STAFF
REPORTS

NO. 1195
MAY 2026

Micro and Macro Cost-Price Dynamics in Normal Times and During Inflation Surges

Luca Gagliardone | Mark Gertler | Simone Lenzu | Joris Tielens

FEDERAL RESERVE BANK of NEW YORK

Micro and Macro Cost-Price Dynamics in Normal Times and During Inflation Surges
Luca Gagliardone, Mark Gertler, Simone Lenzu, and Joris Tielens
Federal Reserve Bank of New York Staff Reports, no. 1195
May 2026
https://doi.org/10.59576/sr.1195

Abstract

We develop a unified approach to studying cost-price dynamics in the cross-section of firms in order to jointly explain the time series of aggregate inflation and the frequency of price changes, both during normal times and inflation surges. A key novelty is the use of microdata on firms' prices and production costs to construct an empirical measure of price gaps—the deviation between a firm's listed and optimal price. Conditional on the path of aggregate cost shocks extracted from the data, a state-dependent pricing model with strategic complementarities accounts well for both the linear cost-price dynamics of the pre-pandemic period and the nonlinear increase in inflation and frequency of price adjustment that followed.

JEL classification: E30, E31, E32, D22
Key words: inflation dynamics, Phillips curve, nominal rigidities, firm-level microdata

---

Lenzu: Federal Reserve Bank of New York, New York University, Stern School of Business, CEPR (email: slenzu@stern.nyu.edu). Gagliardone: Yale University (email: luca.gagliardone@yale.edu). Gertler: New York University, NBER (email: mark.gertler@nyu.edu). Tielens: National Bank of Belgium, KU Leuven (email: joris.tielens@nbb.be). The authors thank Al-Mahdi Ebsim for outstanding research assistance and Andrey Alexandrov, Fernando Alvarez, George-Marios Angeletos, Andres Blanco, Dan Cao, Erwan Gautier, Jonathon Hazell, Peter Karadi, Francesco Lippi, Virgiliu Midrigan, Luigi Paciello, Alari Paulus, John Rust, Luminita Stevens, Ludwig Straub, Raf Wouters, and seminar participants at various institutions, workshops, and conferences for helpful comments and conversations.

1 Introduction

Firms adjust output prices infrequently despite continuously evolving economic conditions, leading their prices to drift from those that maximize flow profits. In all models featuring nominal rigidities, the deviation between listed and optimal prices—the *price gap*—determines firms' pricing behavior, as it reflects the evolution of production costs since the last adjustment. What distinguishes different models is how they map price gaps to price changes, with potentially important implications for aggregate inflation. In time-dependent models (e.g., Taylor 1980; Calvo 1983), prices have a fixed duration or adjust with a fixed probability, and expected price changes are a linear function of price gaps. In contrast, in state-dependent pricing models (e.g. Caballero and Engel 1993, 2007; Golosov and Lucas 2007) expected price changes are a nonlinear function of price gaps, since both the size of adjustments and the frequency of price changes depend endogenously on the gaps.[1]

In normal times, with low inflation, time- and state-dependent pricing models deliver observationally equivalent predictions. The distinction between the two classes of models becomes important, and nonlinearities emerge, when large aggregate shocks hit the economy. The post-pandemic inflation surge well illustrates these mechanisms. Figure 1 plots year-over-year changes in the producer price index for the Belgian manufacturing sector alongside the average frequency of price adjustment from 1999:Q1 to 2023:Q4. Before the pandemic, both inflation and the average frequency remained low and relatively stable, consistent with a linear mapping between expected price changes and price gaps. Starting in early 2021, however, both inflation and the frequency rose sharply—hallmarks of state-dependent pricing and nonlinear cost-price dynamics.[2]

Recent contributions have shown that state-dependent models help explain the recent inflation surge (Blanco et al. 2024a, 2024b; Bunn et al. 2024; Cavallo et al. 2024; Morales-Jiménez and Stevens 2024). A common approach in these studies, as well as in the foundational contributions on which they build (Nakamura and Steinsson 2010; Midrigan 2011; Alvarez et al. 2016), is to rely only on price data along with the model structure

[1]Other seminal contributions include the works by Caplin and Spulber (1987), Caplin and Leahy (1991, 1997), Dotsey et al. (1999), Caballero and Engel (1993), and Caballero and Engel (2007) on state-dependent models and monetary neutrality.

[2]See Blanco et al. (2024b) and Cavallo et al. (2024) for evidence of similar dynamics in the U.S. and other developed economies and Gagliardone and Gertler (2023) and Hazell and Hobler (2024) for a discussion of the drivers of the pandemic-era inflation surge.

Figure 1: Aggregate inflation and frequency of price adjustment

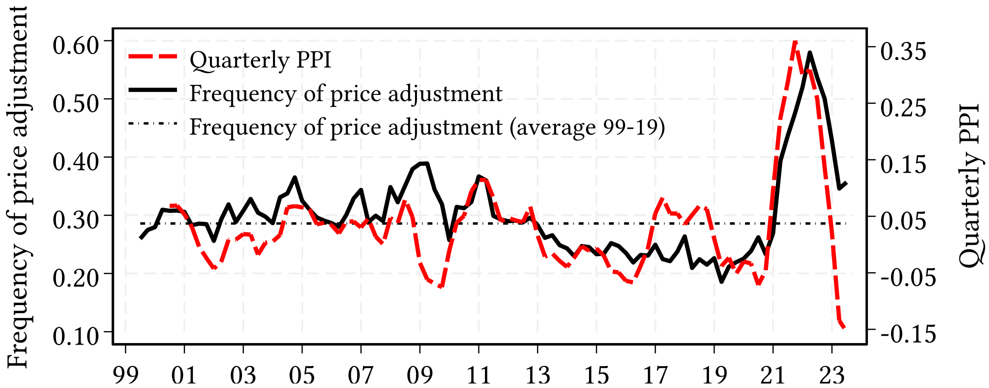

Notes. Year-over-year producer price inflation and frequency of price adjustment, Belgian manufacturing, 1999–2023. The frequency series is a four-quarter trailing average of the quarterly share of firms reporting a price change relative to the preceding quarter. The black dashed line is the average frequency of price adjustment between 1999 and 2019.

to identify state-dependent behavior. In this paper, we use cost and price microdata with the objective to jointly explain micro-level behavior and the aggregate time series. We use administrative records to assemble a rich longitudinal micro dataset covering output prices, output quantities, and production costs at a quarterly frequency for Belgian manufacturing firms. The availability of cost and price microdata allows us to construct a theory-guided, firm-level measure of price gaps. Leveraging the joint distribution of price gaps and price changes, we first nonparametrically test the predictions of a state-dependent pricing model and discipline its calibration with direct micro-evidence. We then show how a menu-cost model with strategic complementarities (Alvarez et al. 2023), calibrated to this micro-evidence and fed with observable cost-shock data, can explain the dynamics of both inflation and the aggregate frequency of price adjustment.

Our analysis provides novel evidence and insights about the passthrough of costs into prices in both the cross-section of firms and aggregate time-series.

Passthrough in the cross-section—The mapping between price gaps and price changes in the microdata provides direct evidence of the state-dependent nature of firms' pricing behavior. First, focusing on the extensive margin, we show that the frequency of price adjustment increases with the absolute value of a firm's price gap and is well approximated by a convex (quadratic) function: an empirical analogy of the generalized hazard function (GHF) postulated by the menu-cost model. Furthermore, by comparing the distribution of price gaps before and during the pandemic, we document how large aggregate shocks shifted the entire distribution, pushing many firms to its tails, where the GHF is steeper. This shift resulted in a sharp rise in the average frequency of price adjustment, consistent with the evidence in Figure 1. Second, on the intensive margin, firms adjust prices to nearly close the price gap. Consistent with the model's prediction, the elasticity of the price change with respect to the gap is close to one.

Finally, combining the extensive and intensive margins, our main proposition shows that average price changes are approximated by a cubic polynomial in the price gap. Consistent with this prediction, the data reveal an S-shaped relationship between price gaps and average price changes—a nonlinear passthrough that depends on shock size. We distinguish two distinct regions. The first region, which we refer to as the "Calvo region," includes firm-quarter observations with small price gaps, i.e., those experiencing modest shocks to their desired prices. In this region, the frequency of price adjustment is approximately constant and near its long-run (steady-state) average, and expected price changes increase linearly with the gap. The second region lies in the tails of the price gap distribution, where large shocks result in substantial deviations from firms' optimal prices (price gaps greater than 20 percent in absolute value). For these firms, the frequency of price adjustment rises more than proportionally with the size of the gap, capturing the nonlinearities induced by state-dependent pricing. Accordingly, the elasticity of average price changes with respect to gaps is approximately twice as large as in the Calvo region.

Passthrough in the time series—Our second set of results leverages our micro-evidence to explain aggregate inflation over time. Using our microdata, we construct a real marginal cost index for the Belgian manufacturing sector (and a corresponding sequence of aggregate cost shocks), with strong predictive power for inflation. Unlike prior studies that rely on unobservable shocks to fit the inflation data, we instead condition on the observable cost-shock series and feed it directly into the model to assess how well the menu-cost framework accounts for inflation dynamics.

The model closely tracks short-run fluctuations in Belgian manufacturing inflation and in the frequency of price adjustment—both during the stable pre-pandemic period and the post-pandemic surge. It captures not only the stable behavior of adjustment frequencies before the pandemic but also the sharp rise in price adjustments that followed, both in timing and magnitude. By contrast, a Calvo model fed with the same sequence of aggregate cost shocks accounts well for inflation during normal times but explains only about two-thirds of the post-pandemic inflation surge. Taken together, these empirical findings support the conclusions of a body of prior theoretical work highlighting that the distinction between time- and state-dependent pricing models is modest in the presence of small shocks (Dias et al. 2007; Gertler and Leahy 2008; Alvarez et al. 2017; Auclert et al. 2024), but emphasize that state dependence becomes essential to rationalize inflation dynamics in the wake of large shocks.

Related literature. Prior research has provided evidence consistent with state-dependent pricing by studying micro-moments extracted from the empirical distribution of price changes and has shown that state-dependent pricing models can rationalize the cross-section of price changes (see Klenow and Kryvtsov 2008, Vavra 2014, Gautier and Saout 2015, and references above). Other contributions include Gagnon (2009) and Alvarez et al. (2019), who present evidence of state-dependent pricing in high-inflation environments.

Our work advances this literature in two directions. First, we construct a firm-level measure of price gaps based on both costs and competitors' prices data. This allows us to go beyond matching moments of the distribution of price changes and directly test how pricing decisions relate to measurable deviations from desired prices. Second, we construct an observable series of cost shocks directly from the data and feed it into the model to explain inflation dynamics. This provides a sharper validation of the model than the traditional approach that relies on unobservable shocks to match inflation.

Similarly to our paper, Campbell and Eden (2014), Eichenbaum et al. (2011), Gagnon et al. (2013), Gautier et al. (2023), and Karadi et al. (2024) provide evidence of state-dependent pricing by showing that the frequency of price adjustment in microdata responds systematically to deviations from a firm's notion of its ideal price. Campbell and Eden (2014) use scanner data from supermarkets to show that the probability of nominal price adjustment is highest when a store's price substantially deviates from the average price charged by competing stores. Eichenbaum et al. (2011) use data on retail and wholesale prices from a large food and drug retailer to construct a "reference price" metric. Gautier et al. (2023) show clean evidence of state-dependent studying the dynamics of gasoline prices at individual gas stations in France. Gagnon et al. (2013) and Karadi et al. (2024) define a reset price as the average price at which competing retailers offer the same product. Like our study, these papers document that the likelihood of price adjustment increases with the absolute deviation from the ideal price.

We depart from this literature in two fundamental ways. First, our empirical measure of price gaps captures not only variation in a firm's own cost structure but also changes in markups arising from competitors' pricing behavior. This measure is anchored in theory and supported by empirical evidence that highlights cost variation and imperfect competition as key determinants of passthrough (Amiti et al. 2019; Gagliardone et al. 2025). Second, the scope and representativeness of our dataset are substantially broader than previous studies. We cover a near-complete sample of domestic manufacturing producers in Belgium—80 to 90 percent of the sector—over a 25-year period that includes both the low-inflation regime and inflationary episodes. This rich time-series variation, combined with the theoretical consistency of our price gap measure and broad cross-sectional coverage, enables a systematic characterization of the nonlinearities of the passthrough arising from shocks of varying magnitudes.

