Dynamic Pricing under Information Frictions: Evidence from Firm-level Subjective Expectations

Luca Gagliardone* Joris Tielens†

October, 2024

Abstract

This paper provides novel theory and evidence on the dynamic pricing behavior of firms, their process of belief formation, and the aggregate implications for monetary non-neutrality. We combine almost three decades of monthly microdata on prices, costs, and survey expectations for the Belgian manufacturing sector to quantify the roles of nominal, real, and information frictions in accounting for the incompleteness of the pass-through of cost shocks into prices. We find that the three rigidities are all quantitatively important, and in particular that firms exhibit a high average monthly discounting of about 0.8 attributable to the information friction. We further show that the discounting is state-dependent, just below one when firms are hit by large shocks and lower in normal times, consistent with a model in which firms update their beliefs *faster* in response to large disturbances. At the aggregate level, these findings imply a non-linear Phillips curve and a central role for heterogeneity across industries via the elasticity of inflation to aggregate shocks. Because the state-dependent information friction operates as an amplification mechanism of cost shocks, the model can explain a larger share of inflation volatility than the full-information benchmark, suggesting a quantitatively relevant channel by which incomplete information *increases* the neutrality of money.

*Gagliardone: Department of Economics, New York University. Email: luca.gagliardone@nyu.edu.
†Tielens: National Bank of Belgium and KU Leuven. Email: joris.tielens@nbb.be.

The views expressed in this paper are those of the authors and do not necessarily reflect the views of the National Bank of Belgium, the Eurosystem, or any other institution with which the authors are affiliated.

1 Introduction

This paper studies how firms form beliefs and set prices dynamically when facing persistent shocks to their production costs. This question is central to understanding the sources of aggregate fluctuations in real variables such as output or unemployment, the welfare costs associated with such gyrations, and how the monetary authority can effectively counteract the effects of disturbances. If the pass-through from costs to prices was complete and production inputs were priced flexibly, money would be neutral. In turn, traditional interventions by the central bank would have no effect on firms' production decisions. However, in the data, one typically observes real effects of monetary policy. This fact has prompted researchers to develop theoretical and empirical frameworks that accommodate forms of price rigidity and therefore an incomplete pass-through of cost shocks into prices.

The literature has identified three key forms of rigidities that possibly underlie the aggregate relation between inflation and measures of real activity, i.e. the Phillips curve.[1] Nominal frictions, which are a key ingredient of standard dynamic stochastic general equilibrium models used in policy analysis, pose a technological constraint on the ability of firms to adjust output prices costlessly in every period. Information frictions limit the ability of firms to acquire or process knowledge regarding the realizations of shocks and therefore factor them into pricing decisions. Real frictions generate strategic motives in pricing behavior and coordinated adjustments for all the competitors within a market.

Nominal, information, and real frictions are all ex-ante plausible candidates for explaining the patterns that we observe in aggregate data. What has often prevented the literature from making conclusive progress in disentangling between these sources of price rigidity is the lack of systematic panel data on costs and prices at the firm level. Moreover, the interaction between nominal and information frictions leads to a prominent role for firms' expectations that implies the necessity to complement the information on realizations with survey data on firms' beliefs (Coibion et al. 2018a).

In this paper, we develop a novel theoretical framework and provide empirical evidence on the relative importance of the three frictions in accounting for the incompleteness of the cost-price pass-through. We construct a unique dataset encompassing almost three decades of monthly observations for firms in the Belgian manufacturing sector with measures of prices, costs, and expectations, and show that the three rigidities all play quantitatively important roles. We further provide evidence of the following three empirical regularities. First, the presence of incomplete information leads to a high discounting at both the micro and macro levels, which reduces the sensitivity of inflation to shocks. Second, firms update their beliefs *faster* in response to large nominal disturbances, behaving as if

[1] Seminal contributions that discussed the importance of nominal, real, and information rigidities are Lucas (1972), Rotemberg (1982), Calvo (1983), Ball and Romer (1990), Woodford (2001), and Mankiw and Reis (2002).

they had approximately complete information. This channel generates a non-linear pass-through, which is higher when shocks are large. Third, high discounting and nonlinearities lead to a key role for heterogeneity across industries, which increases the transmission of cost shocks partially counteracting the aggregate effects of incomplete information.

Our theoretical framework is rich though parsimonious —with one parameter for each of the three rigidities— and nests the standard New Keynesian model with full information as a special case. The nominal rigidity is modeled à la Calvo, whereas the real rigidity allows for standard forms of imperfect competition which include monopolistic competition, Kimball preferences, and static or dynamic oligopolies. To depart from full information, we suppose that costs evolve stochastically over time as the product of a nominal common component and a real idiosyncratic component with different persistence. The information friction is that, though firms can observe their contemporaneous costs and the aggregate price level, they cannot distinguish between the two components of costs. Due to the unobservability of the split, firms solve a signal-extraction problem to forecast the persistence of the overall process for costs. When forming subjective beliefs, the weight that firms attach to their private forecast of the idiosyncratic component —i.e. the Kalman gain— captures the speed at which expectations are updated. In the model, this coefficient may vary endogenously over time whenever firms are surprised by aggregate inflation, and thus infer the presence of a new nominal cost shock.

The model directly implies a moment restriction on the joint intertemporal distribution of prices and costs, i.e. a dynamic pass-through regression. Due to endogeneity concerns and measurement limitations, identification requires a set of instruments. We show that identification of the degrees of nominal and real rigidities is obtained by using lags of observables and appropriate controls for demand. The novel identification argument pertains to the estimation of the Kalman gain. The conceptual difficulty that arises due to the dynamic nature of our setting is that the Kalman gain needs to be identified using variation in expectations that is uncorrelated with the noise in the private forecasts of the idiosyncratic cost components. As expectations depend directly on such noise (away from full information), they do not generally satisfy the exclusion restriction. However, we show that a coarse transformation of expectations, the expected *sign* of a price change, is uncorrelated with the noise in large samples. Therefore, a set of shifters that includes lags of costs and prices and the expected sign of the price change provides identification for the parameters of the model.

We begin by bringing the model to data for the special case of a constant Kalman gain and then inspect time variation and heterogeneity thereafter. The estimated average pass-through of contemporaneous shocks is large compared to standard estimates based on aggregate data (around 1.6% at a monthly frequency), but it is substantially lower for future shocks (around 1.2%). In particular, we find that firms exhibit a high average discounting of about 0.8 attributable to the information friction. Moreover, consistent with earlier findings (Gagliardone et al. 2023), the real rigidity decreases the pass-through by approximately 60% and the nominal friction implies a duration of the price spell of roughly a year. These estimates imply a substantial degree of monetary non-neutrality, consistent with the evidence from aggregate data.

We then move to study the cyclical properties of the Kalman gain and its heterogeneity across industries. We show empirically that the pass-through from costs to prices is state-dependent, high when the economy is hit by a large nominal shock and smaller in normal times. We provide evidence that these nonlinearities are partially attributable to the formation process of beliefs, which displays a faster updating in response to large nominal shocks that can be described via a higher Kalman gain. By extending the dynamic pass-through regression, we obtain quantitative estimates of regime-specific Kalman gains. We find that the estimated coefficient is just below one when industry inflation (in absolute value) is in the top 10% of the distribution, indicating that firms respond to large shocks as if they had almost complete information. On the other hand, in the remaining sample, the estimates are much lower, ranging from 0.69 to 0.28 with large standard errors. We conclude that a state-dependent Kalman gain, which increases in response to nominal cost shocks, is consistent with the evidence that we provide and can partially account for the observed nonlinearities of the pass-through.

The aggregate implications of our findings are discussed by deriving analytically the cost-based New Keynesian Phillips curve (NKPC), which is a description of the relation between inflation and costs at the aggregate level. In the full-information benchmark, aggregate inflation satisfies a standard linear NKPC with exponential discounting of expectations. However, the state-dependent information friction generally leads to nonlinearities in the NKPC and a higher discounting of expectations. In particular, we formally show that, under the information friction, the aggregate discounting fluctuates over time but is generally bounded away from one. Moreover, high discounting and nonlinearities lead to a response of aggregate inflation to nominal cost shocks that increases with the dispersion of inflation rates across industries. Therefore, industry heterogeneity plays a central role in determining aggregate inflation via its elasticity to nominal disturbances.

Finally, we ask how much inflation volatility can the model explain. We derive analytically a "reduced-form" NKPC, which relates expected inflation to aggregate real marginal cost. We can then use this characterization to obtain the model-based time series for aggregate inflation from the realized path of aggregate costs. We show that the nonlinear model lines up well against the data, though not all the realized volatility of inflation can be explained. In particular, the model can account for up to 63% of the volatility of inflation. Nevertheless, this departure from complete information significantly improves the fit. Indeed, when repeating the exercise with the linear NKPC, we find that the full-information benchmark accounts for only 50% of inflation volatility. We conclude from the exercise that the state-dependent information friction operates as an amplification mechanism of nominal cost shocks leading to larger fluctuations in prices, therefore implying *more* monetary neutrality compared to the full-information benchmark.

1.1 Literature Review

This paper is related to four strands of the literature. First and foremost, we contribute to the recent agenda that investigates empirically how beliefs are updated depending on the aggregate environment in which agents operate. Moreover, this paper provides new facts that complement existing survey-based evidence on firm expectations regarding inflation, new estimates of the slope of the NKPC, and novel theoretical insights regarding the implications of incomplete information for the neutrality of money. Compared to the literature, we are the only paper —to the best of our knowledge— that uses firm-level data on both realizations of prices and costs as well as beliefs, which allows us to inform on the mechanism by which expectations are updated and derive the novel theoretical implications.

Evidence of state-dependent information frictions. A novel strand of the literature on incomplete information shows that the process of expectation formation changes depending on aggregate conditions. Weber et al. (2023) use a randomized-control trial to identify the Kalman gain and show that it increases with aggregate inflation. We embed the idea in a dynamic pricing model and provide model-based identification for the parameter from the intertemporal co-movements in costs, prices, and expectations. Afrouzi et al. (2024) model theoretically the information acquisition of a firm that resets prices subject to nominal rigidities. Consistently with their theoretical and empirical findings, we provide evidence of a role for "selection," that increases the pass-through of shocks reducing monetary non-neutrality. Differently from the authors, we model a mechanism via a forward-looking updating of the prior uncertainty, as opposed to the variance of the noise component.[2] Pfäuti (2023) studies empirically the evolution of the Kalman gain using aggregate data on inflation and expectations, and provides estimates for a threshold for inflation above which the Kalman gain increases. Similarly, we also provide estimates for the

[2]This departure from the approach in the rational inattention literature (Sims 2003, Maćkowiak and Wiederholt 2009, Maćkowiak et al. 2023) is motivated by the empirical evidence that we provide and leads to a convenient solution to the problem of relating the (observable) dispersion of posterior beliefs with the (unobservable) Kalman gain. Indeed, whereas the variance of the noise is non-monotonically related to the dispersion in beliefs, the prior uncertainty is monotone in it and thus allows us to infer from data on belief dispersion the behavior of the Kalman gain. See section 5 for a formal treatment.

parameter using a threshold rule on the realized squared industry inflation rate. Finally, this paper provides theory and micro-evidence consistent with the "Attention-Inflation Hypothesis" proposed by Bracha and Tang (2024), according to which attention to inflation varies systematically with the price level. However, in our model, it is the common uncertainty regarding the evolution of the price distribution (prior variance) that increases following a nominal shock. Our mechanism is in line with evidence from the quasi-natural experiment studied in Drenik and Perez (2020) of a systematic positive relation between uncertainty and dispersion in prices.

Evidence on expectations of firms. This paper also contributes to the vast literature that provides survey-based evidence on firm expectations regarding inflation (Coibion et al. 2020, Andrade et al. 2022, Kumar et al. 2023, Candia et al. 2024). In particular, consistent with the findings of Coibion et al. (2018b), we document that firms form beliefs in a seemingly rational way by anticipating accurately whether they are going to increase prices in the next months, and use this fact to construct a relevant shifter for realized future prices. Nevertheless, we discuss evidence of a potential departure from Bayesian updating that is instead consistent with a form of present bias or hyperbolic discounting (Laibson 1994).

Evidence on the slope of the NKPC. Building on our earlier work, we provide novel estimates for the slope of the NKPC at monthly frequency. In line with previous results (Gagliardone et al. 2023), we find that the slope is statistically larger than zero even in normal times. This paper further clarifies the key importance of controlling for a time fixed effect, as discussed in McLeay and Tenreyro (2020), Hazell et al. (2022), and Fitzgerald et al. (2024). In particular, we derive within the model the time fixed effect and show that it indeed controls for a common trend in beliefs which would lead to bias, as the previous literature argued. Finally, our framework provides a constructive explanation for the flattening of the cost-price correlation in normal times by showing that it can be partially attributed to changes in the discount factor.

Aggregate implications of incomplete information. Finally, our paper contributes to the theoretical literature that investigates the role of incomplete information for inflation dynamics in models with information and nominal rigidities. Building from the seminal works of Lucas (1972), Woodford (2001), and Mankiw and Reis (2002), several papers have shown how incomplete information leads to an "anchoring" of expectations to the prior (i.e. a Kalman gain strictly smaller than one) and concluded that incomplete information increases monetary non-neutrality. Some influential contributions to this vast literature are Nimark (2008), Angeletos and La'O (2009), Maćkowiak and Wiederholt (2009), Melosi (2014), Paciello and Wiederholt (2014), Angeletos and Lian (2018), and Angeletos and La'O (2020). Compared to this literature, we argue that the anchoring to the prior has no impact on the *expected* inflation dynamics, though it leads to a slower decay of expectations over time that can rationalize the evidence that we provide on the nonlinear pass-through. Instead, the presence of the state-dependent information friction leads to a reduction in monetary non-neutrality and an increase in inflation volatility in response to nominal shocks compared to the full-information benchmark. The closest paper in this literature is Hellwig and Venkateswaran (2014), which shows that the interaction between the nominal rigidity and imperfect information regarding the persistence of shocks leads to a signal-extraction problem. Compared to that paper, we highlight the role of the state-dependent information friction, which leads to novel aggregate implications.

The paper develops as follows. Section 2 outlines the theoretical framework underlying the empirical analysis. Section 3 describes the data and measurement of the relevant variables for the empirical analysis. Section 4 discusses identification and the baseline estimation results. Section 5 provides evidence and discusses the mechanism underlying the state-dependent information friction. Section 6 derives the aggregate implications of our findings and provides validation for the model using aggregate data. Section 7 concludes.

2 Theoretical framework

This section develops the theoretical framework that underlies the empirical analysis that will follow. We first discuss the primitives of the model, namely demand, costs, and information. We then introduce real, nominal, and information rigidities. The intertemporal problem of firms and the formation of beliefs are then characterized. Finally, we derive the Dynamic Pass-through Regression, which we estimate in section 4 after discussing data and identification.

2.1 Primitives

Demand. The economy is populated by heterogeneous firms, denoted by f, each operating in an industry $i \in [0,1]$. We denote by $\mathcal{F}_i$ the set of competitors in industry i. Let p_{fit} be the log price charged by each firm for a unit of its output, $p_{it} := \int_{\mathcal{F}_i} p_{fit}\, df$ the log Cobb-Douglas industry price index, and Y_{it} the real industry output. For any industry i, we consider a demand system that generates a residual demand function of the following form:[3]

$$\mathcal{D}_{fit} := d(p_{fit},\, p_{it}) \cdot Y_{it} \quad \forall f \in \mathcal{F}_i. \tag{1}$$

[3] We allow for demand shocks in the empirical analysis. See section 4.3.

Denote with μ_{fit} the log Lerner index and with ϵ_{fit} the residual elasticity of demand:

$$\mu_{fit} := \ln\left(\frac{\epsilon_{fit}}{\epsilon_{fit}-1}\right), \qquad\qquad \epsilon_{fit} := -\frac{\partial \ln \mathcal{D}_{fit}}{\partial p_{fit}}.$$

We refer to μ_{fit} as the desired markup, which is the markup that a firm would enjoy absent frictions. Under (1), the desired markup depends only on relative prices, $\mu_{fit} \equiv \mu(p_{fit},\, p_{it})$. If the demand function can be log-linearized around a symmetric steady-state, then:

$$\mu_{fit} - \mu = -\Gamma(p_{fit} - p_{it}), \tag{2}$$

where $\{\mu_{fit} = \mu\}_{f\in\mathcal{F}_i,\ i\in[0,1]}$ at the steady state and $\Gamma > 0$ is the steady-state elasticity of the desired markup with respect to relative prices. As we show in Appendix A.1, this specification nests standard forms of imperfect competition such as monopolistic competition with Kimball preferences (Kimball 1995) or static and dynamic oligopolies (Atkeson and Burstein 2008, Wang and Werning 2022) and has been used in empirical work for special cases of this model (Amiti et al. 2019, Gagliardone et al. 2023). Finally, firms commit to producing enough output to meet demand for any realized price such that markets clear.

Costs. Firms are heterogeneous in their production technologies. We assume that a unit of output Y_{fit} is produced at a nominal marginal cost:

$$MC^n_{fit} = \widetilde{MC}_{fit} \cdot MC^n_{it}, \tag{3}$$

where $MC^n_{it} := \exp(\int_{\mathcal{F}_i} \ln(MC^n_{fit})df)$ denotes the average nominal marginal cost in an industry and $\widetilde{MC}_{fit}$ is the real idiosyncratic component. Whereas the industry marginal cost fluctuates because of nominal industry shocks, the idiosyncratic component moves whenever a firm-specific real shock occurs. For example, a nominal shock to aggregate costs can be a monetary policy shock, and an idiosyncratic real shock can be a productivity shock. To make the forecasting problem not trivial while maintaining tractability, we assume that the two processes are Markov with different persistence. We denote with lowercase letters the costs in logs, e.g. $mc^n_{fit} := \ln(MC^n_{fit})$ is the log nominal marginal cost. In particular, we suppose that the nominal component follows a random walk process and the real component a stationary AR(1) process with Gaussian shocks and time-invariant distribution:

$$\begin{aligned} mc^n_{it+1} &= mc^n_{it} + \varepsilon_{n,it+1} \\ \widetilde{mc}_{fit+1} &= \rho\widetilde{mc}_{fit} + \varepsilon_{r,fit+1}, \quad \varepsilon_{r,fit+1} \overset{i.i.d.}{\sim} \mathcal{N}(0,\sigma_r^2), \quad |\rho| < 1 \end{aligned} \tag{4}$$

We will show in section 3 that these functional forms are supported in the data and, specifically, that we can reject a null hypothesis of unit root for the real component but not for the nominal component. Whereas the assumption of normal real shocks leads to a price distribution that is log-normal within an industry, at this stage we make no assumptions on the distribution of the nominal shocks so that the distribution across industries can depart from normality.

Information. At the beginning of period t, every firm within an industry forms a prior belief regarding the evolution of costs. We assume that the prior belief is in common because based on a common information set $\mathcal{I}_{it}$, which evolves over time as:

$$\mathcal{I}_{it} = \{p_{it}, mc^n_{it-1}\} \cup \mathcal{I}_{it-1}\,.$$

All firms observe the industry price index and the realized nominal component of costs in the previous period. The prior belief that is formed based on this common information set evolves over time generating a common trend in subjective beliefs. We denote by $\mathbb{E}_{it} \equiv \mathbb{E}(\cdot|\mathcal{I}_{it})$ the conditional expectation.

As is standard in the incomplete information literature (e.g. Angeletos and La'O 2009, Angeletos and Lian 2018) we suppose that goods' prices do not fully reveal the distribution of costs. Such information rigidity is reasonable because, in reality, firms produce with complex production functions that include several inputs, whose prices may be imperfectly correlated with the retail prices.[4] Specifically for our setup, we suppose that, whereas firms observe their own costs mc^n_{fit} before making pricing decisions, they cannot disentangle between the idiosyncratic and common components, $\widetilde{mc}_{fit}$ and mc^n_{it}, even after observing p_{it}. As processes are Markov, if firms were able to decompose costs there would be "full information," by which we refer to the usual case in which firms can perfectly forecast future costs up to the realization of new cost shocks.

To model the departure from full information in a tractable way, we introduce a noisy private Gaussian signal s_{fit} which is informative regarding the idiosyncratic component of costs $\widetilde{mc}_{fit}$ and therefore the decomposition. Firms incorporate the signal optimally into their pricing decisions to the extent that it is useful (i.e. they do not know $\widetilde{mc}_{fit}$ already). In section 2.5, we discuss in detail the belief formation process. For the moment, consider the noisy signal as a simple modeling device that introduces noise in expectations, so that expected and realized prices can systematically and persistently differ. In the limiting case of full information, the noise becomes orthogonal to expectations implying that

[4] This friction can be viewed as a cognitive constraint on the ability of firms to forecast aggregate variables, which is arguably a more costly or complex task than making pricing decisions upon observing the aggregate price index, as it involves higher-order reasoning.

the differences between expected and realized prices cannot be systematic (Muth 1961).

Finally, we also allow firms to privately observe histories of their own costs and prices. The private information set evolves over time as:

$$\mathcal{I}_{fit} = \{mc^n_{fit}, p_{fit-1}, s_{fit}\} \cup \mathcal{I}_{it} \cup \mathcal{I}_{fit-1}\,.$$

We denote by $\mathbb{E}_{fit} \equiv \mathbb{E}(\cdot|\mathcal{I}_{fit})$ the associated conditional expectation. Given this information, what prevents firms from forming a perfect forecast of the path of future prices (conditionally on no new shocks occurring) is solely the inability to observe separately the marginal cost components. At the end of period t, after price decisions are taken and prices are realized and observed, firms recover $\widetilde{mc}_{fit}$ from the signal, so that they enter the new period $t+1$ with knowledge of mc^n_{it}.

2.2 Static Pricing

We start by describing how prices are set in a static setting without nominal rigidities. It follows directly from cost minimization that the static optimal log price that firms would set absent nominal rigidities is given by a markup over marginal cost:

$$p^\star_{fit} := \mu_{fit} + mc^n_{fit}.$$

Replacing the log-linearized markup (equation 2) and solving for the price, the static optimal price is given by:

$$\begin{aligned} p^\star_{fit} &= (1-\Omega)(\mu + mc^n_{fit}) + \Omega p_{it} \\ &= \mu_f + (1-\Omega) mc^r_{fit} + p_{it} \end{aligned} \tag{5}$$

where $\Omega := \frac{\Gamma}{1+\Gamma} \in (0,1)$ is the degree of strategic complementarities in price setting, $mc^r_{fit} := mc^n_{fit} - p_{it}$ is the real marginal cost, and $\mu_f := (1-\Omega)\mu$ is the steady-state markup scaled by the degree of strategic complementarities. As standard, approximating the firm's profit function to a second order around the static optimum $p_{fit} = p^\star_{fit}$, the firm's losses from mispricing are given by:

$$\Pi_{fit} \equiv \Pi(p_{fit}, p^\star_{fit}) = -\frac{\epsilon(\epsilon-1)}{2(1-\Omega)}\left(p_{fit} - p^\star_{fit}\right)^2, \tag{6}$$

where ϵ is the residual elasticity of demand at the symmetric steady state. The curvature of the profit function increases in the demand elasticity so that losses from mispricing are larger when facing a more elastic demand. Finally, notice that, absent nominal rigidities that make the problem dynamic, the information friction is immaterial as price setting does not depend on the idiosyncratic versus common components separately. In fact, from equation (5), knowledge of mc^n_{fit} and p_{it} is sufficient to compute $p^\star_{fit}$ regardless of whether cost shocks are real or nominal.

