (2017)).

2.2 Intuition for the Menu Cost Model

In the previous sub-section, we exploited the fact that the distribution of price changes had only two mass points in that particular version of the continuous time model. As we will illustrate, the negative relationships between inflation and dispersion as well as inflation and skewness hold in discrete time models, which will be the focus of our analysis from now on. This confirms that our results do not depend on the discrete nature of the price change distribution in the model analyzed so far. Indeed, the same patterns can be seen in menu cost models for which price changes take a range of values, more in line with the distributions seen in the data. In this subsection, we first present the intuition for the relationship between inflation and higher moments of the price change distribution.

Price change dynamics in the menu cost models that we will now analyze can be thought of in the following way: both idiosyncratic and aggregate nominal shocks to firms' optimal prices yield 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 price changes 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 1.

The presence of idiosyncratic shocks implies variation in firms' desired price changes, and nominal aggregate shocks move the position (average) of the underlying distribution. For example, a positive aggregate shock 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). 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. Consequently, *skewness*, a measure of the asymmetry of a distribution, or the relative sizes of the right and left tails, becomes 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). 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. This implies that the correlation between skewness and inflation is negative.[6] 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.

Another implication is that positive aggregate shocks reduce 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). The dispersion decreases with inflation in the positive region, and increases in the negative region, with the maximum attained at zero inflation (illustrated in Figure 9 in Appendix A). The intuition for this relationship has been applied by Vavra (2013) to explain why, in standard menu cost models, the frequency of price change and dispersion will move in opposite directions in response to aggregate shocks. What we show here is that the same logic leads to an observable relationship with inflation, and that it also applies to the skewness of price changes.

What makes these correlations interesting is that they have to do with the central mechanism of the menu cost model: the selection effect, (i.e. whether the prices that change are those most mis-aligned from their optimal value). Indeed, monetary non-neutrality is low in menu cost models (relative to the Calvo model) because the menu cost ensures that only large prices occur, which makes the average price response to aggregate shocks (and therefore aggregate flexibility) relatively high. In the same way, the fact that only relatively large price changes respond to such shocks is what leads the price change distribution to become less dispersed and more asymmetric

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

(lower skewness) in response. This makes the correlations that we emphasize particularly informative about the presence of the selection effect. Naturally, the selection effect has received much attention in recent research on sticky prices, as it makes a crucial difference to the degree of monetary non-neutrality (reference?). However, the fundamental difficulty in empirically evaluating the strength of the selection effect is that it involves the desired price change of firms. Since most firms' prices do not change in any given month, the desired price change is unobserved in most cases.[7] This makes it impossible to directly test whether the prices that change are those that are most mis-aligned, in line with the selection effect. Instead, one must make an inference based on the implications made by models for realized (and therefore observable) price changes. In this paper, we are presenting and implementing a new way of testing for the strength of the selection effect: the presence of selection in menu cost models implies the negative skewness and dispersion correlations (which are observable) that are the focus of our analysis, and this motivates our focus on these statistics.

2.3 Existing Models

We consider the empirical implications of the selection effect in the existing sticky price models, including the Calvo model, the Golosov and Lucas menu cost model and the variants of it that have appeared since. To do this, we consider models that can be separated into four categories: 1) Calvo, 2) Menu cost, 3) Observation costs, and 4) Rational Inattention. We choose six models in those categories to evaluate, namely the standard Calvo model, Golosov and Lucas (2007), Nakamura and Steinsson (2010), Midrigan (2011), Alvarez et al. (2011) and Woodford (2009).

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

[7] Eichenbaum et al. (2011) study the selection effect using a grocery store dataset containing input price changes. However, the exact desired price change is still missing, and their sample is restricted to grocery products.

tion of price changes more generally) will also be different across models. This is what we are documenting in this section, and below we provide a formal set-up of the models.

First, households maximize expected discounted utility of the following form:

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

There is 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). Following Midrigan (2011) and Vavra (2013), we assume that idiosyncratic shocks arrive infrequently with a Poisson probability p_ϵ, and model the process in the following way:

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

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

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

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

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 (2016) developed a menu cost model with a Taylor-type policy rule, we do not attempt this for the models in this section. Our goal is to show how the price change distribution changes with inflation under different sticky price models, and the aggregate demand process that we use enables us to do this.

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

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

In the standard Golosov and Lucas (2007) menu cost model, the cost χ is fixed for all firms and periods, and can be calibrated to match the frequency of price changes observed in the data. The idiosyncratic shock process is Normal AR(1), so p_ϵ is set to 1, and the standard deviation of shocks is calibrated to match the average size of price changes. This is, in a way, the most "state-dependent" model, as 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).

The first extension to the menu cost model that we consider is the Nakamura and Steinsson (2010) multi-sector menu cost model, in which firms are separated into sectors. Firms in different sectors face a different menu cost and variance of idiosyncratic shocks. Second, we also analyze the model in Midrigan (2011), who introduced other modifications to the standard menu cost model: first by changing the idiosyncratic shock process so that it would feature fat tails (which we described above), and giving firms a motive to make small price changes[8]. In his model, multi-product firms can change the prices of all their products by paying the menu cost. This enables the model to match the considerable fraction of small price changes that are observed in the data, but it also makes the model much more difficult to solve. We follow Vavra (2013) in simplifying the Midrigan model by assuming that, instead of producing multiple products, firms each period are randomly given the possibility of changing their price for free (with a low probability), or by paying a menu cost. The random menu cost structure yields similar results for monetary non-neutrality as introducing multi-product firms. This is also a variation of the CalvoPlus model presented by Nakamura and Steinsson (2010), and adds the probability of drawing a zero menu cost (free price change, p_z) as an additional parameter to calibrate. With the additional parameters in this model, we target the fraction of price changes that are small, as in Midrigan (2011).[9]

We also consider a Calvo model, which has the set-up described above, except that firms 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.

Finally, we also include two models involving imperfect information: the Alvarez et al. (2011) model of observation and menu costs, and the rational inattention model of Woodford (2009). In the former, firms must pay a fixed cost to observe the relevant state (or conduct a "price review"), and a menu cost to change their price. Facing such costs, firms conducting a price review choose the date of the next review, and a price plan until that date. Because the Alvarez et al. (2011)

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

[9]Midrigan (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 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.

model includes a menu cost, it features a high degree of selection. Woodford (2009) considers the same type of price-setting problem, but within the rational inattention framework proposed by Sims (2003): firms face a cost based on how much information they process, and therefore choose to receive limited information based on which they choose when to review prices. In this model, the cost of processing information is a crucial parameter, and both the Calvo model and standard menu cost model are nested as extreme cases of the information cost in this set-up (infinite and zero, respectively). Furthermore, intermediate values of the information cost result in what is 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.

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

To analyze short-run fluctuations (the first case), we solve each model with a fixed value for the parameters of the nominal aggregate demand process (μ and σ_η), and simulate each for a large number of firms and periods. From the simulated price series, we then compute the various price change moments for each period (obtaining a time series for each moment), and look at the relationship with the time series for inflation endogenously derived. Our steady-state analysis (the second case) is more in line with what is doneby other papers, such as Alvarez et al. (2016a). Because much of the variation in inflation throughout our sample period is generally understood to reflect regime changes caused by systematic changes to the conduct of monetary policy, it is important to consider whether the correlations in question are the same when it is steady-state inflation that changes. For this analysis, we solve each model with different values for the steady-state inflation parameter (μ, keeping all other parameters fixed), and for each solution computing the values of the price change moments from the model's stationary distribution. We find that the relationships between inflation and price change moments are qualitatively the same in both cases (that is, with respect to short or long run changes in inflation).