Finally, the results in this paper connect to Gagliardone et al. (2025), which estimates the slope of the cost-based New Keynesian Phillips curve in normal times, i.e. when aggregate cost shocks are small. This paper complements our earlier work by providing a nonparametric characterization of the nonlinearities of the passthrough from cost to prices—which determines the slope of the generalized NKPC (Auclert et al. 2024)—as a function of the magnitude of the cost shocks.

The paper proceeds as follows. Section 2 presents the theoretical framework, which lays the groundwork for a set of testable propositions that we evaluate in Sections 4 and 6. Sections 3 and 5 describe our dataset and the empirical measures of prices, costs, and price gaps, both at the micro and macro-level. Section 7 offers concluding remarks.

2 Theoretical framework

Our theoretical framework is a discrete-time version of Alvarez et al. (2023), which is a menu-cost model with strategic complementarities. Strategic complementarities in price setting arise because goods of different varieties are aggregated à la Kimball (1995). We allow for random menu costs of price adjustment (Dotsey et al. 1999), and random opportunities for free adjustment, which introduce a time-dependent pricing component as in the "CalvoPlus" model of Nakamura and Steinsson (2010). This tractable framework delivers testable predictions about the mapping from price gaps to price changes in the cross-section of firms (section 4), and serves as the basis for a quantitative exercise on fitting the time-series of inflation (section 5).

2.1 Model setup

In each period t, the economy is populated by a continuum of heterogeneous firms $f \in [0, 1]$ selling a single differentiated product under monopolistic competition facing a demand function à la Kimball. Let lowercase letters denote the logarithms of the corresponding uppercase variables. We denote by $p_t(f)$ the firm's price and by p_t the aggregate price index. The price index is given by:

$$p_t = \int_{[0,1]} \Big(p_t(f) - \varphi_t(f)\Big) df + O(\sigma_t^2), \tag{1}$$

where $\varphi_t(f)$ denotes a firm-specific log-taste shock, σ_t^2 is proportional to the variance of prices, and $O(\sigma_t^2)$ denotes an approximation error that goes to zero at the same rate as σ_t^2. Note that this is an approximation error due to aggregation and does not carry through in the firm-level policies characterized below.

Technology. Each firm uses a composite input $l_t(f)$ and a constant return to scale production technology to produce one unit of output, $y_t(f) = z_t(f) + l_t(f)$, where $z_t(f)$ denotes total factor productivity.[3] We assume that the latter evolves as a random walk, $z_t(f) = z_{t-1}(f) + \zeta_t(f)$, where $\zeta_t(f)$ denotes an idiosyncratic shock that is mean zero and i.i.d. over time and across firms. In this setting, firms' nominal marginal cost is given by:

$$mc_t^n(f) = mc_t^n - z_t(f). \tag{2}$$

The term mc_t^n captures an aggregate nominal cost shifter. Consistent with the empirical evidence, we assume that mc_t^n obeys a random walk $mc_t^n = mc_{t-1}^n + g_t$, where g_t is an aggregate shock, that is i.i.d. over time with mean μ_g, drawn from a unimodal, symmetric, and smooth distribution. For analytical tractability, we assume no trend inflation ($\mu_g = 0$) in what follows.[4] We relax this assumption in our empirical and quantitative exercises.

[3] In our empirical analysis, we relax the constant-return-to-scale assumption and account for variation in short-run variable costs driven by both labor and intermediate inputs.

[4] As Nakamura et al. (2018), Alvarez et al. (2019), and Alvarez et al. (2022) show, an economy with zero inflation provides an accurate approximation for economies where inflation is low, as the effect of low trend

Profit maximization. Firms choose prices to maximize the present value of profits, subject to nominal rigidities. Each firm pays a fixed cost $\chi_t(f)$ upon changing its price. As in Caballero and Engel (2007), the fixed cost $\chi_t(f)$ is the realization of a random variable, i.i.d. between firms and time, and uniformly distributed on the $[0, \overline{\chi}]$ interval. As in the CalvoPlus model, with probability $(1 - \theta^o)$ the fixed cost is zero, which implies that the firm can adjust its price for free.

We denote by $p_t^o(f)$ the firm's *static target price*, that is, the price it would choose absent nominal rigidities. Under Kimball preferences, a firm's price elasticity of demand increases with its relative price $(p_t(f) - p_t)$, which makes the desired markup decrease in relative prices. As we show in Appendix A, this implies that $p_t^o(f)$ is given by the sum of the steady-state (log) markup, $\mu(f)$, and a convex combination of the firm's nominal marginal cost and the price index:

$$p_t^o(f) = (1 - \Omega)\Big(\mu(f) + mc_t^n(f)\Big) + \Omega\Big(p_t + \varphi_t(f)\Big) \,, \tag{3}$$

where the scalar $\Omega \in [0, 1)$ captures the strength of strategic complementarities in price setting. The taste shock $\varphi_t(f)$ shows up in the target price as noise.

We assume that $z_t(f) = -\varphi_t(f)$, that is, that preference shocks are inversely proportional to productivity.[5] Following Alvarez et al. (2023), we work with a quadratic approximation of the per-period profit function around the static optimum $p_t(f) = p_t^o(f)$ and normalize it by steady-state profits. This yields the following loss function measuring the cost of deviations of the price from the target,

$$\Pi_t(f) = -\frac{\sigma(\sigma - 1)}{2(1 - \Omega)}\big(p_t(f) - p_t^o(f)\big)^2,$$

where σ is the steady-state price elasticity of demand and steady-state profits are equal to $1/\sigma$. Note that the weight on the loss function increases with the complementarity parameter Ω. By increasing the curvature of the profit function, strategic complementarities raise the firm's desire to keep its price close to the target relative to the cost of adjustment, and reduce price stickiness under both state-dependent and time-dependent pricing.

inflation on firms' decision rules is of second order.

[5]This standard assumption in the literature allows the problem to be described by a single scalar stationary state variable (price gap), and ensures that the policies of the problem do not depend on the identity of the firm. See also Woodford (2009), Midrigan (2011), Alvarez et al. (2016) for other examples.

2.2 Characterization of firms' pricing policies

Let $\mathbb{I}_t(f)$ denote an indicator equal to one if firm f adjusts its price in period t and zero otherwise. Given the time-t information set, the firm chooses a sequence of prices and adjustment decisions to maximize the expected present discounted value of profits net of adjustment costs:

$$V_t(f) = \max_{\{\, p_t(f)\,,\, \mathbb{I}_t(f)\, \}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Big\{ \Pi_t(f) - \chi_t(f) \cdot \mathbb{I}_t(f) \Big\}. \tag{4}$$

The solution to the firm's problem has an "Ss flavor": the pricing rule reduces to choosing a state-dependent probability of price adjustment and, conditional on adjustment, resetting to the frictionless target price, thereby closing the price gap.

Price gap. To characterize optimal pricing policies, we define the (ex-ante) price gap:

$$x_{t-1}(f) \equiv p_t^o(f) - p_{t-1}(f), \tag{5}$$

the percentage difference between the static target price and the price set by the firm in the previous period. Under our framework, the price gap serves as a state of the firm's problem: it is measurable *before* the firm decides whether to adjust its price (hence, "ex-ante"), yet incorporates the realization of all time t shocks through their impact on $p_t^o(f)$. To see this, substitute $p_t^o(f)$ from Equation (3) and $mc_t^n(f)$ from Equation (2) into Equation (5), and use the expressions for the aggregate and idiosyncratic components of $mc_t^n(f)$ to obtain:

$$x_{t-1}(f) = \Big(p_{t-1}^o(f) + (1-\Omega)\big(g_t + \varepsilon_t(f)\big) + \Omega\big(p_t - p_{t-1}\big) \Big) - p_{t-1}(f).$$

Here we defined $\varepsilon_t(f) \equiv \frac{\Omega}{1-\Omega}(\varphi_t(f) - \varphi_{t-1}(f)) - \zeta_t(f)$, a composite idiosyncratic shock, which combines idiosyncratic technology and taste shocks. We assume that $\varepsilon_t(f)$ is drawn from a unimodal, symmetric, and smooth distribution with mean zero and variance σ_ε^2.

When the firm chooses to adjust its price ($\mathbb{I}_t^\star(f) = 1$), the price gap jumps to:

$$x_t^\star \equiv p_t^o(f) - p_t^\star(f),$$

where $p_t^\star$ is the optimal (dynamic) reset price and $x_t^\star$ as the optimal reset gap. Note that the optimal reset gap varies over time due to aggregate shocks, but it does common across firms. (See Appendix A.2 for the derivations.)

Optimal probability of price adjustment. Unlike time-dependent models, the adjustment is an endogenous variable that depends on the price gap $x_{t-1}(f)$.

Let $h_t(x_{t-1}(f)) \equiv \mathbb{E}_\chi\big(\mathbb{I}_t(f) \mid x_{t-1}(f)\big)$ denote the probability of adjustment, ex-ante with respect to the period-t adjustment shocks, for a firm with price gap $x_{t-1}(f)$.[6] Denote by $V_t(x_{t-1}(f))$ the value of a firm that does not adjust and by V_t^a the value of a firm that adjusts to the optimal reset gap $x_t^\star$, characterized below. The value of not adjusting equals current profits at the price gap plus the discounted expected continuation value:

$$V_t(x_{t-1}) = \Pi_t(x_{t-1}) + \beta\, \mathbb{E}_t \left\{h_{t+1}(x_t) \cdot (V_{t+1}^a - \chi_{t+1}(f)) + \big(1 - h_{t+1}(x_t)\big) \cdot V_{t+1}(x_t(f))\right\}, \quad (6)$$

where $x_t(f)$ is the price gap carried into period t+1 by a firm that does not adjust in period t. The probability that a firm adjusts its price increases with the distance between the two value functions. Dropping the firm index to ease notation, given the random menu cost and the random possibility of a free price adjustment, $h_t(x_{t-1})$—also known as the generalized hazard function (GHF)—is given by:

$$\begin{aligned} h_t(x_{t-1}) &= (1-\theta^o) + \theta^o \cdot \Pr(V_t^a - \chi_t(f) \geq V_t(x_{t-1})) \\ &= (1-\theta^o) + \theta^o \cdot \min\left\{\frac{V_t^a - V_t(x_{t-1})}{\overline{\chi}}, 1\right\}, \end{aligned} \quad (7)$$

where the second line uses the assumption that the distribution of the menu cost is uniform. The expression above shows that the probability of price adjustment in a given period, $h_t(x_{t-1})$, depends, among other factors, on its price gap $x_{t-1}(f)$. With no trend inflation (and symmetric profit function), the minimum of the GHF is achieved when $x_{t-1} = 0$ and $h_t(0) = (1-\theta^o)$, the probability of a free price adjustment. Note that, as the upper bound for the menu cost $\overline{\chi}$ approaches infinity, the adjustment frequency becomes exogenous and converges to $(1-\theta^o)$. Thus, as a limiting case, the model nests a time-dependent Calvo model parameterized by θ^o.

The following lemma shows that the GHF can be accurately approximated in a neighborhood of the zero gap by a quadratic function of the price gap.

Lemma 1. *Assume stationarity of the value function, $V_t(x) = V(x)$. Up to a second-order approximation around $x_t^\star = 0$, the GHF is given by:*

$$h(x_{t-1}(f)) = (1-\theta^o) + \phi \cdot (x_{t-1}(f))^2 + o\big((x_{t-1}(f))^2\big), \quad (8)$$

[6] The expectation $\mathbb{E}_\chi$ integrates over all period-t adjustment shocks. In the CalvoPlus specification, the distribution of $\chi_t(f)$ has a mass point at zero with probability $(1-\theta^o)$, capturing the time-dependent (Calvo) component.