2.3 Firm's Problem

Firms are price setters subject to nominal rigidities à la Calvo. Denote by $p^o_{fit} := \ln(P^o_{fit})$ the log reset price, which equals the realization of the price p_{fit} if a price change occurs ($\Delta p_{fit} \neq 0$). Denote with θ the probability that a firm cannot adjust price between any two consecutive periods and let the discount factor be equal to one.[5] We consider the problem of a firm choosing a reset pricing policy at time $t = 0$ that maximizes expected intertemporal profits for every possible history. More precisely, a firm commits to a state-contingent reset pricing policy, which is a mapping from the information set into the price that the firm would like to set, i.e. $p^o_{fit} : \mathcal{I}_{fit} \mapsto \mathbb{R}$. If no reset opportunity occurs, the firm keeps the price constant. The time-0 problem can be written as:

$$\max_{\{p^o_{fit}\}_{t\geq 0}} \mathbb{E}\left[\sum_{t=0}^{\infty} \theta^t \, \Pi(p^o_{fit}, p^\star_{fit}) \;\Big|\; \mathcal{I}_{fi0}\right]. \tag{7}$$

Differentiation with respect to p^o_{fit} and application of the law of iterated expectations for nested sets $\mathcal{I}_{fit} \subseteq \mathcal{I}_{fit+1}$ then permit obtaining the familiar recursive first-order condition for the reset price:

$$\begin{aligned} p^o_{fit} &= (1-\theta)p^\star_{fit} + \theta \mathbb{E}_{fit}\left\{p^o_{fit+1}\right\} \\ &= (1-\theta)\left(\mu_f + (1-\Omega)mc^r_{fit} + p_{it}\right) + \theta \mathbb{E}_{fit}\left\{p^o_{fit+1}\right\}. \end{aligned} \tag{8}$$

The reset price is the expected discounted present value of the static target prices, which are equal to a linear combination of marginal cost and the industry price index with weight equal to the degree of strategic complementarities (equation 5). Because firms have private information, the forecast $\mathbb{E}_{fit}(p^o_{fit+1}) \neq \mathbb{E}_{it}(p^o_{fit+1})$ is possibly different across firms.

2.4 Full-Information Benchmark

We now illustrate how firms would compute the present value if they were able to observe the real and nominal components of costs separately. Recursive substitution of the first-

[5] In the data, we are working at a monthly frequency at which a standard calibration of the discount factor is 0.998. The difference does not quantitatively impact our estimates. We consider the limit for the discount factor approaching one from below for every limiting case of vanishing information friction.

order condition (8) leads to the following formula for the present value:

$$p^o_{fit} = \mu_f + (1-\theta)\sum_{\tau=0}^{\infty}\theta^\tau \mathbb{E}_{fit}\Big\{(1-\Omega)(\widetilde{mc}_{fit+\tau} + mc^n_{it+\tau}) + \Omega p_{it+\tau}\Big\}.$$

Making use of the cost dynamics (4) and the fact that firms observe real shocks $\widetilde{mc}_{fit} \in \mathcal{I}_{fit}$ in the full-information benchmark, the above reduces to:

$$p^o_{fit} = \mu_f + \frac{(1-\Omega)(1-\theta)}{1-\theta\rho}\widetilde{mc}_{fit} + (1-\Omega)mc^n_{it} + \Omega\sum_{\tau=0}^{\infty}(1-\theta)\theta^\tau\, \mathbb{E}_{it}(p_{it+\tau}).$$

The term involving $\widetilde{mc}_{fit}$ is the expected present discounted value of future real shocks conditionally on the value $\widetilde{mc}_{fit}$, which is increasing in the persistence of the shock ρ; similarly, the subsequent term is the present value of the nominal shock, which simplifies due to the random walk assumption. $\widetilde{mc}_{fit}$ and mc^n_{it} enter the reset price separately to the extent that $\rho \neq 1$, which shows that the information friction is binding whenever the real and nominal components have different persistence.

Finally, under full information, the response of the reset price relative to the industry-average reset price (p^o_{it}) is solely a function of the real cost shocks:

$$\frac{\partial(p^o_{fit} - p^o_{it})}{\partial mc^n_{fit}} = \frac{\partial(p^o_{fit} - p^o_{it})}{\partial \widetilde{mc}_{fit}} = \frac{(1-\Omega)(1-\theta)}{1-\theta\rho}$$

and, in particular, the response of the relative reset price to a real shock is constant for constant parameters Ω, θ, and ρ. We will show in section 5.1 that this absence of state dependence in the pass-through from real costs to relative prices is at odds with the data and motivates the following departure from full information.

2.5 Subjective Expectations

We now describe the formation of subjective expectations under the information friction. Firms form a common prior belief conditional on the common information set $\mathcal{I}_{it}$, then receive the private signal s_{fit} and form a posterior combining the prior and signal.

Denote with $\Delta p_{fit+1|t} := (1-\theta)(p^o_{fit+1} - p_{fit}) = (1-\theta)\Delta p^o_{fit+1}$ the growth rate of prices conditional on resetting at time t and no new shocks realized between t and $t+1$, taking into account the probability $(1-\theta)$ of drawing an adjustment opportunity between t and $t+1$.[6] It is immediate to show that $\Delta p_{fit+1|t}$ summarizes the uncertainty of firms given their information at time t. Indeed, subtracting p_{fit-1} from both sides of the first-order

[6]Whereas unconditional price changes are a mixture of the Bernoulli random variable and the shocks to costs, conditional price changes are a linear function of the shocks and thus inherit the Gaussian distribution.

condition (8) and rearranging we obtain:

$$(1-\theta)(p^{o}_{fit}-p_{fit-1}) = \frac{(1-\theta)^2}{1-\theta(1-\theta)}\left(p^{\star}_{fit}-p_{fit-1}\right) + \frac{\theta}{1-\theta(1-\theta)}\mathbb{E}_{fit}\left\{\Delta p_{fit+1|t}\right\}. \tag{9}$$

Differently from the full-information benchmark, the forecast of $\Delta p_{fit+1|t}$ cannot be obtained by directly using the processes for the two components, as the split is unobserved and they have different persistence. Instead, firms compute the forecast as follows. First, they forecast the industry distribution of price changes.[7] We refer to such forecast as a common "prior belief."[8] Because idiosyncratic shocks are Gaussian with finite present values, the belief can be represented with a Gaussian distribution. We denote by $\pi^e_{it} := \mathbb{E}_{it}(\Delta p_{fit+1|t})$ the forecast of the industry inflation rate and by $\sigma^2_{e,it} := \mathbb{V}_{it}(\Delta p_{fit+1|t})$ the forecast of the industry dispersion in conditional price changes:

$$\Delta p_{fit+1|t} \,\Big|\, \mathcal{I}_{it} \sim \mathcal{N}\left(\pi^e_{it},\ \sigma^2_{e,it}\right). \tag{10}$$

Firms realize that there will be a non-degenerate distribution of growth rates around a common trend due to the presence of idiosyncratic shocks and nominal rigidity. The prior variance $\sigma^2_{e,it}$ pins down the expected distribution under the normality assumption. The expected inflation rate captures persistent co-movements of prices within an industry that are due to the nominal component.

Firms generally expect their own path of prices to deviate persistently from the industry inflation rate. In fact, when a firm experiences a large positive realization of the idiosyncratic cost shock, the firm's price increases faster than the industry inflation rate. Eventually, as time passes, the growth rate re-aligns with industry inflation as the idiosyncratic shock is absorbed into the firm's price level. However, even shortly-lived differences in growth rates lead to possibly large fluctuations in prices.

To parsimoniously describe these fluctuations, we make use of the noisy signal. In the spirit of Lucas (1972), we think about a signal as a noisy forecast of the price growth rate of the firm.[9] Absent the noise, firms would be able to perfectly anticipate the growth rate in conditional prices if no new shocks realize, as in the full-information benchmark. Away from full information, firms make use of the forecast of the industry distribution

[7] Beaudry et al. (2024) provides evidence that inflation expectations depend on a common component that can be extracted from disaggregated data. Similarly to this paper, the authors rationalize the evidence with a signal-extraction problem between common and idiosyncratic supply shocks.

[8] We notice that, as the real component washes out when aggregating within an industry, price adjustments are i.i.d. over time, and firms are small in the industry, information regarding the specific history of a firm or the realized signal about the real component are not informative to forecast the industry distribution.

[9] Differently from Lucas (1972), the uncertainty is here on the growth rate rather than the level of the price. Though practically relevant when bringing the model to data, in principle the two models are capturing similar forces. As an extreme example, consider a one-time unanticipated shock starting from an initial distribution with $p_{fit-1} = 0$ for all firms in all industries.

to "filter out" the noise in the signal. Formally, the signal is defined as the conditional expectation given the signal:[10]

$$s_{fit} := \mathbb{E}(\Delta p_{fit+1|t} \mid s_{fit}).$$

We further assume that the signal has a time-invariant Gaussian distribution:

$$s_{fit} = \Delta p_{fit+1|t} + \eta_{fit}, \quad \eta_{fit} \,\Big|\, s_{fit} \overset{i.i.d.}{\sim} \mathcal{N}(0, \sigma_\eta^2). \tag{11}$$

The noise term η_{fit} is defined as the difference between the realization and the conditional expectation. Several standard assumptions are implicit in equation (11). First, the noise is i.i.d. with mean zero; this implies that, when averaging across firms within an industry, the noise "washes out." This implicitly assumes that a law of large numbers applies within the industry and that there is no further uncertainty on top of the one induced by shocks and the information friction. Second, σ_η^2 does not depend on time and is homogeneous across all firms; though one could relax this assumption along the lines of Afrouzi et al. (2024), we focus here on the time variation in beliefs that is induced by the evolution of $\sigma_{e,it}^2$.[11]

Using Bayes rule, the subjective expectation is given by a linear combination between the forecast of the industry inflation rate and the private signal:

$$\mathbb{E}_{fit}(\Delta p_{fit+1|t}) = (1 - \beta_{it})\pi_{it}^e + \beta_{it} s_{fit}, \tag{12}$$

where β_{it} is the Kalman gain given by:

$$\beta_{it} = \frac{\sigma_{e,it}^2}{\sigma_{e,it}^2 + \sigma_\eta^2}.$$

Equation (12) illustrates the solution of the signal-extraction problem. As firms cannot disentangle between common nominal and idiosyncratic real shocks, which have different persistence, the forecast is a linear combination between the forecast of the industry inflation and the forecast of the firm's price growth rate upon observing the signal. The Kalman gain β_{it} captures the effects of the persistence ρ, the dispersion of idiosyncratic shocks σ_r^2, and variance of noise σ_η^2. Moreover, β_{it} fluctuates over time as a function of realized inflation $\pi_{it} \in \mathcal{I}_{it}$, which is observed and thus affects the forecast of the distribution of price

[10] Though the signal appears on both sides of this definition, this normalization can be implemented directly announcing to firms the conditional expectation in place of a signal. This follows from the fact that observing the signal is equivalent to observing the conditional expectation. See e.g. the technical appendix in Vives (2010).

[11] This is motivated by (i) the difficulty of estimating precisely a time-varying volatility of firm-specific forecast errors with our data, and (ii) the empirical observation that the dispersion in price changes varies over time in response to shocks and these movements are largely anticipated by beliefs. See section 5 for details.

changes via $\sigma^2_{e,it}$. We discuss in section 5 the role of the latter channel for the endogenous propagation of shocks. Finally, in the limit case of full information, the prior belief becomes an improper uniform distribution ($\sigma^2_{e,it} \to \infty$), and the forecast error becomes orthogonal to the expectation. In turn, Kalman gain converges to a constant ($\beta_{it} \to 1$) for all industries in all periods.[12]

Identification of the Kalman gain. Before deriving the empirical specification that will be estimated in the data, we provide an intermediate identification result. To understand its relevance, suppose for the moment that the econometrician has data directly on expectations $\mathbb{E}_{fit}(\Delta p_{fit+1|t})$ and the signal s_{fit} in equation (12). Even in this case, identifying the coefficient β_{it} is not immediate, as one generally does not observe π^e_{it}. The next lemma shows that identification of β_{it} is achieved by a regression of posterior beliefs on signals and an industry-by-time fixed effect. This insight will be then used to derive the empirical specification of the model in the next section.

Lemma 1. *Subjective expectations of (unconditional) price changes in deviation from the industry average are proportional to relative inflation rates:*

$$\mathbb{E}_{fit}(\Delta p_{fit+1}) - \int_{\mathcal{F}_i} \mathbb{E}_{jit}(\Delta p_{jit+1})dj = \beta_{it} \cdot \Big(\Delta p_{fit+1|t} - \pi_{it+1} + \eta_{fit}\Big), \tag{13}$$

where β_{it} is the Kalman gain.

Lemma (1) rearranges equation (12) using the fact that the prior is common to replace it with the average belief. Making use of the expected law of motion, it follows from the lemma that a regression of the posterior belief of the price change ($\mathbb{E}_{fit}(\Delta p_{fit+1})$) on an industry-by-time fixed effect (FE_{it}) and the ex-post realization of price changes (Δp_{fit+1}) identifies the same coefficient as that in equation (12), i.e. the Kalman gain:

$$\mathbb{E}_{fit}(\Delta p_{fit+1}) = FE_{it} + \beta_{it}\Big(\Delta p_{fit+1} + \eta_{fit}\Big) + u_{fit+1},$$

where the industry-by-time fixed effect is given by:

$$FE_{it} = \int_{\mathcal{F}_i} \mathbb{E}_{fit}(\Delta p_{fit+1})df - \beta_{it}\pi_{it+1},$$

and $u_{fit+1} := \beta_{it}(\Delta p_{fit+1|t} - \Delta p_{fit+1})$ denotes a mean-zero regression residual due to the nominal rigidity. This error term has a time-varying volatility, function of β_{it}. We make use of the insight of lemma 1 to control for the (unobserved but common) prior belief in the derivations of the Dynamic Pass-through Regression.

[12] We notice that the same limit is equivalently obtained by letting $\sigma^2_\eta \to 0$.

Interpretation of the Kalman gain. The presence of incomplete information leads to a form of discounting (Gabaix and Laibson 2017). In particular, we notice that our recursive representation of the problem (equations 8 and 13) is consistent with the $\delta - \beta$ preferences as in Harris and Laibson (2001), with $\delta \approx 1$ (omitted) parametrizing the exponential decay and $\beta \leq 1$ the hyperbolic discounting. In our model, hyperbolic discounting comes from the imperfect ability of firms to forecast the present value of future markups over costs. In turn, there is no "extra discounting" when forecasting a variable at time $t+1$ compared to a variable at time $t+2$, as both enter the same calculation of the present value ($\Delta p_{fit+1|t}$) from the perspective of a firm that is setting prices at time t. We discuss evidence for this interpretation in section 5.1 and argue that the Kalman gain is identified in the data from the discounting of cost shocks at frequencies shorter than a month.

2.6 Dynamic Pass-through Regression

In section 2.3, we outlined how firms set prices for any given expected present value of markups over marginal costs. Using the recursive representation of the first-order condition, we characterized how the expectation of the future reset price is mapped into the choice of the current reset price. In this sense, equation (8) can be viewed as defining an operator that takes expectations as an input and returns the firm's pricing choice as an outcome. In section 2.5, we described the complementary step that leads to the formation of expectations: for any given collection of data points that results from the choices taken by firms and the realizations of shocks, equation (13) provides the mapping from data into expectations. Accordingly, the equation can be used to define an operator that takes realized choices as inputs and returns expectations as an outcome. The Dynamic Pass-through Regression combines these two operations with the expected law of motion to obtain a mapping from data (future prices) into data (current prices). In turn, the mapping imposes a moment restriction of the joint distribution of current and future prices and costs, which in practice boils down to a linear regression model with coefficient restrictions. Such a formulation and appropriate shifters then permit bringing the model to the data to estimate structural parameters and assess the relative roles of the frictions in shaping the pass-through from costs to prices.

Let us now derive the regression. Denote with a tilde the variables in deviation from their industry average $\widetilde{x}_{fit} := x_{fit} - \int_{\mathcal{F}_i} x_{jit} dj$.

Proposition 1 (Dynamic Pass-through Regression)**.** *Prices and costs satisfy the following regression equation:*

$$\Delta \widetilde{p}_{fit} = \frac{(1-\theta)^2}{\theta} \left[\mu_f + (1-\Omega)\widetilde{mc}_{fit} - \widetilde{p}_{fit} \right] + \beta_{it} \cdot (\Delta \widetilde{p}_{fit+1} + \eta_{fit}) + \widetilde{u}_{fit} \tag{14}$$

where η_{fit} is the noise in signals and $\widetilde{u}_{fit} \perp \mathcal{I}_{fit}$ is a sampling error with cross-sectional mean zero that depends on the realization of the Calvo fairy.

The proof of the proposition is in section A.3 of the Appendix. Let us now describe the regression model (14). All the variables are demeaned with respect to the industry average; as argued in the previous section, this demeaning allows us to remove the variation due to a common trend in beliefs.[13] Demeaned price changes Δp_{fit} are then regressed on: (i) a firm-specific intercept μ_f; (ii) marginal cost mc_{fit} (either real or nominal, demeaned); (iii) the relative price level $p_{it} - p_{fit}$; (iv) future price changes Δp_{fit+1}, demeaned. The coefficient θ is identified from variation in marginal cost and relative prices, with asynchronous movements of the two identifying also Ω. The Kalman gain is identified from expected variation in future price changes. The two error terms, whose orthogonality with respect to an instrument set will be addressed formally in section 4.1, are present because (i) information is incomplete so that expectations depend on the noise in the signal; (ii) when firms are choosing the reset price, the uncertainty regarding whether they can or cannot adjust prices is not yet resolved. Finally, we notice that the presence of the relative price level ($p_{it} - p_{fit}$) as a control follows directly from theory but it is often overlooked in applied work. The reason why this might be the case is that, at the aggregate level, it cancels out from the regression as it averages to zero in the cross-section. However, when estimating pricing equations with micro-data, not only is it necessary to avoid the omitted variable problem, but it is also essential to correctly identify the stickiness parameter θ and the degree of real rigidities Ω (Gagliardone et al. 2023).

3 Data and Measurement

This section assembles a micro-level dataset that covers the manufacturing sector in Belgium between 1995 and 2023 at a monthly frequency. We begin by outlining the various data sources and subsequently discuss the mapping of the theoretical counterparts from section 2 to the data. Finally, we discuss various features of the data that justify aspects of the theoretical framework and are salient for identification of the Dynamic Pass-through Regression.

3.1 Data Sources

Our analysis draws on five confidential data sources and combines detailed firm-level information on monthly production, costs, pricing decisions and subjective expectations on

[13]See e.g. Hazell et al. (2022) for a discussion of the importance of controlling for the trend in beliefs.

future firm-level variables. Our dataset extends and enriches the annual dataset used by Amiti et al. (2019) and the quarterly dataset used by Gagliardone et al. (2023, 2024).

First, the PRODCOM dataset tracks domestic firms' monthly sales and physical quantities sold for each narrowly defined (8-digit PC codes) manufacturing product (e.g., 10.32.14.00 is "Pineapple juice", 10.32.15.00 is "Grape juice" and 10.32.16.00 is "Apple juice"). We use this highly disaggregated information to calculate domestic unit values (sales over quantities) at the firm-product level.[14] To account for the presence of foreign competitors active in Belgian product markets, we obtain similar data from the administrative records of Belgian customs declarations. Specifically, for each manufacturing product sold by a foreign producer to a Belgian buyer, we observe monthly sales and quantity sold for different products (8-digit CN codes, which are mapped to 8-digit PC codes), from which we compute unit values of foreign competitors in the local market (Belgium).

Second, we use detailed administrative data to measure firms' variable production costs. Specifically, we obtain information on firms' monthly purchases of intermediates (materials and services) from their VAT declarations submitted to the tax authorities. Additionally, we draw upon firms' social security declarations to obtain a measure of their labor costs (the wage bill).

Finally, we tap into the monthly NBB Business Survey (NBB-BS). This survey covers a representative panel of the Belgian manufacturing sector and is designed to provide both a timely and bottom-up measure of various business cycle indicators (De Greef and Van Nieuwenhuyze 2009). The NBB-BS includes one module that is backward-looking and another that is forward-looking. The former tracks past realizations of price changes, inventory, production, and demand. The latter captures firms' subjective expectations (for the next three months) such as future demand for its products, production, and price changes. The survey is mainly qualitative in nature — i.e. firms typically have three reply options: "decrease", "no change", "increase". The qualitative nature reduces the reporting burden on the firm (appropriate, as the survey is not mandatory) and reduces measurement error.

3.2 Measurement

We now describe how we construct the variables that enter our empirical strategy. Appendix C provides additional details on the data cleaning procedures.

[14]PRODCOM surveys all Belgian firms involved in manufacturing production with more than 10 employees, covering over 90% of production in each NACE 4-digit industry. The survey does not require firms to distinguish between production and sales to domestic and international customers. Therefore, we recover domestic values and quantities sold by combining information from PRODCOM with international trade data on firms' product-level exports (quantities and sales).

Output prices. The variable of interest is the domestic price charged by firm f in industry i, $P_{fit} := \exp(p_{fit})$. Constructing this variable requires us to aggregate the product dimension in PRODCOM to the firm-industry dimension. To that end, we exploit the fact that the first four digits of this product classification follow the NACE taxonomy, i.e. the official industry nomenclature in the European Union. In keeping with this taxonomy, we define an industry i as the first four digits of the product codes. This classification optimally balances a coherent definition of the industry (which is mostly precise if narrow) with the ability to identify an appropriate set of competitors (both domestic and foreign). This leads to 158 manufacturing industries, distributed across 9 (2-digit) manufacturing sectors (reported in Table A.1).

Due to repeated product code revisions, a consistent 8-digit product code taxonomy does not exist across the entire sample period. Therefore, we compute the sequence of price changes across consecutive time periods (t and $t+1$) by mapping the product codes at $t+1$ to their corresponding codes at t, aggregating them at the firm-industry level, and recovering the time series of the firm-industry price index (in levels) by concatenating quarterly price changes. More precisely, we compute the change in the firm-industry price index, P_{fit}/P_{fit-1}, using the most disaggregated level in the data. For domestic producers, the finest level of aggregation is the firm×8-digit PC product code level. For foreign competitors, it is the importing-firm×source country× 8-digit PC product code level.[15] Approximately half of the domestic firms in our sample are multi-product firms, meaning they produce multiple 8-digit products within the same industry. For these entities, we compute the price change by aggregating changes in product-level prices using a Törnqvist index:

$$P_{fit}/P_{fit-1} = \prod_{p \in \mathcal{P}_{fit}} (P_{pfit}/P_{pfit-1})^{\bar{s}_{pfit}}.$$

In the formula above, $\mathcal{P}_{fit}$ represents the set of 8-digit products manufactured by firm f in industry i, P_{pfit} is the unit value of product p in $\mathcal{P}_{fit}$, and $\bar{s}_{pfit}$ is a Törnqvist weight computed as the average of the sale shares between t and $t-1$: $\bar{s}_{pfit} := \frac{s_{pfit}+s_{pfit-1}}{2}$. Finally, we construct the time series of price levels P_{fit} by concatenating monthly changes.[16]

Using a similar approach, we construct the price index of competitors for each do-

[15] In the raw customs data, products are measured using the more disaggregated CN 8-digit product classification. We map the CN product codes in the customs data to PC product codes used in PRODCOM using the official bridge tables available on the Eurostat web page. See Appendix C.1 for additional details.

[16] Let t^0_{fi} denote the first month when f appears in our data. Starting from a base period P_{fi0}, which we can normalize to one, prices are concatenated using the formula: $P_{fit} = P_{fi0} \prod_{\tau=t^0_{fi}+1}^{t} \left(P_{fi\tau}/P_{fi\tau-1}\right)$. The normalization of the level of the firm's price index in the base period, P_{fi0}, is one rationale for the inclusion of firm fixed effects in our empirical specifications downstream.

mestic firm by concatenating monthly changes as follows:

$$P_{it}^{-f}/P_{it-1}^{-f} = \prod_{k\in\mathcal{F}_i/f} (P_{kt}/P_{kt-1})^{\bar{s}_{kt}^{-f}}. \tag{15}$$

Here, $\bar{s}_{kt}^{-f} := \frac{1}{2}\left(\frac{s_{kt}}{1-s_{fit}} + \frac{s_{kt-1}}{1-s_{fit-1}}\right)$ represents a Törnqvist weight, constructed by averaging the residual revenue share of competitors in the industry at time t (net of firm f revenues) with that at time $t-1$.[17] Note that the set of domestic competitors for each Belgian producer, denoted as $\mathcal{F}_i$, includes not only other Belgian manufacturers operating in the same industry but also foreign manufacturers selling the same goods to Belgian customers.

