Learning To Price

Isaac Baley and Julio A. Blanco*
New York University

February 14, 2013

Abstract

Is monetary policy less effective in increasing real output during uncertain times? We argue that firm uncertainty adds flexibility to the aggregate price level and thus the effects of monetary policy are reduced in times of high uncertainty. To construct the argument we develop a price-setting model where firms have imperfect information about their nominal cost and must forecast it to set their prices. Our measure of uncertainty is firms' forecast variance and it is the result of an endogenous learning process. In normal times, the information friction delays the response of prices to a monetary shock and allows for large and persistent real effects on output. In highly uncertain times, however, the real effects are reduced. The reason is that highly uncertain firms become more responsive to new information -which comes with noise- and adjust their prices more often.

JEL: D8, E3, E5

Keywords: Uncertainty shocks, imperfect information, menu costs, (S,s) policies, learning, monetary policy

*We thank Fernando Álvarez, Anmol Bhandari, Ricardo Lagos, John Leahy, Robert E. Lucas, Rodolfo Manuelli, Virgiliu Midrigan, Simon Mongey, Joseph Mullins, Emi Nakamura, Gastón Navarro, Thomas Sargent, Ennio Stacchetti, Laura Veldkamp, and seminar participants at New York University, Princeton University, Washington University St. Louis, and St. Louis FED for very useful comments and suggestions. Correspondence: isaac.baley@nyu.edu and jab772@nyu.edu, 19 West 4th St, 6th Floor, NY 10012.

1 Introduction

How effective is monetary policy to increase real output during uncertain times? This question has become particularly relevant after the 2007-2009 recession. On one hand, there is an agreement that several measures of uncertainty, particularly at the firm level, increased during the recession[1]. On the other hand, the introduction of unconventional monetary policies by the FED may raise questions regarding the effectiveness of conventional policies during such uncertain times. In this paper we argue that monetary policy is less effective to change real output when firm uncertainty is high because uncertainty adds flexibility to the aggregate price level.

Uncertainty in our model is endogenous and it is generated by an information friction that prevents firms from knowing their nominal costs. Firms need to forecast their costs and the forecast variance is our measure of uncertainty[2]. We show that the frequency of price adjustment is increasing in uncertainty because with higher uncertainty forecasts becomes more responsive to new information, which is noisy. Therefore uncertainty adds flexibility to the aggregate price level. When a monetary shocks arrives in times of low uncertainty, firms need time to learn the presence of the shock in order to reflect it in their prices and this delay allows for positive output effects. However, if the monetary shock arrives in a time of high uncertainty, firms forecasts are very volatile and the frequency of adjustment is high. The flexibility of the price level brought by uncertainty reduces the output effects.

Uncertainty plays a role through three channels. In our model, firms must pay a fixed adjustment cost to change their prices and higher uncertainty increases the option value of waiting to make an adjustment. This first effect, known as "wait and see" effect, calls for wider inaction regions and reduces the probability of adjustment. The second channel refers to the role of uncertainty in the evolution of firms' beliefs about their nominal costs. If uncertainty is high, firms put more weight on new information - which is noisy - relative to their past forecasts. Volatile forecasts push firms faster towards their action regions increasing the probability of adjustment. For small adjustment costs, this second effect dominates the first and overall higher uncertainty is associated with more price changes. Lastly, imperfect information slows the transmission of monetary shocks to markup forecasts, and this delay is a decreasing function of the uncertainty in the economy. When uncertainty is high, the transmission is faster.

We calibrate the model to target several moments of the cross-sectional distribution of price adjustments, including the mean and the dispersion of the price changes, together with the frequency and the hazard rate of price adjustment. Targeting a decreasing hazard rate of price adjustment at the firm level, a fact in the data, allows us increase nominal rigidities. We then implement a monetary shock and find that information frictions double the effect of nominal rigidities: without

[1]See for example Arellano, Bai and Kehoe (2010), Bloom, Bond and Van Reenen (2006), and Schaal (2010).

[2]Other studies with endogenous uncertainty include Cogley and Sargent (2005) and Orlik and Veldkamp (2012). In these models, agents learn about models and parameters and uncertainty is the result of their forecasting process.

information frictions a monetary shock has real output effects up to 10 months, while including the information friction extends the response up to 20 months, which is closer to empirical findings in Christiano, Eichenbaum and Evans (2005).

Additionally, we conduct several experiments that combine the monetary shock with an increase in uncertainty for all firms. These experiments are aimed to capture the response of the economy to a monetary shock in times of high uncertainty, but we remain silent in terms of the origin of this uncertainty. It could come from the monetary policy itself, for example if there is an announcement of intervention without specifying its magnitude, or it could come from sources orthogonal to the monetary shock. We find that the output effects are reduced with higher uncertainty and resemble those in the absence of information frictions.

In a related study, Vavra (2012) documents that the frequency of price changes increases during recessions, and explains this fact in a model of menu costs with exogenous countercyclical volatility shocks to firm productivity. In our model with imperfect information, the frequency of adjustment is also increasing in endogenous uncertainty. Since recessions can be associated with higher firm uncertainty, we can also rationalize this fact. In our setup, however, uncertainty shocks are the result of a forecasting process and not an exogenous process.

Our model also offers a new explanation for a decreasing hazard rate of price adjustment, a fact documented in the data but at odds with standard pricing models. A decreasing hazard rate means that the probability of adjusting the price is lower the longer the price has remained unchanged[3]. Decreasing hazard rates are documented in the data by Nakamura and Steinsson (2008) using monthly BLS data and Campbell and Eden (2005) using weekly scanner data[4]. Most price-setting models, however, cannot generate decreasing hazard rates. For example, the hazard is constant in Calvo (1983), increasing in Golosov and Lucas (2007) or models with information frictions, flat in Gertler and Leahy (2008) and Midrigan (2012), or non-monotonic in Taylor (1979) or Álvarez, Lippi and Paciello (2010). In our model, as firms learn their nominal costs, their uncertainty decreases and they expect to transit from high-uncertainty/frequent-adjustment state to a low-uncertainty/infrequent-adjustment state. These learning dynamics rationalize the decreasing hazard rate, which provides a testable implication of a passive learning process in a price-setting framework[5].

[3]The hazard rate is a function of time and should not be confused with the adjustment hazard that is a function of the state, usually the distance between the target and the current price.

[4]Other studies that document decreasing hazards, but do not control for heterogeneity, are Eden and Jaremsky (2009) who use Dominik's weekly scanner data and control for sale prices; Dhyne et al. (2006) who use monthly CPI data for Euro zone countries and Cavallo (2010) that finds hump-shaped hazards for daily online prices in four Latin American countries.

[5]Alternative explanations to decreasing hazards are heterogeneity and survivor bias as noted in Keifer (1980), sales as in Kehoe and Midrigan (2012), mean reversion together with large shocks as in Nakamura and Steinsson (2008), or experimentation as in Bachmann and Moscarini (2011). The data we use controls for heterogeneity or sales, so the first two alternatives cannot explain the decreasing hazard.

This paper complements the literature on price-setting with menu costs. In these models, monetary neutrality results are very sensitive to the stochastic process assumed for the driving force. We nest different specifications of this process including frequent small shocks as in Golosov and Lucas (2007), infrequent large shocks as in Gertler and Leahy (2008), Kehoe and Midrigan (2012), and Midrigan (2012), and transitory shocks that are the source of the information friction as in Maćkowiak and Wiederholt (2009). The tractability of our framework allows us to characterize analytically pricing policies and price statistics and show the contributions of each type of shock to the neutrality results. Interestingly, under imperfect information, the infrequent large shocks to markups translate directly into uncertainty shocks, which are a key feature of our analysis.

The paper also complements the literature on price-setting with information frictions. Models with information frictions target the delayed response of prices to nominal shocks by reducing the information set of the agents[6]. Our model is based on the same principle with two differences. The first difference is that we do not distinguish the source of uncertainty, idiosyncratic or aggregate, and only total uncertainty matters for price flexibility. A second difference is in terms of calibrating the information friction. Firm uncertainty is at the core of our model but it is not observable in the data. We propose a method to calibrate the parameters governing uncertainty using the shape of the empirical hazard rate of adjustment. This is a novel way to discipline information structures, first suggested in Jovanovic (1979).

Finally, our assumptions about information processing reflects empirical work on information frictions. In particular, Coibion and Gorodnichenko (2012) use surveys to compare different information rigidities and find that learning is the most important in explaining US inflation dynamics. They claim that *'Much of our empirical evidence suggests that noisy information models are likely to be the most appropriate characterization of the expectations formation process for professional forecasters, consumers, firms and central banks.'*. Based on their results, we model firms as fully rational agents that face noisy information flows to be processed.

We proceed as follows. Section 2 presents the pricing problem of a representative firm and characterizes analytically the pricing policy. Section 3 characterizes the price statistics that will allow to calibrate the model in Section 4. Section 5 studies the effects of monetary shocks in economies with different levels of uncertainty[7].

[6] See for example Maćkowiak and Wiederholt (2009), Reis (2006) and Woodford (2008).

[7] The Technical, Numerical and Web Appendix can be found at https://sites.google.com/a/nyu.edu/isaacbaley/Research.

2 Firm Pricing Problem

2.1 Setup

Time is continuous[8]. Consider an infinitely lived firm that maximizes expected discounted profits by choosing prices and is subject to productivity shocks. Her instantaneous return function is given by the distance between the currently charged log price p_t and a target log price p_t^*

$$\Pi(p_t, p_t^*) = -B(p_t - p_t^*)^2$$

where B is a constant. This type of quadratic return function is standard in price setting models (see for example Caplin and Spulber (1987), Caplin and Leahy (1991) and Álvarez, Lippi and Paciello (2010)). The random target price is assumed to be a constant markup times the inverse of productivity, therefore in logs we have $p_t^* = M^* - a_t$, where M^* is a constant and a_t denotes the log of productivity. The firm chooses whether to keep her price or to change it, in which case she must pay a menu cost[9] denoted with θ.

Markup Process Instead of prices, it is more convenient to work with accumulated markup gaps μ_t, defined as the accumulated difference between current and target log markups

$$\mu_t \equiv \int_0^t (M_s - M^*)ds$$

where $M_s \equiv p_s + a_s$ is the current log markup given by the price plus productivity, and the target markup M^* is a constant. With this notation, both the markup gap and price gap are given by $d\mu_t$ and the instantaneous profit function can be written as $\Pi(d\mu_t) = -B(d\mu_t)^2$.