Figure 2: Generalized hazard function and the impact of small and large aggregate shocks

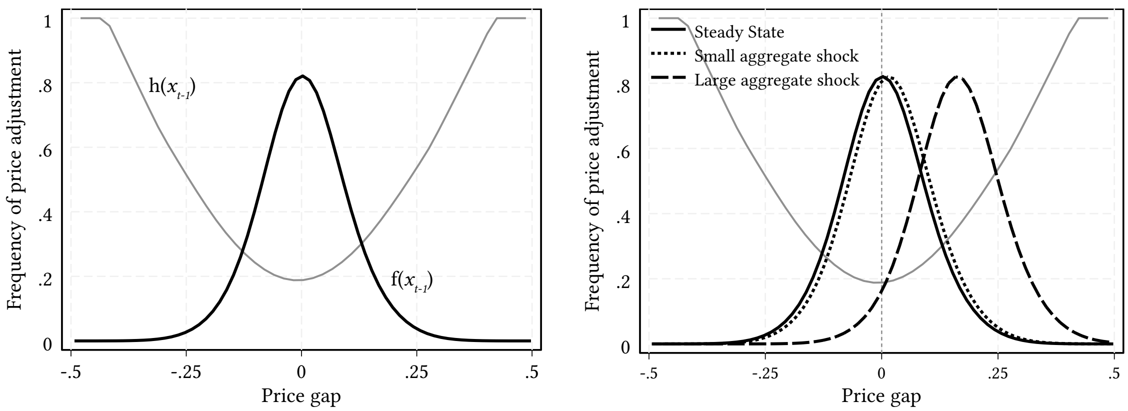

Notes. The left panel shows the GHF $h(x_{t-1})$ (gray curve) as a function of the price gap. The black curve is the cross-sectional density of price gaps in steady state $f(x_{t-1})$. Most firms cluster near zero, where the hazard is low and approximately constant. The right panel shows the displacement in the price gap distribution due to small aggregate shocks (dotted curve) and large aggregate shocks (dashed curve) relative to the steady state (solid curve).

where $o((x(f))^2)$ is an approximation error that goes to zero (an order of magnitude) faster than $(x(f))^2$, and $\phi \equiv -\frac{\theta^o}{2\bar{\chi}} \frac{\partial^2 V(x)}{\partial x^2}\Big|_{x=0}$ controls the sensitivity of the GHF to changes in gaps.

Proof. See Appendix A.3. □

Lemma 1 states that, under the second-order approximation, the GHF is U-shaped and symmetric around $x_{t-1} = 0$ in the zero-trend-inflation case, attaining its minimum $h(0) = (1 - \theta^o)$ at the zero gap; as gaps widen, the adjustment probability rises monotonically, with ϕ governing the "steepness" of the parabola. Figure 2 (left panel) illustrates the stationary equilibrium under the baseline calibration: the horizontal axis shows the simulated steady-state density of price gaps $f(x_{t-1})$—unimodal and bell-shaped—and the vertical axis shows the GHF. Given the realization of the shocks affecting $p_t^o(f)$, firms in the right (left) tail operate with a suboptimally low (high) markup and are thus more likely to increase (decrease) their price relative to the last adjusted price.

When aggregate shocks hit, they shift the entire distribution of price gaps horizontally and increase the dispersion in gaps endogenously, pushing firms into regions where the generalized hazard function is steep and substantially raising the average frequency of price adjustment. Figure 2 (right panel) illustrates this mechanism. In the spirit of the exercise in Cavallo et al. (2024), we start from an economy in its stationary equilibrium and shock it with an unexpected aggregate cost shock of size $g_t > 0$. A small aggregate shock induces only a modest shift in the price gap distribution (dotted curve), resulting in little or no change in the average frequency of price adjustment. In contrast, a large shock shifts the gap distribution toward the tails (dashed curve) and makes the quadratic term in Equation (8) quantitatively relevant. This raises the frequency of price adjustment.

Optimal reset gap. The value of the firm conditional on adjusting is the value in Equation (6) evaluated at the reset price $p_t^\star(f)$:

$$V_t^a = V_t\left(x_t^\star\right).$$

Equivalently, the optimal reset gap $x_t^\star$ solves the first-order condition $\left.\frac{\partial V_t(x)}{\partial x}\right|_{x=x_t^\star} = 0$. As we show in Appendix A, the solution of the problem can be written as:[7]

$$p_t^\star(f) = p_t^o(f) + \Psi_t,$$

where Ψ_t is an aggregate wedge given by the present discounted value of aggregate inflation. Under our assumption of random walk dynamics for marginal cost, the wedge is exactly zero provided no strategic complementarities, $\Omega = 0 \Rightarrow \Psi_t = 0$. With strategic complementarities, the wedge remains negligible, as we verify quantitatively in Appendix B. Importantly, the approximation is more accurate after large aggregate shocks: when the probability of a price adjustment in the near future is high, firms discount distant profit flows heavily (Dotsey and King, 2005), further compressing the wedge between static and dynamic price targets. It follows that, for firms that choose to adjust their prices, the optimal price change is approximately equal to their price gap:

$$\Delta p_t(f) \equiv p_t^\star(f) - p_{t-1}(f) \approx x_{t-1}(f) \qquad \text{if } \mathbb{I}_t^\star(f) = 1. \tag{9}$$

Equation (9) provides a testable implication for the mapping from price gaps to price changes along the intensive margin.

[7] The exact log-linearity of the first-order condition for the reset price follows from the quadratic approximation functional form of the profit function. We verified numerically that departures from this benchmark have little quantitative impact on our quantitative results.

2.3 Aggregate inflation

Given the solution of the firm's problem and using the formula for the price index in Equation (1), we can express aggregate inflation π_t as:

$$\begin{aligned}\pi_t &= \int \Big(p_t(f) - p_{t-1}(f)\Big)\, df + O(\sigma_t^2) = \int h_t(x_{t-1}(f)) \cdot \big(p_t^\star(f) - p_{t-1}(f)\big)\, df + O(\sigma_t^2) \\ &= \int h_t(x_{t-1}(f))\, df \cdot \int \Big(p_t^\star(f) - p_{t-1}(f)\Big)\, df + Cov\Big(h_t(x_{t-1}(f))\,,\ (p_t^\star(f) - p_{t-1}(f))\Big) + O(\sigma_t^2)\end{aligned}$$

The first line expresses aggregate inflation as the frequency-weighted average of firm-level price changes, up to $O(\sigma_t^2)$. The second line further decomposes this into (i) the product of average adjustment frequency and average price adjustment conditional on adjusting, and (ii) the covariance between frequency and conditional adjustment sizes.

With Calvo pricing, the adjustment frequency is orthogonal to the price gap and therefore uniform across firms and independent of shock size; price adjustment is a linear function of the gap, and inflation equals the product of the constant adjustment frequency and the average price gap. With state-dependent pricing, the adjustment frequency is endogenous and increases as a quadratic function of the price gap (Lemma 1). As a result, the magnitude of inflation increases convexly in the size of the aggregate price gap. Moreover, firms with larger gaps are both more likely to adjust and, conditional on adjustment, change their prices more. This "selection effect," captured by the covariance term, further amplifies the inflation response to aggregate shocks.

Below, we derive and test the model's cross-sectional and time-series predictions.

3 Data and measurement

3.1 Measurement of price changes and price gaps

Our dataset extends that in Gagliardone et al. (2025). We combine quarterly administrative records from Belgium on output prices, production costs, and competitors' prices to construct measures of price changes and price gaps. The sample covers a period of low and stable inflation (1999:Q1–2021:Q1) and the recent inflation surge (2021:Q2–2023:Q4). A detailed description of data sources, sample selection, and variable construction is provided in Gagliardone et al. (2025) and its supplemental appendix.

Price changes. Our baseline measure draws on PRODCOM, which is designed to survey all Belgian manufacturing firms with more than 10 employees and to cover at least 90 percent of output in each NACE 4-digit industry—a near census of the manufacturing sector spanning our full sample period from 1999. Using product-level domestic unit values (sales over quantity sold), we construct a Törnqvist price index for each firm (f) in industry (i) pair:[8]

$$P_{ft} = P_{f0} \prod_{\tau=t_f^o+1}^{t} \prod_{p \in \mathcal{P}_{f\tau}} \left(\frac{P_{p\tau}}{P_{p\tau-1}} \right)^{\bar{s}_{p\tau}}, \tag{10}$$

where $\mathcal{P}_{f\tau}$ is the set of 8-digit products (p) manufactured by firm f in period τ and $\bar{s}_{p\tau} \equiv (s_{p\tau} + s_{p\tau-1})/2$ is the Törnqvist weight assigned to each product. P_{f0} is the price level in the base period t_f^o—the first quarter it appears in the data. Log price changes are then:

$$\Delta p_{ft} \equiv \ln(P_{ft}/P_{ft-1}).$$

We demean price changes by removing firm- and industry×calendar-quarter fixed effects. This absorbs trend inflation (approximately 0.6 percent per quarter before the surge), accounts for permanent cross-sectional heterogeneity in pricing behavior, and controls for industry-specific seasonality.

Frequency of price adjustment. Our baseline measure of the frequency of price changes relies on combined micro-level records from PRODCOM and the National Bank of Belgium Business Survey (NBB-BS); the NBB-BS is a large, representative survey of manufacturing firms that asks respondents whether they increased, decreased, or maintained the price of a representative product relative to the previous period. We use the NBB-BS to filter spurious price changes in PRODCOM due to the measurement error in unit values (Eichenbaum et al., 2014; Cavallo and Rigobon, 2016). Specifically, we set firm-specific minimum thresholds for classifying upward and downward unit-value changes as price changes, $\mathbb{I}^+_{ft} = 0$ and $\mathbb{I}^-_{ft} = 0$, calibrated so that the implied frequencies

[8] The PRODCOM and PPI microdata are obtained from Statbel (the Belgian statistical office). We focus on pricing decisions in the domestic market. To that end, we recover domestic values and quantities sold for each firm-product pair by merging product-level PRODCOM data with firms' product-level export quantities and sales from Belgian customs declarations. Product codes are harmonized over time to account for changes in product definitions. The vast majority of firms operate in a single industry, and the main industry accounts for the lion's share of multi-industry firms' sales. This implies that for the majority of our sample, the firm's domestic price index coincides with the firm-industry's domestic price index. See Gagliardone et al. (2025) for further discussion on the data and measurement.

of upward and downward price adjustment match the average upward and downward frequency of price change in the NBB-BS.[9]

While our baseline frequency measure partially relies on unit values, in Appendix C we show robustness using an alternative measure of prices changes from the official PPI microdata. This series is well-suited for computing statistics that are sensitive to the mismeasurement of small price changes—including the frequency and kurtosis of price changes. However, its coverage is restricted across time, firms, and product portfolios.[10]

Price gaps. Guided by the theoretical framework, we construct an empirical counterpart of the ex-ante price gap using high-frequency firm-level price and cost data:

$$x_{ft-1} = p^o_{ft} - p_{ft-1}.$$

Our measure of firms' target prices (Equation (3)) is a convex combination of the firm's nominal marginal cost and its competitors' price index:

$$p^o_{ft} = (1-\Omega) mc^n_{ft} + \Omega p^{-f}_t.$$

To derive a firm-level nominal marginal cost index, mc^n_{ft}, we assume a cost structure in which the nominal marginal cost is proportional to average variable costs (total variable costs over physical output, TVC^n_{ft}/Y_{ft}). In logs:

$$mc^n_{ft} = (tvc^n_{ft} - y_{ft}) + \ln(1+\nu_f).$$

We measure quarterly total variable cost as the sum of intermediate input costs (materials and services, from VAT declarations) and labor costs (wage bill, from social security declarations). The curvature of the short-run cost function, $\nu_f \equiv 1/SRTS_f - 1$ governs the proportionality between marginal and average variable costs; we allow it to vary across firms by absorbing $\ln(1+\nu_f)$ into the firm fixed effects removed from price gaps.

We construct the competitors' price index, p^{-f}_t, combining microdata from

[9]Formally, $\mathbb{I}^+_{ft} = 0$ if $\left(\Delta p_{ft} < \kappa^+_f := \kappa^+ \cdot \mathrm{Var}_f(\Delta p_{ft}) \text{ and } \Delta p_{ft} > 0\right)$, and $\mathbb{I}^-_{ft} = 0$ if $\left(\Delta p_{ft} > -\kappa^-_f := -\kappa^- \cdot \mathrm{Var}_f(\Delta p_{ft}) \text{ and } \Delta p_{ft} < 0\right)$. This approach generalizes the standard practice of setting price changes below a given threshold to zero (Klenow and Kryvtsov 2008; Alvarez et al. 2016), replacing the uniform cutoff (e.g, ±1 percent) with firm-specific thresholds that reflect heterogeneous measurement error across firms and sectors (Hottman et al., 2016). Furthermore, the asymmetric thresholds are consistent with evidence of differential upward and downward nominal rigidity (Karadi et al., 2024; Luo and Villar, 2021).

[10]The PPI-matched sample spans 15 percent of our baseline sample (728 firms from 2003 onward), concentrated among larger firms; moreover, within each firm, the PPI tracks only a subset of firms' high-volume products rather than their full portfolio. See Appendix C for further discussion and a comparison of the NBB-BS and PPI datasets.

