The Price Adjustment Hazard Function: Evidence from High Inflation Periods *

Shaowen Luo†

Department of Economics
Virginia Tech

Daniel Villar‡

Federal Reserve Board of Governors

May 21, 2020

Abstract

The price adjustment hazard function - the probability of a good's price changing as a function of its price misalignment - enables the examination of the relationship between price stickiness and monetary non-neutrality without specifying a micro-founded model, as discussed by Caballero and Engel (1993a, 2007). Using the micro data underlying the U.S. Consumer Price Index going back to the 1970s, we estimate the hazard function relying on empirical patterns from high and low inflation periods. Our estimated hazard function features three properties: 1) the probability of adjustment rises only slowly in the absolute value of the price misalignment; 2) the value at zero (the probability of infinitesimal price changes occurring) must be significantly positive; 3) the hazard function is asymmetric in that price increases are significantly more likely to occur than price decreases (holding the absolute value fixed). Our estimated hazard function therefore implies a high degree of monetary non-neutrality. Finally, estimating hazard functions at the sectoral level, we find substantial heterogeneity in the aggregate flexibility across sectors.

JEL classification codes: E31, E32, E47, E52

1 Introduction

It has long been established that monetary policy has real effects (or is non-neutral) if firms face constraints in changing the prices of their goods and services (i.e. if prices are "sticky"). This has led to a proliferation of micro-founded models of price-setting that feature various nominal

*We would like to thank Emi Nakamura, Jón Steinsson, Ricardo Reis, Michael Woodford, Serena Ng, and seminar participants at the Federal Reserve Board 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

rigidities (Caplin and Spulber, 1987; Golosov and Lucas, 2007; Midrigan, 2011). These studies have notably shown that a key determinant of monetary non-neutrality is the extent to which the prices that change are selected based on their misalignment, the selection effect. While most work in this field has involved micro-founded models, Caballero and Engel (1993a) proposed an alternative, model-free approach to analyze the relation between sticky prices and non-neutrality. They introduced the state-dependent price adjustment hazard function: the probability of an individual price changing, as a function of the gap between a firm's current and desired price. This model-free approach also provides a way to determine monetary non-neutrality through the mapping between the hazard function and the flexibility of the price level (as derived in Caballero and Engel, 2007). In particular, every price setting model implies a certain hazard function, thus this general approach does not require formulating and solving a specific model.

In this paper, we contribute to the debate on non-neutrality by estimating the price adjustment hazard function using the micro data that underlies the U.S. Consumer Price Index (CPI) from 1977 to 2014. While other studies have offered estimates of the hazard function (such as Berger and Vavra, 2018; Caballero and Engel, 1993b, 2006; Gagnon et al., 2012), our estimation is based on a richer data set covering high inflation periods and novel empirical moments, notably the correlation between inflation and the skewness of the price change distribution. The most important property of our estimated hazard function is the small increase of the adjustment probability with the size of the price misalignment from the desired level. In other words, we find that the slope of the hazard function is small. A we will show, using the correlation between inflation and price change skewness is what enables us to determine that the slope of the hazard function is small. This property also implies a high degree of monetary non-neutrality.

The slope of the hazard function is particularly important because of its close relation to the selection effect in price setting. Caballero and Engel (2007), in deriving the mapping between the hazard function and aggregate flexibility, made explicit how the selection effect depends critically on the slope of the hazard function. A large slope means that more misaligned prices have a much higher probability of changing. Consider the standard menu cost model as an extreme state-dependent example: firms decide to change their price if and only if the price misalignment is large enough to justify the fixed adjustment cost. The hazard is zero within the inaction region, and one elsewhere, which implies an infinite slope at the inaction thresholds. Consequently, the model implies a large aggregate price response and small real effects of monetary shocks, given that adjusting firms are the ones that react strongly to nominal shocks. Conversely, a model (such as a Calvo model) with a small hazard function slope generates weak aggregate price responses to nominal shocks, thus high monetary non-neutrality.

In order to infer the slope of the hazard function, we use patterns related to the skewness of price change that are not considered by previous attempts to estimate the hazard function. If the hazard function slope is large, the resulting price dynamics will imply that the skewness of price changes falls with inflation. Price change skewness does not fall with inflation in our data, which leads us to conclude that the hazard function must have a small slope (and that non-neutrality is high). In order to determine the empirical relation between inflation and skewness, it is necessary to observe the skewness in periods of intermediate or high inflation. Most price data sets used in the literature only cover recent periods in which inflation has been low and stable (such as the U.S. CPI data from 1988 onwards, Dominicks and the Nielsen Homescan Dataset), making them unsuitable for our approach. Instead we use the U.S. CPI micro data extended back to 1977, which covers periods of high and intermediate inflation.

The skewness patterns that help identify the hazard function slope has been discussed throughly in Luo and Villar (2017), where we showed that a broad class of menu cost models predict that skewness falls with inflation. Only models that are weakly state-dependent could match the non-negative inflation-skewness correlation, and such models predict a high degree of non-neutrality. While we conclude that non-neutrality is high in both papers, the results in this paper are, in a sense more general. In Luo and Villar (2017), we present a random menu cost model characterized by firm level random menu cost draws from a calibrated distribution in every period. Instead, this paper uses a hazard function approach, which is model-free. Our results here reject hazard functions with a large slope, and thus all models that would imply such a hazard function.

Our estimation establishes three additional features of the hazard function that offer insight into the price setting process. First, we find that the hazard function is asymmetric around zero. Specifically, for an equivalent magnitude, price increases are more likely to occur than price decreases. Second, the function takes a significantly positive value at zero. Thus, even extremely small price imbalances have a positive probability of occurring (so there are no inaction regions). Third, the hazard function makes it clear that the probability of price adjustment stays far below one even for large price misalignments. This is related to the small slope feature of the hazard function and is contrary to what would be obtained from menu cost, or hybrid menu cost models, such as Midrigan (2011). These findings are relevant for attempts to model the constraints faced by firms in setting and changing their prices.

As mentioned before, an advantage of using the hazard function to evaluate non-neutrality is that doing so does not require specifying the price setting constraints faced by firms, nor does it require solving for firms' optimal decisions given those constraints. Once the empirical values required to estimate the function are known, it is a relatively simple procedure to carry out. One way that we exploit the simplicity of this procedure by implement it for individual sectors. (Nakamura and Steinsson, 2008, 2010) showed that the frequency of price change differs markedly across sectors, and it would not be surprising if other aspects of the price setting process (captured by the shape of the hazard function) also differed across sectors. We estimate sector-specific hazard functions by targeting the values of price change moments by sector. According to our results, different sectors have considerably different degrees of state dependence in pricing, and of aggregate flexibility.

The fact that the hazard function involves an unobserved variable (the desired price gap) makes it particularly difficult to estimate. An estimation strategy must make assumptions about the desired price gap to relate different hazard functions to observable facts about prices and price changes. Two recent papers (Berger and Vavra, 2018; Petrella et al., 2018) also attempt to estimate the price adjustment hazard function with assumptions that differ from those in our paper. One difference is that both of those papers estimate a flexible functional form for the distribution of desired price gaps and a quadratic hazard function. In contrast, we estimate a more restricted form for the distribution of desired price gaps, but a more flexible form for the hazard function. Perhaps more importantly, we use a broader set of empirical moments in our estimation. Our approach has the advantage of using information from the variation of price change moments over time (mainly in the form of the skewness correlation previously discussed) to identify the key parameters of the hazard function. As we will show, these moments provide a considerable amount of information that identifies the slope of the hazard function.[1] In addition, our data sample includes the high inflation periods of the U.S. going back to 1977.

This paper (as well as Berger and Vavra, 2018; Petrella et al., 2018) contribute the sticky price literature with non-structual estimates of price changes following Caballero and Engel (1993a, 2007) using high order moments of price changes. An important finding in our work is that an asymmetric hazard function fits the data best, which implies a stronger monetary non-neutrality during deflationary episodes. Berger and Vavra (2018), in contrast, estimate a symmetric quadratic hazard function that can not generate this result. Petrella et al. (2018) estimate a hazard function that is potentially asymmetric, but they do not seem to find a major role for asymmetry.

Before continuing our discussion of the price adjustment hazard function, it is important to draw a distinction with another hazard function that has received attention in the sticky price literature: the time-dependent hazard function. The time-dependent hazard function gives the probability of a price change occurring at time t, given that t periods have passed since the last price adjustment. Different models also imply different types of time-dependent hazard functions, and some work has been done to estimate these using micro price data (such as Klenow and Kryvtsov, 2008; Nakamura and Steinsson, 2008). However, in this paper we only investigate the state-dependent price adjustment hazard function.

[1] Another significant difference is that both Berger and Vavra (2018) and Petrella et al. (2018) estimate the parameters of the hazard function for every period of time (which gives them, among other things, a time-varying for the degree of monetary non-neutrality), while we estimate a hazard function that is static over time. Note that estimating the hazard function period by period makes it impossible to use the correlations of inflation with price change moments over time, and thus discarding the valuable information those correlations provide.

The rest of this paper is organized as follows. Section 2 formalizes the hazard function approach, illustrates it with various examples, and describes our estimating method. Section 3 presents our main results: the general hazard function estimates (parametric and non-parametric), and an illustration of how different key features of the hazard function are identified by our set of moments. In Section 4, we derive results on monetary non-neutrality from the estimated hazard function. Section 5 contains the estimates for sector-specific hazard functions. Finally, we provide some concluding remarks in Section 6.

2 Price Adjustment Hazard Function and Monetary Non-neutrality

The basic framework of the model is that firms adjust their prices infrequent and the adjustment probability is state dependent.

In this section, we propose a simple price adjustment framework similar to that in Caballero and Engel (2007), in which the price adjustment policy of a firm depends on a state-dependent continuous hazard function. We do not explicitly model the underlying microeconomic optimization problem of firms or the general equilibrium aspects. Despite the absence of micro-foundations, this approach can be viewed as an empirical-theoretical strategy to shed some light on and complement the structural sticky price models as Caballero and Engel (1993a) had suggested.

In this economy, a continuum of firms indexed by $i \in [0, 1]$ face idiosyncratic and aggregate shocks. Let p_{it} denote the log of the price for firm i in period t, and p^*_{it} denote the log of the target price. The target is the sum of two components, which changes stochastically as follows,

$$p^*_{it} = z_{it} + m_t,$$

where z_{it} is the idiosyncratic component and m_t is the aggregate component, and both are subject to shocks over time. The idiosyncratic component is important in order to match the fact that, even within a given period, there is a wide variation of price changes (as shown notably by Bils and Klenow (2004), Nakamura and Steinsson (2008), and Klenow and Kryvtsov (2008)). This component follows an AR(1) process, as the idiosyncratic productivity shocks in menu cost models like Golosov and Lucas (2007):

$$z_{it} = \rho z_{i,t-1} + \epsilon_{it},\ \epsilon_{it} \sim N(0, \sigma_\epsilon).$$

The aggregate component of shocks follows a random walk with drift:

$$m_t = \mu + m_{t-1} + \eta_t,\ \eta_t \sim N(0, \sigma_\eta),$$

where the drift parameter (μ) accounts for the fact that inflation is positive, on average (as over the long run, price changes will average μ). The aggregate shocks lead to variations in inflation period by period, and represent monetary, or aggregate demand, shocks.