During periods when the price is fixed, μ_t is driven completely by productivity which we assume satisfies the following system of differential equations (we use μ instead of a to denote the productivity process since they are the equal):

$$\begin{aligned} d\mu_t &= \mu_t^P dt + \gamma dT_t \\ d\mu_t^P &= \sigma_f dF_t + \sigma_u \epsilon_t dq_t \end{aligned}$$

where F_t and T_t are standard Brownian motions, and q_t is a Poisson process of Gaussian shocks $\epsilon_t \sim N(0,1)$ with arrival rate λ. When there is a price adjustment, the markup gap process is reset to some optimal value and then follows again the process above. Notice that the markup process

[8]The continuous time model is the limit of a discrete time problem developed in the Web Appendix. We show that the discreet time problem converges to the problem analyzed here.

[9]Our favored interpretation for the menu cost is the penalty payed for breaking existent implicit or explicit contracts with customers, which are identified as an important explanation for rigid prices according to the firm surveys conducted by Fabiani, *et.al.* (2005) for European firms and by Blinder, *et.al.* (1998) for US firms.

has no deterministic trend, this is, we assume zero inflation. Since we will calibrate our model to recent US data, we believe zero inflation is a good approximation.

Our productivity process nests different specifications used previously in the literature:

1. Permanent *frequent* shocks with low volatility: captured by the Brownian motion F_t with variance σ_f^2. These small frequent shocks are the driving force in standard menu cost models, for example Golosov and Lucas (2007)[10].

2. Permanent *infrequent* shocks with high volatility: modeled through a Poisson process q_t. With arrival rate λ, the permanent markup gap receives a Gaussian shock with mean zero and variance σ_u^2. These types of shocks have been used in Gertler and Leahy (2008), Kehoe and Midrigan (2012) and Midrigan (2012) as a way to capture relevant features of the empirical price change distribution, such as fat tails.

3. *Transitory* shocks: modeled with the Brownian motion T_t with variance γ^2. These shocks have been used in Midrigan (2012) to explain sales behavior and in pricing models with information frictions as Maćkowiak and Wiederholt (2009) where the shocks are identified with signal noise.

A novelty of our specification is to include the three types of shocks in the markup process. Our motivation is that the implications for monetary neutrality are different if we assume only one type of shock and shut down the others. Using a richer setup, we can identify how a particular type of shock affects pricing policies and aggregate price statistics and then measure their relative importance when calibrating the model with data.

Information Sets To highlight the effects of the information friction we solve two versions of the model, one with perfect information and the second with imperfect information. In the first case the firm can distinguish perfectly between permanent and transitory components of markups. Her relevant information set at period t is given by $I_t = \sigma\{d\mu_s^P, dT_s, s \leq t\}$, where $\sigma(\cdot)$ stands for the generated σ-algebra, i.e. she knows the difference between the components of the markup shocks.

Under the second scenario the firm has imperfect information since she cannot distinguish between the permanent and transitory shocks and only observes their sum. We further assume that she knows that there has been an infrequent shock but she does not know its magnitude. Her information set at time t is given by: $I_t \equiv \sigma\{\mu_s, dq_s; s \leq t\}$. The assumptions about infrequent shocks preserve the normality of the problem and, as we show below, these shocks only affect uncertainty upon impact but not the markup estimates since the shock has zero mean.

[10]Golosov and Lucas (2007) use a mean reverting process for productivity instead of a random walk. Still, our results concerning small frequent shocks will be identified with their setup.

We require that the firm makes her adjustment decisions before her information set is updated with the shock information[11]. The objective is to eliminate the effect of the transitory shock on the pricing policy that operates through the profit function and focus only on its effect through the information structure. Under this assumption, policies only depend on permanent shocks (or their forecast) as shown in Propositions 1 and 3 below. With this assumption we gain tractability but constrain the response of prices to transitory shocks. Kehoe and Midrigan (2012) adopt an alternative timing in which firms can respond to transitory shocks and show that they are irrelevant for monetary shocks.

2.2 Firm Problem and Pricing Policies

In this section we establish the firm problem with perfect and imperfect information and characterize the optimal policies as a function of the fundamental parameters. The policies are a stopping time τ when the price is changed and the markup gap reset, as well as the new markup x. Proposition 1 shows that a firm with perfect information only cares about permanent shocks to markups. Transitory shocks enter her value as sunk costs and do not affect the policy rules directly, only though the information structure (which is irrelevant under perfect information).

Proposition 1 ***(Problem With Perfect Information)*** *Let τ be the optimal stopping policy and x be the optimal new markup. Then for a initial condition μ_0^P, τ solves the following problem:*

$$V(\mu_0^P) = \max_{\tau} \mathbb{E}_0 \left[\int_0^{\tau} -e^{-rt} B(\mu_t^P)^2 dt + e^{-r\tau} \Big(-\theta + \max_x V(x) \Big) \right] \tag{1}$$

subject to

$$d\mu_t^P = \sigma_f dF_t + \sigma_u \epsilon_t dq_t$$

and the markup is reset to $\mu_0 = x$ with every price change.

Using results from stopping-time theory, see Øksendal (2002) or Stokey (2009) for the proof, the optimal time to adjust the markup is given by $\tau = \inf\{t > 0 : \mu_t \notin D\}$ where $D = [-\bar{\mu}, \bar{\mu}]$ is the inaction region or Ss band, and $\bar{\mu}$ depends on fundamental parameters. The symmetry of the bands is inherited from the specification of the stochastic processes and the profit function. Symmetry also implies that the optimal reset level of the markup gap is equal to $x = 0$ (recall that we assume no inflation). Proposition 2 approximates the inaction region under two extreme cases of the parameter space: $\bar{\mu}_f$, when we only have frequent shocks ($\lambda = 0$) and $\bar{\mu}_u$ when we only have infrequent shocks ($\sigma_f = 0$). The proofs of this and other results are in the Technical Appendix and an evaluation of analytical approximations can be found in the Numerical Appendix.

[11] In the discrete time model in the Web Appendix, we require μ_t to be $\mathcal{I}_{t-\Delta}$-measurable. The continuous time problem is approached as the limit of the discrete time problem.

Proposition 2 ***(Pricing Policy With Perfect Information)*** *Define* $\bar{\theta} = \frac{\theta}{B}$.

i) For $\lambda = 0$ *and* $\bar{\theta}\sigma_f^2$ *and* r *small:*

$$\bar{\mu}_f = (6\bar{\theta}\sigma_f^2)^{1/4}$$

ii) For $\sigma_f = 0$ *and* $\frac{\lambda\bar{\theta}}{\sigma_u}$ *small:*

$$\bar{\mu}_u = (\bar{\theta}(r+\lambda))^{1/2}$$

In both cases, the optimal new markup gap is $x = 0$.

The inaction regions are increasing in the normalized menu cost $\bar{\theta}$, but with different elasticities with respect to the menu cost. When $\lambda = 0$, the inaction region is increasing in the underlying volatility σ_f (the value of inaction increases with volatility, as in real option theory). When $\sigma_f = 0$ and σ_u is large enough, the inaction region increases with the arrival rate of infrequent shocks λ. Although σ_u does not appear in the formula for $\bar{\mu}_u$, the numerical elasticity of $\bar{\mu}_u$ with respect to σ_u is positive, but quantitatively insignificant.

Forecasting technology With imperfect information, the firm cannot distinguish between permanent and transitory shocks to her markup but it is on her interest to forecast the permanent component. By observing changes on her markup gap, the firm updates her beliefs about how much is a permanent change and how much only is transitory. We assume that the firm forecasts in a Bayesian way by optimally weighting new versus old information. Her learning is passive in the sense that she cannot take any actions to change the amount or speed of information.

Let $I_t = \{\mu_s, dq_s; s \leq t\}$ denote the information set of the firm. Define $\hat{\mu}_t \equiv \mathbb{E}[\mu_t^P | I_t]$ as the firm's best estimator of the permanent markup gap and $\Sigma_t \equiv \mathbb{E}[(\hat{\mu}_t - \mu_t^P)^2 | I_t]$ as the forecast variance. For convenience we also define $\Omega_t \equiv \frac{\Sigma_t}{\gamma}$, which is forecast variance normalized by the standard deviation of the transitory shock and is our measure of firm uncertainty. The state of the firm with imperfect information now has two components: the forecast of the permanent markup gap $\hat{\mu}$ and also her level of uncertainty Ω.

Proposition 3 ***(Problem With Imperfect Information)*** *Let* τ *be the optimal stopping policy of the imperfect information problem with initial forecast markup* $\hat{\mu}_0$ *and initial uncertainty* Ω_0. *Then* τ *solves the following problem:*

$$V(\hat{\mu}_0, \Omega_0) = \max_{\tau} \mathbb{E}\left[\int_0^{\tau} -e^{-rt} B \hat{\mu}_t^2 dt + e^{-r\tau}\Big(-\theta + \max_x V(x, \Omega_\tau)\Big)\right] \tag{2}$$

subject to

$$\begin{aligned} d\hat{\mu}_t &= \Omega_t dW_t \\ d\Omega_t &= \frac{\sigma_f^2 - \Omega_t^2}{\gamma} dt + \frac{\sigma_u^2}{\gamma} dq_t \\ (\hat{\mu}_0, \Omega_0) & \quad \textit{given} \end{aligned} \tag{3}$$

where W_t is a standard Brownian motion. The process is reset to $\hat{\mu}_0 = x$ when the price is changed.

Equation (2) shows that the forecast enters directly in the instantaneous profit function, while uncertainty affects the continuation value. The system of forecasting equations in (3) says that the forecast $\hat{\mu}_t$ follows a random walk process whose instantaneous volatility Ω_t evolves with time. In the absence of infrequent shocks ($\lambda = 0$) uncertainty Ω_t decreases deterministically towards σ_f, the volatility of the true permanent markup gap. This is a result of the learning process: as time goes by, forecasts precision increases until the only volatility left is fundamental. With $\lambda > 0$, when a shock ϵ_t arrives it only shows as a jump in uncertainty. Since these shocks have zero mean, they do not affect the markup estimation and it moves continuously[12]. These infrequent shocks guarantee that the learning process is active in steady state and can be interpreted as uncertainty shocks.