PRODCOM and Belgian customs declarations (covering foreign firms active in the Belgian market). Following a procedure analogous to the firm-level price index, we aggregate the price changes of products sold by domestic and international competitors in the same industry as f ($k \in \mathcal{F}_i \setminus f$):

$$P_{it}^{-f} = P_0^{-f} \prod_{\tau=t_f^o+1}^{t} \prod_{k \in \mathcal{F}_i \setminus f} \left(\frac{P_{k\tau}}{P_{k\tau-1}} \right)^{\bar{s}_{k\tau}^{-f}} .$$

where the Törnqvist weight assigned to competitor k in the industry (excluding firm f's revenues) is given by $\bar{s}_{k\tau}^{-f} \equiv \frac{1}{2}\left(\frac{s_{k\tau}}{1-s_{f\tau}} + \frac{s_{k\tau-1}}{1-s_{f\tau-1}}\right)$ and P_0^{-f} is the level in the base period. Finally, we calibrate $\Omega = 0.5$ to match the strength of strategic complementarities in the microdata (Gagliardone et al. 2025).

As for price changes, we demean price gaps by removing firm- and industry×calendar-quarter intercepts. These absorb components of the theoretical price gap that enter the empirical measure as additive constants: unobserved steady-state markups, μ_f, returns-to-scale curvature, $\ln(1 + \nu_f)$, and the normalizations of base-period price levels (P_{f0} and P_0^{-f}). A discrepancy between the empirical and theoretical price gap measure that cannot be directly addressed is the idiosyncratic taste shock φ_{ft}. As the evidence in Gagliardone et al. (2025) suggests, such shocks are transitory and average out in the cross-section. At the firm level, however, they introduce measurement error that can attenuate the relationship between price gaps and price changes. We return to this point in Section 4.1.

3.2 The empirical distribution of price changes and price gaps

The final dataset tracks 4,974 domestic firms operating across 169 narrowly defined manufacturing industries (4-digit NACE rev.2 codes), for a total of 118,308 observations. Table 1 presents summary statistics of the distribution of price changes (first four columns) and price gaps (last three columns).

Panel a focuses on the pre-pandemic period (1999–2019), characterized by low and stable inflation and—apart from the global financial crisis—the absence of large aggregate shocks. We treat this period as our empirical analogue of the model's ergodic distribution. By construction, after removing firm-industry intercepts and calendar-quarter fixed effects, the harmonized average price change is near zero. The standard deviation of price

Table 1: Summary statistics of price changes and price gaps

Price changes (Δp_{ft})				Price gaps (x_{ft-1})		
Panel a: Time period 2000-2019						
Mean	Std.	Freq. Adj.	Kurt.	Mean	Std.	Kurt.
0.00	0.12	0.29	2.94	-0.00	0.13	2.84
Panel b: Time period 2020-2023						
Mean	Std.	Freq. Adj.	Kurt.	Mean	Std.	Kurt.
0.02	0.11	0.44	2.90	0.01	0.14	2.79

Notes. Summary statistics for price changes and price gaps before and after the inflation surge.

changes is 0.11 and the average frequency of price adjustment is $h = 0.29$, which implies that in a low-inflation environment, firms revise prices every 3 to 4 quarters on average. This is consistent with the evidence in Nakamura and Steinsson (2008) and Gagliardone et al. (2025) and with the quarterly frequency of price adjustment measured over the same time period in the official PPI microdata (0.275).[11]

Panel b reports the same statistics for the pandemic and post-pandemic period (2020–2023), marked by high inflationary pressures followed by a gradual tapering. Quarterly inflation in the manufacturing sector averaged 1 percentage point above the long-run trend, while the frequency of price changes rose by 10 percentage points, both in the NBB-BS and in the official PPI data. The fourth column reports the kurtosis of quarterly price changes (2.94 in normal times and 2.9 during the surge). Because kurtosis is sensitive to small changes in unit values, we compute it from the PPI microdata following the approach in Alvarez et al. (2022) to correct for biases introduced by product aggregation and unobserved heterogeneity.[12]

The last three columns of Table 1 report summary statistics for the empirical distribution of price gaps. This distribution, which is typically unobserved, is of great interest, as it reflects the history of shocks faced by firms and the inefficiencies arising

[11] See Appendix C for additional analysis and literature benchmarks of the frequency of price adjustment.

[12] The kurtosis without bias corrections is 3.26. Calibrating our model to this value requires menu costs above the range documented by Levy et al. (1997) and Zbaracki et al. (2004), and generates insufficient volatility in the frequency of price adjustment (Blanco et al. 2024a).

Figure 3: Empirical distribution of price gaps

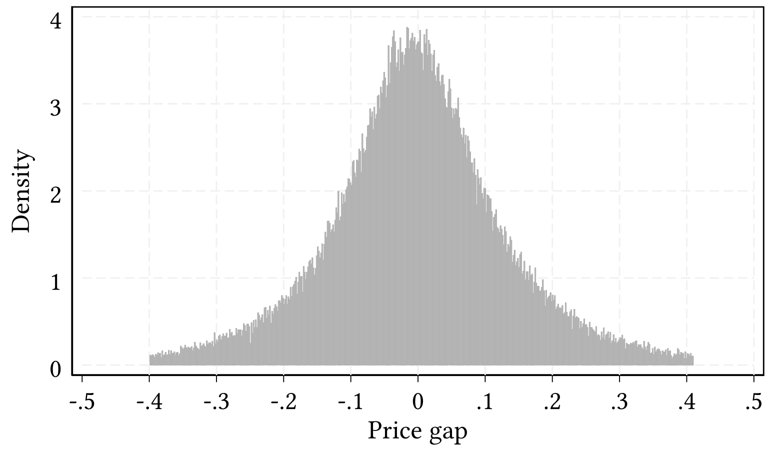

Notes. Empirical probability density function of price gaps for the pre-pandemic period (2000–2019).

from nominal rigidities. The histogram in Figure 3 plots the probability density function of price gaps in the pre-pandemic period. Consistent with our theory, the distribution is unimodal, bell-shaped, and approximately symmetric around zero. Table 1 reports that the average price gap rose by approximately 1 percentage point in the post-pandemic period. In line with model predictions, this increase corresponds closely to the average price change over the same period—a co-movement the model rationalizes through the pass-through of cost shocks into prices, as we show below.

4 State-dependent pricing in the cross-section of firms

Guided by the theoretical framework presented in Section 2, we design empirical tests of key model predictions that nonparametrically link micro-level pricing behavior to the underlying distribution of price gaps. These exercises provide direct evidence of the state-dependent nature of firms' pricing decisions—especially in the presence of large aggregate cost shocks—and yield identification results we use to discipline the calibration of structural parameters.

4.1 The empirical generalized hazard function

We begin by analyzing the mapping from price gaps to the frequency of price adjustment—*the extensive margin.* The following proposition establishes an empirical counterpart of the GHF in Lemma 1, characterizing the convex relationship between the two variables in the cross-section of firms.

Proposition 1. *Partition the price gap distribution into a finite number of narrowly defined bins $b \in \mathcal{B}$ with mass n_b, each corresponding to a quantile of the price gap distribution, ordered so that the lowest bins contain firms with the most negative price gaps. Denote by $x_b \equiv \frac{1}{n_b}\int_{f\in b} x_{t-1}(f)df$ and $\sigma_b^2 \equiv \frac{1}{n_b}\int_{f\in b}(x_{t-1}(f))^2 df - x_b^2$ the within-bin average and variance of price gaps.*
The average frequency of price adjustment in bin b is given by:

$$h_b \equiv \frac{1}{n_b}\int_{f\in b} h(x_{t-1}(f))df = (1-\theta^o) + \phi\Big(x_b^2 + \sigma_b^2\Big) + o(x_b^2). \tag{11}$$

The average frequency across bins in a neighborhood of $x_b = 0$ is a consistent estimator of the free-price adjustment probability $(1-\theta^o)$.

Proof. See Appendix A.3. □

The data are consistent with Proposition 1. In Figure 4, the black solid line depicts the kernel estimate of the empirical probability density function of price gaps. The red line plots the empirical GHF—the curve traced out by the fraction of firms adjusting prices within each bin (500 quantiles) of the price gap distribution. The black dotted line shows the fitted values from a cross-sectional regression of bin-level adjustment frequencies on the square of the bin-level average price gap.

Consistent with state dependence, firms with wider price gaps are more likely to adjust prices, and the relationship is convex: a quadratic polynomial captures well the convex pattern in adjustment frequencies. Proposition 1 further states that, in a neighborhood of $x_b = 0$ (the vertex of the GHF), the average frequency of price adjustment identifies the free-adjustment probability, which we estimate at $(1-\widehat{\theta}^o) = 0.188$. We leverage this identification result to calibrate the quantitative model in Section 6.

Two observations follow. First, the empirical GHF in Figure 4 should not be read as a direct estimate of the theoretical GHF—particularly its slope. As we show in Appendix D, the slope is subject to a downward bias because of measurement error in price gaps and

Figure 4: Empirical GHF and price gaps distribution

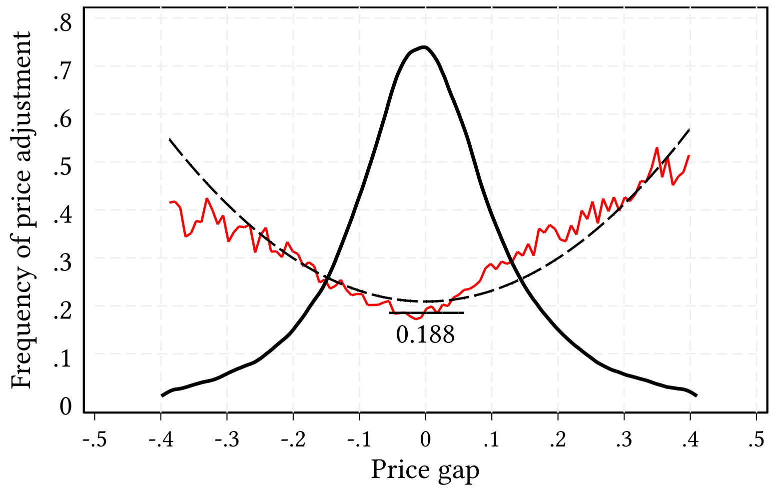

Notes. Empirical probability density function of price gaps, $f(x_b)$ (black line), alongside the empirical GHF, h_b (red line). The black dotted line shows the fitted values from a cross-sectional regression of the adjustment frequency in each bin (b) on a constant and the squared average price gap in the same bin. Each bin is weighted by its number of observations.

heterogeneity across firms or industries.[13] This means that the theoretical GHF should be steeper than Figure 4 suggests. We also show that measurement error might impact the estimate of the intercept. However, this bias is an order of magnitude smaller than the variance of gaps and upward, making our estimate conservative for calibration purposes (Section 6.1). Second, the empirical GHF is steeper for positive than for negative gaps, consistent with downward nominal rigidity (Karadi et al. 2024; Luo and Villar 2021): firms are more willing to raise prices when their markup falls below target than to cut prices when it rises above.

Appendix C constructs the price gap and adjustment frequency using PPI microdata. The PPI-based GHF is noisier—owing to the small sample size—but displays the same convex shape. Importantly, it yields an intercept close to the baseline, corroborating our calibration of the free-adjustment probability. Its slope is subject to downward bias, since measurement error and aggregation bias are present in both data sources.

[13]We are grateful to Erwan Gautier and Luminita Stevens for sharing useful insights on these points. Consistent with this prediction, the estimated slope is 2.56—flatter than those recovered in studies of individual retailers or homogeneous-product industries (Eichenbaum et al. 2011; Gautier et al. 2023), where measurement error is lower and the sample more homogeneous, as the bias result would predict.

