

NBER WORKING PAPER SERIES

WHAT CAN MEASURED BELIEFS TELL US ABOUT MONETARY NON-NEUTRALITY?

Hassan Afrouzi
Joel P. Flynn
Choongryul Yang

Working Paper 32541
<http://www.nber.org/papers/w32541>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2024, Revised January 2026

This paper was formerly circulated under the title: “Selection in Information Acquisition and Monetary Non-Neutrality.” We are grateful to Olivier Coibion and Yuriy Gorodnichenko for their invaluable feedback and for sharing data. We also thank Fernando Alvarez (discussant), Saroj Bhattarai, Mark Dean, Mishel Ghassibe (discussant), Zhen Huo (discussant), Chen Lian, Davide Melcangi (discussant), Giuseppe Moscarini, George Nikolakoudis, Karthik Sastry, Luminita Stevens (discussant), and Michael Woodford as well as seminar participants at Columbia University, University of Notre Dame, LMU Munich, Princeton University, University of Texas at Austin, the Federal Reserve Banks of Atlanta, Cleveland, Richmond, San Francisco, and St. Louis, the Federal Reserve Board, Keio University, University of Tokyo, Kyoto University, the Bank of Japan, the Salento Macro Meetings, the 12th Banco de Portugal Conference on Monetary Economics, the 3rd Workshop on Firm Heterogeneity and Macroeconomics, the 2024 Stanford Institute for Theoretical Economics Conference, and the NBER Monetary Economics conference for thoughtful comments and suggestions. We thank the Center for Advancement of Data and Research in Economics at the Federal Reserve Bank of Kansas City for providing computational resources that contributed to the results reported in this paper. The views expressed herein are those of the authors and do not reflect the views of the Federal Reserve Board, any person associated with the Federal Reserve System, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Hassan Afrouzi, Joel P. Flynn, and Choongryul Yang. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

What Can Measured Beliefs Tell Us About Monetary Non-Neutrality?

Hassan Afrouzi, Joel P. Flynn, and Choongryul Yang

NBER Working Paper No. 32541

June 2024, Revised January 2026

JEL No. E31, E32, E71

ABSTRACT

This paper studies how to use microeconomic data on beliefs and prices to identify the relative contributions of pricing and information frictions to monetary non-neutrality. In a canonical general equilibrium model with both pricing frictions and information acquisition, we analytically characterize how these frictions contribute to non-neutrality. Exploiting this characterization, we show that data on the cross-sectional distributions of uncertainty and pricing durations are necessary and sufficient to identify non-neutrality. Implementing our approach in survey data, we find that: (i) information and pricing frictions are approximately equally important, and (ii) assuming exogenous information would overstate non-neutrality by 60%.

Hassan Afrouzi
Columbia University
Department of Economics
and NBER
ha2475@columbia.edu

Choongryul Yang
Federal Reserve Board of Governors
cryang1224@gmail.com

Joel P. Flynn
Yale University
Department of Economics
and NBER
joel.flynn@yale.edu

1 Introduction

There are two classical hypotheses for why changes in monetary policy have effects on real aggregate outcomes. The first is that firms are subject to *pricing frictions*: even if they would like to change their price, they either cannot do so (as in time-dependent pricing frameworks, see *e.g.*, [Taylor, 1979](#), [Calvo, 1983](#)) or they would have to pay a cost to do so (as in state-dependent pricing frameworks, see *e.g.*, [Sheshinski and Weiss, 1977](#)). The second is that firms are subject to *information frictions*: firms may be perfectly able to adjust their prices, but they may not know the *ex post* correct price to charge ([Lucas, 1972](#), [Woodford, 2003](#)). Understanding the relative contributions of pricing and information frictions to monetary non-neutrality is critical in light of the [Lucas \(1976\)](#) critique: if we do not know why money is non-neutral, then we cannot know how changes in economic circumstances or the conduct and communication of monetary policy will affect non-neutrality.

In the context of the pricing frictions hypothesis, empirical work documents that firms' prices indeed appear to change only infrequently (see *e.g.*, [Bils and Klenow, 2004](#), [Nakamura and Steinsson, 2008](#)), consistent with the notion that firms are subject to pricing frictions. However, this observation alone does not suffice for the presence of substantial monetary non-neutrality ([Caplin and Spulber, 1987](#), [Golosov and Lucas, 2007](#)). For this reason, a seminal body of work studies how to discipline the extent of monetary non-neutrality using such microeconomic data on firms' pricing decisions. For time-dependent models, [Carvalho and Schwartzman \(2015\)](#) show that it suffices to measure the average duration of pricing spells across firms. For state-dependent models, [Alvarez, Le Bihan, and Lippi \(2016\)](#) show that the cross-sectional distribution of price changes—in particular, the frequency of price changes and their kurtosis—pin down the extent of monetary non-neutrality.

To study the information frictions hypothesis, a burgeoning literature has used surveys to understand imperfections in firms' expectations (see [Candia, Coibion, and Gorodnichenko, 2023](#), for a review). This survey evidence demonstrates that the average firm holds highly inaccurate and diffuse beliefs about its economic environment and that there is substantial heterogeneity in these beliefs across firms. However, while these facts are interesting on their own, how these surveys connect to and inform the resultant extent of monetary non-neutrality remains largely ambiguous in both qualitative and quantitative terms.

In this paper, we analyze how to use microeconomic data on firms' pricing decisions and beliefs to assess the relative importance of pricing frictions and information frictions in generating monetary non-neutrality. To do this, we develop a theoretical framework in which firms are subject to pricing frictions and information frictions. In this setting, we: (i) characterize firms' optimal

information acquisition behavior, (ii) explicitly aggregate this behavior to derive a formula for the real effects of a monetary shock that characterizes the roles of pricing and information frictions in generating monetary non-neutrality, and (iii) show how to measure the components of this formula using microeconomic data on firms' pricing behavior and beliefs. Using survey data on firms' expectations from New Zealand (Coibion, Gorodnichenko, and Kumar, 2018), we find that information frictions and pricing frictions are approximately equally important in generating monetary non-neutrality. Moreover, models with exogenous information would overstate the role of information frictions by a factor of two.

Model. We study a canonical general equilibrium monetary economy in which the microfoundations for households and monetary policy follow those of Golosov and Lucas (2007): a representative household has logarithmic preferences over real money balances, linear labor disutility, and constant relative risk aversion preferences over the consumption of a constant elasticity of substitution aggregate of the outputs of a continuum of monopolistically competitive firms. In the model, these firms face geometric Brownian productivity shocks and are subject to both pricing and information frictions. For our main theoretical analysis, firms face arbitrary time-dependent pricing frictions, nesting both the Calvo (1983) and the Taylor (1979) models. We later extend the model to feature state-dependent pricing frictions in the form of random menu costs (following Alvarez, Lippi, and Oskolkov, 2022), strategic complementarities in pricing (following Alvarez, Lippi, and Souganidis, 2023), and richer information structures. In this otherwise standard model, we allow firms to acquire any dynamic information structure about their marginal costs. Firms face information frictions in the form of a cost that is proportional to the flow of the information acquired, which is a canonical way of modeling information costs (see Maćkowiak, Matějka, and Wiederholt, 2023).

Theoretical Results. We begin our analysis by theoretically characterizing firms' optimal information acquisition and pricing decisions. We show that this takes a simple form: acquire information only when changing prices and acquire exactly enough information to reset posterior uncertainty about the optimal price to some state-independent level, U^* . Intuitively, while being better informed reduces the costs of achieving any given level of uncertainty in the future as you need to acquire less information, it does not affect the *marginal cost* of reduced uncertainty. Moreover, we show that: the optimal level of uncertainty is decreasing in the firm's demand elasticity (as this increases the losses from setting the wrong price); increasing in the volatility of marginal costs (as this reduces the value of information acquired today for future decisions); and ambiguously affected by price stickiness (as this both increases the value of information for this pricing spell and decreases the value of information for all future pricing spells).

A key implication of this result is that a firm's uncertainty is increasing in the duration of its pricing spell. This implies that price-setting firms are the least uncertain firms in the economy. We call this phenomenon *selection in information acquisition* as price-setting firms are the most informed in the cross-section at any given point in time. This differentiates our model relative to alternatives with exogenous informational frictions and nominal rigidities (as in [Nimark, 2008](#), [Angeletos and La'O, 2009](#)) or models of endogenous information acquisition without nominal rigidities (e.g., [Sims, 2003](#), [Maćkowiak and Wiederholt, 2009](#)): in both such cases, firms' uncertainty has no relationship with the duration of their pricing spell.

Next, we study the real effects of monetary shocks by characterizing the cumulative impulse response (CIR) of aggregate output to a one-time unexpected increase in nominal marginal costs of firms. Normalizing the shock size so that the impact response of output is 1 percent, we denote the CIR by $\mathcal{M}^b$ and derive a closed-form representation for it as:

$$\mathcal{M}^b = \bar{D} + \frac{U^*}{\sigma^2} \quad (1)$$

where $\bar{D}$ is the average duration of (ongoing) pricing spells as measured in a cross-section of firms, σ^2 is the variance of shocks to firms' idiosyncratic productivity, and U^* is the subjective uncertainty of *price-setting firms* about their marginal costs. In this formula, the first term is the usual one derived in models of time-dependent price stickiness (as in [Carvalho, 2006](#), [Carvalho and Schwartzman, 2015](#)), and captures the notion that (all else equal) monetary non-neutrality increases as firms' nominal prices are stickier for longer. The second term is new to our analysis and captures the lifetime lack of responsiveness of all firms in the economy in resetting their prices in light of the uncertainty that arises due to informational frictions. Intuitively, when price resetting firms are more uncertain, they respond to their current information to a lesser degree and so adjust their prices by less in response to a monetary shock. Moreover, when microeconomic volatility is higher, firms know that their old information is less likely to be useful as things will have since changed by a larger amount; this makes firms more responsive to their information and lowers the extent of monetary non-neutrality. Interestingly, as U^* moves ambiguously when price stickiness changes, increases in price stickiness have a theoretically ambiguous effect on monetary non-neutrality in the presence of endogenous information.

This result establishes that uncertainty amplifies the real effects of monetary shocks relative to a full-information benchmark. However, due to selection in information acquisition, looking at the data through the lens of an exogenous information model would systemically overstate the real effects of monetary shocks by relating them to the uncertainty of the *average* firm.

Our final theoretical results establish that the quantities that determine the CIR (as per Equation 1) can be estimated given cross-sectional data on firms’ uncertainty and the time since they last reset their price. Thus, survey data on these quantities are sufficient to identify the model. Moreover, we find that such data are necessary, in contrast to benchmark models in which firms have full information about their environment—it is well known that cross-sectional data on the distribution of firms’ price changes are sufficient for identifying the real effects of monetary shocks in such models (see *e.g.*, [Carvalho and Schwartzman, 2015](#), for pure time-dependence and [Alvarez, Le Bihan, and Lippi, 2016](#), for state-dependence with random menu costs). This can indeed be seen in Equation 1, as $U^* = 0$ under full information and the formula collapses to $\bar{D}$, which can be measured using data on price changes. However, we show that, in the presence of information costs that imply a positive degree of uncertainty for price-setters, *i.e.*, $U^* > 0$, data on the distribution of price changes *cannot* identify U^* . This is because the firm’s choice of information renders the distribution of price changes *invariant* to U^* .

Using Survey Data to Quantify the Model. Finally, we adopt a “micro-to-macro” approach of combining measured beliefs with the structure of the model to quantify the extent to which imperfect information and endogenous information acquisition matter for monetary non-neutrality.

First, by using a survey of New Zealand firms between Q4 2017 and Q2 2018 ([Coibion, Gorodnichenko, and Kumar, 2018](#), [Coibion, Gorodnichenko, Kumar, and Ryngaert, 2021](#)), we measure firms’ uncertainty about their optimal prices and the durations of their pricing spells.

Second, by applying the estimators for the CIR from the theory to the survey data, we find that pricing and information frictions contribute approximately equally to monetary non-neutrality. Moreover, ignoring the endogeneity of information acquisition would lead us to overstate the importance of information frictions by a factor of two. Thus, we argue that pricing frictions, information frictions, and the endogeneity of information acquisition are quantitatively important.

Finally, by using the firm’s first-order condition for its optimal uncertainty, we can derive and implement estimators of the effect of counterfactually increasing microeconomic volatility and price stickiness on the CIR. We find that greater microeconomic volatility significantly dampens the real effects of monetary policy. This is because the direct effect of reducing firms’ reliance on past information quantitatively dominates the indirect effect that firms optimally choose to be less informed in the face of this increase. We also find that greater price stickiness increases monetary non-neutrality, but by approximately 47% less than with full information. This is because we find that firms would become better informed in the face of increased stickiness. This highlights the importance of jointly modeling pricing and information frictions.

Extensions. In a final section, we generalize our analysis in three directions to probe the robustness of the findings derived in our baseline model. First, we introduce random menu costs that generate state-dependence in both pricing and information acquisition. Calibrating this model to match the fact from [Alvarez, Lippi, and Paciello \(2018\)](#) that firms review prices twice as frequently as they change them, we find that our baseline formula for the effects of incomplete information on the CIR is 99% accurate. This further highlights the value of our approach of modeling pricing frictions as time-dependent: it provides quantitatively accurate guidance while remaining analytically tractable. Second, we show that introducing strategic complementarities in pricing amplifies the effects of pricing frictions on the CIR but leaves the contribution of information frictions unchanged. Third, we generalize our CIR formula to consider richer information structures, including settings in which firms periodically receive free information (perhaps obtained from making other decisions).

Related Literature. Our key contribution is to show how microeconomic data on measured beliefs and pricing can be used to understand and identify the relative importance of information and pricing frictions for monetary non-neutrality. Our findings regarding the necessity and sufficiency of measured beliefs for quantifying the aggregate effects of monetary shocks align with a broader body of literature advocating for measuring attention (for a review, see [Caplin, 2016](#)).

This paper builds on and contributes to several strands of literature that study the real effects of monetary policy shocks under time-dependent price stickiness or informational frictions. First, on the nominal rigidities side, our work builds on the work of [Carvalho and Schwartzman \(2015\)](#) and the general equilibrium model of [Alvarez, Le Bihan, and Lippi \(2016\)](#), who study monetary non-neutrality in time- and state-dependent models, respectively. We contribute to this literature by introducing endogenous information acquisition into time-dependent models and showing that it has a quantitatively important effect on the real effects of monetary shocks.

Second, our work is also related to the literature on pricing models with no nominal rigidities but with either exogenous informational frictions ([Mankiw and Reis, 2002](#), [Woodford, 2003](#), [Nimark, 2008](#), [Angeletos and La'O, 2009](#)) or endogenous information acquisition with rational inattention (*e.g.*, [Sims, 2003](#), [Maćkowiak and Wiederholt, 2009](#)).¹ We build on this literature by investigating the real effects of monetary policy shocks in a unified framework featuring both information acquisition and nominal rigidities and find that their interaction is both qualitatively and quantitatively important. In this context, the most related works are [Woodford \(2009\)](#), [Stevens \(2020\)](#)

¹See also [Moscarini \(2004\)](#), [Sims \(2010\)](#), [Paciello and Wiederholt \(2014\)](#), [Maćkowiak, Matějka, and Wiederholt \(2018\)](#), [Afrouzi and Yang \(2021\)](#). We refer the reader to [Maćkowiak, Matějka, and Wiederholt \(2023\)](#) for a comprehensive review. Another notable contribution beyond the rational inattention framework is [Reis \(2006\)](#), which studies the optimal information acquisition of firms under fixed observation costs but no nominal rigidities.

and [Morales-Jiménez and Stevens \(2024\)](#), who micro-found state- and time-dependent nominal rigidities *through* rational inattention to the timing of pricing decisions. To focus on the degree of endogenous state-dependence, these papers assume that the “reference distribution” for the cost of information is the unconditional time-invariant one that emerges in the steady state, which is necessary to keep the firm’s problem tractable ([Morales-Jiménez and Stevens, 2024](#)). In this paper, it is necessary for us to keep track of these reference distributions as state variables for firms (*i.e.*, the priors of firms about their desired prices), which do vary endogenously over time because firms optimally acquire more information when their priors are more diffuse—a result that is integral to our finding selection in information acquisition and which has significant quantitative implications.

In a broader sense than rational inattention with Shannon entropy costs, our paper is also related to the work that studies nominal rigidities and information costs jointly (see [Gorodnichenko, 2008](#), [Yang, 2022](#), for quantitative papers on this topic). [Alvarez, Lippi, and Paciello \(2011, 2016\)](#), and [Bonomo, Carvalho, Garcia, Malta, and Rigato \(2023\)](#) study models with both menu costs and observational costs, where firms decide when they observe either idiosyncratic shocks or aggregate shocks by paying a fixed cost. In these models, firms can perfectly observe the underlying shocks whenever they pay the fixed cost. As a result, their prior information set at the time of information acquisition becomes irrelevant as they become fully aware of their marginal cost. In our framework, firms decide *how much* information they want to acquire. As a result, firms do not become fully aware of their marginal cost upon acquiring information. This relates their current prices to the full history of their past information sets, which is not *ex ante* trivial to characterize. Despite this, we show in our model that this infinite-dimensional history dependence is captured by a one-dimensional state variable: firms’ uncertainty about their optimal prices. Beyond the technical contribution, this is the novel component of our theory that gives rise to the quantitatively important subjective uncertainty term in Equation 1.²

Outline. The rest of the paper proceeds as follows. Section 2 introduces the model. Section 3 characterizes firm behavior in the model. Section 4 characterizes the roles of information and pricing frictions in generating monetary non-neutrality. Section 5 shows how to identify the role of these frictions using microeconomic data. Section 6 quantifies these roles using survey data. Section 7 studies how changes in microeconomic volatility and price stickiness affect monetary non-neutrality. Section 8 extends our baseline analysis to feature menu costs, strategic complementarities, and richer learning structures. Section 9 concludes.

²In fact, if firms were to become fully aware of their marginal costs upon acquiring information, this term would be zero and the formula would collapse back to the average duration of ongoing spells, $\bar{D}$.

2 A Monetary Economy with Pricing and Information Frictions

We study a general equilibrium monetary economy with endogenous information acquisition by firms that are subject to general, time-dependent pricing frictions. To make the role of information acquisition as clear as possible, the macroeconomic side of the model follows [Goloso and Lucas \(2007\)](#), [Alvarez and Lippi \(2014\)](#), and [Alvarez, Le Bihan, and Lippi \(2016\)](#). Our goal is to understand the relative contributions of pricing and information frictions to monetary non-neutrality.

2.1. Households

Primitives. Time is continuous and indexed by $t \in [0, \infty]$. A representative household has preferences over consumption C_t , real money balances M_t/P_t (where M_t is money and P_t is the price of consumption), and labor L_t given by:

$$\int_0^\infty e^{-rt} \left[\frac{C_t^{1-\gamma} - 1}{1-\gamma} + \log\left(\frac{M_t}{P_t}\right) - \alpha L_t \right] dt \quad (2)$$

where $r > 0$ is the discount rate, $\gamma^{-1} > 0$ is the elasticity of intertemporal substitution, and $\alpha > 0$ indexes the extent of labor disutility. Consumption is a constant elasticity of substitution aggregate of a continuum of varieties, indexed by $i \in [0, 1]$:

$$C_t = \left(\int_0^1 A_{i,t}^{\frac{1}{\eta}} C_{i,t}^{\frac{\eta-1}{\eta}} di \right)^{\frac{\eta}{\eta-1}} \quad (3)$$

where $\eta > 1$ is the elasticity of substitution between varieties and $A_{i,t}$ is a variety-specific taste shock. The household can also trade a risk-free nominal bond in zero net supply that pays a nominal interest rate of R_t . Thus, the household's lifetime budget constraint is:

$$M_0 + \int_0^\infty \exp\left(-\int_0^t R_s ds\right) \left[w_t L_t + \int_0^1 \Pi_{i,t} di - \int_0^1 P_{i,t} C_{i,t} di - R_t M_t \right] dt = 0 \quad (4)$$

where w_t is the wage and $P_{i,t}$ and $\Pi_{i,t}$ are the price and profits of variety i at time t , respectively. The money supply is constant and equals $\bar{M}$. Later, we will shock $\bar{M}$ to $\bar{M} + \delta$ for $\delta \in \mathbb{R}$.

Optimality Conditions. As is well-known, this setup implies the following optimality conditions, which reduce understanding aggregate dynamics to understanding the price-setting decisions of each firm in the economy. First, the household's demand for consumption variety i at time t is:

$$C_{i,t} = A_{i,t} C_t \left(\frac{P_{i,t}}{P_t} \right)^{-\eta} \quad (5)$$

where the aggregate price index is given by:

$$P_t = \left(\int_0^1 A_{i,t} P_{i,t}^{1-\eta} \right)^{\frac{1}{1-\eta}} \quad (6)$$

and the intratemporal Euler equation is given by:

$$C_t^{-\gamma} = \alpha \frac{P_t}{w_t} \quad (7)$$

Moreover, nominal wages and the interest rate, under a constant path for the money supply are:³

$$w_t = \alpha r M_t \quad \text{and} \quad R_t = r \quad (8)$$

2.2. Firms' Production, Pricing, and Profits

Production Technology. Each variety $i \in [0, 1]$ is produced by a firm with the same index. Firms produce output $Y_{i,t}$ according to the linear production technology:

$$Y_{i,t} = \frac{1}{Z_{i,t}} L_{i,t} \quad (9)$$

where $L_{i,t}$ is the labor input and $Z_{i,t}$ is a marginal cost shock to the firm. As in [Alvarez and Lippi \(2014\)](#), we make the simplifying assumption that $Z_{i,t}^{1-\eta} A_{i,t} = 1$, which is irrelevant for our subsequent approximation of firms' profits but ensures that the firm-size distribution is well-behaved under flexible prices.⁴ Moreover, we assume that:

$$Z_{i,t} = \exp\{\sigma W_{i,t}\} \quad (10)$$

where $\{W_{i,t}\}_{t \geq 0}$ is a standard Brownian motion that is independent across $i \in [0, 1]$.

Pricing Frictions. Firms are price setters and subject to time-dependent pricing frictions. Formally, price change opportunities for firm i are governed by the counting process $N_{i,t}$ which is independent across $i \in [0, 1]$. We assume that the distribution of times of the arrival of price reset opportunities ($dN_{i,t} = 1$) is exogenously given by a cumulative distribution function (CDF) G with finite first and second moments. We moreover assume that G admits a density g and define its hazard rate as $\theta(h) \equiv g(h)/(1 - G(h))$. This general model of time-dependent pricing nests important benchmarks, including [Calvo \(1983\)](#) pricing:

Example 1 (Calvo Pricing). *Price reset opportunities arise at a constant rate $\theta(h) = \theta$.* $\triangle$

A more general formulation, in which G does not admit a density, also allows for [Taylor \(1979\)](#) pricing, under which firms reset their prices periodically. All of our results hold in this case:

Example 2 (Taylor Pricing). *Price reset opportunities arise every $k \in \mathbb{R}_+$ periods and so $g = \delta_k$.* $\triangle$

³All of our analysis extends to a case with wage rigidity, i.e., $w_t = \alpha r M_t^\chi$ for some $\chi \geq 0$. In this case, a δ percent change in money would change firms' marginal costs (and optimal prices) by $\chi\delta$ percent. Rigidity could matter for the absolute level of monetary non-neutrality, but it does not matter for the relative contributions of pricing frictions and information frictions to monetary non-neutrality.

⁴None of our results rely on this assumption. However, as we are intentionally following the setup of [Alvarez and Lippi \(2014\)](#), we maintain it for maximum comparability to the existing literature.

Our interpretation of time-dependent pricing frictions is that they represent a tractable stand-in for three relevant and well-established features of firm decisionmaking: (i) firms can only periodically revise contracts (Taylor, 1979), (ii) firms only update decisions periodically (Reis, 2006), and (iii) they are a stand-in for menu costs (Auclert, Rigato, Rognlie, and Straub, 2024). In Section 8 and Appendix C, we explicitly extend this framework to introduce menu costs.

Approximating Firms' Profits. Given their price at a given time, firms commit to hiring enough labor to meet demand at their given price. Define the (log) optimal price of the firm as $q_{i,t} \equiv \log\left(\frac{\eta}{\eta-1} w_t Z_{i,t}\right)$ and the (log) price of the firm as $p_{i,t} \equiv \log P_{i,t}$. Approximating the firm's profit function to second-order around $p_{i,t} = q_{i,t}$, as is well-known, the firm's profits are:

$$\mathcal{L}(p_{i,t}, q_{i,t}) = -\frac{B}{2} (p_{i,t} - q_{i,t})^2 \quad (11)$$

where $B = \eta(\eta - 1)$. Thus, when demand is more elastic, the losses from mispricing are larger. In Section 8 and Appendix D, we extend our framework to feature strategic complementarity or substitutability in price-setting.

2.3. Firms' Costly Information Acquisition

So far, we have followed the textbook model of firm pricing in general equilibrium. We now introduce the novel feature of our analysis: endogenous information acquisition. We assume firms are aware of their price change opportunities, *i.e.*, they observe the process $N_{i,t}$, but cannot directly observe the shock to their marginal costs and acquire information about this process subject to a cost.

Formally, given the joint measure for the process $\{(q_{i,t}, N_{i,t}) : t \geq 0\}$, firm i chooses a joint measure for $\{(q_{i,t}, N_{i,t}, s_{i,t}) : t \geq 0\}$, observes realizations of the process $s_{i,t}$ along with $N_{i,t}$ and makes decisions at time t given the information set $S_i^t \equiv \{(s_{i,h}, N_{i,h}) : h \leq t\} \in \mathcal{S}^t$. We assume that the cost of acquiring information is given by mutual information à la Sims (2003). Formally, given an information structure $\{S_i^t : t \geq 0\}$, we measure the amount of information acquired by firm i up to time t as the mutual information between the history of the optimal price, $\mathcal{Q}_i^t \equiv \{q_{i,h} : h \leq t\}$, and the information set S_i^t . Thus, letting $\mu_{i,t}^{\mathcal{Q}S}$ be the measure for the process $\{(q_{i,h}, s_{i,h}, N_{i,h}) : h \leq t\}$, and $\mu_{i,t}^{\mathcal{Q}} \otimes \mu_{i,t}^S$ be the product measure induced by $\mu_{i,t}^{\mathcal{Q}}$, mutual information is defined by:

$$\mathbb{I}(\mu_{i,t}^{\mathcal{Q}S}) \equiv \int \log \left(\frac{d\mu_{i,t}^{\mathcal{Q}S}}{d(\mu_{i,t}^{\mathcal{Q}} \otimes \mu_{i,t}^S)} \right) d\mu_{i,t}^{\mathcal{Q}S} \quad (12)$$

where the term inside the logarithm is the Radon-Nikodym derivative between the joint measure $\mu_{i,t}^{\mathcal{Q}S}$ and the product measure $\mu_{i,t}^{\mathcal{Q}} \otimes \mu_{i,t}^S$. We also define the amount of information processed in the

time interval $(h, t]$ as $\mathbb{I}(\mu_{i,t}^{\mathcal{Q}S}) - \mathbb{I}(\mu_{i,h}^{\mathcal{Q}S})$ and let $d\mathbb{I}(\mu_{i,t}^{\mathcal{Q}S})$ denote the differential form of this object—*i.e.*, the amount of information processed at the “instant” t .

As is standard in the rational inattention literature (see [Maćkowiak, Matějka, and Wiederholt, 2023](#)), we assume that the cost of information is linear in the information that the firm acquires, with scaling parameter $\omega > 0$. That is, the cost of the information flow is $\omega d\mathbb{I}$. In [Section 8](#) and [Appendix E](#), we extend our framework to accommodate more general learning procedures.

2.4. The Firm's Problem

Putting together the firm's profit function and its information costs, we obtain that the firm's problem is to choose a pricing and information policy to maximize the expected discounted value of its profits net of its information costs. Formally, a pricing policy for the firm is a map that returns the price that the firm charges after each history at each time $\hat{p}_{i,t} : S_i^t \rightarrow \mathbb{R}$. A pricing policy is feasible if it is constant whenever the firm does not receive a price change opportunity. The firm chooses its information policy $\mu_{i,t}^{\mathcal{Q}S}$ along with a feasible pricing policy to maximize its expected discounted profits net of information costs:

$$\sup_{\{\mu_{i,t}^{\mathcal{Q}S}, \hat{p}_{i,t}\}_{t \geq 0}} \mathbb{E} \left[\int_0^\infty e^{-rt} \left(-\frac{B}{2} (p_{i,t} - q_{i,t})^2 dt - \omega d\mathbb{I}_t \right) \middle| S_i^0 \right] \quad (13)$$

where $dp_{i,t} = (\hat{p}_{i,t} - p_{i,t})dN_{i,t}$, *i.e.*, prices change to $\hat{p}$ only when a pricing opportunity arises. Moreover, these pricing opportunities follow the given hazard θ .

2.5. Equilibrium

An equilibrium is a path for all endogenous variables such that the household maximizes its expected utility, the firm maximizes its profits, and all markets clear.

Definition 1 (Equilibrium). *An equilibrium is a sequence of random variables:*

$$\left\{ C_t, P_t, L_t, R_t, w_t, \{\Pi_{i,t}, P_{i,t}, C_{i,t}, L_{i,t}, Y_{i,t}, q_{i,t}\}_{i \in [0,1]} \right\}_{t \in \mathbb{R}_+} \quad (14)$$

and a collection of policy functions $(\mu_{i,t}^{\mathcal{Q}S}, \hat{p}_{i,t})_{i \in [0,1], t \in \mathbb{R}_+}$ such that:

1. *The policy functions solve [Equation 13](#)*
2. *Production occurs according to [Equation 9](#)*
3. *The household optimizes its expected discounted utility ([Equation 2](#)) subject to its intertemporal budget constraint ([Equation 4](#)) and so [Equations 5, 6, 7, 8](#) hold.*
4. *The markets for labor, goods, bonds, and money clear.*

In the following sections, we will study equilibrium firm policies and characterize the resulting implications for monetary non-neutrality.

3 Firms' Information Acquisition

We now solve for firms' optimal pricing and information strategies. Optimal information policies take a striking form: only acquire information when resetting prices and *always* acquire exactly enough information to reset uncertainty about optimal prices to some fixed, state-invariant level.

3.1. Optimal Information Acquisition

We begin by fully characterizing firms' optimal information and pricing policies. Because firms' marginal costs follow a Martingale, they simply set prices equal to the expected optimal price:

$$p_{i,t} = \mathbb{E}[q_{i,t}|S_i^t] \quad (15)$$

Toward characterizing the optimal information policy, define firm i 's posterior uncertainty about its optimal reset price at time t as $U_{i,t} = \mathbb{V}[q_{i,t}|S_i^t]$. We let $U_{i,t-}$ denote the corresponding prior uncertainty about $q_{i,t}$ at time t . The following result characterizes optimal information acquisition.