Marginal costs. To derive an empirical counterpart to marginal cost in equation (3), we follow Gagliardone et al. (2023, 2024) and assume a cost structure in which the nominal marginal cost of a firm is proportional to its average variable costs: $MC_{fit}^{n} = (1+\nu_f)AVC_{fit}$. The coefficient ν_f captures the curvature of the short-run cost function, and it is inversely related to the firm's short-run returns to scale in production ($\nu_f \equiv 1/RS_f - 1$). Using the definition of average variable costs (total variable costs over output, TVC_{fit}^{n}/Y_{fit}) and applying a logarithmic transformation, we have that firm-level log-nominal marginal cost is given by:

$$mc_{fit}^{n} = (tvc_{fit}^{n} - y_{fit}) + \ln(1+\nu_f)$$

In the data, we measure total variable costs as the sum of intermediate costs (materials and services purchased) and labor costs (wage bill). While the former is available at the monthly level, the latter is observed at a quarterly frequency and split equally across the three months of the quarter.[18]

We compute a quantity index by dividing a firm's domestic revenue by its domestic price index.[19] Firm-specific short-run returns to scale are not directly observable in the data. Therefore, to the extent that individual firms' production technologies deviate from constant returns to scale ($\nu_f \neq 0$), our measure of log-marginal costs would be missing an additive constant. Below, we will rely on firm fixed effects to neutralize the impact of this unobserved factor.

[17] As with the firm's price index, the level of the price index of competitors is constructed by normalizing the first period to one and concatenating quarterly changes. Also, in this case, the normalization is immaterial for estimation purposes as our empirical model always includes firm fixed effects.

[18] For multi-industry firms, we allocate the observed total firm-level variable cost to individual firm-industries according using the sales-share of that industry.

[19] Specifically, we compute $Y_{fit} = (PY)_{fit}/\bar{P}_{fit}$, where $\bar{P}_{fit}$ denotes the firm-month domestic price index. For single-industry firms, $\bar{P}_{fit}$ coincides with the firm-industry price index P_{fit}. For multi-industry firms, we construct $\bar{P}_{fit}$ as an average of the different firm-industry price indexes using as weights the firm-specific revenue shares of each industry.

Table 1: Summary statistics

Panel (a): Firm characteristics

	Mean	p. 5	p. 25	p. 50	p. 75	p. 95
Number of employees	102	9	18	32	66	362
Turnover	3037	46	175	439	1245	8368
Number of products	3.26	1	1	2	3	11
Number of industries	1.05	1	1	1	1	1
Revenue share of main industry	97.63	100	100	100	100	100
Number of consecutive months in the sample	99.55	24	45	75	132	261

Panel (b): Product market structure

	Mean	p. 5	p. 25	p. 50	p. 75	p. 95
Market share of the firm within the industry	6.83	0.14	0.55	1.55	5.33	35.19
Market share of the firm within the broad sector	0.58	0.01	0.03	0.09	0.26	1.77
Market share of the firm within manufacturing	0.06	0	0	0.01	0.02	0.16

Panel (c): Distributional moments

	Mean	p. 5	p. 25	p. 50	p. 75	p. 95
Δp_{fit}	0.00	−0.22	−0.02	0.00	0.03	0.23
$\mathbb{V}_t(\Delta p_{fit})$	0.02	0.01	0.02	0.02	0.02	0.03
$\text{Skew}_t(\Delta p_{fit})$	0.00	−0.39	−0.14	−0.01	0.17	0.40
π_{it}	0.00	−0.17	−0.03	0.00	0.04	0.17
$\mathbb{V}_t(\pi_{it})$	0.01	0.01	0.01	0.01	0.02	0.02
$\text{Skew}_t(\pi_{it})$	−0.01	−1.57	−0.57	−0.02	0.65	1.47

Notes. The summary statistics reported in this table refer to the sample of domestic producers in PRODCOM. The sample includes 5,738 firms observed over 344 months (1995:m1—2023:m12), totaling 596,412 observations.

3.3 Summary Statistics

Our final sample includes 5,738 firms observed over 344 months (1995:m1—2023:m12), resulting in a total of 596,412 observations. Table 1 presents summary statistics of our dataset and documents (i) characteristics of the firms in our dataset, (ii) product market features and (iii) distributional moments of the variables of interest.

Firm characteristics. Our dataset covers the lion's share of domestic manufacturing production in Belgium. Table 1 Panel (a) documents that the average firm in our dataset employs 102 employees (measured in full-time equivalents) and has a domestic turnover (sales) of €3 million per month. Approximately half of the domestic firms in our sample are multi-product firms, meaning they produce multiple 8-digit products within the same industry. Nonetheless, the vast majority of the firms in our sample specializes in only one manufacturing industry. Even for those firms that are active in multiple industries, the contribution of the main industry (i.e. the industry from which most revenue accrues) to total firm revenues is, on average, 98% (median 100%). For the few multi-industry firms, we treat each industry as a separate firm.[20] Finally, our data allow us to track individual firms over a long time span. On average, we observe firms for approximately 8.5 consecutive years (100 months). This feature of the data is particularly important for identification purposes as a long time series enables us to include firm fixed effects in our empirical models to control for time-invariant confounding factors without suffering from the classical Nickell bias (Nickell 1981) that may complicate estimation with dynamic panel models.

Product market structure. Our demand framework in section 2 onboards a role for strategic complementarities within industries. Table 1 panel (b) highlights that this modeling approach is warranted. The typical sector is characterized by a large number of firms with small market shares —the median within-industry share is 1.55%— and a few relatively large producers. To the extent that large firms internalize the effect of their pricing and production decisions on industry aggregates and strategically react to the pricing decisions of their competitors, the monopolistic competition benchmark would effectively be a poor approximation.

At the same time, our framework does not feature strategic interactions across the boundary of the industry. This is justified, as the largest firms are small compared to the volume of economic activity of their macro sector (e.g., "textile manufacturing" or "chemicals and chemical products") and, even more so, compared to the volume of economic ac-

[20] Because most firms operate in only one industry, and the main industry accounts for the lion's share of sales of multi-industry firms, all our results presented below are essentially unchanged if we restrict the sample to the main industry for each firm.

tivity in the whole manufacturing sector in Belgium. It is therefore reasonable to assume that even the largest producers do not internalize the effect of their pricing and production decisions beyond the perimeter of their own industry.

Distributional moments. Table 1 panel (c) reports the percentiles of moments of the distribution of price changes. First, our sample displays a trend inflation of approximately zero at a monthly frequency. Second, on average the dispersion in price changes is approximately 0.14, but displays some variation over time (0.1−0.17) within the 95%-range. Third, the skewness of price changes is zero on average but ranges between negative and positive values. Finally, as indicated by the last three rows of the panel, a significant portion — though not all— of the variation in price changes and the moments of the distributions is due to differences across industries.

Survey responses. Of the baseline sample documented in Table 1, 712 firms also respond to the NBB-BS. For this subsample, we leverage four survey questions which will serve different purposes in our analysis in section 4. In particular, we exploit the *realized* direction of the price change in the current period ($\text{sign}(\Delta p_{fit})$) and the *expected* direction of this price change in the preceding period ($\mathbb{E}_{fit-1}(\text{sign}(\Delta p_{fit}))$). Similarly, we exploit today's *realized* variation in demand ($\text{sign}(\Delta \mathcal{D}_{fit})$) as well as the *expected* direction of the change ($\mathbb{E}_{fit-1}(\text{sign}(\Delta \mathcal{D}_{fit}))$).[21] Panel (a) of Table 2 reveals that, on average, 79.3% of firms keeps prices fixed compared to the preceding month (consistent with U.S. evidence for producer prices from Nakamura and Steinsson 2008). Similarly, on average firms expect not to update prices 74.55% of the time. Around 70.79% of the firms expects demand to remain the same in the upcoming period. In practice, however, realized demand varies more than anticipated (decreases 26.47% of the times and increases 20.75% of the times). The fact that firms systematically underestimate the probability of a change in demand can be interpreted as a sign of incomplete information, with expectations "moving less" than realizations (Angeletos et al. 2021).

Though panel (a) suggests the presence of incomplete information, panel (b) shows consistency with rational behavior. In the latter, we compare the realized changes in prices/demand at time t with the shift that was expected in the preceding month (time $t-1$). The table reveals that firms typically follow up on their expected pricing actions. E.g., for 84.30% of observations where prices are kept constant, the status quo was anticipated the month before. A price increase (decrease) is anticipated in 61.99% (50.75%) of

[21] We use the question regarding realized price changes from the NBB-BS to address measurement concerns regarding p_{fit}, which is based on product-level unit values. Indeed, as shown by Eichenbaum et al. (2014), monthly variation in unit values tends to overstate the frequency of small price changes even with small measurement errors. We circumvent this measurement concern by relying on the NBB-BS to identify price changes.

Table 2: Summary statistics: NBB-BS

Panel (a): Moments					
	Decrease	Unchanged	Increase	s.d.	skewness
$\text{sign}(\Delta p_{fit})$	9.03	79.31	11.66	0.45	0.11
$\mathbb{E}_{fit}(\text{sign}(\Delta p_{fit+1}))$	8.69	74.55	16.76	0.50	0.16
$\text{sign}(\Delta \mathcal{D}_{fit})$	26.47	52.78	20.75	0.68	0.07
$\mathbb{E}_{fit}(\text{sign}(\Delta \mathcal{D}_{fit+1}))$	15.79	70.79	13.42	0.54	-0.02

Panel (b): Expectations vs. realizations				
		$\text{sign}(\Delta p_{fit})$		
		Decrease	Unchanged	Increase
$\mathbb{E}_{fit-1}(\text{sign}(\Delta p_{fit}))$	Decrease	50.75	4.79	3.68
	Unchanged	41.27	84.30	34.33
	Increase	7.98	10.91	61.99
		$\text{sign}(\Delta \mathcal{D}_{fit})$		
		Decrease	Unchanged	Increase
$\mathbb{E}_{fit-1}(\text{sign}(\Delta \mathcal{D}_{fit}))$	Decrease	32.34	10.34	14.32
	Unchanged	55.70	79.14	60.51
	Increase	11.96	10.52	25.18

Notes. The summary statistics reported in this table refer to the sample of domestic producers in PRODCOM that also participate to the NBB-BS. The subsample includes 712 firms and represents a total of 59,731 observations. sign(·) variables are encoded as -1 (“decrease”), 0 (“unchanged”), or +1 (“increase”). Columns 2—4 sum to 100%. Panel (b) depicts a tabulate of expectations vs. realizations of demand and prices. Individuals columns sum to 100%.

Table 3: Persistence of the cost components

	mc_t^n	mc_{it}^n	$\widetilde{mc}_{fit}$
ρ	0.997	0.994	0.926
	(0.018)	(0.009)	(0.036)
P-value	0.733	0.316	0.039

Notes. Estimates of autoregressive processes of order one. Newey-West standard errors are in brackets. P-values for the null hypothesis of a unitary coefficient are reported. Similarly to the tests that we report, a Dickey-Fuller test for the aggregate time series cannot reject the null hypothesis of unit root.

cases. The instances in which they expect to increase (decrease) the price but end up decreasing (increasing) are rare: 7.98% (3.68%) of cases. This observed pattern confirms that firms seem rational in the way they form expectations, with differences from expectations and realizations likely attributable to new information revealed within the month. In fact, in the vast majority of cases, firms correctly forecast the direction of their price changes. Similar conclusions hold for demand realizations and expectations, albeit with less pronounced patterns. Therefore, expected price and demand changes are strong predictors of realized price and demand changes, a feature that we will exploit in section 4.

Marginal costs dynamics. In Table 3 we offer empirical evidence supporting the processes for marginal cost mc_{it}^n and $\widetilde{mc}_{fit}$ posited in equation 4. In particular, we regress our marginal cost measure on its one-month lag, instrumenting the latter with a two-month lag to reduce downward bias due to measurement error. We find strong evidence in favor of a random walk assumption for mc_{it}^n and an AR(1) process with $\rho < 1$ for $\widetilde{mc}_{it}$. The aggregate marginal cost index (constructed as a weighted average of firm-level costs), mc_t^n, also follows a random walk.

4 Econometric Framework

In this section, we discuss the identification and estimation of three parameters of the regression (14), namely the degree of nominal rigidities θ, the degree of strategic complementarities Ω, and the Kalman gain β_{it}. As the presence of a time-varying coefficient β_{it} leads to additional challenges for estimation, we take the route of first studying the regression with constant coefficients, which can be interpreted as a local average treatment effect. This exercise is informative also to confirm a departure from the full-information rational-expectation (FIRE) benchmark, which would predict an estimate of $\beta_{it} = 1$ for all industries in all periods. We then explore the cyclical properties of β_{it} in section 5.

4.1 Parameter Identification

We now discuss formally how identification of the parameters is achieved using data on subjective expectations and lags of observables. From the Dynamic Pass-through Regression (equation 14), identification requires finding a set of instrumental variables $\mathcal{I}_{fit}^{IV}$ that are relevant and satisfy the exclusion restriction:

$$\mathbb{E}\left(\eta_{fit} + \widetilde{u}_{fit} \mid \mathcal{I}_{fit}^{IV}\right) = 0,$$

where η_{fit} is the noise in the signal and $\widetilde{u}_{fit}$ is the sampling error due to the nominal rigidity. First, we notice that, by proposition 1, $\widetilde{u}_{fit}$ is uncorrelated with any variable in the information set $\mathcal{I}_{fit}$. Second, η_{fit} is i.i.d. over time and therefore uncorrelated with any variable that is lagged. Therefore, we include lags of observables which help identification while satisfying the exclusion restriction as discussed in more detail in the next subsection.

Nevertheless, to identify the Kalman gain one needs to use variation that comes from news that is not contained in lags. We address this challenge by leveraging the data on subjective expectations from the NBB-BS. However, isolating variation in realized prices that comes from expectations *and* is orthogonal to the noise in the signal is a daunting task. The reason is that expectations are directly a function of the noise in the signal from equation (12). Consequently, the expectation itself is generally not a good shifter of realizations (under incomplete information) because it is correlated with the error term of the Dynamic Pass-through Regression. However, we can show that a coarser transformation of the expectation —the expected sign of a price change— is a valid shifter under a weak condition of symmetry on the shock distribution that is satisfied, for example, under our assumption of Gaussian cost shocks.

Proposition 2 (Identification)**.** *Suppose that the unconditional distribution of price changes is symmetric, i.e.* $\mathbb{P}(\Delta p_{fit+1}^{o} > 0) = \mathbb{P}(\Delta p_{fit+1}^{o} < 0)$*. Then the expected sign of a price change satisfies the exclusion restriction:*

$$\mathrm{Cov}\left(\eta_{fit},\, \mathbb{E}_{fit}(\mathrm{sign}(\Delta p_{fit+1}))\right) = 0.$$

The intuition for the proof (in Appendix A.3) is that, as we are assuming no trend inflation and shocks are Gaussian, the *unconditional* probability of a price increase equals that of a price decrease. Because the subjective expectation of the sign is given by the difference between the two *conditional* probabilities, and firms can recover the noise in the signal after setting prices, the unconditional covariance with the error term is exactly zero by the law of iterated expectations. In other words, when estimating the regression across a sufficiently vast sample size that includes both expansion and recession periods, the noise is correlated with the expected "size" of the adjustments, but not with the expected "sign."

Proposition 2 is a useful result. Generally, even with perfect data on expectations, identification of the Kalman gain is a complex task. However, a simpler statistics like the expected direction of a price change, which is arguably easier to elicit in surveys and less prone to measurement error, is sufficient to identify the parameters of the moment condition that relates current and future price changes from data on realized prices.

In the next section, we discuss the details of the implementation and results.

4.2 Baseline Analysis

Specification. We begin by bringing the Dynamic Pass-through Regression (14) to the data for the special case of constant and homogeneous coefficients $\beta_{it} = \beta$ for all i and t. This leads to the baseline specification:

$$\Delta p_{fit} = FE_f + FE_{it} + \frac{(1-\theta)^2}{\theta}\Big((1-\Omega)mc^n_{fit} - p_{fit}\Big) + \beta\Delta p_{fit+1} + \varepsilon_{fit}, \tag{16}$$

where FE_f is a firm fixed effect, FE_{it} is an industry-by-time fixed effect, and ε_{fit} is a composite regression residual. Under the benchmark model with full information and rational expectations (FIRE), $\beta = 1$ and the coefficient is redundant. Therefore we can test for the presence of information frictions by checking whether the coefficient is statistically smaller than one.

Instrument set. As shown in proposition 1, lags of endogenous variables are orthogonal to both the noise in the signal, which is i.i.d. over time, and the sampling error due to the nominal rigidity, which is orthogonal to the information set $\mathcal{I}_{fit}$. Moreover, as we further discuss in the next paragraph, they provide power for identification. Therefore, we include lagged marginal cost (in the previous quarter) and lagged price level (in the previous month) in the instrument set. These lags deal with a number of potential identification concerns which include measurement error in marginal cost, which is arguably the main concern regarding marginal cost (Amiti et al. 2019), as well as the Nickell bias which can arise in the context of dynamic panel models (Nickell 1981).

In addition, we include the contemporaneous expected sign of a price change and three lags of it, which is a valid instrument under the assumption of proposition 2. This leads to the following baseline instrument set:

$$\mathcal{I}^V_{fit} = \left\{ \left\{ \mathbb{E}_{fit-\tau}(\text{sign}(\Delta p_{fit-\tau+1})) \right\}^3_{\tau=0}, \left\{ mc^n_{fit-\tau} \right\}^6_{\tau=3}, p_{fit-2} \right\}.$$

Sources of identification. We now discuss how the instrument set and set of fixed effects provide identification power for the parameters of interest. First, from the equation for the static target price (5), the degree of strategic complementarities (Ω) is identified from contemporaneous movements in marginal cost (mc^n_{fit}) relative to the industry price index (p_{it}), which captures movements in markups. Therefore, Ω is identified only from industry cross-sectional variation *within* the period t. Such variation is leveraged by our specification by including industry-by-time fixed effects and lagged marginal cost as a shifter of current marginal cost, which is highly persistent.

Second, from the expected law of motion, the degree of nominal rigidities (θ) is identified from movements in the current price level p_{fit} due to changes in the target price p^o_{fit} relatively to the price level in the previous period p_{fit-1}. Hence, θ is identified from time-series variation *across* periods t and $t-1$. This variation is produced by including lags of both marginal cost and the firm price level in the instrument set.

Third, from the equation for the reset price (8), the Kalman gain β is identified from variation in the expected present value of future marginal costs relative to the current value of marginal cost. Therefore, β is identified from time-series variation *across* periods t and $t+1$. We leverage this variation by including both realized marginal cost and expectation of future price changes, which convey information regarding the expected evolution of future costs.

Implementation. Whereas industries are narrowly defined at the four-digit NACE level, we allow for correlated shocks across industries in the computation of standard errors. In turn, we cluster standard errors at the two-digit level so that firms that are in different industries but in the same sector can be exposed to common shocks.

We estimate the baseline regression (16) via the generalized method of moments (Hansen 2010). This allows us to impose the coefficient restrictions from theory and strengthen the identification. We employ a two-step estimator with a weighting matrix at the same level as the clustering (two-digit NACE).

Finally, all the regressions are weighted using Törnqvist weights, i.e. a weighted average between current and (one-month) lagged revenue weights. This weighting scheme replicates accurately the construction of the aggregate price indexes (section 3.2).

Baseline results. Results from estimating regression (16) are reported in Table 4, column A. First, the estimate of the degree of nominal rigidities $\theta = 0.832$ (0.008) is precisely estimated. This number is consistent with a substantial role for the nominal rigidity in reducing the pass-through at a monthly frequency as it implies an average duration of a price spell of roughly a year. As price changes are infrequent, firms' behavior is forward looking.[22]

Second, the estimate of the degree of real rigidities $\Omega = 0.597$ (0.022) is also remarkably precise. This number implies that real rigidities alone reduce the pass-through by roughly 60%, that is firms adjust markups in response to changes in costs by absorbing a large fraction of the shock into their profits. This is in line with estimates obtained by previous literature that employ special cases of our Dynamic Pass-through Regression (Amiti et al. 2019 and Gagliardone et al. 2023).

To the best of our knowledge, this is the first paper providing an estimate for β using data on firms' costs, prices, and expectations. The estimate $\beta = 0.823$ (0.027) is remarkably precise. As the test for the full-information rational expectation hypothesis shows, this value implies a statistically significant departure from the full-information benchmark of $\beta = 1$ at any standard confidence level. Moreover, the estimate indicates a substantially higher "discounting" by about 20% than what is used in standard calibrations. Consistently with evidence from previous literature (e.g. Angeletos et al. 2021), expectations of future price changes seem to move significantly less than realizations. This leads to a form of present bias/myopia or hyperbolic discounting in price setting in response to shocks (e.g. Farhi and Werning 2019, García-Schmidt and Woodford 2019), which can be rationalized by information frictions (e.g. Gabaix and Laibson 2017, Angeletos and Lian 2018, Angeletos and Huo 2021). Notably, our evidence shows that the additional discounting is present at the micro-level and therefore is not (only) due to aggregation but (also) to individual behavior.

Finally, Hansen's J overidentification test does not reject the null hypothesis of valid instruments for any standard confidence level, confirming that instruments are plausibly exogenous in line with the theory.

The estimates of the structural parameters imply an estimate for the overall pass-through from costs to prices. In particular, the effect of a contemporaneous shock to marginal cost is $\frac{(1-\theta)^2}{\theta}(1-\Omega) \approx 1.6\%$, which is a large number at monthly frequency. To put this number in context, we will show in section 6 that this coefficient maps under empirically plausible assumptions directly into the slope of the marginal cost-based Phillips curve, and implies an estimate of the slope at a quarterly frequency above 5%.[23] Though this number is consistent with our earlier work (Gagliardone et al. 2023), it is an order of

[22] To provide additional confidence regarding this estimate, we compute the fraction of firms that do not change price in a month using additional PPI micro-data and obtain an estimate of 0.8 (0.023), which aligns closely with the regression result.

[23] The quarterly slope is computed as $\frac{(1-.7)^2}{.7}(1-\Omega) \approx 7.7\%$ because Ω is independent of the frequency.

Table 4: Estimation of the Baseline Regression

	(A)	(B)	(C)
θ	0.832	0.813	0.814
	(0.008)	(0.011)	(0.008)
Ω	0.597	0.636	0.645
	(0.022)	(0.028)	(0.016)
β	0.823	0.642	0.654
	(0.027)	(0.054)	(0.041)
Firm FE	y	y	y
Ind × time FE	y	y	y
Demand instrument	n	y	n
Demand control	n	n	y
	Hansen's J overidentification test		
χ^2-stat	5.925	5.911	5.520
P-value	0.431	0.433	0.477
	Test of FIRE $\beta = 1$		
Z-stat	6.44	6.67	8.37
P-value	0.000	0.000	0.000

Notes. This table provides GMM estimates for the baseline regression (16) at monthly frequency:

$$\Delta p_{fit} = FE_f + FE_{it} + \frac{(1-\theta)^2}{\theta}\left((1-\Omega)mc^n_{fit} - p_{fit}\right) + \beta \Delta p_{fit+1} + \varepsilon_{fit}.$$

Column A's orthogonality condition is implemented by using the baseline instrument set. Column B augments the instrument set with expectations of future demand growth. Column C generalizes previous demand systems by allowing for firm demand shocks, which leads to an additional control for $\mathbb{E}_{fit}(\text{sign}(\Delta\mathcal{D}_{fit+1}))$ in the population regression (14). Regressions are weighted using Törnqvist weights. GMM is implemented using a two-step estimator with a weighting matrix and standard error clustering at the sector level.

magnitude larger than previous estimates from the macroeconomic literature (e.g. McLeay and Tenreyro 2020, Hazell et al. 2022) and suggests a relatively steep cost-based Phillips curve even in "normal times." A novel lesson of this paper is that there is a substantial dampening of the pass-through of news regarding future costs. In particular, the effect of the latter is given by $\beta\frac{(1-\theta)^2}{\theta}(1-\Omega)\approx 1.2\%$. This number is surprisingly low as it implies that a monetary policy announcement of a change in policy that will be implemented only a month later leads to a 20% lower effect on inflation (and thus a much larger effect on output and unemployment) than the same policy if implemented immediately. Moreover, this estimate provides empirical validation for theories that introduce departures from full-information rational expectations in the Phillips curve as a way to address the forward guidance puzzle (e.g. Gabaix 2020) or connect with experimental evidence on inflation expectations (Coibion et al. 2018a, Coibion et al. 2020).

4.3 Robustness to Demand

In this section, we extend the baseline analysis to incorporate information about the expected demand growth, which is arguably a factor influencing firms' price setting.