Figure 3: Simulated moments and inflation from different models

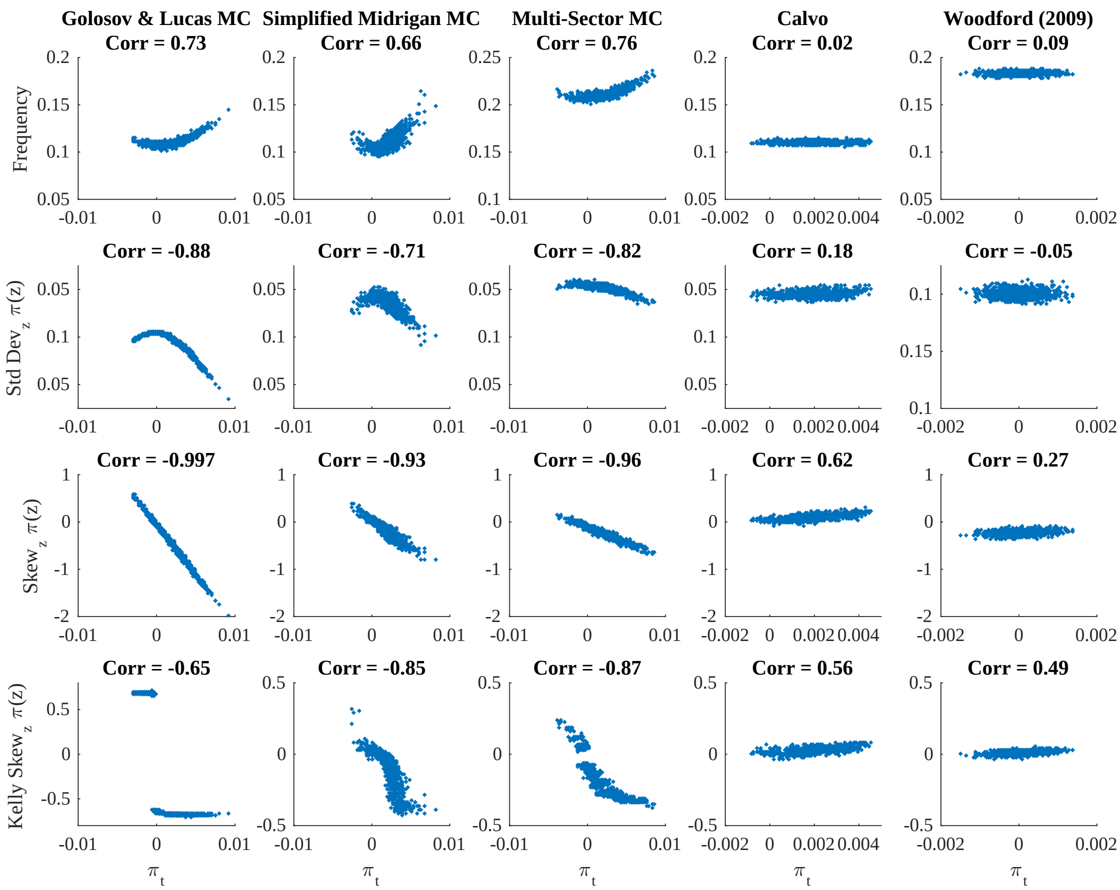

In order to further illustrate these results, we present scatter plots between inflation and the different moments from the simulations (based on 1,000 months and 50,000 firms) corresponding to the short-run analysis. Figure 3 shows the correlations for the frequency of price change, the dispersion and skewness of price changes, with a point representing a time period in the simulations.[10] 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

[10]The Alvarez et al. (2011) model contains no aggregate shocks. Therefore, the "short-run" analysis of this model is excluded. Strictly speaking, the Woodford (2009) model cannot be solved with aggregate nominal disturbances. Nonetheless, we take a simplified approach following Section 5 of Woodford (2009). We simulate the model with the dynamics of aggregate nominal expenditure being i.i.d. and mean zero to conduct the "short-run" analysis (refer to appendix A for detail). The "long-run" analysis of this model is excluded.

(2007) model for the dispersion): the skewness of price change falls very sharply with inflation in menu cost models, as does the dispersion for positive values of inflation (as explained above, the inflation-dispersion relationship is non-monotonic). In contrast, the same relations in the Calvo and imperfect information models are not so strong. However, the Calvo and rational inattention models feature weakly positive relationships for price change skewness and dispersion. That is because price changes are not selected in the Calvo model, so the mechanism described earlier is entirely absent.

The intuition for the correlations is easiest to explain in the case of the "standard" Golosov and Lucas model, as in subsection 2.2, yet it also applies to the other menu cost models. 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. Our simplified version of the Midrigan model involves firms randomly facing either a positive or zero menu cost. This weakens the selection effect, because there is now a positive probability that a firm will change its price even if it will be a small change, so that price changes are not entirely "selected" based on how out of line the original price is. However, the selection effect is still present to a certain extent, because it is only relatively large price changes that will happen with certainty (as those will be the only ones for which a firm will be willing to pay the positive menu cost, when it is faced). The tails of the price change distribution will therefore be very sensitive to the aggregate shocks that drive inflation in the model, leading to the same relationships for price change dispersion and skewness as in the Golosov and Lucas model.

Although the relationships come out very clearly in these simulations, it could be a concern that the higher moments that we are estimating might not be well defined in the distributions that we are working with. In addition, estimates of higher moments are very sensitive to outliers, which would be of concern particularly when we estimate from the data. That is why we also consider alternative measures for the dispersion and skewness of price change: the inter-quartile range (for dispersion) and Kelly's coefficient of skewness (as opposed to "moment skewness", which is what we have been estimating so far).[11] Since these statistics are quantile-based, they are well-defined

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

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). The last row of Figure 3 shows scatter plots of Kelly Skewness in the different models[12].

Figure 4: Simulated long run statistics from different models

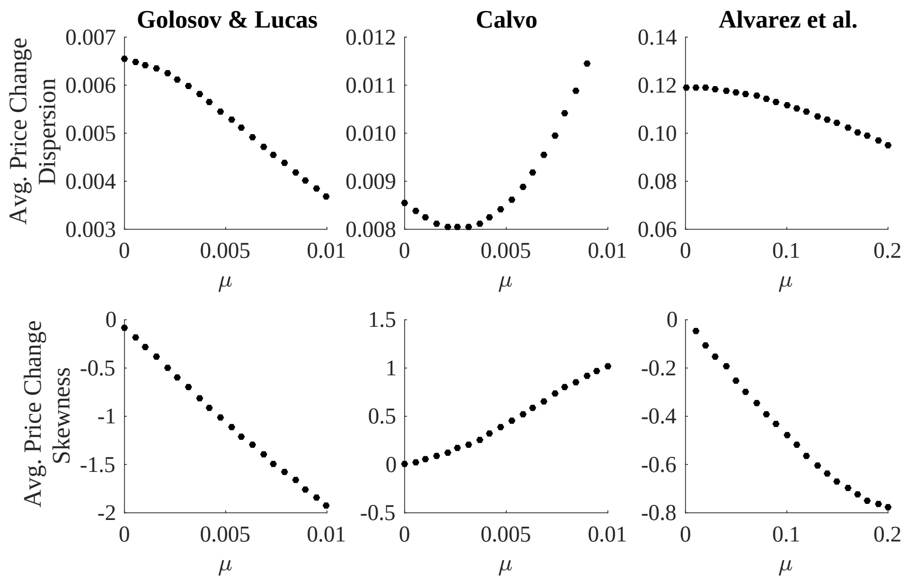

In Figure 4, we plot the results for the long-run analysis, in which we vary the value of steady-state inflation. For each model solution, we can construct a stationary distribution of price changes, from which we can then compute the stationary value for the different price change moments, and these are the values plotted in the figure.

What the scatter plots show is that, as in the "short-run" analysis, the dispersion and skewness of price changes fall with trend inflation in the menu cost model (we are only plotting results for the Golosov and Lucas (2007) model, but the same pattern holds for the other menu cost models). As in the short-run analysis, the Calvo model predicts weak positive relations for both moments with respect to steady-state inflation. 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

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

have the same sign in the four menu cost models (Golosov and Lucas (2007), Nakamura and Steinsson (2010), Midrigan (2011), and Alvarez et al. (2011)). The scatter plots show that the values taken by moments we report do vary across the models (for example, in the Golosov and Lucas (2007) 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 (2010) and Midrigan (2011) menu cost models were developed as extensions of the Golosov and Lucas (2007) model to make it match new empirical facts, and the changes made considerably weakened the selection effect that reduces the importance of monetary shocks. However, what we find here is that, despite the important changes made, they all have the same implications along the dimensions that we are considering.

3 Empirical Evidence from High Inflation Periods

In the previous section, we documented the predictions made by various sticky price models on the behaviour of price changes at different inflation rates. In this section, we present the data set that we use, and the empirical results that test the model predictions of the previous section.

3.1 Data Set and Construction of Statistics

Along with much of the sticky price literature, we make use of the micro data that underlies the U.S. Consumer Price Index (CPI). 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. (2017), and that paper also describes in detail just how the data set was re-constructed. 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). In addition, the data set contains information on when any given price is a temporary sale, or an imputation (not properly collected). Because of this, we are confident that we are observing the price changes of identical products and services, with the price being actually observed; and all of this with the same standards throughout the sample period.

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

Nakamura and Steinsson (2008, 2010) have shown that there is significant heterogeneity of price change statistics across sectors. We use their method to report the average overall frequency of price change: estimate the frequency of price change for each ELI, and then take a weighted average of the ELI frequencies (using the expenditure weights that go into the CPI). For the frequency of price change we consider both the aggregate weighted median and mean frequency.[13] For the dispersion and 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 has another advantage for testing the model predictions that we focus on. Indeed, the models do not allow for differences across sectors, such as sector-specific shocks. These 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 dis-

[13] Nakamura and Steinsson (2008) highlight the difference between the mean and the median, arising from the fact that the distribution of frequencies by ELI is very skewed to the right, with a few ELI's having very high frequencies. They argue that the median is a better measure of the average frequency 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.

tribution. 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 types of effects by computing statistics sector by sector.