Theorem 1 (Optimal Dynamic Information Policy). *The firm only acquires information when it changes its price. Moreover, there exists a threshold level of uncertainty U^* such that:*

1. *If $U_{i,t-} \leq U^*$, then the firm acquires no information and $U_{i,t} = U_{i,t-}$.*
2. *If $U_{i,t-} > U^*$, then the firm acquires a Gaussian signal of its optimal price such that its posterior uncertainty is $U_{i,t} = U^*$.*

Furthermore, U^* is the unique solution to:

$$\underbrace{\frac{\omega}{U^*} - \mathbb{E}^h \left[e^{-rh} \frac{\omega}{U^* + \sigma^2 h} \right]}_{\text{marginal cost of information}} = \underbrace{B \left(\frac{1 - \mathbb{E}^h[e^{-rh}]}{r} \right)}_{\text{marginal benefit of information}} \quad (16)$$

Proof. See Appendix A.1. ■

We prove this result in three steps. First, we show that firms should only wish to acquire information when they change their prices. This is because, by acquiring information only when it is used, the firm pushes information acquisition further into the future and so never acquires information that becomes stale.

Second, we show that the firm should always acquire Gaussian signals when it resets its prices. Intuitively, as the firm sets $p_{i,t} = \mathbb{E}_{i,t}[q_{i,t}|S_i^t]$, the firm's expected per period loss until it resets its price is proportional to $\mathbb{V}[q_{i,t}|S_i^t]$. Thus, the firm's payoffs depend only on a sequence of conditional variances of a Gaussian random variable. Under mutual information, the cheapest way to achieve such a sequence is with a sequence of signals that maximizes entropy. The highest entropy distribution for any expected variance-covariance matrix is the Gaussian one. Combining this observation

with the fact that the best predictor of future optimal prices is the current optimal price, we obtain that the firm should always acquire a Gaussian signal of its current optimal price.

Third, we characterize the optimal noise in signals. To do this, we observe that the firm's posterior variance about optimal reset prices is a sufficient statistic for the firm's dynamic problem. Thus, letting $U_{i,t-}$ be the firm i 's prior uncertainty in period t , we have that firms solve:

$$V(U_{i,t-}) = \max_{U_{i,t} \leq U_{i,t-}} -U_{i,t} \frac{B}{2} \mathbb{E}^h \left[\int_0^h e^{-r\tau} d\tau \right] + \mathbb{E}^h \left[e^{-rh} V(U_{i,t} + \sigma^2 h) \right] + \frac{\omega}{2} \ln \left(\frac{U_{i,t}}{U_{i,t-}} \right) \quad (17)$$

The first term is the expected loss from mispricing, which is $U_{i,t} \times \frac{B}{2}$ per period, for the expected discounted duration of the pricing spell. The second term is the continuation value. If you reset your price in h periods, uncertainty at that point is your posterior uncertainty today plus the volatility of the ideal price multiplied by h . These two terms give rise to a trade-off: information today is more valuable the more likely it is that you reset your price soon, because you will have better information the next time you set your price, but losses from mispricing are lower if you reset your prices sooner. The final term is simply the cost of achieving a given level of posterior uncertainty given the mutual information form of costs. These trade-offs yield the claimed first-order condition.

Importantly, the optimal level of *posterior* uncertainty does not depend on *prior* uncertainty when firms come to reset prices. Intuitively, having better prior information reduces the cost of obtaining better posterior information. However, under mutual information, it does not change the *marginal cost* of better information, and so the optimal policy is invariant to $U_{i,t-}$.

3.2. The Economic Forces That Shape Optimal Uncertainty

We now study how changes in the economic environment affect the optimal level of uncertainty.

Corollary 1 (Comparative Statics). *The optimal level of uncertainty upon resetting the price, U^* , is:*


1. *Decreasing in the price elasticity of demand, η*
2. *Increasing in the cost of information, ω*
3. *Increasing in the volatility of marginal costs, σ^2*
4. *Increasing in the discount rate, r .*

Changes in the distribution of price reset opportunities, G , in the sense of first-order stochastic dominance (FOSD), have an ambiguous effect on U^ .*

Proof. See Appendix A.2 ■

We illustrate these comparative statics in Figure 1. Intuitively, a greater price elasticity of demand increases the profit losses from mispricing and leads firms to acquire more precise information. Increased information costs increase optimal uncertainty for obvious reasons. Moreover, when

Figure 1: Comparative Statics of Optimal Reset Uncertainty in Model Parameters



marginal costs become more volatile, it becomes more expensive to target a given level of uncertainty, and the benefits do not change. Thus, when marginal cost volatility increases, so too does optimal uncertainty. When the discount rate increases, future losses from mispricing become smaller, and the value of information for future decisions is smaller. Thus, higher discount rates lead to greater uncertainty.

The most surprising aspect of this result is that stickier prices could cause a firm to wish to be less informed about its optimal price. Changes in the flexibility of prices have ambiguous impacts because of two countervailing effects. On the one hand, as price reset opportunities become less frequent, the value of information until you next reset prices is greater because you keep your price fixed based on this information for a longer period of time. On the other hand, when price adjustment is less frequent, information acquired today is less valuable for future price resetting opportunities because marginal costs are likely to have changed by more when you next come to reset your prices. Which of these effects dominates depends on the other parameters of the problem, and the total effect of price flexibility on optimal uncertainty is ambiguous.

Special Cases and Bounds on Uncertainty. To illustrate these results, it is informative to consider the special case of Taylor pricing, in which optimal uncertainty can be solved in closed form.

Corollary 2 (Optimal Uncertainty Under Taylor Pricing). *Under Taylor pricing, i.e., firms reset prices every $k \in \mathbb{R}_+$ periods, optimal reset uncertainty is given by:*

$$U^* = \frac{-\left(B \frac{1-e^{-rk}}{r} \sigma^2 k - \omega (1 - e^{-rk})\right) + \sqrt{\left(B \frac{1-e^{-rk}}{r} \sigma^2 k - \omega (1 - e^{-rk})\right)^2 + 4B \frac{1-e^{-rk}}{r} \omega \sigma^2 k}}{2B \frac{1-e^{-rk}}{r}} \quad (18)$$

Moreover, in the special case in which discounting is zero, we have that:

$$\lim_{r \rightarrow 0} U^* = -\frac{\sigma^2 k}{2} + \sqrt{\left(\frac{\sigma^2 k}{2}\right)^2 + \omega \frac{\sigma^2}{B}} \quad (19)$$

Proof. See Appendix A.3. ■

In the special case of no discounting, the general comparative statics from Corollary 1 are particularly simple to observe: increases in ω and σ^2 and decreases in B increase U^* . Moreover, in the Taylor special case with no discounting, we observe that U^* is decreasing in k . This is because the effect of using information for longer dominates the effect that information acquired today is less useful the next time that the firm resets its price.⁵

Finally, even when optimal uncertainty does not admit an explicit solution, tight bounds on its value can be attained by considering the special limit cases in which marginal costs are infinitely volatile and marginal costs are constant over time. In these cases, we can solve for the optimal level of uncertainty in closed form. As per our earlier comparative statics, these cases also provide upper and lower bounds on firms' optimal uncertainty.

Corollary 3 (Special Cases and Bounds for Optimal Uncertainty). *In the limit of infinite volatility, optimal reset uncertainty is:*

$$\lim_{\sigma^2 \rightarrow \infty} U^* = \frac{\omega r}{B} \frac{1}{1 - \mathbb{E}^h[e^{-rh}]} \equiv U^{Max} \quad (20)$$

In the limit of zero volatility, optimal reset uncertainty is:

$$\lim_{\sigma^2 \rightarrow 0} U^* = \frac{\omega r}{B} \equiv U^{Min} \quad (21)$$

Moreover, any optimal reset uncertainty is such that $U^{Min} \leq U^* \leq U^{Max}$.

Proof. Immediate from Theorem 1 and Corollary 1. ■

⁵This domination is a quantitative consequence of the functional form of Taylor pricing. Indeed, it is simple to construct examples under which prices become stickier, but the firm's optimal uncertainty increases (because today's information is less useful for future decisions). Suppose that $h = \varepsilon$ with probability ψ and $h = k$ with probability $1 - \psi$. The basic Taylor case corresponds to $\varepsilon = 0$. Elementary algebra shows that $U^*(\varepsilon)$ is strictly increasing on some non-empty interval $I = (0, \bar{\varepsilon})$ (in the $r \rightarrow 0$ limit). That is, prices become stickier in the sense of first-order stochastic dominance, but optimal uncertainty increases.

Intuitively, when marginal costs are infinitely volatile, information acquired today has no value in making future price-setting decisions because the current state of marginal costs is completely uninformative about the future state of marginal costs. In this case, as price adjustment becomes more frequent, firms' optimal uncertainty increases. Because information today has no continuation value, the only effect of more frequent price adjustments is that losses from mispricing based on information today occur for fewer periods. This makes information today less valuable, increasing the optimal level of uncertainty. As this case minimizes the continuation value of information, it also places an upper bound on the optimal uncertainty that a firm will choose.

Conversely, when marginal costs are close to constant, information today is equally useful today as it will be when the firm resets prices. Thus, the frequency of price adjustment is irrelevant to optimal uncertainty. As this case maximizes the continuation value of information, this case places a lower bound on firms' optimal uncertainty.

3.3. Selection and Uncertainty

Our model of endogenous information acquisition implies an important property: firms that are setting prices are the least uncertain. An important implication of this fact is that it is not average uncertainty that is relevant for the price-setting decisions of firms, but rather the optimal reset level of uncertainty. We call this phenomenon *selection in information acquisition*: it is the price-setting firm whose uncertainty matters, and as these are the firms that most recently acquired information, they are the least uncertain firms.

Corollary 4 (Uncertainty and Time Since Changing Price). *Consider a firm i at time t that changed its price h periods ago. The firm's uncertainty about its optimal price follows:*

$$U_{i,t} = U^* + \sigma^2 h \tag{22}$$

Proof. See Appendix A.4. ■

This relationship between a firm's uncertainty and the duration of its pricing spell distinguishes our theory from models with exogenous information processing capacity or Gaussian signals with constant precision. This is because, in models with Gaussian signals or constant capacity, the firm's beliefs follow a Kalman-Bucy filter in which $U_{i,t}$ converges to a constant. Thus, under either model, the firm's level of uncertainty is constant and does not depend on the time since the firm reset its price. As we will shortly see, the fact that the extent of firms' uncertainty depends on the duration of their pricing spell has important qualitative implications for monetary non-neutrality.

4 How Pricing and Information Frictions Affect Monetary Non-Neutrality

Having characterized firms' optimal dynamic information policies, we now explore the implications of endogenous information acquisition for the propagation of monetary shocks. We find that uncertainty affects the cumulative impulse response of output to a monetary shock in a surprisingly simple way: it is equal to the benchmark with perfect information plus the ratio of the uncertainty of price-setting firms to the instantaneous variance of their marginal costs. This highlights the importance of the selection mechanism: it is not average uncertainty that matters, it is the uncertainty of price setters. Thus, the effects of a monetary shock with endogenous information acquisition always lie between those with perfect information and the benchmark under exogenously given imperfect information.

4.1. From Firm-Level Price Gaps to the Aggregate Output Gap

We begin by decomposing the aggregate response to shocks into firm-level responses to shocks. From the household's optimality conditions (Equations 7 and 8), aggregate output follows:

$$y_t = \frac{1}{\gamma}(m_t - p_t) \quad (23)$$

where $y_t \equiv \log Y_t - \log Y_0$, $m_t \equiv \log M_t - \log M_0$, and $p_t \equiv \log P_t - \log P_0$. Following the literature on the propagation of monetary shocks (see *e.g.*, Alvarez and Lippi, 2014), we will primarily be interested in studying the cumulative impulse response (CIR) of output to a monetary shock from the steady state at time $t = 0$:

$$\mathcal{M} = \int_0^\infty y_t dt \quad (24)$$

More generally, this object describes how a shock to firms' marginal costs generates real aggregate fluctuations. Thus, in richer models of the labor market in which the elasticity of wages to monetary shocks is not unitary, our results can be understood as characterizing how shocks to wages pass into prices and output. Moreover, even if monetary shocks are not truly one-time, permanent events, more complicated monetary shocks can be decomposed into a series of one-time shocks. Once decomposed, our CIR formula applies to each shock and thereby characterizes the effects of more complicated shock processes. Of course, more generally, one might care about the entire impulse response of aggregate output. In Appendix D, we characterize the entire path of aggregate output following a monetary shock.⁶

To compute this CIR, we can re-express the aggregate output gap as an integral of firm-level

⁶Lemma D.3 does this with strategic complementarities. Setting $\beta = 0$ in that formula provides the IRF of output in our baseline model.

output gaps and then integrate this over time. Formally, by log-linearizing the ideal price index:

$$p_t = \int_0^1 p_{i,t} di \quad (25)$$

Thus, we decompose the aggregate output gap as the integral of firm-level output gaps, $y_t = \int_0^1 y_{i,t} di$, where firm-level output gaps follow:

$$y_{i,t} = -\frac{1}{\gamma}(p_{i,t} - q_{i,t}) \quad (26)$$

Hence, to characterize the response to monetary shocks, we need only consider how firms' prices respond to the shock. To do this, we decompose firms' output gaps into two components. The first is the belief gap, $y_{i,t}^b = \frac{1}{\gamma}(q_{i,t} - \mathbb{E}_{i,t}[q_{i,t}])$, which measures the output effects of firms' errors in pricing from having incorrect information. The second is the perceived gap, $y_{i,t}^x = -\frac{1}{\gamma}(p_{i,t} - \mathbb{E}_{i,t}[q_{i,t}])$, which arises from a firm's price not having adjusted since it receives information. For a firm that last changed its price h periods ago and that has an initial belief gap y^b , perceived gap y^x , we define the firm-level cumulative output gap as

$$Y(y^b, y^x, h) = \mathbb{E} \left[\int_0^\infty y_{i,t} dt \mid y_{i,0}^b = y^b, y_{i,0}^x = y^x, D_{i,0} = h \right] \quad (27)$$

Following the monetary shock, we define the initial joint distribution of changes in belief gaps and perceived gaps and the lengths of pricing spells as $\mathcal{F} \in \Delta(\mathbb{R}^3)$. Moreover, we define the respective marginal distributions as $\mathcal{F}^b$, $\mathcal{F}^x$, and $\mathcal{F}^h$. As pricing is time-dependent, the distribution of pricing durations is exogenous to any monetary shock. Thus, $\mathcal{F}^h = F$, which is the distribution of pricing spell lengths in the cross-section of firms, and y^b and y^x are independent of h . We therefore have that the CIR is given by:

$$\mathcal{M}(\mathcal{F}) = \int_{\mathbb{R}^3} Y(y^b, y^x, h) d\mathcal{F}(y^b, y^x, h) \quad (28)$$

This reduces the question of how monetary shocks affect output to answering two questions. First, how do firms' lifetime output gaps depend on their initial belief gap, initial perceived gap, and the time since they last changed their price via Y ? Second, how do we aggregate firms' lifetime output gaps to compute the CIR?

4.2. Characterization of Lifetime Output Gaps

We first characterize a firm's expected lifetime output gap. To do this, we make use of the following definitions. We define the average conditional duration as $\bar{D}_h = \mathbb{E}_g^{h'}[h'|h]$, which is simply how long a firm that reset its price h periods ago expects to wait before resetting its price. By Theorem 1, we have that the Kalman gain for a firm that resets its price τ periods after last resetting its price

is $\kappa_\tau = \frac{\sigma^2_\tau}{U^* + \sigma^2_\tau}$. We define the average conditional Kalman gain as $\bar{\kappa}_h = \mathbb{E}_g^{h'}[\kappa_{h'+h}|h]$, which is the expected Kalman gain at the next price reset opportunity for a firm that last reset its price h periods ago. With these objects in hand, the following Proposition characterizes the expected lifetime output gap of a firm:

Proposition 1 (Lifetime Output Gap Characterization). *The expected lifetime output gap of a firm with initial pricing duration h , initial belief gap y^b , and initial perceived gap y^x is given by:*


$$Y(y^b, y^x, h) = \bar{D}_h y^x + \left(\bar{D}_h + \bar{D}_0 \frac{1 - \bar{\kappa}_h}{\bar{\kappa}_0} \right) y^b \quad (29)$$

Proof. See Appendix A.5 ■


To understand this result, consider first the lifetime output effect of a perceived gap. Importantly, as the firm knows its perceived gap, it persists only until the firm can reset its price in h' periods, at which point any perceived gap is reset to zero. We illustrate this in panel (a) of Figure 2. Thus, as the firm on average will take $\bar{D}_h$ periods to reset its price, the lifetime effect of a perceived gap y^x is simply $\bar{D}_h y^x$.

Second, in contrast to perceived gaps, belief gaps persist forever. We illustrate the dynamics for a sample path of price adjustment following a monetary shock that affects belief gaps in panel (b) of Figure 2. Initially, a belief gap operates in much the same way as a perceived gap. Until the firm next resets its price, in expectation its belief gap remains y^b . Thus, until the first price reset, a belief gap also contributes $\bar{D}_h y^b$ to the expected lifetime output gap of the firm. After this point, its behavior becomes more complicated. In particular, when a firm that reset its price h periods ago comes to reset its price in h' periods, Theorem 1 implies that it acquires a Gaussian signal of its marginal costs with a Kalman gain of $\kappa_{h+h'}$. Hence, if this firm had a belief gap of y^b at time t , it would have an expected belief gap of $\mathbb{E}_g^{h'}[1 - \kappa_{h+h'}|h]y^b = (1 - \bar{\kappa}_h)y^b$ at time $t + h'$. Moreover, on average, this belief gap persists for $\bar{D}_0$ periods before the firm's next price reset opportunity. Thus, between the first price reset and the second, the expected total output gap of a firm is $\bar{D}_0(1 - \bar{\kappa}_h)y^b$. After this point, if a further h'' periods elapse before the firm resets its price next, its Kalman gain at that point would be $\kappa_{h''}$ and so the firm's expected output gap at the second price reset opportunity would be $\mathbb{E}_g^{h''}[1 - \kappa_{h''}]\mathbb{E}_g^{h'}[1 - \kappa_{h+h'}|h]y^b = (1 - \bar{\kappa}_0)(1 - \bar{\kappa}_h)y^b$. Thus, once again integrating over the expected duration of the third pricing spell, this period contributes $\bar{D}_0(1 - \bar{\kappa}_0)(1 - \bar{\kappa}_h)y^b$ to the expected lifetime output gap. The same process now happens *ad infinitum* for all future spells: the initial belief gap gets down-weighted by $1 - \bar{\kappa}_0$ because of the acquisition of new information, and each spell lasts $\bar{D}_0$ periods on average. Hence, the total effect of the belief gap on the lifetime output

Figure 2: Contribution of a Single Firm to Monetary Non-Neutrality



(a) Monetary Non-Neutrality with Full Information



(b) Monetary Non-Neutrality with Costly Information

gap is given by the following geometric series:

$$\bar{D}_h y^b + \sum_{k=0}^{\infty} \bar{D}_0 (1 - \bar{\kappa}_0)^k (1 - \bar{\kappa}_h) y^b = \bar{D}_h y^b + \bar{D}_0 y^b \frac{1 - \bar{\kappa}_h}{\bar{\kappa}_0} \quad (30)$$

which collapses to the claimed expression in Proposition 1.

4.3. The Propagation of Monetary Shocks

We now characterize the propagation of monetary shocks conditional on the distribution of output gaps that they induce on impact. This is simply the integral of the expected lifetime output gaps of firms over the joint distribution of price gaps and pricing spells. As price gaps and spell duration are independent, Proposition 1 immediately implies that:

$$\mathcal{M}(\mathcal{F}) = \mathbb{E}_{\mathcal{F}}[y^x] \bar{D} + \mathbb{E}_{\mathcal{F}}[y^b] \left(\bar{D} + \bar{D}_0 \frac{1 - \bar{\kappa}_h}{\bar{\kappa}_0} \right) \quad (31)$$

where $\bar{D} = \mathbb{E}_f^h[\bar{D}_h]$ is the average pricing duration in the population and $\bar{\kappa} = \mathbb{E}_f^h[\bar{\kappa}_h]$ is the average across all firms of the expected Kalman gain when they next reset their prices. These objects are in principle quite complicated: they are double integrals of Kalman gains and durations with respect to two different distributions—the conditional distribution of price reset opportunities G and the cross-sectional distribution of pricing spell durations F . However, Theorem 2 shows that they collapse to a simple formula in terms of only the uncertainty of price-setters U^* and the instantaneous variance of marginal costs σ^2 :

Theorem 2 (CIR Characterization). *Given an initial distribution $\mathcal{F} \in \Delta(\mathbb{R}^3)$, the CIR is given by:*

$$\mathcal{M}(\mathcal{F}) = \mathbb{E}_{\mathcal{F}}[y^x] \bar{D} + \mathbb{E}_{\mathcal{F}}[y^b] \left(\bar{D} + \frac{U^*}{\sigma^2} \right) \quad (32)$$

Proof. See Appendix A.6. ■

This result follows from showing that the net present value of the average Kalman gain in the cross-section is given by the ratio of price-setters' uncertainty to the instantaneous variance of marginal costs. Moreover, it has two important implications: imperfect information about monetary shocks amplifies their real effects, and selection effects in information acquisition dampen the importance of imperfect information.

Imperfect Information Amplifies Monetary Non-Neutrality. Theorem 2 highlights that the effects of a monetary policy shock hinge on whether monetary policy shocks are observed (thus affecting perceived gaps) or unobserved (thus affecting belief gaps). Concretely, if there is a permanent monetary expansion of amount $m = \log M_t - \log M_0$ and it is unobserved, then all firms' initial belief gaps change by $y_m^b = \frac{m}{\gamma}$. We let the normalized CIR in this case be given by $\mathcal{M}^b = \mathcal{M}(\delta_0, \delta_{\frac{m}{\gamma}}, F) / \frac{m}{\gamma}$. By contrast, if the monetary shock m is observed, then $y^x = \frac{m}{\gamma}$ and no firm's belief gap changes. We let the normalized CIR in this case be given by $\mathcal{M}^x = \mathcal{M}(\delta_{\frac{m}{\gamma}}, \delta_0, F) / \frac{m}{\gamma}$. The following corollary characterizes the relative expansion of the economy under these two scenarios:

Corollary 5 (Imperfect Information Amplifies Monetary Non-Neutrality). *The difference between the normalized CIRs to a permanent and unobserved monetary shock and a permanent and observed monetary shock of the same size is:*

$$\Delta^{Info} \equiv \mathcal{M}^b - \mathcal{M}^x = \frac{U^*}{\sigma^2} > 0 \quad (33)$$

Proof. Immediate from Theorem 2. ■

The intuition for this result is simple: if firms are more sluggish in their adjustment of prices, then monetary policy has larger effects. Moreover, when firms have imperfect information, they are slower to adjust because they only learn about the shock over time.

Selection Dampens Monetary Non-Neutrality. Importantly, Theorem 2 shows that it is the uncertainty of price-setters alone that determines the non-neutrality of shocks and not the average uncertainty in the population. We let $\mathcal{M}^{exo}$ be the CIR of an unobserved monetary shock when firms' uncertainty is exogenously fixed at some level $\bar{U}$. The following result characterizes the importance of selection or the fact that price-setters' uncertainty is what matters and not the average level of uncertainty in the population:

Corollary 6 (Selection Dampens Monetary Non-Neutrality). *The difference between the normalized CIRs to permanent and unobserved monetary shocks under exogenous uncertainty and endogenous uncertainty is given by:*

$$\Delta^{Select} \equiv \mathcal{M}^{exo} - \mathcal{M}^b = \frac{\bar{U} - U^*}{\sigma^2} > 0 \quad (34)$$

Proof. Immediate from Theorem 2. ■

Intuitively, as uncertainty is lowest for price-setters by Theorem 1, and greater uncertainty amplifies monetary non-neutrality, it is immediate that selection effects in information acquisition dampen monetary non-neutrality relative to a benchmark model in which all firms have exogenous uncertainty equal to some level $\bar{U}$. Moreover, our characterization from Theorem 2 gives us a simple formula by which selection effects can be quantified in the data.

4.4. Comparative Statics for Monetary Non-Neutrality

Finally, we study how changes in uncertainty, microeconomic volatility, and price stickiness affect the CIR. To aid intuition, we also provide a formula for the CIR in the special case of Taylor pricing.

Uncertainty Shocks Dampen Monetary Non-Neutrality. First, we gauge how changes in firms' uncertainty affect the propagation of monetary policy shocks. Concretely, suppose that at time $t = 0$, each firm is subject to a shock that increases their prior uncertainty about their optimal reset price by $\tilde{U} > 0$. By computing the changes in the profile of Kalman gains across firms, we find the following formula for the effect of an uncertainty shock on the CIR:

Proposition 2 (Uncertainty Shocks Dampen Monetary Non-Neutrality). *The effect of an uncertainty shock $\tilde{U} > 0$ on the CIR is given by:*

$$\left. \frac{\partial^+ \mathcal{M}^b}{\partial^+ \tilde{U}} \right|_{\tilde{U}=0} = -\frac{1}{\bar{\kappa}_0} \frac{U^*}{\sigma^2} \mathbb{E}_g^h \left[\frac{\kappa_h^2}{\sigma^2 h} \right] < 0 \quad (35)$$

where ∂^+ denotes the right partial derivative of a function.

Proof. See Appendix A.7. ■

Intuitively, if firms are more uncertain, they rely less on their prior information and so update their prices more aggressively in response to the information they acquire. As a result, prices adjust more rapidly, and the real effects of monetary policy are dampened following an uncertainty shock.

Greater Microeconomic Volatility Dampens Monetary Non-Neutrality. So far, we have seen how exogenous uncertainty shocks affect monetary non-neutrality. We now study the more complicated question of how changes in microeconomic volatility affect monetary non-neutrality. This is more subtle because while σ^2 decreases the CIR all else equal, we know from Corollary 1 that U^* will increase in response to an increase in σ^2 . This potential ambiguity notwithstanding, by combining Theorems 1 and 2, we find that increases in microeconomic volatility always dampen monetary non-neutrality:

Proposition 3 (Microeconomic Volatility Dampens Monetary Non-Neutrality). *The effect of greater microeconomic volatility σ^2 on the CIR is given by:*

$$\frac{\partial \mathcal{M}^b}{\partial \sigma^2} = -\frac{U^*}{\sigma^4} \frac{1 - \mathbb{E}_g^h[e^{-rh}(1 - \kappa_h)]}{1 - \mathbb{E}_g^h[e^{-rh}(1 - \kappa_h)^2]} < 0 \quad (36)$$

Proof. See Appendix A.8. ■

Intuitively, the direct effect of greater microeconomic volatility making firms rely less on prior information always dominates the fact that firms acquire less information when microeconomic volatility rises. This is qualitatively different from the role of microeconomic volatility in models that do not feature nominal rigidities, such as Moscarini (2004), in which marginal cost volatility can have non-monotone effects on price-responsiveness.

Price Stickiness Has an Ambiguous Effect on Monetary Non-Neutrality. Moreover, we observe in the following result that changes in the stickiness of prices have an *ambiguous* effect on the real effects of a monetary shock:

Proposition 4 (Ambiguous Effects of Price Stickiness on Monetary Non-Neutrality). *For $\varepsilon > 0$, let $G_\varepsilon(h) \equiv G(h - \varepsilon)$, $\forall h \geq \varepsilon$ denote a distribution that increases the duration of all price spells by ε . The effect of greater price stickiness on the CIR is given by:*

$$\left. \frac{\partial^+ \mathcal{M}^b}{\partial^+ \varepsilon} \right|_{\varepsilon=0} = \frac{1}{1 - \mathbb{E}_g^h[e^{-rh}(1 - \kappa_h)^2]} \left[1 - r \frac{U^*}{\sigma^2} \left(\frac{U^*}{U^{Min}} - 1 \right) \right] - \frac{\bar{D}}{\bar{D}_0} \geq 0 \quad (37)$$

Proof. See Appendix A.9. ■

This ambiguity arises because there are two (potentially) opposing forces at play. First, more sticky prices increase the average duration of pricing spells $\bar{D}$, which increases the real effects of monetary shocks. Second, more sticky prices affect firms' optimal choice of uncertainty U^* . As

we saw in Corollary 2 for the special case of Taylor pricing, more sticky prices can decrease firms' optimal uncertainty. This is quite intuitive: if the price is stuck for longer, it's more important to make that price a good one, so it's better to acquire more information. Thus, how changes in the stickiness of prices affect the CIR is a quantitative question to which we will return in Section 7.

The CIR Under Taylor Pricing. Finally, to aid intuition for the economic forces that shape the CIR, we solve in closed form for the CIR in the special case of Taylor pricing with no discounting:

Corollary 7. *Under Taylor pricing and zero discounting, i.e., firms reset prices every $k \in \mathbb{R}_+$ periods and $r = 0$, the CIR is given by:*

$$\mathcal{M}^b = \sqrt{\left(\frac{k}{2}\right)^2 + \frac{\omega}{B\sigma^2}} \quad (38)$$

Proof. Immediate from combining Theorem 2 and Corollary 2. ■

It is immediate from this formula that: increases in microeconomic volatility lower the CIR; increases in the price elasticity of demand lower the CIR; increases in the cost of information increase the CIR; and increases in price stickiness increase the CIR, but by less than one-for-one because of the endogenous response of firms acquiring more information when stickiness increases.

5 Identification of Monetary Non-Neutrality from Microeconomic Data

In our final theoretical results, we turn to which data are necessary and sufficient to identify the effects of pricing and information frictions on monetary non-neutrality in our model. We show that the cross-sectional distributions of uncertainty and pricing durations across firms are sufficient to identify the CIR. Moreover, we show that access to standard data on price changes is insufficient to identify the component of CIR that stems from firms' subjective uncertainty.

5.1. The Distributions of Uncertainty and Pricing Durations Are Sufficient for Identification

We first show how data on firms' uncertainty about their optimal reset prices and the duration of their pricing spells are sufficient for identification. Formally, let l be the density of firms' uncertainty.

Proposition 5 (Characterization of the Distribution of Uncertainty). *The cross-sectional density of uncertainty about optimal reset prices $l \in \Delta(\mathbb{R}_+)$ is given by:*

$$l(z) = \begin{cases} 0, & z < U^*, \\ \frac{1}{\sigma^2} f\left(\frac{z-U^*}{\sigma^2}\right), & z \geq U^*. \end{cases} \quad (39)$$

where $f(\cdot) = \frac{1}{D_0}(1 - G(\cdot))$ is the density of ongoing spell lengths in the cross-section.

Proof. See Appendix A.10 ■

This result tells us that knowledge of the distribution of uncertainty l and the length of ongoing pricing spells f is sufficient to identify the uncertainty of price-setters U^* , the instantaneous variance of marginal costs σ^2 , and the average expected duration of pricing spell $\bar{D}$, which in turn identify the CIR $\mathcal{M}(\mathcal{F})$ for any $\mathcal{F}$.

From Identification to Estimation. Moreover, this result suggests a simple methodology by which U^* and σ^2 can be estimated from data. First, observe that the uncertainty of price-setters is given by the mode of the uncertainty distribution $U^* = \text{mode}_l[U]$. Thus, given an empirical estimate of the uncertainty distribution $\hat{l}$, we obtain the following estimator for U^* :

$$\hat{U}^* = \text{mode}_{\hat{l}}[U] \quad (40)$$

Second, given an empirical estimate $\hat{f}$ of the distribution of ongoing spell lengths and our estimate of the uncertainty of price-setters $\hat{U}^*$, by Proposition 5 we can determine the model implied uncertainty distribution as:

$$l^M(z; \sigma^2) = I_{[z \geq \hat{U}^*]} \frac{1}{\sigma^2} \hat{f}\left(\frac{z - \hat{U}^*}{\sigma^2}\right) \quad (41)$$

which depends on a single parameter, the volatility of marginal costs σ^2 . We can then therefore estimate σ^2 by minimizing the distance between $l^M(\sigma^2)$ and $\hat{l}$:

$$\hat{\sigma}^2 \in \argmin \int_{\hat{U}^*}^{\infty} (\hat{l}(z) - l^M(z; \sigma^2))^2 dz \quad (42)$$

We now have a practical method to estimate the CIR using data on uncertainty and pricing durations.