Firms adjust the prices of their products infrequently due to adjustment costs (e.g. menu cost or information cost, which are not modeled here) and follows an adjustment hazard function. Furthermore, when a price adjusts, it is set to its current optimal level (p^*_{it}), such that the desired price gap (or price imbalance) is closed. Thus, the hazard function gives the probability of a firm's price adjusting given the difference between the current price and the optimal price, i.e.

$$H(x) = P(\Delta p_{it} = x | x \equiv p^*_{it} - p_{i,t-1}).$$

This set-up is extremely general, and nests existing sticky price models. For instance, in the frictionless case, it would be optimal for firms to set their price equal to the target at all times (i.e. $H(x) = 1, \forall x$).

While we are not modelling firms' optimal response to these shocks, the hazard function (along with the distribution of idiosyncratic shocks) will determine how the aggregate price level responds to them. It is in this way that the hazard function allows us to assess the degree of aggregate price flexibility, or the inverse of monetary non-neutrality. In what follows, we denote the cross-sectional density of price imbalances (x) at time t as $f_t(x)$. Although we have made assumptions about the distribution of the imbalance's components, the theoretical results apply to a general distribution. The change in the aggregate price level (or the average price change, or inflation) can therefore be expressed as:

$$\Delta p_t = \int x H(x) f_t(x) dx.$$

In particular, we are interested in how this change will depend on the aggregate shock, and so we consider the change in the price level as a function of the change in m_t. As presented by Caballero and Engel (2007), the average inflation response to an aggregate shock deviation from its average rate ($\Delta m_t = 0$) is,

$$\Delta p_t(\Delta m_t) = \int (x + \Delta m_t) H(x + \Delta m_t) f_t(x) dx.$$

Aggregate flexibility is then defined as the derivative of inflation with respect to Δm, evaluated at ($\Delta m = 0$), which is equal to:[2]

$$\Delta p'_t(\Delta m = 0) = \underbrace{\int H(x) f_t(x) dx}_{\text{frequency of price change}} + \underbrace{\int x H'(x) f_t(x) dx}_{\text{extensive margin}}. \tag{1}$$

This illustrates the importance of having a precise estimate of the hazard function, and of knowing what shape it takes, in particular. The first term in this expression is simply the frequency of price

[2]Note that the average sensitivity of the price level to the change in the aggregate component takes the same expression, with the ergodic density of price imbalances, $f_E(x)$ replacing $f_t(x)$.

Table 1: Examples From Existing Models

	$H(x)$	Extensive margin
Calvo	f	None
Menu cost	$\begin{cases} 0, & x \in (L,U) \\ 1, & \text{otherwise} \end{cases}$	$\lvert L \rvert f_t(L) + U f_t(U)$
Calvo & menu cost hybrid	$\begin{cases} p_z, & x \in (L,U) \\ 1, & \text{otherwise} \end{cases}$	$(1-p_z)[\lvert L \rvert f_t(L) + U f_t(U)]$

change in period t. Although the hazard function clearly plays a role, the frequency of price change can be observed directly from price micro data, and can therefore be assessed independently of the hazard function. However, the second term does not have a simple relation to anything measurable in the data. Caballero and Engel (2007) refer to this term as the role of the extensive margin, and show that it is typically twice as large as the frequency, in a wide variety of models associated with increasing hazard functions. This incorporates the role for the ”selection effect” induced by state-dependent pricing, which makes the price level more responsive to aggregate shocks. Since the derivative of the hazard function enters the extensive margin, its shape has an important influence on monetary non-neutrality. Also important is whether the density of price imbalances is high in the regions where the slope is large.

Table 1 illustrates the importance of these concepts with a few examples of hazard functions implied by different sticky price models. The simplest and perhaps most used sticky price model, Calvo (1983), has a constant hazard function with f representing the frequency of price change. Consequently, the derivative is zero everywhere, meaning that the contribution of the extensive margin to aggregate flexibility is zero. At the other extreme in terms of flexibility is the class of menu cost, or Ss, models (such as Golosov and Lucas (2007)). When firms have to pay a fixed cost to change their nominal price, they optimally choose a threshold rule, under which they re-set their optimal price if and only if their price mis-alignment is outside of some interval (L, U). Thus, its extensive margin term is clearly positive as shown in the second row. In addition, some advanced sticky price models are close to a hybrid of Calvo and menu cost, such as Midrigan (2011)'s menu cost model or the CalvoPlus model presented by Nakamura and Steinsson (2010). In each case, firms will occasionally and randomly receive the opportunity to change their price for free or for a relatively low menu cost. The common feature shared by these models is that firms occasionally have the chance to change their price for free or low cost with a hazard function and the extensive margin as presented in the last row.[3]

A comparison of the extensive margin generated by various models clearly infers their levels of monetary non-neutrality. A menu cost has an extreme extensive margin, because only the most mis-aligned prices adjust. Since they adjust by a large amount, the average price response to monetary shocks is very large, leading to high aggregate flexibility and monetary non-neutrality. On the other hand, the Calvo model has a much lower degree of monetary non-neutrality than the menu cost.

The models that we have presented so far yield hazard functions that are extremely simple. However, richer and more interesting hazard functions can be obtained from models of imperfect information (e.g. Alvarez et al., 2011b; Costain and Nakov, 2014; Woodford, 2009). For instance, the hazard function of the rational inattention model in Woodford (2009) change drastically for different levels of tightness of the information constraint.[4] Indeed, the model can nest the Calvo model (in the no information case) and the menu cost model (free information, or no constraint), as well as intermediate cases. This can be seen in the different hazard functions implied by different values for the cost of information (denoted by θ), shown in Appendix A Figure 6. Thus, the extensive margin effect depends on the exact shape of the hazard function.

As these examples illustrate, various sticky price models can be analyzed under the hazard function approach, which makes it possible to clearly see and understand the differences in the degree of aggregate flexibility that they imply. In the following section, we present the estimation of the hazard function, which will then be used to revisit the aggregate flexibility results of the different models.

3 Hazard Function Estimation

3.1 The set of hazard functions

In this section, we describe the estimation of the hazard function. The estimation is carried out by comparing the price change moments implied by different guesses of the hazard function (obtained from simulations) with the moments from the data, and choosing the hazard function that matches these moments best. For this to be feasible, the hazard function is characterized by a finite number

[3] In the Midrigan (2011) model, each firm produces multiple products, but once it pays the menu cost, it can change the price of all its products. Once a particular product's price is mis-aligned beyond a certain amount, the firm pays the menu cost to re-set the price. When this occurs, however, the firm's other prices will also be re-set, essentially for free, and these will not necessarily be very mis-aligned. This means that all prices will have a positive probability of adjusting, no matter how small their imbalance. In the CalvoPlus model, every period firms face either a high menu cost or a low menu cost, with a fixed probability (typically higher for the high menu cost).

[4] In this model, firms face a cost to processing relevant information (measured in terms of entropy reduction) to making their pricing decisions. A related discrete-choice model is presented in (Costain and Nakov, 2014, 2015), in which price stickiness is the result of errors due to costly decision-making.

of parameters. One way to do this is to impose a functional form on the hazard function, and to estimate the parameters that characterize it. This can be applied to the models in section 2, for example, as most of them imply hazard functions that are summarized by a small number of parameters. The resulting estimates do not recover the structural parameters of the models, but they do make it possible to assess the model's implications for various empirical patterns, and the degree of aggregate flexibility.

Our focus is on producing more flexible estimates that are not tied to a particular model. The first way in which we will carry this out will be by proposing flexible parametric forms for the hazard function, and estimating the parameters of the parametric form. One such form that we consider is an asymmetric quadratic function with discrete jumps in the probability of price change:

$$H(x)^{\text{quad}} = \begin{cases} 1, & \text{if} \quad x < c^- \\ p_0 + a^- x + b^- x^2, & \text{if} \quad c^- \leq x < 0 \\ p_0, & \text{if} \quad x = 0 \\ p_0 + a^+ x + b^+ x^2, & \text{if} \quad 0 < x \leq c^+ \\ 1, & \text{if} \quad x > c^+ \end{cases} . \tag{2}$$

The discrete jumps in probability enable this functional form to nest the menu cost and CalvoSs hazard functions. The coefficients on x determine the slope of the hazard function, and therefore to a large extent the size of the extensive margin effect. As in the CalvoSs hazard function, this form allows for a positive probability at an output gap of 0, which also nests the Calvo hazard function. The positive probability at 0 and the thresholds at which the probability can generalize the simple quadratic form that Caballero and Engel (2006) estimated.

A second functional form that we use is the following logistic function:

$$H(x)^{\text{logit}} = \begin{cases} \frac{a^- \cdot \exp(b^-(|x|-c^-))}{1+\exp(b^-(|x|-c^-))}, & \text{if} \quad x < 0; \\ \frac{a^+ \cdot \exp(b^+(|x|-c^+))}{1+\exp(b^+(|x|-c^+))}, & \text{if} \quad x \geq 0. \end{cases} \tag{3}$$

This functional form has the interesting feature that a discrete choice problem in a rational inattention framework implies a logit specification like this in general (see Woodford, 2008, 2009; Yang, 2015).[5]

Last but not least, we estimate the hazard function in a non-parametric way as an alternative. Specifically, we select nine "points" on the grid of price imbalances, equally spaced between -0.5 and 0.5. We then search for the value of the candidate function at the nine grid points, and assign the values between the grid points by linearly interpolating between them. In other words, the non-

[5]The discrete choice problem (as presented by (see Woodford, 2008, 2009; Yang, 2015) departs from the linear-quadratic-Guassian framework used by (Sims, 1998, 2003).

parametric hazard functions we are considering are a subset of the space of real-valued functions spanned by a basis consisting of the nine grid points on the space of price imbalances. The hazard function is then based on the candidate function, $h(x)$ in the following way:

$$H(x)^{\text{non-param}} = \begin{cases} h(-0.5), & \text{if} \quad x < -0.5 \\ h(x), & \text{if} \quad x \in [-0.5, 0.5] \\ h(0.5) & \text{if} \quad x > 0.5 \end{cases} . \tag{4}$$

We further assume that $h'(x) \leq 0$ if $x < 0$ and $h'(x) \geq 0$ if $x \geq 0$. However, we notably do not impose that the hazard function be symmetric around zero. While the functional forms that we are considering are very flexible, the non-parametric approach naturally imposes even fewer restrictions on the shape of the function.

3.2 Data and Moments

As in Luo and Villar (2017), we use the micro data underlying the U.S. Consumer Price Index for the period 1977-2014. Since it contains a very large number of individual prices tracked over time, this data set enables us to construct statistics related to individual price changes. These statistics can then be compared with the results of simulating firms adjusting prices according to various hazard functions.