3.2 Results

The goal of our empirical work is to determine whether the theoretical patterns documented above are borne out by the data. As in 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 monthly core PCE inflation. Sharp changes in headline inflation tend to be driven by the global market prices of food and commodities, which would not be well described by the price-setting models that we are working with, making core inflation preferable for us. However, we also compute correlations 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. All of these additional results can be found in Appendix C.

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. Quarterly and annual inflation averages also have the advantage of containing less noise than monthly inflation series, so we will focus on presenting results using quarterly series (although we include all the monthly and annual results in the appendix). Figure 9 in Appendix plots the quarterly time series that we construct for the Inter-Quartile Range and Skewness of price changes.

In the next subsection, we present the correlation results in two ways: first, with raw correlations and scatter plots (which are reported in Figure 5), as with the models. Secondly, we estimate these relationships with regressions (allowing us to test for significance and to include controls, which are reported in Table 1).

Figure 5: Moments of Price Change and Inflation, Quarterly
Source: Authors' calculations from BLS CPI Research Database

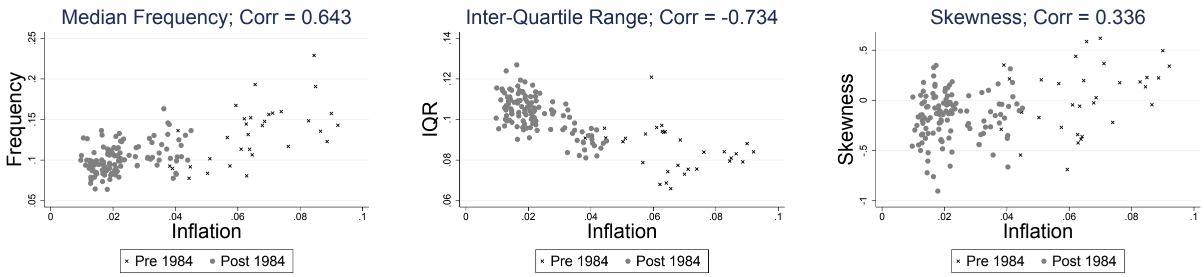

3.2.1 Correlations

We first verify that the frequency of price change rises with inflation, as found by Gagnon (2009) and Alvarez et al. (2016a). We present scatter plots using the quarterly moment and inflation series (the empirical counterpart to the simulation scatter plots from the previous section). Correlation values are reported in Tables 9-12 in Appendix C. Figure 5 confirms that there is a positive association between the frequency and inflation. 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.

Next, we look at the results for the moments that our discussion has focused on: the dispersion and skewness of price changes. 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). Skewness, while varying over time, does not change with inflation in a systematic way for low levels of inflation (although there does seem to be a positive relationship when inflation is high). 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 (as defined by Eichenbaum et al. (2013)) or apply seasonal adjustment and smoothing to the data series. Next, we formalize this analysis with linear regressions.

3.2.2 Regressions

Table 1: Coefficients on Inflation for Price Change Moments - Using CPI Data
Source: Authors' calculations from BLS CPI Research Database

	1977-2014			1985-2014		
	All	Fed Dummies	Inflation Only	All	Fed Dummies	Inflation Only
Frequency	0.708	0.728	0.771	0.777	0.810	0.587
	(0.071)	(0.095)	(0.237)	(0.224)	(0.208)	(0.252)
IQR	-0.296	-0.186	-0.257	-0.428	-0.414	-0.222
	(0.042)	(0.038)	(0.089)	(0.070)	(0.077)	(0.086)
Skewness	3.936	4.309	2.665	1.732	1.541	3.634
	(0.827)	(1.012)	(2.788)	(1.641)	(1.857)	(3.279)
Kelly Skewness	2.499	2.439	1.658	0.320	0.710	0.942
	(0.354)	(0.363)	(0.948)	(0.454)	(0.423)	(0.595)

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 what controls are used. Standard errors that are consistent for heteroskedasticity and auto-correlation of the residuals (Newey-West) are reported.

We now turn to regressions to determine whether these correlations are statistically significant, and to consider different control variables. The question of interest about the coefficients on inflation is not merely whether they are statistically significantly different from zero, but also whether they are significantly different from what the models predict. To do this, we estimate regressions of the frequency, dispersion (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:

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

where y_t denotes the different price change moments (frequency, dispersion, and skewness). Controls are included to address the fact that many important changes occurred in the U.S. monetary environment over our sample period, which could conceivably have a direct effect on the price change distribution. For example, expected inflation could affect firms' price setting decisions separately from present realized nominal shocks, so we include expected inflation (measured by the University of Michigan Survey of Consumers) as a control. 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. Table 1 show the estimates for β from these different specifications.[14]

These results support what the correlations showed: the frequency of price change rises with inflation and 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). These results confirm 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 showed in Figure 4.

Table 2: Coefficients on Inflation for Price Change Moment - Using Simulated Data

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
Calvo	-0.003	0.040	2.93	1.00
Rational Inattention	0.020	0.029	2.87	1.00
BLS CPI Data	0.708	-0.296	3.94	2.50

Table 2 presents the coefficients on inflation from regressions of the same type, but run on simulated data from the different models. The last row presents the coefficients using CPI data, which replicates the first column of Table 1. 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

[14] Regression results excluding certain small price changes based on Eichenbaum et al. (2013) are presented in Table 21 in appendix C.

intervals that we estimate for the skewness coefficients under all specifications, and the same is true for Kelly skewness[15].

As our empirical results have shown, while the dispersion of price changes falls as inflation rises, there is no evidence for such a relationship with the skewness of price. This goes against the predictions of a broad class of menu cost models and calls into question the importance of the selection effect that these models emphasize. We will show in the next section how weakening the selection effect can reconcile menu cost models with the data, while also raising the implied degree of monetary non-neutrality. However, we first compare our approach to that of Alvarez et al. (2016b), who show that in a broad class of sticky price models non-neutrality can be measured by the kurtosis of price changes.

The theoretical result of Alvarez et al. (2016b) is important, because it shows how monetary non-neutrality can be related to observeable quantities. However, an important advantage of our approach lies in its practical implementation of the empirical estimation. The key statistic in Alvarez et al. (2016b) is the kurtosis of price changes, and the result on the degree of monetary non-neutrality will be only as reliable as the estimate of price change kurtosis is precise. Obtaining precise estimates for the kurtosis is difficult for a few reasons, not least of which is the simple fact that moment estimates are less precise in finite samples the higher is the order of the moment. Perhaps more important in practice, however, is the fact that estimates of kurtosis are potentially very sensitive to just how they are computed. For example, Alvarez et al. (2016b) point out that in order to compute the kurtosis from a combined sample of potentially heterogeneous disributions, it is necessary to normalize each price change observation in order to obtain an unbiased estimate. As we show in **?**, the kurtosis of price change estimated from the U.S. CPI data varies considerably depending on just how this normalization is done, and over time. Natuarlly, estimates of the skewness are also subject to such concerns (although to a slightly lower degree because of the lower order of the moment). However, the advantage of our approach is that it emphasizes correctly estimating the sign of the relationship between skewness and inflation, and not the precise value of skewness. As we have shown, the various specifications that we use lead to the same results

[15] The 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.

in what matters for non-neutrality and for evaluating the models, which is that skewness does not fall with inflation. In this sense, our approach yields more consistent results when implemented in practice.

4 A Generalized Menu Cost Model

In this section, we present a menu cost model that has a similar setup as the menu cost models presented in Section 2: the demand system and technology faced by the firm are 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:

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

The difference with the Golosov and Lucas model is that here 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.[16] The 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 now.[17]

4.1 Background on Random Menu Costs

In addition to nesting the existing menu cost models considered thus far, our 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 functionof the deviation of the current price from its optimal value (p^*):

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

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

[17] The 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.

Any of the models we have considered will imply a price adjustment hazard function. 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 approaches, and we show in a separate paper (Luo and Villar (2016)) that the same data and empirical patterns can be used to estimate the price adjustment hazard function.

A more structural approach to price stickiness that is also related to ours is Woodford (2009)'s model of rational inattention. 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. Although not provided, we believe that a decision-theoretic justification for this random menu cost model can be derived based on the rational inattention framework. A menu cost model with inattention as a source of randomized discrete adjustment is observationally equivalent to a random-menu-cost model (see (Woodford, 2008, 2009)).[18].