Demand expectations in the instrument set. Under the demand system (1), knowledge by firms of future demand $\mathcal{D}_{fit+\tau}$ for $\tau>0$ improves the ability of the firm to forecast future marginal costs to the extent that returns to scale are not constant ($\nu_f\neq 0$), and thus their ability to choose a reset price. That is, any available information regarding future demand should be incorporated into pricing decisions. Nevertheless, forecasts of future prices $\mathbb{E}_{fit}(\Delta p_{fit+1})$ should already summarize all the relevant information for pricing decisions, including that related to future demand. If this is not the case in practice, it is useful to directly include measures of expected demand growth, $\mathbb{E}_{fit}(\mathrm{sign}(\Delta\mathcal{D}_{fit+1}))$, in the instrument set $\mathcal{I}^V_{fit}$. Accordingly, in column B of Table 4, we extend the instrument set to include subjective forecasts of demand growth from the survey:

$$\mathcal{I}^V_{fit}=\left\{\left\{\mathbb{E}_{fit-\tau}(\mathrm{sign}(\Delta p_{fit-\tau+1}))\right\}_{\tau=0}^{3},\ \left\{\mathbb{E}_{fit-\tau}(\mathrm{sign}(\Delta\mathcal{D}_{fit-\tau+1}))\right\}_{\tau=0}^{3},\ \left\{mc^n_{fit-\tau}\right\}_{\tau=3}^{6},\ p_{fit-2}\right\}.$$

The estimates of both θ and Ω remain remarkably stable, showing robustness to the inclusion of the additional instrument. The estimate of $\beta=0.642$ (0.054) is lower than that in column A, but not statistically different at 95% confidence level. Hansen's J test is again robustly passed and, moreover, the test of FIRE is soundly rejected at any standard confidence level. Overall, this exercise indicates that including a demand shifter in the instrument set does not alter the conclusion of departure from the full-information rational-expectation benchmark.

Demand shocks. Alternatively, we can replace the assumption regarding the demand system (1) and allow for firm-specific i.i.d. demand shocks, $\varphi_{fit+1} \notin \mathcal{I}_{fit}$, as follows:

$$\mathcal{D}_{fit} := d(p_{fit},\, p_{it},\, \varphi_{fit+1}) \cdot Y_{it} \quad \forall f \in \mathcal{F}_i. \tag{17}$$

Under the stochastic demand system (17), the loglinear markup varies also because of firm demand shocks and the static target price (equation 5) becomes:

$$p^{\star}_{fit} = \mu_f + (1-\Omega)(mc^n_{fit} + \varphi_{fit+1}) + \Omega p_{it}.$$

Following similar steps as before, the baseline regression (16) under the stochastic demand system becomes:

$$\Delta p_{fit} = FE_f + FE_{it} + \frac{(1-\theta)^2}{\theta}\Big[(1-\Omega)\Big(mc^n_{fit} + \mathbb{E}_{fit}(\varphi_{fit+1})\Big) - p_{fit}\Big] + \beta \Delta p_{fit+1} + \varepsilon_{fit}, \tag{18}$$

Here $\mathbb{E}_{fit}\,\varphi_{fit+1}$ is a relevant control: failure to observe demand expectations would result in an omitted variable bias if demand expectations are systematically correlated with the present value of marginal cost. We measure $\mathbb{E}_{fit}(\varphi_{fit+1})$ with $\mathbb{E}_{fit}(\text{sign}(\Delta\mathcal{D}_{fit+1}))$ and use the same instrument set as specification A. Column C of Table 4 shows that this systematic correlation seems modest: when controlling for demand expectations, estimates of all the parameters are within the confidence bands of the baseline specification at 95% confidence level. Moreover, the additional robustness leads to similar results as including the expectation in the instrument set (column B), and thus does not affect the conclusion of a departure from full information.

5 State Dependence of the Pass-through

In the aftermath of the inflation surge that occurred between 2021 and 2023, the macroeconomic literature has devoted significant effort to reconciling the pre-pandemic evidence of the low aggregate pass-through from measures of real activity to inflation with the steep rise in prices that followed the sudden spike in energy costs. One explanation that seems able to explain the recent events is a state-dependent pass-through, which is low in "normal times" but increases as soon as large shocks hit the economy. Such a nonlinear response of prices would have key implications for the quantification of monetary non-neutrality because it would imply that the same interest rate change has different effects depending on the pre-policy state of the economy, as well as the mix of shocks that are driving the business-cycle fluctuations.

However, the channels through which such nonlinearities are generated are still largely subject to debate.[24] In this section, we provide reduced-form micro-evidence that the pass-through is indeed state-dependent. We will then argue that our model under full information and nominal rigidities à la Calvo would not be able to explain such evidence, but the model with information frictions can rationalize it. Finally, we provide empirical validation of the mechanism through information frictions by extending the Dynamic Pass-through Regression and using data on the dispersion in beliefs that speak directly to the specific features of the model. We argue that, though information frictions hardly account for the entire time variation of the pass-through, they can nevertheless explain a relevant share of its fluctuations. This suggests that the interaction between information frictions and other possible sources of state dependence is likely to be quantitatively important.

5.1 Reduced-form Evidence of State-Dependent Pass-through

We now discuss evidence that the pass-through is state-dependent. We aim to provide such evidence in reduced form without relying heavily on the structure of the model from section 2. Nevertheless, the regression that we run is consistent with the structural approach of the Dynamic Pass-through Regression (equation 14), and results can be interpreted through the lenses of the model. In section 5.3, we then impose additional structure and discuss identification more formally.

The basic idea behind our reduced-form approach is that, under nominal frictions, prices are forward-looking and therefore correlated with leads of marginal cost. This correlation should persist until prices are reset, after which price changes become independent of past cost shocks. The decline in the correlation should reflect not only the effect of the nominal friction but also the persistence of the shocks and the discounting that firms use to compute the present value. With real rigidities, this correlation is also mediated by movements in the industry price index, which controls for changes in markups due to movements in relative prices. The above reasoning leads to a regression of relative price changes (change over time of firm price relative to the industry price) on relative cost shocks (change over time of firm costs relative to the industry average), which can be implemented including fixed effects as follows:

[24] Cavallo et al. (2023), Blanco et al. (2024a), Blanco et al. (2024b), and Gagliardone et al. (2024) provide evidence that the frequency of price changes is state-dependent and rationalize it using menu cost models as in Golosov and Lucas (2007) and Nakamura and Steinsson (2008). On the other hand, Pfäuti (2023), Weber et al. (2023), Bracha and Tang (2024), and Afrouzi et al. (2024) provide evidence that attention is state-dependent and rationalize it using models with incomplete information. As the frequency of price changes acts in the menu cost model effectively as a time-varying discount factor (see equation 6 and related discussion in Dotsey and King 2005, or Alvarez and Lippi 2022), it is particularly challenging to disentangle the two stories using data. We provide evidence in section 5.3 that even accounting for the observed time variation of the frequency of price changes still leaves room for state dependence of the estimated Kalman gain.

Figure 1: Reduced-form Evidence of State-dependent Pass-through

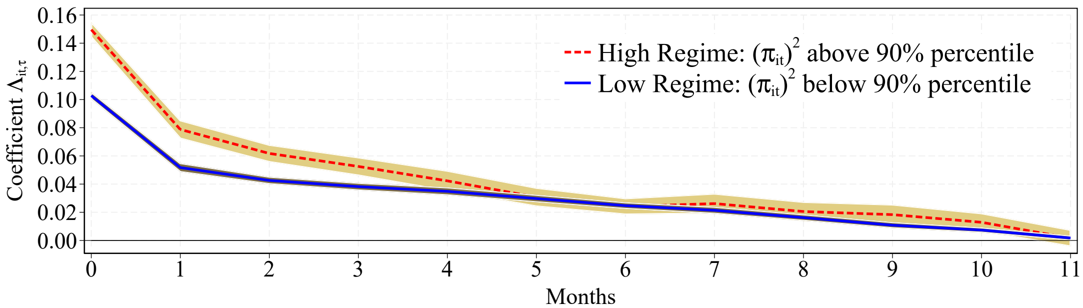

Notes. Estimates of the reduced-form coefficients of regression (19). The red dashed line (with yellow bands) is computed on a subsample corresponding to all the firm-level observations belonging to an industry with an inflation rate in the top 10% of all the industry inflation realizations (corresponding to about 8%). The black line (with blue bands) is estimated on the remaining subsample. Confidence bands at 90% are computed using robust standard error. Regressions are weighted using Törnqvist weights.

$$\Delta p_{fit} = FE_f + FE_{it} + \sum_{\tau \geq 0} \Lambda_{it,\tau} \cdot \Delta mc_{fit+\tau} + \Lambda_0 p_{fit-1} + \varepsilon_{fit}. \tag{19}$$

As before, FE_f is a firm-specific intercept and FE_{it} is an industry-by-time fixed effect, which implies that all the variables are demeaned with respect to the industry average. The lagged price is included as a control for shocks up to time $t-1$.

The pass-through coefficient at horizon τ is given by $\Lambda_{it,\tau}$, which can in principle be heterogeneous across industries and time-varying. It follows from the derivations in section 2.4 that $\Lambda_{it,\tau}$ does not depend on time in the full-information benchmark to the extent that parameters are not time-varying:

$$\Lambda_{it,\tau} = \Lambda_\tau = \frac{(1-\Omega)(1-\theta)^2}{1-\theta\rho}\rho^\tau.$$

We now argue that this prediction of a constant pass-through of the full-information benchmark is at odds with the reduced-form evidence. Figure 1 plots the coefficients of the regression estimated on the two subsamples: the red dashed line corresponds to a subsample of industries-month pairs in which industry inflation was "large" in absolute value, corresponding to the top 10 percent of all realizations of industry inflation rates (in absolute value, approximately 8% or higher), whereas the solid black line is estimated over the remaining part of the sample. The two subsamples reflect disproportionately historical periods of large inflation or deflation such as the surge in 2021-2023 or the recession in 2009-2010, but also exploit industry-level variation in inflation which has substantially more volatility than aggregate inflation.

Let us now analyze the estimates.[25] First, consistently with the presence of a nominal friction, contemporaneous price changes are correlated with idiosyncratic cost shocks up to about a year. After that, coefficients are not statistically significant. In turn, current prices do not reflect innovations beyond the one-year horizon. This is in line with our structural estimates from Table 4 that imply an average duration of a price spell of about a year.

Second, there is a steep drop in the correlation when moving from a contemporaneous cost shock to the innovation one month ahead, after which the correlation declines more smoothly. This is consistent with a role for a form of present bias or hyperbolic discounting (Gabaix and Laibson 2017) in price setting as discussed in section 2.5. In other words, the steep drop suggests that the precision of the information regarding future costs has a discontinuity at very short horizons. This feature of the cost-price correlation is what identifies the Kalman gain in the structural estimation.

Third, the two subsamples lead to statistically different pass-through estimates. We interpret this as evidence of a dependence of the pass-through on industry variables, and, in particular, industry inflation. In fact, we notice that there is no a-priori reason why the estimates over the different subsamples should be statistically different, with the estimated pass-through being larger in high-inflation industries. In particular, *relative* price changes need not mechanically correlate more with *idiosyncratic* cost shocks in industries displaying large inflation rates, whose direct effect is already soaked up by the industry-time fixed effect. This is indeed the correct intuition for the full information benchmark, in which large industry shocks lead to high industry inflation, but that effect only generates an increase in the common industry trend and no effect on the elasticity of firms' relative prices. Under incomplete information, firms cannot forecast the exact persistence of the shocks; as they observe an increase in their costs and an associated large increase in industry inflation, they conclude that the shock is nominal and thus persistent. In turn, a high perceived persistence leads to a large price response to the shock. Accordingly, when selecting subsamples with large industry shocks, i.e. sampling more from the distribution of price changers, one would detect in the data a statistically larger elasticity of relative prices to idiosyncratic cost shocks compared to "normal times," consistently with the reduced-from evidence.

In the next section, we model a mechanism via information frictions that leads to such state-dependent behavior of the pass-through.

[25]We refrain from making quantitative statements on the overall pass-through for concerns regarding downward bias due to measurement error, which is not addressed by ordinary least squares. We report in Figure A.1 of Appendix B the scatterplots of price changes against cost changes for different horizons, which also indicate a bias due to —possibly mismeasured— outliers. Nevertheless, the qualitative patterns are informative as a first step. In section 5.3, we address measurement error among other concerns via a structural estimation approach.

5.2 State Dependence via Information Frictions

The mechanism is modeled as follows. A large persistent industry cost shock has two effects. First, a direct effect that increases the dispersion in prices and dispersion in price growth rates —keeping fixed the Kalman gain— because firms are subject to a nominal rigidity. Intuitively, when the shock hits, some firms adjust prices whereas others do not adjust, leading to an increase in the dispersion. This is a standard channel in any model with nominal frictions. Second, there is an indirect effect that amplifies the price response under incomplete information. The increase in dispersion through the first channel leads to an increase in the Kalman gain, i.e. an increase in the *speed* at which beliefs are updated. The latter channel generates an extra boost in inflation, further price dispersion via the direct channel, and therefore additional rounds of adjustments. At the end of this process, inflation, dispersion, and the pass-through are all higher compared to the case in which the Kalman gain is fixed at the pre-shock level.

Let us now break down the mechanism into steps and provide evidence.

Fact (1): Large shocks increase price dispersion. First, as we show in Appendix A.2, we can decompose the variance of price changes under Calvo adjustments in the following direct and indirect effects:

$$\mathbb{V}(\Delta p_{fit}) = \underbrace{\frac{\theta}{1-\theta}\pi_{it}^2}_{\text{Direct Effect}} + \underbrace{(1-\theta)\mathbb{V}(\Delta p_{fit}^o)}_{\substack{\text{Indirect Effect} \\ \text{through change in } \beta_{it}}} \tag{20}$$

The first term is the direct effect, which is present simply because of the nominal rigidity. The second term captures the indirect effect, which depends on the variance of conditional price changes. In particular, the latter depends on industry inflation (in absolute value) only through the change in β_{it}. Hence, the indirect effect is absent if the Kalman gain is fixed. Therefore industry inflation should be positively correlated with the dispersion of price changes via both channels provided that β_{it} is increasing in $|\pi_{it}|$.

To corroborate this argument with data, we compute the industry dispersion of price changes and the industry inflation rate in every month and industry. We then demean variables with a sector-by-time fixed effect (defined at the 2-digit NACE level), which serves the purpose of isolating industry variation by removing aggregate trends from the variables. Figure 2 is the scatterplot of these detrended variables. As we can see, whenever industry inflation is large (in absolute terms), industry dispersion is large. This confirms that the dispersion in price changes is indeed positively correlated with the absolute value of industry inflation.[26]

[26] Similar patterns of inflation and dispersion are documented for retail prices by Sara-Zaror (2022).

Figure 2: Evidence that large shocks increase price dispersion

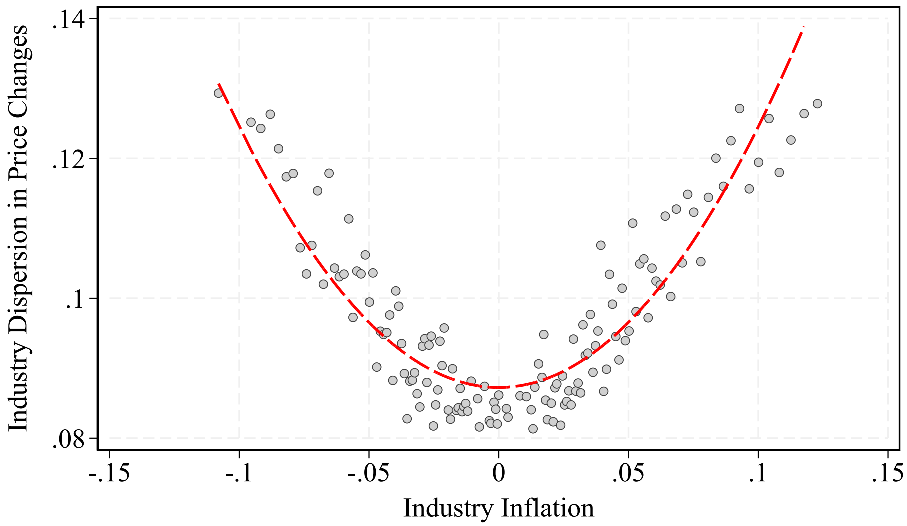

Notes. Scatterplot of monthly industry inflation against the monthly industry dispersion of price changes. The scatterplot removes a sector-by-time fixed effect (2-digit NACE). Observations are weighted using the Törnqvist weights. The regression line (red dashed line) is obtained by regressing the dispersion on the square of inflation.

Fact (2): Belief dispersion predicts price dispersion. We now want to establish that the Kalman gain β_{it} is increasing in the variance of price changes $\mathbb{V}(\Delta p_{fit+1})$. This will conclude the argument as it implies that both the direct and indirect channels of equation (20) operate in the same direction, leading to a feedback loop that generates a high pass-through whenever inflation is high and persistent.

Though we cannot directly observe the Kalman gain, as we show in Appendix A.2, β_{it} is a monotone increasing function of the industry dispersion in posterior beliefs:

$$\mathbb{V}(\mathbb{E}_{fit}(\Delta p_{fit+1})) = \beta_{it} \cdot \sigma^2_{e,it} \tag{21}$$

As β_{it} is an increasing function of the prior variance $\sigma^2_{e,it}$ (equation 12), and the dispersion in posterior beliefs is high whenever the prior variance is high (equation 21), we can then use the observed dispersion in expectations as a proxy for the unobserved Kalman gain to asses qualitative patterns driven by fluctuations of the prior variance.

We measure the industry dispersion of beliefs with the industry dispersion of our survey measure of expected price changes and correlate it against the dispersion in price changes occurring in the year after the belief is reported. We consider these longer changes because, as we argued, the duration of a price spell is roughly a year. As before, we detrend variables with a sector-by-time fixed effect and report the scatterplot of the data.

Figure 3: Evidence that belief dispersion predicts price dispersion.

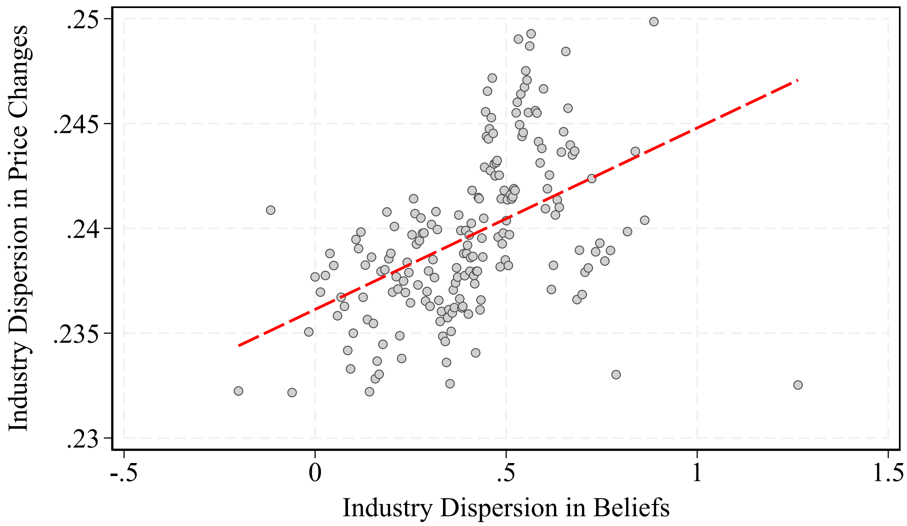

Notes. Scatterplot of the industry dispersion in beliefs against the industry dispersion in annual price changes. The dispersion in beliefs is constructed as the industry dispersion of the survey expectation of the sign of a price change. The scatterplot removes a sector-by-time fixed effect (2-digit NACE). Observations are weighted using the Törnqvist weights. The regression line (red dashed line) is obtained by regressing the dispersion of price changes on the dispersion in beliefs.

As Figure 3 shows, a high industry dispersion in beliefs predicts a high dispersion in price changes over the subsequent year. This provides evidence for a mechanism in which the prior variance $\sigma^2_{e,it} := \mathbb{V}_{it}(\Delta p_{fit+1|t})$ contains information on the (ex-post) realized dispersion $\mathbb{V}(\Delta p_{fit+1})$, and therefore, in response to a persistent aggregate shock which increases the dispersion (fact 1), firms' uncertainty regarding the evolution of the industry distribution increases (high $\sigma^2_{e,it}$) leading to a larger Kalman gain (high β_{it}) and a higher sensitivity to shocks. The higher pass-through then leads to a further increase of the price dispersion, which in turn increases even more the Kalman gain and so on. Also, we notice that this mechanism can rationalize the reduced-form evidence discussed in the previous section. Finally, the mechanism described here is consistent with the evidence from the natural experiment studied in Drenik and Perez (2020).

Fact (3): Belief dispersion is high when inflation volatility is high. Putting together the two previous facts, we expect to find a positive correlation between aggregate dispersion in beliefs and the aggregate inflation rate (in absolute value). To show this, we compute a centered annual moving average of the dispersion in the survey measure of beliefs to capture business-cycle fluctuations of the variable. We then compute aggregate annual

Figure 4: Evidence that belief dispersion is positively correlated with inflation volatility.

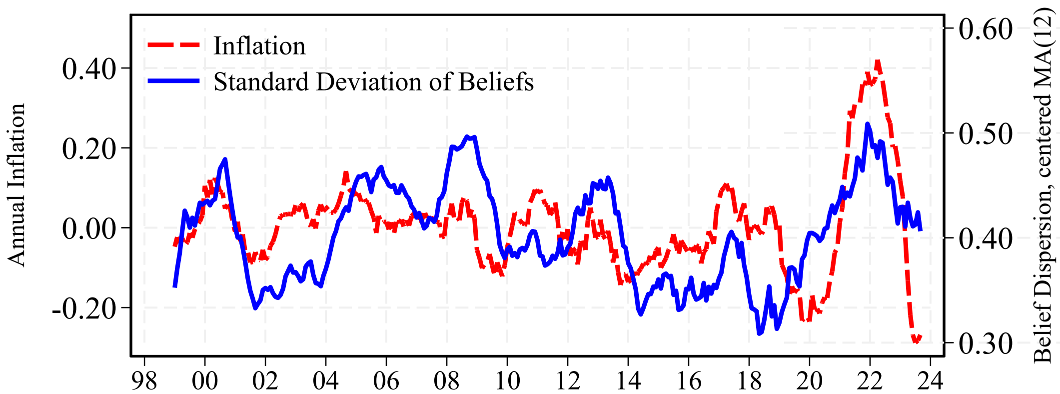

Notes. Time series of annual aggregate inflation and aggregate standard deviation of beliefs. All variables are aggregated using the Törnqvist weights. Inflation is detrended by removing from the aggregate price level a piece-wise linear trend with one break. The standard deviation is smoothed using a centered 12-month moving average.

inflation and remove the linear trend of the price level to be able to interpret the inflation rate as log-deviations from trend inflation consistently with the model.

Figure 4 reports the two time series for inflation and the standard deviation of beliefs. As the plot shows, in periods in which inflation has been above trend, the dispersion in beliefs has also been high. In particular, the dispersion in beliefs has been high at the beginning of the sample (1999-2001), in the period 2005-2008, and during the inflation surge in 2021-2023. The dispersion instead decreased in periods of moderate inflation such as the 2002-2004 period, the years that followed the 2009-2010 recession, and the period of moderation that occurred between 2014 and 2017. Finally, in the two episodes in our sample of deflation corresponding to the sovereign debt crisis in 2012-2014 and the pandemic in 2019-2020, the dispersion rose.

This aggregate correlation is not new to the literature. In fact, Mankiw et al. (2003) documented similar patterns using data for the US, from which the authors also concluded that a model with information frictions is able to rationalize the evidence. Overall, these three correlations suggest a possible quantitative role for movements in higher-order moments of the price and belief distributions in explaining inflation, i.e. the first moment. We provide a more formal identification of this quantitative role via the Dynamic Pass-through Regression in the next section.