The first key moment in the estimation is the average frequency of price change. As shown in equation 1, monetary non-neutrality depends on the frequency of price change, and on the extensive margin effect. The latter has to do with the extent to which prices that are more misaligned are more likely to change, but the size of this effect cannot be observed. That is why a model, or a hazard function, is needed, and this is the focus of our exercise. However, the frequency of price change can be directly observed and estimated with the micro data, and we ensure that our estimated hazard function matches its correct value. The second moment used is the average absolute value of price changes (measured as a percentage change, conditional on a non-zero change occurring). While this statistic does not enter directly into the expressions for aggregate flexibility, it is key to discipline the variance of the idiosyncratic component of desired price changes. It has also been consistently found that price changes are large on average (e.g. Klenow and Kryvtsov, 2008; Nakamura and Steinsson, 2008).[6]

These are the basic moments that all sticky price models or hazard functions should match. And while they are necessary to place restrictions on the parameters of these models (such as the price adjustment probability, or the width of the inaction bands), it is also worth noting that almost

[6] Caballero and Engel (2006) only used the frequency and size moments to estimate the hazard function. They were able to do this by restricting the function to a quadratic function that is potentially asymmetric around 0.

any model can match these. Another moment that can provide additional information is the fraction of price changes that are smaller than a certain threshold. Midrigan (2011) had shown that the standard menu cost model cannot match this moment, as under a fixed menu cost firms will never be willing to pay the cost to carry out small changes. Midrigan's model provides additional flexibility that allows it to match this, and this leads it to predict a much higher degree of monetary non-neutrality than the standard menu cost model. Our hazard function will also match this fraction, with the threshold for a price change being small set at 1% (in absolute value). All the moments described thus far (the frequency, size, and fraction of small price changes) are what we denote "unconditional" moments, as they are all averages across time, and do not capture any time variation in price change behavior.

In Luo and Villar (2017), we had argued that the unconditional moments do not provide enough information to conclusively discriminate between models that predict very different degrees of monetary non-neutrality. However, conditional moments that describe the relationship between inflation and the shape of the price change distribution can be much more informative. Indeed, Luo and Villar (2017) showed that menu cost models predict that the distribution of price changes should become less dispersed and more negatively skewed as inflation increases, while other models (such as the Calvo model, or the rational inattention model under certain parameter values for the cost of information) do not make these predictions. Our estimation therefore includes these correlations among the moments to be matched, and we will show that it is precisely these moments that allow us to reject several types of hazard functions. In addition, we will include the correlation between inflation and the frequency of price change, as it creates a very clear and simple distinction between state-dependent models (in which more firms choose to change their prices when inflation is high) and the Calvo model (in which the same fraction of firms change their prices in every period, by assumption).

All of the moments mentioned are used as targets to estimate the hazard function. We also include the average over time of the dispersion and skewness of price changes, to add further discipline to the shape of the hazard function. Our main results will show that while our different hazard function estimates match the unconditional moments that have been considered by the literature before us, they also match the correlations that we have emphasized here and in Luo and Villar (2017), in a way that no existing model has. In Table 2, we list the moments that we will target, as well as their values in the data.

Using these moments, we will estimate the hazard function according to the functional forms specified above, and the non-parametric form. We do this by running simulations for 60,000 firms and 1,000 periods. Based on the simulated shocks for the firms, a candidate hazard function determines the simulated price changes. We then calculate moments based on the price change distribution, and compare them to their empirical counterparts. The estimated parameters of the

Table 2: Target Moments

	Full sample	1977-1984	1985 onwards
Avg. Frequency	10.7%	13.1%	10.1%
Avg. Dispersion (IQR)	0.099	0.085	0.103
Avg. Skewness	-0.14	-0.02	-0.171
Avg. absolute price change	7.5%	7.6%	7.5%
Fraction of Small Changes	13.2%	14.8%	12.8%
Avg. Frequency of Increases	7.64%	10.3%	6.93%
Avg. Frequency of Decreases	2.97%	2.5%	3.1%
Avg. Size of Increases	7.2%	7.5%	7.1%
Avg. Size of Decreases	7.9%	6.7%	8.2%
Dispersion of Price Increases	0.077	0.073	0.078
Dispersion of Price Decreases	0.087	0.084	0.088
Corr(Frequency,π)	0.69	0.713	0.389
Corr(IQR,π)	-0.676	-0.432	-0.498
Corr(Skewness,π)	0.361	0.425	0.023

hazard function then solve the following optimization problem:

$$\min_{H(x)} \sum_{t \in \{\text{Moments}\}} \left(1 - \frac{t}{t^*}\right)^2, \tag{5}$$

where t is the value of a given moment by a particular hazard function $H(x)$ and t^* is the empirical value. [7]

Next, we will use a subset of the moments to estimate parametric functions related to some of the existing models. Although the main result of this paper is the estimates not based on models, the model estimates illustrate what moments can be matched by hazard functions with different restrictions. Furthermore, they can be used to evaluate the degree of monetary non-neutrality that they imply (with which we compare the non-neutrality from our general estimates).

3.3 General Hazard Function

Here we present the main result of our paper: the price adjustment hazard function estimated using both unconditional and conditional price change moments.

[7] When implementing this procedure, we place a higher weight on the frequency and size of price change moments. The reason for this is that we are particularly interested in making the hazard function imply values of these moments that match the empirical values as accurately as possible. In contrast, for the correlations between inflation and the different price change moments, we are more interested in matching the sign and the general magnitude. This is because we believe that it is the signs of the correlations that are informative about the price setting process, while the exact values are likely to be affected by variables that we do not consider in our framework.

In all the different cases (the quadratic and logistic forms, and the non-parametric form), this is implemented by searching for parameters that lead to the best between the simulated and empirical moments. This is done with a two-step procedure. In the first step, the idiosyncratic shock process parameters (p_ϵ and σ_ϵ) are held constant at a chosen value, and then we solve the minimization problem in equation 5 under various functional forms of $H(x)$ (according to equations 2 - 4). Specifically, we use the pattern search optimization procedure, which is intended for optimization problems in which the gradient of the objective is not defined. This is exactly the case for our problem, as the values of the model-implied moments can only be computed by simulation. Davidon (1991) describes this specific procedure in more detail, and we stop the process once the value of the objective function changes by less than 10^{-6} across successive iterations. With the values of the hazard function fixed, in the second step we adjust the values of p_ϵ and σ_ϵ manually to match the average size and fraction of small price changes. Since there is a very close relationship between those two parameters and moments, this step is relatively simple to implement manually. However, if necessary (if the simulated moment values are still considerably off), we re-run the optimization. Finally, we obtain the the values of p_ϵ and σ_ϵ as well as the parameters in the hazard functions (equations 2 - 4).

Our procedure is in practice a calibration, as we cannot be certain that the parameter values that we settle on are the only ones that match the data. Nevertheless, in searching for the values that best match the moments that we target, we can rule out several types of hazard functions. In the following sub-section we will explain how the moments that we target, including the sign restrictions on the moment correlations, make it possible to set values for the parameters of the hazard function.

In Figure 1, we plot the three different (quadratic, logistic and non-parametric) hazard function estimates that best matches the moments we had set as targets.[8][9] According to our estimate, there is a significant probability of price adjustment (7%) at a price gap of 0, which gives the hazard function a strong ”Calvo feature”. The same was true for the CalvoSs hazard function, but the estimated probability of zero was lower in that case (because a higher probability at 0 helps to match the non-negative correlation between inflation and price change skewness). The function is also considerably asymmetric around 0: for a given absolute value of the price gap, a price increase is considerably more likely to occur than a price decrease (the slope over the positive region is much larger than over the negative region). For example, a desired price cut of 10% only

[8] The quadratic function goes with the following parameter values for the idiosyncratic shock process: $\rho = 0.7$, $\sigma_\epsilon = 0.058$, $p_\epsilon = 0.485$.

[9] Even though they are estimated with different specifications, the functions are quite similar. While there are some noticeable differences in the values over certain regions, this occurs at values of the price gap that have a very low probability density. As we will show, this means that the different functions will imply very similar degrees of monetary non-neutrality, which is high in all of them.

has a 10% probability of occurring, while a 10% price increase has a probability of over 20% in the quadratic hazard function.

What is also striking about this function is that the probability of price adjustment stays relatively low even for large desired price adjustments. Our specification allows for a steep hazard function, and even for discrete jumps in the probability of price adjustment, but the data rejects parameter values that yielded high probabilities of price adjustment for small or intermediate values of the price gap. It is also worth noting that because of the variance of the idiosyncratic shocks, the price gap will lie between -0.2 and 0.2 in the vast majority of cases, and in this region the probability of price adjustment never exceeds 0.4. This is in stark contrast to all menu cost models, under which the probability of adjustment rises to 1 very quickly. For example, according to our estimate of the hazard function associated with the CalvoSs model, the thresholds are -9% and 4.4%. As we will show, this feature of the estimate results from the restriction that the correlation between inflation and price change skewness be non-negative. Furthermore, it leads the hazard function to imply a very high degree of monetary non-neutrality.

Figure 2 plots the hazard function estimates for the two sub-samples, 1977-1984 and 1985 onwards. Although price changes (especially price increases) were more frequent before the Great Moderation, the shape of the estimated hazard functions for the two sub-periods is quite similar. This is at least in part because the values of the moment correlations are also similar.

In summary, the main features of our estimated hazard function are therefore: 1) a significant probability of price adjustment at 0 (there is no inaction region), 2) asymmetry around 0, making price cuts less likely than price increases, and 3) positive slope (implying some state-dependence), but small enough that the adjustment probability remains low in the region over which price gaps are concentrated.

Table 3 presents the moments implied by the estimated hazard functions. They match the unconditional price change statistics as well as the menu cost hazard functions, but succeed in no longer predicting a negative inflation-skewness correlation.

The three different estimates match almost all the moments very closely. The main shortcoming is in matching the large difference between the frequency of price increases and decreases. However, the estimates all imply a frequency of increases that is much larger than the frequency of decreases.[10] Next, we explain in more detail how the key parameters of the hazard function are identified by different moments.

[10] The target moments corresponding to this that were included in the estimation were the frequency of price increases and the difference between the frequency of increases and decreases

Figure 1: All Estimates of the Price Adjustment Hazard Function

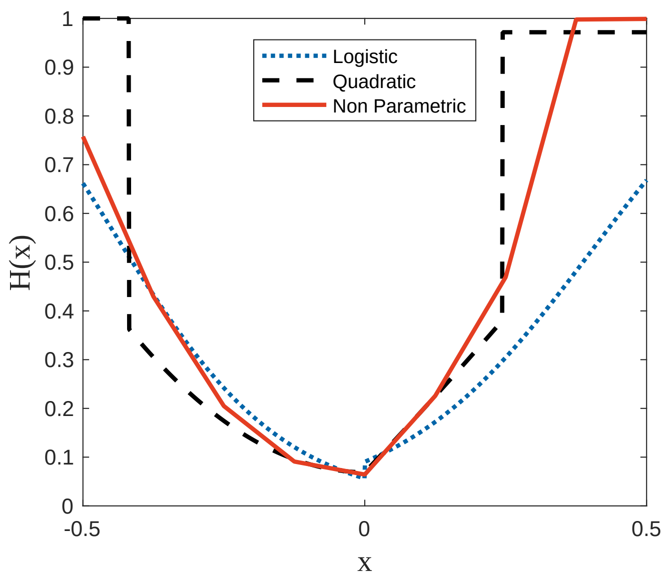

Note: we obtain $p_0 = 0.069$, $a^+ = 1.224$, $b^+ = 0.208$, $c^+ = 0.244$, $a^- = -0.005$, $b^- = 1.665$, $c^- = -0.412$, $\sigma_\epsilon = 0.058$, $p_\epsilon = 0.485$ in the estimated quadratic hazard function.