Uncertainty shocks arise as part of the learning process. Firms markups are infrequently hit by a first moment shock with unknown magnitude. Although these shocks do not affect markup estimations, they do affect the variance of such forecasts. Imperfect information translates first moment shocks into second moment shocks. In this sense, our model takes a stand in the causality of the shocks following the ideas in Bachmann an Moscarini (2011) and Bachmann and Bayer (2011).

As before, the solution to this type of problem is characterized by an inaction region, but in this case the firm makes decisions based on markup forecasts and the inaction region denoted by $\bar{\mu}(\Omega)$ is also a function of uncertainty. Thus the policy becomes:

$$\text{change price if} \quad (\hat{\mu}_t, \Omega_t) \notin D \equiv \{(\mu, \Omega) : |\mu| \leq \bar{\mu}(\Omega)\}$$

It is convenient to define $\Omega^* = \sqrt{\sigma_f^2 + \lambda \sigma_u^2}$. This is the level of uncertainty such that the expected decrease in uncertainty is zero, i.e. $\mathbb{E}[d\Omega_t | \Omega_t = \Omega^*] = 0$.

[12]Suppose that at time t, $dq_t = 1$ then uncertainty becomes $\Sigma(t) = \mathbb{E}[(\mu_t^P - \hat{\mu}_t)^2 | I_t] = \mathbb{E}[(\mu_{t-}^P + \sigma_u \epsilon_t - \hat{\mu}_t)^2 | I_t] = \mathbb{E}[(\mu_{t-}^P - \hat{\mu}_{t-})^2 | I_t] + \sigma_u^2 = \Sigma(t-) + \sigma_u^2$. At time t, the ϵ shock only affects the trend of μ_t and $\mu_t = \mu_{t-}$. This makes $\hat{\mu}_t$ a continuous variable: $\hat{\mu}_{t-} = \hat{\mu}_t$ on $\{w : dq_t = 1\}$.

Proposition 4 ***(Pricing Policy with Imperfect Information)*** *Let* $\Omega^* = \sqrt{\sigma_f^2 + \lambda\sigma_u^2}$*. For small* r*, the inaction region is characterized by*

$$\bar{\mu}(\Omega_t) = \underbrace{\left(6\bar{\theta}\Omega_t^2\right)^{1/4}}_{option}\underbrace{\mathcal{L}_{\bar{\mu}}(\Omega_t)}_{learning} \quad where \quad \mathcal{L}_{\bar{\mu}}(\Omega_t) \equiv \left[1 + \frac{24\bar{\theta}^2}{\gamma}\left(\frac{\Omega_t}{\Omega^*} - 1\right)\right]^{-1/4}$$

and the new markup gap is set to zero: $x = 0$*. The elasticity of the band with respect to* Ω *is* $\mathcal{E} \equiv \frac{\partial \ln \bar{\mu}(\Omega)}{\partial \ln \Omega} = \frac{1}{2} - \frac{6\bar{\theta}^2}{\gamma}\frac{\Omega}{\Omega^*}$

The inaction region has an option and a learning effect that we now explain.

Option effect: reflects the standard result from real option theory that the value of inaction is increasing in the underlying volatility. It is also present in the perfect information case and is increasing in Ω_t. With higher volatility, the continuation value after adjustment decreases because the probability of higher markups gap increases, whereas the continuation value if inactivity is chosen increases because the firm still owns the option of adjustment. Together, higher volatility increases the value of waiting and pushes towards wider inaction regions. This is the $\frac{1}{2}$ term in the elasticity that reflects the option value owned by the firm.

Learning effect $\mathcal{L}_{\bar{\mu}}(\Omega_t)$: captures the effect of the learning process in the policy rule. The current option value depends both on current uncertainty (Ω_t) and also its future path. Since learning affects the path of uncertainty, it also affects the size of future option values and current SS bands. When there is high uncertainty ($\Omega_t > \Omega^*$), it is expected to decrease ($\mathbb{E}[d\Omega_t] < 0$) and therefore future option values also decrease. In turn, this decreases the current option value and shrinks the inaction region by a factor $\mathcal{L}_{\bar{\mu}}(\Omega_t) < 1$. This is a surprising result since one would expect that with high uncertainty the value of inaction would always be higher. This is the $-\frac{6\bar{\theta}^2}{\gamma}\frac{\Omega}{\Omega^*}$ term in the elasticity.

The amplifying effect of learning effect and thus the overall effect of uncertainty on the policy is determined by transitory volatility γ. As we will show in the next section, price statistics will also be affected by this parameter, in particular the hazard rate of price adjustment. In this way, we are able to map the parameter that controls information frictions into observable statistics.

Figure 1 shows the inaction regions in steady state without infrequent shocks (observationally equivalent to perfect information) and with infrequent shocks. All the figures in the paper are computed using the benchmark calibration detailed in Section 4 unless otherwise noted. Observe that without infrequent shocks ($\lambda = 0$, dotted line), the steady state inaction region is given by $\bar{\mu}(\sigma_f)$ since $\Omega_t \to \sigma_f$ as $t \to \infty$, and the distributions of price changes and stopping times converge to degenerate distributions. All firms will eventually be able to track the process with the same variance and the model collapses to that in Golosov and Lucas (2007). By including infrequent shocks ($\lambda > 0$), the learning process remains active in steady state and the ergodic distribution of price changes and stopping times become empirically relevant.

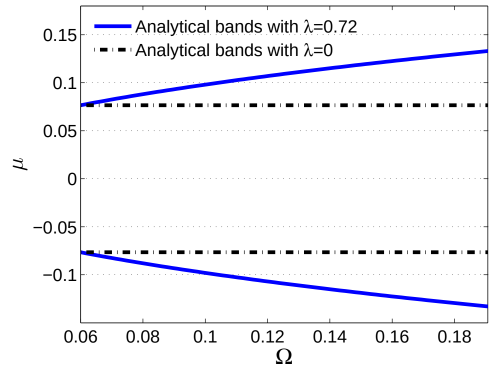

Figure 1: Inaction region in steady state with imperfect information: with infrequent shocks ($\lambda = 0.72$) and without infrequent shocks ($\lambda = 0$) (observationally equivalent to perfect information).

Figure 2 shows one sample path for a firm. The left panel plots the forecast of the markup gap (solid) and the inaction region (dashed). The evolution of the inaction region follows the evolution of uncertainty in the center panel. Uncertainty decreases monotonically with learning until a Poisson shock arrives and makes uncertainty jump up, then learning brings uncertainty down again. The dashed horizontal line is Ω^*, the level of uncertainty that makes expected learning equal to zero. Observe that when the inaction region is wider, the forecast is also more volatile. This is the key mechanism to generate additional price flexibility: higher uncertainty brings more price changes. This mechanism is also explored in Vavra (2012) for exogenous volatility shocks. Finally in the right panel we mark with black the times when there is a price change. These changes are triggered when the markup forecast touches the frontier of the inaction region.

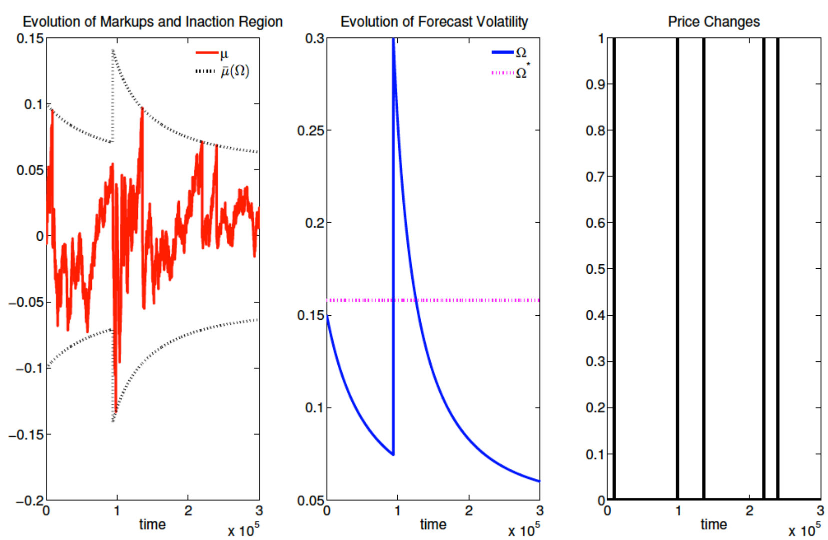

Figure 2: Sample paths for one firm: markup gap forecast and inaction region (left), forecast volatility (center) and price changes (right).

3 Price Statistics

In order to calibrate the structural parameters of the model we use statistics from the cross-sectional distributions of price adjustments and stopping times. In this section we characterize these price statistics for both perfect and imperfect information and for the different underlying stochastic processes nested in our specification to make explicit how each type of shock enters the statistics. The characterizations use approximations and the Numerical Appendix shows that these approximations are accurate enough by comparing them with the exact counterparts computed numerically.

We show three population statistics related to distribution of the absolute value of price changes: the expected size of price change $\mathbb{E}[|\Delta p|]$, the variance of price changes $\mathbb{V}[|\Delta p|]$, and the minimum price change $\min[|\Delta p|]$; and two statistics related to the distribution of stopping times: the expected time between price change $\mathbb{E}[\tau]$ and the hazard rate of price adjustment $h(t)$. The hazard rate of price adjustment $h(t)$ is defined as the probability of changing the price at t given that it has not been changed up to that time, this is $h(t) \equiv \frac{f(t)}{\int_t^\infty f(s)ds}$, where $f(s)$ is the distribution of stopping times. Recall that we are after a decreasing hazard rate as in the data, and we will explain why the specifications with perfect information cannot get it.

3.1 Price statistics with perfect information

Propositions 5 and 6 characterize analytically the price statistics with perfect information in the extreme cases with $\lambda = 0$ and $\sigma_f = 0$, respectively. With perfect information, the individual and average or population statistics are equivalent because there is no heterogeneity.