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. An extreme case is perfect price flexibility, or firms being free to change their prices every period without facing any kind of cost for doing so (this is ruled out for being inconsistent with the fact that most prices do not change in any given month). 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

[18] As Woodford (2009) also points out, the direct empirical evidence on the actual costs of price adjustement 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 explicitly modelling of the informational constraints or consumer considerations that underly it to future research.

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 set the degree of exogeneity.

The distribution of menu costs will need two important features: first, a positive probability of the menu cost being zero (of a free price change), which eliminates the 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 Midrigran model already includes this, and 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. Indeed, 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:

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

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 20 in appendix D shows how the shape of the cumulative distribution function changes with α.

For our purposes, what is important is that for low values of α, the probability of very low menu costs is relatively high, but the probability of very high menu costs is also quite high. When α is high, these extreme probabilities are low, and as α rises, the density 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 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.

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

Second, and in line with the focus of our paper, we will target the signs of the correlations between inflation and the different price change moments. As previous studies had shown (and we confirmed), the correlation between inflation and the frequency of price change is positive, so our model also matches this fact. In addition, our model will imply a strongly negative correlation between inflation and the dispersion of price changes (as menu cost models do). The novelty will be that the implied correlation between inflation and the skewness of price changes will be non-negative, as in the data. [19]

Table 3 presents the parameter values that we choose to match these moments, and Table 4 shows the moments attained by the model, compared to their empirical values. The first two moments are matched almost exactly. For the empirical value of the correlations (illustrated by the scatter plots in Figure 6, 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.

Table 3: Parameter values

Parameter	Description	Value
λ	Inv. average menu cost	0.177
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	0.0967

Table 4: Simulation results

Moment	Model	Data
Avg. Frequency	10.7%	10.7%
Avg. Size	7.6%	7.5%
Corr(IQR,π)	-0.67	-0.70
Corr(Skew,π)	0.19	0.39
Corr(Freq,π)	0.67	0.63

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%)[20], the correlation also appears to be close to zero in the data. We carry out the same "long-run" analysis as in Figure 4: solving the model for different values of trend inflation. We find that for higher steady-state inflation, the average level of skewness in the price change distribution rises, and the correlation between period-by-period price change

[19] In this exercise, we do not target the relationship between the average size of price changes and inflation, which is an important part of the analysis in Nakamura et al. (2017). That paper highlights that the average size of price changes was no higher during high inflation periods, using the same data. As they emphasize, this is consistent with the menu cost model, and not with the Calvo model. Nonetheless, in Luo and Villar (2016) we estimate a hazard function of price adjustment, and find that a hazard function with both Calvo and menu cost features can match the signs of all the correlations, including the average size of price changes being constant with respect to inflation. Given that the hazard function approach and the random menu cost approach are closely related as discussed in Luo and Villar (2016), we believe that the random menu cost model presented in this paper can capture the zero inflation-size correlation if that moment is targeted. However, without targeting a zero inflation-size correlation, the random menu cost model estimated in this section predicts that the average size of price changes rises somewhat with inflation, but considerably less so than the Calvo model.

[20] Inflation is less volatile and moves within a narrower range in our generalized model than in the other menu cost models, even though the parameters of the nominal aggregate demand process are the same. This is a direct result of the differences in monetary 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.

Figure 6: Scatter plots, Generalized MC Model

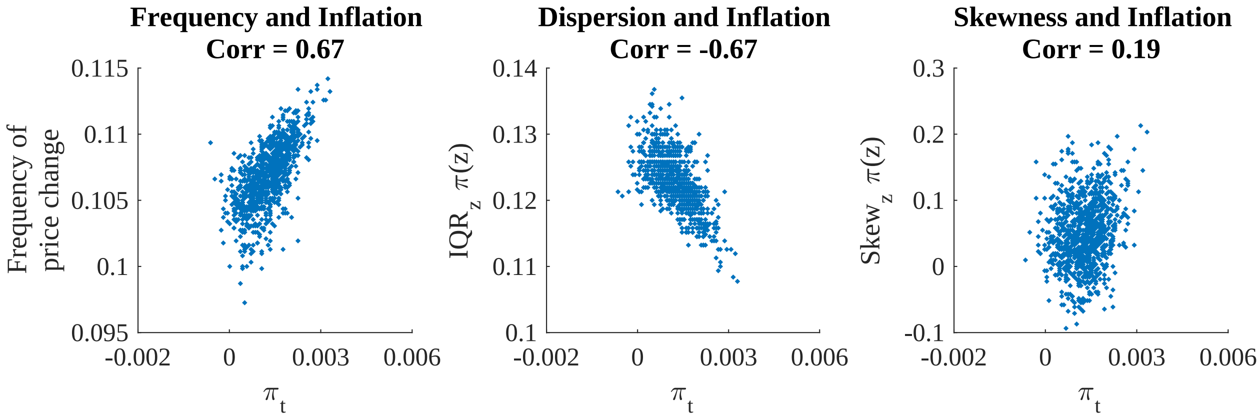

skewness and inflation (the same correlations we have been focusing on so far) also rises. This result makes our model even more consistent with the data, as it shows that when steady-state inflation is higher (as it surely was in the early, high-inflation part of our sample), we should expect to see the skewness rising with inflation. 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 7 below shows this clearly by plotting the steady-state skewness correlations for the Midrigan model (as an example) and our heterogeneous menu cost model separately.

Figure 7: Steady-State Skewness Correlation

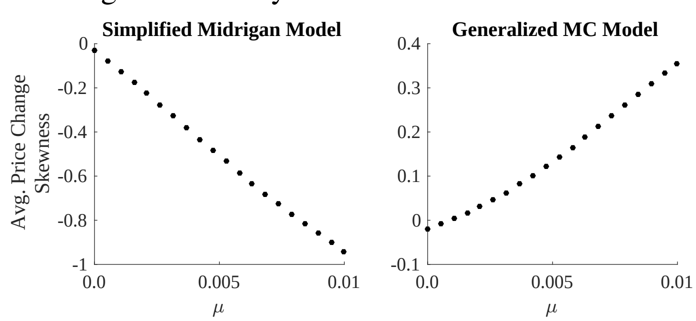

This pattern highlights how trend inflation plays an important role behind our model's non-negative skewness correlation. Indeed, positive trend inflation leads firms to expect positive future inflation when considering whether to re-set their prices. This will lead them to be less likely to cut their 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 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 counter-factual 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 costs random in the way that we have proposed weakens the selection effect that is at work in menu cost models, so it is to be expected that this model would imply a greater degree of monetary non-neutrality. Table 5 below provides a quantitative illustration of this.

Table 5: Monetary Non-Neutrality

Model	$\text{Var}(C_t) * 10^4$
Golosov and Lucas	0.05704
Midrigan	0.17718
Generalized Menu Cost	0.35094
Calvo	0.52517

Figure 8: Impulse Responses in Models

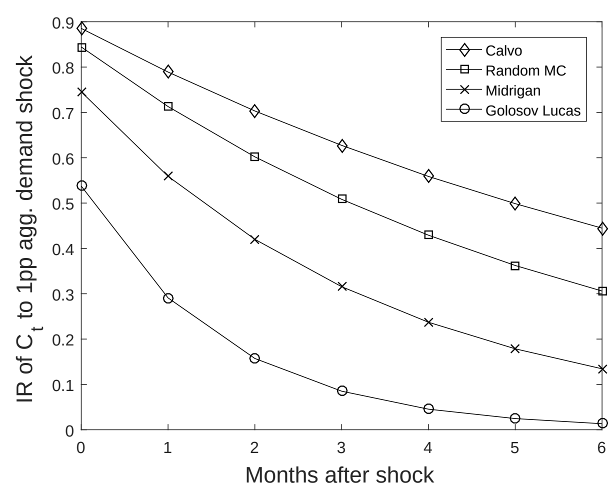

As Golosov and Lucas (2007) had famously shown, their model features a trivial amount of monetary non-neutrality compared to the Calvo model. Between the menu cost models, the major difference is between the baseline (Golosov and Lucas) and the others. Allowing for small price changes, as the Midrigan 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, 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 8 we plot the impulse response of real aggregate consumption to a one percentage point increase in nominal aggregate demand in the same four models.

The effect on real activity is not only large, but also quite persistent in our model, and much more so than in the menu cost models. In this sense, our model is also much closer to the Calvo model.

5 Conclusion

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

In the context of the debate between time-dependent and state-dependent pricing models, we follow Woodford (2009) in presenting the distinction between time- and state-dependent models as a continuum, or spectrum. Woodford (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, our contribution is to determine what point on the spectrum is most consistent with the data. Future research could combine these two approaches to gain a better understanding into the nature and importance of the informational constraints that underly price rigidity.