5.3 Identification via the Dynamic Pass-through Regression

Now that we established the correlations underlying a possible mechanism of state-dependent pass-through via information frictions, let us address the identification more formally via the Dynamic Pass-through Regression. We extend the analysis from section 4 to include in regression (16) an interaction term between the Kalman gain and the squared industry inflation rate, as follows:

$$\begin{aligned}\Delta p_{fit} =& FE_f + FE_{it} + \frac{(1-\theta)^2}{\theta}\Big((1-\Omega)mc^n_{fit} - p_{fit}\Big) + \beta \cdot \mathbb{I}(\pi^2_{it} \geq 90\%) \cdot \Delta p_{fit+1} \\ &+ \beta \cdot \mathbb{I}(\pi^2_{it} < 90\%) \cdot \Delta p_{fit+1} + \varepsilon_{fit},\end{aligned} \tag{22}$$

where $\mathbb{I}$ denotes an indicator function for the industry inflation regime.

As before, we are regressing price changes on a firm and industry-by-time fixed effect, marginal cost, the relative price level, and future price changes. Differently from before, we allow for two separate Kalman gain coefficients depending on the industry inflation regime, effectively splitting the estimation sample for this parameter into two depending on whether industry inflation is large (top 10% of all the inflation realizations) or low (remaining 90%). As firms observe the price level (section 2.1), adding this interaction term does not change the requirements for parameter identification (proposition 2). Therefore we proceed by using the baseline instrument set and perform robustness including demand expectations as an additional instrument.[27]

Columns A and B of Table 5 report the results. The estimates of θ and Ω are in line with the baseline analysis, suggesting that accounting for the state dependence of the pass-through does not significantly affect the estimates of these parameters. On the other hand, we estimate different Kalman gains for the two inflation regimes, with $\beta_{it} = 0.984$ (0.079) in the high inflation regime and $\beta_{it} = 0.690$ (0.067) in the low inflation regime when using the baseline instrument set, and $\beta_{it} = 0.950$ (0.080) in the high inflation regime and $\beta_{it} = 0.270$ (0.228) in the low inflation regime when including also the demand instrument. Though the split of the sample leads to larger robust standard errors, we can nevertheless reject a null hypothesis of constant coefficients at a 10% confidence level. Interestingly, our estimates imply that the Kalman gain in the high-inflation regime becomes close to one, that is firms behave as if they had approximately complete information whenever inflation is high. This result is in line with experimental evidence from Weber et al. (2023), which argues that firms privately acquire information in high-inflation regimes and therefore respond less to exogenously provided information regarding inflation.

To address potential confounding mechanisms via the state-dependent frequency of

[27] Results are also robust to including demand expectations as a control.

Table 5: Estimation of the Regression including the Interaction

	(A)	(B)	(C)	(D)
θ	0.821	0.789		
	(0.011)	(0.017)		
Ω	0.621	0.650	0.569	0.600
	(0.024)	(0.030)	(0.013)	(0.019)
$\beta \cdot \mathbb{I}(\pi^2_{it} \geq 90\%)$	0.984	0.950	0.645	0.516
	(0.079)	(0.080)	(0.021)	(0.061)
$\beta \cdot \mathbb{I}(\pi^2_{it} < 90\%)$	0.690	0.279	0.477	0.317
	(0.067)	(0.228)	(0.021)	(0.022)
Demand instrument	n	y	n	y
Time-varying θ_t	n	n	y	y
	Hansen's J overidentification test			
χ^2-stat	5.817	5.438	6.062	6.084
P-value	0.324	0.364	0.416	0.413
	Test of constant coefficients $\beta \cdot \mathbb{I}(\pi^2_{it} \geq 90\%) = \beta \cdot \mathbb{I}(\pi^2_{it} < 90\%)$			
Z-stat	2.16	3.15	1.77	1.86
P-value	0.031	0.002	0.077	0.063

Notes. This table provides GMM estimates for the baseline regression (22) at monthly frequency:

$$\Delta p_{fit} = FE_f + FE_{it} + \frac{(1-\theta)^2}{\theta}\left((1-\Omega)mc^n_{fit} - p_{fit}\right) + \beta \cdot \mathbb{I}(\pi^2_{it} \geq 90\%) \cdot \Delta p_{fit+1} + \beta \cdot \mathbb{I}(\pi^2_{it} < 90\%) \cdot \Delta p_{fit+1} + \varepsilon_{fit}.$$

Column A's orthogonality condition is implemented by using the baseline instrument set. Column B augments the instrument set with expectations of future demand growth. Columns C and D include a time-varying measure of the aggregate frequency of price changes. Regressions are weighted using Törnqvist weights. GMM is implemented using a two-step estimator with a weighting matrix and standard error clustering at the sector level.

price changes in menu cost models (Cavallo et al. 2023, Blanco et al. 2022, Gagliardone et al. 2024), we perform an additional robustness exercise in which we use data on the aggregate realized frequency of price changes. To avoid measurement error due to our measure of prices as unit values (footnote 21), we use sign(Δp_{fit}) from the NBB-BS. We then compute the fraction of firms that report having changed prices and use this measure in place of the coefficient θ, effectively allowing for an additional control in the regression in the form of a time-varying θ_t. We report a moving average of this time series in Figure A.2 of Appendix B.

Columns C and D of Table 5 report the results. First, the estimate of Ω is consistent with our previous results, further consolidating confidence with respect to the identification of this parameter. Second, the estimates of the two interaction terms are both lower than the estimates from columns A and B and become closer to each other, suggesting that there is indeed a rich interaction between nominal and information rigidities. In particular, we obtain estimates of $\beta_{it} = 0.645$ (0.021) in the high inflation regime and $\beta_{it} = 0.477$ (0.021) in the low inflation regime when using the baseline instrument set, and $\beta_{it} = 0.516$ (0.061) in the high inflation regime and $\beta_{it} = 0.317$ (0.022) in the low inflation regime when including also the demand instrument. Notably, the standard errors for the two coefficients become tighter than in the previous estimation, indicating that adding the control for the time-varying frequency strengthens the inference. Nevertheless, we can still reject the null hypothesis of constant coefficients at a 10% confidence level.

To sum up, from the quantitative exercise of this section, we conclude that state-dependent information frictions can be detected in the data and the evidence survives even when accounting for the fluctuations in the frequency of price changes. In high-inflation periods, firms behave as if they were more “forward-looking,” responding more to idiosyncratic cost shocks. We now draw the aggregate implications of these findings.

6 Aggregate Implications

This section draws the aggregate implications of our findings of state-dependent information frictions. First, we derive the New Keynesian Phillips curve for the setting. We then illustrate how incomplete information leads to high discounting, nonlinear cost-price pass-through, and a correlation between higher-order moments of the price distribution and aggregate inflation. Finally, we show that our model can account for a larger share of inflation volatility compared to the full-information benchmark.

6.1 Nonlinear Phillips Curve under Information Frictions

Let us now derive the New Keynesian Phillips curve for the framework. As standard, the NKPC is obtained by aggregating the first-order conditions of the firms (equation 8) and applying the law of large numbers. The next propositions and corollaries clarify how the presence of information frictions leads to a nonlinear amplification of the response of inflation to changes in aggregate real marginal cost. Denote by $\pi_t := \int_{[0,1]} \int_{\mathcal{F}_i} \Delta p_{fit} df \; di$ the aggregate inflation rate, $mc_t^r := \int_{[0,1]} \int_{\mathcal{F}_i} mc_{fit}^r df \; di$ the average real marginal cost, and by $\beta_t := \int_{[0,1]} \beta_{it} di$ the average Kalman gain. We assume, consistently with the empirical evidence (Lenzu et al. 2023), that short-term average returns to scale are constant.

Proposition 3 (NKPC)**.** *The New Keynesian Phillips curve under information frictions approximately satisfies:*

$$\pi_t = \lambda(\mu + mc_t^r) + \mathbb{E}_t(\pi_{t+1}) + \mathrm{Cov}(\beta_{it}, \pi_{it+1}) + \epsilon_{t+1} \tag{23}$$

where μ is the steady-state markup, the slope of the Phillips curve is:

$$\lambda := \frac{(1-\theta)^2}{\theta}(1-\Omega)$$

The aggregate "cost-push shock:"

$$\epsilon_{t+1} := \beta_t(\pi_{t+1} - \mathbb{E}_t(\pi_{t+1})),$$

is a conditionally mean-zero Gaussian expectational error. Cov *denotes the cross-sectional covariance across industries.*

The NKPC in equation (23) is a direct generalization of the Phillips curve obtained in standard models with Calvo frictions and follows immediately from aggregating the Dynamic Pass-through Regression. Inflation is an increasing function of the aggregate real marginal cost, in deviation from its steady state value. Under the assumption of constant aggregate returns to scale, the slope of the cost-based NKPC corresponds to the pass-through of contemporaneous cost shocks as estimated in section 4.[28] Due to the presence of nominal rigidities, inflation responds to expectations of future inflation with an — approximately— unitary coefficient. In the full-information benchmark ($\beta_t \to 1, \mathrm{Cov} \to 0$), the expectation of future inflation is discounted at an exponential rate of $\delta \approx 1$. For expectations to decay so rapidly, observation of the price level must reveal all the uncertainty relevant to the forecasting problem. Shocks are then unforecastable, i.e. orthogonal to the

[28] As discussed in Gagliardone et al. (2023), accounting for an empirically reasonable departure from constant aggregate returns reduces the estimates of the slope by about 6% at quarterly frequency, which is less than half a percentage point.

information set of the firm.[29] A useful analogy is to think about the exponential discounting of expectations as information becoming rapidly obsolete, leading to low persistence of expectations over time.

However, we documented evidence that the pass-through from expectations to prices has a steep drop at short horizons, consistently with a hyperbolic form of discounting (section 5.1). As the hyperbolic decay is slower than exponential, information persists over time leading to slow fluctuations of expectations. In the data, we indeed detect correlations over time between prices and beliefs, as well as persistence of expectations. Through the lenses of the model, the covariance in equation (23) incorporates these comovements between realizations of shocks and information of firms as mediated by the Kalman gain, which parametrizes the hyperbolic decay.

If the Kalman gain is state-dependent, large shocks increase β_{it}. This generates inflation volatility because industries with large shocks are also mostly reactive. Indeed, when averaging across industries, the covariance term has the same sign of the mean of the shocks, i.e. the same sign of inflation, contributing positively to inflation volatility.

Let us now formalize these insights.

High discounting. We discussed in section 2.5 how the presence of information frictions implies a form of discounting of cost shocks at the firm level. In turn, when aggregating across firms, the high micro-discounting of shocks leads to a high macro-discounting of expectations. The next corollary formalizes the connection between information frictions and discounting at the aggregate level.

Corollary 1. *Suppose that the covariance varies over time. Then β_t is bounded away from one.*

The proof of the corollary in Appendix A.3 shows that the discounting can be tightly bounded above by the nonlinear effect of an unanticipated nominal cost shock $c_{t+1} \notin \mathcal{I}_t$:

$$\beta_t \leq 1 - \frac{\partial \text{Cov}(\beta_{it}, \pi_{it+1})}{\partial c_{t+1}} < 1.$$

We argued in the previous section that β_t is high when shocks are large, i.e. the covariance increases in response to shocks. As the pass-through of any nominal shock is at most complete, movements in the covariance lead to a loose upper bound. Tightness of the bound follows from showing that the covariance is either constant at zero or strictly positive (except with a zero shock). This leads to the result that, as the covariance varies over time, the

[29] For the cost-push shock to become orthogonal to the information set under complete information, the prior belief must converge to a improper uniform distribution. Expectations of the error term are nevertheless well-defined because β_t and $\pi_{t+1} - \mathbb{E}_t(\pi_{t+1})$ become uncorrelated, leading to a zero average.

discount rate must be strictly smaller than one. Hence, under the information friction, the pass-through of future costs has a slower decay than exponential.[30]

To conclude, the effect on contemporaneous inflation of a nominal unanticipated shock that raises inflation tomorrow by one percentage point —to a first order— is equal to β_t%, which is generally smaller than one percent. Hence, the response of inflation to such news is discounted. Similarly, the effect of an anticipated shock is given by $(1-\beta_t)$%, which is also strictly smaller than one percent. Therefore both types of shocks do not lead to large deviations of inflation from trend to the extent that nonlinear effects are also small.

Nonlinear pass-through. We now impose additional structure on beliefs to deliver a closed-form characterization of the nonlinear term that provides intuition and can then be tested with data in the next section. The practical purpose of the assumption is to link beliefs, i.e. the prior variance $\sigma^2_{e,it} := \mathbb{V}_{it}(\Delta p_{fit+1|t})$, to the realized dispersion $\mathbb{V}(\Delta p_{fit+1})$ in a time-invariant manner. Formally, we require that the elasticity of the prior variance with respect to the realized dispersion conditionally on a nominal shock is constant over time and bounded, as follows:

$$\zeta \equiv \zeta_{it} := \frac{\partial \sigma^2_{e,it}}{\partial \mathbb{V}(\Delta p_{fit+1})} \in [0,1].$$

This assumption can be interpreted as firms forming beliefs in a stationary environment, with ζ parametrizing the extent to which they anticipate future shocks and incorporate that knowledge into beliefs. In particular, if $\zeta = 0$, beliefs are uncorrelated with ex-post realizations. We will show shortly that this "lower bound" must bind with full information, implying that the covariance in equation (23) is zero. On the other hand, if $\zeta > 0$, beliefs predict realizations, that is the Kalman gain increases *in anticipation* of shocks. In turn, firms' response increases with the size of shocks, due to the pass-through into beliefs via the instantaneous change in β_{it}. If $\zeta = 1$, beliefs perfectly comove with realizations, that is the Kalman gain increases one-to-one in anticipation of shocks. In turn, firms are mostly responsive leading to an "upper bound" for inflation volatility. In the quantitative analysis of the next section, we provide a comparison of the two benchmarks ($\zeta \in \{0,1\}$).

The next proposition characterizes the resulting covariance in terms of observable moments of the distribution of price changes.

Proposition 4 (Nonlinearities of the NKPC). *Denote by $\sigma^2_t, \gamma_t < \infty$ the variance and skewness of the cross-sectional distribution F of industry inflation rates:*

$$\pi_{it+1} \sim F(\pi_{t+1}, \sigma^2_t, \gamma_t).$$

[30] Such a feature of "high discounting" in macroeconomic models has been discussed to be desirable (Gabaix 2020) to resolve several puzzles in the literature, including the forward-guidance puzzle.

Up to a second-order approximation of the Kalman gain around the symmetric steady state, the covariance simplifies to:

$$\mathrm{Cov}(\beta_{it}, \pi_{it+1}) = \zeta \frac{\theta}{1-\theta} \frac{\sigma_t^2}{\sigma_\eta^2} \left(\pi_{t+1} + \frac{1}{2} \gamma_t \sigma_t \right). \tag{24}$$

where σ_t^2/σ_η^2 measures the aggregate signal-to-noise ratio.

Proposition 4 provides a characterization of the nonlinear terms of the NKPC in equation (23). When the elasticity ζ is zero, the covariance term is zero. In particular, this must be the case provided that inflation has finite moments in the full-information benchmark, which is characterized by a constant Kalman gain ($\beta_{it} \to 1$) and improper prior ($\sigma_{e,it}^2 \to \infty$). With a positive elasticity, the variance of price changes comoves with inflation (equation 20) at rate $\theta/(1-\theta)$. Hence, the nonlinear effects can be detected in the data only at a high frequency ($\theta > 0$). The aggregate "signal-to-noise ratio" also shows up in the covariance as a multiplicative term, scaling endogenously up or down the response of the covariance to shocks that affect inflation (π_{t+1}) and the skewness (γ_t).

The next corollary clarifies how the covariance leads to a nonlinear pass-through even when shocks affect only the first moment of the distribution.

Corollary 2. *Suppose that $\gamma_t = 0$. The pass-through of an unanticipated nominal cost shock $c_{t+1} \notin \mathcal{I}_t$ is nonlinear and given by:*

$$\frac{\partial \pi_t}{\partial c_{t+1}} = \left(\beta_t + \zeta \frac{\theta}{1-\theta} \frac{\sigma_t^2}{\sigma_\eta^2} \right) \cdot \frac{\partial \pi_{t+1}}{\partial c_{t+1}}.$$

In a neighborhood of the symmetric steady state, σ_t^2 is small, hence the first-order effect (β_t) dominates over the higher-order effect (σ_t^2/σ_η^2). Away from the steady state, the nonlinear term leads to amplification of the effect of shocks.

Effects of heterogeneity. Finally, we focus now on the role that the skewness of the distribution of price changes plays in equation (24). The skewness is positive when the distribution of industry inflation rates displays a long right tail, i.e. few industries have large positive shocks. On the other hand, when a large share of industries experiences high inflation, the skewness is negative. If shocks are independently distributed across industries and follow the same distribution, aggregate inflation is symmetrically distributed and there is no skewness. Away from this benchmark, asymmetric industry shocks induce correlations between the first and third moments of the distribution of price changes.

In Figure 5 we report the time series for annual inflation and the skewness of beliefs, which we constructed from the survey measure of expected price changes ($\mathbb{E}_{fit-1}(\mathrm{sign}(\Delta p_{fit}))$)

Figure 5: Evidence for the role of higher-order moments of the belief distribution.

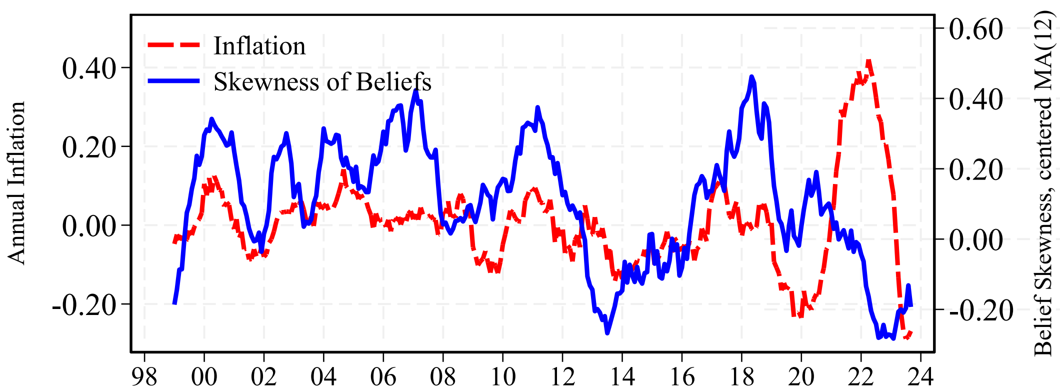

Notes. Time series of annual aggregate inflation and aggregate skewness of beliefs. All variables are aggregated using the Törnqvist weights. Inflation is detrended by removing from the aggregate price level a piece-wise linear trend with one break. The skewness is smoothed using a centered 12-month moving average.

and smoothed similarly to Figure 4. As the plot suggests, firms anticipate few extreme realizations of industry inflation when inflation is stable, leading to a positive skewness. In particular, this is the case for the periods 1999-2008, 2010-2011, and 2016-2018 in which the economy was arguably in normal times. The skewness decreases mildly during the 2009-2010 recession relative to the previous period and became negative during the sovereign debt crisis (2012-2014) and the inflation surge in 2021-2023, which can be arguably considered moments of high price instability. Through the lenses of the model, the decline in the skewness in high inflation regimes operates as a stabilizing force for inflation, preventing it from drifting further away from trend.

6.2 How much Inflation Volatility can the Model Explain?

Combining propositions (3) and (4), we now derive testable implications for the model at the aggregate level. With symmetric shocks ($\gamma_t = 0$), the New Keynesian Phillips curve with state-dependent information frictions simplifies to:

$$\pi_t = \lambda(\mu + mc_t^r) + \mathbb{E}_t(\pi_{t+1}) + \zeta \frac{\theta}{1-\theta} \frac{\sigma_t^2}{\sigma_\eta^2} \pi_{t+1} + \epsilon_{t+1}.$$

Taking expectations of both sides with respect to the average prior belief $\mathcal{I}_t$ and treating σ_t^2/σ_η^2 as predetermined (or a time-varying parameter), we obtain a difference equation for the inflation rate predicted by the model (π_t^e) as a function of the aggregate real marginal cost:[31]

$$\pi_t^e = \lambda(\mu + mc_t^r) + \left(1 + \zeta \frac{\theta}{1-\theta} \frac{\sigma_t^2}{\sigma_\eta^2}\right) \mathbb{E}_t(\pi_{t+1}^e), \tag{25}$$

where the Kalman gain β_t does not show up directly because ϵ_{t+1} is mean zero conditionally on $\mathcal{I}_t$. Nevertheless, the presence of the information friction amplifies the pass-through of shocks via the nonlinear term $\zeta \frac{\theta}{1-\theta} \sigma_t^2 / \sigma_\eta^2$.

The solution of this difference equation leads to a relation between current inflation and the real marginal cost with a time-varying coefficient. In turn, this relation pins down the share of volatility that the model can generate for any given path of real marginal cost. The next proposition derives the solution of equation (25) obtained using the method of undetermined coefficients.

Proposition 5 (Inflation Dynamics)**.** *Assume that the distribution of inflation rates is symmetric ($\gamma_t = 0$). The expected dynamics of aggregate inflation (π_t^e) as a function of the innovation to the aggregate nominal marginal cost ($\varepsilon_{n,t}$) is given by:*

$$\pi_t^e(\varepsilon_{n,t}) = \Psi_t \cdot (mc_{t-1}^r + \varepsilon_{n,t}), \tag{26}$$

where $\Psi_t \equiv \Psi(\varepsilon_{n,t}^2)$ is an increasing function of the size of the shock.

Equation (26) is a "reduced-form NKPC," with a time-varying coefficient that reflects the state dependence of the pass-through. In particular, Ψ_t is strictly increasing in $\varepsilon_{n,t}^2$ only if $\zeta > 0$ is strictly positive. This means that the reduced-form pass-through displays state dependence only if beliefs —to some extent— predict realizations. In particular, when $\zeta = 1$, Ψ_t provides an upper bound for the response of prices to cost shocks. The lower bound is obtained by setting $\zeta = 0$, i.e. by imposing that the covariance in the NKPC is zero (equation 23).

Proposition 5 characterizes how much inflation volatility can be generated by the model, depending on the calibration of parameters and assumptions on beliefs. To test the performance of the model at business-cycle frequency, we use equation (26) and construct the year-over-year aggregate inflation that the model predicts as a function of the realized path of aggregate marginal cost during the year. Following Gagliardone et al. (2023) and Gagliardone et al. (2024), we measure the latter as the Törnqvist-weighted average of firm-level marginal costs, which mirrors the construction of the aggregate price index from firm-level prices. Finally, we detrend both the price level (as before) and the aggregate nominal marginal cost removing a linear trend. Parameters are calibrated at our baseline

[31] It is indeed possible to treat σ_t^2/σ_η^2 as a predetermined variable because it shows up only via the prior variance ($\sigma_{e,it}^2$), which is predetermined, and the elasticity ζ, which is non-stochastic by our assumption on beliefs.

Figure 6: Inflation (year-over-year) in data and model

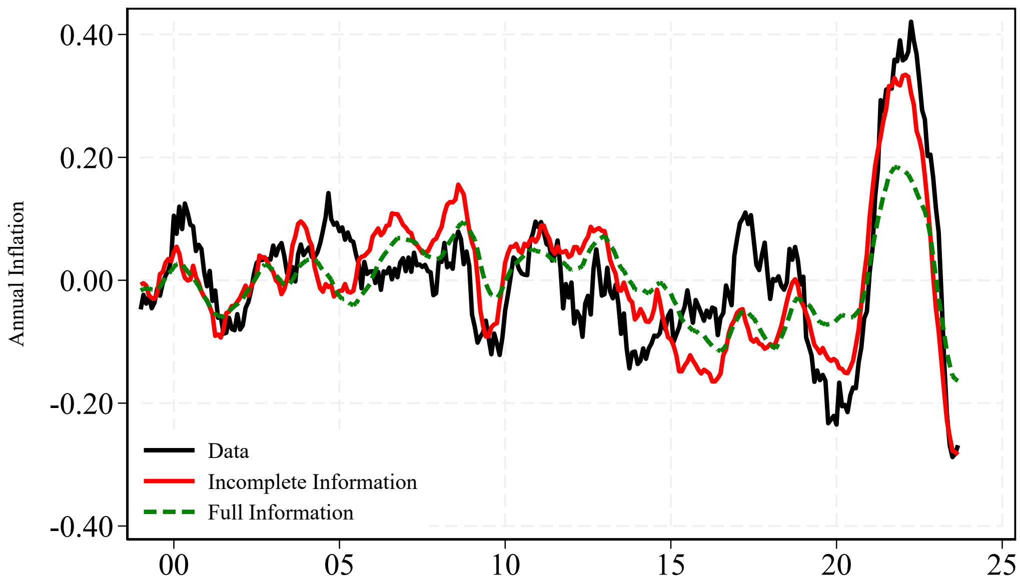

Notes. Time series of annual aggregate inflation and model-based inflation. All variables are aggregated using Törnqvist weights. Inflation is detrended by removing from the aggregate price level a piece-wise linear trend with one break. The red full line is the model-based equation for the nonlinear model with incomplete information, obtained from equation (26) for the case $\zeta = 1$. The green dashed line is the model-based equation for the linear model with complete information, obtained from equation (26) for the case $\zeta = 0$.

estimates (Column A of Table 4).[32]

Figure 6 reports the results of this exercise. The black line is year-over-year inflation in the data. We then feed the aggregate path of marginal cost and the dispersion of industry inflation rates into the model via equation (26). We obtain the model-based inflation, which is plotted in red. As the figure shows, the model tracks remarkably well the time series of inflation considering that we are only feeding aggregate information on costs and dispersion across industries. In detail, regressing data on model (black on red), we obtain an $R^2 = 0.63$. In turn, the nonlinear model can account for about two-thirds of the overall volatility of inflation. This number provides an upper bound for the share of the volatility of inflation that can be explained via our model of state-dependent information frictions.