5.2. Data on Price Changes Are Insufficient for Identification

We finally show that data on uncertainty is *necessary* in the sense that data on price changes and pricing durations are *insufficient* to identify the CIR in the absence of information about uncertainty. As is well known (see *e.g.*, Alvarez, Le Bihan, and Lippi, 2016), data on price changes are sufficient to identify the CIR in many models with both state-dependent pricing and time-dependent pricing frictions. Thus, it is natural to ask if data on price changes are sufficient to identify the CIR in the presence of information frictions. The following result answers this question in the negative:

Theorem 3 (Invariance to Uncertainty of the Distribution of Price Changes). *The distribution of price changes conditional on a firm changing its price $H \in \Delta(\mathbb{R})$ is invariant to U^* and follows:*

$$H(\Delta p) = \int_0^{\infty} \Phi\left(\frac{\Delta p}{\sigma\sqrt{h}}\right) dG(h) \quad (43)$$

where Φ is the standard normal CDF

Proof. See Appendix A.11. ■

We prove this result by first deriving the conditional distribution of price changes conditional on a firm's last pricing spell lasting h periods and conditional on a firm's information set at the beginning of its last pricing spell S_i^{t-h} . We show that the conditional variance of such price changes is invariant to the information set of the firm. Intuitively, the nature of the firm's optimal information acquisition makes its price change independent of the prices that it previously charged. Moreover, from the form of the firm's optimal information policy derived in Theorem 1, the conditional variance of price changes depends only on the volatility of marginal costs σ and the length of the pricing spell h and is given by $\sigma^2 h$. By mixing this distribution over the distribution of pricing durations, we obtain the distribution of price changes.

The important upshot of this result is that data on price changes, even in conjunction with data on pricing durations, are insufficient to identify U^* and, therefore, the real effects of monetary policy when there is endogenous information acquisition. Thus, data on uncertainty are not only sufficient for identifying U^* , but they are also necessary.

6 Using Survey Data to Quantify and Test the Model

We have shown how to identify the effects of uncertainty on the real effects of monetary shocks, given information about firms' uncertainty and the volatility of their marginal costs. We now show how to use survey microdata on firms' uncertainty and the duration of their pricing spells to identify and estimate these quantities. Using a survey of New Zealand firms from Coibion, Gorodnichenko, Kumar, and Ryngaert (2021), we perform this estimation. Turning to the CIR to a monetary shock, we find that the effect of uncertainty is of comparable magnitude to the effect of price stickiness itself and that the effect of selection is of a comparable magnitude again. From this, we conclude that uncertainty is critical for understanding the real effects of monetary policy and that the endogeneity of information acquisition is equally important.

6.1. Survey Data on Firms' Uncertainty and Pricing Duration

Motivated by our identification results, we need data on firms' uncertainty about their optimal reset prices and how long ago they last reset their price. To obtain these data, we use the survey of firm managers in New Zealand described in Coibion, Gorodnichenko, Kumar, and Ryngaert (2021), implemented between 2018Q1 and 2018Q2. The survey included 515 firms with six or more employees. These firms were a random sample of firms in New Zealand with broad sectoral coverage.⁷

⁷Previous works have used the survey data to characterize how firms form their expectations. For example, Afrouzi (2024) shows that strategic complementarity decreases with competition and reports that firms with more competitors

In Appendix B, we provide the full details on exactly how these data were collected and how we process the raw data. To summarize, these data contain two questions that allow us to measure the key objects of interest. First, firms are asked about their subjective uncertainty about their ideal prices:

*Q1: If your firm was free to change its price (i.e. suppose there was no cost to renegotiating contracts with clients, no costs of reprinting catalogues, etc...) today, what probability would you assign to each of the following categories of possible price changes the firm would make? Please provide a percentage answer.*⁸

As the survey was conducted by phone, firms' answers are consistent in that they feature no probabilities below zero and all probabilities sum to one. To compute an estimate of the firm's uncertainty, we first compute an estimate of the firm's expectation of its optimal price by taking the midpoint of each bin and computing its expected value under the probabilities the firm manager provides. Then, we construct an estimate of the firm's uncertainty by computing the variance under the elicited probability distribution.⁹ This gives us a measure U_i of firm i 's uncertainty about its optimal reset price for each of the firms in our sample.

Second, firms are asked the time that has elapsed since they last changed their price:

Q2: When did your firm last change its price (in months) and by how much (in % change)?

This straightforwardly gives us a measure D_i of the duration of firm i 's pricing spell.

6.2. The Quantitative Impact of Uncertainty and Selection

We now use these data to quantify the importance of both uncertainty and selection for monetary non-neutrality. We first estimate the density of pricing durations and uncertainty using standard kernel density methods to obtain $\hat{f}$ and $\hat{l}$.¹⁰ We then obtain $\hat{U}^*$ and $\hat{\sigma}^2$ using our estimators from


have more certain posteriors about aggregate inflation. Also, Coibion, Gorodnichenko, Kumar, and Ryngaert (2021) evaluate the relation between first-order and higher-order expectations of firms, including how they adjust their beliefs in response to a variety of information treatments. Yang (2022) shows that firms producing more goods have both better information about inflation and more frequent but smaller price changes. See Coibion, Gorodnichenko, Kumar, and Ryngaert (2021) for a comprehensive description of the survey.

⁸Firms assigned probabilities to the following 16 bins: less than -25%, from -25% to -15%, from -15% to -10%, from -10% to -8%, from -8% to -6%, from -6% to -4%, from -4% to -2%, from -2% to 0%, from 0% to 2%, from 2% to 4%, from 4% to 6%, from 6% to 8%, from 8% to 10%, from 10% to 15%, from 15% to 25%, more than 25%.

⁹When we calculate the variance, we assume a uniform distribution within each bin. For example, if a firm assigns 100% on the bin "2-4 percent", then the implied variance is $\frac{1}{12}(4-2)^2 = 1/3$.

¹⁰We estimate $\hat{l}$ using a standard kernel density function with a bandwidth of 0.34 on $[0, 50]$. We then obtain $\hat{U}^*$ as the mode of $\hat{l}$ and reestimate the kernel density on $[\hat{U}^*, 50]$. We estimate $\hat{f}$ with a bandwidth of 2.4 on $[0, 80]$.

Figure 3: Distributions of Firms' Subjective Uncertainty in the Data and the Model




Notes: This figure shows the distribution of firms' subjective uncertainty about their ideal prices. The black vertical solid line shows the mode of the empirical distribution of subjective uncertainty ($\hat{U}^*$) and the black vertical dashed line shows the mean of the subjective uncertainty observed in the survey data. The blue solid line is the empirical distribution of uncertainty $\hat{l}(z)$. The red dashed line shows the estimated distribution of uncertainty ($l^M(z)$) from Equation 41 using the empirical distribution of time since the last price changes ($\hat{f}$) and the estimated uncertainty of shocks ($\hat{\sigma}^2$).

Equations 40 and 42. For all estimated objects, we construct standard errors using the bootstrap.¹¹ From this exercise, we obtain that $\hat{U}^* = 1.16$ (S.E.: 0.03) and $\hat{\sigma}^2 = 0.21$ (S.E.: 0.02). In Figure 3, we plot the estimated uncertainty distribution (in red) alongside the empirical uncertainty distribution (in blue). The fit, while not perfect, is surprisingly good given we only have one degree of freedom (the volatility of marginal costs σ^2) to match the entire distribution. In Appendix Figure F1, we plot the estimated conditional durations of pricing spells $\bar{D}_h$ as well as the estimated conditional Kalman gains $\bar{\kappa}_h$ that these estimates imply. In Appendix Figure F2, we plot the estimated distribution of price reset opportunities G and the corresponding hazard function θ , which is increasing in the duration of the pricing spell.

Using Theorem 2, we now estimate the extent to which uncertainty affects monetary non-neutrality as well as the extent to which selection effects in information acquisition matter. Figure 4 shows the monthly CIR of a 1 percentage point (pp) shock to output gaps under different scenarios (*i.e.*, to obtain the annual CIRs, simply divide the following numbers by 12). First, we recall as

¹¹Formally, for $d = 1, \dots, 10,000$, we uniformly resample $N = 515$ data points (the number of observations in the survey data). We re-estimate $\hat{f}_d$ and $\hat{l}_d$ using these data. We then re-estimate any model-implied quantities under these distributions and compute the distribution of the resulting estimates over the 10,000 bootstrap samples. We then compute the standard error as the standard deviation of the resulting distribution.

Figure 4: Estimated Monthly Cumulative Impulse Responses to an Initial 1 Percentage Point Output Gap under Different Scenarios



Notes: This figure shows the output effects of a 1 percentage point shock to perceived gaps (left bar), to belief gaps (middle bar), and belief gaps ignoring the selection effect (right bar). The output effect of a 1pp perceived gap is the average duration of firms' pricing spells $\Delta^{\text{Sticky}} = \bar{D}$, the effect of a 1pp belief gap is the effect of a perceived gap plus $\Delta^{\text{Info}} = \frac{U^*}{\sigma^2}$, and the effect of 1pp belief gap without selection effect is $\Delta^{\text{Sticky}} + \Delta^{\text{Info}}$ plus $\Delta^{\text{Select}} = \frac{\bar{U} - U^*}{\sigma^2}$. We present 95% confidence intervals as black vertical lines.

a baseline that the output effect of a 1pp perceived gap is simply the average duration of firms' pricing spells $\Delta^{\text{Sticky}} = \bar{D}$, which we estimate to be 5.95pp (S.E.: 0.17). The effect of a 1pp belief gap is the effect of a perceived gap plus $\Delta^{\text{Info}} = \frac{U^*}{\sigma^2}$, which we estimate to be 5.63pp (S.E.: 0.36). Thus, accounting for uncertainty is approximately as important for monetary non-neutrality as accounting for the mechanical effects of price stickiness. We also estimate the importance of selection $\Delta^{\text{Select}} = \frac{\bar{U} - U^*}{\sigma^2}$, which is the error in what we would have estimated Δ^{Info} to be if we naively used firms' average uncertainty rather than the uncertainty of price-setters, which we find to be 6.88pp (S.E.: 0.76). Thus, explicitly accounting for uncertainty is about as important as accounting for price stickiness itself. Moreover, accounting for selection is slightly more important than accounting for price stickiness itself. Indeed, computing the effects of shocks ignoring selection would massively overstate the non-neutrality of monetary shocks.

Further Model Predictions in the Data. Given that endogenous and exogenous information models have significantly different implications for monetary policy, we discuss other features of the data that are consistent with the endogenous information model.

First, one significant prediction of the endogenous information model is that the distribution of uncertainty should inherit the shape of the distribution of pricing durations up to a scaling factor.

As we saw in Figure 3, these two estimated distributions from the survey data are quite close, despite the fact that our model only allows for one free parameter to relate them. We suggest that such a close fit would be unlikely to be obtained in a model with exogenous information, as there should be no relationship between the two distributions.

Second, a stark quantitatively testable implication of the theory is that the magnitude of selection effects should always be equal to the average duration of firms' pricing spells. To see this, we can combine Corollary 6 with Corollary 4 to observe that:

$$\Delta^{\text{Select}} = \frac{\bar{U} - U^*}{\sigma^2} = \frac{\mathbb{E}_f[U^* + \sigma^2 h] - U^*}{\sigma^2} = \mathbb{E}_f[h] = \bar{D} \quad (44)$$

As nothing in our estimation approach imposes such a relationship, the prediction that $\Delta^{\text{Select}} = \bar{D}$ represents a strong overidentifying test of the theory. We estimate that $\widehat{\Delta^{\text{Select}}} - \bar{D} = 0.93$ with a 95% confidence interval (computed via the bootstrap) of $(-0.67, 2.17)$. We plot these estimates in Appendix Figure E5. The t -statistic against the null that $\Delta^{\text{Select}} = \bar{D}$ is 1.273. Thus, we cannot reject this overidentifying restriction at any conventional level of statistical significance. This provides additional evidence in favor of the theory.

Finally, the endogenous information model also implies an upward-sloping relationship between uncertainty and time since the last price change at the firm level. We show in Appendix Figure B.2 and Appendix Table F.1 that the survey evidence is consistent with this prediction. We do not wish to over-emphasize this result, as many factors that vary at the firm level could potentially drive such a result. That said, it provides further suggestive evidence in favor of the model.¹²

6.3. Robustness: Heterogeneity, Measurement Error, and General Time-Dependence

In three further analyses, we first probe the quantitative robustness of our findings when firms are heterogeneous in their nominal rigidities and marginal cost volatilities. Second, we perform a deconvolution analysis to explicitly account for the potential impact that measurement error in the survey might have on our findings. Finally, we examine the importance of allowing for time-dependent pricing frictions that are more general than those of Calvo (1983).

Ex Ante Heterogeneity. We have assumed in our analysis that all firms are *ex ante* identical and differ only because they experience different productivity shocks and pricing spells. Of course, firms may be heterogeneous in several respects, and this could matter for the propagation of monetary shocks. However, Theorem 2 tells us how heterogeneity can matter in very precise ways. In particular, if we

¹²Given our dataset is limited by the number of observations that we have in the survey, while we do find that firms that reset their prices more than one year ago are more uncertain, alternative specifications yield unsurprisingly noisy estimates.

Table 1: Estimates of Sectoral Heterogeneity in Uncertainty and Marginal Cost Volatility

	GDP Share	Obs.	$\hat{U}^*$	$\hat{\sigma}^2$
Manufacturing and Construction	0.284	195	1.210	0.152
Trade, Transportation, Accommodation, and Food Services	0.290	150	1.150	0.283
FIRE and Professional Services	0.426	170	1.090	0.240
GDP-Weighted Average of Three Sectors	1	515	1.142	0.228
All sector (Baseline)	1	515	1.160	0.206

Notes: This table shows the estimation results for $\hat{U}^*$ and $\hat{\sigma}^2$ for three groups of sectors. We also present the GDP-weighted average of these estimates, as well as the baseline estimates for all sectors. The GDP share is computed using the 2018 New Zealand GDP by sectors. FIRE stands for Financial Activities, Information, and Real Estate services sectors.

augment the model to allow for arbitrary cross-firm heterogeneity in all relevant primitives (pricing durations G_i , the costs of mispricing B_i , the costs of information acquisition ω_i , and the volatility of marginal costs σ_i), we have that the CIR to a belief shock is given by:

$$\mathcal{M}^b = \mathbb{E}[\bar{D}_i] + \mathbb{E}\left[\frac{U_i^*}{\sigma_i^2}\right] \quad (45)$$

where $\bar{D}_i$ is the average expected duration implied by G_i and U_i^* is the posterior uncertainty of price setter i . Moreover, as $\mathbb{E}[\bar{D}_i] = \bar{D}$, heterogeneity does not matter for the mechanical term coming from price stickiness. Heterogeneity therefore matters precisely insofar as there is heterogeneity in $\frac{U_i^*}{\sigma_i^2}$. Moreover, by allowing for unrestricted heterogeneity in pricing hazards across firms, this formula holds under many recently developed extensions of the simple Calvo model, such as the mixed proportional hazard model proposed by [Alvarez, Borovičková, and Shimer \(2021\)](#).

To gauge the potential importance of such heterogeneity, we re-estimate U_i^* and σ_i^2 across different sectors, which are potentially quite likely to differ along each of the possible margins highlighted above. We present the results of this analysis in Table 1. We find estimates of U^* that are very similar across sectors, ranging between 1.09 and 1.21, while finding more substantial heterogeneity in the instantaneous variance of marginal costs, ranging between 0.15 and 0.28. Weighting each sector by its GDP contribution, we find that $\Delta^{\text{Info}} = \hat{\mathbb{E}}\left[\frac{\hat{U}_i^*}{\hat{\sigma}_i^2}\right] = 5.37$ (S.E.: 0.48), which is close to our baseline estimate of 5.63 without sectoral heterogeneity. More advanced modeling of heterogeneous pricing hazards across firms, such as that performed by [Alvarez, Borovičková, and Shimer \(2021\)](#), would require panel data to which we do not have access from this survey. Extending the analysis to account for heterogeneity of this sort is an interesting avenue for future work.

Measurement Error. As uncertainty is a complex variable to elicit, it is of course possible that firms' measured uncertainty is contaminated with measurement error. To examine the robustness of our results to the possibility of measurement error in firms' uncertainty, we use a standard deconvolution approach.

Formally, we assume that measurement error is additive in logarithms:

$$\log U_i = \log U'_i + \zeta_i \quad (46)$$

where U_i is the uncertainty that we measure, U'_i is true uncertainty for firm i , and $\zeta_i \sim N(0, \sigma_\zeta^2)$ is measurement error with mean zero and variance σ_ζ^2 . We then estimate the distribution of firms' true uncertainty l' by using the deconvolution kernel density approach of [Stefanski and Carroll \(1990\)](#) and selecting the theoretically optimal bandwidth for a Gaussian distribution from the observed data. From the estimated distribution of true uncertainty $\hat{l}'(\sigma_\zeta^2)$, we compute its mode as our estimate of the optimal reset uncertainty, $\hat{U}^*(\sigma_\zeta^2) = \text{Mode}_{\hat{l}'}[U']$. Following [Proposition 5](#), we then have that the model-implied uncertainty distribution is given by:

$$l^M(z; \sigma^2, \sigma_\zeta^2) = \mathbb{I}[z \geq \hat{U}^*(\sigma_\zeta^2)] \frac{1}{\sigma^2} \hat{f}\left(\frac{z - \hat{U}^*(\sigma_\zeta^2)}{\sigma^2}\right) \quad (47)$$

We can then estimate the variance of marginal costs σ^2 and the extent of measurement error σ_ζ^2 by minimizing the distance between the model-implied uncertainty distribution and the estimated distribution of true uncertainty:

$$(\hat{\sigma}^2, \hat{\sigma}_\zeta^2) \in \arg \min \int_{\hat{U}^*(\sigma_\zeta^2)}^{\infty} \left(l'(z; \sigma_\zeta^2) - l^M(z; \sigma^2, \sigma_\zeta^2) \right)^2 dz \quad (48)$$

The estimated variance of measurement error is $\hat{\sigma}_\zeta^2 = 0.48$, and the estimated variance of marginal cost is $\hat{\sigma}^2 = 0.23$, which is larger than the baseline estimate of 0.21. Moreover, $\hat{U}^*(\hat{\sigma}_\zeta^2) = 0.87$ is smaller than the baseline estimate of 1.16. In [Appendix Figure F3](#), we compare the estimated uncertainty distributions with and without measurement error. In [Appendix Figure F6](#), we show the quantitative effects of accounting for measurement error on the CIR. We find a smaller, but quantitatively similar, value of Δ^{Info} and a larger, but quantitatively similar, value of Δ^{Select} .

Performance of the Model with [Calvo \(1983\)](#) Pricing. We have allowed for general time-dependence of pricing. It is interesting to inspect the extent to which restricting to [Calvo \(1983\)](#) pricing from the outset would have affected our quantitative conclusions. Concretely, we estimate the frequency of price-setting as $\hat{\theta} = \frac{1}{\hat{D}}$. We then take $\hat{f}$ as an exponential distribution with hazard $\hat{\theta}$ and re-estimate the CIR as before. In [Appendix Figure F4](#), compare the estimated uncertainty distribution to the empirical one. In [Appendix Figure F6](#), we show the effects of restricting to Calvo pricing on the CIR

and find that this has minimal quantitative effects. As a result, we argue that restricting to [Calvo \(1983\)](#) pricing is with little quantitative loss for quantifying the real effects of monetary policy in this empirical setting.

7 Counterfactuals: How Microeconomic Volatility and Price Stickiness Affect Monetary Non-Neutrality

In a final quantitative analysis, we study how changes in microeconomic volatility and price stickiness affect the degree of monetary non-neutrality. We leverage our empirical estimates to both sign and quantify the extent to which greater microeconomic volatility and price stickiness would affect the efficacy of monetary policy.

7.1. Microeconomic Volatility


We first use the model and data to ask how changes in microeconomic volatility matter for the propagation of monetary shocks. As evidence from [Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry \(2018\)](#) shows that microeconomic volatility is significantly higher in recessions, this allows us to gauge the implications of this fact, through the lens of our model, for the relative efficacy of monetary policy in booms versus recessions.

The effect of σ^2 on the output CIR is regulated by two opposing forces. First, there is a direct effect of increasing the volatility of firms' marginal costs. This makes them pay less attention to their priors as they know that their past information is less accurate. This means that firms pay more attention to their information, which dampens the real effects of monetary shocks. Second, there is an indirect effect on firms' optimal information choice. By Proposition 3, we know that the first effect theoretically dominates and the CIR is always decreasing with σ^2 ; however, since this effect is mitigated by the optimal choice of U^* , our objective here is to quantify the net effect of σ^2 .

Given our identification results, we can estimate both the sign and magnitude of $\frac{\partial \mathcal{M}^b}{\partial \sigma^2}$ by using the structure of our model and our estimates of firms' pricing durations $\hat{g}$, optimal uncertainty $\hat{U}^*$, and microeconomic volatility $\hat{\sigma}^2$. Together, this information pins down the effects of microeconomic volatility on the CIR up to a single parameter, the discount rate of firms r .

The left panel of Figure 5 shows the value of the CIR for different values of microeconomic volatility σ^2 , where we have calibrated the value of r to 0.0034 to match an annual interest rate of 4 percent. We observe that doubling microeconomic volatility decreases the CIR from its benchmark value of 11.58pp to around 9pp; *i.e.*, a monetary shock that increases the value of output by one percent on impact has around 2.5 percentage points less impact on CIR when microeconomic

Figure 5: Microeconomic Volatility, Price Stickiness, and Monetary Non-Neutrality



Notes: This figure shows two counterfactual analyses of how micro uncertainty and price stickiness affect monetary non-neutrality. The left panel shows the effect of microeconomic uncertainty on monetary non-neutrality. The right panel shows the effect of price stickiness on monetary non-neutrality. Red stars show our baseline estimates $\hat{\sigma}^2 = 0.21$ and $\epsilon = 0$. We present 95% confidence intervals as blue dashed lines.

volatility is doubled. This effect is not symmetric as the relationship is convex: cutting microeconomic volatility to half its estimated value of 0.21 increases the CIR to around 15.4pp, increasing the real effects of monetary shocks by around 4 percentage points. Finally, since r is the only externally calibrated parameter in this setting, the top panels of Appendix Figure F.7 show that these effects are robust and largely insensitive to alternative calibrations of the discount rate r .

Thus, we find that higher microeconomic volatility significantly dampens the real effects of monetary policy. Given the fact that microeconomic volatility is elevated during recessions, this model prediction is consistent with evidence from [Tenreyro and Thwaites \(2016\)](#) that monetary policy has less powerful real effects during recessions. This parallels a similar point that has been made in the context of models of lumpy adjustment (see *e.g.*, [Vavra, 2014](#)), in which higher volatility affects the frequency of price adjustments. However, the mechanism that underlies this result in our model is entirely different and independent of its effect on the frequency of price adjustments. In our setting, frequency is governed by the time-dependent arrival hazard that is not affected by volatility. Instead, this effect follows because firms pay less attention to prior information when marginal costs are more volatile and are therefore more responsive to current information and monetary shocks.

Moreover, the current literature on monetary non-neutrality with informational frictions (see *e.g.*, Afrouzi and Yang, 2021) largely emphasizes the role of macroeconomic volatility for monetary non-neutrality, while this result emphasizes the importance of microeconomic volatility (see *e.g.*, Lucas, 1972, Flynn, Nikolakoudis, and Sastry, 2023).

7.2. Price Stickiness

We now use the model to analyze how changes in price stickiness affect monetary non-neutrality. As firms' information acquisition decisions depend on the stickiness of prices, there is a non-trivial interaction between pricing frictions and information frictions in generating monetary non-neutrality. Indeed, our theoretical analysis showed that because of the endogenous response of firm information acquisition behavior, increases in price-stickiness have dampened effects on monetary non-neutrality. Here, we quantify the magnitude of this dampening mechanism.

As we have modeled general time-dependent pricing, there are many ways to perturb the stickiness of prices. For this exercise, to maximize transparency, we simply increase the duration of all pricing spells by a constant amount $\varepsilon > 0$, *i.e.*, a firm that would have reset its price at time h now resets its price at time $h + \varepsilon$. More formally, the distribution of price reset times changes from G to $\tilde{G}$, where $\tilde{G}(x) = G(x - \varepsilon)$ for all $x \geq \varepsilon$. Theorem 2 then implies that the effects on monetary non-neutrality of such an increase in price stickiness are given by:

$$\mathcal{M}^b(\varepsilon) = \bar{D}(\varepsilon) + \frac{U^*(\varepsilon)}{\sigma^2} \quad (49)$$

where the first term is the direct effect of an increase in stickiness. The second term is the indirect effect, which comes from how price stickiness affects the optimal level of uncertainty. Theorem 1 implies that this indirect effect has a theoretically ambiguous sign because of two countervailing effects of ε on U^* (as per Proposition 4). First, longer pricing durations make information more valuable for the current pricing spell by increasing its duration. Thus, a marginally better pricing decision now yields higher profits for a longer time. This encourages a lower level of optimal uncertainty. Second, longer pricing durations make information less valuable for all future pricing spells because today's information is less valuable further into the future. This serves to increase the marginal cost of information in the future and encourages a higher level of optimal uncertainty.

Despite this theoretical ambiguity, we can estimate both the sign and magnitude of these effects using our data, up to a calibration of the discount rate r . This requires us to estimate the ratio of the losses from mispricing parameter B to the information cost parameter ω , $\widehat{\left(\frac{B}{\omega}\right)}$, which we can do by finding the value of $\frac{B}{\omega}$ that rationalizes the U^* we see in the data. That is, we find the exact value of $\frac{B}{\omega}$ that solves the firm's first-order condition for the optimal choice of U^* (from Theorem 1) given

the U^* , σ^2 and g that we see in the data and any fixed value for r :

$$\left(\frac{B}{\omega}\right)(r) = \frac{r}{1 - \mathbb{E}_{\hat{g}}^h[e^{-rh}]} \left(\frac{1}{\hat{U}^*} - \mathbb{E}_{\hat{g}}^h \left[e^{-rh} \frac{1}{\hat{U}^* + \hat{\sigma}^2 h} \right] \right) \quad (50)$$

With this in hand, as ε moves, we can use Theorem 1 to solve for $U^*(\varepsilon)$ and then use Theorem 2 to compute how the CIR depends on ε .

The right panel of Figure 5 plots CIR as a function of ε under the calibrated value of r . We find that U^* decreases with ε but that the direct effect of price duration mostly dominates the sign of these changes on CIR; *e.g.*, increasing the duration of pricing spells by 4 months increases the CIR from its calibrated value of 11.58pp to around 12.47pp. If uncertainty were not endogenous to stickiness, the effect on the CIR would have been the direct effect on average durations, and so the CIR would have risen to 13.28pp. Thus, the fact that firms optimally acquire more information in the face of increased stickiness offsets 47% of the mechanical effect of price stickiness on the CIR. Finally, to probe robustness to the sole externally calibrated parameter, we plot the results of this exercise as we vary r from 0 to 0.02 in Figure F8. The results are not sensitive to this choice.

Taken together, this analysis shows that price stickiness and information frictions have important interactions: ignoring the response of firms' uncertainty to changing price stickiness would substantially overstate the impact of stickier prices on monetary non-neutrality.

8 Extensions and Robustness

In this final section, we extend our framework in three directions and document the qualitative and quantitative robustness of the results derived from our baseline model.

8.1. State-Dependent Pricing with Menu Costs

Our baseline model implies that price changes and information acquisition always occur at the same time (as per Theorem 1). This is a consequence of our assumption that firms face a purely time-dependent pricing friction. This framework thus captures the value of information in getting the size of price changes right, but does not endogenize the value of information in deciding *whether* to change prices. Additionally, taken literally, our baseline model is not consistent with evidence from Alvarez, Lippi, and Paciello (2018) that firms review their prices (*i.e.*, acquire information) about twice as frequently as they adjust prices.

Menu Cost Model and Solution Method. In this section, we relax these assumptions and consider an extension of the model that allows for state-dependence of price changes by adding a random menu cost structure for changing prices. In this model, firms' menu costs ϕ evolve according

to a Markov kernel, as is the state-of-the-art in modern analytical (Alvarez, Lippi, and Oskolkov, 2022) and quantitative (Blanco, Boar, Jones, and Midrigan, 2024) work.¹³ For our exact quantitative specification, we follow Alvarez, Lippi, and Oskolkov (2022) in considering a random menu cost structure in which menu costs are often infinite, but finite costs $\phi < \infty$ are drawn from a probability measure Λ according to a Poisson process with arrival rate θ . Our goal in adopting this model is to work with a canonical and well-understood existing specification for menu costs and to see if the results obtained in our time-dependent model are robust. The rest of the model follows our maintained assumptions, with one exception: to overcome the infinite-dimensionality of arbitrary firm beliefs as a state variable in the firm’s dynamic problems, we restrict the firms to acquire only Gaussian signals of their optimal reset prices. We provide the full details of the corresponding firm problem, value functions, and policies in Appendix C.1.

Solving this model is a non-trivial computational task as firms face a complicated decision problem—whether and how much information to acquire, whether to change prices, and what prices to charge conditional on reset—and must keep track of a three-dimensional state—their perceived gap, their uncertainty about their perceived gap, and the level of the menu cost. Moreover, aggregating this behavior to compute the response to a monetary shock requires keeping track of the evolution of the joint distribution of the three aforementioned states with the true price gap. Leveraging GPU processing, deriving an analytical four-step decomposition of the model’s Markov transition, and exploiting the sparsity of each of these four steps (though not the composite Markov matrix) renders the problem computationally feasible. As calibration targets, we use the empirical uncertainty and price duration distributions from our previous analysis alongside the relative frequency of price reviews and price changes from Alvarez, Lippi, and Paciello (2018). In Appendix C.2, we provide the exact details of our novel solution approach. The calibration and fit of this model are provided in Appendix Tables C.1 and C.2, respectively.


We now explore how this extension affects the main conclusions of our analysis.

Optimal Dynamic Information and Pricing Policies. We begin by exploring the quantitative analog of Theorem 1, which described firms’ optimal information acquisition and pricing policies. Figure 6 shows the inaction regions for information and pricing policies in the estimated menu cost model.

First, consider the pricing inaction boundaries (blue line). As uncertainty about optimal prices increases, the price inaction bands become wider. This is intuitive: with greater uncertainty, firms need to believe their perceived gap is larger to justify the risk of paying the menu cost when they

¹³The general setting nests discrete-time menu cost models by setting $\phi = \infty$ at all times except those that are integers. When time is an integer, the menu cost can be drawn IID from some distribution Λ .