References

Alvarez, Fernando, Martin Beraja, Martin Gonzalez-Rozada, and Pablo Adres Neumeyer (2016a), "From hyperinflation to stable prices: Argentina's evidence on menu cost models." Working Paper.

Alvarez, Fernando, Herve Le Bihan, and Francesco Lippi (2016b), "The real effects of monetary shocks in sticky price models: A sufficient statistics approach." *American Economic Review*, 106, 2817–2851.

Alvarez, Fernando, Francesco Lippi, and Luigi Paciello (2011), "Optimal price setting with observation and menu costs." *Quarterly Journal of Economics*, 126, 1909–1960.

Anderson, Eric T., Benjamin A. Malin, Emi Nakamura, Duncan I. Simester, and Jn Steinsson (2015), "Informational rigidities and the stickiness of temporary sales." Working Paper.

Anderson, Eric T. and Duncan I. Simester (2010), "Price stickiness and customer antagonism." *Quarterly Journal of Economics*, 125, 729–765.

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. (2016), "Optimal inflation target in an economy with menu costs and zero lower bound." Working Paper.

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, 655–690.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo (2011), "Reference prices, costs, and nominal rigidities." *American Economic Review*, 101, 234–262.

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, 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, 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 (2016), “The state-dependent price adjustment hazard function: Evidence from high inflation environments.” Working Paper.

Midrigan, Virgiliu (2011), “Menu costs, multi-product firms and aggregate fluctuations.” *Econometrica*, 79, 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.

Nakamura, Emi, Jón Steinsson, Patrick Sun, and Daniel Villar (2017), “The elusive costs of inlfation: Price dispersion during the u.s. great inflation.” Working Paper.

Sims, Chrisotpher 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 (2013), “Inflation dynamics and time-varying volatility: New evidence and an ss interpretation.” *Quarterly Journal of Economics*, 129, 215–58.

Woodford, Michael (2008), “Inattention as a source of randomized discrete adjustment.” Working Paper.

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.

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

Shaowen Luo
Virginia Tech

Daniel Villar
Federal Reserve Board of Governors

A Analytical Study

We focus on the distribution of non-zero price changes. The average of this distribution is:

$$\mu = \frac{\lambda^+\Delta^+ - \lambda^-\Delta^-}{\lambda^+ + \lambda^-}$$

To evaluate the moments of the price change distribution, we consider the normalized price change, which can take two discrete values:

$$\Delta p - \mu = \begin{cases} \Delta^+ - \frac{\lambda^+\Delta^+ - \lambda^-\Delta^-}{\lambda^+ + \lambda^-} & , \text{ with probability } \frac{\lambda^+}{\lambda^+ + \lambda^-} \\ -\Delta^- - \frac{\lambda^+\Delta^+ - \lambda^-\Delta^-}{\lambda^+ + \lambda^-} & , \text{ with probability } \frac{\lambda^-}{\lambda^+ + \lambda^-} \end{cases}$$

$$= \begin{cases} \frac{\lambda^-(\Delta^+ + \Delta^-)}{\lambda^+ + \lambda^-} & , \text{ with probability } \frac{\lambda^+}{\lambda^+ + \lambda^-} \\ \frac{-\lambda^+(\Delta^+ + \Delta^-)}{\lambda^+ + \lambda^-} & , \text{ with probability } \frac{\lambda^-}{\lambda^+ + \lambda^-} \end{cases}$$

We define the share of price changes that are increases as $\beta \equiv \frac{\lambda^+}{\lambda^+ + \lambda^-}$. The variance of price changes is then:

$$\begin{aligned} E(\Delta p - \mu)^2 &= \beta(\frac{\lambda^{-2}(\Delta^+ + \Delta^-)^2}{(\lambda^+ + \lambda^-)^2}) + (1-\beta)(\frac{\lambda^{+2}(\Delta^+ + \Delta^-)^2}{(\lambda^+ + \lambda^-)^2}) \\ &= (\beta(1-\beta)^2 + (1-\beta)\beta^2)(\Delta^+ + \Delta^-)^2 \\ &= \beta(1-\beta)(\Delta^+ + \Delta^-)^2 \end{aligned}$$

To evaluate the relationship of this moment with inflaiton, we use the fact, shown by Alvarez et al. (2016a), that the price change thresholds chosen by firms (which determine Δ^+ and Δ^-) do not vary with inflation for low and intermediate values of inflation (the range that we are considering in our analysis). Changes in inflation will therefore result in changes in the fraction of price increases (β): rising inflation will be associated with rising β . The variance of price changes will rise with β (and with inflation) when

$\beta < 0.5$ (generally, negative inflation periods), and will fall with inflation when $\beta \geq 0.5$. This is also what our model simulations show.

In order to derive an expression for price change skewness, we evaluate:

$$\begin{aligned} E(\Delta p - \mu)^3 &= \beta(\frac{\lambda^{-3}(\Delta^+ + \Delta^-)^3}{(\lambda^+ + \lambda^-)^3}) + (1-\beta)(-\frac{\lambda^{+3}(\Delta^+ + \Delta^-)^3}{(\lambda^+ + \lambda^-)^3}) \\ &= (\beta(1-\beta)^3 - (1-\beta)\beta^3)(\Delta^+ + \Delta^-)^3 \\ &= \beta(1-2\beta)(1-\beta)(\Delta^+ + \Delta^-)^3 \end{aligned}$$

The Skewness is then:

$$\text{Skewness} = \frac{E(\Delta p - \mu)^3}{[E(\Delta p - \mu)^2]^{\frac{3}{2}}} = \frac{\beta(1-2\beta)(1-\beta)}{\beta^{\frac{3}{2}}(1-\beta)^{\frac{3}{2}}} = \frac{1-2\beta}{\beta^{\frac{1}{2}}(1-\beta)^{\frac{1}{2}}}$$

It can be seen that the skewness will be negative when $\beta > \frac{1}{2}$ (when inflation is positive), and positive when $\beta < \frac{1}{2}$ (when inflation is negative). Furthermore, the derivative of skewness with respect to β can be shown to be negative:

$$\frac{\partial \text{Skewness}}{\partial \beta} = \frac{-2 - \frac{(1-2\beta)^2}{2\beta(1-\beta)}}{\beta^{\frac{1}{2}}(1-\beta)^{\frac{1}{2}}} < 0$$

Due to the relationship between inflation and the fraction of price increases, this also means that the skewness of price change and inflation are negatively related, as shown by our simulation results.

Finally, the Kurtosis of price changes is:

$$\begin{aligned} \text{Kurtosis} &= \frac{E(x-\mu)^4}{(E(x-\mu)^2)^2} = \frac{\beta(1-\beta)[(1-\beta)^3 - \beta^3]}{\beta^2(1-\beta)^2} \\ &= \frac{(1-2\beta)^2}{\beta(1-\beta)} + \frac{\beta(1+\beta)(1-2\beta)}{\beta(1-\beta)} \\ &= \text{Skewness}^2 + \frac{\beta(1+\beta)}{\beta^{1/2}(1-\beta)^{1/2}}\text{Skewness} \end{aligned}$$

B Computational Procedure and Calibration of Sticky Price Models

We solve the standard Calvo model, the Golosov and Lucas (2007) model, the Nakamura and Steinsson (2010) model, and the Midrigan (2011) model mentioned above 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

Figure 9: Intuition for the Menu Cost Model: correlations

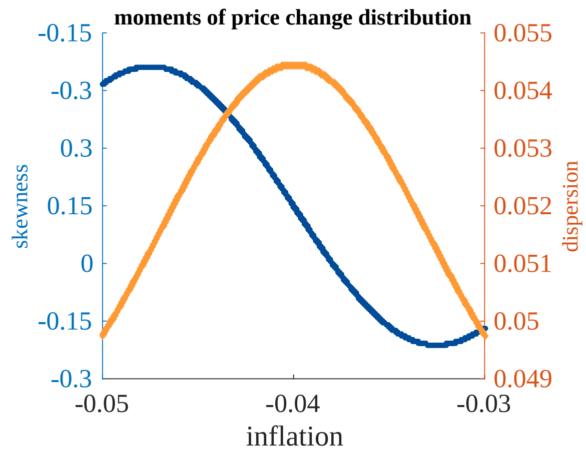

Note: This figure plots skewness (blue curve) and dispersion (orange curve) of the realized price change distribution as a function of the level of inflation. Desired price changes follow $\mathcal{N}(\mu, 0.05^2)$, "S" band at 0.01, "s" band at -0.01, while varying μ.

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:

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

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

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

where $V(\cdot)$ is firm z's value function, $\Pi_t^R(z)$[21] 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[22], by value function iteration, using the proposed $\Gamma(\cdot)$ function, the stochastic processes for the exogenous variables (applied

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

[22] Because 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.