To illustrate the role of nonlinearities, we compute the lower bound which corresponds to the full-information benchmark. In this case, the response of inflation to shocks is linear in the cost shocks. This lower bound is plotted as the green dashed line. We notice that the model is still correlated with data, due to the path of marginal cost being corre-

[32] The variance of the noise (σ_η^2) is calibrated to match the estimate of the Kalman gain (β) at the average value of the dispersion of industry inflation rates (σ_t^2).

lated with inflation. However, the full-information benchmark displays less volatility than the incomplete-information model, and therefore less volatility than the data. Regressing the black line on the green, we find an $R^2 = 0.5$, showing that indeed the model with state-dependent information frictions can rationalize a larger share of inflation volatility.

7 Conclusions

In this paper, we discussed theory and evidence on the dynamic pricing behavior of firms. We argued that the data calls for a model that includes real, nominal, and information frictions in order to match the intertemporal co-movements of prices, costs, and expectations.

We provided three novel insights. First, firms apply a high discounting to cost shocks at the micro-level. When aggregating across firms, the high discounting leads to a decay of inflation expectations which is slower than exponential. New information is only slowly incorporated into beliefs and hence expectations persist over time. In turn, unanticipated cost shocks have a smaller impact on inflation compared to a full-information benchmark.

Second, the speed at which beliefs are updated, i.e. the Kalman gain, is state-dependent. This coefficient increases in response to nominal cost shocks and increases more with large shocks. In turn, in industries that are hit by large nominal disturbances, firms behave as if they had approximately complete information, updating their beliefs rapidly in response to news. The quicker revision of expectations leads to a stronger contemporaneous reaction to shocks, which scales up with the magnitude of the disturbance. Therefore, the cost-price pass-through is nonlinear and increases with the size of shocks.

Third, high discounting and nonlinearities imply a correlation between the first and higher-order moments of the price distribution. In particular, the response of inflation to shocks increases with the dispersion of inflation rates across industries. Hence heterogeneity matters because it affects the elasticity of inflation to nominal disturbances.

We conclude from the analysis that the state dependence of the information friction serves as an amplifying mechanism of the propagation of cost shocks. Specifically, the resulting inflation volatility is greater than in an alternative model where information is complete or the Kalman gain remains fixed. As a result, our model of state-dependent information frictions demonstrates more monetary neutrality compared to the full-information benchmark. This reveals a channel that operates in the opposite direction of what is typically understood to be the aggregate effect of incomplete information.

References

Hassan Afrouzi, Joel P Flynn, and Choongryul Yang. What can measured beliefs tell us about monetary non-neutrality? Technical report, National Bureau of Economic Research, 2024.

Fernando Alvarez and Francesco Lippi. The analytic theory of a monetary shock. *Econometrica*, 90(4):1655–1680, 2022.

Mary Amiti, Oleg Itskhoki, and Jozef Konings. International shocks, variable markups, and domestic prices. *The Review of Economic Studies*, 86(6):2356–2402, 2019.

Philippe Andrade, Olivier Coibion, Erwan Gautier, and Yuriy Gorodnichenko. No firm is an island? how industry conditions shape firms' expectations. *Journal of Monetary Economics*, 125:40–56, 2022.

George-Marios Angeletos and Zhen Huo. Myopia and anchoring. *American Economic Review*, 111(4):1166–1200, 2021.

George-Marios Angeletos and Jennifer La'O. Incomplete information, higher-order beliefs and price inertia. *Journal of Monetary Economics*, 56:S19–S37, 2009.

George-Marios Angeletos and Jennifer La'O. Optimal monetary policy with informational frictions. *Journal of Political Economy*, 128(3):1027–1064, 2020.

George-Marios Angeletos and Chen Lian. Forward guidance without common knowledge. *American Economic Review*, 108(9):2477–2512, 2018.

George-Marios Angeletos, Zhen Huo, and Karthik A Sastry. Imperfect macroeconomic expectations: Evidence and theory. *NBER Macroeconomics Annual*, 35(1):1–86, 2021.

Andrew Atkeson and Ariel Burstein. Pricing-to-market, trade costs, and international relative prices. *American Economic Review*, 98(5):1998–2031, 2008.

Laurence Ball and David Romer. Real rigidities and the non-neutrality of money. *The Review of Economic Studies*, 57(2):183–203, 1990.

Paul Beaudry, Chenyu Hou, and Franck Portier. The dominant role of expectations and broad-based supply shocks in driving inflation. Technical report, National Bureau of Economic Research, 2024.

Andres Blanco, Corina Boar, Callum Jones, and Virgiliu Midrigan. Nonlinear inflation dynamics in menu cost economies. *Manuscript, December*, 2022.

Andres Blanco, Corina Boar, Callum J Jones, and Virgiliu Midrigan. Non-linear inflation dynamics in menu cost economies. Technical report, National Bureau of Economic Research, 2024a.

Andrés Blanco, Corina Boar, Callum J Jones, and Virgiliu Midrigan. Nonlinear dynamics in menu cost economies? evidence from us data. Technical report, National Bureau of Economic Research, 2024b.

Anat Bracha and Jenny Tang. Inflation levels and (in) attention. *Review of Economic Studies*, 2024.

Guillermo A Calvo. Staggered prices in a utility-maximizing framework. *Journal of monetary Economics*, 12(3):383–398, 1983.

Bernardo Candia, Olivier Coibion, and Yuriy Gorodnichenko. The inflation expectations of us firms: Evidence from a new survey. *Journal of Monetary Economics*, page 103569, 2024.

Alberto Cavallo, Francesco Lippi, and Ken Miyahara. Large shocks travel fast. Technical report, National Bureau of Economic Research, 2023.

Olivier Coibion, Yuriy Gorodnichenko, and Rupal Kamdar. The formation of expectations, inflation, and the phillips curve. *Journal of Economic Literature*, 56(4):1447–1491, 2018a.

Olivier Coibion, Yuriy Gorodnichenko, and Saten Kumar. How do firms form their expectations? new survey evidence. *American Economic Review*, 108(9):2671–2713, 2018b.

Olivier Coibion, Yuriy Gorodnichenko, and Tiziano Ropele. Inflation expectations and firm decisions: New causal evidence. *The Quarterly Journal of Economics*, 135(1):165–219, 2020.

Isabelle De Greef and Christophe Van Nieuwenhuyze. The national bank of belgium's new business survey indicator. *NBB Economic Review*, June:31–51, 2009.

Michael Dotsey and Robert G King. Implications of state-dependent pricing for dynamic macroeconomic models. *Journal of Monetary Economics*, 52(1):213–242, 2005.

Andres Drenik and Diego J Perez. Price setting under uncertainty about inflation. *Journal of Monetary Economics*, 116:23–38, 2020.

Martin Eichenbaum, Nir Jaimovich, Sergio Rebelo, and Josephine Smith. How frequent are small price changes? *American Economic Journal: Macroeconomics*, 6(2):137–155, 2014.

European Commission. Quality report on prodcom 2014 annual data. PRODCOM Working Group, 2014.

Emmanuel Farhi and Iván Werning. Monetary policy, bounded rationality, and incomplete markets. *American Economic Review*, 109(11):3887–3928, 2019.

Terry Fitzgerald, Callum Jones, Mariano Kulish, and Juan Pablo Nicolini. Is there a stable relationship between unemployment and future inflation? *American Economic Journal: Macroeconomics*, 16(4):114–142, 2024.

Xavier Gabaix. A behavioral new keynesian model. *American Economic Review*, 110(8): 2271–2327, 2020.

Xavier Gabaix and David Laibson. Myopia and discounting. Technical report, National bureau of economic research, 2017.

Luca Gagliardone, Mark Gertler, Simone Lenzu, and Joris Tielens. Anatomy of the phillips curve: Micro evidence and macro implications. Technical report, National Bureau of Economic Research, 2023.

Luca Gagliardone, Mark Gertler, Simone Lenzu, and Joris Tielens. Micro and macro cost-price dynamics across inflation regimes. Technical report, 2024.

Mariana García-Schmidt and Michael Woodford. Are low interest rates deflationary? a paradox of perfect-foresight analysis. *American Economic Review*, 109(1):86–120, 2019.

Mikhail Golosov and Robert E Lucas. Menu costs and phillips curves. *Journal of Political Economy*, 115(2):171–199, 2007.

Lars Peter Hansen. Generalized method of moments estimation. *Macroeconometrics and time series analysis*, pages 105–118, 2010.

Christopher Harris and David Laibson. Dynamic choices of hyperbolic consumers. *Econometrica*, 69(4):935–957, 2001.

Jonathon Hazell, Juan Herreno, Emi Nakamura, and Jón Steinsson. The slope of the phillips curve: evidence from us states. *The Quarterly Journal of Economics*, 137(3):1299–1344, 2022.

Christian Hellwig and Venky Venkateswaran. Dispersed information, sticky prices and monetary business cycles: A hayekian perspective. *New York University, Mimeo*, 2014.

Miles S Kimball. The quantitative analytics of the basic neomonetarist model. *Journal of Money, Credit, and Banking*, 27(4):1241–1277, 1995.

Saten Kumar, Yuriy Gorodnichenko, and Olivier Coibion. The effect of macroeconomic uncertainty on firm decisions. *Econometrica*, 91(4):1297–1332, 2023.

David Laibson. Self-control and saving. *Massachusets Institute of Technology mimeo*, 1994.

Simone Lenzu, David Rivers, and Joris Tielens. Financial shocks, productivity, and prices. *Working paper*, 2023.

Jr Robert E Lucas. Expectations and the neutrality of money. *Journal of economic theory*, 4 (2):103–124, 1972.

Bartosz Maćkowiak and Mirko Wiederholt. Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803, 2009.

Bartosz Maćkowiak, Filip Matějka, and Mirko Wiederholt. Rational inattention: A review. *Journal of Economic Literature*, 61(1):226–273, 2023.

N Gregory Mankiw and Ricardo Reis. Sticky information versus sticky prices: a proposal to replace the new keynesian phillips curve. *The Quarterly Journal of Economics*, 117(4): 1295–1328, 2002.

N Gregory Mankiw, Ricardo Reis, and Justin Wolfers. Disagreement about inflation expectations. *NBER macroeconomics annual*, 18:209–248, 2003.

Michael McLeay and Silvana Tenreyro. Optimal inflation and the identification of the phillips curve. *NBER Macroeconomics Annual*, 34(1):199–255, 2020.

Leonardo Melosi. Estimating models with dispersed information. *American Economic Journal: Macroeconomics*, 6(1):1–31, 2014.

John F Muth. Rational expectations and the theory of price movements. *Econometrica: journal of the Econometric Society*, pages 315–335, 1961.

Emi Nakamura and Jón Steinsson. Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics*, 123(4):1415–1464, 2008.

Stephen Nickell. Biases in dynamic models with fixed effects. *Econometrica*, 49(6):1417–1426, 1981.

Kristoffer Nimark. Dynamic pricing and imperfect common knowledge. *Journal of monetary Economics*, 55(2):365–382, 2008.

Luigi Paciello and Mirko Wiederholt. Exogenous information, endogenous information, and optimal monetary policy. *Review of Economic Studies*, 81(1):356–388, 2014.

Oliver Pfäuti. The inflation attention threshold and inflation surges. *arXiv preprint arXiv:2308.09480*, 2023.

Julio J Rotemberg. Sticky prices in the united states. *Journal of political economy*, 90(6):1187–1211, 1982.

Francisca Sara-Zaror. *Expected inflation and welfare: The role of consumer search*. SSRN, 2022.

Christopher A Sims. Implications of rational inattention. *Journal of monetary Economics*, 50(3):665–690, 2003.

Xavier Vives. *Information and learning in markets: the impact of market microstructure*. Princeton University Press, 2010.

Olivier Wang and Iván Werning. Dynamic oligopoly and price stickiness. *American Economic Review*, 112(8):2815–49, 2022.

Michael Weber, Bernardo Candia, Hassan Afrouzi, Tiziano Ropele, Rodrigo Lluberas, Serafin Frache, Brent H Meyer, Saten Kumar, Yuriy Gorodnichenko, Dimitris Georgarakos, et al. Tell me something i don't already know: Learning in low and high-inflation settings. Technical report, National Bureau of Economic Research, 2023.

Michael Woodford. Imperfect common knowledge and the effects of monetary policy, 2001.

Dynamic Pricing under Information Frictions: Evidence from Firm-level Subjective Expectations

L. Gagliardone J. Tielens

Appendix

A Derivations

A.1 Derivation of markup function

Dynamic oligopoly with nested CES preferences

Assume that there is a continuum of industries (indexed by i) and a finite number of firms N within each industry. Each firm is indexed by f (or j). Within each industry, firms compete à la Bertrand. In this environment, the price indexes for each industry P_{it} and the aggregate price index P_t are defined, respectively, as:

$$P_{it} := \left(\frac{1}{N} \sum_{f=1}^{N} P_{fit}^{1-\gamma} \right)^{\frac{1}{1-\gamma}} \; ; \; P_t := \left(\int_{i \in I} P_{it}^{1-\sigma} di \right)^{\frac{1}{1-\sigma}},$$

Log-linearization of the price indexes leads to the approximate Cobb-Douglas price index as in the text. The demand function for firm $f \in \mathcal{F}_i$ takes a nested CES form, with the elasticity of substitution across industries $\sigma > 1$ and the elasticity of substitution within industries $\gamma > \sigma$:

$$\mathcal{D}_{fit} = \left(\frac{P_{fit}^o}{P_{it}} \right)^{-\gamma} \left(\frac{P_{it}}{P_t} \right)^{-\sigma} Y_t. \tag{A.1}$$

Firms internalize the dynamic effect of their choices on the industry price index and on industry demand. Therefore, the residual elasticity of demand faced by firm f takes the following form:

$$\epsilon_{fit} := -\frac{\partial \ln \mathcal{D}_{fit}}{\partial \ln P_{fit}^o} = \gamma - (\gamma - \sigma) \frac{\partial p_{it}}{\partial p_{fit}^o}. \tag{A.2}$$

We can further characterize the derivative above. First, the price index of competitors of firm f is defined as:

$$P_{it}^{-f} := \left(\frac{1}{N-1} \sum_{j \neq f}^{N-1} P_{jit}^{1-\gamma} \right)^{\frac{1}{1-\gamma}}.$$

It follows that $P_{it}^{1-\gamma} = \frac{N-1}{N}\left(P_{it}^{-f}\right)^{1-\gamma} + \frac{1}{N}\left(P_{fit}^{o}\right)^{1-\gamma}$. Next, we can express the derivative of the price index in period t with respect to the firms' reset price in period t as follows:

$$\frac{\partial P_{it}}{\partial P_{fit}^{o}} = P_{it}^{\gamma}\left[\left(\frac{N-1}{N}\right)(P_{it}^{-f})^{-\gamma}\frac{\partial P_{it}^{-f}}{\partial P_{fit}^{o}} + \left(\frac{1}{N}\right)(P_{fit}^{o})^{-\gamma}\right].$$

Multiplying both sides by $\frac{P_{fit}^{o}}{P_{it}}$, and defining the competitors' reaction function $\zeta_{fit} := \frac{\partial p_{it}^{-f}}{\partial p_{fit}^{o}}$, we obtain:

$$\begin{aligned}\frac{\partial p_{it}}{\partial p_{fit}^{o}} &= \zeta_{fit}\left(\frac{N-1}{N}\right)\left(\frac{P_{it}^{-f}}{P_{it}}\right)^{1-\gamma} + \frac{1}{N}\left(\frac{P_{fit}^{o}}{P_{it}}\right)^{1-\gamma} \\ &= \zeta_{fit}(1-s_{fit}) + s_{fit},\end{aligned}$$

where $s_{fit} := \frac{1}{N}\frac{P_{fit}^{o}\mathcal{D}_{fit}}{P_{it}Y_{it}} = \frac{1}{N}\left(\frac{P_{fit}^{o}}{P_{it}}\right)^{1-\gamma}$ denotes the within-industry revenue share of firm f, and $Y_{it} := \left(\frac{P_{it}}{P_t}\right)^{-\sigma} Y_t$ is the industry demand. Replacing the expression for $\frac{\partial p_{it}}{\partial p_{fit}^{o}}$ into equation (A.2), we find that the within-industry elasticity of demand faced by firm f is given by:

$$\epsilon_{fit} = \gamma - (\gamma-\sigma)\left[\zeta_{fit}(1-s_{fit}) + s_{fit}\right]. \tag{A.3}$$

The intuition behind this expression is straightforward. The stronger the reaction of competitors to a firm's price change—-captured by ζ_{fit}—-the lower the residual elasticity of demand. A low residual elasticity of demand, in turn, implies that the firm can sustain a higher markup in equilibrium. This result mirrors the one in the dynamic oligopoly environment in Wang and Werning (2022) and it nests a number of static environments featuring imperfectly competitive firms. In Atkeson and Burstein (2008) static Nash oligopoly, $\epsilon_{fit} = 0$ and $\zeta_{fit} = 0$. Under monopolistic competition, $N \to \infty$, which implies $\zeta_{fit} \to 0$ and $s_{fit} \to 0$.

We now use this result to derive the expression for the log-linearized desired markup in equation (2) in the paper. We log-linearize around a symmetric Nash steady state. Log-linearizing the elasticity in (A.3) around the steady state, we obtain the steady state residual demand elasticity:

$$\epsilon = \gamma - (\gamma-\sigma)\frac{1}{N},$$

which corresponds to the expression in Atkeson and Burstein (2008). In this model, the desired markup is given by the Lerner index $\mu_{fit} := \ln(\epsilon_{fit}/(\epsilon_{fit}-1))$. Log-linearizing this expression and substituting the expression for steady-state residual demand elasticity, we obtain the expression for the log-linearized desired markup (in deviation from steady state) in equation (2.1):

$$\mu_{fit} - \mu = -\Gamma\left(p^{o}_{fit} - p^{-f}_{it}\right) + u^{\mu}_{fit},$$

where $\Gamma := \frac{(\gamma-\sigma)(\gamma-1)}{\epsilon(\epsilon-1)}\frac{N-1}{N} > 0$ denotes the markup elasticity with respect to prices, and:

$$u^{\mu}_{fit} := \frac{\gamma-\sigma}{\epsilon(\epsilon-1)}\frac{N-1}{N}\zeta_{fit}, \tag{A.4}$$

captures residual variation in the markup that depends on the changes in the slope of competitors' reaction function. Finally, for sufficiently small deviations from steady state, $\zeta_{fit} \approx 0$ implies the result. See Gagliardone et al. (2023) for a discussion of the role of $\zeta_{fit+\tau}$ when $\tau > 0$ in the identification arguments.

Monopolistic competition with Kimball preferences

Assume that the industry output Y_{it} is produced by a unitary measure of perfectly competitive firms using a bundle of differentiated intermediate inputs Y_{fit}, $f \in i$. The bundle of inputs is assembled into final goods using the Kimball aggregator:

$$\int_0^1 \Upsilon\left(\frac{Y_{fit}}{Y_{it}}\right) df = 1,$$

where $\Upsilon(\cdot)$ is strictly increasing, strictly concave, and satisfies $\Upsilon(1) = 1$.

Taking as given the industry demand Y_{it}, each firm minimizes costs subject to the aggregate constraint:

$$\min_{Y_{fit}} \int_0^1 P_{fit} Y_{fit} df \qquad \text{s.t.} \int_0^1 \Upsilon\left(\frac{Y_{fit}}{Y_{it}}\right) df = 1.$$

Denoting by ψ the Lagrange multiplier of the constraint, the first-order condition of the problem is:

$$P_{fit} = \psi \Upsilon'\left(\frac{Y_{fit}}{Y_{it}}\right)\frac{1}{Y_{it}} \tag{A.5}$$

Define implicitly the industry price index P_{it} as:

$$\int_0^1 \beta\left(\Upsilon'(1)\frac{P_{fit}}{P_{it}}\right) df = 1$$

where $\beta := \Upsilon \circ (\Upsilon')^{-1}$. Evaluating the first-order condition (A.5) at symmetric prices, $P_{fit} =$

P_{it}, we get $\psi = \frac{P_{it}Y_{it}}{\Upsilon'(1)}$. Replacing for ψ, we recover the demand function:

$$\frac{P_{fit}}{P_{it}} = \frac{1}{\Upsilon'(1)}\Upsilon'\left(\frac{Y_{fit}}{Y_{it}}\right). \tag{A.6}$$

Therefore, the demand function faced by firms when resetting prices is:

$$\mathcal{D}_{fit} = \left[(\Upsilon')^{-1}\left(\Upsilon'(1)\frac{P^o_{fit}}{P_{it}}\right)\right]\left(\frac{P_{it}}{P_t}\right)^{-\sigma} Y_t$$

Taking logs of equation (A.1) and differentiating, we obtain the following expression for the residual elasticity of demand:

$$\epsilon_{fit} := -\frac{\partial \ln \mathcal{D}_{fit}}{\partial \ln P^o_{fit}} = -\frac{\Upsilon'\left(\frac{Y_{fit}}{Y_{it}}\right)}{\Upsilon''\left(\frac{Y_{fit}}{Y_{it}}\right)\cdot\left(\frac{Y_{fit}}{Y_{it}}\right)} \tag{A.7}$$

We now use this result to derive the expression for the log-linearized desired markup in equation (2.1) in the paper, under monopolistic competition with Kimball preferences. As before, we focus on small departures from the symmetric steady state. Denote the steady-state residual demand elasticity by $\epsilon = -\frac{\Upsilon'(1)}{\Upsilon''(1)}$. Then the derivative of the residual demand elasticity ϵ_{fit} in (A.7) with respect to $\frac{Y_{fit}}{Y_{it}}$, evaluated at the steady state, is given by:

$$\epsilon' = \frac{\Upsilon'(1)(\Upsilon'''(1)+\Upsilon''(1)) - (\Upsilon''(1))^2}{(\Upsilon''(1))^2} \leq 0. \tag{A.8}$$

The equation above holds with equality if the elasticity is constant (e.g., under CES preferences). Also in this model, the desired markup is given by the Lerner index. Log-linearizing the Lerner index around the steady state and using equation (A.8), we have that, up to a first-order approximation, the log-markup (in deviation from the steady state) is equal to:

$$\mu_{fit} - \mu = \frac{\epsilon'}{\epsilon(\epsilon-1)}\left(y_{fit} - y_{it}\right)$$

Finally, log-linearizing the demand function (A.1) and using it to replace the log difference in output, we obtain:

$$\mu_{fit} - \mu = -\Gamma\left(p^o_{fit} - p_{it}\right)$$

where, in the case of Kimball preferences, the sensitivity of the markup to the relative price is given by $\Gamma := \frac{\epsilon'}{\epsilon(\epsilon-1)}\frac{1}{\Upsilon''(1)}$. Notice that, because there there is a continuum of firms within an industry, we have that $p_{it} = p^{-f}_{it}$ without loss of generality.