Figure 5: Impact of aggregate cost shocks on the price gap distribution and average frequency of price adjustment

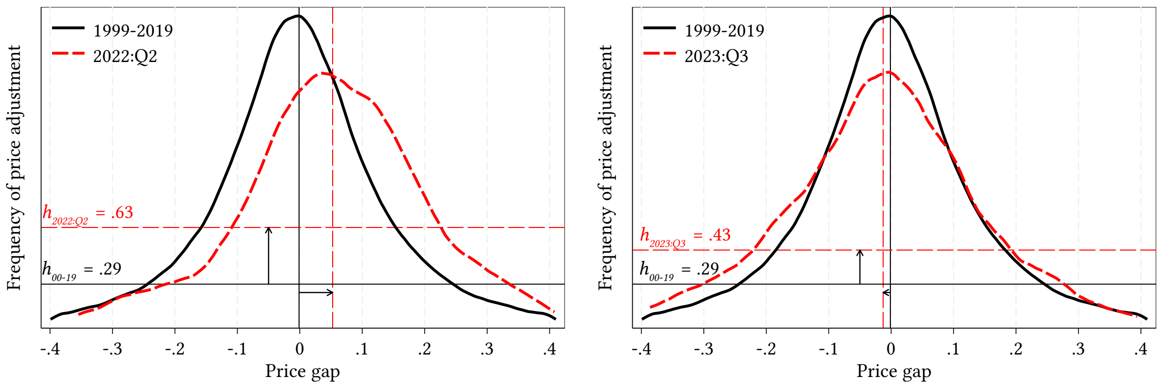

Notes. Empirical probability density function of price gaps during the pre-pandemic period (1999–2019, black solid line) and two post-pandemic snapshots: 2022:Q2 and 2023:Q3 (red dashed lines). The solid and dashed vertical lines mark the average price gap of the different distributions. The solid and dashed horizontal lines report the average frequency of price adjustment in the pre-pandemic period (black solid line) and in 2022:Q2 and 2023:Q3 (red dashed lines).

4.2 The impact of large aggregate shocks

In Section 2, we characterized how idiosyncratic shocks generate dispersion in price gaps while aggregate shocks shift the entire distribution, raising the fraction of firms that adjust prices (Figure 2). The pronounced increase and subsequent normalization of production costs in the post-pandemic period offer a sharp illustration of this mechanism.

In Figure 5, the solid black line shows the distribution of price gaps in the pre-pandemic period. In panel (a), we overlay the distribution in 2022:Q2 (red dashed line), a quarter in which firms' marginal costs increased by an average of 6.2 percent. Consistent with the model, the distribution shifts to the right, displacing many firms from their optimal price levels, and the variance of price gaps rises endogenously. The resulting compression in the average markup as well as the increase in markup dispersion raises the cost of inaction, leading to a sharp increase in the average frequency of price adjustment, from a pre-pandemic average of 0.29 to a peak of 0.63 in 2022:Q2. Panel (b) repeats the analysis for 2023:Q3, when firms' marginal costs declined by an average of 3.8 percent. Consistent with model predictions, the distribution shifts to the left, and adjustment frequency rises again—though by less than in the up-episode, reflecting both the smaller magnitude of the cost shock and the asymmetry of adjustment costs.

Figure 6: Price changes and price gaps, conditional on adjusting

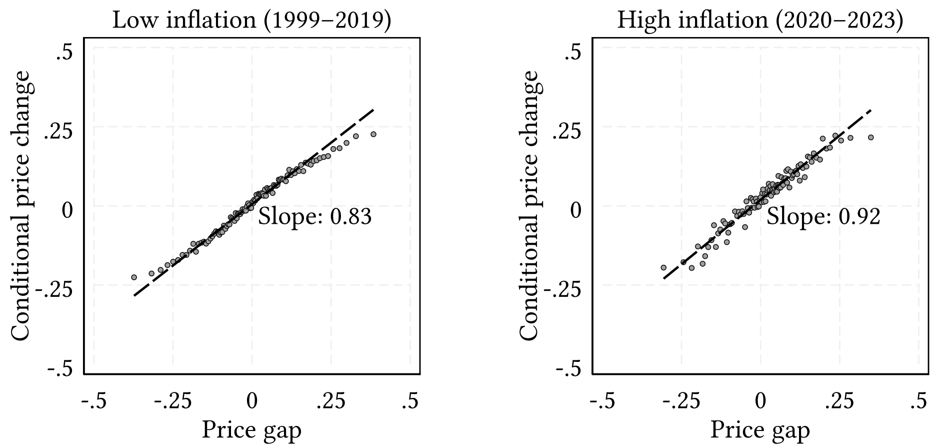

Notes. Binned scatterplot of log price changes against their price gaps focusing on the subsample of firms that adjusts their prices. Each dot marks the average price gap of a given percentile of the price gap distribution (x-axis) and the corresponding average percentage change in prices of firms in the same percentile (y-axis). The black dashed line depicts a linear fit of price changes on price gaps across the percentiles of the distribution of price gaps. The regression sample excludes the bottom and top 5 percentiles of the price gap distribution, to minimize the impact of outliers.

4.3 Price gaps and price changes conditional on adjustment

We now move to the mapping from price gaps to price changes, focusing on firms that adjust their prices—the *intensive margin*. Under our assumptions, $\Delta p_{ft} \approx x_{ft-1}$ if $\Delta p_{ft} \neq 0$, as formalized in Equation (9). It follows that, conditional on adjusting, the elasticity of price changes with respect to price gaps should be approximately one.

Testing this prediction, Figure 6 presents binned scatterplots of average price gaps (x-axis) against average price changes (y-axis) for each percentile of the price-gap distribution, focusing on the subset of firms that adjust prices. The left panel focuses on the pre-pandemic period, and the right panel covers the pandemic and post-pandemic period. In both panels, the black dashed line shows a linear fit of price changes on price gaps across percentiles. The data reveal that the gradient between price changes and price gaps is not only positive but also close to one—ranging from 0.8 to 0.9 depending on the subsample—as predicted by the theory. This finding suggests that, on average, the intensive-margin evidence corroborates the model's assumption that firm-level nominal marginal costs follow approximately random-walk dynamics. Note that this prediction is stated conditional on adjusting, so no correction for extensive-margin selection is needed to test it, though adjusters are of course selected by the hazard function.

A few factors contribute to explaining why the observed passthrough from price gaps to conditional price changes is lower than unity. First, classical measurement error in price gaps and deviations from random walk dynamics attenuates the estimated elasticity. Second, in the presence of strategic complementarities in price setting, approximating the dynamic reset price $p_t^{\star}(f)$ with the static target $p_t^{o}(f)$ introduces a downward bias. As shown in 2, nominal rigidities drive a positive wedge (a forward-looking premium) between the two price targets. Third, the model assumes complete information; to the extent that firms underreact to fundamentals due to inattention (Morales-Jiménez and Stevens 2024; Gagliardone and Tielens 2025), this would attenuate the measured elasticity further. We abstract from this channel and leave its quantification to future research.

4.4 The nonlinear passthrough from price gaps to price changes

We now document the combined cross-sectional mapping from price gaps to price changes, which integrates the extensive and intensive margins under the model's structure. The following proposition characterizes the (nonlinear) passthrough of price gaps into price changes as a function of the shock magnitude in the cross-section of firms.

Proposition 2. *Partition the price gap distribution into a finite number of narrowly defined bins $b \in \mathcal{B}$, each corresponding to a quantile of the price gap distribution, ordered so that the lowest bins contain firms with the most negative price gaps. Denote by x_b and σ_b^2 the within-bin average price gap and within-bin dispersion of gaps (defined in Proposition 1) and by $\gamma_b \equiv \frac{1}{n_b}\int_{f\in b}((x(f)-x_b)/\sigma_b)^3 df$ the skewness of price gaps within a bin.*
The inflation rate within a bin is given by:

$$\begin{aligned}\pi_b &\equiv \frac{1}{n_b}\int_{f\in b}\big(p_t(f)-p_{t-1}(f)\big)df = \frac{1}{n_b}\int_{f\in b}\big(h_t(f)\cdot x_{t-1}(f)\big)df \\ &= \phi_b^o \cdot x_b + \phi\cdot x_b^3 + \omega_b.\end{aligned} \tag{12}$$

where $\phi_b^o \equiv 1-\theta^o+3\phi\sigma_b^2$ is a bin-specific coefficient and $\omega_b \equiv \phi\gamma_b\sigma_b^3 + o(x_b^3)$ is small when the bins are sufficiently narrowly defined.

Proof. See Appendix A.3. □

Proposition 2 states that the inflation rate within a bin (i.e., the average within-bin price change) can be approximated by a third-order polynomial in the average price gap of

Figure 7: Nonlinear inflation dynamics

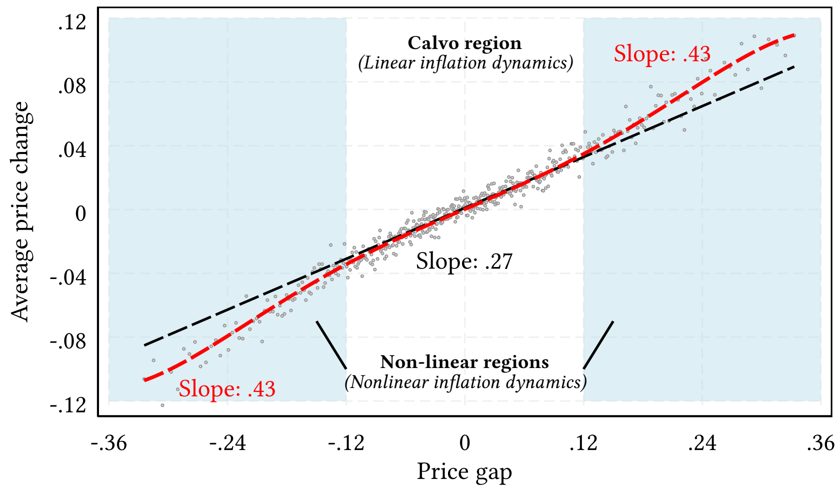

Notes. Binned scatterplot of the average price gap in each bin of the price gap distribution, x_b, against the corresponding average inflation (log price change), π_b. Bins are constructed using 500 quantiles. The black dashed line represents a linear fit of price changes on price gaps, $\widehat{\pi}_b = \widehat{a}_1 \cdot x_b$, estimated on the subsample of bins covering firms between the 25th and 75th percentiles of the price gap distribution, with the estimated slope ($\widehat{a}_1$) reported in black. The red dashed line represents the fit of a third-order (cubic) polynomial in price gaps, $\widehat{\pi}_b = \widehat{b}_1 \cdot x_b + \widehat{b}_2 \cdot x_b^3$, estimated using bins across the entire price gap distribution. The average slope of the third-order polynomial fit in the tails of the distribution (below the 25th and above the 75th percentiles) is reported in red.

the bin. The intuition is as follows. Inflation reflects the average of individual firms' price adjustments, which arise from the interaction between the extensive margin (the frequency of adjustment) and the intensive margin (the size of adjustment, conditional on adjusting).[14] As discussed above, the extensive margin can be approximated by a second-order polynomial in the price gap, while the intensive margin is approximately linear. Their interaction yields a cubic response of the average price changes, which captures both the linear passthrough observed in response to small shocks as well as the nonlinear response to large shocks.

When price gaps are close to zero, the third-order term becomes negligible, and the average price change is proportional to the average price gap. As a result, firms operating near their optimal prices exhibit near-linear pricing dynamics in both state- and time-dependent models. The mapping between price gaps and average price changes becomes nonlinear at the tails of the price gap distribution. Firms that are in these regions

[14]Note that, because bins are defined to be arbitrarily narrow, the aggregation error of the Kimball price index, $O(\sigma_t^2)$, does not show up in the formula for inflation within a bin.

experienced large shocks that moved them substantially away from their optimal markup. As a result, their frequency of adjustment increases more than proportionally with the size of the gap (Proposition 1) and their average price response is larger than proportional to the average price gap. Such nonlinearities are the hallmark of state dependence.

Figure 7 provides empirical support to the predictions of Proposition 2. We sort observations into quantiles (bins) spanning the price gap distribution. We chose 500 quantiles, but the results are robust to the number of bins. We then plot the average price change within each bin (π_b, y-axis) against the corresponding average price gap (x_b, x-axis). The red dashed line shows the fitted values from a cross-sectional polynomial regression of bin-level average inflation on bin-level average price gaps, estimated over the full support of the price gap distribution: $\pi_b = b_1 \cdot x_b + b_2 \cdot x_b^3 + \eta_b$. The black dashed line shows the fitted values from a linear regression of the same two variables, estimated on the interquartile range of the price gap distribution.

The relationship between the two is S-shaped—the shape predicted by the theory—and the cubic polynomial fits the full distribution well. We report the slope of the fitted polynomial—the reduced form passthrough from price gaps to (unconditional) price changes—for two distinct regions.