Figure 6: Optimal Information and Pricing Policies in the Menu Cost Model




Notes: This figure shows the firm's optimal information and pricing policies over the state space of the perceived price gap (x) and normalized prior uncertainty (U/σ^2), conditional on a price review opportunity. The blue solid line denotes the pricing inaction boundary, while the red solid line denotes the information inaction boundary. The green shaded area indicates the region where firms both adjust prices and acquire information ($|\Delta p| > 0$ and $U' < U$). The blue shaded area corresponds to price adjustments without information acquisition ($|\Delta p| > 0$ and $U' = U$). The red shaded area, centered around the origin, depicts the region of total inaction where neither prices change nor information is acquired ($|\Delta p| = 0$ and $U' = U$). The white area represents the region where firms acquire information but do not adjust prices ($|\Delta p| = 0$ and $U' < U$). These policies are computed using a finer grid resolution than the baseline quantitative exercise, with $dt = 0.006$, $\Delta_x = 0.0352$, and $\Delta_U = 0.0012$.

did not actually need to do so.

Second, consider the information inaction boundary (red line). When uncertainty lies above this boundary, the firm will acquire information when it is possible for it to reset its price. On each side of $x = 0$ (about which the boundary is symmetric), the boundary is non-monotonic. Intuitively, when firms perceive a very small gap, they know they are less likely to pay a menu cost after they acquire information. As the absolute size of the perceived gap increases, firms become less tolerant of uncertainty about their optimal price. Intuitively, they perceive that they are more likely to be willing to pay the menu cost, and they would like to know if it is worth doing so. Finally, as the absolute size of the perceived gap becomes large, firms become more tolerant of uncertainty again (though still less tolerant than if they perceived their price gap as very small). These firms know they will pay the menu cost after acquiring information with very high probability, and so they do not need to use information to work out if it is worth paying the menu cost. At the same time, as they are very likely to pay the menu cost, they also know they will be setting a price. This is why

Figure 7: Comparative Statics with Respect to Menu Costs (ϕ)



Notes: This figure shows the sensitivity of key model statistics and CIRs to variations in the menu cost parameter, ϕ . The left panel shows the CIR to a perceived gap shock, $\mathcal{M}^x$ (blue line), plotted alongside the average duration of price changes, $\bar{D}$ (orange line). The middle panel compares the difference between the CIR to a belief shock and the CIR to a perceived gap shock, $\mathcal{M}^b - \mathcal{M}^x$ (blue line) against the baseline formula U^*/σ^2 (orange line), where U^* is the average posterior uncertainty of price-adjusting firms. The right panel shows the ratio of the frequency of realized price changes to the frequency of price reviews ($N(|\Delta p| > 0)/N(Review)$). In all panels, the vertical green dashed line indicates the value of ϕ used in the baseline calibration.

they demand more information than the firms with very small perceived gaps.

Thus, there are four regions of the state space: (i) some firms acquire no information and do not change their prices (red region), (ii) some firms acquire information but do not change their prices (white region), (iii) some firms do not acquire information but do change their price (blue region), and (iv) some firms both acquire information and then change their price (green region). In our baseline time-dependent model, Theorem 1 implies that only events in region (iv) could take place in the steady state.

Monetary Non-Neutrality. We now explore the quantitative analog of Theorem 2, which characterized the CIR to a one-time monetary shock. Figure 7 shows how the CIR depends on the menu cost parameter, ϕ , and how the simple formula from our baseline model (Theorem 2) performs.

First, consider how state-dependent pricing affects the CIR to perfectly observed monetary policy shocks (the first panel of Figure 7). When the menu cost is 0 (and this model collapses to our baseline model), observe that $\mathcal{M}^x = \bar{D}$, as the formula from Theorem 2 showed. As menu costs increase, we observe that $\mathcal{M}^x < \bar{D}$. This is because of the well-known selection effect: those price changes that do happen are larger than the average price gap. Thus, the time-dependent model, as is well-known (see e.g., Golosov and Lucas, 2007) *overstates* the importance of pricing frictions (as

$\mathcal{M}^x < \bar{D}$) if firms are truly subject to state-dependent pricing frictions.


Second, consider how state-dependent pricing affects the role of information frictions, as described by the effect of incomplete information on the CIR (the second panel of Figure 7). At $\phi = 0$, our theoretical results show that $\mathcal{M}^b - \mathcal{M}^x = U^* / \sigma^2$. As we increase the size of menu costs, we observe that this simple formula remains extremely quantitatively accurate. In particular, if we benchmark the size of menu costs to match a relative frequency of price changes to price reviews (the third panel of Figure 7) of one half (as suggested by the work of Alvarez, Lippi, and Paciello, 2018), then the simple formula from the time-dependent model is 99% accurate. Even if we double menu costs from this baseline value, the simple formula remains almost perfectly accurate. Thus, we conclude that the incorporation of state-dependent pricing frictions has no bearing on the accuracy of our formula for the effects of incomplete information. In Appendix C.4, we show that this conclusion remains true for a wide range of values of steady state inflation and for large shocks to the money supply. Intuitively, while firms do acquire information between realized price changes in this model, most information is still acquired close to the times at which firms are engaged in resetting prices. This is driven by the fact that—except when the perceived gap is very close to the inaction band, which occurs with low probability in the stationary distribution—the decision of whether to change a price is a binary decision that is cheap to optimally learn about while the decision of what price to charge is a real-valued decision that is expensive to optimally learn about. This highlights the value of our time-dependent analysis: it provides an analytically tractable yet quantitatively accurate description of how incomplete information affects monetary non-neutrality.

Taken together, as U^* / σ^2 remains extremely accurate for measuring the effects of information frictions and $\bar{D}$ overstates the effects of pricing frictions, our baseline time-dependent analysis differs from this state-dependent analysis by *understating* the relative importance of information frictions for monetary non-neutrality.

(Non-)Identification. Next, we investigate the quantitative analog of Theorem 3, which showed that the price change distribution is invariant to firms' uncertainty. The key upshot of this result was that the price change distribution does not identify the CIR. We now investigate if these conclusions translate to the quantitative menu cost model.

Figure 8 shows the distribution of price changes for price-adjusting firms across varying levels of the information cost parameter, ω . In our baseline model, this distribution is completely invariant to ω (Theorem 3). In the quantitative model, large changes in ω have imperceptible effects on the price change distribution. However, changes in ω have large effects on the CIR: increasing ω from

Figure 8: Distribution of Price Changes by Information Cost (ω)



Notes: This figure shows the distribution of price changes for price-adjusting firms across varying levels of the information cost parameter, ω . The blue bars represent the histogram for a low information cost specification ($\omega = 105.88$), the green bars represent the baseline specification ($\omega = 151.25$), and the red bars represent a high information cost specification ($\omega = 196.62$). The legend reports the corresponding average duration of price changes ($\bar{D}$) and the difference between the CIR to a belief gap shock and the CIR to a perceived gap shock ($\mathcal{M}^b - \mathcal{M}^x$) for each case.

106 to 197 increases the effects of incomplete information by approximately 50%.¹⁴

These results show that our non-identification result (Theorem 3) continues to provide accurate guidance: economies with more costly information have larger CIRs than those with less costly information, but the price change distribution is close to identical in both cases.

8.2. Strategic Complementarity in Pricing

In our baseline model, we made assumptions on preferences such that there is no strategic complementarity in price-setting. As is well-known, however, strategic complementarity or substitutability can arise from either the presence of income effects in labor supply and/or money demand, or non-isoelastic demand (*e.g.*, [Kimball \(1995\)](#) demand). To understand how the presence of strategic complementarities affects our results, we generalize our setting to allow for them following [Alvarez,](#)

¹⁴Appendix figure F.9 more systematically shows the sensitivity of key model statistics and CIRs to variations in the information cost parameter, ω . As is intuitive, greater costs of information increase the effects of unobserved monetary policy shocks as firms' uncertainty increases.

Lippi, and Souganidis (2023), and model the process of firms' optimal (log) prices as:

$$\tilde{q}_{i,t} = (1 - \beta)q_{i,t} + \beta p_t \quad (51)$$

where $\beta \in (-\infty, 1)$. If $\beta \in (0, 1)$, then price-setting decisions are strategic complements. If $\beta < 0$, then they are strategic substitutes. Our baseline model is nested as $\beta = 0$. It is well-known that this reduced-form structure arises from Kimball (1995) demand (see *e.g.*, Alvarez, Lippi, and Souganidis, 2023). We assume that firms can observe the aggregate price level when they make their price-setting decisions. The rest of our model is as in Section 2.

We analyze this model in Appendix D. We show that the path of real output following a one-time monetary shock obeys a non-autonomous and linear integral (or Volterra) equation (Lemma D.3). By analyzing this equation, we derive that (Theorem 4):

$$\mathcal{M}^b = \frac{\bar{D}}{1 - \beta} + \frac{U^*}{\sigma^2} \quad (52)$$

Thus, strategic complementarities interact with nominal rigidities in the usual way: they slow price adjustment and increase the real effects of monetary policy. However, quite surprisingly, there is no interplay between them and the effects of incomplete information. That is, it remains true that $\mathcal{M}^b - \mathcal{M}^x = U^* / \sigma^2$, no matter the extent of strategic complementarity.

8.3. Richer Learning Structures

Our baseline model had a particularly simple structure for firms' optimal information acquisition: acquire information only when changing prices. However, it is natural to wish to understand how different learning structures would affect the resulting monetary non-neutrality. Our baseline model implied that the effective Kalman gain (*i.e.*, the fraction of the belief gap in firms' prices that is closed when they reset their prices) for a firm resetting its price h' time after the shock that had a pricing duration of h when the monetary shock happened is given by $\kappa(h, h') = \sigma^2(h + h') / (U^* + \sigma^2(h + h'))$.

In Appendix E, we explore how different structures of κ affect monetary non-neutrality. As we show, this allows us to capture situations in which firms receive free information, perhaps because they learned from making other decisions or were exposed to information without actively seeking it out. We show that the effect of incomplete information on the CIR is given by (Theorem 5):

$$\mathcal{M}^b - \mathcal{M}^x = \frac{\int_0^\infty g(\tau) \left[\int_0^\tau (1 - \kappa(h, \tau - h)) dh \right] d\tau}{\int_0^\infty g(\tau) \kappa(0, \tau) d\tau} \quad (53)$$

To show the utility of this formula, we apply it to two variant information structures: (i) the firm receives free Gaussian signals about its optimal price in continuous time, and (ii) the firm receives these free signals and can acquire additional information, as in our baseline model. This

free information could be interpreted as information that spills over from other decisions made in the same firm. In both cases, we solve for the implied κ , plug it into the general formula, and find that $\mathcal{M}^b - \mathcal{M}^x = U^* / \sigma^2$. Thus, our formula from Theorem 2 exactly applies under these information structures. More generally, in Corollary E.1, we provide conditions under which our baseline formula represents a lower bound, *i.e.*, $\mathcal{M}^b - \mathcal{M}^x \geq U^* / \sigma^2$. These conditions are intuitive (i) firms make better pricing decisions when they have been able to learn about the true underlying state for longer ($\kappa(h, h') \leq \kappa(0, h + h')$), and (ii) firms are less responsive to information acquired over the life of the spell than they would be if all information were acquired at the end of the spell ($\kappa(0, h') \leq \sigma^2 h' / (U^* + \sigma^2 h')$).

9 Conclusion

In this paper, we study how to use microeconomic data on beliefs and pricing decisions to quantify the relative contributions of pricing and information frictions to monetary non-neutrality. In a canonical general equilibrium monetary economy with both such frictions, we analytically characterized: (i) firms' optimal dynamic information acquisition behavior (Theorem 1), (ii) how pricing and information frictions contribute to monetary non-neutrality (Theorem 2), and (iii) how data on firms uncertainty and pricing frictions—but, critically, not price changes alone—are sufficient to identify these contributions (Theorem 3 and Proposition 5). Moreover, while we find that information frictions amplify monetary non-neutrality, a key implication of this analysis was that there is selection in information acquisition: the price-setting firms are the most informed in the cross-section at any given time, and it is their beliefs that ultimately determine the degree of monetary non-neutrality. Thus, models with exogenous information are liable to overstate the role of information frictions in driving monetary non-neutrality.

Implementing our identification and estimation approach in a survey of firms' beliefs in New Zealand, we estimate that endogenous information acquisition doubles the degree of monetary non-neutrality relative to the benchmark model with no information costs, while a model with exogenous information would overstate monetary non-neutrality by approximately 60%. In estimated counterfactuals, we find that increases in microeconomic uncertainty substantially reduce monetary non-neutrality. This implies that times of high *microeconomic* uncertainty, such as recessions (see Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry, 2018), may also be times in which monetary policy is less effective in stimulating output. This prediction is moreover consistent with time series evidence that monetary policy is less powerful during recessions (Tenreyro and Thwaites, 2016). Moreover, we find that increased price stickiness increases monetary non-neutrality, but by

about 47% less than when information frictions are absent, because of the endogenous response of firms to acquire information. This result highlights the quantitative importance of modeling both information and pricing frictions jointly.

More broadly, our framework has implications for how measured beliefs (*e.g.*, from surveys) can be used to uncover the macroeconomic impacts of imperfect and endogenous information. This is useful because it is *ex ante* unclear whose beliefs, and which aspects of those beliefs, matter for any given outcome. For instance, within a standard general model of price-setting with endogenous information acquisition, we showed that the relevant moment of beliefs for monetary non-neutrality is *price-setters' uncertainty* about their optimal prices. This highlights how, for a given outcome of interest, one can use theory to narrow down whose beliefs to measure, what aspects of these beliefs to measure, and how to use these measured beliefs to understand macroeconomic phenomena at both quantitative and qualitative levels. Interestingly, in our case, our results imply that the ideal survey would use a *selected sample* of price-setters—as opposed to a representative sample of *all* firms, which is usually the targeted pool for firm surveys—and measure their *uncertainty about their desired prices*. Further work to elicit this object from firms could be highly informative for understanding the efficacy of monetary policy. We believe this implication should also hold in some form for settings where economic agents make infrequent decisions, such as households buying houses or other durable goods or firms making lumpy investment decisions. Indeed, interesting subsequent work has explored exactly this question in the context of household decisionmaking (de Silva and Mei, 2025). In all such settings, agents might prefer to acquire information when the decision is relevant, and so averages of uncertainty from representative samples might exaggerate the degree of information rigidities that are relevant for macroeconomic outcomes.

References

- AFROUZI, H. (2024): “Strategic inattention, inflation dynamics, and the non-neutrality of money,” *Journal of Political Economy*, Forthcoming.
- AFROUZI, H., AND C. YANG (2021): “Dynamic rational inattention and the Phillips curve,” CESifo Working Paper.
- ALVAREZ, F., H. LE BIHAN, AND F. LIPPI (2016): “The real effects of monetary shocks in sticky price models: A sufficient statistic approach,” *American Economic Review*, 106(10), 2817–2851.
- ALVAREZ, F., AND F. LIPPI (2014): “Price setting with menu cost for multiproduct firms,” *Econometrica*, 82(1), 89–135.
- ALVAREZ, F., F. LIPPI, AND A. OSKOLKOV (2022): “The macroeconomics of sticky prices with generalized hazard functions,” *The Quarterly Journal of Economics*, 137(2), 989–1038.
- ALVAREZ, F., F. LIPPI, AND L. PACIELLO (2018): “Monetary shocks in models with observation and menu costs,” *Journal of the European Economic Association*, 16(2), 353–382.
- ALVAREZ, F., F. LIPPI, AND P. SOUGANIDIS (2023): “Price setting with strategic complementarities as a mean field game,” *Econometrica*, 91(6), 2005–2039.

- ALVAREZ, F. E., K. BOROVÍČKOVÁ, AND R. SHIMER (2021): “Consistent evidence on duration dependence of price changes,” NBER Working Paper No. 29112.
- ALVAREZ, F. E., F. LIPPI, AND L. PACIELLO (2011): “Optimal price setting with observation and menu costs,” *The Quarterly Journal of Economics*, 126(4), 1909–1960.
- (2016): “Monetary shocks in models with inattentive producers,” *The Review of Economic Studies*, 83(2), 421–459.
- ANGELETOS, G.-M., AND J. LA’O (2009): “Incomplete information, higher-order beliefs and price inertia,” *Journal of Monetary Economics*, 56, S19–S37.
- AUCLERT, A., R. RIGATO, M. ROGNLIE, AND L. STRAUB (2024): “New pricing models, same old Phillips curves?,” *The Quarterly Journal of Economics*, 139(1), 121–186.
- BILS, M., AND P. J. KLENOW (2004): “Some evidence on the importance of sticky prices,” *Journal of political economy*, 112(5), 947–985.
- BLANCO, A., C. BOAR, C. JONES, AND V. MIDRIGAN (2024): *Non-linear inflation dynamics in menu cost economies*, no. w32094. National Bureau of Economic Research.
- BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. J. TERRY (2018): “Really uncertain business cycles,” *Econometrica*, 86(3), 1031–1065.
- BONOMO, M., C. CARVALHO, R. GARCIA, V. MALTA, AND R. RIGATO (2023): “Persistent Monetary Non-neutrality in an Estimated Menu Cost Model with Partially Costly Information,” *American Economic Journal: Macroeconomics*, 15(2), 466–505.
- CALVO, G. A. (1983): “Staggered prices in a utility-maximizing framework,” *Journal of Monetary Economics*, 12(3), 383–398.
- CANDIA, B., O. COIBION, AND Y. GORODNICHENKO (2023): “The macroeconomic expectations of firms,” in *Handbook of Economic Expectations*, pp. 321–353. Elsevier.
- CAPLIN, A. (2016): “Measuring and modeling attention,” *Annual Review of Economics*, 8, 379–403.
- CAPLIN, A. S., AND D. F. SPULBER (1987): “Menu costs and the neutrality of money,” *The Quarterly Journal of Economics*, 102(4), 703–725.
- CARVALHO, C. (2006): “Heterogeneity in price stickiness and the real effects of monetary shocks,” *Frontiers in Macroeconomics*, 2(1).
- CARVALHO, C., AND F. SCHWARTZMAN (2015): “Selection and monetary non-neutrality in time-dependent pricing models,” *Journal of Monetary Economics*, 76, 141–156.
- COIBION, O., Y. GORODNICHENKO, AND S. KUMAR (2018): “How do firms form their expectations? New survey evidence,” *American Economic Review*, 108(9), 2671–2713.
- COIBION, O., Y. GORODNICHENKO, S. KUMAR, AND J. RYNGAERT (2021): “Do you know that I know that you know...? Higher-order beliefs in survey data,” *The Quarterly Journal of Economics*, 136(3), 1387–1446.
- COVER, T. M., AND J. A. THOMAS (1991): *Elements of information theory*. John Wiley & Sons.
- DE SILVA, T., AND P. MEI (2025): “Selective Inattention to Interest Rates,” *Available at SSRN 5291837*.
- FLYNN, J. P., G. NIKOLAKOUDIS, AND K. A. SASTRY (2023): “A Theory of Supply Function Choice and Aggregate Supply,” Manuscript.
- GOLOSOV, M., AND R. E. LUCAS (2007): “Menu costs and Phillips curves,” *Journal of Political Economy*, 115(2), 171–199.
- GORODNICHENKO, Y. (2008): “Endogenous information, menu costs and inflation persistence,” NBER Working Paper No. 14184.
- KIMBALL, M. S. (1995): “The quantitative analytics of the basic neomonetarist model,” .
- LUCAS, R. E. (1972): “Expectations and the neutrality of money,” *Journal of Economic Theory*, 4(2), 103–124.
- (1976): “Econometric policy evaluation: A critique,” in *Carnegie-Rochester conference series on public policy*, vol. 1, pp. 19–46. North-Holland.
- MAĆKOWIAK, B., F. MATĚJKA, AND M. WIEDERHOLT (2018): “Dynamic rational inattention: Analytical results,” *Journal of Economic Theory*, 176, 650–692.
- (2023): “Rational inattention: A review,” *Journal of Economic Literature*, 61(1), 226–273.
- MAĆKOWIAK, B., AND M. WIEDERHOLT (2009): “Optimal sticky prices under rational inattention,” *American Economic Review*, 99(3), 769–803.
- MANKIW, N. G., AND R. REIS (2002): “Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve,” *The Quarterly Journal of Economics*, 117(4), 1295–1328.

- MORALES-JIMÉNEZ, C., AND L. STEVENS (2024): “Price Rigidities in US Business Cycles,” Manuscript.
- MOSCARINI, G. (2004): “Limited information capacity as a source of inertia,” *Journal of Economic Dynamics and control*, 28(10), 2003–2035.
- NAKAMURA, E., AND J. STEINSSON (2008): “Five facts about prices: A reevaluation of menu cost models,” *The Quarterly Journal of Economics*, 123(4), 1415–1464.
- NIMARK, K. (2008): “Dynamic pricing and imperfect common knowledge,” *Journal of Monetary Economics*, 55(2), 365–382.
- PACIELLO, L., AND M. WIEDERHOLT (2014): “Exogenous information, endogenous information, and optimal monetary policy,” *Review of Economic Studies*, 81(1), 356–388.
- REIS, R. (2006): “Inattentive producers,” *The Review of Economic Studies*, 73(3), 793–821.
- SHESHINSKI, E., AND Y. WEISS (1977): “Inflation and costs of price adjustment,” *The Review of Economic Studies*, 44(2), 287–303.
- SIMS, C. A. (2003): “Implications of rational inattention,” *Journal of Monetary Economics*, 50(3), 665–690.
- (2010): “Rational inattention and monetary economics,” in *Handbook of monetary economics*, vol. 3, pp. 155–181. Elsevier.
- STEFANSKI, L. A., AND R. J. CARROLL (1990): “Deconvolving kernel density estimators,” *Statistics*, 21(2), 169–184.
- STEVENS, L. (2020): “Coarse pricing policies,” *The Review of Economic Studies*, 87(1), 420–453.
- TAYLOR, J. B. (1979): “Staggered wage setting in a macro model,” *American Economic Review*, 69(2), 108–113.
- TENREYRO, S., AND G. THWAITES (2016): “Pushing on a string: US monetary policy is less powerful in recessions,” *American Economic Journal: Macroeconomics*, 8(4), 43–74.
- VAVRA, J. (2014): “Inflation dynamics and time-varying volatility: New evidence and an ss interpretation,” *The Quarterly Journal of Economics*, 129(1), 215–258.
- WOODFORD, M. (2003): “Imperfect common knowledge and the effects of monetary policy,” in *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, ed. by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, pp. 25–58. Princeton University Press.
- (2009): “Information-constrained state-dependent pricing,” *Journal of Monetary Economics*, 56, S100–S124.
- YANG, C. (2022): “Rational inattention, menu costs, and multi-product firms: Micro evidence and aggregate implications,” *Journal of Monetary Economics*, 128, 105–123.

Appendices

A Proofs

A.1. Proof of Theorem 1

Proof. We first characterize optimal pricing conditional on an arbitrary information policy $\mu_{\mathcal{W}S}^{i,t}$. Let $v_{i,t}$ be the firm's belief regarding W_{it} at time t . Suppose that the firm has received a pricing opportunity at some date t . The firm's price policy problem is given by:

$$J(v_{i,t}) = \sup_p \mathbb{E}^h \left[\int_0^h e^{-r\tau} \left[-\frac{B}{2} (p - q_{i,\tau})^2 \right] d\tau + e^{-rh} J(v_{i,t+h}) \mid v_{it} \right] \quad (\text{A.1})$$

Thus, any optimal price solves:

$$p_{i,t}(v_{i,t}) \mathbb{E}^h \left[\int_0^h e^{-r\tau} d\tau \right] = \mathbb{E}^h \left[\int_0^h e^{-r\tau} \mathbb{E}[q_{i,\tau} \mid v_{i,t}] d\tau \right] \quad (\text{A.2})$$

Using the fact that $\mathbb{E}[q_{i,\tau} \mid v_{i,t}] = \bar{q} + \sigma \mathbb{E}[W_{i,t} \mid v_{i,t}]$, we obtain:

$$p_{it}(v_{it}) = \bar{q} + \sigma \mathbb{E}[W_{i,t} \mid v_{i,t}] \quad (\text{A.3})$$

We can therefore compute the value function $J(v_{it})$ as:

$$J(v_{i,t}) = \mathbb{E}^h \left[\int_0^h e^{-r\tau} \left[-\frac{B}{2} \sigma^2 \mathbb{V}[W_{i,\tau} \mid v_{i,t}] \right] d\tau \right] + \mathbb{E} \left[e^{-rh} J(v_{i,t+h}) \mid v_{i,t} \right] \quad (\text{A.4})$$

We now show that the firm only acquires information when it changes its price. Fix a time t at which the firm cannot change its price. The value of a given information policy is given by:

$$\tilde{V}(v_{i,t}) = \mathbb{E}^h \left[-\omega \int_0^h e^{-r\tau} \frac{d\mathbb{l}_{i,\tau}}{d\tau} d\tau + e^{-rh} J(v_{i,t+h}) \mid v_{i,t} \right] \quad (\text{A.5})$$

Fix the horizon at which the firm next adjusts its price h . For each such h , suppose that the information policy yields $v_{i,t+h}$ and let the information be $\mathbb{l}_{i,\tau}$ under this policy. Consider instead an information policy that acquires no information until time $t+h$ and achieves the same $v_{i,t+h}$ and let the information be $\tilde{\mathbb{l}}_{i,\tau}$ under the policy. As both policies attain the same posterior at the next price-setting opportunity, the difference in the values of these policies is just the difference in the information costs. Moreover, we have that this difference in information costs satisfies:

$$\begin{aligned} \omega \left(\int_0^h e^{-r\tau} \frac{d\mathbb{l}_{i,\tau}}{d\tau} d\tau - e^{-rh} (\tilde{\mathbb{l}}_{i,t+h} - \tilde{\mathbb{l}}_{i,t}) \right) &\geq \omega e^{-rh} \left(\int_0^h \frac{d\mathbb{l}_{i,\tau}}{d\tau} d\tau - (\tilde{\mathbb{l}}_{i,t+h} - \tilde{\mathbb{l}}_{i,t}) \right) \\ &= \omega e^{-rh} ((\mathbb{l}_{i,t+h} - \mathbb{l}_{i,t}) - (\tilde{\mathbb{l}}_{i,t+h} - \tilde{\mathbb{l}}_{i,t})) \\ &= \omega e^{-rh} (\mathbb{l}_{i,t+h} - \tilde{\mathbb{l}}_{i,t+h}) \end{aligned} \quad (\text{A.6})$$

where the inequality follows as $e^{-rh} \leq e^{-r\tau}$ for $\tau \leq h$, the first equality follows by the fundamental theorem of calculus, and the final equality follows as the initial information under both policies is the same. Thus, acquiring information only when there is a price reset opportunity yields a higher value if this policy leads to acquiring less information in total. Consider the following garbling of

the signals obtained under the baseline information policy: receive a perfect signal about $v_{i,t+h}$, *i.e.*, garble $\{s_{i,\tau}\}_{\tau \in [t,t+h]}$ into the induced posterior at time $t+h$. As this is a garbling, and mutual information is monotone in the Blackwell order, we have that $\mathbb{I}_{i,t+h} \geq \bar{\mathbb{I}}_{i,t+h}$.

It remains to characterize optimal information acquisition when firms reset their price. First, we show that any optimal information structure is Gaussian. Fix a path of price reset times $\mathcal{R}$, let such a reset time be t , and let $v_{i,t-}$ be the belief at the start of time t . We have that $v_{i,0} = N(0, \sigma_0^2)$. Let $\{p_t\}_{t \in \mathcal{R}}$ be the sequence of random variables corresponding to the firm's reset prices at each reset date and let S^t be the information set implied by this price sequence. Now define a sequence of Gaussian random variables $\{\hat{p}_t\}_{t \in \mathcal{R}}$ such that for all $t \in \mathcal{R}$: $\mathbb{V}[W_{i,t}|\hat{p}_t] = \mathbb{E}[\mathbb{V}[W_{i,t}|S^t]]$. The expected nominal profits of the firm are the same under both policies. Thus, $\{\hat{p}_t\}_{t \in \mathcal{R}}$ yields a payoff improvement if and only if its total mutual information is lesser. This is immediate as, for any given expected variance-covariance matrix, the Gaussian random variable maximizes entropy (see Chapter 12 in [Cover and Thomas, 1991](#)). Thus, as $\mathcal{R}$ was arbitrary, the firm should acquire a Gaussian signal at each price reset opportunity regardless of the sequence of price reset times.

Second, we write their dynamic optimization problem using this structure. We observe that, $W_{i,t+h} = W_{i,t+h} - W_{i,t} + W_{i,t} - W_{i,0}$, $W_{i,t+h} - W_{i,t} \perp W_{i,t} - W_{i,0}$, and $W_{i,t+h} - W_{i,t} | v_{i,t} \sim N(0, h)$. Thus, $v_{i,t+h-}$ is the convolution measure of $v_{i,t}$ with $N(0, h)$, which we will denote by $v_{i,t} * N(0, h)$. Moreover, we know that $\mathbb{V}[W_{i,t}|v_{i,t}] = \tau + \mathbb{V}[W_{i,t}|v_{i,t}]$. As the firm acquires a Gaussian signal, we have that their problem reduces to:

$$V(U_{i,t-}) = \max_{U_{i,t} \leq U_{i,t-}} -U_{i,t} \frac{B}{2} \mathbb{E}^h \left[\int_0^h e^{-r\tau} d\tau \right] + \mathbb{E}^h \left[e^{-rh} V(U_{i,t} + \sigma^2 h) \right] + \frac{\omega}{2} \ln \left(\frac{U_{i,t}}{U_{i,t-}} \right) \quad (\text{A.7})$$

Taking the first-order condition we have that (if the constraint that $U_{i,t} \leq U_{i,t-}$ is slack):

$$0 = -\frac{B}{2} \mathbb{E}^h \left[\int_0^h e^{-r\tau} d\tau \right] + \mathbb{E}^h \left[e^{-rh} V'(U_{i,t} + \sigma^2 h) \right] + \frac{\omega}{2} \frac{1}{U_{i,t}} \quad (\text{A.8})$$

By the envelope theorem, we also have that:

$$V'(U_{i,t} + \sigma^2 h) = -\frac{\omega}{2} \frac{1}{U_{i,t} + \sigma^2 h} \quad (\text{A.9})$$

Thus, we obtain the following condition for the optimality of $U_{i,t}$:

$$\frac{\omega}{U_{i,t}} - \mathbb{E}^h \left[e^{-rh} \frac{\omega}{U_{i,t} + \sigma^2 h} \right] = B \mathbb{E}^h \left[\int_0^h e^{-r\tau} d\tau \right] \quad (\text{A.10})$$

To see that this equation has a unique solution, we rewrite it as:

$$1 - U_{i,t} \frac{B}{\omega} \mathbb{E}^h \left[\int_0^h e^{-r\tau} d\tau \right] = \mathbb{E}^h \left[e^{-rh} \frac{U_{i,t}}{U_{i,t} + \sigma^2 h} \right] \quad (\text{A.11})$$

The right-hand side is a strictly positive and strictly increasing function of $U_{i,t}$ and the left-hand side is a strictly decreasing function that attains a value of 1 at $U_{i,t} = 0$ and attains a value of 0 at $\bar{z} = \frac{1}{\frac{B}{\omega} \mathbb{E}^h \left[\int_0^h e^{-r\tau} d\tau \right]}$. Thus, this equation has a unique solution U^* , which moreover satisfies $U^* \leq \bar{z}$.