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 simulate Alvarez et al. (2011) and Woodford (2009)) using the replication files provided by the authors. We use the same methods and parameter values used in the original papers (Alvarez et al. (2011) for the observation costs model; Woodford (2009) for the rational inattention model), and use the policy functions to simulate the models. The kind of analysis that we can carry out with these models is more restricted than for the perfect information models. Indeed, the Alvarez et al. (2011) model contains no aggregate shocks (which in the other models drive period-by-period inflation movements). Therefore, we exclude this model from the "short-run" analaysis, in which the trend inflation parameter is fixed but there is no aggregate disturbance in the model. Instead, we conduct the "long-run" analysis by varying the parameters of trend inflation (from $\mu = 0$ to $\mu = 0.2$).[23] For each level of trend inflation, we compute the average dispersion and skewness of price change and plot them against the level of trend inflation. Finally, the Woodford (2009) model contains no trend inflation. Strictly speaking, the model cannot be solved with aggregate nominal disturbances. Nonetheless, we take a simplified approach following Section 5 of Woodford (2009): an aggregate nominal shock, which shifts the desired price of firms by the same amount, would affect each firm's price-review decision the same way as in the stationary equilibrium with only idiosyncratic shocks.[24] Therefore, we take the hazard function for the case $\theta = 5$ (unit information cost) as given and simulate the dynamics of price change for 1,000 periods and 40,000 firms with the dynamics of aggregate nominal expenditure being i.i.d. and mean zero. We use the simulated data to conduct the "short-run" analysis. The "long-run" analysis of this model is excluded.

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 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 10.7%, and the average size of price change is 7.5%. For the Midrigan model (as well as our random menu cost model), we also target the fraction of price changes

[23]The range of the trend inflation is much wider in this "long-run" study (from 0 to 0.2) than in the study of the other models (from 0 to 0.01), because the Alvarez et al. (2011) model is less sensitive to the level of trend inflation than the other models.

[24]Woodford (2009) uses this simplified approach to study the monetary non-neutrality of the model.

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 13.2%. Table 6 below shows the model-implied moments for the Golosov and Lucas, Midrigan, and Calvo models, as well as the random menu cost model from section 4, and compares them to their empirical values:

Table 6: Model implied moments

Model	Average Frequency (%)	Average Size (%)	Fraction Small (%)
Golosov and Lucas	10.7	7.6	0
Midrigan	10.6	7.6	12.3
Calvo	10.7	7.6	8.3
Random MC	10.7	7.6	12.6
Data	10.7	7.5	13.2

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 7 below shows the parameter values that we choose for these models.

Table 7: Parameter values for models

Parameter	Golosov and Lucas	Value
χ	Menu cost (as share of steady state revenue)	0.0178
σ_ϵ	Std. dev. of idiosyncratic tech. shocks	0.038
	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	0.107
σ_ϵ	Std. dev. of idiosyncratic tech. shocks	0.194

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 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. 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 Appendix A.2., 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 10 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.

Figure 10: Number of observations by month

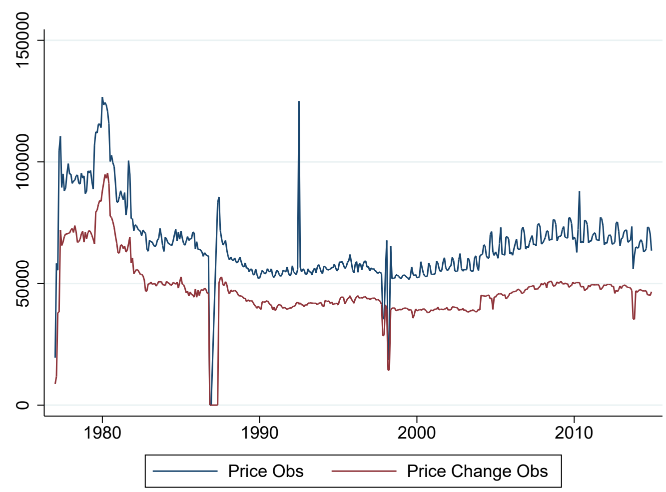

We provide here an explanation for the restrictions that we make on our sample of price changes. The empirical literature on price setting has emphasized the importance of identifying "pure" or regular price changes, as opposed to price changes coming from temporary sales or substitutions. The reason 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. The BLS also identifies whenever a product substitution occurs, or when a new "version" of a particular product is introduced. We treat a new version as an entirely new product, and only compute price changes by comparing price changes within identical versions.

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. We compute price changes as the difference of the log price, or:

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

As discussed previously, 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). 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 Kryvtsov (2008); Alvarez et al. (2016b)), (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. We label estimates constructed with this restriction wtih ”EJRS”.

For the dispersion and skewness statistics, we first separate observations into categories that we label major groups. There are thirteen of these, and table 8 below provides a list, along with the share of expenditure weight that they represent.

Table 8: CPI group weight

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.

Figure 11: IQR and Skewness of Price Change Distribution, Quarterly
Source: Authors' calculations from BLS CPI Research Database

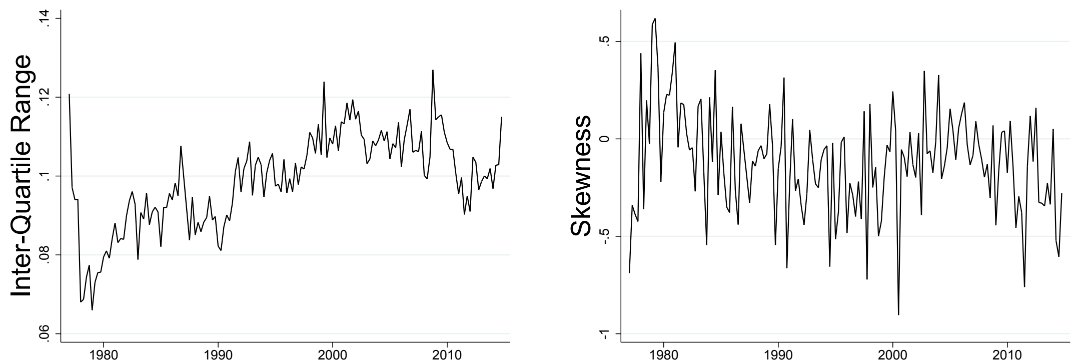

D Additional Empirical Results

In Section 2, we presented results on the empirical result between inflation and various price change moments, using both scatter plots and regressions. We provide additional empirical results that support the main message of 2: that the dispersion of price change falls with inflation, and that price change skewness does not. We start with the correlation values between inflation and the different moments, at various frequencies, and for excluding and including the high inflation period.

Table 9: 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	-	

Next, we present scatter plots in which the dispersion and skewness measures were computed by excluding small price changes in the ELI's pointed out by Eichenbaum et al. (2013).

The measure of inflation that we had used in the scatter plots and regressions 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 12 below, we plot the twelve month log change for both indexes. They both

Table 10: 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 11: 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 12: 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	-	

co-move very strongly, although the peak is much higher for the CPI.

Figure 12: Inflation

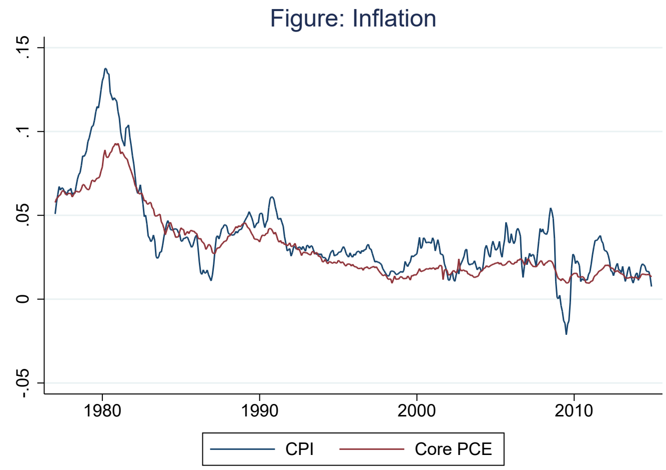

In this section we show that our results do not depend on which inflation measure we use, so we present scatter plots with CPI inflation, and regression results with Core PCE inflation as the regressor. The only difference that this makes is that in the regressions, the absolute value of the coefficients 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 13-16 below present scatter plots of the smoothed moment and inflation series.

Figures 17-20 are scatter plots using CPI inflation.

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.