The first region comprises bins in the central part of the price gap distribution (25th to 75th percentiles), where deviations from optimal prices are relatively small. We refer to this range as the "*Calvo region*." There, the mapping between price gaps and expected price changes is approximately linear, consistent with the predictions of the Calvo model and characteristic of "normal times," when shocks are modest and inflation remains low and stable.[15] In this region, the estimated passthrough is 0.27—nearly identical to the average price adjustment frequency observed in the pre-pandemic period, as the Calvo model predicts. This result corroborates the theoretical results in Gertler and Leahy (2008), Alvarez et al. (2017), and Auclert et al. (2024), showing a near-equivalence of state-dependent and time-dependent pricing models when shocks are small.

The second region covers the tails of the price gap distribution—firms hit by shocks that generate a price gap of 20 percent or more in absolute value. As shown above,

[15]In a Calvo model, $\mathbb{E}[p_t(f) \mid \mathcal{I}_t(f)] = \mathbb{E}[p_t(f) \mid p_t^\star(f), p_{t-1}(f)] = (1-\theta^c)p_t^\star(f) + \theta^c p_{t-1}(f)$, where $\mathcal{I}_t(f)$ denotes the information set of a firm entering period t. Using the approximation $p_t^\star(f) \approx p_t^o(f)$ and rearranging, we obtain an analog of Equation (12) for a time-dependent Calvo model: a linear mapping $\pi_b = (1-\theta^c) \cdot x_b$, with an exogenous hazard rate $h^c := (1-\theta^c)$.

in this part of the distribution, the frequency of price adjustment increases more than proportionally with the size of the gap. Consequently, the sensitivity of price changes to price gaps rises sharply: the gradient steepens by nearly 60 percent, from 0.27 in the Calvo region to 0.43 in the tails. The cubic term in Proposition 2 captures these nonlinearities, which arise from the interaction between the extensive and intensive margins in response to large shocks.

5 Aggregate cost-price dynamics in the time-series

In this section, we shift focus from cross-sectional to aggregate cost-price dynamics in the time series. We show that state-dependent pricing at the firm level gives rise to nonlinear inflation dynamics in the aggregate, with the passthrough of cost shocks depending systematically on their magnitude.

5.1 Aggregate inflation and aggregate costs

We use our microdata to compute domestic producer price inflation and an index of nominal marginal costs for the Belgian manufacturing sector. Following the standard approach used by national statistical agencies, we calculate domestic PPI inflation as a Törnqvist price index, averaging firms' quarterly price changes weighted by smoothed sales shares:

$$\pi_t = \sum_{f \in \mathcal{F}} \bar{s}_{ft} \cdot \Delta p_{ft} \qquad \bar{s}_{ft} \equiv \frac{s_{ft} + s_{ft-1}}{2}.$$

Similarly, we construct an aggregate nominal cost index by aggregating and concatenating average changes in firm-level nominal marginal costs across producers (Δmc_t^n):

$$mc_t^n = \sum_{\tau=1999:Q2}^{t} \Delta mc_\tau^n$$
$$\Delta mc_t^n = \sum_{f \in \mathcal{F}} \bar{s}_{ft} \cdot \Delta mc_{ft}^n,$$

where the value of the index in the first quarter of our sample is normalized to zero.

According to our model, firms set prices based on current and anticipated movements in marginal costs. Accordingly, year-over-year inflation, $p_t - p_{t-4}$, should depend on nominal marginal cost at t, relative to the price level at $t-4$. The logarithmic

Figure 8: Inflation, cost, and markup dynamics

Panel a: Aggregate inflation and real marginal cost

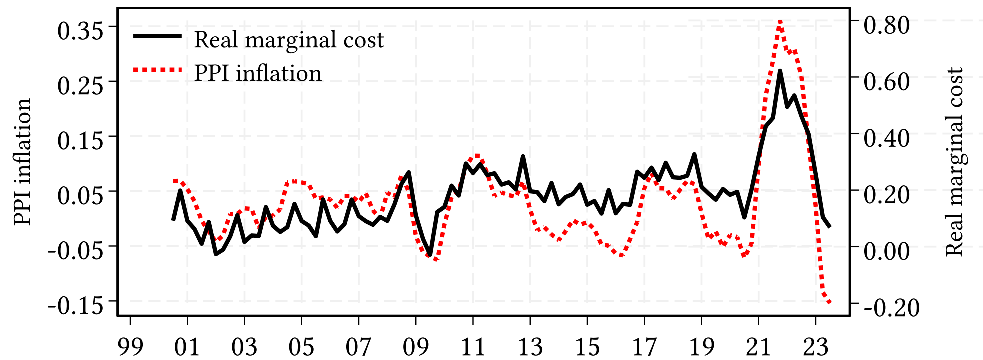

Panel b: Aggregate inflation and markup

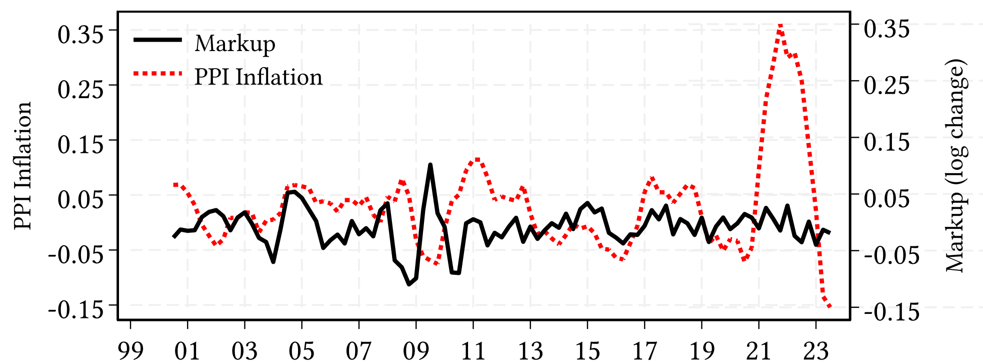

Notes. Time series of year-over-year PPI inflation ($p_t - p_{t-4}$), plotted alongside the time series of the real marginal cost index ($mc_t^n - p_{t-4}$, Panel a) and the log change in average realized markups ($\Delta \ln(\text{Markup}_t)$, Panel b) for the Belgian manufacturing sector.

difference between these variables, $mc_t^n - p_{t-4}$, provides a measure of real marginal cost. Panel a of Figure 8 plots the evolution of manufacturing PPI inflation (red dashed line, left axis) against the time series of real marginal costs (black line, right axis) throughout our sample period. Note that the scales of the two axes differ for the variables.

Two patterns emerge. First, consistent with the theory, inflation co-moves closely with real marginal costs throughout the sample, with less than one-for-one co-movement—consistent with nominal rigidities dampening the short-run passthrough of cost shocks into prices. Second, the sharp rise and subsequent fall in inflation during the pandemic period mirrors the large swings in real marginal costs.

To further illustrate the contribution of cost passthrough to inflation, panel b of

Figure 9: Decomposition of aggregate nominal marginal cost index

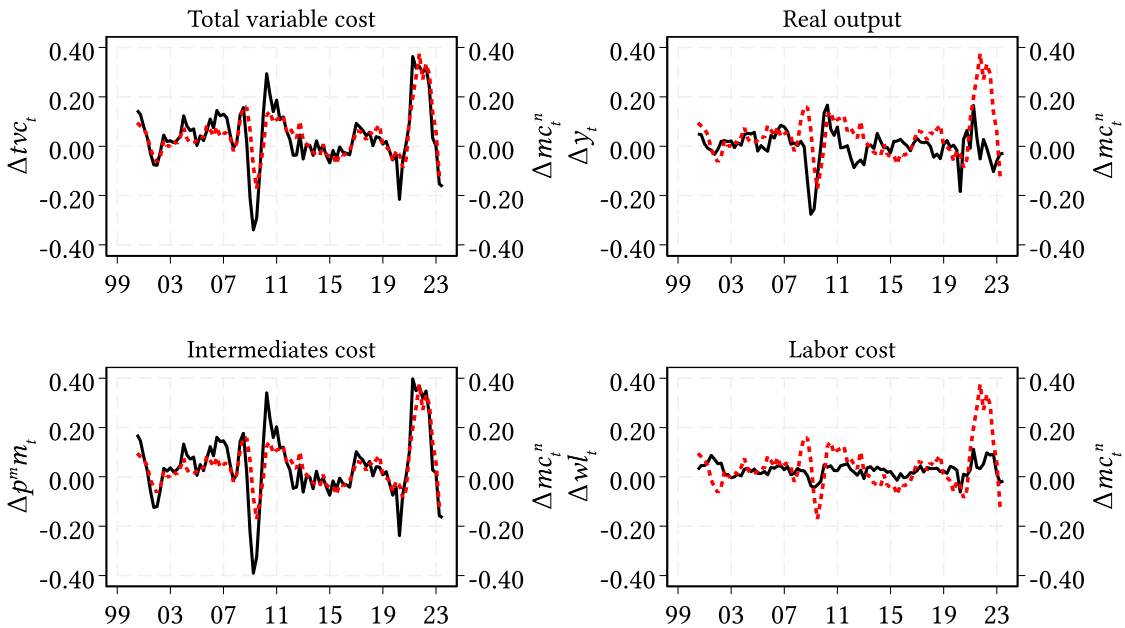

Notes. Decomposition of the log change in our nominal aggregate marginal cost index (Δmc_t^n) into the log change in total variable costs (Δtvc_t, top left panel), real output (Δy_t, top right panel), cost of intermediate goods ($\Delta p^m m_t$, bottom left panel), and labor costs (Δwl_t, bottom right panel).

Figure 8 contrasts inflation with changes in average realized markups, computed from the accounting identity $\Delta \ln(\text{Markup}_t) \equiv \pi_t - \Delta mc_t^n$. Restating the finding from panel a, and consistent with the evidence in Alvarez et al. (2024b), cost movements rather than markup expansion account for the bulk of the inflation surge.

Finally, to get a sense of what drove the fluctuations in nominal costs, Figure 9 presents a decomposition of our aggregate cost index into its different components. The top left panel plots the growth rates of total variable costs and real output (black lines), alongside the growth rate of the nominal marginal cost index (red dashed line). The two panels make clear that throughout the sample, and in particular during the recent inflation surge, fluctuations in total variable costs are the main drivers of the time series evolution of nominal marginal cost.

Figure 10: Passthrough of costs into inflation

Panel a: Nonlinear passthrough and frequency of price adjustment

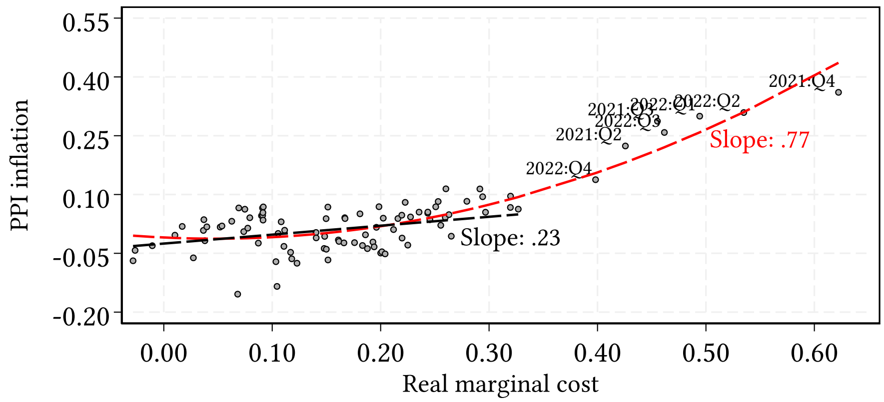

Panel b: Frequency of price adjustment and inflation

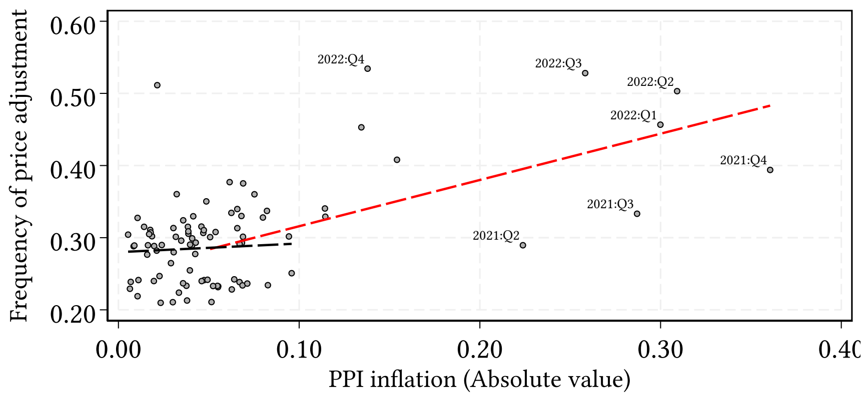

Notes. Dynamics of real marginal costs, aggregate inflation, and the average frequency of price adjustment. Panel a plots year-over-year manufacturing PPI inflation against our aggregate real marginal cost index, with quarters sorted by the latter. Panel b plots the average frequency of price adjustment against inflation, with quarters sorted by the latter. The average frequency of price adjustment is a rolling average of the quarterly frequency of price adjustment over the previous four quarters. In both panels, the black dashed line represents the linear fit of the variable on the y-axis based on the values of the variables on the x-axis during periods of low inflation (below 10 percent year-over-year); the red dashed line represents a fit across both high- and low-inflation periods.