Moreover, computing the second derivative of the objective function, we obtain:

$$-\frac{\omega}{2} \left(\frac{1}{U_{i,t}^2} - \mathbb{E}^h \left[e^{-rh} \frac{1}{(U_{i,t} + \sigma^2 h)^2} \right] \right) < -\frac{\omega}{2} \mathbb{E}^h \left[e^{-rh} \left(\frac{1}{U_{i,t}^2} - \frac{1}{(U_{i,t} + \sigma^2 h)^2} \right) \right] \leq 0 \quad (\text{A.12})$$

Thus, as the problem is strictly concave, we have this solution is simply the minimum between $U_{i,t-}$ and U^* . As a result, if $U_{i,t-} \leq U^*$ the firm acquires no information, and if $U_{i,t-} > U^*$, the firm acquires a Gaussian signal of $W_{i,t}$ that resets its posterior uncertainty about Z_{it} to U^* . ■

A.2. Proof of Corollary 1

Proof. By Theorem 1, the optimal level of uncertainty solves:

$$\begin{aligned} \text{LHS}(U^*; B, \omega, r, G) &\equiv 1 - U^* \frac{B}{\omega} \mathbb{E}^h \left[\int_0^h e^{-r\tau} d\tau \right] \\ &= \mathbb{E}^h \left[e^{-rh} \frac{U^*}{U^* + \sigma^2 h} \right] \equiv \text{RHS}(U^*; r, \sigma^2, G) \end{aligned} \quad (\text{A.13})$$

Given the existence of a unique solution U^* (from Theorem 1), the results are immediate from the observations that: LHS is decreasing in U^* , B (which is increasing in η), and G (in the sense of first-order stochastic dominance) and increasing in ω and r , and RHS is decreasing in σ^2 , r , and G and increasing in U^* . See Proposition 4 for an explicit example of how U^* moves ambiguously with respect to FOSD changes in G . ■

A.3. Proof of Corollary 2

Proof. Combining Equation 16 with the assumption of Taylor pricing, we have that:

$$\frac{\omega}{U^*} - e^{-rk} \frac{\omega}{U^* + \sigma^2 k} = B \left(\frac{1 - e^{-rk}}{r} \right) \quad (\text{A.14})$$

Rewriting this equation, we obtain that:

$$U^{*2} \left[B \left(\frac{1 - e^{-rk}}{r} \right) \right] + U^* \left[B \left(\frac{1 - e^{-rk}}{r} \right) \sigma^2 k - \omega (1 - e^{-rk}) \right] - \omega \sigma^2 k = 0 \quad (\text{A.15})$$

Application of the quadratic formula and noting that only the greater solution is valid completes the proof. ■

A.4. Proof of Corollary 4

Proof. By Theorem 1, the firm's uncertainty at a price-setting opportunity is reset to U^* and they acquire no information between price-setting opportunities. Thus, in h periods, their uncertainty is given by:

$$\begin{aligned} U_{i,t} &= \mathbb{V}[q_{i,t} | S_i^t] = \mathbb{V}[q_{i,t} | S_i^{t-h}] = \mathbb{V} \left[\sigma (W_{i,t} - W_{i,t-h}) + \sigma W_{i,t-h} | S_i^{t-h} \right] \\ &= \sigma^2 h + \mathbb{V} \left[\sigma W_{i,t-h} | S_i^{t-h} \right] = \sigma^2 h + U^* \end{aligned} \quad (\text{A.16})$$

as claimed. ■

A.5. Proof of Proposition 1

Proof. From Theorem 1, we know that firms do not acquire information between price resetting opportunities. Thus, $\mathbb{E}_{i,t}[q_{i,t}] = \mathbb{E}_{i,0}[q_{i,0}]$ until the firm next resets its price, which we will suppose happens in h' periods. As firms' marginal costs follow a martingale, this implies that the firm's expected belief gap until period h' is simply the firm's initial belief gap, y^b . From Theorem 1, we have that when firms reset their prices, they acquire a Gaussian signal of their marginal costs with a signal noise $\tilde{\sigma}_{h+h'}$ that resets their posterior uncertainty to U^* :

$$s_{i,t+h'} = W_{i,t+h'} + \tilde{\sigma}_{h+h'} \varepsilon_{i,t+h'} \quad (\text{A.17})$$

where $\varepsilon_{i,t+h'} \sim N(0, 1)$. Because of this, a resetting firm has a conditional expectation of the random component of their marginal costs that is given by:

$$\begin{aligned} \mathbb{E}_{i,t+h'}[W_{i,t+h'}] &= \kappa_{h+h'} s_{i,t+h'} + (1 - \kappa_{h+h'}) \mathbb{E}_{i,t}[W_{i,t}] \\ &= W_{i,t+h'} + (1 - \kappa_{h+h'}) (\mathbb{E}_{i,t}[W_{i,t}] - W_{i,t+h'}) + \kappa_{h+h'} \tilde{\sigma}_{h+h'} \varepsilon_{i,t+h'} \end{aligned} \quad (\text{A.18})$$

This implies that the belief gap is given by:

$$\begin{aligned} y_{i,t+h'}^b &= (1 - \kappa_{h+h'}) y_{i,t}^b + (1 - \kappa_{h+h'}) (W_{i,t+h'} - W_{i,t}) \frac{1}{\gamma} - \frac{\sigma}{\gamma} \kappa_{h+h'} \tilde{\sigma}_{h+h'} \varepsilon_{i,t+h'} \\ &= (1 - \kappa_{h+h'}) y_{i,t}^b + Z_{i,t+h'} \end{aligned} \quad (\text{A.19})$$

where $Z_{i,t+h'} \sim N(0, \hat{\sigma}_{h+h'}^2)$.

We can then proceed recursively to characterize expected lifetime output gaps by observing that:

$$Y(y^b, y^x, h) = \mathbb{E}^{h', Z} \left[\int_0^{h'} y^b d\tau + \int_0^{h'} y^x d\tau + Y \left((1 - \kappa_{h+h'}) y^b + Z_{h'}, 0, 0 \right) \right] \quad (\text{A.20})$$

We now guess and verify that $Y(y^b, y^x, h) = \beta(h) y^x + m(h) y^b$. Plugging this guess into Equation A.20 and matching coefficients, we obtain that $\beta(h)$ and $m(h)$ must satisfy:

$$\beta(h) = \mathbb{E}_g[h' | h] = \bar{D}_h \quad (\text{A.21})$$

$$m(h) = \mathbb{E}_g[h' | h] + m(0) \mathbb{E}_g^{h'} [1 - \kappa_{h+h'} | h] = \bar{D}_h + m(0) (1 - \bar{\kappa}_h) \quad (\text{A.22})$$

$$m(0) = \frac{\mathbb{E}_g[h']}{1 - \mathbb{E}_g[1 - \kappa_{h'}]} = \bar{D}_0 \frac{1}{\bar{\kappa}_0} \quad (\text{A.23})$$

completing the proof. ■

A.6. Proof of Theorem 2

Proof. First, by Proposition 1, we have that the CIR is given by Equation 31. We now show that

$\bar{D}_0 \frac{1-\bar{\kappa}}{\bar{\kappa}_0} = \frac{U^*}{\sigma^2}$. By definition, we have that:

$$\begin{aligned} 1 - \bar{\kappa} &= \mathbb{E}_f[1 - \bar{\kappa}_h] = \mathbb{E}_f \left[1 - \mathbb{E}_g^{h'} \left[\frac{\sigma^2(h+h')}{U^* + \sigma^2(h+h')} | h \right] \right] = \mathbb{E}_f \left[\mathbb{E}_g^{h'} \left[\frac{U^*}{U^* + \sigma^2(h+h')} | h \right] \right] \\ &= \mathbb{E}_f \left[\mathbb{E}_g^{h'} \left[\frac{\frac{U^*}{\sigma^2}}{\frac{U^*}{\sigma^2} + (h+h')} | h \right] \right] = \int_0^\infty \left[\int_h^\infty \frac{\frac{U^*}{\sigma^2}}{\tau + \frac{U^*}{\sigma^2}} \frac{g(\tau)}{1 - G(h)} d\tau \right] f(h) dh \end{aligned} \quad (\text{A.24})$$

We now state and prove an ancillary result that characterizes the cross-sectional distribution of durations in terms of the expected duration of a price-setting firm and the distribution of price-setting opportunities.¹⁵

Lemma A.1. *The distribution of pricing durations in the cross-section is given by:*

$$f(h) = \frac{1}{\bar{D}_0} (1 - G(h)) \quad (\text{A.27})$$

Proof. To derive f , define $p_h = \mathbb{P}[\tilde{h} \in [h-\delta, h]]$ and observe that $p_h = p_{h-\delta} \times (1 - \mathbb{P}[\text{Reset between } h-\delta \text{ and } h | \text{Not reset by } h-\delta])$. Thus, we have that:

$$p_h - p_{h-\delta} = -p_{h-\delta} \frac{G(h) - G(h-\delta)}{1 - G(h-\delta)} \quad (\text{A.28})$$

dividing by δ and taking the limit $\delta \rightarrow 0$, we obtain:

$$f'(h) = -f(h)\theta(h) \quad (\text{A.29})$$

Integrating this expression yields:

$$f(h) \propto \exp \left\{ - \int_0^h \theta(s) ds \right\} = \exp \left\{ - \int_0^h \frac{g(s)}{1 - G(s)} ds \right\} = 1 - G(h) \quad (\text{A.30})$$

Using the fact that $G(0) = 0$, we then have that $f(h) = f(0)(1 - G(h))$. Integrating both sides of this expression, we then have that:

$$1 = \int_0^\infty f(h) dh = f(0) \int_0^\infty (1 - G(h)) dh = f(0) \mathbb{E}_g[h] = f(0) \bar{D}_0 \quad (\text{A.31})$$

which implies that $f(h) = \frac{1}{\bar{D}_0} (1 - G(h))$, as claimed. ■

¹⁵As this result uses the fact that G admits a density, it does not nest Taylor pricing. However, our result still goes through. Concretely, we observe that $h' = k - h$ and f is uniform over $[0, k]$. Thus, we have that:

$$\mathbb{E}_f^h \left[\mathbb{E}_g^{h'} [h' | h] \right] = \mathbb{E}_f^h [k - h] = \frac{k}{2} \quad (\text{A.25})$$

Moreover, we have that:

$$\mathbb{E}_g[h' | h = 0] = \frac{\mathbb{E}_f^h \left[\mathbb{E}_g^{h'} \left[\frac{U^*}{U^* + \sigma^2(h+h')} | h \right] \right]}{1 - \mathbb{E}_g^{h'} \left[\frac{U^*}{U^* + \sigma^2 h'} | h = 0 \right]} = k \frac{\frac{U^*}{U^* + \sigma^2 k}}{1 - \frac{U^*}{U^* + \sigma^2 k}} = \frac{U^*}{\sigma^2} \quad (\text{A.26})$$

And the conclusion of Theorem 2 still holds.

Combining Equations A.24 and A.27, we obtain that:

$$\begin{aligned}
1 - \bar{\kappa} &= \int_0^\infty \left[\int_h^\infty \frac{\frac{U^*}{\sigma^2}}{\tau + \frac{U^*}{\sigma^2}} \frac{g(\tau)}{1 - G(h)} d\tau \right] \frac{1}{\bar{D}_0} (1 - G(h)) dh \\
&= \frac{1}{\bar{D}_0} \int_0^\infty \int_h^\infty \frac{\frac{U^*}{\sigma^2}}{\tau + \frac{U^*}{\sigma^2}} g(\tau) d\tau dh = \frac{1}{\bar{D}_0} \int_0^\infty \left[\int_0^\tau \frac{\frac{U^*}{\sigma^2}}{\tau + \frac{U^*}{\sigma^2}} g(\tau) dh \right] d\tau \\
&= \frac{1}{\bar{D}_0} \int_0^\infty \frac{\frac{U^*}{\sigma^2} \tau}{\tau + \frac{U^*}{\sigma^2}} g(\tau) d\tau = \frac{1}{\bar{D}_0} \frac{U^*}{\sigma^2} \int_0^\infty \frac{\tau}{\tau + \frac{U^*}{\sigma^2}} g(\tau) d\tau \\
&= \frac{1}{\bar{D}_0} \frac{U^*}{\sigma^2} \bar{\kappa}_0
\end{aligned} \tag{A.32}$$

which implies that $\bar{D}_0 \frac{1 - \bar{\kappa}}{\bar{\kappa}_0} = \frac{U^*}{\sigma^2}$. Substituting this into Equation 31 yields the result. $\blacksquare$

A.7. Proof of Proposition 2

Proof. After an uncertainty shock of $\tilde{U} > 0$, we have that a firm with pricing duration of h now has a prior uncertainty of $U^* + \tilde{U} + \sigma^2 h$ at the time the monetary shock hits. Moreover, by Theorem 1, we have that at the firm's next price reset opportunity, it will reset its posterior uncertainty to U^* . Thus, its Kalman gain must solve $U^* = (1 - \kappa_h(\tilde{U}))(U^* + \tilde{U} + \sigma^2 h)$ and so:

$$\kappa_h(\tilde{U}) = \frac{\tilde{U} + \sigma^2 h}{U^* + \tilde{U} + \sigma^2 h} \tag{A.33}$$

By adapting the arguments of Proposition 1, we then obtain that the CIR is given by:

$$\mathcal{M}^b = \bar{D} + \bar{D}_0 \frac{1 - \bar{\kappa}(\tilde{U})}{\bar{\kappa}_0} \tag{A.34}$$

Thus, the impact of an uncertainty shock is given by:

$$\left. \frac{\partial^+ \mathcal{M}^b}{\partial^+ \tilde{U}} \right|_{\tilde{U}=0} = -\frac{\bar{D}_0}{\bar{\kappa}_0} \bar{\kappa}'(\tilde{U})|_{\tilde{U}=0} \tag{A.35}$$

where we have that:

$$\begin{aligned}
\bar{\kappa}'(\tilde{U})|_{\tilde{U}=0} &= \mathbb{E}_f^h \left[\mathbb{E}_g^{h'} \left[\kappa'_{h+h'}(\tilde{U})|_{\tilde{U}=0} | h \right] \right] = \mathbb{E}_f^h \left[\mathbb{E}_g^{h'} \left[\frac{U^*}{(U^* + \sigma^2(h + h'))^2} | h \right] \right] \\
&= \int_0^\infty \left[\int_h^\infty \frac{U^*}{(U^* + \sigma^2 \tau)^2} \frac{g(\tau)}{1 - G(h)} d\tau \right] f(h) dh = \frac{1}{\bar{D}_0} \int_0^\infty \left[\int_h^\infty \frac{U^*}{(U^* + \sigma^2 \tau)^2} g(\tau) d\tau \right] dh \\
&= \frac{1}{\bar{D}_0} \int_0^\infty \left[\int_0^\tau \frac{U^*}{(U^* + \sigma^2 \tau)^2} g(\tau) dh \right] d\tau = \frac{1}{\bar{D}_0} \int_0^\infty \frac{U^* \tau}{(U^* + \sigma^2 \tau)^2} g(\tau) d\tau \\
&= \frac{1}{\bar{D}_0} \frac{U^*}{\sigma^2} \mathbb{E}_g^h \left[\frac{\kappa_h^2}{\sigma^2 h} \right]
\end{aligned} \tag{A.36}$$

Completing the proof. $\blacksquare$

A.8. Proof of Proposition 3

Proof. By Theorem 2, we have that:

$$\frac{\partial \mathcal{M}^b}{\partial \sigma^2} = \frac{\frac{\partial U^*}{\partial \sigma^2} - \frac{U^*}{\sigma^2}}{\sigma^2} \quad (\text{A.37})$$

Moreover, implicitly differentiating Equation 16 from Theorem 1, we obtain that:

$$\begin{aligned} 0 &= \omega \left(-\frac{1}{U^{*2}} \frac{\partial U^*}{\partial \sigma^2} + \mathbb{E}^h \left[e^{-rh} \left(\frac{\partial U^*}{\partial \sigma^2} + h \right) \frac{1}{(U^* + \sigma^2 h)^2} \right] \right) \\ &= \frac{\partial U^*}{\partial \sigma^2} \left(\mathbb{E}^h \left[e^{-rh} \frac{1}{(U^* + \sigma^2 h)^2} \right] - \frac{1}{U^{*2}} \right) + \mathbb{E}^h \left[e^{-rh} \frac{h}{(U^* + \sigma^2 h)^2} \right] \end{aligned} \quad (\text{A.38})$$

Thus, we can write:

$$\frac{\partial U^*}{\partial \sigma^2} = \frac{\mathbb{E}^h \left[e^{-rh} \frac{h}{(U^* + \sigma^2 h)^2} \right]}{\frac{1}{U^{*2}} - \mathbb{E}^h \left[e^{-rh} \frac{1}{(U^* + \sigma^2 h)^2} \right]} = \frac{U^*}{\sigma^2} \frac{\mathbb{E}^h \left[e^{-rh} \frac{U^*}{U^* + \sigma^2 h} \frac{\sigma^2 h}{U^* + \sigma^2 h} \right]}{1 - \mathbb{E}^h \left[e^{-rh} \left(\frac{U^*}{U^* + \sigma^2 h} \right)^2 \right]} = \frac{U^*}{\sigma^2} \frac{\mathbb{E}^h [e^{-rh} \kappa_h (1 - \kappa_h)]}{1 - \mathbb{E}^h [e^{-rh} (1 - \kappa_h)^2]} \quad (\text{A.39})$$

Combining this with Equation A.37, we obtain that:

$$\frac{\partial \mathcal{M}^b}{\partial \sigma^2} = \frac{U^*}{\sigma^4} \left(\frac{\mathbb{E}^h [e^{-rh} \kappa_h (1 - \kappa_h)]}{1 - \mathbb{E}^h [e^{-rh} (1 - \kappa_h)^2]} - 1 \right) = -\frac{U^*}{\sigma^4} \frac{1 - \mathbb{E}^h [e^{-rh} (1 - \kappa_h)]}{1 - \mathbb{E}^h [e^{-rh} (1 - \kappa_h)^2]} \quad (\text{A.40})$$

Observing that $e^{-rh} \leq 1$ and $\kappa_h \in [0, 1)$ for all $h \in \mathbb{R}_+$, we obtain that $\mathbb{E}^h [e^{-rh} (1 - \kappa_h)] < 1$ and $\mathbb{E}^h [e^{-rh} (1 - \kappa_h)^2] < 1$ and so $\frac{\partial \mathcal{M}^b}{\partial \sigma^2} < 0$. ■

A.9. Proof of Proposition 4

Proof. It follows from $\mathcal{M}_b = \bar{D} + \frac{U^*}{\sigma^2}$ that

$$\frac{\partial \mathcal{M}_b}{\partial \varepsilon} \Big|_{\varepsilon=0} = \frac{\partial \bar{D}}{\partial \varepsilon} \Big|_{\varepsilon=0} + \frac{1}{\sigma^2} \frac{\partial U^*}{\partial \varepsilon} \Big|_{\varepsilon=0}$$

To calculate $\frac{\partial \bar{D}}{\partial \varepsilon} \Big|_{\varepsilon=0}$. Let us define $\bar{D}(\varepsilon)$ as the average duration under $G_\varepsilon(h)$, which is given by:

$$\begin{aligned} \bar{D}(\varepsilon) &= \frac{1}{\varepsilon + \int_0^\infty (1 - G(h')) dh'} \int_0^\varepsilon h dh + \int_0^\infty (h + \varepsilon) \frac{(1 - G(h))}{\varepsilon + \int_0^\infty (1 - G(h')) dh'} dh \\ &= \frac{1}{\varepsilon + \bar{D}_0} \frac{\varepsilon^2}{2} + \frac{\bar{D}_0}{\bar{D}_0 + \varepsilon} \int_0^\infty (h + \varepsilon) \frac{(1 - G(h))}{\int_0^\infty (1 - G(h')) dh'} dh \\ &= \frac{1}{\varepsilon + \bar{D}_0} \frac{\varepsilon^2}{2} + \frac{\bar{D}_0}{\bar{D}_0 + \varepsilon} (\bar{D} + \varepsilon) \end{aligned} \quad (\text{A.41})$$

Thus, we have that:

$$\frac{\partial \bar{D}}{\partial \varepsilon} \Big|_{\varepsilon=0} = 1 - \frac{\bar{D}}{\bar{D}_0} \quad (\text{A.42})$$

As for the second term, let $U^*(\varepsilon)$ be the reset uncertainty defined under $G_\varepsilon(h)$. Then, by the

definition of $G_\varepsilon(h)$, Theorem 1 implies that $U^*(\varepsilon)$ solves:

$$\underbrace{\frac{\omega}{U^*(\varepsilon)} - \mathbb{E}^h \left[e^{-r(h+\varepsilon)} \frac{\omega}{U^*(\varepsilon) + \sigma^2(h+\varepsilon)} \right]}_{\equiv MC(\varepsilon)} = \underbrace{B \left(\frac{1 - \mathbb{E}^h[e^{-r(h+\varepsilon)}]}{r} \right)}_{\equiv MB(\varepsilon)}$$

Differentiating each side with respect to ε and evaluating at $\varepsilon = 0$ we have:

$$\begin{aligned} \left. \frac{\partial MC(\varepsilon)}{\partial \varepsilon} \right|_{\varepsilon=0} &= -\frac{\omega}{U^{*2}} \left. \frac{\partial U^*}{\partial \varepsilon} \right|_{\varepsilon=0} + \mathbb{E}^h \left[e^{-rh} \frac{\omega}{(U^* + \sigma^2 h)^2} \right] \left(\left. \frac{\partial U^*}{\partial \varepsilon} \right|_{\varepsilon=0} + \sigma^2 \right) + r \mathbb{E}^h \left[e^{-rh} \frac{\omega}{U^* + \sigma^2 h} \right] \\ \left. \frac{\partial MB(\varepsilon)}{\partial \varepsilon} \right|_{\varepsilon=0} &= B \mathbb{E}^h[e^{-rh}] \end{aligned}$$

Equating these two equations, we arrive at

$$\begin{aligned} \left. \frac{\partial U^*}{\partial \varepsilon} \right|_{\varepsilon=0} &= \frac{\sigma^2 \mathbb{E}^h \left[e^{-rh} \left(\frac{U^*}{U^* + \sigma^2 h} \right)^2 \right] + r \frac{U^{*2}}{\omega} \left(\mathbb{E}^h \left[e^{-rh} \frac{\omega}{U^* + \sigma^2 h} \right] - \frac{B}{r} \mathbb{E}^h[e^{-rh}] \right)}{1 - \mathbb{E}^h \left[e^{-rh} \left(\frac{U^*}{U^* + \sigma^2 h} \right)^2 \right]} \\ &= \frac{\sigma^2 \mathbb{E}^h \left[e^{-rh} (1 - \kappa_h)^2 \right] - r U^* \left(\frac{B U^*}{r \omega} - 1 \right)}{1 - \mathbb{E}^h \left[e^{-rh} (1 - \kappa_h)^2 \right]} \end{aligned}$$

Thus, we have

$$\begin{aligned} \left. \frac{\partial \mathcal{M}_b}{\partial \varepsilon} \right|_{\varepsilon=0} &= \left. \frac{\partial \bar{D}}{\partial \varepsilon} \right|_{\varepsilon=0} + \frac{1}{\sigma^2} \left. \frac{\partial U^*}{\partial \varepsilon} \right|_{\varepsilon=0} \\ &= \frac{1 - r \frac{U^*}{\sigma^2} \left(\frac{U^*}{U^{\text{Min}}} - 1 \right)}{1 - \mathbb{E}^h \left[e^{-rh} (1 - \kappa_h)^2 \right]} - \frac{\bar{D}}{\bar{D}_0} \end{aligned}$$

where $U^{\text{Min}} = \frac{\omega r}{B}$ is the minimum uncertainty as defined in the main text, so that $\frac{U^*}{U^{\text{Min}}} - 1 > 0$. ■

A.10. Proof of Proposition 5

Proof. By Corollary 4, a firm's uncertainty h periods after changing its price is $U = U^* + \sigma^2 h \geq U^*$. Thus, $L(z) = \mathbb{P}[U \leq z] = \mathbb{P} \left[h \leq \frac{z - U^*}{\sigma^2} \right] = F \left(\frac{z - U^*}{\sigma^2} \right)$. Differentiating this expression yields the claimed formula for $l(z)$. ■

A.11. Proof of Theorem 3

Proof. To derive the distribution of price changes, we start by finding the conditional distribution of price changes for firms that had a given duration of h periods who had a fixed information set at their last price change opportunity. We then marginalize over the distribution of price durations and information sets to obtain the price change distribution. To this end, consider a firm i that is changing its price at time t , which changed its price h periods ago and define:

$$\Delta^h p_{i,t} \equiv p_{i,t} - p_{i,t-h} = \sigma \left(\mathbb{E}_{i,t}[W_{i,t}] - \mathbb{E}_{i,t-h}[W_{i,t-h}] \right) \quad (\text{A.43})$$

Moreover, we have that:

$$\mathbb{E}[W_{i,t}] = \kappa_h s_{i,t} + (1 - \kappa_h) \mathbb{E}_{i,t-h}[W_{i,t-h}] \quad (\text{A.44})$$

where:

$$s_{i,t} = W_{i,t} + \tilde{\sigma}_h \varepsilon_{i,t} \quad (\text{A.45})$$

Combining these equations, we can write:

$$\Delta^h p_{i,t} = \sigma \kappa_h (W_{i,t} + \tilde{\sigma}_h \varepsilon_{i,t} - \mathbb{E}_{i,t-h}[W_{i,t-h}]) \quad (\text{A.46})$$

Therefore, we have that:

$$\Delta^h p_{i,t} | S_i^{t-h} \sim N(0, \check{\sigma}^2(S_i^{t-h})) \quad (\text{A.47})$$

where:

$$\check{\sigma}^2(S_i^{t-h}) = \kappa_h^2 \mathbb{V}[\sigma W_{i,t} + \sigma \tilde{\sigma}_h \varepsilon_{i,t} | S_i^{t-h}] \quad (\text{A.48})$$

where we know that:

$$\mathbb{V}[\sigma W_{i,t} | S_i^{t-h}] = \mathbb{V}[\sigma(W_{i,t} - W_{i,t-h}) + \sigma W_{i,t-h} | S_i^{t-h}] = \sigma^2 h + U^* \quad (\text{A.49})$$

as, by Theorem 1, we have that at a time of price reset (which $t - h$ is by assumption) the firm's posterior uncertainty is always equal to $\mathbb{V}[\sigma W_{i,t-h} | S_i^{t-h}] = U^*$. Thus, we have that:

$$\check{\sigma}^2(S_i^{t-h}) = \kappa_h^2 (\sigma^2 h + U^* + \sigma^2 \tilde{\sigma}_h^2) \quad (\text{A.50})$$

Moreover, the signal noise $\tilde{\sigma}_h^2$ that achieves the Kalman gain κ_h solves:

$$\sigma^2 \tilde{\sigma}_h^2 = (U^* + \sigma^2 h) \frac{1 - \kappa_h}{\kappa_h} \quad (\text{A.51})$$

and so we have that:

$$\check{\sigma}^2(S_i^{t-h}) = \kappa_h^2 (U^* + \sigma^2 h) \left(1 + \frac{1 - \kappa_h}{\kappa_h}\right) = \kappa_h (U^* + \sigma^2 h) = \sigma^2 h \quad (\text{A.52})$$

Thus, we have that conditioning on the firm's information set is irrelevant and the conditional distribution of price changes is the marginal distribution of price changes:

$$\Delta^h p_{i,t} | S_i^{t-h} \sim N(0, \sigma^2 h) \implies \Delta^h p_{i,t} \sim N(0, \sigma^2 h) \quad (\text{A.53})$$

Finally, integrating over the distribution of price durations, G , we obtain that the distribution of price changes is:

$$\begin{aligned} H(\Delta p) &= \mathbb{P}[\Delta p_{i,t} \leq \Delta p | \Delta p_{i,t} \neq 0] = \mathbb{E}_g^h \left[\mathbb{P}[\Delta^h p_{i,t} \leq \Delta p | \Delta p_{i,t} \neq 0] \right] \\ &= \mathbb{E}_g^h \left[\Phi \left(\frac{\Delta p}{\sigma \sqrt{h}} \right) \right] = \int_0^\infty \Phi \left(\frac{\Delta p}{\sigma \sqrt{h}} \right) dG(h) \end{aligned} \quad (\text{A.54})$$

which depends on σ and G but does not depend on U^* . ■

B New Zealand Firm-Level Survey Data

Data Source and Sampling. Our analysis relies on a novel survey of firms in New Zealand described in Coibion, Gorodnichenko, Kumar, and Ryngaert (2021). While that study focused on higher-order beliefs, we leverage a distinct set of questions from the same survey wave to analyze the relationship between pricing frictions, uncertainty, and information acquisition. Specifically, we utilize the second wave of the survey, which was conducted between 2018Q1 and 2018Q2.

The sampling frame and protocol are detailed in Section II of Coibion, Gorodnichenko, Kumar, and Ryngaert (2021). The population of firms was drawn from Kompass New Zealand and Equifax, classified by broad industry using the Australia and New Zealand Standard Industrial Classification 2006 (ANZSIC06). To focus on established business enterprises, the authors state that they explicitly “focus on firms with six or more employees”. Furthermore, the sample excluded “industries related to the government, community service, agriculture, fishing and mining, energy, gas, and water,” noting that these sectors “are often dominated by a few extensively regulated firms or by very small firms”.

To ensure the sample was representative of the New Zealand economy, the survey design “targeted for two-thirds of the sample to come from professional and financial services and manufacturing because these industries account for relatively large shares of New Zealand’s GDP.” The remaining sample was drawn from trade, construction, communication, and transportation. The survey also “oversampled firms with 50-99 workers and 100+ workers” to better capture the behavior of major economic players.

Regarding data collection, Coibion, Gorodnichenko, Kumar, and Ryngaert (2021) report that “responses were collected over the phone” by research assistants who contacted general managers directly. This method ensures high data quality, as “an independent RA then listened to the recording and confirmed the accuracy of the handwritten responses.” Our final sample consists of the 515 firms that participated in this specific wave of the survey.

Measuring Time-Dependent Frictions (Time Since Last Change). To measure the duration of price spells—a key statistic in time-dependent pricing models—we utilize a direct recall question regarding the firm’s pricing history. Firms were asked (Question 8):

“When did your firm last change its price (in months) and by how much (in % change)?”

This question provides a direct measure of the backward-looking duration of the current price spell for the firm’s main product.

Measuring Subjective Uncertainty and Optimal Reset Prices. To measure the firm’s subjective uncertainty about its optimal reset price, we utilize a hypothetical question designed to elicit the firm’s desired price adjustment if menu costs were *momentarily* removed. Firms were asked (Question 7):

“If your firm was free to change its price (i.e. suppose there was no cost to renegotiating contracts with clients, no costs of reprinting catalogues, etc...) today, what probability

would you assign to each of the following categories of possible price changes the firm would make?”