Figure 13: Frequency of price change & inflation smoothed, quarterly

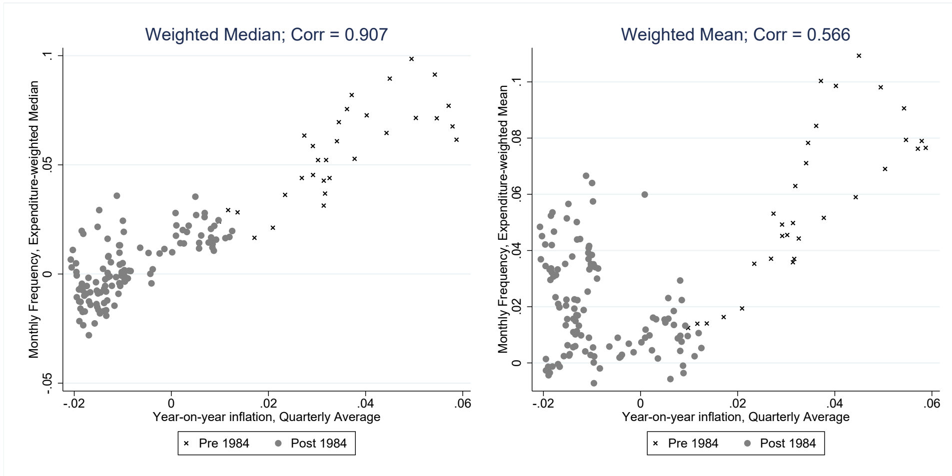

Figure 14: IQR of price change & inflation smoothed, quarterly

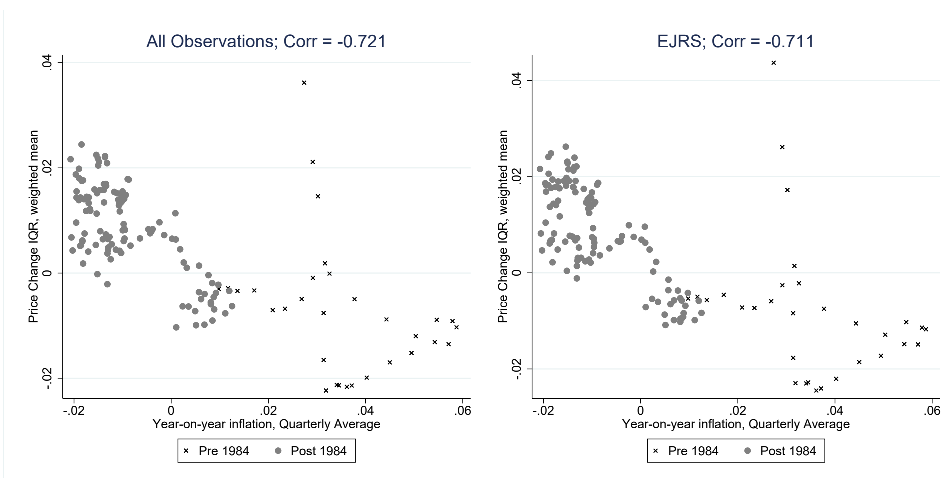

Table 13: Core inflation as regressor - frequency

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

Figure 15: Skewness & inflation smoothed, quarterly

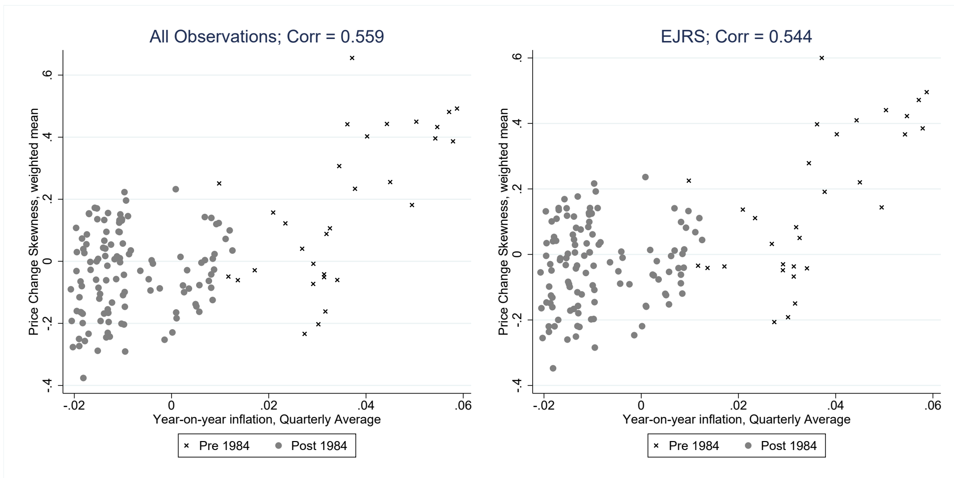

Figure 16: Kelly skewness & inflation smoothed, quarterly, corr=0.734

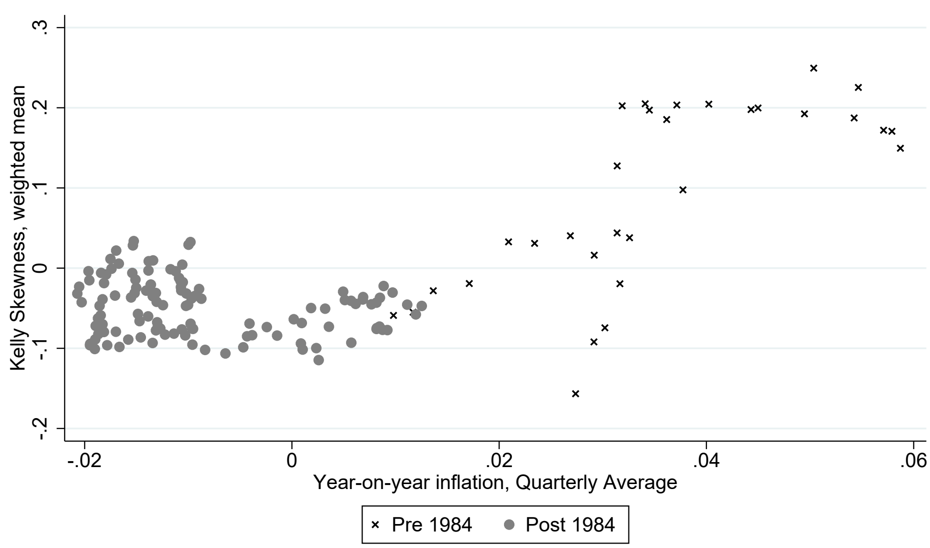

Figure 17: Frequency of price change & CPI inflation, quarterly

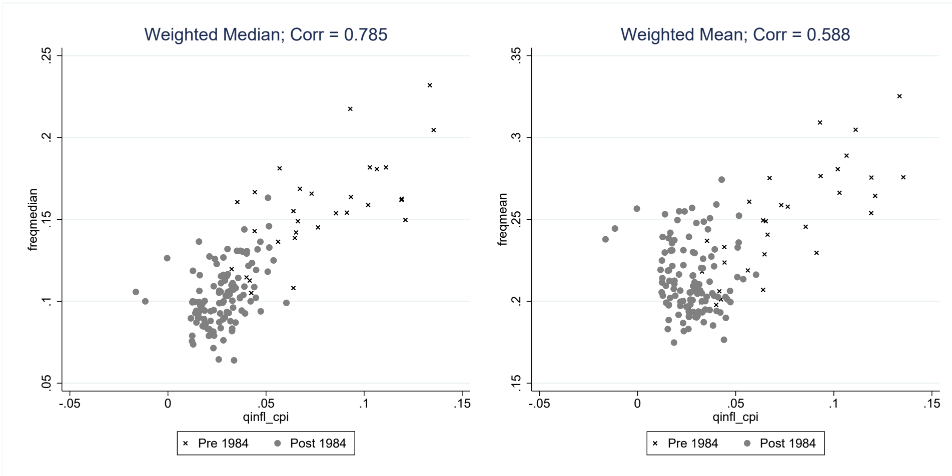

Figure 18: IQR & CPI inflation, quarterly

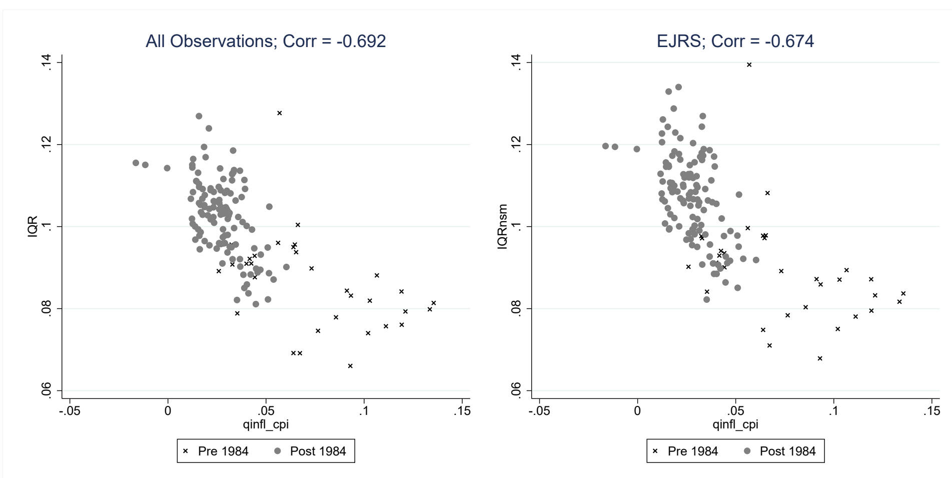

Table 14: Smoothed and seasonal adjusted series - frequency

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

Figure 19: Skewness & CPI inflation, quarterly

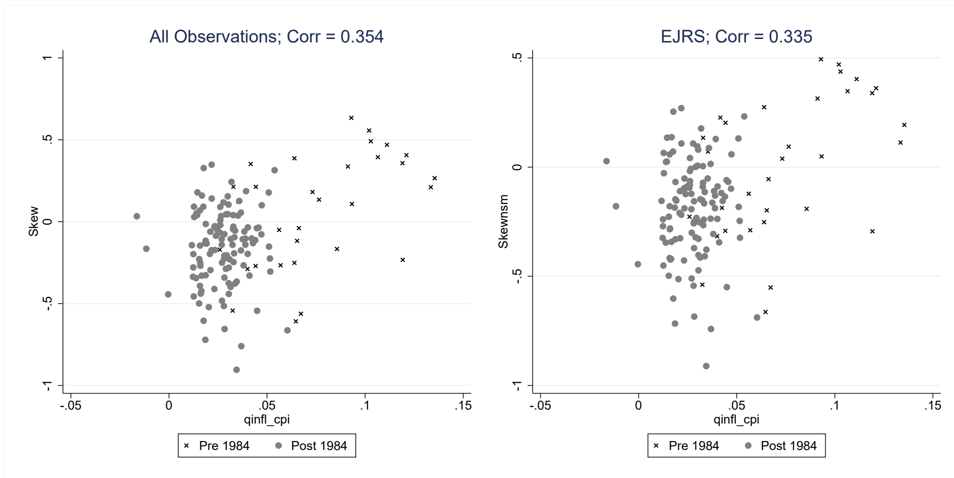

Figure 20: Kelly skewness & CPI inflation, quarterly, corr=0.674

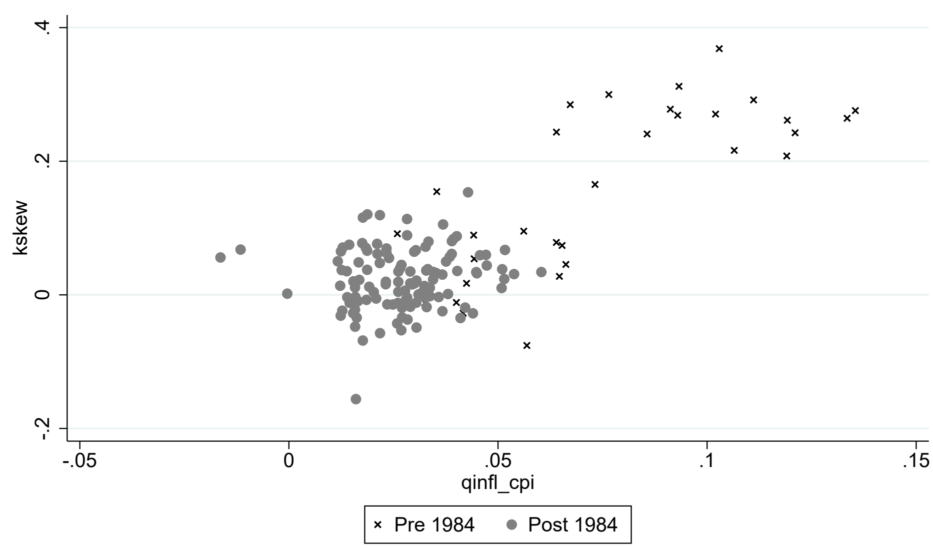

Table 15: Core inflation as regressor - IQR

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

Table 16: Smoothed and seasonal adjusted series - IQR

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

Table 17: Core inflation as regressor - skewness

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

Table 18: Smoothed and seasonal adjusted series - skewness

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