5.2 Nonlinear aggregate passthrough

The two bottom panels of Figure 9 further decompose total variable costs into intermediate input costs (purchases of materials, services, and energy) and labor costs. As we can see, both cost components rose during the post-pandemic period. However, the increase in intermediate input costs was roughly four times larger. On average, this component alone accounts for approximately 70 percent of manufacturing firms' revenues. Moreover, over 80 percent of intermediate input costs are paid to foreign suppliers. These figures underscore that the surge in the cost of foreign-supplied intermediates—rather than labor costs—was the main driver of inflation between 2021 and 2023, at least in our sample.

Panel a of Figure 10 plots PPI inflation against the real marginal cost index, with quarters sorted by the latter. The black dashed line shows a linear fit estimated over low-inflation periods (below 10 percent year-over-year), while the red dashed line shows a quadratic fit estimated over the full sample. The slopes of these curves provide descriptive evidence on how aggregate cost passthrough varies with the magnitude of shocks to marginal cost.

Consistent with the micro-level dynamics in Section 4, we find a linear relationship between aggregate inflation and nominal costs during periods of low inflation. As discussed, a linear passthrough is consistent with the predictions of both the Calvo and menu-cost models when aggregate shocks are small. The estimated slope is 0.23, in line with the aggregate cost-passthrough estimated by Gagliardone et al. (2025) for a low-inflation environment.

During the recent surge, however, this linear relationship breaks down: the co-movement between inflation and costs rises sharply, revealing pronounced nonlinear cost-price dynamics. At the core of this nonlinearity is the endogenous movement of the average frequency of price adjustment. Panel b plots this variable against PPI manufacturing inflation. During low-inflation periods, the average frequency is essentially stable—consistent with the Calvo approximation. In contrast, during high-inflation periods inflation and the adjustment frequency co-move strongly (Alvarez et al. 2019; Cavallo et al. 2023; Blanco et al. 2024a).

6 Quantitative implications

We use the microdata to calibrate the menu-cost model presented in Section 2 and conduct two quantitative exercises. The first simulates the model's response to cost shocks of varying magnitudes and compares it to its nested Calvo limit. The second feeds the model a sequence of aggregate cost shocks extracted from the data and evaluates the empirical fit of the model-implied inflation dynamics.

6.1 Calibration

The calibration involves seven parameters. Four are set to standard values from the literature. We calibrate the elasticity of substitution between goods σ, to 6, implying a 20 percent markup in the symmetric steady-state equilibrium. The firm's risk-neutral discount factor, β, is set to 0.99. As in our empirical analysis, we calibrate $\Omega = 0.5$ to properly account for the strength of strategic complementarities in the microdata (Amiti et al. 2019; Gagliardone et al. 2025). To align model and data, we allow for a drift in the aggregate component of nominal marginal cost, $\mu_g = 0.005$, consistent with 1.6 percent annual trend inflation.

Calibration of the menu-cost model. We need to calibrate the parameters θ^o, σ_ϵ^2, and $\bar{\chi}$, which govern the degree of nominal rigidity and the state dependence of price adjustment. To do so, we leverage moments from the joint distribution of price changes and price gaps:

Step 1. We calibrate the free-adjustment probability to the average frequency of price adjustment around the intercept of the empirical GHF in Section 4.1, which yields $(1 - \hat{\theta}^o) = 0.188$. As we discuss in Appendix D, this parameter should be interpreted as a weighted average of free-adjustment probabilities across firms and industries—equivalently, the free-adjustment probability of a representative firm—which is the relevant calibration target in our single-sector framework. This calibration is conservative: measurement error may introduce a small upward bias—an order of magnitude smaller than the variance of gaps—that reduces the degree of state dependence of the model (see Appendix D).

Step 2. With low trend inflation and i.i.d. Gaussian idiosyncratic shocks, Alvarez et al. (2016) show that the following identity links the steady-state variance of shocks, σ_ε^2, to the average frequency and variance of price changes:

$$\sigma_\epsilon^2 = h \cdot \mathrm{Var}_{ss}(p_t(f) - p_{t-1}(f)).$$

Using this identity, we calibrate σ_ε^2 to 0.0036 to match the product of the average frequency and the variance of price changes reported in panel a of Table 1.

Step 3. Given σ_ε^2 and θ^o, we calibrate $\bar{\chi}$—the upper bound of the uniform distribution of menu costs—to 0.61 to match the frequency of price changes in the pre-pandemic period.

Table 2: Calibration: Data vs. model

	Price change ($p_{ft} - p_{ft-1}$)				Price gap (x_{ft-1})			Share MC
	Mean	Std.	Freq. Adj.	Kurt.	Mean	Std.	Kurt.	Mean (%)
Data	0.00	0.12	0.29	2.94	-0.00	0.13	2.86	1.22
Menu cost	0.00	0.12	0.29	2.62	0.00	0.09	3.30	1.70
Calvo	0.00	0.12	0.29	5.21	0.00	0.12	5.21	

Notes. Moments of the price-change and price-gap distributions for the data (2000–2019), the menu-cost model, and the Calvo model (both in steady state). The last column reports the average menu cost as a share of firm revenues; the data estimate is from Zbaracki et al. (2004).

Calibration of the Calvo model. In our quantitative exercises, we also consider a standard time-dependent Calvo model. As discussed in Section 2, our menu-cost model nests the Calvo model as a special case. We calibrate the Calvo model by setting the maximum menu cost, $\bar{\chi}$, to an arbitrarily large number and adjusting the free adjustment probability to match the steady-state frequency, $(1 - \theta^c = h)$.

Untargeted moments. Table 2 reports moments for both models alongside the data. The targeted moments are matched by construction. The untargeted moments—menu cost size and the fatness of tails—provide sharp validation benchmarks. Menu costs paid upon adjustment represent 1.7 percent of steady-state firm revenues, consistent with the estimates in Levy et al. (1997), Zbaracki et al. (2004), and Gautier and Le Bihan (2022). The calibrated model also displays a fatness of tails consistent with the data (kurtosis of 2.94 versus 2.62), in contrast to the Calvo model, which exhibits substantial leptokurtosis well above the empirical value.[16]

6.2 Impulse-responses to small and large aggregate shocks

We use our calibrated menu-cost and Calvo models to study price dynamics in response to both small and large shocks under state- and time-dependent pricing. Starting from an economy in steady state, we shock the system with permanent and unanticipated

[16]See Alvarez et al. (2016) for a theoretical treatment of the kurtosis of the price-change distribution as a sufficient statistic for monetary non-neutrality, and Alvarez et al. (2024a) for an empirical investigation.

Figure 11: Impact of aggregate cost shocks in state- and time-dependent models

Panel a: State-dependent pricing (Menu costs)

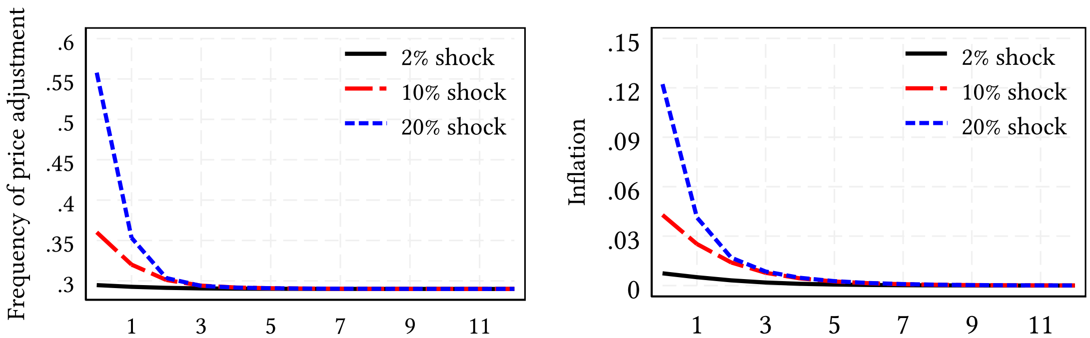

Panel b: Time-dependent pricing (Calvo)

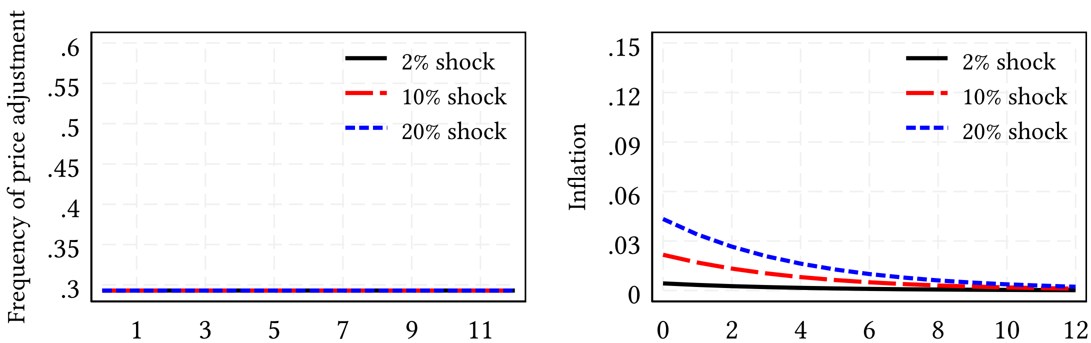

Notes. Impulse responses of inflation and the frequency of price adjustment to aggregate cost shocks of varying magnitudes. Panel a shows results for the menu-cost model (state-dependent), and Panel b for the Calvo model (time-dependent). The x-axis reports quarters since the shock.

aggregate cost shocks of different magnitudes, $g_t = 2$, 10, and 20 percent.

Figure 11 plots the impulse response functions (IRFs) of the frequency of price adjustment (left panels) and aggregate inflation (right panels) for both models. In all cases, the shocks raise the optimal reset price, shifting the distribution of price gaps to the right and triggering an inflation response. However, as discussed in Section 2 and shown empirically in Section 4, large shocks in the menu-cost model displace firms into regions of the price gap distribution where the GHF is steep. This leads to a spike in the adjustment frequency and, in turn, a sharper and faster rise in inflation relative to the Calvo benchmark.

Comparing IRFs across shocks of varying magnitudes highlights the nonlinearities inherent in state-dependent pricing, which become increasingly pronounced as shock size

Figure 12: Persistence of inflation in state- and time-dependent models

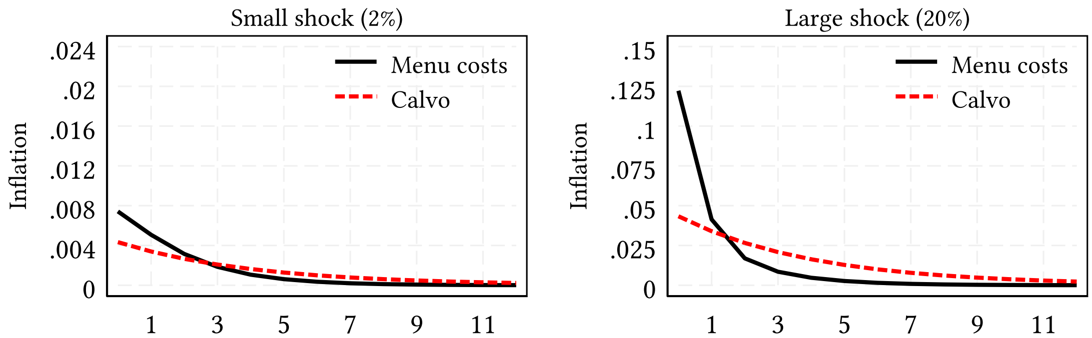

Notes. Inflation impulse responses to marginal cost shocks of different magnitudes under the menu-cost and Calvo models. The x-axis shows quarters since the shock.

grows. On impact, a 20-percent shock generates an inflation response and adjustment frequency nearly three times greater than that of a 10-percent shock, even though its magnitude is only twice as large. This convexity is robust to perturbations in $\bar{\chi}$, σ_{ε}^2, and Ω. In contrast, under the Calvo model, inflation rises proportionally with shock size.