Range of “Costless” Price Changes today	Probability (in %)
More than 25%:	 %
From 15 to 25%:	 %
From 10 to 15%:	 %
From 8 to 10%:	 %
From 6 to 8%:	 %
From 4 to 6%:	 %
From 2 to 4%:	 %
From 0 to 2%:	 %
From -2 to 0%:	 %
From -4 to -2%:	 %
From -6 to -4%:	 %
From -8 to -6%:	 %
From -10 to -8%:	 %
From -15 to -10%:	 %
From -25 to -15%:	 %
Less than -25%:	 %
Total (the column should sum to 100%):	100%

The phrasing of the hypothetical scenario in Question 7 specifies being free to change the price “today” and explicitly lists examples of adjustment costs (renegotiating contracts, reprinting catalogues, etc) that are removed for this decision. This phrasing aligns with a “once-off” opportunity to adjust prices (similar to the arrival of a Calvo fairy or a momentary setting of menu costs to zero) rather than a permanent regime change to flexible pricing. The question effectively elicits the distribution of the gap between the firm’s current price and their optimal reset price at the time of the survey.

We construct an estimate of the firm’s subjective uncertainty by computing the variance of the elicited probability distribution from Question 7. To calculate this variance, we assume the probability mass reported for each bin is distributed uniformly within that bin. Let J be the set of bins provided in the questionnaire. For each bin $j \in J$, let $[L_j, H_j]$ define the lower and upper bounds of the bin, and let $p_{i,j}$ be the probability assigned to bin j by firm i . The bins provided in the survey are: $(-\infty, -25)$, $[-25, -15)$, $[-15, -10)$, $\dots$, $[25, \infty)$. For the open-ended bins (“More than 25%” and “Less than -25%”), we assign a width of 10 percentage points, consistent with the widest interior bins, setting the bounds to $[25, 35]$ and $[-35, -25]$, respectively. Under the assumption

of a uniform distribution within each bin, the conditional mean of bin j is $\mu_j = \frac{L_j + H_j}{2}$ and the conditional variance is $\sigma_j^2 = \frac{(H_j - L_j)^2}{12}$. For example, if a firm assigns 100% probability to the bin “2 to 4 percent”, the implied variance is $\frac{1}{12}(4 - 2)^2 = 1/3$. The firm-level mean expected price change, $E_i[\Delta p^*]$, and the firm-level uncertainty (variance), $Var_i[\Delta p^*]$, are computed as:

$$E_i[\Delta p^*] = \sum_{j \in J} p_{i,j} \mu_j$$


$$Var_i[\Delta p^*] = \sum_{j \in J} p_{i,j} \left(\sigma_j^2 + (\mu_j - E_i[\Delta p^*])^2 \right)$$

This measure $Var_i[\Delta p^*]$ constitutes our primary empirical metric for the firm’s subjective uncertainty about the optimal price.

Figure B.1 displays the distribution of firm-level subjective uncertainty regarding optimal prices ($Var_i[\Delta p^*]$). The distribution is right-skewed, indicating that while a significant mass of firms holds a moderate degree of uncertainty (clustered near our baseline estimate of $\hat{U}^* = 1.16$), there exists a long tail of firms with very high subjective uncertainty. This heterogeneity highlights the importance of using firm-specific measures rather than aggregate proxies.


In addition to the marginal distributions, we examine the joint distribution of uncertainty and pricing history. Figure B.2 presents a binned scatter plot of the firm’s subjective uncertainty against the time elapsed since their last price change. As shown in the figure, there is a distinct positive correlation between the duration of the price spell and the firm’s uncertainty about the optimal price. Firms that have not adjusted prices for a longer period exhibit systematically higher variance in their beliefs about the optimal adjustment.

Figure B.1: Distribution of Subjective Uncertainty in the New Zealand Survey Data



Notes: This figure displays the histogram of subjective uncertainty regarding firms' desired price changes. We estimate each firm's subjective uncertainty by computing the variance of the elicited probability distribution from Question 7 in the 2018Q1 survey. To calculate this variance, we assume the probability mass reported for each bin is distributed uniformly within that bin. The histogram is constructed using 60 bins. The vertical red dashed line indicates the baseline estimated value $U^* = 1.16$. See Section 6.2 for the details of the estimate.

Figure B.2: Firms That Recently Changed Their Prices Are Less Uncertain



Notes: This figure plots the time elapsed since firms' last price changes versus firms' subjective uncertainty about their ideal price changes. The x-axis represents the time since the last price change (in months), derived from Question 8 of the survey. The y-axis represents the firm's subjective uncertainty, calculated as the variance of the probability distribution reported in Question 7. The black line is a linear fitted line and the shaded area is the 66% confidence interval. We drop outliers with implied subjective uncertainty greater than 20. The size of the bins represents the average employment of firms in each percentile.

C Extension with Menu Costs and State-Dependent Pricing

In this appendix, we extend our baseline model to feature state-dependent pricing frictions in the form of menu costs. We investigate the robustness of our theoretical and quantitative conclusions from our baseline time-dependent model to the introduction of menu costs and drift in the money supply. The main results from this analysis are that: (i) our baseline U^*/σ^2 formula (from Theorem 2) for the effects of incomplete information on the CIR remain extremely quantitatively accurate and (ii) our non-identification result (Theorem 3) continues to provide accurate guidance: economies with different costs of acquiring information have very different CIRs but quantitatively indistinguishable price change distributions. Thus, even with menu costs, data on price changes are not sufficient to identify the CIR; data on beliefs remain necessary.

C.1. The Extended Model with (Random) Menu Costs

We consider the baseline model from Section 3 and we make one major modification: instead of firms having an arbitrary time-dependent hazard rate at which they reset their prices, we consider a model with a random menu cost. In particular, we assume that there is a finite set of menu costs $\phi \in \mathcal{C}$ and that the firm's Markov transition between these states is governed by a Markov kernel $\Phi(\cdot, dt) : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{R}_+$. In this setting, we specialize firms' information acquisition by requiring them to acquire a Gaussian signal of their optimal reset price at any instant in time. This reduces the dimension of the endogenous state variable from the entire distribution governing the firms' beliefs about their optimal reset price to two state variables: the firms' perceived gap x and their uncertainty about their optimal reset price U . Because of this structure, the firm's dynamic problem can be broken down into three stages: information acquisition, pricing, and the dynamic state transition.

Stage I: Information Acquisition. Consider a firm that has a current perceived gap of $x \in \mathcal{X}$, has uncertainty about its optimal reset price of $U \in \mathcal{U}$, and at this moment in time has a menu cost of $\phi \in \mathcal{C}$. The firm's problem of how much information to acquire can be written as:

$$l(x, U, \phi) = \min_{U' \leq U} \left\{ \mathbb{E} \left[\tilde{L}(x', U', \phi) \mid x, U, U' \right] + \frac{B}{2} U' dt + \frac{\omega}{2} \ln \left(\frac{U}{U'} \right) \right\} \quad (\text{C.1})$$

where we have that:

$$x' \mid x, U, U' \sim N(x, U - U') \quad (\text{C.2})$$

Intuitively, $l : \mathcal{X} \times \mathcal{U} \times \mathcal{C} \rightarrow \mathbb{R}$ is the value of information acquisition. The solution to this problem gives rise to a policy function $u : \mathcal{X} \times \mathcal{U} \times \mathcal{C} \rightarrow \mathbb{R}_+$. Moreover, $\tilde{L} : \mathcal{X} \times \mathcal{U} \times \mathcal{C} \rightarrow \mathbb{R}$ is the value of having perceived gap x' , uncertainty U' , and menu cost ϕ when it chooses: (i) whether to reset its price and (ii) to what price it resets if it elects to reset its price.

Stage II: Pricing. Having potentially acquired information, the firm has perceived gap x , uncertainty U , and menu cost ϕ . The firm must now decide if it changes its price and what price it elects to

choose. Here, the firm faces the following problem:

$$\tilde{L}(x, U, \phi) = \min \left\{ \phi + \min_y \left\{ \frac{B}{2} y^2 dt + (1 - \rho dt) L(y, U, \phi) \right\}, \frac{B}{2} x^2 dt + (1 - \rho dt) L(x, U, \phi) \right\} \quad (\text{C.3})$$

That is, the firm either does not pay the menu cost and it suffers the loss from its current perceived gap and moves to the next period—with value $L : \mathcal{X} \times \mathcal{U} \times \mathcal{C} \rightarrow \mathbb{R}$ —or it pays the menu cost and chooses its price and sets its most preferred perceived gap. This gives rise to two policy functions $y : \mathcal{U} \times \mathcal{C} \rightarrow \mathbb{R}$ governing the reset price if the price is reset (which does not depend on the current perceived gap) and $r : \mathcal{X} \times \mathcal{U} \times \mathcal{C} \rightarrow \{0, 1\}$, a dummy variable for whether the price is reset.

Stage III: Dynamic State Transition. A firm that has finished the period by acquiring information and setting its price now has perceived gap x , uncertainty U , and menu cost ϕ . The perceived gap evolves according to $x' = x - \mu dt$, uncertainty evolves according to $U' = U + \sigma^2 dt$, and the menu cost evolves according to its Markovian structure:

$$L(x, U, \phi) = \sum_{\phi' \in \Phi} \Phi(\phi, \phi', dt) l(x - \mu dt, U + \sigma^2 dt, \phi') \quad (\text{C.4})$$

Quantitative Version. For our main quantitative analysis, we adopt the specialization described by [Alvarez, Lippi, and Oskolkov \(2022\)](#). In this framework, price reviews occur at some exogenous rate θ . Between reviews, price changes cannot take place, but at the time of the review the menu cost parameter ϕ is drawn from a distribution Λ . In this case, the evolution equation becomes:

$$L(x, U, \phi) = (1 - \theta dt) l(x - \mu dt, U + \sigma^2 dt, \infty) + \theta dt \sum_{\phi \in \Phi} \Lambda(\phi) l(x - \mu dt, U + \sigma^2 dt, \phi) \quad (\text{C.5})$$

Alternative hazards for price reviews can also be considered. Indeed, if the price review occurs every k periods and then a random menu cost is drawn, then this model would correspond to a discrete time random menu cost model, as studied by [Blanco, Boar, Jones, and Midrigan \(2024\)](#).

C.2. Solution Method

This section details the numerical algorithm used to solve the firm's optimal dynamic information and pricing policies. We solve the continuous-time model described by Equations [C.1](#), [C.2](#), and [C.5](#) using Value Function Iteration on a discretized state space. A solution to the model is a triple comprising the policy functions $u(x, U, \phi)$, $r(x, U, \phi)$, and $y(U, \phi)$. These policy functions, along with the exogenous stochastic process for the optimal reset price, determine the Markov transition of the model.

The algorithm is implemented in Julia using CUDA.jl to leverage Graphics Processing Unit (GPU) acceleration for the large-scale matrix operations required by the information choice convolution and the Markov transition matrix construction.

C.2.1. Discretization. The state space consists of the perceived price gap x , the posterior uncertainty U , and the menu cost state ϕ . We discretize the continuous state space into finite grids consistent with the model's stochastic laws of motion.

1. Price Gap Grid ($\mathcal{X}$): We construct a finite grid $\mathcal{X}$ for the perceived gap x ranging from x_{min} to x_{max} with N_x points. The grid spacing Δ_x is chosen to be consistent with the model's laws of motion. The relationship between Δ_x and the time step dt is derived from the variance of the stochastic component of the x^* process, σdW_t . The variance of the continuous process is $E[(\sigma dW_t)^2] = \sigma^2 dt$. We approximate this with a simple mean-zero binomial process which has a variance of Δ_x^2 . Equating these gives the fundamental x -grid relationship:

$$\Delta_x^2 = \sigma^2 dt \quad \text{or, equivalently} \quad \Delta_x = \sigma \sqrt{dt}$$

In the computational implementation, we utilize a full grid symmetric around zero, typically spanning $[x_{min}, x_{max}]$.

2. Uncertainty Grid ($\mathcal{U}$): The uncertainty process U has a deterministic drift of $\sigma^2 dt$ per time step. We set the uncertainty grid spacing Δ_U to be exactly equal to this drift. This links our two grid spacings:

$$\Delta_U = \sigma^2 dt \quad \text{which implies} \quad \Delta_U = \Delta_x^2 \quad (\text{C.6})$$

3. Menu Cost Grid (ϕ): In our numerical implementation, we specialize the general Markov structure to a “Calvo-plus” framework with a fixed menu cost. We define a grid consisting of two discrete states: $\Phi = \{\phi, \infty\}$. The first state, ϕ , represents the fixed menu cost the firm faces when a price review opportunity arrives. The second state is set to a sufficiently large value (approximating infinity) to effectively preclude price adjustments between reviews. The transitions between these states are governed by the exogenous Calvo arrival rate θ , such that in every period dt , the firm accesses the review state ϕ with probability θdt and defaults to the non-review state otherwise.

C.2.2. Solving the Optimal Policy Functions. We solve for the value functions and policy functions using an iterative procedure. Let $L^{(n)}(x, U, \phi)$ denote the value function at iteration n . The algorithm proceeds in the following steps within each iteration until convergence.

Step 1: Solving the Optimal Reset Price. : First, we compute the value of resetting the price. If a firm chooses to pay the menu cost, it resets its perceived gap to an optimal target $y(U, \phi)$ to minimize the sum of the immediate flow loss and the continuation value. The value of resetting, denoted $V_{\text{reset}}(U, \phi)$, is given by:

$$V_{\text{reset}}(U, \phi) = \min_{y \in \mathcal{X}} \left\{ \frac{B}{2} y^2 dt + (1 - \rho dt) L^{(n)}(y, U, \phi) \right\}$$

Computationally, this is performed by a grid search over the discretized domain of x for each uncertainty state U .

Step 2: The Pricing Decision (Constructing $\tilde{L}$). : Next, we evaluate the pricing stage value function $\tilde{L}(x, U, \phi)$, which corresponds to the firm's decision to pay the menu cost or wait. For each state (x, U, ϕ) , the firm compares the cost of adjusting (the menu cost plus the optimal reset value

calculated in Step 1) against the continuation value of maintaining the current gap:

$$\tilde{L}(x, U, \phi) = \min\{\phi + V_{\text{reset}}(U, \phi), \frac{B}{2}x^2 dt + (1 - \rho dt)L^{(n)}(x, U, \phi)\}$$

This step yields the pricing policy $r(x, U, \phi)$ and the interim value function $\tilde{L}$ used in the subsequent information stage.

Step 3: The Information Choice. : The core of the rational inattention problem is solving Equation C.1, where the firm chooses a posterior uncertainty $U' \leq U$. This involves calculating the expected value of $\tilde{L}$ over the future distribution of the perceived gap x' . According to Equation C.2, the transition of the perceived gap given a choice U' is Gaussian. To compute the expectation $E[\tilde{L}(x', U', \phi)|x, U, U']$ efficiently, we precompute a transition matrix $\tilde{Q}$ where its element $\tilde{Q}_{ij}(U, U')$ represents the probability of transitioning from gap x_i to x_j given the variance reduction $U - U'$. The expectation is computed via matrix-vector multiplication (convolution) for every feasible pair of (U, U') .

The value function for the information stage is then:

$$l(x, U, \phi) = \min_{U' \leq U} \left\{ \sum_{x'} \tilde{Q}(x, x'|U - U') \tilde{L}(x', U', \phi) + \frac{B}{2}U' dt + \frac{\omega}{2} \ln\left(\frac{U}{U'}\right) \right\}$$

The algorithm searches over all feasible U' on the uncertainty grid, $\mathcal{U}$, to find the minimum. This step is computationally intensive and is parallelized on the GPU.

Step 4: Dynamic State Transition and Drift. : Finally, we update the value function $L^{(n+1)}$ by accounting for the deterministic drift of the state variables and the stochastic arrival of price reviews.

1. **Drift:** The perceived gap x drifts by $-\mu dt$, and uncertainty U increases by $\sigma^2 dt$.

First, regarding uncertainty, as shown in Equation C.6, the grid is constructed such that the deterministic increase $\sigma^2 dt$ corresponds exactly to a one-step shift in the grid index (from U_j to U_{j+1}).

Second, regarding the perceived gap, we handle the continuous drift $-\mu dt$ using a one-sided stochastic approximation to maintain numerical stability and ensure consistency with the transition matrix construction (described in Step III of Section C.2.3). Specifically, the value at grid point x_i is computed as a weighted average of the values at x_i (staying) and x_{i-1} (drifting down). The weight on staying, denoted by p , is computed to match the first moment of the continuous drift:

$$\mathbb{E}[\Delta x] = 0 \cdot p + (-\Delta_x) \cdot (1 - p) = -\Delta_x(1 - p)$$

Equating this to the continuous drift $-\mu dt$ implies:

$$-\Delta_x(1 - p) = -\mu dt \implies p = 1 - \frac{\mu dt}{\Delta_x}$$

Substituting the grid relationship $dt = \Delta_x^2 / \sigma^2$, we obtain the probability used in the numerical

implementation:

$$p = 1 - \frac{\mu \Delta_x}{\sigma^2}. \quad (\text{C.7})$$

2. **Price Review:** We simultaneously account for the exogenous arrival of price review opportunities. Combining the drift logic derived above with the Calvo arrival rate θ , the total update equation is:

$$L^{(n+1)}(x_i, U_j) = \theta dt V_{\text{drift}}(x_i, U_{j+1}, \phi) + (1 - \theta dt) V_{\text{drift}}(x_i, U_{j+1}, \infty),$$

where $V_{\text{drift}}(x_i, U_{j+1}, \phi)$ represents the expected value after the drift step for a given menu cost state ϕ :

$$V_{\text{drift}}(x_i, U_{j+1}, \phi) = p \cdot l(x_i, U_{j+1}, \phi) + (1 - p) \cdot l(x_{i-1}, U_{j+1}, \phi)$$

Implementation and Convergence. The algorithm is initialized with a guess $L^{(0)}$. We iterate Steps 1 through 4 until the maximum absolute difference between consecutive value functions falls below a tolerance level of 10^{-5} . Given the high dimensionality of the convolution in Step 3, the code is executed on a GPU using the CUDA.jl package. The parallelization occurs over the state space (x, U, ϕ) , allowing for simultaneous computation of the optimal information choice for all states.

C.2.3. Computing Markov Transition Matrix. We now construct the Markov transition of the model. To do so, observe that we must keep track of the joint distribution of not only the firms' state variables (x, U, ϕ) but rather the joint distribution that includes the hidden state x^* , the true price gap. We therefore denote the state by $s = (x^*, x, U, \phi) \in \mathbf{S}$. With this grid structure, we have that the discretized state space is $\mathbf{S} = \mathcal{X} \times \mathcal{X} \times \mathcal{U} \times \Phi$.

Our goal is to compute the Markov transition matrix T , with elements:

$$T_{ij} = \mathbb{P}[s_{t+dt} = s_j | s_t = s_i] \quad (\text{C.8})$$

To perform this computation, we will factorize this Markov transition into four steps:

- **Step I (Information Acquisition, P):** Mirrors the firm's information acquisition. Only x and U adjust according to the information policy.
- **Step II (Pricing, Q):** Mirrors the firm's pricing decision. x and x^* adjust according to the adjustment and pricing policies.
- **Step III (Deterministic State Transition, R):** Captures the deterministic "drift" components of the exogenous transition. x and x^* drift by $-\mu dt$, and U increases by $\sigma^2 dt$.
- **Step IV (Stochastic State Transition, S):** Captures the stochastic "shock" components. x^* is hit by the mean-zero noise σdW_t , and the menu cost ϕ evolves according to its own Markov process.

The full Markov transition matrix is the product of these four steps: $T = P \times Q \times R \times S$. We now compute each of these matrices.

Step I: Information Acquisition. The firm arrives at Stage I with state $s_t = s_i$. We now see how x and U adjust at this stage:

$$\begin{aligned}
P_{ij} &= \mathbb{P}[s_t^I = s_j \mid s_t = s_i] \\
&= \mathbb{P}[x^{*I} = x^*(s_j), x^I = x(s_j), U^I = U(s_j), \phi^I = \phi(s_j) \mid s_t = s_i] \\
&= \mathbb{P}[U^I = U(s_j) \mid s_t = s_i] \mathbb{P}[x^{*I} = x^*(s_j), x^I = x(s_j), \phi^I = \phi(s_j) \mid s_t = s_i, U_t^I = U(s_j)] \\
&= \mathbb{P}[U^I = U(s_j) \mid s_t = s_i] \mathbb{P}[x^{*I} = x^*(s_j), \phi^I = \phi(s_j) \mid s_t = s_i, U_t^I = U(s_j)] \\
&\quad \times \mathbb{P}[x_t^I = x(s_j) \mid s_t = s_i, U_t^I = U(s_j), x^{*I} = x^*(s_j), \phi^I = \phi(s_j)] \\
&= \mathbb{P}[U(s_j) = u(x(s_i), U(s_i), \phi(s_i)), x^*(s_j) = x^*(s_i), \phi(s_j) = \phi(s_i)] \\
&\quad \times \mathbb{P}[x_t^I = x(s_j) \mid s_t = s_i, U_t^I = U(s_j), x^{*I} = x^*(s_j), \phi^I = \phi(s_j)]
\end{aligned}$$

We also know that $x_t^I \mid x_t, U_t, U_t^I \sim N(x_t, U_t - U_t^I)$. We need to compute the distribution of $x_t^I \mid x_t, U_t, U_t^I, x_t^*$. To do this, recall that $x_t^* = p_t - q_t$ and $x_t = p_t - \mathbb{E}_t[q_t]$, and the firm observes the signal $s_t = q_t + \tilde{\sigma}_t \varepsilon_t$, where ε_t is $N(0, 1)$ and independent of all other variables and $\tilde{\sigma}_t^2 = \frac{1}{\frac{1}{U_t^I} - \frac{1}{U_t}}$. Thus, we can observe that:

$$s_t - \mathbb{E}_t[q_t] = q_t - \mathbb{E}_t[q_t] + \tilde{\sigma}_t \varepsilon_t = x_t - x_t^* + \tilde{\sigma}_t \varepsilon_t$$

Moreover, we have that:

$$\mathbb{E}_t^I[q_t] = \mathbb{E}_t[q_t] + \kappa_t(s_t - \mathbb{E}_t[q_t]) = \mathbb{E}_t[q_t] + \kappa_t(x_t - x_t^* + \tilde{\sigma}_t \varepsilon_t)$$

where $\kappa_t = 1 - U_t^I/U_t$. And so we have that:

$$\begin{aligned}
x_t^I &= p_t - \mathbb{E}_t^I[q_t] = p_t - \mathbb{E}_t[q_t] + \kappa_t(x_t - x_t^* + \tilde{\sigma}_t \varepsilon_t) = x_t + \kappa_t(x_t^* - x_t - \tilde{\sigma}_t \varepsilon_t) \\
&= (1 - \kappa_t)x_t + \kappa_t x_t^* - \kappa_t \tilde{\sigma}_t \varepsilon_t
\end{aligned}$$

Thus, we have:

$$x_t^I \mid x_t, U_t, U_t^I, x_t^* \sim N((1 - \kappa_t)x_t + \kappa_t x_t^*, \kappa_t^2 \tilde{\sigma}_t^2)$$

We can further simplify this using our formulae for κ_t and $\tilde{\sigma}_t^2$ to obtain:

$$x_t^I \mid x_t, U_t, U_t^I, x_t^* \sim N\left(\frac{U_t^I}{U_t} x_t + \left(1 - \frac{U_t^I}{U_t}\right) x_t^*, U_t^{I^2} \left(\frac{1}{U_t^I} - \frac{1}{U_t}\right)\right)$$

We discretize the exact distribution using our grid. Define $E_{ij} = \frac{U(s_j)}{U(s_i)} x(s_i) + \left(1 - \frac{U(s_j)}{U(s_i)}\right) x^*(s_j)$ and $V_{ij} = U(s_j) \sqrt{U(s_j)^{-1} - U(s_i)^{-1}}$. We want to calculate:

$$\mathbb{P}[x_t^I = x(s_j) \mid s_t = s_i, U_t^I = U(s_j), x^{*I} = x^*(s_j), \phi^I = \phi(s_j)] = \mathbb{P}\left[E_{ij} + V_{ij} Z \in \left[x(s_j) \pm \frac{\Delta_x}{2}\right]\right]$$

where Z is $N(0, 1)$. This reduces to:

$$N_{ij} \equiv \mathbb{P} \left[E_{ij} + V_{ij}Z \in \left[x(s_j) \pm \frac{\Delta x}{2} \right] \right] = \mathbb{P} \left[V \in \left[\frac{x(s_j) - E_{ij} \pm \frac{\Delta x}{2}}{V_{ij}} \right] \right]$$

$$= \begin{cases} 1 - \Phi \left(\frac{\bar{x} - E_{ij} - \frac{\Delta x}{2}}{V_{ij}} \right), & x(s_j) = \bar{x}, \\ \Phi \left(\frac{x - E_{ij} + \frac{\Delta x}{2}}{V_{ij}} \right), & x(s_j) = \underline{x}, \\ \Phi \left(\frac{x(s_j) - E_{ij} + \frac{\Delta x}{2}}{V_{ij}} \right) - \Phi \left(\frac{x(s_j) - E_{ij} - \frac{\Delta x}{2}}{V_{ij}} \right), & \text{otherwise} \end{cases}$$

And so we have found that:

$$P_{ij} = \mathbb{I}[U(s_j) = u(x(s_i), U(s_i), \phi(s_i)), x^*(s_j) = x^*(s_i), \phi(s_j) = \phi(s_i)] N_{ij} \quad (\text{C.9})$$

Step II: Pricing. The firm arrives at Stage II in state $s_t^I = s_i$. We now see how x and x^* adjust:

$$\begin{aligned} Q_{ij} &= \mathbb{P}[s_t^{II} = s_j \mid s_t^I = s_i] \\ &= \mathbb{P}[x^{*II} = x^*(s_j), x^{II} = x(s_j), U^{II} = U(s_j), \phi^{II} = \phi(s_j) \mid s_t^I = s_i] \\ &= \mathbb{P}[U^{II} = U(s_j), \phi^{II} = \phi(s_j) \mid s_t^I = s_i] \mathbb{P}[x^{*II} = x^*(s_j), x^{II} = x(s_j) \mid s_t^I = s_i, U^{II} = U(s_j), \phi^{II} = \phi(s_j)] \\ &= \mathbb{I}[U^{II} = U(s_i), \phi^{II} = \phi(s_i)] \mathbb{P}[x^{*II} = x^*(s_j), x^{II} = x(s_j) \mid s_t^I = s_i, U^{II} = U(s_j), \phi^{II} = \phi(s_j)] \end{aligned}$$

If the firm does not adjust its price $r(x, U, \phi) = 0$, then both the true price gap and the perceived price gap remain unchanged. Otherwise, if $r(x, U, \phi) = 1$, then the price gap resets to $y(U, \phi)$. Moreover, the true gap changes by the amount that the perceived gap changes. That is: $x_t^{II} = y(U_t^{II}, \phi_t^{II})$ and $x_t^{*II} = x_t^{*I} + x_t^{II} - x_t^I$. With this, we have that:

$$\begin{aligned} &\mathbb{P}[x^{*II} = x^*(s_j), x^{II} = x(s_j) \mid s_t^I = s_i, U^{II} = U(s_j), \phi^{II} = \phi(s_j)] = \\ &\quad \mathbb{I}[x^*(s_j) = x^*(s_i), x(s_j) = x(s_i)] \mathbb{I}[r(x(s_i), U(s_j), \phi(s_j)) = 0] \\ &\quad + \mathbb{I}[x^*(s_j) = x^*(s_i) + x(s_j) - x(s_i), x(s_j) = y(U(s_j), \phi(s_j))] \mathbb{I}[r(x(s_i), U(s_j), \phi(s_j)) = 1] \end{aligned}$$

Putting this together, we have that Q_{ij} is given by:

$$\begin{aligned} Q_{ij} &= \mathbb{I}[U(s_j) = U(s_i), \phi(s_j) = \phi(s_i)] \left(\mathbb{I}[x^*(s_j) = x^*(s_i), x(s_j) = x(s_i)] \mathbb{I}[r(x(s_i), U(s_j), \phi(s_j)) = 0] \right. \\ &\quad \left. + \mathbb{I}[x^*(s_j) = x^*(s_i) + x(s_j) - x(s_i), x(s_j) = y(U(s_j), \phi(s_j))] \mathbb{I}[r(x(s_i), U(s_j), \phi(s_j)) = 1] \right) \end{aligned} \quad (\text{C.10})$$

Step III: Deterministic State Transition (Matrix R). The firm arrives in Stage III in state $s_t^{II} = s_i = (x_i^*, x_i, U_i, \phi_i)$. This stage computes the transition resulting from the deterministic components of the state variables' laws of motion. Specifically, the uncertainty U increases by $\sigma^2 dt$, while both the perceived gap x and the true gap x^* drift by $-\mu dt$. The menu cost state ϕ remains unchanged.

We approximate these transitions on our discrete grid using the same logic employed in the value function iteration (Step 4 in Section C.2.2).

- **Uncertainty (U):** As defined in our grid construction, the spacing $\Delta_U = \sigma^2 dt$. Therefore, the uncertainty evolution is exact: the state transitions from U to the next grid point $U' = \min\{U + \Delta_U, U_{max}\}$.
- **Gaps (x, x^*):** We utilize the one-sided stochastic approximation derived in Equation C.7. The system stays at the current gap values with probability $p = 1 - \frac{\mu\Delta_x}{\sigma^2}$, or drifts down by one grid point Δ_x with probability $1 - p$.

Formally, let the current state be $s_i = (x^*, x, U, \phi)$. The elements of the transition matrix $R_{ij} = \mathbb{P}[s_t^{III} = s_j | s_t^{II} = s_i]$ are defined as follows:

$$R_{ij} = \begin{cases} p & \text{if } s_j = (x^*, x, U', \phi) \\ 1 - p & \text{if } s_j = (x^* - \Delta_x, x - \Delta_x, U', \phi) \\ 0 & \text{otherwise} \end{cases}$$

where $U' = \min(U + \Delta_U, U_{max})$ handles the deterministic increase in uncertainty (bounded by the grid maximum), and boundary conditions are applied such that gap values below x_{min} are absorbed at x_{min} .

Step IV: Stochastic State Transition (Matrix S). The firm arrives in Stage IV in state $s_t^{III} = s_i = (x_i^*, x_i, U_i, \phi_i)$. This final stage applies all stochastic shocks to the system to determine the state at $t + dt$. In this step, the perceived gap x and uncertainty U remain constant. The transition $S_{ij} = \mathbb{P}[s_{t+1} = s_j | s_t^{III} = s_i]$ is determined by two independent stochastic events:

- **True Gap (x^*):** The true gap is hit by the mean-zero noise σdW_t . The new state x_j^* is drawn from the distribution of $x_i^* + \sigma dW_t$. We discretize the resulting Normal distribution, $N(x_i^*, \sigma^2 dt)$, directly onto our grid. The probability of transitioning from x_i^* to a grid point x_j^* is found by integrating this Normal distribution over the cell centered at x_j^* . For interior points, this is given by:

$$\mathbb{P}(x_j^* | x_i^*) = \Phi\left(\frac{x_j^* - x_i^* + \Delta_x/2}{\sigma\sqrt{dt}}\right) - \Phi\left(\frac{x_j^* - x_i^* - \Delta_x/2}{\sigma\sqrt{dt}}\right)$$

where Φ is the standard Normal CDF. For the boundary points x_{min} and x_{max} , we integrate the tails of the distribution $((-\infty, x_{min} + \Delta_x/2]$ and $[x_{max} - \Delta_x/2, \infty)$, respectively) to preserve total probability mass.