Proposition 5 ***(Statistics With Perfect Info: Only Permanent Frequent Shocks)*** *Assume $\lambda = 0$, then*

$$
\begin{aligned}
\mathbb{E}[|\Delta p|] &= \bar{\mu}_f && (4)\\
\mathbb{V}[|\Delta p|] &= 0 && (5)\\
\min[|\Delta p|] &= \bar{\mu}_f && (6)\\
\mathbb{E}[\tau] &= \left(\frac{\bar{\mu}_f}{\sigma_f}\right)^2 && (7)\\
h(t) &= \frac{1}{\mathbb{E}[\tau]}\frac{\pi^2}{8}\Psi(t/\mathbb{E}[\tau]) && (8)
\end{aligned}
$$

where

$$\Psi(x) = \frac{\sum_{j=0}^{\infty}(2j+1)(-1)^j \exp\left(-\frac{(2j+1)^2\pi^2}{8}x\right)}{\sum_{j=0}^{\infty}(2j+1)^{-1}(-1)^j \exp\left(-\frac{(2j+1)^2\pi^2}{8}x\right)}$$

is increasing, first convex then concave, $\Psi(0) = 0$ and $\lim_{x\to\infty}\Psi(x) = 1$.

Identification From an identification point of view, the model's parameters $(\bar{\theta}, \sigma_f)$ can be calibrated using $\mathbb{E}[|\Delta p|]$ and $\mathbb{E}[\tau]$ as follows:

$$\begin{aligned}\bar{\theta} &= \frac{\mathbb{E}[\tau]\mathbb{E}\left[|\Delta p|\right]^2}{6} \\ \sigma_f &= \frac{\mathbb{E}\left[|\Delta p|\right]}{\mathbb{E}[\tau]^{1/2}}\end{aligned}$$

This model produces an increasing hazard rate: at the moment of a price adjustment the probability to further adjust is zero ($h(0) = 0$) because the markup gap is at its optimum, and as time goes by, the probability increases in an convex-concave way. More importantly, the hazard rate is only a function of the expected time. Increasing the expected time decreases the level and shrinks the hazard rate in the time dimension, but cannot change its slope. Therefore, there is no calibration that could generate a decreasing hazard rate from this model.

Proposition 6 characterizes the statistics when the firms only receive sporadic shocks. Let $\phi(x)$ and $\Phi(x)$ denote the density and distribution of a standard Normal variable X. We express the statistics in terms of the conditional expectation $\varphi(x) \equiv \mathbb{E}[X|X > x] = \frac{\phi(x)}{1-\Phi(x)}$ (inverse Mills ratio) and the conditional variance $\eta(x) \equiv \mathbb{V}[X|X > x] = \varphi(x)(\varphi(x) - x)$.

Proposition 6 ***(Statistics With Perfect Info: Only Permanent Infrequent Shocks)*** *Assume $\sigma_f = 0$ and let $p(x) \equiv 1 - \left[\Phi(\frac{\bar{\mu}_u - x}{\sigma_u}) - \Phi(\frac{-\bar{\mu}_u - x}{\sigma_u})\right]$ be the probability of falling outside the inaction region given a current markup of x. Then the statistics are given by:*

$$\begin{aligned}\mathbb{E}[|\Delta p|] &= \sigma_u \left[p(0)\varphi\left(\frac{\bar{\mu}_u}{\sigma_u}\right) + (1 - p(0)) \int_{-\bar{\mu}_u}^{\bar{\mu}_u} \varphi\left(\frac{\bar{\mu}_u - \epsilon}{\sigma_u}\right) \tilde{\alpha}(\epsilon|\sigma_u) d\epsilon\right] \quad (9) \\ \min[|\Delta p|] &= \bar{\mu}_u \quad (10) \\ \mathbb{E}[\tau] &= \frac{1}{\lambda p(0)} + e \quad (11)\end{aligned}$$

where $\tilde{\alpha}(\epsilon|\sigma_u)$ is a probability distribution with support in $[-\bar{\mu}_u, \bar{\mu}_u]$ that converges to a uniform distribution as $\sigma_u \to \infty$ and $|e| < \left((1 - p(0))\frac{p(0)-p(\bar{\mu}_u)}{p(\bar{\mu}_u)}\right)$ is a bounded error term.

Let $\{H^n(\mu)\}_{n=0}^{\infty}$ be a sequence of functions defined by:

$$H^n(\mu) = \lambda(1 - p(\mu))\mathbb{E}[H^{n-1}(\epsilon)\big|\bar{\mu}_u \geq |\epsilon|], \text{ with } H^0(\mu) = p(\mu)$$

Then the hazard rate is given by

$$h(t) = \frac{\sum_{n=0}^{\infty} H^n(0)\frac{\lambda e^{-\lambda t} t^n}{n!}}{\sum_{n=0}^{\infty} H^n(0)\frac{\lambda\Gamma(n+1,\lambda t)}{n!}} \quad (12)$$

where Γ is the incomplete gamma function.

Identification In this case, there are three parameters in the model $(\bar{\theta}, \sigma_u, \lambda)$ to identify. We use $\mathbb{E}[\tau]$ to calibrate λ and given that the price statistics are independent of λ we use two price statistics to identify $(\sigma_u, \bar{\mu}_u)$: the mean $\mathbb{E}[|\Delta p|]$ and the normalized variance $\frac{Var[|\Delta p|]}{\mathbb{E}[\Delta p]}$ We check numerically (Figure 3) that both statistics are monotonic with respect σ_u and $\bar{\mu}_u$.

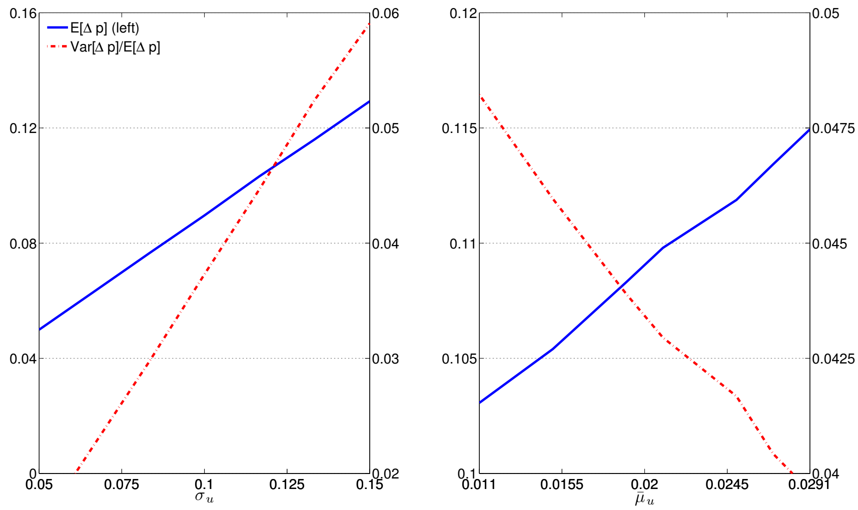

Figure 3: Mean and standardized variance of price adjustments with only infrequent shocks for different values of σ_u and $\bar{\mu}_u$

Regarding the shape of the hazard rate, it depends mainly on σ_u and λ that are pinned down using price statistics. In particular, when σ_u is big enough such that the probability to fall outside the inaction region does not depend on the current markup (i.e. $p(\mu) \approx p(0)$ for all $|\mu| < \frac{\bar{\mu}_u}{\sigma_u}$) then $H^n(\mu) = \lambda^n(1 - p(0)^n p(0)$ and $h(t) = \lambda p(0)$. In other words, in the presence of fat tail shocks the hazard rate is almost flat and this model becomes observationally equivalent to Calvo pricing. This is the type of hazard rate featured in Gertler and Leahy (2008) and Midrigan (2012) and it is associated with monetary non-neutrality.

3.2 Conditional price statistics with imperfect information

Now we turn into the statistics with imperfect information that are conditional on firm uncertainty Ω_t. Proposition 7 characterizes the expected size and minimum of price changes as well as the expected time to the next adjustment. The hazard rate is characterized separately to emphasize its importance as an identification device for transitory volatility γ.

Proposition 7 ***(Price Statistics With Imperfect Information)*** *Let r and $\bar{\theta}$ be small. Then*

$$\mathbb{E}\left[|\Delta p|\big|\Omega_t\right] = \underbrace{\bar{\mu}(\Omega_t)}_{option} - \underbrace{\frac{\bar{\mu}(\Omega_t)^2}{\Omega_t^2}\Big(\frac{\Omega_t}{\Omega^*} - 1\Big) C_1}_{learning} \tag{13}$$

$$\min\left[|\Delta p|\big|\Omega_t\right] = \bar{\mu}(\sigma_f) \tag{14}$$

$$\mathbb{E}[\tau\big|\Omega_t] = \underbrace{\left(\frac{\bar{\mu}(\Omega_t)}{\Omega_t}\right)^2}_{option} \underbrace{\left[1 + \Big(\frac{\Omega_t}{\Omega^*} - 1\Big) C_2\right]}_{learning} \tag{15}$$

where the constants are $C_1 = \bar{\mu}(\Omega^*)\frac{\Omega^*}{3\gamma}\left(\frac{3\gamma - \bar{\theta}^2}{\gamma + 2\bar{\theta}^2}\right)$ and $C_2 = \frac{3\gamma + \bar{\theta}^2}{\gamma + 2\bar{\theta}^2}\frac{24\bar{\theta}^2}{\gamma}$.

The first term of the expected size of the price change $\bar{\mu}(\Omega)$ is the border of the inaction region and is related to the option value of inaction; this is analogous to the perfect information case where $\mathbb{E}(|\Delta p|) = \bar{\mu}_f$. The second term is a result of the learning process that introduces new effects. Under our calibration (see Section 4), the expected size is an increasing function of Ω: a large expected price change is associated with high uncertainty.

The expected time for the next adjustment has two components. An option value component given by the border of the inaction region divided by the volatility. This term is decreasing in Ω: forecasts with high variance are expected to reach the band sooner. The second is a learning component that is increasing in Ω and rescales the option term. The overall effect of uncertainty on the duration of prices depends on its level. A firm with high uncertainty is going to change the price more frequently than a firm with low uncertainty. The main reason is that a firm with high uncertainty is going to put more weight to the new observations compared to her priors, and given that the observation comes with noise, this effect reduces the expected time between price changes.

Identification In this model we need to calibrate 5 parameters: $(\bar{\theta}, \sigma_f, \sigma_u, \lambda, \gamma)$. $\bar{\theta}$, σ_u and λ are set to match $\mathbb{E}[|\Delta p|]$ and $\mathbb{E}[\tau]$. The ratio given by $R = \frac{\lambda\sigma_u}{\sigma_f}$ determines the mode of the distribution. When $R \to 0$, the mode equals the mean, as in Golosov and Lucas (2007). With $R \to \infty$ the mode equals the minimum of the distribution. We can choose σ_f to deliver the desired configuration.