Figure 2: Sub-sample Estimates of the Price Adjustment Hazard Function

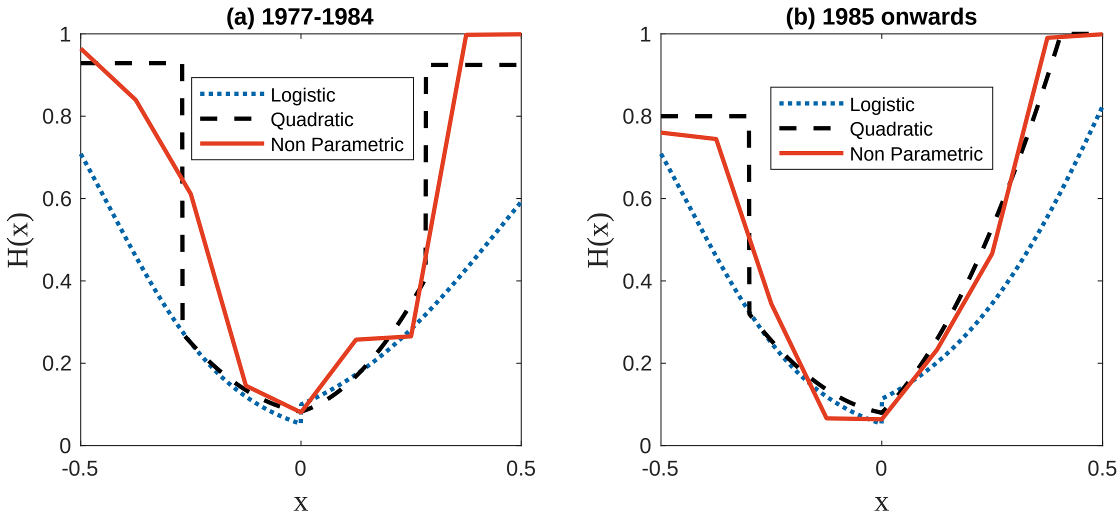

Table 3: Empirical and Simulated Moments

Moment	Data	Quadratic	Logistic	Non-Parametric
Avg. Frequency	0.107	0.107	0.108	0.107
Avg. Frequency of Increases	0.076	0.066	0.068	0.068
Avg. Frequency of Decreases	0.030	0.041	0.040	0.039
Avg. Size	0.075	0.075	0.077	0.077
Avg. Size of Increases	0.072	0.075	0.076	0.076
Avg. Size of Decreases	0.079	0.076	0.080	0.078
Fraction of Small Changes	0.132	0.125	0.130	0.135
Avg. Dispersion (IQR)	0.099	0.12	0.12	0.12
Dispersion of Increases	0.077	0.081	0.089	0.088
Dispersion of Decreases	0.087	0.084	0.084	0.087
Avg. Skewness	-0.14	-0.24	-0.14	-0.27
Corr(Frequency,π)	0.70	0.91	0.82	0.91
Corr(IQR,π)	-0.68	-0.91	-0.88	-0.92
Corr(Skewness,π)	0.36	0.16	0.31	0.11

3.4 Identification of Key Parameters

As presented, the most important features of our estimated hazard function are a positive probability of free price changes, the asymmetry around zero and a small (but positive) slope on both sides. In this subsection, we explain how the moments enable us to set values for the parameters of the hazard function, and why the hazard function must satisfy the aforementioned properties in order to be consistent with the data.[11] We will focus on the flexible quadratic form of the hazard function, as the parameters have a clear intuitive interpretation in that case.

As the most important moment that distinguishes our estimation procedure from existing ones is the inflation-skewness correlation, we mainly focus on the discussion about function characteristics that deliver a positive inflation-skewness correlation. To disentangle the effect of asymmetry and positive probability of free adjustments from other features, consider a linear version of the hazard function 2 (i.e. set $b^+ = b^- = 0$). Figure 3 plots the inflation-skewness correlation under different parameter combinations. The first row illustrates the level of inflation-skewness correlation through color plots, in which blue represents negative values and yellow represents positive

[11] The value of the moments implied by a specific hazard function and idiosyncratic shock process (and how those values change with the parameters of the hazard function and shock process) can only be evaluated by simulation. It is therefore not possible to analytically derive the relationship between the moments and parameters. That is why, to illustrate how are approach is able to pin down the key parameters, we show here how the moments vary in simulations for different values of the parameters.

values. The second row illustrates corresponding hazard functions of the parameter combinations in the first row, while fixing one x-variable or one y-variable (with the color of the function again indicating the value it implies for the skewness correlation).

Figure 3: Effect of varying asymmetry and free price adjustment opportunity

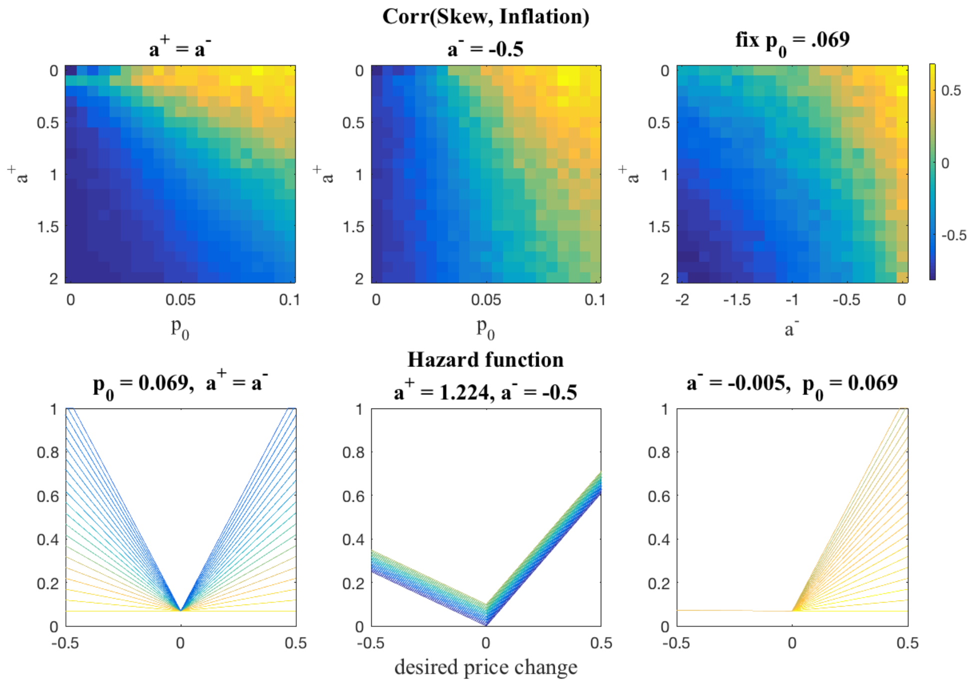

In the first two columns of Figure 3, the probability at 0 (p_0) varies along the x-axis, and the positive linear term (a^+) varies along the y-axis. The negative linear term (a^-) is set to vary with a^+ in the first column, such that hazard functions are symmetric. As illustrated in the figure, the inflation-skewness correlation rises as the linear terms fall or p_0 rises. It is relatively difficult to generate a positive inflation-skewness correlation with symmetric hazard functions unless the adjustment policy has strong Calvo feature (i.e. the linear terms are very small). However, these features would, in turn, generate a very low frequency of price change and a very large fraction of small price changes. In the second column, a^- is fixed at a constant value, -0.5, so that the slope of the hazard function is asymmetric. Here, an intermediate value of p_0 makes it possible to match the non-negative correlation, as the non-negative region has been considerably expanded. This can also be seen in the last column, where a^- is also allowed to vary (allowing for asymmetry) and p_0 is set at the value that we estimate (0.069). Here, if a^- is set to be small, a^+ can be quite high (making it easier to match the overall frequency of price change) while maintaining a non-negative correlation. That is why our estimation procedure settles on an asymmetric hazard function with a positive probability of free price change. In addition, this is consistent with the fact that the frequency of price increases is much larger than that of price decreases in the data.

As presented in the appendix, other patterns in addition to the frequency and size of price change respond to changes in the hazard function parameter values and are important determinants of monetary non-neutrality, such as the fraction of small price changes and the inflation-frequency or inflation-dispersion correlations which provide information on the slope of the hazard function. For instance, a^+ raises the frequency of price increases, but a higher a^+ also decreases the average size of price change (because the density of x is highest for relatively small values). Non-linear parameters, b^+ and b^-, on the other hand, raise the slope of the hazard function, but more so for larger values of x and affect the probability of large price changes occurring, the inflation-skewness correlation is quite sensitive to them. Finally, the patterns for c^+ and c^- show that the skewness correlation is negative whenever these parameters are set to small or intermediate values (less than 0.15 in absolute value). This means that the probability of price adjustment cannot rise sharply for low values of the price gap (as it does in menu cost models) to be consistent with the data.

This analysis has shown how the different parameters of the hazard function can be identified from the moments that we use. In particular, matching the fact that the inflation-skewness correlation is non-negative, while simultaneously matching the empirical value of the frequency and size of price change, places strong restrictions on the slope of the hazard function, and on how high the probability of price adjustment can be for small and intermediate values of the desired price gap. These restrictions are such that the slope of the hazard function cannot be too high, and that there cannot be a large jump in the hazard function for relatively small values of the price gap (as in menu cost models). Finally, it is worth noting that these restrictions can only be inferred using the information from the inflation-skewness correlation. Conversely, disregarding the inflation-skewness correlation, and focusing only on unconditional moments (like the average frequency or average size of price change) would mean that a much larger range of hazard function slopes could be consistent with those moments. That can be seen, for example, in the fact CalvoSs or quadratic Caballero and Engel style hazard functions can match many of these unconditional moments, and yet imply much higher slopes than our estimated hazard function. The ability to use the inflation-skewness correlation to rule out hazard functions with a large slope is, we believe, a significant advantage of our estimation approach, relative to those that estimate a hazard function for each time period and thus cannot use the correlation of moments across time (such as (Berger and Vavra, 2018; Petrella et al., 2018)). Of course, our approach has the disadvantage of not being implementable period by period.

In Appendix C, we present two hazard functions with different shapes and different implications for non-neutrality that match the same set of static moments for a particular time period. This illustrates how the static moments alone may not be sufficient to pin down the key features of the hazard function.

3.5 A Hazard Function Approach of Existing Models

To illustrate the connection between the hazard function and sticky price models, we take the hazard function approach to five existing models, namely Calvo, menu cost, the CalvoSS or CalvoPlus model of Nakamura and Steinsson (2010), Woodford (2009) and Caballero and Engel (2006). Detailed discussions on these estimates are provided in Appendix A.