- **Menu Cost (ϕ):** The menu cost state for the next period, ϕ_{t+1} , is drawn from the distribution governed by the arrival rate θ , independent of the current state ϕ_t . The firm transitions to the “Price Review” state ($\phi_{t+1} = \phi$) with probability θdt , and to the “No Price Review” state ($\phi_{t+1} = \infty$) with probability $1 - \theta dt$.

The full transition probability S_{ij} for a destination state s_j (where $x(s_j) = x(s_i)$ and $U(s_j) = U(s_i)$) is the product of these two probabilities. Specifically, if s_j corresponds to a review state, $S_{ij} = \mathbb{P}(x_j^* | x_i^*) \times \theta dt$. If s_j corresponds to a no-review state, $S_{ij} = \mathbb{P}(x_j^* | x_i^*) \times (1 - \theta dt)$.

Step V: Putting it Together. Theoretically, the full Markov transition matrix T is the product of the four component matrices derived in the previous steps:

$$T = P \times Q \times R \times S. \quad (\text{C.11})$$

However, constructing T explicitly is computationally infeasible due to the high dimensionality of the state space. The total number of grid points is $N_{grid} = 2 \times N_x \times N_x \times N_U$, which in our calibration implies a transition matrix of massive dimension. While the component matrices (P, Q, R, S) are highly sparse, their product T becomes significantly denser, making storage and direct operations prohibitively expensive.

To overcome this, we employ a matrix-free approach utilizing GPU acceleration. Instead of forming T , we compute the sparse transposes of the component matrices (P^T, Q^T, R^T, S^T) and load them into GPU memory using the `CUDA.jl` and `CUDA.CUSPARSE` libraries.

- **Stationary Distribution:** The stationary distribution m^* is the eigenvector of T^T (transpose of T) associated with the unit eigenvalue. We compute this using the Power Iteration method directly on the GPU. In each iteration k , rather than performing a single multiplication with T^T , we apply the component matrices sequentially to the distribution vector m_k :

$$m_{k+1} = S^T \times (R^T \times (Q^T \times (P^T \times m_k)))$$

This sequence exploits the sparsity of each individual matrix, significantly reducing memory bandwidth and computational time. We iterate until the norm of the difference between consecutive distributions falls below a tolerance of 1×10^{-5} .

- **Impulse Response Functions (IRFs):** We compute the dynamic response of the aggregate price level to shocks by iterating the distribution forward from a perturbed initial state. Let m^* be the stationary distribution.
 - **Belief Gap Shock:** We define the initial perturbed distribution m_0 by shifting the mass of m^* along the true gap dimension (x^*) by an amount δ . This corresponds to an instantaneous change in the fundamental x^* while beliefs x remain fixed.
 - **Perceived Gap Shock:** We define m_0 by shifting the mass of m^* along both the true gap (x^*) and perceived gap (x) dimensions by δ . This captures a shock where firms perfectly observe the change in the fundamental.

Starting from m_0 , the distribution evolves according to $m_t = T^T m_{t-dt}$, computed via the same sequential GPU matrix-vector multiplication described above. The impulse response of the aggregate price gap Y_t is calculated as the deviation of the mean gap from the steady state:

$$Y_t = \int x^* dm_t(x^*) - \int x^* dm^*(x^*)$$

- **Cumulative Impulse Responses (CIR):** Finally, the CIR, which measures the total degree of monetary non-neutrality, is computed by integrating the IRF over time. We approximate this

integral using the trapezoidal rule:

$$\text{CIR} \approx \sum_{t=0}^{T_{hor}} \left(\frac{Y_t + Y_{t+dt}}{2} \right) dt$$

C.3. Calibration

We calibrate the model at a monthly frequency. The model parameters are divided into two categories: those assigned *a priori* based on standard values or our baseline estimation, and those calibrated to match specific moments of the data involving price setting and uncertainty.

Assigned Parameters:. Table C.1 summarizes the fixed parameters. We set the time discount rate ρ to be consistent with an annual discount factor of 0.96. The curvature of the profit function is set to $B = 12$, which corresponds to a steady-state markup of 33% for a CES demand elasticity of $\eta = 4$. We assume zero drift in the fundamental price process ($\mu = 0$) for the baseline calibration.

The volatility of the fundamental shocks, σ^2 , is set to 0.206, consistent with our baseline estimation in Section 6.2. The time step for the discretization is set to $dt = 0.05$. The state space for the perceived gap x is bounded by $x_{max} = 7.5\%$ with a grid spacing of $\Delta_x = \sigma\sqrt{dt} \approx 0.101\%$ ($N_x = 147$). The grid for uncertainty U spans $[\sigma^2, 35\sigma^2]$ (up to a maximum of 7.2) with a spacing of $\Delta_U = \sigma^2 dt \approx 0.0103$ ($N_U = 681$).

Calibrated Parameters:. There are three remaining parameters to calibrate: the information cost parameter ω , the Calvo arrival rate of price reviews θ , and the fixed menu cost ϕ . We calibrate these parameters to minimize the distance between model-generated moments and their empirical counterparts. Specifically, we target the following four statistics:

- Average Uncertainty at Reset (U^*): We target an empirical value of 1.16.
- Average Duration of Price Changes: We target a duration of 5.95 months.
- Price Adjustment Probability: We target a ratio of price change frequency to price review frequency of 0.5, consistent with the findings of Alvarez, Lippi, and Paciello (2018).
- Distribution of Uncertainty: We minimize the L^2 distance between the model-implied distribution of uncertainty and the empirical distribution derived from the survey data.

The resulting calibrated values are $\omega = 151.25$, $\theta = 0.3267$, and $\phi = 1.76$.

Simulation Procedure:. To compute the model moments for the calibration, we simulate a panel of $N = 10,000$ firms over a horizon of $T = 60$ months. The simulation proceeds in two steps:

- **Burn-in:** We initialize firms and simulate the economy for a sufficiently long period (minimum 2000 time steps) to eliminate the influence of initial conditions and ensure the distribution of firms has converged to the stationary distribution.
- **Measurement:** We compute cross-sectional statistics (duration, uncertainty, adjustment frequency) over the final 60-month window. The firms' uncertainty U evolves deterministically between information updates and jumps upon optimal information acquisition. The perceived gap x follows the Kalman filter dynamics. Price reviews occur according to the

Table C.1: Model Calibration

Parameter	Value	Description	Source/Target
<i>Panel A. Assigned Parameters</i>			
ρ	0.0034	Discount Rate	Annual discount of 4%
B	12.0	Profit curvature	33% steady-state markup
μ	0.0	Drift	Baseline assumption
σ^2	0.206	Volatility	Baseline value
<i>Panel B. Calibrated Parameters</i>			
ω	151.25	Information cost	U^* and uncertainty distribution
θ	0.3267	Review rate	Ratio of price review and change frequencies
ϕ	1.76	Menu cost	Price change duration

Notes: This table shows the quantitative model calibration with menu costs and state-dependent pricing.

Table C.2: Model Fit

Moment	Data	Model
Ratio of price review and change frequencies ($N(\Delta p)/N(Review)$)	0.50	0.50
Price change duration ($\bar{D}$)	5.95	5.95
Optimal reset uncertainty (U^*)	1.16	1.16
P25 uncertainty distribution	1.08	1.39
P50 uncertainty distribution	1.33	1.66
P75 uncertainty distribution	2.84	2.10

Notes: The table shows the set of moments that were targeted for calibration.

Poisson rate θ . Upon a review, firms evaluate the option to pay the menu cost ϕ and reset prices based on the value functions solved in Section C.2.2.

C.4. Additional Quantitative Results

In this section, we provide additional quantitative results on the menu cost model with drift $\mu > 0$ and correspondingly positive steady-state inflation that we did not have space to summarize in the main text.

We continue to find that the basic formula from Theorem 2 remains highly quantitatively accurate for a wide range of drift parameters and even for large monetary shocks. Specifically, Figure C.1 shows the sensitivity of key model statistics and CIRs to variations in the menu cost parameter, ϕ , within a model featuring a positive drift. Figure C.2 shows the same comparative statics as we vary the drift parameter. Figure C.3 shows the sensitivity of the model's CIRs to variations in the magnitude of the aggregate shock, δ . Figure C.4 shows the sensitivity of the model's CIRs to variations in the magnitude of the aggregate shock, δ , within a model featuring a positive drift.




Figure C.1: Comparative Statics with Respect to Menu Cost (ϕ) under Positive Drift

Notes: This figure shows the sensitivity of key model statistics and CIRs to variations in the menu cost parameter, ϕ , within a model featuring a positive drift $\mu = 2.0/12$, corresponding to a 2% annual trend inflation rate. The left panel shows the CIR to a perceived gap shock, $\mathcal{M}^x$ (blue line), plotted alongside the average duration of price changes, $\bar{D}$ (orange line). The middle panel compares the difference between the CIR to a belief shock and the CIR to a perceived gap shock, $\mathcal{M}^b - \mathcal{M}^x$ (blue line) against the baseline formula U^*/σ^2 (orange line), where U^* is the average posterior uncertainty of price-adjusting firms. The right panel shows the ratio of the frequency of realized price changes to the frequency of price reviews ($N(|\Delta p| > 0)/N(\text{Review})$).




Figure C.2: Comparative Statics with Respect to Trend Inflation (μ)

Notes: This figure shows the sensitivity of key model statistics and CIRs to variations in the annual trend inflation rate (determined by the drift parameter μ). The left panel shows the CIR to a perceived gap shock, $\mathcal{M}^x$ (blue line), plotted alongside the average duration of price changes, $\bar{D}$ (orange line). The middle panel compares the difference between the CIR to a belief shock and the CIR to a perceived gap shock, $\mathcal{M}^b - \mathcal{M}^x$ (blue line) against the baseline formula U^*/σ^2 (orange line), where U^* is the average posterior uncertainty of price-adjusting firms. The right panel shows the ratio of the frequency of realized price changes to the frequency of price reviews ($N(|\Delta p| > 0)/N(\text{Review})$).




Figure C.3: Comparative Statics with Respect to Shock Size (δ)

Notes: This figure shows the sensitivity of the model's CIRs to variations in the magnitude of the aggregate shock, δ . The left panel plots the CIR to a perceived gap shock, $\mathcal{M}^x$. The middle panel plots the CIR to a belief gap shock, $\mathcal{M}^b$. The right panel compares the difference between the belief shock CIR and the perceived gap shock CIR, $\mathcal{M}^b - \mathcal{M}^x$ (blue line) against the baseline formula U^* / σ^2 (orange line), where U^* is the average posterior uncertainty of price-adjusting firms.




Figure C.4: Comparative Statics with Respect to Shock Size (δ) under Positive Drift

Notes: This figure shows the sensitivity of the model's CIRs to variations in the magnitude of the aggregate shock, δ , within a model featuring a positive drift $\mu = 2.0/12$, corresponding to a 2% annual trend inflation rate. The left panel plots the CIR to a perceived gap shock, $\mathcal{M}^x$. The middle panel plots the CIR to a belief gap shock, $\mathcal{M}^b$. The right panel compares the difference between the belief shock CIR and the perceived gap shock CIR, $\mathcal{M}^b - \mathcal{M}^x$ (blue line) against the baseline formula U^* / σ^2 (orange line), where U^* is the average posterior uncertainty of price-adjusting firms.

D Extension to Allow for Strategic Complementarities

In our baseline model, we made assumptions on preferences such that there is no strategic complementarity in price-setting. As is well-known, however, strategic complementarity can arise from: (i) the presence of income effects in labor supply and/or money demand, (ii) non-isoelastic demand (*e.g.*, [Kimball \(1995\)](#) demand). To understand how the presence of strategic complementarities affects our results, we generalize our setting to allow for them.

D.1. Extended Model with Strategic Complementarities

We introduce strategic complementarities following [Alvarez, Lippi, and Souganidis \(2023\)](#) and model the process of firms' optimal (log) prices as:

$$\tilde{q}_{i,t} = (1 - \beta) q_{i,t} + \beta p_t \quad (\text{D.1})$$

where $\beta \in (-\infty, 1)$ indexes the degree of strategic complementarities in price-setting. Our baseline model is nested as $\beta = 0$. It is well-known that this reduced-form structure arises from [Kimball \(1995\)](#) demand. In the interests of space, we omit the exact derivations and refer the reader to [Alvarez, Lippi, and Souganidis \(2023\)](#). We assume that firms can observe the aggregate price level when they make their price-setting decisions. The rest of our model is as in Section 2.

D.2. The Result

In what follows we show, while not obvious, that strategic complementarities have no interaction with the effects of incomplete information. That is, in Theorem 4, we establish that the CIR to a one-time monetary policy shock at date 0 of size δ is given by:

$$\text{CIR}(\delta) = \delta \left(\frac{\bar{D}}{1 - \beta} + \frac{U^*}{\sigma^2} \right) \quad (\text{D.2})$$

Thus, strategic complementarities interact with nominal rigidities in the usual way: they slow price adjustment and increase the real effects of monetary policy. However, quite surprisingly, there is no interplay between them and the effects of incomplete information.¹⁶

D.3. Analysis

As in our main analysis, we are interested in the CIR to a permanent, one-time, and unanticipated monetary shock of size $\delta \in \mathbb{R}$ at time t . The presence of strategic complementarities complicates things as we now need to solve for the entire impulse response function of prices following the shock. We break this analysis into two parts. First, we solve for the price path of an individual firm

¹⁶As an additional note, we observe that we have an amplification factor $1/(1 - \beta)$ rather than the factor $1/\sqrt{1 - \beta}$. The reason for this is simple: we are considering an unobserved “MIT” shock to the money supply. Thus, firms have rational expectations that the price will be unchanged going forward. Hence, their behavior in the face of strategic complementarities is backward looking. All of the analysis we have done here could also be done assuming perfect foresight of the future price path and we would obtain the factor $1/\sqrt{1 - \beta}$. However, we emphasize that for this counterfactual with incomplete information, such a perfect foresight solution is not the rational expectations solution.

and aggregate it to obtain a fixed point equation for the aggregate price path. Second, we solve for the fixed point to find the equilibrium price path.

Step I: The Fixed Point Equation for the Aggregate Price Path. Consider a generic firm i . At time t , they last reset their price at time h^0 . The number of price resets they have performed up to time $\tau > t$ is given by $N_{i,\tau}$. We let h^j be the time of the j -th price reset, where $j \in \{1, \dots, N_{i,\tau}\}$.

Lemma D.1. *For an unobserved monetary policy shock at time t of size δ , the aggregate price path solves the following equation:*

$$p_\tau - p_t = (1 - \beta)\delta \mathbb{E}_i \left[\sum_{j=1}^{N_{i,\tau}} \kappa_{h^j - h^{j-1}} \prod_{k=1}^{j-1} (1 - \kappa_{h^k - h^{k-1}}) \right] + \beta \mathbb{E}_i [p_{h^{N_{i,\tau}}} - p_t] \quad (\text{D.3})$$

Proof. We have, by definition, that:

$$p_{i,\tau} - p_{i,t} = \sum_{j=1}^{N_{i,\tau}} (p_{i,h^j} - p_{i,h^{j-1}}) \quad (\text{D.4})$$

To describe this object, we therefore need to describe $p_{i,h^j} = \mathbb{E}_{i,h^j}[\tilde{q}_{i,h^j}]$. As the monetary shock is an unanticipated “MIT” shock, firms attach zero probability to it. Thus, firms’ optimal information and pricing policies are as in Theorem 1 with the minor modification that firms set their price equal to the conditional expectation of $\tilde{q}$ (incorporating strategic complementarities) rather than q $p_{i,h^j} = \mathbb{E}_{i,h^j}[\tilde{q}_{i,t}]$. Moreover, as firms observe the aggregate price level, adapting the arguments of Proposition 1 we have that:

$$\begin{aligned} p_{i,h^j} &= (1 - \beta)\mathbb{E}_{i,h^j}[q_{i,h^j}] + \beta p_{h^j} \\ &= (1 - \beta)(\kappa_{h^j - h^{j-1}} q_{i,h^j} + (1 - \kappa_{h^j - h^{j-1}})\mathbb{E}_{i,h^{j-1}}[q_{i,h^{j-1}}]) + \beta p_{h^j} + \epsilon_{i,h^j} \end{aligned} \quad (\text{D.5})$$

where ϵ_{i,h^j} is a zero-mean random variable. We now rearrange this object to solve for firms’ price changes:

$$\begin{aligned} p_{i,h^j} - p_{i,h^{j-1}} &= (1 - \beta)(\kappa_{h^j - h^{j-1}} q_{i,h^j} + (1 - \kappa_{h^j - h^{j-1}})\mathbb{E}_{i,h^{j-1}}[q_{i,h^{j-1}}] - \mathbb{E}_{i,h^{j-1}}[q_{i,h^{j-1}}]) + \beta(p_{h^j} - p_{h^{j-1}}) + \tilde{\epsilon}_{i,h^j} \\ &= (1 - \beta)\kappa_{h^j - h^{j-1}}(q_{i,h^j} - \mathbb{E}_{i,h^{j-1}}[q_{i,h^{j-1}}]) + \beta(p_{h^j} - p_{h^{j-1}}) + \tilde{\epsilon}_{i,h^j} \\ &= (1 - \beta)\kappa_{h^j - h^{j-1}}(q_{i,h^j} - \mathbb{E}_{i,h^j}[q_{i,h^j}] + \mathbb{E}_{i,h^j}[q_{i,h^j}] - \mathbb{E}_{i,h^{j-1}}[q_{i,h^{j-1}}]) + \beta(p_{h^j} - p_{h^{j-1}}) + \tilde{\epsilon}_{i,h^j} \end{aligned} \quad (\text{D.6})$$

where $\tilde{\epsilon}_{i,h^j}$ is a zero-mean random variable. We note that:

$$(1 - \beta)(\mathbb{E}_{i,h^j}[q_{i,h^j}] - \mathbb{E}_{i,h^{j-1}}[q_{i,h^{j-1}}]) = p_{i,h^j} - p_{i,h^{j-1}} - \beta(p_{h^j} - p_{h^{j-1}}) \quad (\text{D.7})$$

and so we have that:

$$(p_{i,h^j} - p_{i,h^{j-1}})(1 - \kappa_{h^j - h^{j-1}}) = (1 - \beta)\kappa_{h^j - h^{j-1}}(q_{i,h^j} - \mathbb{E}_{i,h^j}[q_{i,h^j}]) + \beta(1 - \kappa_{h^j - h^{j-1}})(p_{h^j} - p_{h^{j-1}}) + \tilde{\epsilon}_{i,h^j} \quad (\text{D.8})$$

which implies that:

$$p_{i,hj} - p_{i,hj-1} = (1 - \beta) \frac{\kappa_{hj-hj-1}}{1 - \kappa_{hj-hj-1}} b_{i,hj} + \beta(p_{hj} - p_{hj-1}) + \hat{\epsilon}_{i,hj} \quad (D.9)$$

where $\hat{\epsilon}_{i,hj}$ is a zero-mean random variable and $b_{i,hj} = q_{i,hj} - \mathbb{E}_{i,hj}[q_{i,hj}]$ is the belief gap.

We now describe the law of motion of the belief gap. As before, we have that:

$$\mathbb{E}_{i,hj}[q_{i,hj}] = \kappa_{hj-hj-1} q_{i,hj} + (1 - \kappa_{hj-hj-1}) \mathbb{E}_{i,hj-1}[q_{hj-1}] + \varepsilon_{i,hj} \quad (D.10)$$

where $\varepsilon_{i,hj}$ is a zero-mean random variable. Thus, we obtain:

$$\begin{aligned} b_{i,hj} &= q_{i,hj} - \mathbb{E}_{i,hj}[q_{i,hj}] = q_{i,hj} - \kappa_{hj-hj-1} q_{i,hj} - (1 - \kappa_{hj-hj-1}) \mathbb{E}_{i,hj-1}[q_{hj-1}] - \varepsilon_{i,hj} \\ &= (1 - \kappa_{hj-hj-1}) (q_{i,hj} - \mathbb{E}_{i,hj-1}[q_{hj-1}]) - \varepsilon_{i,hj} \\ &= (1 - \kappa_{hj-hj-1}) b_{i,hj-1} - \varepsilon_{i,hj} + (1 - \kappa_{hj-hj-1}) (q_{i,hj} - q_{i,hj-1}) \end{aligned} \quad (D.11)$$

Thus, we have that:

$$b_{i,hj} = b_{i,t} \prod_{k=1}^j (1 - \kappa_{h^k-h^{k-1}}) + \hat{\epsilon}_{i,hj} \quad (D.12)$$

where $\hat{\epsilon}_{i,hj}$ is a zero-mean noise term (as $q_{i,hj} - q_{i,hj-1}$ is the increment of a Brownian motion, it is zero mean). Combining Equations D.9 and D.12, we obtain that:

$$\begin{aligned} p_{i,hj} - p_{i,hj-1} &= (1 - \beta) b_{i,t} \frac{\kappa_{hj-hj-1}}{1 - \kappa_{hj-hj-1}} \prod_{k=1}^j (1 - \kappa_{h^k-h^{k-1}}) + \beta(p_{hj} - p_{hj-1}) + \check{\epsilon}_{i,hj} \\ &= (1 - \beta) b_{i,t} \left[\kappa_{hj-hj-1} \prod_{k=1}^{j-1} (1 - \kappa_{h^k-h^{k-1}}) \right] + \beta(p_{hj} - p_{hj-1}) + \check{\epsilon}_{i,hj} \end{aligned} \quad (D.13)$$

where $\check{\epsilon}_{i,hj}$ is a zero mean noise term. Combining this with Equation D.4, we have that:

$$\begin{aligned} p_{i,\tau} - p_{i,t} &= \sum_{j=1}^{N_{i,\tau}} \left((1 - \beta) b_{i,t} \left[\kappa_{hj-hj-1} \prod_{k=1}^{j-1} (1 - \kappa_{h^k-h^{k-1}}) \right] + \beta(p_{hj} - p_{hj-1}) \right) + \epsilon_{i,hj}^\circ \\ &= (1 - \beta) b_{i,t} \sum_{j=1}^{N_{i,\tau}} \kappa_{hj-hj-1} \prod_{k=1}^{j-1} (1 - \kappa_{h^k-h^{k-1}}) + \beta(p_{h^{N_{i,\tau}}} - p_t) + \epsilon_{i,hj}^\circ \end{aligned} \quad (D.14)$$

where $\epsilon_{i,hj}^\circ$ is a zero-mean noise term. Thus, taking the mean across firms, we have that:

$$p_\tau - p_t = (1 - \beta) \mathbb{E}_i \left[b_{i,t} \sum_{j=1}^{N_{i,\tau}} \kappa_{hj-hj-1} \prod_{k=1}^{j-1} (1 - \kappa_{h^k-h^{k-1}}) \right] + \beta \mathbb{E}_i[p_{h^{N_{i,\tau}}} - p_t] \quad (D.15)$$

For an unobserved monetary shock we have that $\mathbb{E}_i[b_{i,t}] = \delta$ and $b_{i,t}$ is uncorrelated with $\sum_{j=1}^{N_{i,\tau}} \kappa_{hj-hj-1} \prod_{k=1}^{j-1} (1 - \kappa_{h^k-h^{k-1}})$. This completes the proof. $\blacksquare$

Step II: Solving the Fixed Point Equation. We now solve the fixed point equation derived in Lemma D.1. To do this, we define the object:

$$\Omega(\tau) = \mathbb{E}_i \left[\sum_{j=1}^{N_{i,\tau}} \kappa_{hj-hj-1} \prod_{k=1}^{j-1} (1 - \kappa_{h^k-h^{k-1}}) \right] \quad (D.16)$$

which captures the aggregate fraction of the initial belief gap induced by the monetary policy shock that has been eliminated by time τ . We then have that the aggregate price path solves a non-autonomous integral equation:

Lemma D.2. *For an unobserved monetary policy shock at time t of size δ , the aggregate price path solves the following non-autonomous and linear integral (Volterra) equation:*

$$p_\tau - p_t = (1 - \beta)\delta\Omega(\tau) + \beta \int_t^\tau (p_k - p_t)f(\tau - k)dk \quad (\text{D.17})$$

where p_t is given.

Proof. We begin with Lemma D.1 and exploit the fact that $\tau - h^{N_{i,\tau}}$ equals the duration of the price of firm i at time τ , which has distribution F with density f . Thus:

$$\begin{aligned} \mathbb{E}_i [p_{h^{N_{i,\tau}}} - p_t] &= \int_0^\infty p_{\tau-D}f(D)dD - p_t = \int_{-\infty}^\tau p_k f(\tau - k)dk - p_t \\ &= \int_t^\tau (p_k - p_t)f(\tau - k)dk \end{aligned} \quad (\text{D.18})$$

Using the definition of Ω and plugging back into Lemma D.1 yields the result. $\blacksquare$

We can now additionally use this fact to solve for the path of real GDP: $y_\tau - y_t = \delta - (p_\tau - p_t)$. In what follows, we normalize $t = 0$ and $y_0 = p_0 = 0$. We obtain the following:

Lemma D.3. *For an unobserved monetary policy shock at time 0 of size δ , the path of real GDP solves the following non-autonomous and linear integral (Volterra) equation:*

$$y_\tau = \delta \left((1 - \beta)(1 - \Omega(\tau)) + \beta(1 - F(\tau)) \right) + \beta \int_0^\tau y_k f(\tau - k)dk \quad (\text{D.19})$$

Proof. We use the fact that $y_\tau = \delta - p_\tau$ and Lemma D.2 to write:

$$\begin{aligned} y_\tau &= \delta - (1 - \beta)\delta\Omega(\tau) - \beta \int_0^\tau p_k f(\tau - k)dk \\ &= (1 - \beta)\delta(1 - \Omega(\tau)) + \beta \left(\delta - \int_0^\tau p_k f(\tau - k)dk \right) \\ &= (1 - \beta)\delta(1 - \Omega(\tau)) + \beta \left(\delta \left[1 - \int_0^\tau f(\tau - k)dk \right] + \int_0^\tau (\delta - p_k)f(\tau - k)dk \right) \\ &= \delta \left((1 - \beta)(1 - \Omega(\tau)) + \beta(1 - F(\tau)) \right) + \beta \int_0^\tau y_k f(\tau - k)dk \end{aligned} \quad (\text{D.20})$$

Completing the proof. $\blacksquare$

We now use this result to solve for the cumulative impulse response to an unobserved and permanent monetary shock of size δ at time 0: $\text{CIR}(\delta) = \int_0^\infty y_\tau d\tau$.

Theorem 4. *The following formula holds:*

$$\text{CIR}(\delta) = \delta \left(\frac{\bar{D}}{1 - \beta} + \frac{U^*}{\sigma^2} \right) \quad (\text{D.21})$$

Proof. We begin with Lemma D.3 and integrate to obtain:

$$\text{CIR}(\delta) = \delta(1 - \beta) \int_0^\infty (1 - \Omega(\tau)) d\tau + \delta\beta \int_0^\infty (1 - F(\tau)) d\tau + \beta \int_0^\infty \int_0^\tau y_k f(\tau - k) dk d\tau \quad (\text{D.22})$$

By Theorem 2, we have that $\int_0^\infty (1 - \Omega(\tau)) d\tau = \bar{D} + \frac{U^*}{\sigma^2}$. We also have for any CDF that the area above the CDF equals the mean of the distribution: $\int_0^\infty (1 - F(\tau)) d\tau = \bar{D}$. We can also compute the final double integral as:

$$\begin{aligned} \int_0^\infty \int_0^\tau y_k f(\tau - k) dk d\tau &= \int_0^\infty \int_k^\infty y_k f(\tau - k) d\tau dk = \int_0^\infty y_k \int_k^\infty f(\tau - k) d\tau dk \\ &= \int_0^\infty y_k dk = \text{CIR}(\delta) \end{aligned} \quad (\text{D.23})$$

And so we have that:

$$\text{CIR}(\delta) = \delta(1 - \beta) \left(\bar{D} + \frac{U^*}{\sigma^2} \right) + \delta\beta\bar{D} + \beta\text{CIR}(\delta) \quad (\text{D.24})$$

Re-arranging for $\text{CIR}(\delta)$ then yields the claimed formula. $\blacksquare$

With this, we have generalized our main analysis to allow for strategic complementarities. We have also found that the contribution of incomplete information to the CIR is completely unchanged, highlighting the robustness of the U^*/σ^2 formula to this extension.

E Extension to Allow for Richer Learning Procedures

In our baseline model, we assumed that firms receive no free information and that their cost of information acquisition is linear in the flow of the mutual information induced by the signals that they observe. These assumptions led to Theorem 1 and had the consequence that firms: (i) never learn between price-setting events and (ii) always reset to a state-invariant level of posterior uncertainty U^* when they reset their price. When proceeding to our analysis of the effects of monetary policy, this led to particularly simple analysis because it implied that the Kalman gains were always of the form $\kappa_\tau = \sigma^2 \tau / (U^* + \sigma^2 \tau)$. In this appendix, we generalize our results to a family of richer learning procedures and we characterize the CIR for such more general technologies.

E.1. The Extended Model with Flexible Learning Rules

As throughout our analysis, we need to characterize how firms' beliefs evolve following an unanticipated and permanent monetary shock of size δ at time t . Rather than derive the process for firms' belief evolution from specific microfoundations, as we did in the general analysis, we instead adopt a reduced-form perspective in which firms have a generalized linear learning technology. We take this reduced form perspective to understand how the mechanics of firms' belief updating influences the CIR.

Specifically, we assume that firm belief updating follows the process:

$$\mathbb{E}_{i,t+h'}[q_{i,t+h'}] - \mathbb{E}_{i,t}[q_{i,t}] = \kappa(h, h')(q_{i,t+h'} - \mathbb{E}_{i,t}[q_{i,t}]) + \varepsilon_i \quad (\text{E.1})$$

for each firm i who first reset their price h' time after the monetary shock and whose pricing spell duration was h at the time the monetary shock took place, where ε_i is zero mean. For all subsequent spells, the relevant Kalman gain is $\kappa(0, h^n - h^{n-1})$ where $t + h^n$ is the time of the $n - th$ price change after the monetary shock. The Kalman gain $\kappa(h, h')$ represents how quickly firms' beliefs update to reflect the true state of the world after the monetary shock hits.