A.2 Derivations for Section 5

Derivation of equation (20). Denote by $\Delta p^o_{fit} := p^o_{fit} - p_{fit-1}$. Let us first compute $\mathbb{V}(\Delta p_{fit})$ and show that it is increasing in π^2_{it}. Using the law of motion of prices, industry inflation is given by:

$$\pi_{it} = (1-\theta)(p^o_{it} - p_{it-1})$$

Using the fact that price changes are a Bernoulli random variable (second line):

$$\begin{aligned}
\mathbb{V}(\Delta p_{fit}) =& \int_{\mathcal{F}_i} (\Delta p_{fit})^2 df - \left(\int_{\mathcal{F}_i} \Delta p_{fit} df\right)^2 \\
=& \theta \int_{\mathcal{F}_i} (p_{fit-1} - p_{fit-1})^2 df + (1-\theta)\int_{\mathcal{F}_i} (\Delta p^o_{fit})^2 df - \pi^2_{it} \\
=& (1-\theta)\int_{\mathcal{F}_i} (\Delta p^o_{fit})^2 df - \pi^2_{it} \\
=& (1-\theta)\left(\mathbb{V}(\Delta p^o_{fit}) + \left(\frac{1}{1-\theta}\pi_{it}\right)^2\right) - \pi^2_{it} \\
=& (1-\theta)\,\mathbb{V}(\Delta p^o_{fit}) + \frac{\theta}{1-\theta}\pi^2_{it}
\end{aligned}$$

Therefore keeping fixed $\mathbb{V}(\Delta p^o_{fit})$, $\mathbb{V}(\Delta p_{fit})$ is increasing in π^2_{it}. We now show that also $\mathbb{V}(\Delta p^o_{fit})$ is increasing in π^2_{it} through the Kalman gain. Rearranging the first-order condition and making use of lemma 1:

$$\Delta \widetilde{p}^o_{fit} = \frac{1-\theta}{1-\theta(1-\theta)}\left[(1-\Omega)\widetilde{mc}_{fit} - \widetilde{p}_{fit-1}\right] + \frac{\beta_{it}\theta}{1-\theta(1-\theta)}(\Delta \widetilde{p}^o_{fit+1} + \eta_{fit+1})$$

The above is a differential equation in $\Delta\widetilde{p}^o$ that can be solved backward as a function of the histories of real component of costs $\widetilde{mc}_{fit}$, lagged relative prices $\widetilde{p}_{fit-1}$, and noise η_{fit+1}. As all these variables are in deviation from the industry average hence none of them depends directly on the aggregate inflation rate, $\Delta\widetilde{p}^o_{fit+1}$ can only depend on π_{it+1} through β_{it}. Moreover, it is immediate to see that for any given values of $\{\widetilde{mc}_{fi\tau}, \widetilde{p}_{fi\tau-1}, \eta_{fi\tau+1}\}^t_{\tau=0}$, the solution for $\Delta\widetilde{p}^o_{fit+1}$ must be increasing in β_{it}. It follows that an increase in β_{it} leads to an increase in $\mathbb{V}(\Delta p^o_{fit+1})$ and therefore $\mathbb{V}(\Delta p_{fit+1})$ is increasing in π^2_{it+1} both directly because of the term $\frac{\theta}{1-\theta}\pi^2_{it}$ and through the change in β_{it}.

Derivation of equation (21). We now derive $\mathbb{V}(\mathbb{E}_{fit}(\Delta p_{fit+1}))$ as a function of the Kalman gain. The projection onto the signal space is given by:

$$\mathbb{E}_{fit}(\Delta p_{fit+1|t}) = \Delta p_{fit+1|t} + \beta_{it}\eta_{fit+1} + (1-\beta_{it})(\pi^e_{it} - \Delta p_{fit+1|t}). \tag{A.9}$$

The regression coefficient (Kalman gain) is:

$$\beta_{it} = \frac{\text{Cov}(\Delta p_{fit+1|t},\, s_{fit})}{\mathbb{V}(s_{fit})} = \frac{\sigma^2_{e,it}}{\sigma^2_{e,it} + \sigma^2_{\eta}}.$$

The industry mean belief is obtained by averaging equation (A.9) and making use of the fact that noise is mean zero in the cross-section:

$$\int_{\mathcal{F}_i} \mathbb{E}_{fit}(\Delta p_{fit+1|t})df = \beta_{it}\pi_{it+1|t} + (1-\beta_{it})\pi^e_{it},$$

where $\pi_{it+1|t}$ denotes the cross-sectional average of $\Delta p_{fit+1|t}$. Using equation (A.9), the industry dispersion in beliefs can be obtained as:

$$\begin{aligned}
\mathbb{V}(\mathbb{E}_{fit}(\Delta p_{fit+1|t})) = &\int_{\mathcal{F}_i} \left(\mathbb{E}_{fit}(\Delta p_{fit+1|t}) - \int_{\mathcal{F}_i} \mathbb{E}_{jit}(\Delta p_{jit+1|t})dj \right)^2 df \\
=&\beta^2_{it}(\sigma^2_{e,it} + \sigma^2_{\eta}) \\
=&\beta_{it}\sigma^2_{e,it}.
\end{aligned}$$

Finally, we notice that $\mathbb{E}_{fit}(\Delta p_{fit+1|t}) = \mathbb{E}_{fit}(\Delta p_{fit+1})$ when a firm is resetting the price because innovations have a zero mean, which concludes the proof.

A.3 Proofs for the Text

Proof of Lemma 1.

Proof. Denote by $\Delta p^o_{fit} := p^o_{fit} - p_{fit-1}$ the scaled conditional price change. We notice that, as there is no information regarding the identity of firm f in the common-knowledge information set, the prior mean is identical for all firms within the same industry $\mathbb{E}_{it}(\Delta p^o_{fit+1}) = \mathbb{E}_{it}(\Delta p^o_{it+1})$, where p^o_{it} denotes the industry average of p^o_{fit}. Averaging equation (12) in the industry cross-section:

$$\int_{\mathcal{F}_i} \mathbb{E}_{fit}(\Delta p^o_{fit+1})df = (1-\beta_{it})\,\mathbb{E}_{it}(\Delta p^o_{it+1}) + \beta_{it}\Delta p^o_{it+1},$$

as the noise averages out in the cross-section and the prior is common. Because price adjustments are a Bernoulli random variable, by the law of large numbers industry inflation satisfies:

$$\pi_{it+1} = \int_{\mathcal{F}_i} \Delta p_{fit+1} df = (1-\theta)\Delta p^o_{it+1}.$$

Similarly, using the law of iterated expectations and the expected law of motion $\mathbb{E}_{fit}(\Delta p_{fit}) = (1-\theta)(p^o_{fit} - p_{fit-1})$:

$$\mathbb{E}_{fit}(\Delta p_{fit+1}) = (1-\theta)\,\mathbb{E}_{fit}(\Delta p^o_{fit+1}).$$

Replacing the above into equation (12) substituting the prior belief $(1-\beta_{it})\,\mathbb{E}_{it}(\Delta p^o_{it+1})$ leads to the result. □

Proof of Proposition 1

Proof. Denote by $\Delta p^o_{fit} := p^o_{fit} - p_{fit-1}$ the scaled conditional price change. Starting from equation (8) subtract p_{fit-1} from both sides and multiply by $(1-\theta)$:

$$(1-\theta)\Delta p^o_{fit} = \mathbb{E}_{fit}\Big\{(1-\theta)^2\Big[\mu_f + (1-\Omega)mc^r_{fit} + p_{it} - p_{fit-1}\Big] + \theta(1-\theta)(\Delta p^o_{fit+1} + \Delta p_{fit})\Big\}.$$

Using the expected law of motion $\mathbb{E}_{fit}(\Delta p_{fit}) = (1-\theta)\Delta p^o_{fit}$ and the law of iterated expectations for nested sets $\mathcal{I}_{fit} \subseteq \mathcal{I}_{fit+1}$:

$$\mathbb{E}_{fit}(\Delta p_{fit+1}) = (1-\theta)\,\mathbb{E}_{fit}(\Delta p^o_{fit+1}).$$

Replacing into the above:

$$\mathbb{E}_{fit}(\Delta p_{fit}) = \mathbb{E}_{fit}\Big\{(1-\theta)^2\Big[\mu_f + (1-\Omega)mc^r_{fit} + p_{it} - p_{fit} + \Delta p_{fit}\Big] + \theta\Delta p_{fit+1} + \theta(1-\theta)\Delta p_{fit}\Big\}.$$

Collecting $\mathbb{E}_{fit}(\Delta p_{fit})$ and taking the terms that are observed out of the expectations:

$$\mathbb{E}_{fit}(\Delta p_{fit}) = \frac{(1-\theta)^2}{\theta}\Big[\mu_f + (1-\Omega)mc^r_{fit} + (p_{it} - p_{fit})\Big] + \mathbb{E}_{fit}\Big\{\Delta p_{fit+1}\Big\}.$$

Define now a sampling error $\omega_{fit} := (1-\theta)\Delta p^o_{fit} - \Delta p_{fit}$. By construction ω_{fit} is mean zero in the cross-section (as both terms integrate to industry inflation π_{it}). In terms of the sampling error:

$$\Delta p_{fit} = \frac{(1-\theta)^2}{\theta}\Big[\mu_f + (1-\Omega)mc^r_{fit} + (p_{it} - p_{fit})\Big] + \mathbb{E}_{fit}\Big\{\Delta p_{fit+1}\Big\} + \omega_{fit}.$$

Now subtract from both sides the industry average (i.e. an industry-by-time fixed effect):

$$\Delta\widetilde{p}_{fit} = \frac{(1-\theta)^2}{\theta}\Big[\mu_f + (1-\Omega)\widetilde{mc}_{fit} - \widetilde{p}_{fit}\Big] + \mathbb{E}_{fit}(\Delta p_{fit+1}) - \int_{\mathcal{F}_i}\mathbb{E}_{jit}(\Delta p_{jit+1})dj + \omega_{fit}.$$

Making use of lemma 1:

$$\begin{aligned}
\Delta\widetilde{p}_{fit} &= \frac{(1-\theta)^2}{\theta}\left[\mu_f + (1-\Omega)\widetilde{mc}_{fit} - \widetilde{p}_{fit}\right] + \beta_{it}\cdot\left((1-\theta)\Delta p^o_{fit+1} - \pi_{it+1} + \eta_{fit}\right) + \omega_{fit} \\
&= \frac{(1-\theta)^2}{\theta}\left[\mu_f + (1-\Omega)\widetilde{mc}_{fit} - \widetilde{p}_{fit}\right] + \beta_{it}\cdot\left(\Delta\widetilde{p}_{fit+1} + \eta_{fit}\right) + u_{fit+1} + \omega_{fit}
\end{aligned}$$

where $u_{fit+1} := \beta_{it}((1-\theta)\Delta p^o_{fit+1} - \Delta p_{fit+1})$ as before. Collecting the two residuals into one we get the result. □

Proof of Proposition 2

Proof. Using the fact that Δp_{fit+1} is a mixture between a Bernoulli and a continuously distributed random variable:

$$\mathbb{E}\left(\text{sign}(\Delta p_{fit+1}) \mid \mathcal{I}_{fit}\right) = (1-\theta)\left(1\cdot\mathbb{P}(\Delta p^o_{fit+1} > 0 \mid \mathcal{I}_{fit}) - 1\cdot\mathbb{P}(\Delta p^o_{fit+1} < 0 \mid \mathcal{I}_{fit})\right).$$

Ignoring the multiplicative constant, we want to compute the covariance:

$$\text{Cov}\left\{\eta_{fit},\ \mathbb{E}_{fit}\,\text{sign}(\Delta p^o_{fit+1}))\right\} = \mathbb{E}\left\{\eta_{fit}\cdot\mathbb{E}_{fit}(\text{sign}(\Delta p^o_{fit+1}))\right\} - \mathbb{E}\left\{\eta_{fit}\right\}\mathbb{E}\left\{\mathbb{E}_{fit}(\text{sign}(\Delta p^o_{fit+1}))\right\}.$$

First, the expectation of the expected sign is zero by the law of iterated expectation and the symmetry assumption on the distribution of conditional price changes:

$$\begin{aligned}
\mathbb{E}\left\{\mathbb{E}_{fit}(\text{sign}(\Delta p^o_{fit+1}))\right\} &= \mathbb{E}(\text{sign}(\Delta p^o_{fit+1})) \\
&= \mathbb{P}(\Delta p^o_{fit+1} > 0) - \mathbb{P}(\Delta p^o_{fit+1} < 0) \\
&= \mathbb{P}(\Delta p^o_{fit+1} > 0) - \mathbb{P}(\Delta p^o_{fit+1} > 0) = 0
\end{aligned}$$

Let us now focus on the other term:

$$\mathbb{E}\left(\eta_{fit}\cdot\left(\mathbb{P}(\Delta p^o_{fit+1} > 0 \mid \mathcal{I}_{fit}) - \mathbb{P}(\Delta p^o_{fit+1} < 0 \mid \mathcal{I}_{fit})\right)\right).$$

Denote by $\mathbb{I}(A)$ the indicator function for the event A. The result would follow immediately if $\eta_{fit} \in \mathcal{I}_{fit}$ by applying the law of iterated expectations:

$$\begin{aligned}
\mathbb{E}\Big\{\eta_{fit}\cdot\Big[\mathbb{P}(\Delta p^o_{fit+1}>0 \mid \mathcal{I}_{fit}) - \mathbb{P}(\Delta p^o_{fit+1}<0 \mid \mathcal{I}_{fit})\Big]\Big\} &=\\
\mathbb{E}\Big\{\eta_{fit}\cdot\Big[\mathbb{E}\Big(\mathbb{I}(\Delta p^o_{fit+1}>0) \mid \mathcal{I}_{fit}\Big) - \mathbb{E}\Big(\mathbb{I}(\Delta p^o_{fit+1}<0) \mid \mathcal{I}_{fit}\Big)\Big]\Big\} &=\\
\mathbb{E}\Big\{\mathbb{E}\Big(\eta_{fit}\cdot\mathbb{I}(\Delta p^o_{fit+1}>0) \mid \mathcal{I}_{fit}\Big) - \mathbb{E}\Big(\eta_{fit}\cdot\mathbb{I}(\Delta p^o_{fit+1}<0) \mid \mathcal{I}_{fit}\Big)\Big\} &=\\
\mathbb{E}\Big\{\eta_{fit}\cdot\Big(\mathbb{I}(\Delta p^o_{fit+1}>0) - \mathbb{I}(\Delta p^o_{fit+1}<0)\Big)\Big\} &=\\
\mathbb{E}\Big\{\eta_{fit}\cdot\Big(\mathbb{I}(\Delta p^o_{fit+1}>0) - \mathbb{I}(\Delta p^o_{fit+1}>0)\Big)\Big\} &=0.
\end{aligned}$$

Finally, we notice that η_{fit} can be recovered upon observing the private signal s_{fit} and setting the price. Because firms are committing at time zero to a pricing rule for every realization of the information set $\mathcal{I}_{fit}$, in any given period t the entire path of future reset prices $\{p^o_{fit+\tau}\}_{\tau\geq 0}$ is directly observable as it is only a function of $\mathcal{I}_{fit}$. Therefore from equation (11), upon observing s_{fit} and having set p^o_{fit} and therefore $\{p^o_{fit+\tau}\}_{\tau\geq 1}$, firms can recover the noise η_{fit}. □

Proof of Proposition 3

Proof. Denote the average industry belief of an unconditional price change with $\bar{E}_{it}(\pi_{it+1}) := \int_{\mathcal{F}_i}\mathbb{E}_{fit}(\Delta p_{fit+1})df$. The average belief is given by $\bar{E}_{it}(\pi_{it+1}) = \mathbb{E}_{it}(\Delta p_{fit+1}) + \beta_{it}(\pi_{it+1} - \mathbb{E}_{it}(\Delta p_{fit+1}))$, where $\mathbb{E}_{it}(\Delta p_{fit+1})$ is the prior mean. We notice that the prior belief contains no information about the identity of firm f, hence $\mathbb{E}_{it}(\Delta p_{fit+1}) = \mathbb{E}_{it}(\pi_{it+1})$. Aggregating the belief from lemma 1 across all $f\in\mathcal{F}_i$ and industries $i\in[0,1]$:

$$\int_{[0,1]}\int_{\mathcal{F}_i}\mathbb{E}_{fit}(\Delta p_{fit+1})df\,di = \int_{[0,1]}\int_{\mathcal{F}_i}\Big(\bar{E}_{it}(\pi_{it+1}) + \beta_{it}((1-\theta)\Delta p^o_{fit+1} - \pi_{it+1}) + \beta_{it}\eta_{fit+1}\Big)df\,di$$

Because β_{it} is constant within an industry and $(1-\theta)\Delta p^o_{fit+1}$ averages to π_{it+1}:

$$\int_{\mathcal{F}_i}\beta_{it}((1-\theta)\Delta p^o_{fit+1} - \pi_{it+1})df = 0.$$

Moreover, because the noise is mean zero in the cross-section, $\int_{\mathcal{F}_i}\beta_{it}\eta_{fit+1}df = 0$. We now use the fact that the information set $\mathcal{I}_{it}$ contains only information regarding regarding industry i, so that $\mathbb{E}_{it}(\pi_{it+1}) = \mathbb{E}_{it}(\pi_{t+1})$. Denoting by $\mathbb{E}_t(\pi_{t+1}) := \int_{[0,1]}\mathbb{E}_{it}(\pi_{t+1})di$ the average prior belief, the average posterior belief can be rewritten as:

$$
\begin{aligned}
\int_{[0,1]}\int_{\mathcal{F}_i}\mathbb{E}_{fit}(\Delta p_{fit+1})df\ di &= \int_{[0,1]}[\mathbb{E}_{it}(\pi_{it+1})+\beta_{it}(\pi_{it+1}-\mathbb{E}_{it}(\pi_{it+1}))]di \\
&= \mathbb{E}_t(\pi_{t+1})+\left(\int_{[0,1]}\beta_{it}di\right)\cdot(\pi_{t+1}-\mathbb{E}_t(\pi_{t+1}))+\mathrm{Cov}(\beta_{it},\ \pi_{it+1}-\mathbb{E}_{it}(\pi_{it+1})) \\
&= \mathbb{E}_t(\pi_{t+1})+\left(\int_{[0,1]}\beta_{it}di\right)\cdot(\pi_{t+1}-\mathbb{E}_t(\pi_{t+1}))+\mathrm{Cov}(\beta_{it},\pi_{it+1}).
\end{aligned}
$$

The last line follows from $\mathrm{Cov}(\beta_{it},\mathbb{E}_{it}(\pi_{it+1}))=0$ because β_{it} is only a function of the prior variance which is independent of the prior mean under normal distributions (Basu's theorem). Denote by $\beta_t := \int_{[0,1]}\beta_{it}di$. Conditionally on the information set $\mathcal{I}_{it}$, the second term is a mean-zero Gaussian expectational error because it is a linear combination of Gaussian mean-zero random variables. Starting from the industry-level error term:

$$
\pi_{it+1}-\mathbb{E}_{it}(\pi_{it+1})=\int_{\mathcal{F}_i}(\Delta p_{fit+1|t}-\mathbb{E}_{it}(\Delta p_{fit+1|t}))\ df \sim \mathcal{N}.
$$

Therefore averaging once more, we obtain $\pi_{t+1}-\mathbb{E}_t(\pi_{t+1})\sim\mathcal{N}$. As $\beta_{it}\in\mathcal{I}_{it}$ can be treated as a non-stochastic parameter, it follows that the product between the error and the Kalman gain is a Gaussian random variable with mean zero and a time-varying volatility.

We are now ready to derive the NKPC under information frictions. Rearranging the first-order condition (8) we obtain equation (14):

$$
\Delta p_{fit}=\frac{(1-\theta)^2}{\theta}\left[(1-\Omega)(\mu+mc^r_{fit})-\widetilde{p}_{fit}\right]+\mathbb{E}_{fit}(\Delta p_{fit+1})+u_{fit},
$$

where u_{fit} is an i.i.d. error that satisfies $\int_{[0,1]}\int_{\mathcal{F}_i}u_{fit}df\ di=0$ by the law of large numbers. Integrating both sides across all firms and industries and making use of the assumption on constant short-run returns to scale in the aggregate:

$$
\pi_t=\frac{(1-\theta)^2}{\theta}(1-\Omega)(\mu+mc^r_t)+\int_{i\in[0,1]}\int_{\mathcal{F}_i}\left(\mathbb{E}_{fit}(\Delta p_{fit+1})\right)df\ di
$$

Replacing the average belief leads to the desired formula for the NKPC.

□

Proof of Corollary 1

Proof. Denote an unanticipated nominal cost shock with $c_{t+1}\notin\mathcal{I}_t$ that leads to a unitary increase in inflation tomorrow. The increase in current inflation due to the cost shock is bounded above by one, which is the upper bound for the pass-through coefficient of any nominal shock. Therefore:

$$1 \geq \frac{\partial \pi_t}{\partial c_{t+1}} = \beta_t + \frac{\partial \text{Cov}(\beta_{it}, \pi_{it+1})}{\partial c_{t+1}}$$

The Kalman gain β_{it} is a function only of information that is common knowledge, but the shock is unanticipated. Hence it has no effect on it. For small deviations from the symmetric steady state, we can apply the dominated convergence theorem to switch integration and derivative:

$$\beta_t \leq 1 - \frac{\partial \text{Cov}(\beta_{it}, \pi_{it+1})}{\partial c_{t+1}} = 1 - \text{Cov}\left(\beta_{it}, \frac{\partial \pi_{it+1}}{\partial c_{t+1}}\right) \leq 1.$$

We now notice that the covariance cannot be negative and moreover it is an even function:

$$\text{Cov}\left(\beta_{it}, \frac{\partial \pi_{it+1}}{\partial c_{t+1}}\right) = \text{Cov}\left(\beta_{it}, -\frac{\partial \pi_{it+1}}{\partial c_{t+1}}\right),$$

because β_{it} does not depend on the sign of the shock. In addition, the covariance is smooth for small shocks:

$$\lim_{c_{t+1} \downarrow 0} \text{Cov}\left(\beta_{it}, \frac{\partial \pi_{it+1}}{\partial c_{t+1}}\right) = \lim_{c_{t+1} \uparrow 0} \text{Cov}\left(\beta_{it}, \frac{\partial \pi_{it+1}}{\partial c_{t+1}}\right)$$

Hence the limit exists, and the covariance function admits a zero at $\text{Cov}(\cdot, 0) = 0$. Because we are considering small deviations from the steady state, then there is either only one zero or the covariance is constant at zero. Therefore, as the covariance is not constant by assumption, it must be that:

$$\frac{\partial \text{Cov}(\beta_{it}, \pi_{it+1})}{\partial c_{t+1}} > 0 \;\; a.s. \;\; c_{t+1} \neq 0,$$

which concludes the proof. □

Proof of Proposition 4.

Proof. Up to a second-order approximation around the symmetric steady state with no inflation and no dispersion ($\sigma_{e,i}^2 = 0$), the Kalman gain is proportional to the signal-to-noise ratio:

$$\beta_{it} = \left(1 + \frac{\sigma_\eta^2}{\sigma_{e,it}^2}\right)^{-1} \approx \left(1 + \frac{\sigma_\eta^2}{\sigma_{e,i}^2}\right)^{-1} + \frac{1}{2}\left(1 + \frac{\sigma_\eta^2}{\sigma_{e,i}^2}\right)^{-2}\left(\frac{\sigma_\eta^2}{\sigma_{e,i}^4}\right) \cdot (\sigma_{e,it}^2 - \sigma_{e,i}^2) = \frac{1}{2}\frac{\sigma_{e,it}^2}{\sigma_\eta^2}.$$

Using the assumption of constant elasticity and linearity, the covariance simplifies to:

$$\mathrm{Cov}(\beta_{it}, \pi_{it+1}) \approx \frac{1}{2\sigma_\eta^2} \frac{\partial \sigma_{e,it}^2}{\partial \mathbb{V}(\Delta p_{fit+1})} \frac{\partial \mathbb{V}(\Delta p_{fit+1})}{\partial \pi_{it+1}^2} \mathrm{Cov}(\pi_{it+1}^2, \pi_{it+1}).$$

Using formula (20), because the indirect effect is of higher order as it depends on the change in β_{it}, we obtain:

$$\frac{\partial \mathbb{V}(\Delta p_{fit+1})}{\partial \pi_{it+1}^2} \approx \frac{1-\theta}{\theta}.$$

Using the relationship between the raw moments and the cumulants:

$$\begin{aligned}\mathrm{Cov}(\pi_{it+1}^2, \pi_{it+1}) &= \int_{[0,1]} \pi_{it+1}^3 di - \left(\int_{[0,1]} \pi_{it+1}^2 di\right)\left(\int_{[0,1]} \pi_{it+1} di\right) \\ &= (\pi_{t+1}^3 + 3\pi_{t+1}\sigma_t^2 + \gamma_t \sigma_t^3) - \pi_{t+1}(\pi_{t+1}^2 + \sigma_t^2) \\ &= 2\pi_{t+1}\sigma_t^2 + \gamma_t \sigma_t^3.\end{aligned}$$

Putting all together leads to the formula in the text. □

Proof of Proposition 5.