Figure 4: Existing Model Hazard Functions

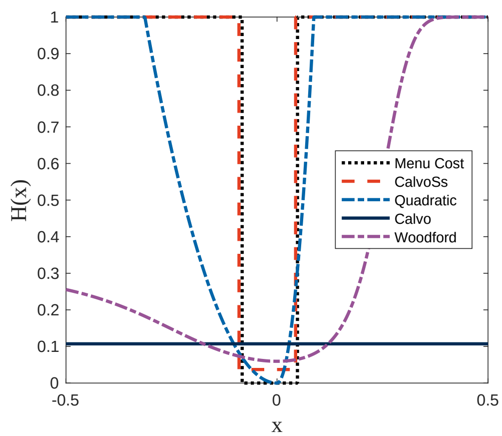

Figure 4 illustrates the hazard functions estimated. Table 4 reports selected price change moments from the simulations under these estimates. The complete simulation results of the targeted moments of each model are listed in Appendix Tables 7-11. As the results make clear, these hazard functions are quite successful at matching the targeted unconditional moments, which have been the focus of most of the literature on sticky prices until now. Indeed, they can all exactly match the overall frequency and absolute value of price changes, and it is mostly in matching the averages decomposed into increases and decreases that there are marginal differences between the hazard functions (as well as the fraction of small price changes). In addition, all the hazard functions (especially for the constant function corresponding to Calvo price setting) feature a very high degree of state-dependence, which can be seen in the large slope of the hazard functions for small values of the price gap. However, following our findings in Luo and Villar (2017), the correlations between inflation and various price change statistics show striking differences between some of the hazard functions.

This general insight has already been known before our results. Indeed, a common criticism of the Calvo model is the assumption that firms are randomly assigned the opportunity to change

Table 4: Moments for Existing Model Hazard Functions

	Avg. Freq	Size	Corr(Frequency,π)	Corr(IQR,π)	Corr(Skewness,π)
Calvo	0.107		0.03	0.49	0.61
Golosov & Lucas	0.108	0.07	0.85	-0.77	-0.99
CalvoSs	0.106	0.074	0.88	-0.94	-0.98
Woodford	0.104	0.152	0.167	-0.109	-0.29
Caballero & Engel	0.107	0.077	0.95	-0.92	-0.99
Data	0.107	0.075	0.70	-0.68	0.36

their price, with a constant probability of price adjustment. While the Calvo model can easily match the frequency of price change, and even the average size of price changes (once the model is augmented with idiosyncratic shocks, as we have done), it was understood that a simple way to reject this assumption would be to show that the frequency of price change rises with inflation. This has indeed been shown by Gagnon (2009) and Alvarez et al. (2011a), among others. In this way, the implications of models at different rates of inflation can provide important information about how they work and how plausible the assumptions underlying them are. The result on the frequency-inflation correlation has, on its own, indicated the general class of menu cost models by highlighting the necessity for state-dependence. However, we show that by looking at the higher moments of the price change distribution, in particular the skewness, we can find important information on how much state-dependence is realistic.

In Table 4, we report the values of the correlations from the simulations under the different hazard functions. One point related to the derivation of the correlations bears clarifying. These correlations are based on the period-by-period variations in inflation and price change moments in the simulations. The variations are the result of the shock (η_t to the aggregate component m_t, and all occur with the trend in the aggregate component μ held constant). That is why we think of these as short-run variations in inflation induced by temporary shocks around a stable trend. This is the simplest way to evaluate the relationship between inflation and the price change moments. However, we verify that these relationships are qualitatively the same if, instead of comparing different periods within the same inflation regime (constant μ), we instead compare the values of the price change moments across different inflation regimes. We implement the latter by running simulations under different values of μ (in each case, simulating the economy for many firms and periods). Under this analysis, it is again true that the dispersion and skewness of price change are lower in simulations with higher values of μ under the menu cost and CalvoSs hazard functions. That is, the hazard functions implied by menu cost models also imply that price change skewness and dispersion will be on average lower in periods of high average inflation. Therefore, whether the variations in inflation in the data are driven by short-run fluctuations or different regimes (with the latter seeming very likely when comparing the period before and after the Great Moderation), the fact that we observe a positive inflation-skewness correlation allows us to rule out the menu-cost based hazard functions.

None of the hazard functions are able to match all three correlations. In particular, while the menu cost hazard functions match the positive frequency correlation, and the negative dispersion correlation, they imply a counter-factual skewness correlation. As we explain in Luo and Villar (2017), this has to do with the fact that, in these models, prices adjust with certainty once they reach a threshold for the mis-alignment. This is a natural consequence of a menu cost, as there will always be a point beyond which it is worth paying the fixed menu cost to adjust a price. This means that as the average of the desired price change distribution rises, a big share of the mass of realized price changes concentrate right beyond the edges of the positive adjustment threshold, inducing more negative skewness. The negative inflation-skewness correlation predicted is not supported by the data. The Calvo model does not feature this kind of effect, as the hazard function implied by it is flat. While this allows it to match the right sign of the skewness correlation, it fails (by assumption) to match the fact that the frequency rises with inflation.

4 Monetary Non-Neutrality

The degree of monetary non-neutrality, or aggregate flexibility, can be computed given a hazard function and parameters for the shock processes. In Section 2, we presented the analytical expressions for aggregate flexibility based on the hazard function and the distribution of price imbalances. We prefer to compute monetary non-neutrality by simulation, as it comes directly from the simulations carried out and does not require computing the stationary distribution of the price misalignments. To do this, we compute the variance (across time) of log real consumption, where log real consumption is defined as $c_t = m_t - p_t$ [12]. The aggregate component of the desired price, m_t, simply follows the random walk process described above, and the aggregate price level is solved for using the hazard function. This is the measure for monetary non-neutrality (the inverse of aggregate flexibility) most commonly used in the sticky price literature (e.g. Golosov and Lucas (2007), Nakamura and Steinsson (2010), and Midrigan (2011)), as it measures the variation in real activity induced aggregate nominal shocks. With full price flexibility, real activity should not vary as prices would respond one-for-one to aggregate shocks. In the hazard function framework, that would be the case if the probability of price adjustment was always 1. In Table 5, we present the results for the Calvo and menu cost hazard functions, the asymmetric quadratic hazard function based on Caballero and Engel (2006) as well as our non-parametric estimate.

The degree of monetary non-neutrality implied by our estimated hazard function is relatively

[12] This would be consistent with thinking about m_t as the money supply, or as nominal aggregate demand

Table 5: Monetary Non-Neutrality

Hazard Function	$Var(c_t) \times 10^4$
Calvo	0.537
Non-Parametric	0.312
Logistic	0.334
Flexible Quadratic	0.329
CalvoSs	0.195
Caballero & Engel	0.176
Golosov & Lucas	0.064

high: it is considerably higher than those based on menu cost models, and it is about half as high as the Calvo The degree of monetary non-neutrality implied by our estimated hazard function is relatively high: it is considerably higher than those based on menu cost models, and it is about half as high as that of the Calvo hazard function. These results are generally in line with our findings in Luo and Villar (2017), which showed that the non-neutrality predicted by the random menu cost model was also between that in the Calvo and Midrigan models. Our results here reiterate the fact that taking into account how the shape of the price change distribution varies with inflation provides evidence in favor of greater non-neutrality than would be expected by simply looking at unconditional moments.

While we do not analytically evaluate the extensive margin component of aggregate flexibility from our estimated hazard function (and the distribution of underlying shocks), the expression for this term in equation 1 helps make sense of the results in Table 5. The extensive margin term is relatively small under our estimated function, because these hazard functions are relatively flat (with a small $H'(x)$) at the imbalances that have the most density (which are mostly those such that $|x| \leq 0.1$). In contrast, the "Ss" type hazard functions feature a very large increase in the price adjustment probability at smaller values of x, for which there is a high density, giving them a very strong extensive margin effect. This is why monetary non-neutrality is high under the estimated hazard functions, and this result comes from the features of the hazard function that are captured by our estimation.

5 Hazard Function of Specific Sectors

In this section, we present the estimated hazard function of four sectors: 1) Processed Food, 2) Unprocessed Food , 3) Household Furnishings and 4) Services. As presented in (Nakamura and Steinsson, 2008, 2010) and Luo and Villar (2017), there is significant heterogeneity of price change statistics across sectors. Given the heterogeneity in various moments of price change distribution across sectors, it is important to estimate and compare hazard functions of various sectors. Table 6 presents the empirical as well as simulated moments of the four sectors. Figure 5 plots the estimated hazard function correspondingly.

Figure 5: Sector-Specific Hazard Functions (Non-parametric)

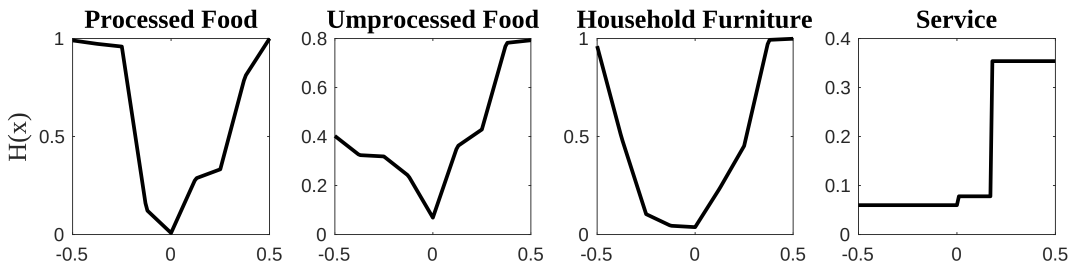

Note: Food 0: $\sigma_\epsilon = 0.065$, $p_\epsilon = 0.55$. Food 1: $\sigma_\epsilon = 0.1$, $p_\epsilon = 1$. Household furniture: $\sigma_\epsilon = 0.06$, $p_\epsilon = 0.4$. Services: $\sigma_\epsilon = 0.07$, $p_\epsilon = 0.15$.

Table 6: Empirical and Simulated Moments - Specific Sectors (Non-parametric)

	Proc Food		UnProc Food		House Furn		Services	
	data	fitted	data	fitted	data	fitted	data	fitted
Avg. Frequency	0.119	0.109	0.287	0.240	0.075	0.078	0.070	0.076
Avg. Size	0.107	0.105	0.123	0.151	0.088	0.085	0.072	0.064
Fraction of Small Changes	0.027	0.024	0.017	0.019	0.097	0.095	0.146	0.242
Corr(Frequency,π)	0.80	0.765	0.700	0.566	0.72	0.948	0.59	0.872
Corr(Skewness,π)	-0.33	-0.968	-0.2	-0.641	0.01	0.022	0.46	0.560
Corr(IQR,π)	-0.69	-0.957	-0.46	-0.659	-0.73	-0.951	0.45	0.312
Avg. Frequency of Increases	0.081	0.068	0.170	0.130	0.052	0.053	0.055	0.053
Avg. Frequency of Decreases	0.037	0.041	0.111	0.110	0.022	0.025	0.011	0.023
Avg. Skewness	-0.309	-0.331	-0.169	-0.151	-0.366	-0.561	0.062	0.171
Avg. Dispersion (IQR)	0.126	0.184	0.250	0.268	0.119	0.124	0.069	0.074
Avg. Dispersion of Increases	0.075	0.077	0.163	0.121	0.085	0.082	0.068	0.090
Avg. Dispersion of Decreases	0.101	0.097	0.178	0.132	0.115	0.100	0.061	0.079
Avg. Size of Increases	0.099	0.099	0.119	0.147	0.082	0.082	0.071	0.065
Avg. Size of Decreases	0.124	0.116	0.127	0.156	0.095	0.091	0.072	0.066

Based on the data and the fitted hazard function, it is clear that price adjustment in the food sectors is more flexible than in the household furnishings or service sectors. This is true in the sense that the frequency of price change is higher in the food sectors (as already shown by Nakamura and Steinsson (2008)). However, we find that this is also true in terms of the degree of state-dependence, and of the role of the extensive margin in price setting. Indeed, the hazard functions for the food sectors have steeper slopes for small values of the price imbalance, and resemble the hazard functions associated with menu cost models. In contrast, the hazard function for services, in particular, is very flat for negative values of the price imbalance. As in the analysis that we have presented so far, the differences in the shape of the hazard functions are inferred from the different signs of the moment correlations across the sectors.