The identification of γ is done through the hazard rate and it is analyzed below. The basic idea is that the probability of price change is positively correlated with the level of uncertainty, and in turn the dynamic behavior of uncertainty is shaped by the transitory volatility. Therefore the dynamic behavior of price changes (the hazard rate) has to be shaped by transitory this volatility.

Hazard Rate With Imperfect Information We characterize the hazard rate of price adjustment *conditional* on an initial condition $\Omega_0 \in (\sigma_f, \infty)$. The hazard rate is a function of two elements: (i) the dynamic behavior of unconditional forecast volatility (Lemma 1); and (ii) the dynamic behavior of inaction regions given the information available at $t = 0$. We cannot provide an analytical expression for the hazard rate with jumps in uncertainty or with non-linear SS bands, so we assume $\lambda = 0$ and constant SS bands to characterize the hazard rate between uncertainty shocks. In the numerical computation we show that these it is good approximation: uncertainty jumps only change the hazard rate at low probability events, where the probability of price change is small.For convenience, we define ignorance as $\kappa \equiv \frac{\Omega_0}{\sigma_f} > 1$, the ratio of initial forecast volatility to permanent volatility σ_f.

Lemma 1 ***(Dynamics of Unconditional Forecast Variance)*** *Assume $\lambda = 0$ and define $\kappa \equiv \frac{\Omega_0}{\sigma_f}$. Then the unconditional forecast (with information set $\mathcal{I}_0$) is given by $\hat{\mu}_t | \mathcal{I}_0 \sim \mathcal{N}(0, p(t))$ and the unconditional forecast variance $p(t)$ follows*

$$p(t) = \underbrace{\sigma_f^2 t}_{option} + \underbrace{\gamma \mathcal{L}_p(t)}_{learning} \tag{16}$$

where $\mathcal{L}_p(t) \equiv \Omega_0 - \sigma_f coth(\log(|\frac{\sigma_f + \Omega_0}{\sigma_f - \Omega_0}|^{1/2}) + \frac{\sigma_f}{\gamma} t)$ is a learning component with the following properties:

- $\mathcal{L}_p'(t) > 0$, $\mathcal{L}_p''(t) < 0$ *and* $\lim_{t\to\infty} \mathcal{L}_p(t) = \Omega_0$.
- *For all t, $\mathcal{L}_p(t)$ is increasing in Ω_0 and increasing in γ.*
- *For all t, if $\sigma_f \to 0$ $p(t) = \gamma(\Omega_0 - \frac{\gamma}{2t})$ with $p''(t) = -\gamma^2 t^{-3}$.*

Unconditional forecast variance $p(t)$ has two components. The option part refers to the linear time trend that comes from the markup estimation following a Brownian motion. The second is the learning component $\mathcal{L}_p(t)$, which is a function only of Ω_0 and (σ_f, γ). Although learning is transitory (with $\lambda = 0$) and eventually dies out, it has a crucial role in shaping its concavity. In turn, the concavity of $p(t)$ will be one of the determinants of the shape of the hazard function.

A higher initial ignorance κ increases $\mathcal{L}_p(t)$, $\mathcal{L}_p'(t)$ and $|\mathcal{L}_p''(t)|$ since the returns to learning will be very high at the beginning and thus the path of unconditional forecast variance will be highly concave. The effect of σ_f/γ on learning is not monotonic or straight forward as with κ, but we can analyze two polar cases in which learning does not have effects. If $\sigma_f/\gamma \to 0$, there is too much noise and learning is slow, therefore $\mathcal{L}_p''(t)$ will be zero. If $\sigma_f/\gamma \to \infty$ for all $t > 0$ $\mathcal{L}_p''(t) = 0$, since learning is too fast. Therefore, if γ is too small or too large, the economy collapses to the perfect information case.

Proposition 8 ***(Hazard Rate Slopes)*** *Let $h_{PI}(t), h_{II}(t)$ be the hazard rates with perfect information and imperfect information, respectively. Then their slopes are related as follows:*

$$h_{II}'(t) \approx \underbrace{[1 + \mathcal{L}_p'(t)]^2}_{>0} \underbrace{h_{PI}'(t + \mathcal{L}_p(t))}_{>0, \to 0} + \underbrace{\mathcal{L}_p''(t)}_{<0} \underbrace{h_{PI}(t + \mathcal{L}_p(t))}_{>0, \to constant}$$

Moreover, if κ is large enough, then the hazard rate is decreasing: $h'_{II}(t) < 0$.

If $\kappa = 1$ (no initial ignorance) then $\mathcal{L}_p(t) = \mathcal{L}'_p(t) = \mathcal{L}''_p(t) = 0$, which implies that $h'_{PI}(t) = h'_{II}(t)$. Recall from Proposition 5 that $h_{PI}(t) \to \frac{\pi^2}{8}\frac{\sigma_f}{\theta}$ and $h'_{PI}(t) \to 0$. Consequently, in this case the hazard rate will be increasing and then constant as in the perfect information case. However, if $k > 1$, the slope of the imperfect information hazard is given by the slope of the perfect information hazard multiplied by the positive term $[1 + \mathcal{L}'_p(t)]^2$ plus the perfect information hazard multiplied by a negative term $\mathcal{L}''_p(t) < 0$, a combination that may yield a negative slope. If k is large enough, the negative effect dominates and the slope of the hazard rate becomes negative.

The key to generating a decreasing hazard rate ($h'_{II}(t) < 0$) is the dynamic behavior of uncertainty. Although learning reduces uncertainty with time, this reduction is not transmitted fully to the policy, in other words, learning has decreasing returns: volatility decreases faster than the inaction regions. Firms with high uncertainty change price more often than firms with lower uncertainty. With learning, a firm is expected to transit from the first to the second group of firms. Furthermore, high initial uncertainty makes the learning process very powerful and makes the slope of the hazard rate more negative.

3.3 Unconditional Price Statistics

In the data we observe population or unconditional statistics. These moments are equal to the weighted average of the conditional statistics, where the weights are given by the renewal distribution of forecast volatility. The renewal distribution $R(\Omega)$ with density $r(\Omega)$ is the stationary distribution of Ω *conditional* on price adjustment: it is the uncertainty faced by adjusting firms. Such distribution is different to the *unconditional* steady state distribution of Ω denoted $S(\Omega)$, which is the uncertainty in the whole cross-section. Propositions 9 and 10 characterize these distributions. Since the learning process is passive, these distributions do not depend on $\hat{\mu}$.

Proposition 9 *Let $S(\Omega)$ be the steady state distribution Ω, with density $s(\Omega)$. Then $S(\Omega)$ solves the Kolmogorov Forward equation*

$$0 = \begin{cases} \frac{d}{d\Omega}\left(\frac{\Omega^2-\sigma_f^2}{\gamma}s(\Omega)\right) - \lambda s(\Omega) & \text{if } \Omega \in (\sigma_f, \sigma_f + \sigma_u^2/\gamma] \\ \frac{d}{d\Omega}\left(\frac{\Omega^2-\sigma_f^2}{\gamma}s(\Omega)\right) + \lambda(s(\Omega - \sigma_u^2/\gamma) - s(\Omega)) & \text{otherwise} \end{cases} \tag{17}$$

making sure that $\int_{\sigma_f}^{\infty} s(\Omega)d\Omega = 1$ and that it is continuous in all its domain.

Proposition 10 *Let $R(\Omega)$ be the renewal distribution with density $r(\Omega)$. Also let $F(\Omega'|0,\Omega)$ be the transition distribution from the current state $(0,\Omega)$ to future Ω':*

$$F(\Omega'|0,\Omega) \equiv Pr[\omega \leq \Omega'|0,\Omega]$$

Then $R(\Omega)$ is characterized by the fixed point of the following expression:

$$R(\Omega) = \int_{\sigma_f}^{\infty} R(\Omega') dF(\Omega'|0, \Omega) \tag{18}$$

Figure 4 plots the densities $s(\Omega)$ and $r(\Omega)$ over the relevant domain for the benchmark calibration in continuous time (see next section for details). We observe that firms' uncertainty is concentrated at low values (s puts most of the weight at low Ω), but the uncertainty of *adjusting* firms is higher (as r puts more weight at high Ω).

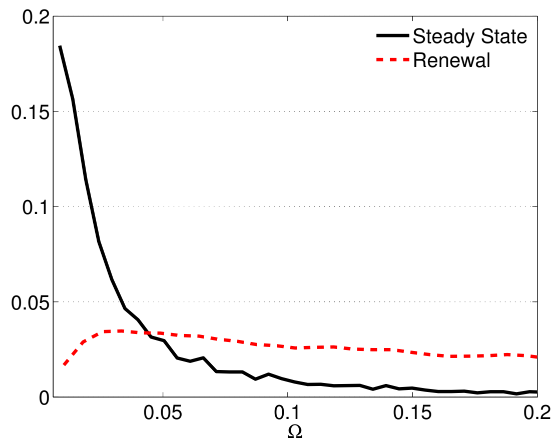

Figure 4: Steady state $s(\Omega)$ and renewal densities $r(\Omega)$ of Ω. Benchmark calibration in continuous time.

With the renewal distribution at hand, we can take any conditional moment $M(x|\Omega)$ and compute the unconditional one as:

$$M(x) = \int_{\sigma_f}^{\infty} M(x|\Omega) r(\Omega) d\Omega$$

For example, the population expected size becomes

$$\mathbb{E}[|\Delta p|] = \int_{\sigma_f}^{\infty} \mathbb{E}\left[|\Delta p| \big| \Omega\right] r(\Omega) d\Omega$$

and the population expected time is

$$\mathbb{E}[\tau] = \int_{\sigma_f}^{\infty} \mathbb{E}[\tau \big| \Omega] r(\Omega) d\Omega$$

Since the renewal density puts more weights on high values of Ω, the conditional behavior of firms with high uncertainty has a larger influence on population statistics. For example, a decreasing hazard rate in the aggregate reflects the behavior of firms with high uncertainty.