Table 19: Core inflation as regressor - Kelly skewness

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

Table 20: Smoothed and seasonal adjusted series - Kelly skewness

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

Figure 21: Moments of Price Change and Inflation, Quarterly

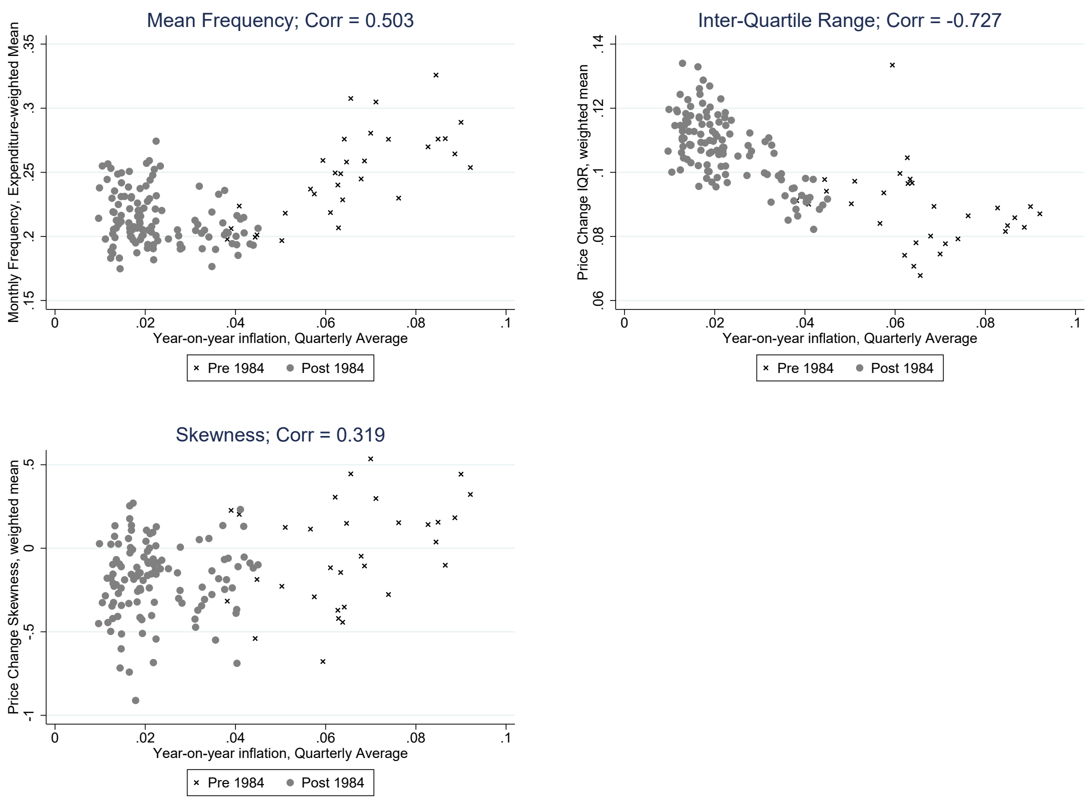

Table 21: Coefficients on Inflation for Price Change Moments - Using CPI Data Excluding Small Price Changes

	1977-2014			1985-2014		
	All	Fed Dummies	Inflation Only	All	Fed Dummies	Inflation Only
Frequency	0.164	0.686***	0.438***	0.018	0.339**	-0.087
	(0.203)	(0.104)	(0.108)	(0.196)	(0.167)	(0.236)
IQR	-0.327***	-0.204***	-0.261***	-0.491***	-0.476***	-0.224***
	(0.046)	(0.044)	(0.095)	(0.082)	(0.089)	(0.092)
Skewness	3.501***	3.928***	1.947	1.108	1.130	2.963
	(0.828)	(0.966)	(2.538)	(1.534)	(1.705)	(2.985)

Note: Significant *** at 1% level (** at 5% level; * at 10% level). This table reports the regression coefficients on inflation from regressions of the weighted average mean frequency of price changes, as well as weighted mean price change IQR and skewness, excluding certain small price changes based on Eichenbaum et al. (2013). 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 what controls are used. exclusion of small price changes. Standard errors that are consistent for heteroskedasticity and auto-correlation of the residuals (Newey-West) are reported.

E Random Menu Cost

Figure 22: Shape of Menu Cost CDF for Different α