6 Conclusion

As has been shown by Caballero and Engel (2007), the shape of the price adjustment hazard function is closely related to, and provides important information on, the degree of aggregate flexibility (or monetary non-neutrality) implied by micro-level price stickiness. While the question of the significance of monetary non-neutrality has been extensively studied using sticky price models, less attention has been paid to the hazard function approach to this question. This may be in part due to the face that, since it is not grounded in optimizing firm-level behavior, there are very few restrictions that can be placed on the shape of the hazard function a priori. Furthermore, while Caballero and Engel (2007) have derived the exact relationship between the hazard function and aggregate flexibility, they did not consider what empirical patterns could be used to discipline the key features of the hazard function. In this paper, we have attempted to fill this gap, by showing which moments can be used to estimate this function. In particular, we have emphasized that the relationship between inflation and the shape of the price change distribution provides a great amount of information on what shape the hazard function can take, and how much aggregate flexibility it can realistically imply.

We have found that while "Ss" type hazard functions (featuring an inaction region, and a threshold beyond which desired price changes occur with certainty) can successfully match statistics related to the average frequency and size of price changes, they imply a very strong, and counterfactual, negative relationship between inflation and price change skewness. Starting from a very general form for the hazard function, we find one that is able to match both the average size and frequency moments, and the correlations with inflation. In order to match the correlations, and the non-negative inflation-skewness correlation in particular, the hazard function has to include three important properties. First, the probability of a price adjustment at a price imbalance of zero must be positive. Second, even for relatively large price imbalances (of up to 20%), the probability of price adjustment must be considerably lower than 50%. Put differently, this means that the threshold beyond which price changes are very likely is high. Finally, price increases are overall somewhat more likely than price decreases, for an equal size of the price imbalance. The first two properties, in particular, imply that aggregate flexibility is relatively low, and much lower than what would be predicted by “Ss” type hazard functions. The degree of aggregate flexibility is, moreover, broadly consistent with what we find in a random menu cost model in Luo and Villar (2017).

The main contribution of this paper has been to provide a new estimate for the price adjustment hazard function using a richer set of data and empirical moments than in Caballero and Engel (2006), yielding different results on aggregate flexibility. While the hazard function framework that we have been working under is very flexible, there are several variations to our estimation procedure that we could attempt. Indeed, a richer set of processes for the idiosyncratic shocks could be considered, as it would be helpful to know how sensitive the hazard function estimates are to changes in the shock process. Specifically, we have shown that the skewness of the price change distribution can provide important information on what shape the hazard function should take, so it would make sense to work with asymmetric distribution of the desired price change distribution to see what that could mean for the results. Finally, one could also derive the hazard functions implied by other sticky price models (in particular imperfect information models other than the rational inattention model) that have been proposed, and use these to empirically evaluate the models.

References

Alvarez, Fernando, Martin Gonzalez-Rozada, Andy Neumeyer, and Martin Beraja (2011a), "From hyperinflation to stable prices: Argentina's evidence on menu cost models." Working Paper, University of Chicago.

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

Berger, David and Joseph Vavra (2018), "Dynamics of the u.s. price distribution." *European Economic Review*, 103, 60–82.

Bils, Mark and Peter J. Klenow (2004), "Some evidence on the importance of sticky prices." *Journal of Political Economy*, 112, 947–985.

Caballero, Ricardo J. and Eduardo Engel (1993a), "Microeconomic adjustment hazards and aggregate dynamics." *Quarterly Journal of Economics*, 108.

Caballero, Ricardo J. and Eduardo Engel (1993b), "Microeconomic rigidities and aggregate price dynamics." *European Economic Review*, 37, 697–717.

Caballero, Ricardo J. and Eduardo Engel (2006), "Price stickiness in ss models: Basic properties." Working Paper.

Caballero, Ricardo J. and Eduardo Engel (2007), "Price stickiness in ss models: New interpretations of old results." *Journal of Monetary Economics*, 54.

Calvo, Guillermo A. (1983), "Staggered prices in a utility-maximizing framework." *Journal of Monetary Economics*, 12, 383–398.

Caplin, Andrew and Daniel Spulber (1987), "Menu costs and the neutrality of money." *Quarterly Journal of Economics*, 102, 703–725.

Costain, James and Anton Nakov (2014), "Logit price dynamics." Working Paper.

Costain, James and Anton Nakov (2015), "Precautionary price stickiness." *Journal of Economic Dynamics and Control*, 58, 218–234.

Davidon, William C. (1991), "Variable metric method for minimization." *SIAM Journal on Optimization*, 1, 1–17.

Gagnon, Etienne (2009), "Price setting under low and high inflation: Evidence from mexico." *Quarterly Journal of Economics*, 124.

Gagnon, Etienne, David Lopez-Salido, and Nicolas Vincent (2012), "Individual price adjustment along the extensive margin." *NBER Macroeconomics Annual 2012*.

Golosov, Mikhail and Robert E. Lucas (2007), "Menu costs and phillips curves." *Journal of Political Economy*, 115, 171–199.

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.

Luo, Shaowen and Daniel Villar (2017), "The skewness of the price change distribution: A new touchstone for sticky price models." Finance and Economics Discussion Series 2017-028. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2017.028.

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

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 multi-sector menu cost model." *Quarterly Journal of Economics*, 125, 961–1013.

Petrella, Ivan, Emiliano Santoro, and Lasse de la Porte Simonsen (2018), "Time-verying price flexibility and inflation dynamics." Working Paper.

Sims, Chrisotpher A. (1998), "Stickiness." *Carnegie-Rochester Conference Series on Public Policy*, 49.

Sims, Chrisotpher A. (2003), "Implications of rational inattention." *Journal of Monetary Economics*, 50.

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.

Yang, Ming (2015), "Coordination with flexible information acquisition." *Journal of Economic Theory*, 158, 721–738.

Table 7: Calvo Hazard Function

Parameter	Value	Moments	
p	0.107	Avg. frequency	0.107
σ_ϵ	0.055	Avg. absolute price change	0.074
ρ	0.6		

Appendix: The Price Adjustment Hazard Function: Evidence from High Inflation Periods

Shaowen Luo   Daniel Villar

Virginia Tech   Federal Reserve Board of Governors

A Estimated Hazard Function of Existing Models

The simplest sticky price model is the Calvo model, which yields a constant hazard function. The only parameter of the hazard function that needs to be set is the adjustment probability. This can simply be set equal to the overall price adjustment frequency, which in our data is 0.107. While this fully describes the hazard function, the shock parameters must also be specified. We set the aggregate shock process to have a drift parameter (μ) of 0.002 and a standard deviation of 0.0037 (to match the time series properties of U.S. nominal GDP). There is no clear reference to calibrate the parameters of the idiosyncratic shock process (the persistence ρ, and the standard deviation σ_ϵ), so we set them to match the average size of price changes and the ratio between price increases and decreases. The size of price changes, in particular, is largely determined by σ_ϵ. In Table 7, we show the parameter values for the Calvo hazard function, and the moments that they imply.

This hazard function can easily match the overall frequency of price change, and the average absolute value of all price changes (increases and decreases). However, when the frequency and average size are decomposed into increases and decreases, the match is no longer as good.

We next consider the hazard function corresponding to the Golosov and Lucas menu cost model, featuring an inaction region. The parameters to estimate here are the bounds of the inaction region (L and U), and again the idiosyncratic shock parameters. Table 8 shows the parameters we set, and the moments obtained.

This hazard function matches the overall frequency and average size of all increases quite closely. The frequency and size of price decreases are slightly too high, but the fact that price increases are considerably more frequent and slightly smaller on average is captured. Note that in

Table 8: Ss Hazard Function

Parameter	Value	Moments	
L(lower bound)	-0.0827	Avg. frequency of increases	0.077
U(upper bound)	0.0485	Avg. frequency of decreases	0.032
σ_ϵ	0.028	Avg. size of increases	0.07
ρ	0.7	Avg. size of decreases	0.1
		Avg. frequency	0.108
		Avg. absolute price change	0.07

Table 9: CalvoSs Hazard Function

Parameter	Value	Moments	
L (left bound)	-0.09	Avg. frequency	0.106
U (right bound)	0.044	Avg. absolute price change	0.074
p_z	0.037	Fraction of small changes	0.128
p_ϵ	0.15	Avg. frequency of increases	0.075
σ_ϵ	0.071	Avg. frequency of decreases	0.032
ρ	0.7	Avg. size of increases	0.067
		Avg. size of decreases	0.093

order to achieve this, the inaction region is very asymmetric around zero, which is also a common feature in the solution to menu cost models.

We also present the CalvoSs hazard function. This function has one additional parameter relative to the previous one, as the function features a positive probability of price adjustment for small price imbalances, which will lead the model to predict the occurrence of small price changes. Another extension introduced in Midrigan (2011)'s model is that the idiosyncratic variable (firm productivity, in the model) follows a Poisson process. We introduce this modification by using the following process for the idiosyncratic component of the desired price:

$$z_{it} = \begin{cases} \rho z_{i,t-1} + \epsilon_{it}, & \text{Probability} = p_\epsilon \\ z_{i,t-1}, & \text{Probability} = 1 - p_\epsilon \end{cases}.$$

This extension allows the model to match the fact that the distribution of price changes in the data has fat tails, and extending the hazard function set-up in this way has the same effect. This also adds one parameter to set: the probability of a shock occurring (p_ϵ). Table 9 shows the parameter values resulting from our calibration, and the implied moments.

This hazard function matches the frequency of price increases and decreases quite closely, and the fraction of price changes that are small (less than 1% in absolute value). In order to achieve this, the bounds beyond which price changes are certain to occur are asymmetric around zero. The difference between the average size of price increases and decrease is again larger than in the data, but the average size of all price changes is as in the data. In addition, the kurtosis of price changes (around 4 on average) is closer to the empirical value than the Golosov and Lucas simulations.

Next, we consider the hazard function corresponding to the rational inattention model by Woodford (2009). An infinite cost of information corresponds to a Calvo-like hazard function, while free information leads to a trough-shaped one, as in the standard menu cost model (firms must also pay a fixed cost to conducting price reviews, which means that even when processing information is free, they will not change prices every period, and there will be an inaction region). Intermediate values for the cost of information yield smooth functions increasing in the absolute value of the price imbalance, as illustrated in Figure fig: woodford HF. It is also noteworthy that these functions are asymmetric around zero, so that for a given size of the price imbalance, a price increase is more likely than a decrease. In this model, this is due to the asymmetry of the profit function, which makes it more costly to the firm to have its price be too low. This is consistent with the hazard function estimated by Caballero and Engel (2006).