4 Bringing the model to data

The macroeconomic consequences of a monetary shock are driven by the type of idiosyncratic shocks assumed. We will compare our model with imperfect information with two benchmarks from the literature as point of references: Golosov and Lucas (2007) [GoLu], and Gertler and Leahy (2008) [GeLe]. We solve numerically the models to get the pricing policies and we simulate the steady state of an economy with $N = 100,000$ firms assuming that all shocks are independent across firms. The only difference between firms is their initial uncertainty which is drawn from its ergodic distribution. We ensure stationarity by checking the convergence of a set of moments. Lastly, we calibrate the parameters by matching the price statistics produced by the simulation with the price statistics found in the data.

Data sources We use statistics reported in Nakamura and Steinsson (2008), who use BLS monthly data for a sample that is representative of consumer and producer prices except services, controlling for heterogeneity and sales. We restrict to *regular* price changes or with sales filtered out. We use the mean price change $\mathbb{E}[|\Delta p|] = 0.11$, the normalized variance of price changes $\frac{\mathbb{V}[|\Delta p|]}{\mathbb{E}[|\Delta p|]} = 0.058$, the expected time between price changes $\mathbb{E}[\tau] = 8$ months, as well as the hazard rate. These moments are also consistent with Dominick's database reported in Midrigan (2012).

Calibration The frequency considered is weekly and we set the discount factor to $\beta = 0.96^{1/52}$ to match an annual interest risk free rate of 4%. The normalized menu cost $\bar{\theta}$ and parameters of the stochastic processes $(\sigma_f, \gamma, \sigma_u, \lambda)$ will be calibrated for each model in order to match the same price statistics across models.

- Perfect info, GoLu ($\gamma = \sigma_u = \lambda = 0$): Two parameters $\bar{\theta} = 0.055$ and $\sigma_f = 0.0178$ match two moments $\mathbb{E}[|\Delta p|]$ and $\mathbb{E}[\tau]$.

- Perfect info, GeLe ($\gamma = \sigma_f = 0$): Three parameters $\bar{\theta} = 0.0167$, $\sigma_u = 0.12$, $\lambda = 0.035$ match three moments: $\mathbb{E}[|\Delta p|]$ and $\mathbb{E}[\tau]$, and $\frac{\mathbb{V}[|\Delta p|]}{\mathbb{E}[|\Delta p|]}$.

- Imperfect info: $(\sigma_u, \bar{\theta}_u) =$ to match $(\mathbb{E}[|\Delta p|], \mathbb{V}[|\Delta p|]/\mathbb{E}[|\Delta p|])$ and $\lambda = 0.025$ to match $\mathbb{E}[\tau]$. We choose $\sigma_f = 0.0005$ and $\sigma_f = 0.155$ to deliver a large ratio $\frac{\lambda \sigma_u}{\sigma_f}$ and set the mode of the distribution equal to the minimum as in the data. We calibrate transitory volatility $\gamma = 0.046$ to match the shape of the hazard rate in the data. The right panel of Figure 5 shows how the slope of the unconditional hazard rate varies with different choices of γ. Notice that the price change distribution (left panel) is not affected by this parameter.

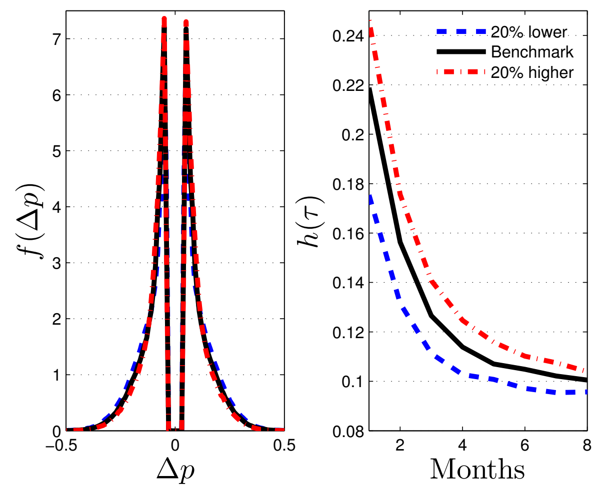

Figure 5: Price change distribution and hazard rate for different choices of transitory volatility γ.

Price Change Distribution The left panel in Figure 6 shows the ergodic distribution of price changes for the three models and the data. Its symmetry comes from the assumption of zero inflation and the symmetry of the problem. The model is not able to match small price changes because the minimum is bounded by the size of the menu cost. For additional features that could generate small price changes see Midrigan (2012) and Alvarez and Lippi (2012) that use multi-product firms.

The model with only small frequent shocks [dotted line, GoLu] generates a price distribution concentrated at the Ss band. The models with large frequent shocks [dashed line, GeLe] and imperfect information [dashed dotted line] are able to match better the data with fat tails and larger dispersion.

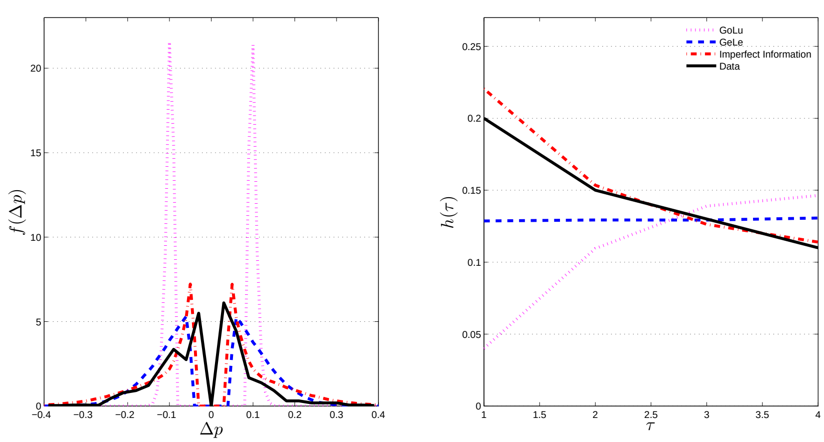

Figure 6: Distribution of price changes and hazard rate of price adjustments for the three models and the data. The data comes from Nakamura and Steinsson (2008).

Hazard Rate The right panel in Figure 6 plots the hazard rate of price adjustment in the models and the data. Infrequent shocks with perfect information [GeLe] generate a non-decreasing almost flat hazard. Perfect information and only frequent shocks [GoLu] and imperfect information without infrequent shocks are observationally equivalent in steady state and feature an increasing hazard rate. The decreasing hazard rate in the imperfect information model results from the interaction of infrequent shocks and the learning process, and it is able to match the shape of the hazard rate in the data.

Cross-sectional distribution of firm uncertainty Uncertainty is unobservable in the data, but we can compute its steady state density and the renewal density from the simulation, pictured in Figure 7. Consistent with the continuous time version of the model (recall Figure 4), the renewal distribution puts most of the weight on large Ω.

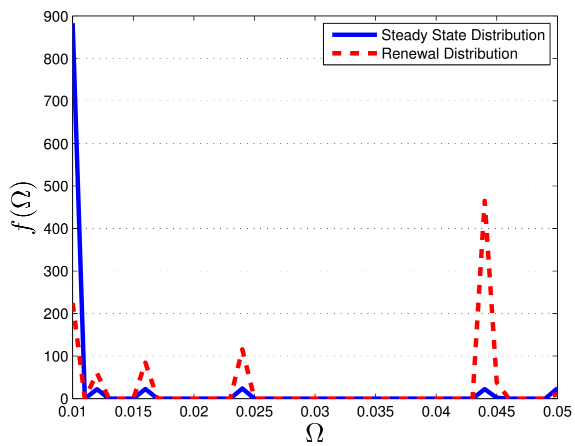

Figure 7: Steady state (solid) and renewal (dashed) distributions of firm uncertainty. The steady state distributions puts most of the weight at low values of uncertainty, while the renewal distribution puts most of the weights at high values of uncertainty.

When we compute price statistics, the relevant distribution of uncertainty is that of firms that change their price (the renewal distribution). This distribution assigns a lot of weight to high uncertainty levels, which means that the behavior of aggregate statistics will mostly mimic the behavior of high uncertain firms. At the same time, the cross-sectional distribution of uncertainty places most of the weight on low levels of uncertainty. Together, this implies that it is possible to have an economy with average low uncertainty but statistics that reflect high uncertainty. We interpret this result as a selection effect in observable actions.

5 Macroeconomic consequences of uncertainty

How does an economy with information frictions respond to an aggregate nominal shock? In this section we will show that imperfect information amplifies nominal rigidities in normal times, while it reduces nominal rigidities if uncertainty is sufficiently high. For this end, we will study the response of output to a monetary shock in a general equilibrium framework as the one in Golosov and Lucas (2007) or Álvarez and Lippi (2012), where we have added imperfect information[13]. For normal times we refer to steady state levels of uncertainty as calibrated from the price statistics.

We implement an aggregate nominal shock as a permanent decrease in markups to all firms of size δ. This shock is unanticipated because firms assign zero probability to it. Using Proposition 8 of Álvarez and Lippi (2012) that shows how the feedback from aggregate markups to individual policies is irrelevant, we compute the impulse response of output to the shock with the aggregate formula:

$$\breve{Y}_t = \delta - \breve{P}_t$$

where δ is the size of the aggregate shock and $\breve{X}$ denotes deviations from steady state of aggregate variable X in logs. In a frictionless world, all firms would increase their price in δ to reflect the higher marginal costs and the shock would have no output effects. As long as the price level fails to reflect the higher costs, there will be real effects on output.

5.1 Information Frictions Increase Nominal Rigidities

The first experiment we conduct is aimed to show how information frictions interact with nominal rigidities. We assume that there is a negative aggregate markup shock of $\delta = 2\%$ but firms don't know that the shock has arrived. We label this experiment as *undisclosed shock*. Next, we compute the response of output under imperfect information and under the two benchmarks of perfect information of Golosov and Lucas (2007) and Gertler and Leahy (2009). Figure 8 compares the impulse-response functions (IRF) of output for the three models. We plot the median IRF obtained after 10,000 simulations and the 95% confidence interval.

The response of output is smallest in an economy with small frequent shocks [GoLu], followed by an economy with large infrequent shocks [GeLe], and lastly with imperfect information. We explain the large response of output in the imperfect information model in the light of three effects.