This specification nests that of our baseline model, with $\kappa(h, h') = \sigma^2(h + h') / (U^* + \sigma^2(h + h'))$, and further accommodates a simple alternative model in which firms receive free Gaussian signals between price changes, so long as this free information never pushes the firm's posterior uncertainty below U^* .¹⁷

E.2. Analysis

We now generalize the arguments underlying Theorem 2 to characterize the effects of incomplete information under this general learning process:

Theorem 5. *Under the learning process described by Equation E.1, we have that:*

$$\mathcal{M}^b - \mathcal{M}^x = \frac{\int_0^\infty g(\tau) \left[\int_0^\tau (1 - \kappa(h, \tau - h)) dh \right] d\tau}{\int_0^\infty g(\tau) \kappa(0, \tau) d\tau} \quad (\text{E.2})$$

Proof. By definition, we have that the belief gap at the time of the price reset $t + h'$ is given for firm i by:

$$\begin{aligned} y_{i,t+h'}^b &= \frac{1}{\gamma} [q_{i,t+h'} - \mathbb{E}_{i,t+h'}[q_{i,t+h'}]] \\ &= \frac{1}{\gamma} [q_{i,t+h'} - q_{i,t} + q_{i,t} - \mathbb{E}_{i,t}[q_{i,t}] + \mathbb{E}_{i,t}[q_{i,t}] - \mathbb{E}_{i,t+h'}[q_{i,t+h'}]] \\ &= y_{i,t}^b - \frac{1}{\gamma} (\mathbb{E}_{i,t+h'}[q_{i,t+h'}] - \mathbb{E}_{i,t}[q_{i,t}]) + \frac{1}{\gamma} (q_{i,t+h'} - q_{i,t}) \end{aligned} \quad (\text{E.3})$$

By using Equation E.1, we can rewrite the second term in the above as follows:

$$\begin{aligned} \mathbb{E}_{i,t+h'}[q_{i,t+h'}] - \mathbb{E}_{i,t}[q_{i,t}] &= \kappa(h, h') (q_{i,t+h'} - \mathbb{E}_{i,t}[q_{i,t}]) + \varepsilon_{i,t+h'} \\ &= \kappa(h, h') (q_{i,t+h'} - q_{i,t} + q_{i,t} - \mathbb{E}_{i,t}[q_{i,t}]) + \varepsilon_{i,t+h'} \\ &= \kappa(h, h') (q_{i,t} - \mathbb{E}_{i,t}[q_{i,t}]) + \kappa(h, h') (q_{i,t+h'} - q_{i,t}) + \varepsilon_{i,t+h'} \end{aligned} \quad (\text{E.4})$$

Replacing this expression, we then have that:

$$y_{i,t+h'}^b = (1 - \kappa(h, h')) y_{i,t}^b + Z_{i,t+h'} \quad (\text{E.5})$$

where $Z_{i,t+h'} = \frac{1}{\gamma} (1 - \kappa(h, h')) (q_{i,t+h'} - q_{i,t}) - \frac{1}{\gamma} \varepsilon_{i,t+h'}$. Using the fact that $\mathbb{E}[q_{i,t+h'} - q_{i,t} | h, h', \mathbb{E}_{i,t}[q_{i,t}], q_{i,t}] = 0$, we have that $\mathbb{E}[Z_{i,t+h'} | h, h', \mathbb{E}_{i,t}[q_{i,t}], q_{i,t}] = 0$.

¹⁷This is because, by the same arguments as Theorem 1, the firm always resets to U^* regardless of its posterior uncertainty, so long as uncertainty remains above U^* .

We can now proceed to recursively characterize expected lifetime output gaps:

$$Y(y^b, y^x, h) = \mathbb{E}^{h', Z} \left[\int_0^{h'} y^b d\tau + \int_0^{h'} y^x d\tau + Y \left((1 - \kappa(h, h')) y^b + Z_{h'}, 0, 0 \right) \right] \quad (\text{E.6})$$

We now guess and verify that $Y(y^b, y^x, h) = \beta(h) y^x + m(h) y^b$. Plugging this guess into Equation E.6 and matching coefficients, we obtain that $\beta(h)$ and $m(h)$ must satisfy:

$$\beta(h) = \mathbb{E}_g[h' | h] = \bar{D}_h \quad (\text{E.7})$$

$$m(h) = \mathbb{E}_g[h' | h] + m(0) \mathbb{E}_g^{h'}[1 - \kappa(h, h') | h] = \bar{D}_h + m(0)(1 - \bar{\kappa}_h) \quad (\text{E.8})$$

$$m(0) = \frac{\mathbb{E}_g[h']}{1 - \mathbb{E}_g[1 - \kappa(h, h')]} = \bar{D}_0 \frac{1}{\bar{\kappa}_0} \quad (\text{E.9})$$

We then have that:

$$\mathcal{M}(\mathcal{F}) = \mathbb{E}_{\mathcal{F}}[y^x] \bar{D} + \mathbb{E}_{\mathcal{F}}[y^b] \left(\bar{D} + \bar{D}_0 \frac{1 - \bar{\kappa}}{\bar{\kappa}_0} \right) \quad (\text{E.10})$$

and so we have that:

$$\begin{aligned} \mathcal{M}^b - \mathcal{M}^x &= \bar{D}_0 \frac{1 - \bar{\kappa}}{\bar{\kappa}_0} \\ &= \bar{D}_0 \frac{\int_0^\infty \int_0^\infty (1 - \kappa(h, h')) \frac{g(h+h')}{1-G(h)} dh' f(h) dh}{\int_0^\infty \kappa(0, h') g(h') dh'} \end{aligned} \quad (\text{E.11})$$

Using the fact that $f(h) = \frac{1}{\bar{D}_0} (1 - G(h))$ and applying a change of variable, we can then simplify this:

$$\begin{aligned} \mathcal{M}^b - \mathcal{M}^x &= \frac{\int_0^\infty \int_0^\infty (1 - \kappa(h, h')) g(h+h') dh' dh}{\int_0^\infty \kappa(0, h') g(h') dh'} \\ &= \frac{\int_0^\infty g(\tau) \left[\int_0^\tau (1 - \kappa(h, \tau - h)) dh \right] d\tau}{\int_0^\infty g(\tau) \kappa(0, \tau) d\tau} \end{aligned} \quad (\text{E.12})$$

As claimed. ■

Theorem 2 follows as a direct corollary from this result. To see this, if we replace $\kappa(h, h') = \sigma^2(h + h') / (U^* + \sigma^2(h + h'))$, then we have that:

$$\begin{aligned} \mathcal{M}^b - \mathcal{M}^x &= \frac{\int_0^\infty g(\tau) \left[\int_0^\tau (1 - \kappa(h, \tau - h)) dh \right] d\tau}{\int_0^\infty g(\tau) \kappa(0, \tau) d\tau} \\ &= \frac{\int_0^\infty g(\tau) \left[\int_0^\tau \frac{U^*}{U^* + \sigma^2 \tau} dh \right] d\tau}{\int_0^\infty g(\tau) \frac{\sigma^2 \tau}{U^* + \sigma^2 \tau} d\tau} \\ &= \frac{U^*}{\sigma^2} \frac{\int_0^\infty g(\tau) \frac{\tau}{U^* + \sigma^2 \tau} d\tau}{\int_0^\infty g(\tau) \frac{\tau}{U^* + \sigma^2 \tau} d\tau} = \frac{U^*}{\sigma^2} \end{aligned} \quad (\text{E.13})$$

More generally, this Theorem clarifies that identification of the CIR rests on knowledge of g , the distribution of price reset opportunities, and κ , the rate at which firms learn. Our more specific modeling assumptions on preferences and information acquisition technology further simplify

our informational needs to just the quantities U^* , and σ^2 . However, as this result makes clear, the qualitative economics are preserved under such general learning technology.

More generally, under a rather weak condition on the learning technology, we have that our baseline specification provides a lower bound on the effects of incomplete information on the CIR:

Corollary E.1. *If $\kappa(h, h') \leq \kappa(0, h + h')$, then:*

$$\mathcal{M}^b - \mathcal{M}^x \geq \frac{\int_0^\infty g(\tau) \tau (1 - \kappa(0, \tau)) d\tau}{\int_0^\infty g(\tau) \kappa(0, \tau) d\tau} \quad (\text{E.14})$$

If, moreover, we have that $\kappa(0, h') \leq \sigma^2 h' / (U^ + \sigma^2 h')$, then:*

$$\mathcal{M}^b - \mathcal{M}^x \geq \frac{U^*}{\sigma^2} \quad (\text{E.15})$$

Proof. By Theorem 5, we have that:

$$\begin{aligned} \mathcal{M}^b - \mathcal{M}^x &= \frac{\int_0^\infty g(\tau) \left[\int_0^\tau (1 - \kappa(h, \tau - h)) dh \right] d\tau}{\int_0^\infty g(\tau) \kappa(0, \tau) d\tau} \\ &\geq \frac{\int_0^\infty g(\tau) \left[\int_0^\tau (1 - \kappa(0, \tau)) dh \right] d\tau}{\int_0^\infty g(\tau) \kappa(0, \tau) d\tau} \\ &= \frac{\int_0^\infty g(\tau) \tau (1 - \kappa(0, \tau)) d\tau}{\int_0^\infty g(\tau) \kappa(0, \tau) d\tau} \end{aligned} \quad (\text{E.16})$$

where the first inequality follows by the hypothesis that $\kappa(h, h') \leq \kappa(0, h + h')$, proving the first statement. Moreover, we have that:

$$\frac{\int_0^\infty g(\tau) \tau (1 - \kappa(0, \tau)) d\tau}{\int_0^\infty g(\tau) \kappa(0, \tau) d\tau} \geq \frac{U^*}{\sigma^2} \frac{\int_0^\infty g(\tau) \frac{\tau}{U^* + \sigma^2 \tau} d\tau}{\int_0^\infty g(\tau) \frac{\tau}{U^* + \sigma^2 \tau} d\tau} = \frac{U^*}{\sigma^2} \quad (\text{E.17})$$

where the inequality follows by the hypothesis that $\kappa(0, h') \leq \sigma^2 h' / (U^* + \sigma^2 h')$, proving the second statement. ■

E.3. Examples

We conclude this analysis by considering some important special cases for which Theorem 5 allows for an explicit computation of the effects of incomplete information on the CIR. The first example derives the CIR under exogenous information and the second derives the CIR under an information structure where exogenous information arrives for free but the firm can freely acquire additional information by paying a mutual information cost. In both cases, the effect of incomplete information on the CIR remains $\mathcal{M}^b - \mathcal{M}^x = U^* / \sigma^2$, as per Theorem 2.

Example 3 (Exogenous Information). *Consider an exogenous information structure where firm i observes the process $ds_{i,t} = q_{i,t} dt + \sigma_s dW_{s,i,t}$ over time where $W_{s,i,t}$ is a Wiener process that is i.i.d. across firms. It follows that, starting from the stationary uncertainty distribution which is a mass point on the particular $U \equiv \sigma \sigma_s$, the firm's belief about $q_{i,t}$ evolves according to the Kalman-Bucy*

filter where

$$d\mathbb{E}_{i,t}[q_{i,t}] = \lambda(ds_{i,t} - \mathbb{E}_{i,t}[q_{i,t}]dt), \quad \lambda \equiv \frac{\sigma}{\sigma_s} \quad (\text{E.18})$$

Integrating this forward to $t + h'$, we observe that

$$\begin{aligned} \mathbb{E}_{i,t+h'}[q_{i,t+h'}] - \mathbb{E}_{i,t}[q_{i,t}] &= \int_0^{h'} \lambda e^{-\lambda(h'-\tau)} (ds_{i,t+\tau} - \mathbb{E}_{i,t}[q_{i,t}]d\tau) \\ &= (1 - e^{-\lambda h'}) (q_{i,t} - \mathbb{E}_{i,t}[q_{i,t}]) + \int_0^{h'} \lambda e^{-\lambda(h'-\tau)} ((q_{i,t+\tau} - q_{i,t})d\tau + \sigma_s dW_{i,s,\tau}) \\ &= (1 - e^{-\lambda h'}) (q_{i,t+h'} - \mathbb{E}_{i,t}[q_{i,t}]) + \varepsilon_i \end{aligned}$$

where ε_i is mean zero (but not necessarily independent of $q_{i,t+h'}$ as we never relied on independence in our argument above). Thus, this model is nested as a special case of the general learning process considered in this appendix with $\kappa(h, h') = 1 - e^{-\lambda h'}$. We observe that $\kappa(h, h') \leq \kappa(0, h + h')$.

We can compute in this case, by using Theorem 5, that:

$$\mathcal{M}^b - \mathcal{M}^x = \frac{\int_0^\infty g(\tau) \left[\int_0^\tau e^{-\lambda(\tau-h)} dh \right] d\tau}{\int_0^\infty g(\tau) (1 - e^{-\lambda\tau}) d\tau} = \frac{\int_0^\infty g(\tau) \frac{1 - e^{-\lambda\tau}}{\lambda} d\tau}{\int_0^\infty g(\tau) (1 - e^{-\lambda\tau}) d\tau} = \frac{1}{\lambda} = U/\sigma^2 \quad (\text{E.19})$$

Hence, under the exogenous information case, our formula collapses back to U/σ^2 . Of course, if one measured U in the data as $\bar{U}$, one would overstate monetary non-neutrality relative to our benchmark. If U were measured as U^* , then this would coincide exactly with our baseline formula. $\triangle$

Example 4 (Free Information with Optimal Information Acquisition). Suppose that the firm has access to an exogenous information structure where firm i observes the process $ds_{i,t} = q_{i,t}dt + \sigma_s dW_{s,i,t}$ over time where $W_{s,i,t}$ is a Wiener process that is i.i.d. across firms. Moreover, as in our baseline model, they can acquire additional information about $q_{i,t}$ subject to the mutual information cost. So long as $U^* \leq \sigma\sigma_s$, Theorem 1 applies in this setting and the optimal information policy is to acquire no additional information until the price reset opportunity arrives and to then acquire enough information to reset posterior uncertainty to U^* .

We now derive the as-if Kalman gain representation of this information process. In what follows, we suppress i subscripts for ease of reading. We start with the fact that between h and $h + h'$ (before $h + h'^+$ when the costly information acquisition happens), the mean belief m evolves as:

$$dm_t = \lambda_t (ds_t - m_t dt) \quad (\text{E.20})$$

where $\lambda_t = \frac{U_t}{\sigma_s^2}$. We can represent m_t by following a few steps. First, we observe that:

$$d\left(e^{\int_0^t \lambda_\tau d\tau} m_t\right) = e^{\int_0^t \lambda_\tau d\tau} (dm_t + \lambda_t m_t dt) = \lambda_t e^{\int_0^t \lambda_\tau d\tau} ds_t \quad (\text{E.21})$$

and so:

$$e^{\int_0^{h+h'} \lambda_\tau d\tau} m_{h+h'} - e^{\int_0^h \lambda_\tau d\tau} m_h = \int_h^{h+h'} \lambda_t e^{\int_0^t \lambda_\tau d\tau} ds_t \quad (\text{E.22})$$

which implies that:

$$m_{h+h'} - m_h = - \left(1 - e^{-\int_h^{h+h'} \lambda_\tau d\tau} \right) m_h + \int_h^{h+h'} \lambda_t e^{-\int_t^{h+h'} \lambda_\tau d\tau} ds_t \quad (\text{E.23})$$

Recall that $ds_t = q_t dt + \sigma_s dW_{s,t} = q_{h+h'} dt + \underbrace{(q_t - q_{h+h'}) dt + \sigma_s dW_{s,t}}_{\equiv d\epsilon_t}$ which gives

$$\begin{aligned} m_{h+h'} - m_h &= - \left(1 - e^{-\int_h^{h+h'} \lambda_\tau d\tau} \right) m_h + \int_h^{h+h'} \lambda_t e^{-\int_t^{h+h'} \lambda_\tau d\tau} (q_{h+h'} dt + d\epsilon_t) \\ &= - \left(1 - e^{-\int_h^{h+h'} \lambda_\tau d\tau} \right) m_h + q_{h+h'} \underbrace{\int_h^{h+h'} \lambda_t e^{-\int_t^{h+h'} \lambda_\tau d\tau} dt}_{= \int_h^{h+h'} d \left(e^{-\int_t^{h+h'} \lambda_\tau d\tau} \right) = 1 - e^{-\int_h^{h+h'} \lambda_\tau d\tau}} + \underbrace{\int_h^{h+h'} \lambda_t e^{-\int_t^{h+h'} \lambda_\tau d\tau} d\epsilon_t}_{\epsilon_{h+h'}} \\ &= \left(1 - e^{-\int_h^{h+h'} \lambda_\tau d\tau} \right) (q_{h+h'} - m_h) + \epsilon_{h+h'} \end{aligned} \quad (\text{E.24})$$

where $\epsilon_{h+h'}$ is mean zero (across firms) but not necessarily independent of other terms. As for the evolution of the posterior variance, we have

$$dU_t = \left(\sigma^2 - \frac{U_t^2}{\sigma_s^2} \right) dt \quad (\text{E.25})$$

which implies that

$$U(\tau) = \sigma \sigma_s \tanh \left(\frac{\sigma (\sigma_s^2 c_1 + \tau)}{\sigma_s} \right) \quad (\text{E.26})$$

with $U(0) = U^*$ so the constant c_1 is given by

$$U^* = \sigma \sigma_s \tanh(\sigma \sigma_s c_1) \implies c_1 = \frac{1}{\sigma \sigma_s} \tanh^{-1} \left(\frac{U^*}{\sigma \sigma_s} \right) \quad (\text{E.27})$$

And so we have that posterior variance evolves over the pricing spell according to:

$$U(\tau) = \sigma \sigma_s \tanh \left(\tanh^{-1} \left(\frac{U^*}{\sigma \sigma_s} \right) + \frac{\sigma}{\sigma_s} \tau \right) \quad (\text{E.28})$$

It remains to describe what happens at the instant of the price change, where the firm's uncertainty jumps from $U(h+h')$ to U^* . At this point, the posterior mean must obey:

$$m_{h+h'+} = \frac{U^*}{U(h+h')} m_{h+h'} + \left(1 - \frac{U^*}{U(h+h')} \right) (q_{h+h'} + \epsilon'_{h+h'}) \quad (\text{E.29})$$

Rearranging this and substituting our earlier value of $m_{h+h'} - m_h$, we obtain that:

$$\begin{aligned}
m_{h+h'} - m_h &= \frac{U^*}{U(h+h')} (m_{h+h'} - m_h) + \left(1 - \frac{U^*}{U(h+h')}\right) (q_{h+h'} - m_h + \varepsilon'_{h+h'}) \\
&= \frac{U^*}{U(h+h')} \left(\left(1 - e^{-\int_h^{h+h'} \lambda_\tau d\tau}\right) (q_{h+h'} - m_h) + \varepsilon_{h+h'} \right) + \left(1 - \frac{U^*}{U(h+h')}\right) (q_{h+h'} - m_h + \varepsilon'_{h+h'}) \\
&= \left(\frac{U^*}{U(h+h')} \left(1 - e^{-\int_h^{h+h'} \lambda_\tau d\tau}\right) + 1 - \frac{U^*}{U(h+h')} \right) (q_{h+h'} - m_h) + \tilde{\varepsilon}_{h+h'} \\
&= \left(1 - \frac{U^*}{U(h+h')}\right) e^{-\int_h^{h+h'} \lambda_\tau d\tau} (q_{h+h'} - m_h) + \tilde{\varepsilon}_{h+h'}
\end{aligned} \tag{E.30}$$

where $\tilde{\varepsilon}_{h+h'}$ is mean zero (across firms) but not necessarily independent of other terms. This implies that this information structure is a specific example of that considered in this appendix. Indeed, we have that the as-if Kalman gain is given by:

$$\kappa(h, h') = 1 - \frac{U^*}{U(h+h')} e^{-\int_h^{h+h'} \lambda_\tau d\tau} \tag{E.31}$$

We can further simplify this expression by replacing our solution for $U(\tau)$ and $\lambda_\tau = U(\tau)/\sigma_s^2$. Doing this, we obtain:

$$\begin{aligned}
\kappa(h, h') &= 1 - \frac{U^*}{\sigma\sigma_s \tanh\left(\tanh^{-1}\left(\frac{U^*}{\sigma\sigma_s}\right) + \frac{\sigma}{\sigma_s}(h+h')\right)} \exp\left(-\frac{\sigma}{\sigma_s} \int_h^{h+h'} \tanh\left(\tanh^{-1}\left(\frac{U^*}{\sigma\sigma_s}\right) + \frac{\sigma}{\sigma_s}\tau\right) d\tau\right) \\
&= 1 - \frac{U^*}{\sigma\sigma_s \tanh\left(\tanh^{-1}\left(\frac{U^*}{\sigma\sigma_s}\right) + \frac{\sigma}{\sigma_s}(h+h')\right)} \exp\left(-\frac{\sigma}{\sigma_s} \left[\log\left(\cosh\left(\tanh^{-1}\left(\frac{U^*}{\sigma\sigma_s}\right) + \frac{\sigma}{\sigma_s}\tau\right)\right) \right]_h^{h+h'}\right) \\
&= 1 - \frac{U^*}{\sigma\sigma_s \tanh\left(\tanh^{-1}\left(\frac{U^*}{\sigma\sigma_s}\right) + \frac{\sigma}{\sigma_s}(h+h')\right)} \exp\left(-\frac{\sigma}{\sigma_s} \frac{\log\left(\frac{\cosh\left(\tanh^{-1}\left(\frac{U^*}{\sigma\sigma_s}\right) + \frac{\sigma}{\sigma_s}(h+h')\right)}{\cosh\left(\tanh^{-1}\left(\frac{U^*}{\sigma\sigma_s}\right) + \frac{\sigma}{\sigma_s}h\right)}\right)}{\frac{\sigma}{\sigma_s}}\right) \\
&= 1 - \frac{U^*}{\sigma\sigma_s} \frac{\cosh\left(\tanh^{-1}\left(\frac{U^*}{\sigma\sigma_s}\right) + \frac{\sigma}{\sigma_s}h\right)}{\sinh\left(\tanh^{-1}\left(\frac{U^*}{\sigma\sigma_s}\right) + \frac{\sigma}{\sigma_s}(h+h')\right)}
\end{aligned} \tag{E.32}$$

To summarize, we have derived that this information structure lies within the general class considered in this section with:

$$\kappa(h, h') = 1 - \frac{U^*}{\sigma\sigma_s} \frac{\cosh\left(\tanh^{-1}\left(\frac{U^*}{\sigma\sigma_s}\right) + \frac{\sigma}{\sigma_s}h\right)}{\sinh\left(\tanh^{-1}\left(\frac{U^*}{\sigma\sigma_s}\right) + \frac{\sigma}{\sigma_s}(h+h')\right)} \tag{E.33}$$

To compute the CIR under this information structure, it remains to plug it into the formula from

Theorem 5:

$$\begin{aligned}
\mathcal{M}^b - \mathcal{M}^x &= \frac{\int_0^\infty g(\tau) \left[\int_0^\tau (1 - \kappa(h, \tau - h)) dh \right] d\tau}{\int_0^\infty g(\tau) \kappa(0, \tau) d\tau} \\
&= \frac{\int_0^\infty g(\tau) \int_0^\tau \frac{U^*}{\sigma \sigma_s} \frac{\cosh\left(\tanh^{-1}\left(\frac{U^*}{\sigma \sigma_s}\right) + \frac{\sigma}{\sigma_s} h\right)}{\sinh\left(\tanh^{-1}\left(\frac{U^*}{\sigma \sigma_s}\right) + \frac{\sigma}{\sigma_s} \tau\right)} dh d\tau}{\int_0^\infty g(\tau) \kappa(0, \tau) d\tau} \\
&= \frac{U^* \int_0^\infty g(\tau) \frac{\sinh\left(\tanh^{-1}\left(\frac{U^*}{\sigma \sigma_s}\right) + \frac{\sigma}{\sigma_s} \tau\right) - \sinh\left(\tanh^{-1}\left(\frac{U^*}{\sigma \sigma_s}\right)\right)}{\sinh\left(\tanh^{-1}\left(\frac{U^*}{\sigma \sigma_s}\right) + \frac{\sigma}{\sigma_s} \tau\right)} d\tau}{\sigma^2 \int_0^\infty g(\tau) \kappa(0, \tau) d\tau} \\
&= \frac{U^* \int_0^\infty g(\tau) \frac{\sinh\left(\tanh^{-1}\left(\frac{U^*}{\sigma \sigma_s}\right) + \frac{\sigma}{\sigma_s} \tau\right) - \sinh\left(\tanh^{-1}\left(\frac{U^*}{\sigma \sigma_s}\right)\right)}{\sinh\left(\tanh^{-1}\left(\frac{U^*}{\sigma \sigma_s}\right) + \frac{\sigma}{\sigma_s} \tau\right)} d\tau}{\sigma^2 \int_0^\infty g(\tau) \frac{\sinh\left(\tanh^{-1}\left(\frac{U^*}{\sigma \sigma_s}\right) + \frac{\sigma}{\sigma_s} \tau\right) - \sinh\left(\tanh^{-1}\left(\frac{U^*}{\sigma \sigma_s}\right)\right)}{\sinh\left(\tanh^{-1}\left(\frac{U^*}{\sigma \sigma_s}\right) + \frac{\sigma}{\sigma_s} \tau\right)} d\tau} \\
&= \frac{U^*}{\sigma^2}
\end{aligned} \tag{E.34}$$

Thus, we have derived that this information structure gives rise to an identical CIR as our baseline information structure. $\triangle$

F Additional Figures and Tables




Figure F.1: Expected Duration of Next Price Changes and Kalman Gains

Notes: The left panel shows the average conditional duration, $\bar{D}_h = \mathbb{E}_g^{h'}[h'|h]$, which is how long a firm that reset its price h periods ago expects to wait before resetting its price (blue solid line), as well as the average duration, $\bar{D} = \mathbb{E}_f^h[\bar{D}_h]$, which is how long the firms expect to wait *on average* before resetting their prices (blue dashed line). The right panel shows the average conditional Kalman gain, $\bar{\kappa}_h = \mathbb{E}_g^{h'}[\kappa_{h'+h}|h]$, which is the expected Kalman gain at the next price reset opportunity for a firm that last reset its price h periods ago (blue solid line), as well as the average Kalman gain, $\bar{\kappa} = \mathbb{E}_f^h[\bar{\kappa}_h]$, which is the average across all firms of the expected Kalman gain when they next reset their prices (blue dashed line)




Figure F.2: Distribution of Price Reset Opportunities and the Hazard Rate

Notes: The left panel shows the empirically estimated distribution of price reset opportunities G , given by $G(h) = 1 - \hat{f}(h)/\hat{f}(0)$ where $\hat{f}$ is the empirical distribution of time since firms' last price changes. g is the density function. The right panel shows the hazard rate, $\theta(h) = g(h)/(1 - G(h))$.




Figure F3: Uncertainty Distribution with Measurement Errors

Notes: This figure shows the distribution of firms' subjective uncertainty about their ideal prices under our baseline approach and the approach to account for measurement error that we describe in Section 6.3. The labeling follows Figure 3.




Figure F4: Uncertainty Distribution under Calvo

Notes: This figure shows the distribution of firms' subjective uncertainty about their ideal prices under our baseline approach and when we impose that the pricing hazard is constant as in Calvo (1983) that we describe in Section 6.3. The labeling follows Figure 3.

Table F.1: The Relationship Between Uncertainty and Time Since Changing Price

	(1)	(2)	(3)
<i>Dependent variable: Subjective uncertainty about firms' ideal price changes</i>			
Dummy for price changes in the last 3 months	0.0495 (0.0862)		
Dummy for price changes in the last 6 months		0.0306 (0.0850)	
Dummy for price changes in the last 12 months			-0.643*** (0.151)
Observations	467	467	467
R-squared	0.114	0.114	0.153
Industry Controls	Yes	Yes	Yes
Firm-level Controls	Yes	Yes	Yes
Manager Controls	Yes	Yes	Yes

Notes: This table reports results for the Huber robust regression. The dependent variable is the subjective uncertainty about firms' ideal price changes in the 2018Q1 survey, which is measured by the variance implied by each firm's reported probability distribution over different outcomes of their ideal price changes if firms are free to change their prices. Each Column uses different thresholds for the dummy for the last price changes. Industry fixed effects include dummies for 13 sub-industries. Firm-level controls include a log of firms' age, a log of firms' employment, foreign trade share, number of competitors, the slope of the profit function, firms' expected size of price changes in 3 months, and firms' subjective uncertainty about their ideal prices in next three months reported in the 2017Q4 survey. Manager controls include the age, education, and tenure at the firm of the respondent (each firm's manager). Sample weights are applied to all specifications. Robust standard errors (clustered at the 3-digit Australian and New Zealand Standard Industrial Classification level) are reported in parentheses. *** denotes statistical significance at the 1% level.




Figure E5: Overidentification Test for $\Delta^{\text{Select}} = \bar{D}$

Notes: This figure shows the baseline estimates of the information selection effect (Δ^{Select}) and the average pricing duration ($\bar{D}$). We also present the difference $\Delta^{\text{Select}} - \bar{D}$ to implement the overidentification test of the theory derived in Equation 44. We present 95% confidence intervals as black vertical lines.




Figure F6: CIR Decomposition with Measurement Errors and Calvo Pricing

Notes: This figure shows the output effects of a 1 percentage point shock to perceived gaps (left bar), to belief gaps (middle bar), and to belief gaps ignoring the selection effect (right bar). We compare our baseline estimates (blue bars) to the estimates that we obtain when we account for measurement error (gray bars) and impose that firms have a constant pricing hazard as in Calvo (1983) (red bars). With Calvo pricing, the output effect of a shock to belief gaps is 11.24pp (middle red bar) and the output effect of a shock to belief gaps ignoring the selection effect is 17.70pp (right red bar). After accounting for measurement error, the output effect of a shock to belief gaps is 9.65pp (middle gray bar) and the output effect of a shock to belief gaps ignoring the selection effect is 16.94pp (right gray bar). The estimation approaches for these two comparisons are described in Section 6.3. The labeling follows Figure 4.




Figure F.7: Microeconomic Volatility, Price Stickiness, Discount Rate, and Monetary Non-Neutrality

Notes: This figure shows counterfactual analyses on how micro uncertainty, price stickiness, and the discount rate affect monetary non-neutrality. The two left panels show the effect of price stickiness and the discount rate on firms' optimal reset uncertainty (U^*) as a function of the volatility of marginal cost. The two right panels show the effect of price stickiness and discount rate on monetary non-neutrality ($\mathcal{M}^b$) as a function of the volatility of marginal cost. Red stars show the estimates with $\varepsilon = 0$, $\hat{\sigma}^2 = 0.21$, and the baseline discount rate $r = 0.0034$, which implies an annual discount factor of approximately 0.96.




Figure E8: Microeconomic Volatility, Price Stickiness, and Monetary Non-Neutrality

Notes: This figure shows two counterfactual analyses on how micro uncertainty and price stickiness affect monetary non-neutrality. The left panel shows the effect of microeconomic uncertainty on monetary non-neutrality induced by information friction, $\partial \mathcal{M}^b / \partial \sigma^2$ as a function of the discount rate (r). The right panel shows the effect of price stickiness on monetary non-neutrality, $\partial \mathcal{M}^b / \partial \varepsilon|_{\varepsilon=0}$ as a function of the discount rate (r). Red stars show the estimates with the baseline discount rate $r = 0.0034$, which implies an annual discount factor of approximately 0.96. We present 95% confidence intervals as blue dashed lines.




Figure E9: Comparative Statics with Respect to Information Cost (ω)

Notes: This figure shows the sensitivity of key model statistics and CIRs to variations in the information cost parameter, ω . The left panel shows the CIR to a perceived gap shock, $\mathcal{M}^x$ (blue line), plotted alongside the average duration of price changes, $\bar{D}$ (orange line). The middle panel compares the difference between the CIR to a belief shock and the CIR to a perceived gap shock, $\mathcal{M}^b - \mathcal{M}^x$ (blue line) against the baseline formula U^*/σ^2 (orange line), where U^* is the average posterior uncertainty of price-adjusting firms. The right panel shows the ratio of the frequency of realized price changes to the frequency of price reviews ($N(|\Delta p| > 0)/N(Review)$). In all panels, the vertical green dashed line indicates the value of ω used in the baseline calibration.