Figure 6: Rational Inattention Hazard Functions

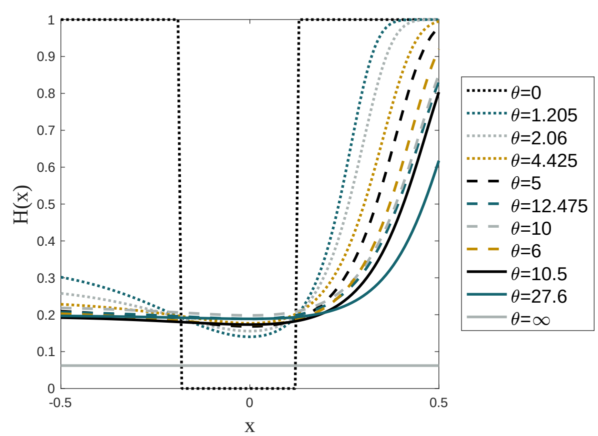

The parameters to estimate are the menu cost variable κ, a cost per unit of information θ and the standard deviation of the idiosyncratic shock σ_ϵ. Table 10 shows the estimated parameters and the moments obtained. While the overall frequency of price change is matched quite closely, the moments related to the size of price changes are not.

Finally, we re-visit the hazard function estimate of Caballero and Engel (2006). Their approach was to use the frequency and average size of price increases and decreases to estimate a simple

Table 10: Woodford Hazard Function

Parameter	Value	Moments	
Information cost θ	1	Avg. frequency	0.104
Menu cost κ	0.6	Avg. absolute price change	0.152
σ_ϵ	0.07	Fraction of small changes	0.04

Table 11: Caballero-Engel Hazard Function

Parameter	Value	Moments	
λ_n	10.2	Avg. frequency	0.107
λ_p	131.5	Avg. absolute price change	0.077
σ_ϵ	0.065	Avg. frequency of increases	0.077
ρ	0.7	Avg. frequency of decreases	0.03
		Avg. size of increases	0.07
		Avg. size of decreases	0.10

asymmetric, quadratic hazard function:

$$H(x)^{\text{CE}} = \begin{cases} \lambda_n x^2 & , \quad \text{if} \quad x \leq 0 \\ \lambda_p x^2 & , \quad \text{if} \quad x > 0 \end{cases}.$$

This form of the hazard function is estimated using the relevant moments from Table 2 (the values are slightly different than in their paper because their values were based on a shorter time period), and the results are in the table below. As the previous hazard functions, this one is reasonably successful in matching the frequency and size statistics. It is clear that the function has to be very strongly asymmetric around zero (with price increases being more likely) in order to match the considerably higher fraction of price increases. Caballero and Engel (2006) had obtained similar results with their original estimate.

B Key Parameters

Figure 7 shows how various moments change when the value of a^{pos} (the linear term in the hazard function for the positive region) varies from 0 to 2. The other parameters of the hazard function are held constant at the values recovered from the estimation.

The patterns displayed are intuitive. A larger a^{pos} raises the probability of price adjustment of price adjustment for positive values of x, which raises the frequency of price increases (and

Figure 7: Effect of varying a^{pos}

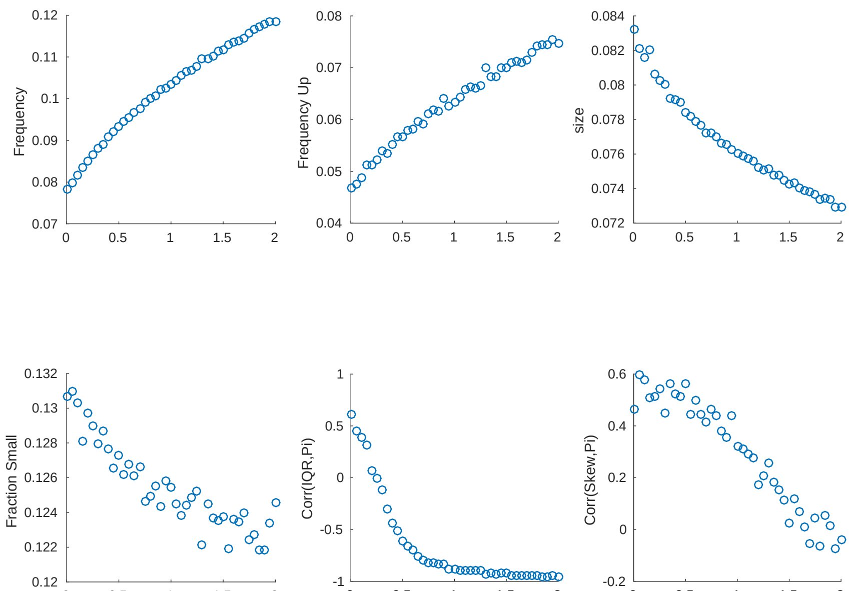

of overall price changes). However, a higher a^{pos} also raises the slope of the hazard function, leading to a stronger selection effect (the extent to which larger price imbalances have a higher probability of adjustment is greater). That is why a smaller share of price changes are smaller than 1% in absolute value. However, the average size of price changes falls because the density of price gaps (x) is highest for relatively small values. This means that as a^{pos} rises, there are proportionally more small and intermediate price changes (and larger than 1% in absolute value). Finally, a greater slope clearly lowers both the dispersion and skewness correlations. For the dispersion correlation, the effect is greatest at relatively small values for the slope parameter (the correlation becomes almost perfectly negative at around $a^{pos} = 1$), while for the skewness correlation the effect is quite consistent throughout the region considered.

Figure 8: Effect of varying a^{neg}

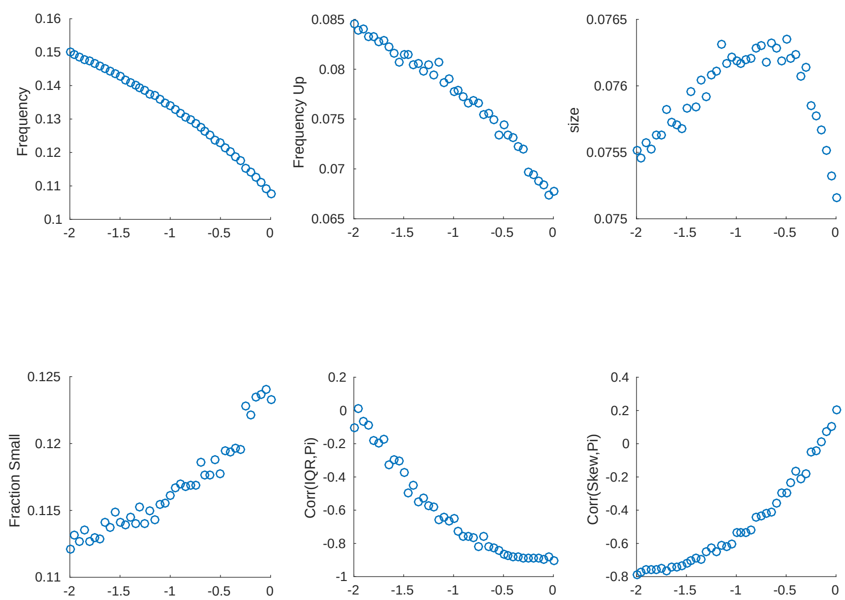

In Figure 8, the same patterns can generally be seen for the slope of the hazard function in the negative region (a^{neg}). As the absolute value of a^{neg} rises, the frequency of price change and decreases rise, while the fraction of small price changes falls. In addition, the skewness correlation falls sharply, and only very small values of a^{neg} are consistent with a non-negative correlation.

Figure 9 below shows the patterns for the parameter setting the probability of price adjustment

Figure 9: Effect of varying p_0

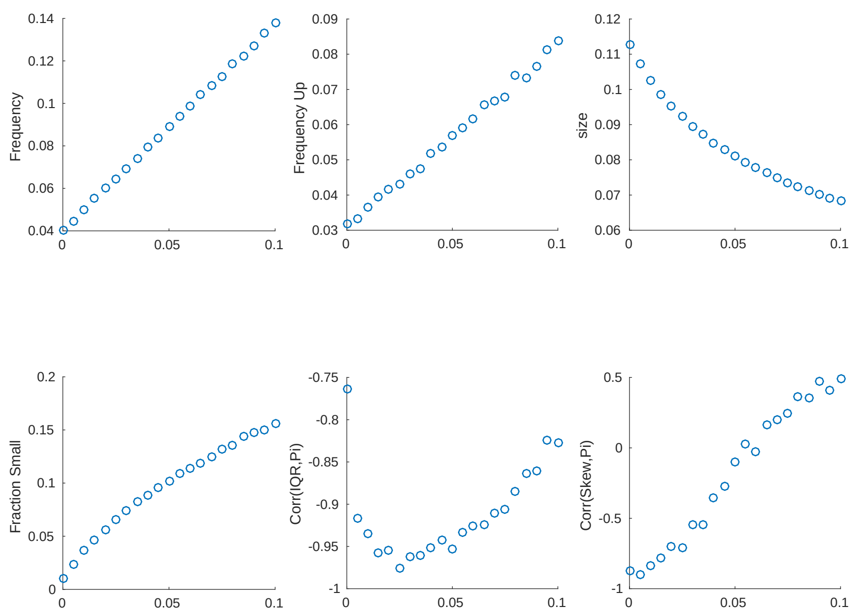

at 0 (p_0). As expected, the frequency of price change and fraction of small price changes rise with $p0$, while the average size of price changes falls. There is also a very strong positive relationship with the skewness correlation, indicating that $p0$must be above 0.05 to attain a positive correlation as in the data. The relation with the dispersion correlation is less clear, and this correlation remains significantly negative for all values of $p0$.

Figure 10: Effect of varying b^{pos}

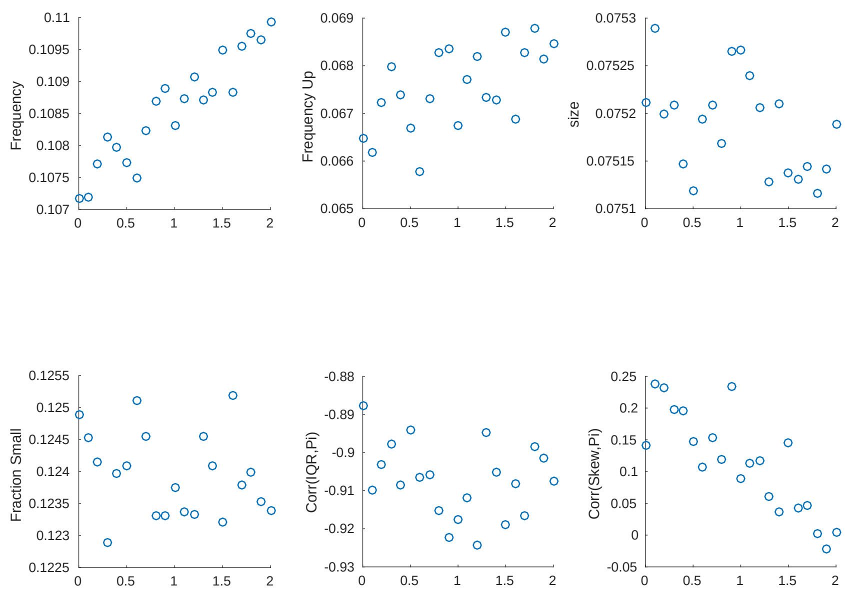

Figure 10 focuses on the effect of varying the quadratic coefficient b^{pos} in the positive region. The relations here are not as clear, but it can still be seen that the frequency of price change rises, and that the size of price change and the moment correlations fall as b^{pos} rises. This parameter raises the slope of the hazard function, but more so for larger values of x. Since the density of x is concentrated around relatively small values of x, b^{pos} does not have such a strong effect on most observable moments. However, since it affects the probability of large price changes occurring, the skewness correlation is quite sensitive to it.