(i) **Absence of Extensive Margin Effect:** This is the main difference between GoLu and the two other models. After a positive monetary shock, the distribution of markup gaps shifts to the left and on impact, the mass of firms below the SS band will increase the price in the size of the SS band (in the figure, it appears as the initial drop in the IRF). The crucial

[13] In the Web Appendix, we develop such model and explain how to compute the impulse response of output. Features of this framework include monopolistic competition and a representative household with money in the utility function.

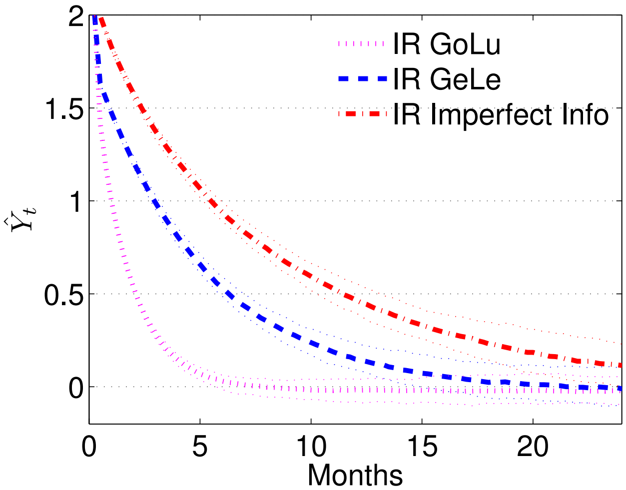

Figure 8: Impulse-response of output to an aggregate negative shock of markups of $\delta = 2\%$.

difference between models lies in the subsequent price changes. In GoLu, the probability of adjusting for firms near the SS band remains high and they make large adjustments in short time, eliminating the effect of the monetary shock. This is the selection or extensive margin effect highlighted in the literature. In GeLe and imperfect information, the probability of adjusting is primarily governed by the arrival rate of infrequent shocks λ (because σ_f is either zero or very small), so there are no more changes triggered by the money shock. The presence of large infrequent shocks, either to markups directly or to uncertainty, eliminates this effect.

(ii) **Decreasing Hazard rate Effect:** Both the models of GeLe and imperfect information have infrequent shocks, but they differ in the shape of the hazard rate. As remarked by Carvalho (2006), what matters for measuring the impact of monetary shocks is the first price change after the shock. Given that σ_f is small in GeLe or very small with imperfect information, the arrival rate of the first price change is given by the probability of receiving an infrequent shock λ.

In the imperfect information model, a decreasing hazard rate at the individual level implies correlated price changes after the first price change. Since the expected time between price changes is fixed across models, the extra price changes triggered by the first change allow us to calibrate a lower arrival rate of the shock. The arrival rate with perfect information is 0.035 and in the model with imperfect information is 0.025. Therefore, the arrival rate of the first price change is lower in the imperfect information setting. By calibrating a smaller arrival rate, the imperfect information model is able to produce a larger response of output than GeLe.

(iii) **Learning Effect:** Firms' learning process generates a further delay in the response of prices to the shock. After a monetary shock, firms gradually adjust their expectation about the permanent component of their marginal cost and therefore their actions are also gradual. The shock will be fully reflected in price changes until learning about the shock is completed.

In our model the only object of interest to forecast is nominal marginal cost. Given that firms know their price, the individual forecast errors are given by $\hat{\mu}_t^i - \mu_t^i$ and the aggregate forecast error is the average $\mathbb{E}^i[\hat{\mu}_t^i - \mu_t^i]$. An undisclosed monetary shock generates a positive aggregate forecast error that decreases over time. This is a common feature of several models with information rigidities such as rational inattention, inattentiveness, among others. The monetary shock, however, does not change the cross-sectional dispersion of forecast errors, it remains constant. Constant cross-sectional dispersion of forecast errors following a structural shock is a feature exclusive to learning models as discussed in Coibion and Gorodnichenko (2012).

To illustrate the dynamics of forecast errors we compare the response of output to the *undisclosed shock* above with the response to a *fully disclosed shock*, a scenario where firms know that the shock arrived and also its exact magnitude. In this second exercise, we shut down the learning effect completely. The response of output under the two experiments is plotted in Figure 9. The left panel shows the IRF, the center panel shows the cross-sectional average of forecast errors $\mathbb{E}^i[\hat{\mu}_t^i - \mu_t^i]$ and the right panel shows the cross-sectional average of markup gap forecasts $\mathbb{E}^i[\hat{\mu}_t^i]$. We do not show forecast dispersion because it remains constant at its steady state level.

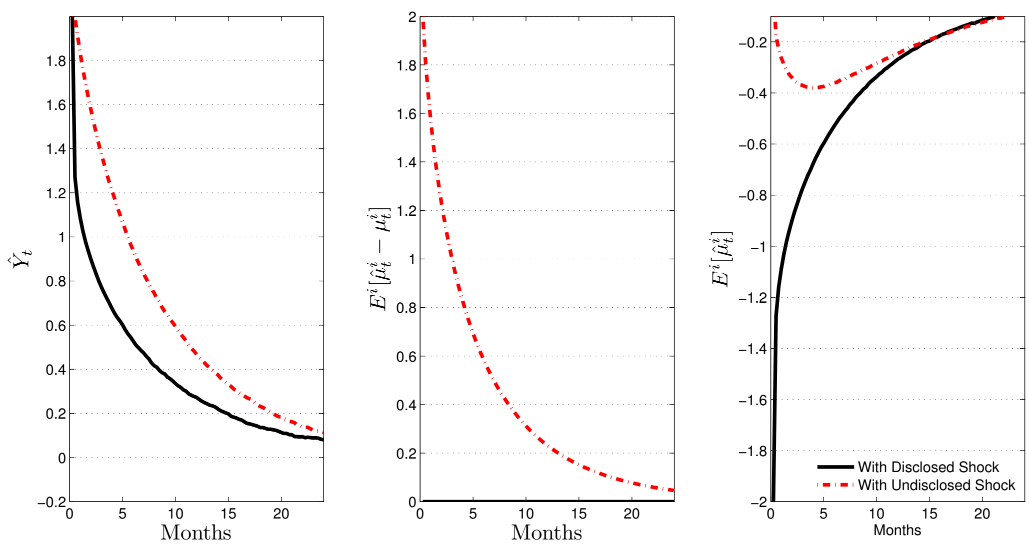

Figure 9: Responses to aggregate negative shock in markups of $\delta = 2\%$. Left panel: Impulse-response of output, Center panel: Cross-sectional average forecast error, Right panel: Cross-sectional average of markup gap forecast.

With the fully disclosed shock (in solid line) firms fully adjust their expectations since there is no learning effect and average forecast errors are zero. The average forecasts are adjusted in the size of the shock δ and go back slowly to zero as the infrequent shocks start to hit and firms adjust their prices. With the undisclosed shock (dashed line) there is a larger output response. On impact, the shock creates a bias in firms' forecasts, but forecast errors converge to the disclosure case as firms learn that the shock is permanent. Note that the evolution of markup forecasts is hump-shaped, reflecting the delay caused by learning.

In conclusion, imperfect information adds rigidities to the aggregate price level that manifest though the three effects (absence of extensive margin, decreasing hazard rate and learning) and important monetary non-neutrality arises.

5.2 High Uncertainty Reduces Nominal Rigidities

In this section, we show that if uncertainty is very high, imperfect information actually reduces nominal rigidities and monetary shocks do not have significant effects. To do this, we interact the monetary shock with a correlated shock of uncertainty across firms. We implement the shocks by moving the markups of all firms in δ and increasing their uncertainty Ω in $n\delta^2$ units, where n takes different values across experiments $n \in \{0, 2, 4\}$. The firms know that a shock has arrived but they do not know its size, they only know that the shock is within $\sqrt{n}\delta$ standard deviations away from zero. We do not take a stand on the source of the correlated uncertainty shock. In principle, it could come from a change in structural conditions of the economy or from monetary policy itself.

The experiments are explained in Figure 10. We plot the cross-sectional distribution of uncertainty in the left panel and the cross-sectional distribution of markup forecasts normalized by the SS band (which is shown as vertical lines) in the right panel.

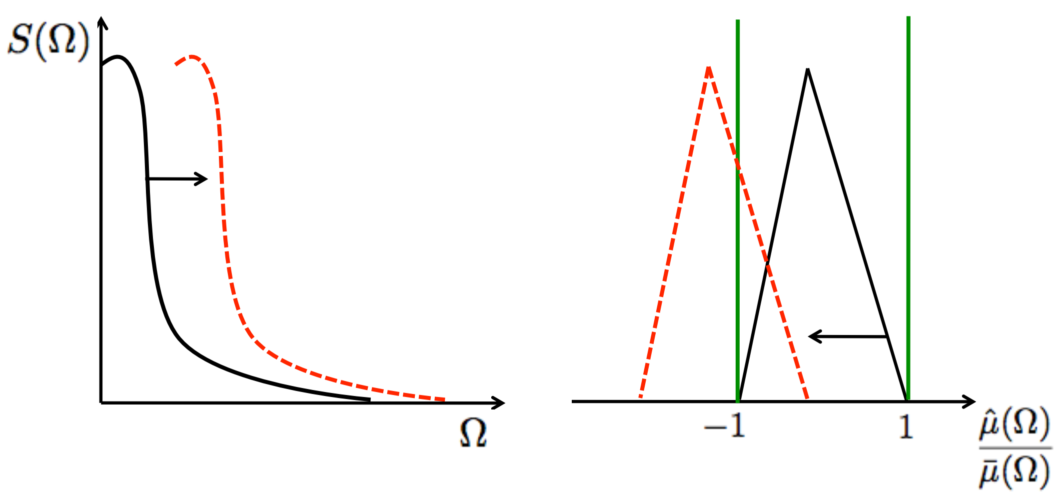

Figure 10: Left panel: Steady state distribution of uncertainty $S(\Omega)$. Right panel: Steady state distribution of markup forecasts, normalized by Ss band.

When there is no uncertainty shock ($n = 0$), the uncertainty distribution $S(\Omega)$ always remains fixed and the distribution of markup forecasts also stays fixed on impact but slowly moves to the left as firms incorporate the shock. When there is a correlated uncertainty shock ($n = 2, 4$), the uncertainty distribution shifts to the right to the dashed line. Higher uncertainty increases the optimal weight on new information and the markup distribution approaches the dashed line faster as n increases. This is a reduction in the learning effect since forecast errors converge faster to zero. Furthermore, since the weight is higher on new information and this information comes with noise, the frequency of price changes is increased even more. Therefore firms learn faster and also change their prices more often, decreasing the effect of the monetary shock.