Proof. To derive a solution, it is more convenient to aggregate the first-order conditions starting from equation (9), which has on the RHS the lag of the price rather than the current price. Aggregating and following the same steps as before we obtain:

$$\pi_t^e = \frac{(1-\theta)^2}{1-\theta(1-\theta)}[(1-\Omega)(mc_t^n - p_{t-1}) + \Omega \pi_t^e] + \frac{\theta}{1-\theta(1-\theta)}\left(1 + \zeta \frac{\theta}{1-\theta}\frac{\sigma_t^2}{\sigma_\eta^2}\right)\mathbb{E}_t(\pi_{t+1}^e). \quad \text{(A.10)}$$

Equivalently, we can obtain the above by rearranging equation (25).

Fix the price level at p_{t-1} and consider a shock (ε_t) to aggregate nominal marginal cost ($mc_t^n = mc_{t-1}^n + \varepsilon_t$) that leads to the realized dispersion σ_t^2. We look for a solution with time-varying coefficient Ψ_t of the following form:

$$\pi_t^e = \Psi_t(mc_t^n - p_{t-1}) = \Psi_t(mc_{t-1}^r + \varepsilon_t).$$

Using the random walk assumption of marginal cost:

$$\begin{aligned}\mathbb{E}_t(\pi_{t+1}^e) &= \mathbb{E}_t\left(\zeta_{t+1}(mc_{t+1}^n - p_t^e)\right) \\ &= \Psi_t\left(mc_t^n - \Psi_t(mc_t^n - p_{t-1}) + p_{t-1}\right) \\ &= \Psi_t(1-\Psi_t)(mc_t^n - p_{t-1})\end{aligned}$$

Replacing the guess and the above into equation (A.10) leads to the following restriction on the undetermined coefficient:

$$\Psi_t = \frac{(1-\theta)^2}{1-\theta(1-\theta)}\left[1-\Omega+\Psi_t\Omega\right] + \frac{\theta}{1-\theta(1-\theta)}\left(1+\zeta\frac{\theta}{1-\theta}\frac{\sigma_t^2}{\sigma_\eta^2}\right)\Psi_t(1-\Psi_t). \tag{A.11}$$

For σ_t^2 sufficiently close to the symmetric steady state (with no price dispersion $\sigma^2 = 0$), equation (A.11) has two roots with opposite signs. We look for an analytical solution as a function of σ_t^2 that is positive at $\sigma^2 = 0$. Rewrite the formula as follows:

$$A_t\Psi_t^2 + B_t\Psi_t + C = 0$$

where

$$A_t := \frac{\theta}{1-\theta(1-\theta)}\left(1+\zeta\frac{\theta}{1-\theta}\frac{\sigma_t^2}{\sigma_\eta^2}\right),$$

$$B_t := 1-\Omega\frac{(1-\theta)^2}{1-\theta(1-\theta)} - A_t,$$

$$C := -(1-\Omega)\frac{(1-\theta)^2}{1-\theta(1-\theta)}.$$

The positive root is given by:

$$\Psi_t = \frac{-B_t+\sqrt{B_t^2-4A_tC}}{2A_t}.$$

The quadratic formula has a solution for $\sigma_t^2 = 0$ and is continuous at $\sigma_t^2 = 0$, thus a solution exists for a sufficiently small σ_t^2/σ_η^2.

Finally, it is immediate to verify from equation (A.11) that Ψ_t is increasing in σ_t^2. Therefore, as σ_t^2 is increasing in ε_t^2, Ψ_t is increasing in ε_t^2.

□

B Additional Figures

Figure A.1: Reduced-form Evidence of State-dependent Pass-through

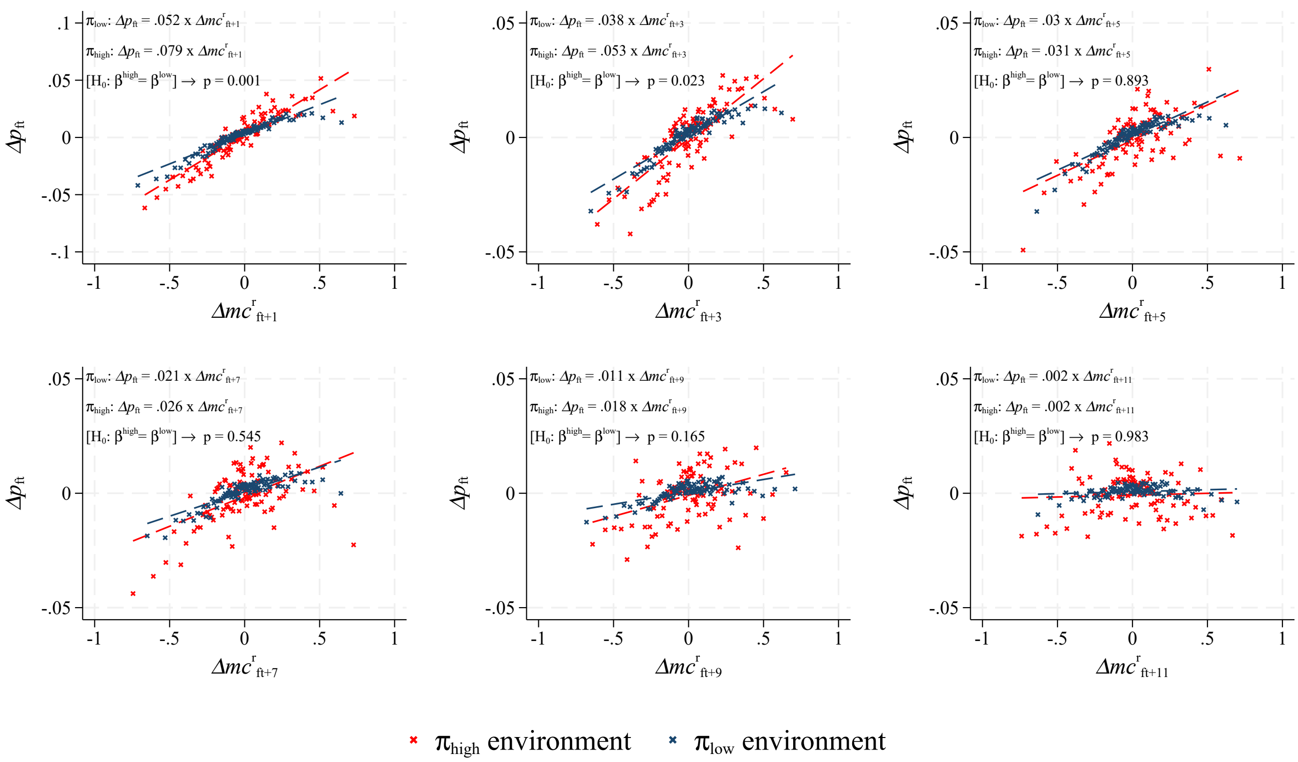

Notes. Estimates of the reduced-form coefficients of regression 19. The red dots are data from a subsample corresponding to all the firm-level observations belonging to an industry with an inflation rate in the top 10% of all the industry inflation realizations (corresponding to about 8%). The blue dots are data on the remaining subsample.

Figure A.2: Fraction of firms that did not adjust Price in Previous Month

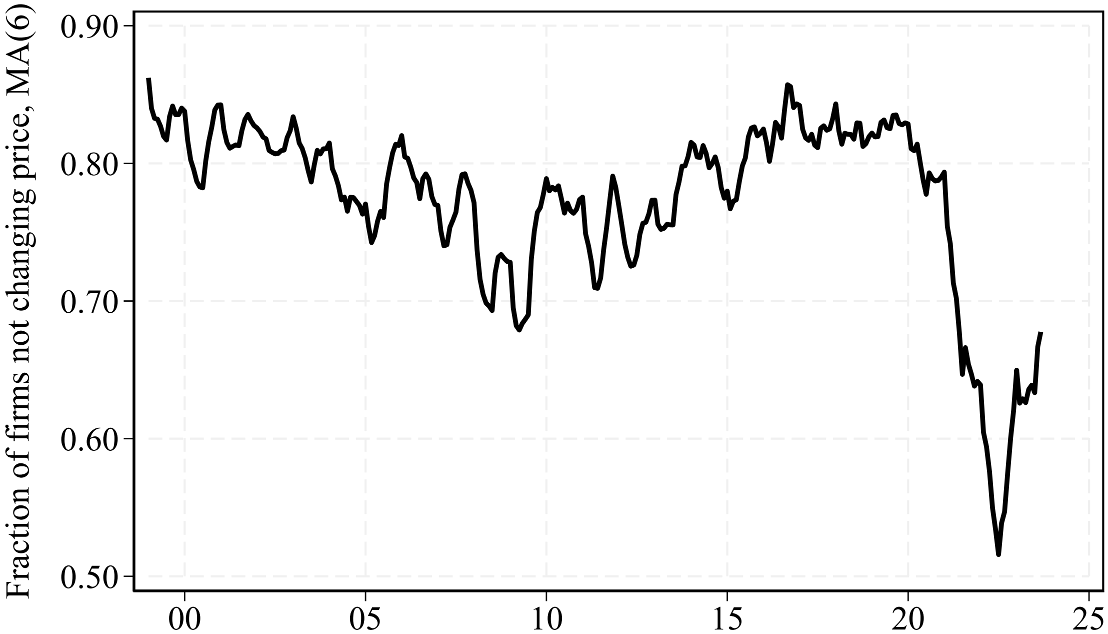

Notes. This plot reflects the fraction of firms that changed price as reported in the NBB-BS (6 month moving average).

C Data and Measurement

In this section we provide additional details on the various administrative data sources and data cleaning procedures that underlies the operationalization of our empirical framework. Most of the material in this section draws on Gagliardone et al. (2023, 2024), amended to account for the higher frequency of the estimation and additional data sources.

C.1 Data Sources and Data Cleaning

We use information from PRODCOM to compute the monthly change in product- and firm-level prices and to define the boundaries of markets (industries) in which firms compete. PRODCOM is a large-scale survey commissioned by Eurostat and administered in Belgium by the National Statistical Office. The PRODCOM sampling strategy is designed to cover at least 90% of domestic production value within each manufacturing industry (4-digit NACE codes) by surveying all firms operating in the country with (i) a minimum of 20 employees or (ii) total revenue above 4.5 million euros (European Commission 2014). Firms are required to disclose, on a monthly basis, product-specific physical quantities (e.g., volume, kg., m^2, etc.) of production sold and the value of production sold (in euros) for all their manufacturing products. Products are defined in PRODCOM by an 8-digit PC code (e.g., 15.20.13.61 is ”Mens's sandals”, 15.20.13.62 is ”Womans's sandals”, and 15.20.13.63 is ”Children's sandals”). Industries are defined by the first four digits of the product code (e.g., 15.20 is ”Manufacture of footwear”). Sectors, i.e. more aggregate structures within the economy, are defined by the first two digits of the product codes (e.g., CB is ”Textile, apparel, leather and related products”).

In the raw data, there are approximately $4,000$ product headings distributed across 13 manufacturing sectors. The PC product codes have been revised several times between 1995 and 2023, with a substantial overhaul in 2008. We use the conversion tables provided by Eurostat and firm-specific information on firms' product portfolios to harmonize the 8-digit product codes across consecutive months and harmonize 4-digit industry codes over time.[33] In most cases, the conversion tables provide a unique mapping of the 8-digit product codes across consecutive years. In a limited number of cases, the mapping is many-to-one, one-to-many, or many-to-many. The many-to-one mapping is straightforward, while the one-to-many and many-to-many mappings could be problematic. We are able to handle most of these cases using information on the basket of products produced by each firm.[34] In a limited number of cases, we do not have sufficient information to re-

[33] The official conversion tables are available at https://ec.europa.eu/eurostat/ramon. The harmonization of the industry code essentially consists of harmonizing the NACE Rev.1 industry, used before 2008, to the NACE Rev.2 industry codes, used from 2008.

[34] For example, consider a case where the official mapping indicates that product 11.11.11.11 in year t could

solve the uncertainty regarding the mapping. We drop these observations from the sample. Table A.1 reports the list of manufacturing sectors and their 2-digit PC codes.

We construct product-level prices (unit values) by dividing product-level sales by product-level quantities sold. As explained in the body of the paper, we are interested in domestic prices, i.e., prices charged by producers in Belgium. PRODCOM does not require firms to separately report production and sales to domestic and international customers. Therefore, we recover domestic values and quantities sold by combining information from PRODCOM with data on firms' product-level exports (quantities and sales) available through Belgian Customs (for extra-EU trade, "Extrastat") and the Intrastat Inquiry (for intra-EU trade).[35] We use the official conversion tables provided by Eurostat to map the CN product code classification used in the international trade data to the PRODCOM product code classification. In most cases, the CN-to-PC conversion involves either a one-to-one or many-to-one mapping, which poses no issues. We drop observations that involve one-to-many and many-to-many mappings.

We apply the following filters and data manipulations to the PRODCOM data set. First, we retain firms' observations if there was positive production reported for at least one product in all months throughout the year. In the rare cases where a firm reports positive values but quantities are missing, we impute the quantity sold from the average value-to-quantity ratio in the months where both values and quantities are reported. Second, we require firms to file VAT declarations and Social Security declarations (as explained below). These two data sources allow us to measure firms' marginal costs.

The second important use of international trade data is to obtain information on international competitors selling manufacturing products in Belgium. For each domestic firm, the merged Extrastat-Intrastat data reports the quantity purchased (in kg.) and sales (converted to euros) of different manufacturing products (about 10,000 distinct CN product headings) purchased by Belgian firms from each foreign country. As is standard when dealing with customs data, we define a foreign competitor as a foreign country-domestic buyer pair. For each foreign competitor, we aggregate the product-level sales and quantity sold at the monthly level and compute monthly prices (unit values) by taking the ratio of the two.[36]

---

map to either 22.22.22.21 or 22.22.22.22 in year $t+1$. Suppose two firms, f_1 and f_2, report in period t sales of product 11.11.11.11 in year t. If f_1 reports *only* sales of 22.22.22.21 and f_2 *only* reports sales of 22.22.22.22 in year $t+1$ we infer that we should map 11.11.11.11 to 22.22.22.21 for the former and 11.11.11.11 to 22.22.22.22 for the latter.

[35] In constructing our measure of domestic sales, we address issues related to carry-along trade, which might overstate the amount of production by firms that import products destined for immediate sales.

[36] Some CN codes change over time (although to a lesser extent than PC codes). We use the official conversion tables, available on the Eurostat website, to map CN product codes across consecutive years. We make adjustments only if the code change is one-to-one between two years. We do not account for changes in PC codes that involve splitting into multiple codes or multiple PC codes combining into one code. Effectively, these changes in the PC codes are treated as if new products are generated.

We leverage data from two administrative sources to measure firms' total production (turnover) and variable production costs on a monthly basis. First, Belgian firms file VAT declarations to the tax authority that contain information on the total sales of the enterprise as well as information on purchases of raw materials and other goods and services that entail VAT-liable transactions, including domestic and international transactions. While this submission requirement applies to quasi all Belgian firms, small firms — i.e. firms having an annual turnover < 2.5M euros excl. VAT — are allowed to submit the information on a quarterly frequency. We restrict our analysis to the (larger) firms that submit on a monthly frequency. In view of the the stratified sampling strategy of PRODCOM (supra), this has a small impact on the composition of our sample.

Second, we obtain information on employment and labor costs (wage bill) from the Social Security declarations filed quarterly by each Belgian firm with the Department of Social Security of Belgium. We smooth the quarterly wage mass equally across the three months of the quarter (and assume an equal headcount in each month). We qualify this as an innocuous assumption in view of tight labor market protection laws in Belgium, such as restrictions on hiring or firing, and work rules. Moreover, in manufacturing, the segment of a firms' labor force that can most easily be scaled up/down in the short run is often employed through temporary employment agencies. Importantly, in such a contracting setup, the wage expense shows up as a component of the monthly VAT declaration of the manufacturing firm (as a payment to the employment agency, who, in turn pays a wage that is tracked through Social Security declarations).

We sum firm-monthly level expenses on intermediates and labor to obtain a measure of total variable costs, which we use to construct firms' marginal costs. We multiply these costs by the ratio of total manufacturing sales (from PRODCOM) to total sales (from the VAT declarations) to adjust for the fact that some firms also have production outside manufacturing.[37]

Finally, we apply the following data-cleaning steps to address missing values and outliers. (i) We focus on manufacturing industries defined by the NACE 2-digit codes 15—36, dropping from our sample all product headings that correspond to mining and quarrying, and all product codes corresponding to industrial services. (ii) As is standard, we exclude firms that operate in the "Coke and refined petroleum products" sector and the "Pharmaceuticals, medicinal chemical, and botanical products" sector, whose output prices are often privately bargained or determined in international (spot) markets. We also exclude firms operating in the "Other manufacturing and repair and installation of machinery and equipment" sector, a residual grouping that encompasses firms produc-

[37] We drop observations referring to firms whose sales from manufacturing products (as measured in PRODCOM) are lower than seventy percent of total firm-level sales (as reported in the VAT declarations). This ensures that our sample includes firms whose real activity is primarily, if not entirely, in manufacturing.

ing diverse and varied products for which it is difficult to define an appropriate set of competitors. (iii) We keep only observations for which we are able to compute product-level price indexes, the corresponding quantity indexes, competitors' price indexes, and marginal costs. (iv) We drop observations for which the month-to-month change of either the firm-level price index or marginal costs is greater than 100% in absolute value. (v) Finally, for each firm-industry pair that enters our dataset discontinuously, we keep only the longest continuous time spell. This ensures that each time series used in the estimation has no gaps, which would otherwise force us to interpolate by making assumptions about prices and marginal costs when the data is not recorded. We only preserve firm-industry observations for which the time spell exceeds 24 months.

Table A.1: List of manufacturing sectors

Sector	Sector definition	NACE 2-digits codes
CA	Food products, beverages and tobacco products	10–12
CB	Textiles, apparel, leather and related products	13–15
CC	Wood and paper products, and printing	16–18
CE	Chemicals and chemical products	20
CG	Rubber and plastics products, and other non-metallic mineral products	22–23
CH	Basic metals and fabricated metal products, except machinery and equipment	24–25
CI	Computer, electronic and optical products	26
CJ	Electrical equipment	27
CK	Machinery and equipment n.e.c.	28
CL	Transport equipment	29–30

Notes. This table reports the list of manufacturing sectors in our sample and the corresponding 2-digit NACE codes.

C.2 Construction of price indexes

We construct a set of indexes that capture price changes in manufacturing goods at various levels of aggregation (firm-industry, firm, industry, individual manufacturing sector, and whole manufacturing sector).

Firm-industry price index. The main variable of interest is the price of domestically sold manufacturing products at the firm-industry level, P_{fit}, for both domestic and foreign producers. We construct this variable using information on price changes at the most disaggregated level allowed by the data.

Due to repeated product code revisions, a consistent 8-digit product code taxonomy does not exist across the entire sample period.[38] Therefore, we compute the sequence of price changes across consecutive time periods (t and $t+1$) by mapping the product codes at $t+1$ to their corresponding codes at t, aggregating them at the firm-industry level, and recovering the time series of the firm-industry price index (in levels) by concatenating monthly price changes.

Specifically, denote by $\mathcal{P}_{fit}$ the set of products manufactured by firm f and by P_{pt} the price (unit value) of a given product $p \in \mathcal{P}_{fit}$. We first compute the gross price change for each product, P_{pt}/P_{pt-1}. In doing so, we appropriately account for any changes in product codes and drop product-level observations with abnormally large price jumps in a given month ($P_{pt}/P_{pt-1} > 3$ or $P_{pt}/P_{pt-1} < 1/3$). We then construct the Törnqvist index, which measures the firm-industry price change:

$$P_{fit}/P_{fit-1} = \prod_{p \in \mathcal{P}_{fit}} (P_{pt}/P_{pt-1})^{\bar{s}_{pt}}, \tag{A.12}$$

where $\bar{s}_{pt}$ is a Törnqvist weight computed as the average of the sale shares between t and $t-1$: $\bar{s}_{pt} := \frac{s_{pt}+s_{pt-1}}{2}$.[39] Finally, we use the sequence of monthly price changes to construct the time series of firm-industry prices (in levels):

$$P_{fit} = P_{f0} \prod_{\tau=t_f^0+1}^{t} \left(P_{f\tau}/P_{f\tau-1}\right), \tag{A.13}$$

where t_f^0 denotes the first month when f appears in our data, and P_{f0} is the price level in that month. We normalize P_{f0} to one for all firm-industry pairs f in our dataset. As discussed in the paper, this normalization is immaterial for our empirical analysis, as any level effects are absorbed by the firm-industry fixed effects included in all our empirical specifications.

Firm price index. As discussed in the paper, the vast majority of firms in our data operate in only one (4-digit) industry, implying that the firm-industry price index, P_{fit}, and the firm price index, $\bar{P}_{fit}$, coincide. However, in a limited number of cases, it becomes necessary to construct a firm's price index that aggregates across different firm-industry price indexes. In doing this, we construct the firm-level price index $\bar{P}_{fit}$ following a method similar to the one described above. Specifically, we construct a Törnqvist index that aggregates across price changes of the individual (4-digit) industry bundles $i \in I_f$ produced

[38] See Appendix C.1 for additional information on the data.

[39] This index accounts for the presence of multi-product firms by averaging across products produced by the same firm in a given industry. The Törnqvist weights, $\bar{s}_{pt}$, give larger weights to those products that account for a larger share of the firm's turnover.

by firm f in month t: $\bar{P}_{fit}/\bar{P}_{fit-1} = \prod_{i \in I_f} (P_{fit}/P_{fit-1})^{\bar{s}_{fit}}$, with Törnqvist weights defined as $\bar{s}_{fit} := (s_{fit} + s_{fit-1})/2$, where s_{fit} is the share of sales of industry i in the firms' total sales (across manufacturing industries). We then concatenate the monthly price changes above to obtain the price index $\bar{P}_{fit}$, normalizing the level of the price index to one in the first month when the firm first appears in our dataset. Note that for single-industry firms the price index $\bar{P}_{fit}$ coincides with the firm-industry price index P_{fit} in (A.13).

Competitors price index. Using a similar approach, we construct the competitors' price index for each domestic firm. We start by computing monthly price changes: $P_{it}^{-f}/P_{it-1}^{-f} = \prod_{k \in \mathcal{F}_i/f} (P_{kt}/P_{kt-1})^{\bar{s}_{kt}^{-f}}$, with $\bar{s}_{kt}^{-f} := \frac{1}{2}\left(\frac{s_{kt}}{1-s_{fit}} + \frac{s_{kt-1}}{1-s_{fit-1}}\right)$ denoting a Törnqvist weight constructed by averaging the residual revenue share of competitors in the industry at time t (net of firm f revenues) with that at time $t-1$. We then concatenate the changes, normalizing the level of the price index in the first period to one. Also, in this case, the normalization is immaterial for estimation purposes as our empirical model always includes firm fixed effects. Note that the set of domestic competitors for each Belgian producer, denoted in the paper by $\mathcal{F}_i$, includes not only other Belgian manufacturers operating in the same industry but also foreign manufacturers that belong to the same industry and sell to Belgian customers.

Industry, sector, and aggregate price index. We construct the industry-level, sector-level, and aggregate (manufacturing) price indexes by aggregating monthly firm-level price changes. The formula to construct the percentage change in these price indexes is analogous to the one in (A.12), where now the Törnqvist weights assigned to each firm-industry price change, P_{fit}/P_{fit-1}, capture the (weighted) average market shares of the firm in its own industry, sector, or manufacturing, respectively. Once again, the level of the indexes is constructed by concatenating changes and normalizing the level of the price index to one for the first observation in the time series.