For b^{neg} (Figure 11, these patterns come out somewhat more clearly. That is because the linear parameter for the negative region (a^{neg}) is set to be close to zero (a the value that we estimate), so raising b^{neg} makes a greater difference to the slope of the hazard function in the negative region.

Figure 11: Effect of varying b^{neg}

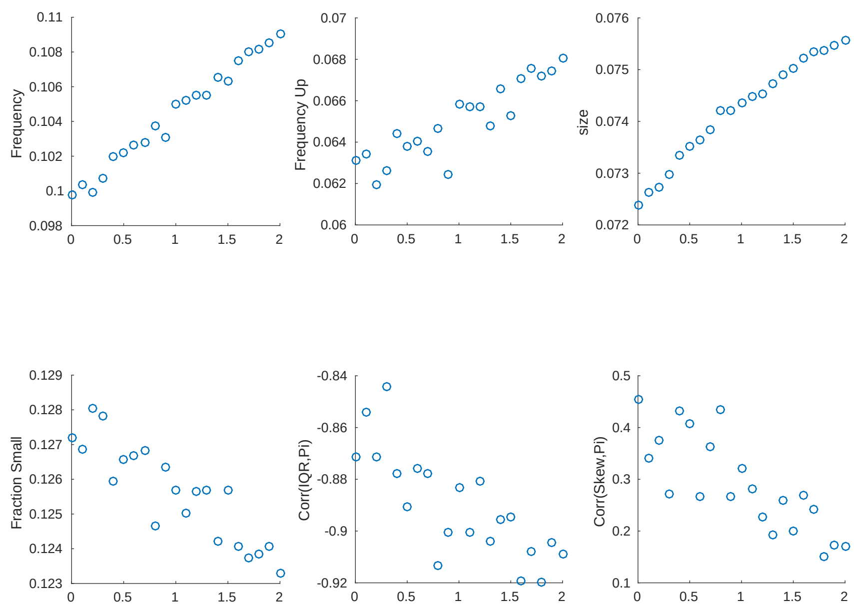

Figure 12: Effect of varying c^{pos}

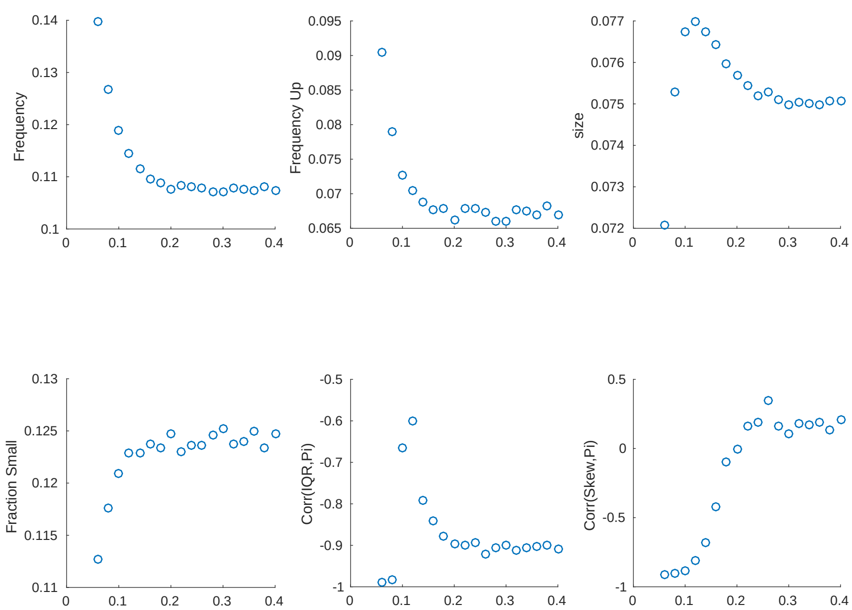

The final set of parameters to consider are the thresholds at which the probability of price adjustment jumps to 1 (c^{pos}&c^{neg}). Figure **??** shows the results for the positive threshold. What stands out is the fact that most moments are fairly constant until the value of c^{pos} reaches 10-20%, after which there are sharp movements. This is again due to the fact that the values of the price gap are concentrated below these values. As expected, the frequency of price change rises sharply once the threshold goes below 15%, and the skewness correlation becomes sharply negative. The patterns for the negative threshold ($c^{negative}$) are generally the same.

Figure 13: Effect of varying c^{neg}

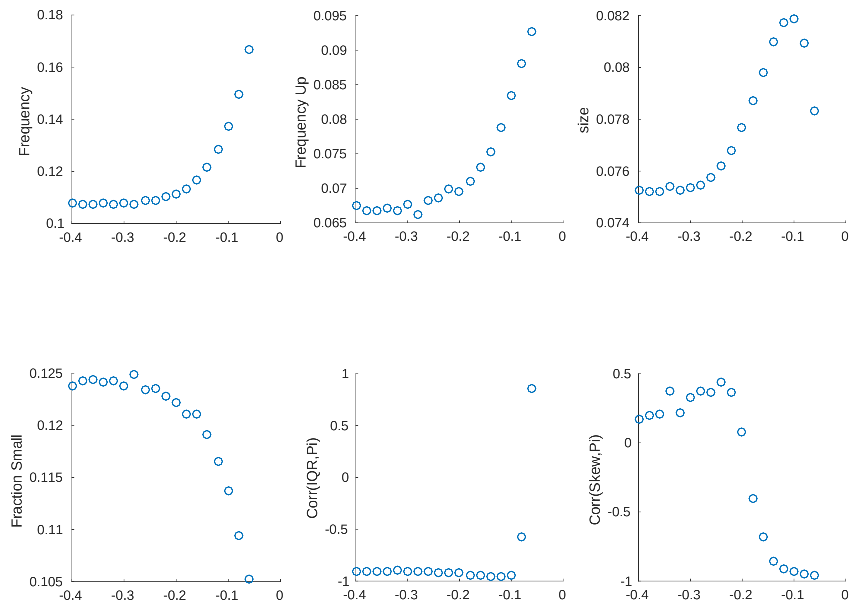

C The Value of Using Moment Correlations

As discussed in section 3, using the inflation-skewness correlation as a target moment in the estimation provides considerable information on the shape of the hazard function. Using this moment in the estimation prevents us from estimating a hazard function for each time period, as is done by (Berger and Vavra, 2018; Petrella et al., 2018), as we cannot estimate the value of the correlation period by period. In this appendix, we illustrate an important benefit of using the inflation-

Table 12: 1984Q2 Moments

Moment	Data	Quadratic	CalvoSS
targeted moments			
Avg. Frequency	9.2%	9.0%	8.9%
Avg. Dispersion (IQR)	0.077	0.080	0.075
Avg. Skewness	-0.15	-0.15	-0.12
Avg. absolute price change	7.2	7.1%	7.3%
Avg. Frequency of Increases	6.5%	6.4%	6.2%
Avg. Frequency of Decreases	2.3%	2.6%	2.6%
untargeted moment			
Corr(Skewness,π)	0.36	0.10	-0.88%

skewness correlation, namely that it allows us to discriminate between different hazard functions that would appear plausible in a generic period if we only relied on other moments.

We estimate two different types of hazard function: a flexible quadratic function as in equation 2, and a CalvoSS-type function that features a discrete jump in the probability of adjustment beyond an inaction region. These two functions are estimated using the same set of target moments from one particular period in our sample[13]. The moments used are: the average frequency of price change, the frequency of price increases, frequency of price decreases, average absolute value of price changes, and the average (over time) inter-quantile range and skewness of price changes. These moments are similar to the ones used by (Berger and Vavra, 2018; Petrella et al., 2018), and their values (in the data, and the values implied by the two different hazard functions) are shown below:

The two hazard functions estimated with these target moments are shown in figure 14.

Although the two hazard functions have similar values for small price gaps, the probability of price adjustment according to the CalvoSS-type function rises rapidly for smaller values of the price gap. An important reason why the estimated functions are different, despite matching the same static moments, is that the CalvoSS function implies a negative inflation-skewness correlation, while the flexible quadratic function implies a flat inflation-skewness correlation[14]. All of this is consistent with the logic laid out at the end of section 3: that a flat or non-negative inflation-skewness correlation is not consistent with a hazard function with rapidly increasing probability of

[13]The period is 1984Q2, and it was chosen for this exercise because its values for the selected moments were broadly similar to those of the average over time.

[14]Although the inflation-skewness correlation is not one of the target moments in this estimation, the estimated quadratic hazard function implies a flat inflation-skewness hazard function because the function is quite similar to the one estimated in section 3 based on all the moments, and the function estimated in 3 implies a flat inflation-skewness correlation.

Figure 14: Alternative Hazard Functions for 1984Q2

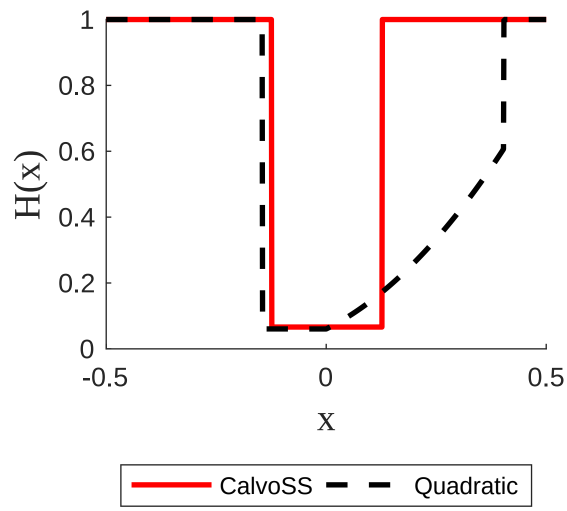

Table 13: Monetary Non-Neutrality for 1984Q2 Hazard Functions

Hazard Function	$Var(c_t) \times 10^4$
Quadratic	0.420
CalvoSS	0.332

price adjustment for small values of the price gap.

This exercise illustrates how not considering the inflation-skewness correlation can result in different estimates for the hazard function. This is quantitatively important, because these functions imply different degrees of monetary non-neutrality, as suggested by the difference in their shape.

The values of non-neutrality implied by each function are shown in table 13. Although these differences are not as large as those between the different types of hazard functions shown in our main results (Table 5), the differences are still substantial.