Figure 11 shows the results from the uncertainty experiments. The first panel shows the impulse response of output. As expected, higher uncertainty decreases the effect of a monetary shock. The second and third panels plot the frequency of price changes measured as $\mathbb{E}_i i[\mathbb{1}_{\{|\Delta p|>0\}}]$ and the average forecast errors $\mathbb{E}_i[\hat{\mu}_t^i - \mu_t^i]$. The frequency increases with uncertainty (bringing higher flexibility) and the forecast errors decrease with uncertainty (measuring faster learning). The larger flexibility of the aggregate price level, given by faster learning and faster price adjustment, reduces the output effects.

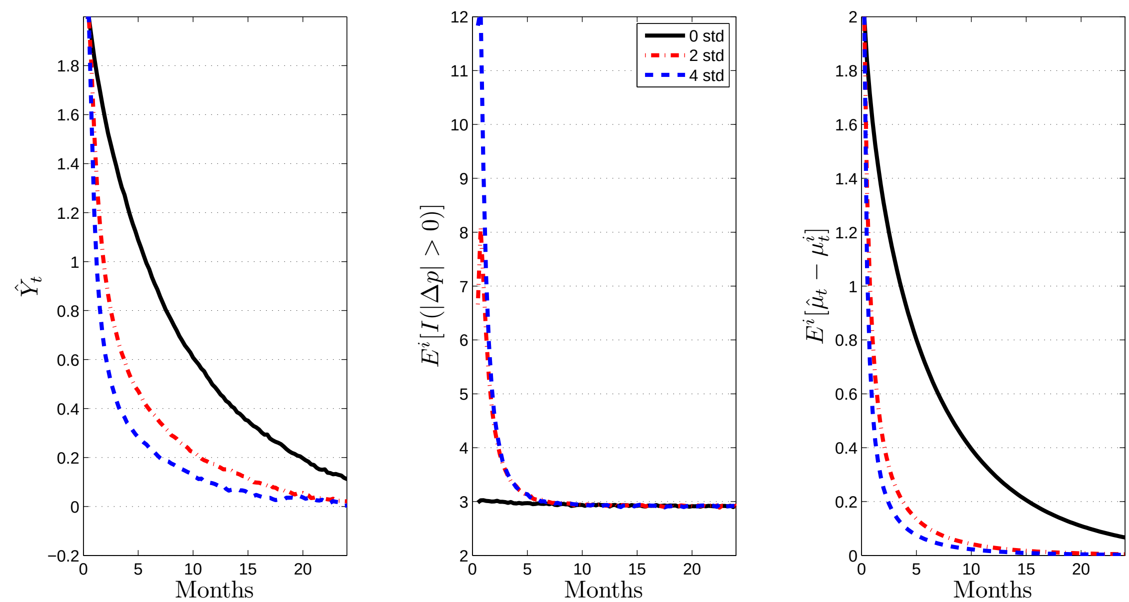

Figure 11: Responses to aggregate negative shock in markups of $\delta = 2\%$. Left panel: Impulse-response of output, Center panel: Frequency of price change, Right panel: Average forecast errors.

6 Conclusions

What are the effects of monetary policy when information frictions make firms uncertain about their environment? In normal times, information frictions delay the response of prices to monetary shocks and amplify output effects. In contrast, if uncertainty is high, the response of prices to shocks is accelerated and output effects are reduced. The reason is that forecasts adjust faster to the new environment because they weight new information more with respect to their priors.

We have explored the effect of firm uncertainty in price statistics and aggregate price flexibility. Our results have interesting monetary policy consequences, as they suggest that policy effectiveness may be weakened in times when it is most needed. We have remained agnostic about the source of uncertainty shocks. In principle, they could come from the monetary policy itself or from a completely orthogonal source. It would be interesting to further study the interactions between the source of uncertainty, for example idiosyncratic vs. aggregate, and price flexibility.

References

[1] **Álvarez, Burriel and Hernando (2010)** Price Setting Behaviour in Spain: Evidence From Micro PPI Data. *Managerial and Decision Economics*, 31

[2] **Álvarez and Lippi (2012)** Price Setting With Menu Costs For Multi-product Firms. NBER Working Paper 17923.

[3] **Álvarez, Lippi and Paciello (2010)** Optimal price setting with observation and menu costs. *Quarterly Journal of Economics*, 126(4)

[4] **Arellano, Bai and Kehoe (2012)** Financial Markets and Fluctuations in Volatility. Mimeo.

[5] **Bachmann and Moscarini (2011)** Business Cycles and Endogenous Uncertainty.

[6] **Bachmann and Bayer (2011)** Uncertainty Business Cycles - Really?, NBER Working Paper 16861.

[7] **Blinder, Canetti, Lebow and Rudd (1998)** Asking About Prices: A New Approach to Understanding Price Stickiness. Russell Sage Foundation.

[8] **Bloom (2009)** The Impact of Uncertainty Shocks. *Econometrica*, 77.

[9] **Bloom, Bond and van Reenen (2006)** Uncertainty and Investment Dynamics. *Review of Economic Studies*, 74(2).

[10] **Bonomo, Carvalho and Garcia (2010)** State-Dependent Pricing under Infrequent Information: A Unified Framework. Staff Reports 455, Federal Reserve Bank of New York.

[11] **Calvo (1983)** Staggered prices in a utility-maximizing framework, *Journal of Monetary Economics*, Vol. 12(3).

[12] **Campbell and Eden (2010)** Rigid Prices: Evidence from U.S. Scanner Data. Federal Reserve Bank of Chicago.

[13] **Caplin and Spulber (1987)** Menu Costs and the Neutrality of Money, *Quarterly Journal of Economics.*

[14] **Caplin and Leahy (1991)** State-Dependent Pricing and the Dynamics of Money and Output, *Quarterly Journal of Economics*, Vol. 106(3).

[15] **Carvalho (2006)** Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks. *The B.E. Journal of Macroeconomics*, Vol.2, No. 1.

[16] **Cavallo (2010)** Scraped Data and Sticky Prices: Frequency, Hazards, and Synchronization. Mimeo.

[17] **Christiano, Eichenbaum and Evans (2005)** Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *[Journal of Political Economy*,Vol. 113, No. 1.

[18] **Cogley and Sargent (2005)** The conquest of US inflation: Learning and robustness to model uncertainty, *Review of Economic Dynamics*, Vol. 8(2)

[19] **Coibion and Gorodnichenko (2012)** What Can Survey Forecasts Tell Us about Information Rigidities?, Journal of Political Economy, 120(1).

[20] **Dhyne, *et. al.* (2006)** Price Changes in the Euro Area and the United States: Some Facts from Individual Consumer Price Data. *Journal of Economic Perspectives*, 20.

[21] **Eden and Jaremski (2009)** Rigidity, Discreteness and Dispersion in Chain Prices. Mimeo.

[22] **Fabiani *et.al.* (2005)** The Pricing Behavior of Firms in the Euro Area, New Survey Evidence. *Eurosystem Inflation Persistence Network*, Working Paper 535.

[23] **Gali and Gertler (1999)** Inflation Dynamics: A Structural Econometric Approach, *Journal of Monetary Economics*, 44(2).

[24] **Gertler and Leahy (2008)** A Phillips Curve with an Ss Foundation. *Journal of Political Economy* 116(3).

[25] **Golosov and Lucas (2007)** Menu Costs and Phillips Curves. *Journal of Political Economy* Vol. 115.

[26] **Gourio (2008)** Estimating firm-level risk. Boston University, Mimeo.

[27] **Jovanovic (1979)** Job Matching and the Theory of Turnover, **Journal of Political Economy**, Vol. 87(5).

[28] **Karatzas and Shreve (2001)** Brownian Motion and Stochastic Calculus. Second Edition.

[29] **Kehoe and Midrigan (2012)** Prices Are Sticky After All. Federal Reserve Bank of Minneapolis Research Department Staff Report 413.

[30] **Kiefer (1988)** Economic Duration Data and Hazard Functions, *Journal of Economic Literature*, 26.

[31] **Klenow and Kryvtsov (2008)** State-Dependent or Time-Dependent Pricing: Does It Matter For Recent U.S. Inflation?, *Quarterly Journal of Economics*, 123.

[32] **Klenow and Malin (2010)** Microeconomic Evidence on Price-Setting, *Handbook of Monetary Economics*, 3.

[33] **Kolkiewicz (2002)** Pricing and hedging more general double-barrier options. Journal of Computational Finance 5(3).

[34] **Levy, Bergen, Dutta and Venable (1997)** The Magnitude of Menu Costs: Direct Evidence from Large U.S. Supermarket Chains, Quarterly Journal of Economics, 112(3).

[35] **Maćkowiak and Wiederholt (2009)** Optimal Sticky Prices under Rational Inattention, *American Economic Review*, 99.

[36] **Midrigan (2012)** Menu Costs, Multi-Product Firms, and Aggregate Fluctuations. *Econometrica.* (Forthcoming)

[37] **Nakamura and Steinsson (2008)** Five Facts about Prices: A Reevaluation of Menu Cost Models. *Quarterly Journal of Economics*, 123(4).

[38] **Øksendal (2002)** Stochastic Differential Equations: An introduction with Applications. Sixth Edition.

[39] **Orlik and Veldkamp (2012)** Origins of Uncertainty Shocks. Mimeo.

[40] **Schaal (2012)** Uncertainty, Productivity and Unemployment in the Great Recession, Mimeo.

[41] **Stokey (2009)** The Economics of Inaction, *Princeton University Press.*

[42] **Taylor (1979)** Estimation and control of an econometric model with rational expectations, *Econometrica*, 47.

[43] **Vavra (2012)** Inflation Dynamics and Time-Varying Volatility: New Evidence and an SS Interpretation. Mimeo.

[44] **Veldkamp (2012)** Information Frictions in Macroeconomics and Finance. *Princeton University Press.*

[45] **Woodford (2008)** Information-Constrained State-Dependent Pricing, Mimeo.

[46] **Zbaracki, Ritson, Levy, Dutta and Bergen (2004)** Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets, *Review of Economics and Statistics*, 86(2).