A second key difference concerns the speed at which permanent cost shocks are absorbed into prices. Figure 12 overlays inflation IRFs from both models in response to small and large shocks. While passthrough dynamics are nearly identical for small shocks, the menu-cost model exhibits materially faster passthrough in response to large shocks. This differential response reflects the endogenous increase in adjustment frequency and the selection effect characteristic of state-dependent pricing.

Appendix B presents two additional quantitative exercises. In the first exercise, we examine how cost shocks of different magnitudes affect both the static target price, p_{ft}^{o}, and the dynamic optimal price, $p_{ft}^{\star}$. We show that the gap between the two prices is negligible for small cost shocks and remains modest even when the shock is larger.

Static and dynamic reset prices are more closely aligned in the menu-cost model than in the Calvo model, consistent with the model's assumptions and in line with empirical findings from Section 4, particularly in response to large shocks. As discussed in Section 2, this reflects the fact that, in the menu-cost model, the discount factor is time-varying and state-dependent, fluctuating with the current and expected probability of price adjustment. By contrast, the Calvo model features a discounting that is constant across any two consecutive periods.

The second exercise examines the role of strategic complementarities. We compare inflation dynamics under high and low cost shocks in models with and without strategic complementarities ($\Omega = 0$ vs. $\Omega = 0.5$). Strategic complementarities dampen passthrough in both modeling frameworks. However, due to the greater curvature of the value function under state dependence, the difference in IRFs with and without complementarities is smaller in the menu-cost model.

6.3 Explaining the time series of inflation

We now evaluate the ability of time- and state-dependent pricing models to account for the observed time series of aggregate inflation. In our quantitative exercise, we feed the models the sequence of aggregate cost shocks extracted from the data and simulate their implications for aggregate inflation and the frequency of price adjustment.

Assuming the economy is in steady state in 1999:Q1, we feed the model a shock to the aggregate component of nominal marginal cost implied by the change in our nominal cost index, Δmc_t^n, between 1999:Q1 and 1999:Q2. In doing so, we maintain the model's assumption that aggregate component of firms' marginal costs follows a random walk with drift. Given this shock, we solve the model and compute the updated distribution of price gaps, the resulting inflation response, and the frequency of price adjustment, assuming all future aggregate shocks are unanticipated, as in a standard impulse-response exercise. Using the updated price gap distribution as the new equilibrium, we repeat this procedure for all subsequent quarters through the end of our sample.

Figure 13 compares the model simulations to the data for three series: quarterly inflation, year-over-year inflation, and the quarterly frequency of price adjustment. Panels a and b show that the menu-cost model closely tracks the broad pattern of manufacturing inflation, during the low-inflation period pre-pandemic period and during the pandemic surge. During the pre-pandemic period, the menu-cost model is nearly indistinguishable from the Calvo model—as the theory predicts, since aggregate shocks are small and most firms' price gaps stay within the Calvo region (Section 4). By contrast, the Calvo model accounts for a modest portion of the inflation surge in the data (about two-thirds) and fails to replicate the rapid disinflation of late 2023, despite being fed the same cost sequence as the menu-cost model. In short, it responds too little and for too long.

Finally, Panel c compares the quarterly frequency of price adjustment in the model

Figure 13: Inflation and frequency of price adjustment: Model versus data

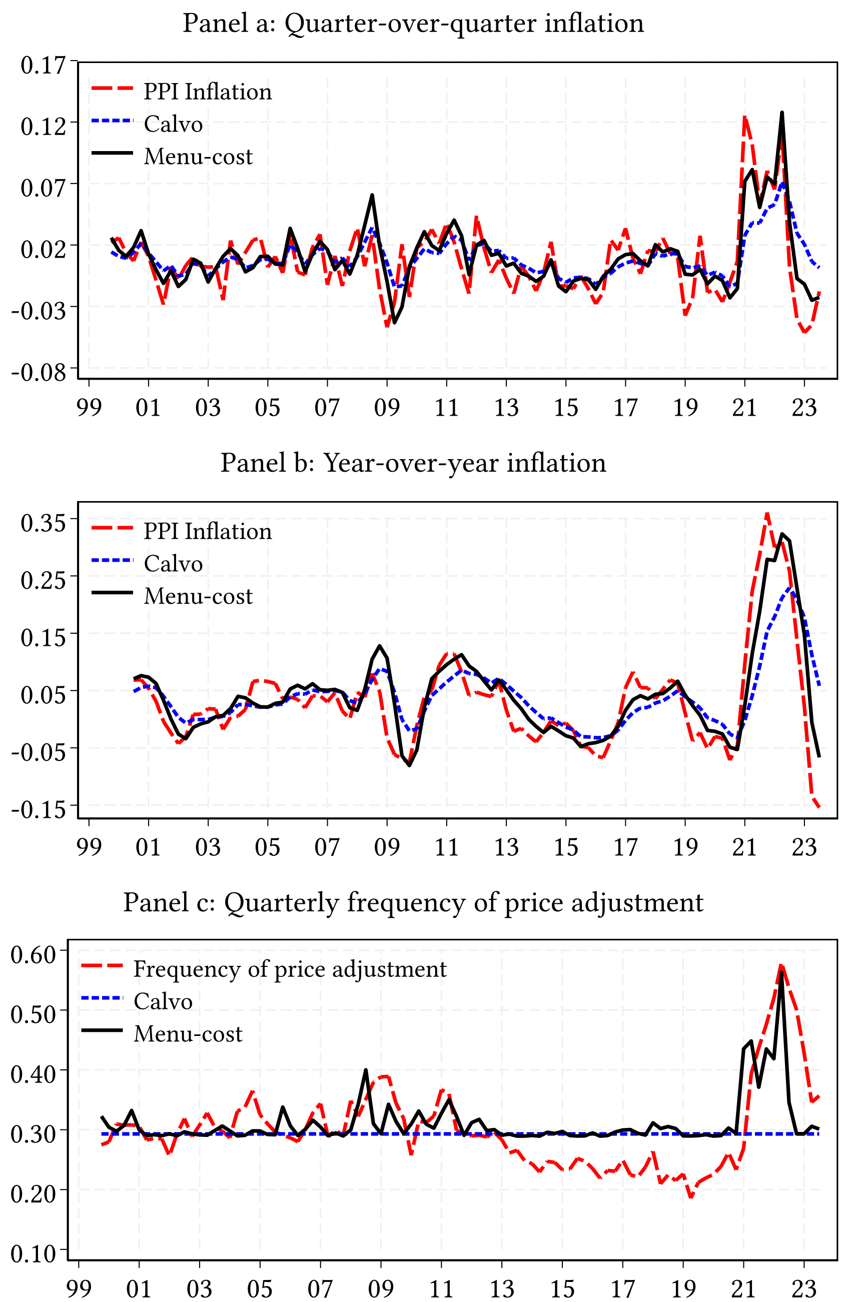

Notes. PPI manufacturing inflation and the frequency of price adjustment in the data (red dashed line), alongside simulated series from the Calvo model (blue dashed line) and the menu-cost model (black solid line). Both models are fed the same time series of aggregate cost shocks recovered from the data.

to that in the data. Our calibrated menu-cost model tracks the stable adjustment frequency observed in the pre-pandemic period—though it misses the gradual downward trend between 2012 and 2019—and captures the sharp increase in frequency following the onset of the pandemic, both in timing and magnitude. We note, however, that as inflation fell in the second half of 2023, the model-implied adjustment frequency fell more rapidly than in the data. This discrepancy may reflect firms anticipating the mean reversion in nominal marginal costs more accurately than implied by our random walk assumption.

7 Concluding remarks

Guided by theory, we use microdata on prices and costs to produce novel nonparametric evidence on the state-dependent nature of firms' pricing decisions, both in the cross-section of firms and in the aggregate time series. We then show how a menu cost model with strategic complementarities, calibrated to match our cross-sectional evidence, can explain the dynamics of both inflation and frequency at the aggregate level.

The use of cost data in addition to prices allows us to depart from the previous literature in two ways. First, instead of recovering the unobserved distribution of fundamentals from observed prices, we use observable cost shocks and competitors' prices to test the model's predictions. Consistent with state dependence, we document a nonlinear (cubic) passthrough from costs into prices, which captures firms' intensive and extensive margins of price adjustment. Second, we use cost data to explain the time series, in normal times and inflation surges, without relying on unobservable shocks.

We see two promising directions for future research. First, firm-level cost microdata of the kind we exploit are, to our knowledge, not yet available for services and retail. Yet these sectors also featured sharp increases in the frequency of price adjustment during the 2021–2023 surge (Blanco et al. 2024b; Cavallo et al. 2024; Morales-Jiménez and Stevens 2024; Dedola et al. 2025), consistent with a role for state-dependent pricing. As comparable cost microdata become available, testing whether the price-gap mechanism extends to services would be particularly valuable given the secular shift toward a service-based economy and the role services now play in aggregate inflation dynamics.

A second direction is to refine the modeling of marginal cost and its connection to real activity. The standard New Keynesian framework (e.g., Galí 2015), in its baseline form, models labor as the sole variable input, so that the labor share proxies for marginal cost. The microdata show, however, that intermediate input costs— primary commodities and energy—were the dominant source of marginal cost variation during the inflation surge. A joint study of price and cost microdata would offer a new lens for tracing how sectoral shocks pass through—possibly nonlinearly—to aggregate inflation.

References

Fernando Alvarez, Hervé Le Bihan, and Francesco Lippi. The real effects of monetary shocks in sticky price models: a sufficient statistic approach. *American Economic Review*, 106(10):2817–2851, 2016.

Fernando Alvarez, Francesco Lippi, and Juan Passadore. Are state-and time-dependent models really different? *NBER Macroeconomics Annual*, 31(1):379–457, 2017.

Fernando Alvarez, Martin Beraja, Martin Gonzalez-Rozada, and Pablo Andrés Neumeyer. From hyperinflation to stable prices: Argentina's evidence on menu cost models. *The Quarterly Journal of Economics*, 134(1):451–505, 2019.

Fernando Alvarez, Francesco Lippi, and Aleksei Oskolkov. The macroeconomics of sticky prices with generalized hazard functions. *The Quarterly Journal of Economics*, 137(2): 989–1038, 2022.

Fernando Alvarez, Francesco Lippi, and Panagiotis Souganidis. Price setting with strategic complementarities as a mean field game. *Econometrica*, 91(6):2005–2039, 2023.

Fernando Alvarez, Andrea Ferrara, Erwan Gautier, Hervé Le Bihan, and Francesco Lippi. Empirical investigation of a sufficient statistic for monetary shocks. *Review of Economic Studies*, page rdae082, 2024a.

Santiago Alvarez, Alberto Cavallo, Alexander MacKay, and Paolo Mengano. Markups and cost pass-through along the supply chain. *Unpublished manuscript, Harvard Business School*, 2024b.

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

Adrien Auclert, Rodolfo Rigato, Matthew Rognlie, and Ludwig Straub. New pricing models, same old phillips curves? *The Quarterly Journal of Economics*, 139(1):121–186, 2024.

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

Andrés Blanco, Corina Boar, Callum J Jones, and Virgiliu Midrigan. The inflation accelerator. Technical report, National Bureau of Economic Research, 2024b.

Philip Bunn, Lena Anayi, Nicholas Bloom, Paul Mizen, Gregory Thwaites, and Ivan Yotzov. How curvy is the phillips curve? Technical report, National Bureau of Economic Research, 2024.

Ricardo J Caballero and Eduardo MRA Engel. Microeconomic rigidities and aggregate price dynamics. *European Economic Review*, 37(4):697–711, 1993.

Ricardo J Caballero and Eduardo MRA Engel. Price stickiness in ss models: New interpretations of old results. *Journal of Monetary Economics*, 54:100–121, 2007.

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

Jeffrey R Campbell and Benjamin Eden. Rigid prices: Evidence from us scanner data. *International Economic Review*, 55(2):423–442, 2014.

Andrew Caplin and John Leahy. State-dependent pricing and the dynamics of money and output. *The Quarterly Journal of Economics*, 106(3):683–708, 1991.

Andrew Caplin and John Leahy. Aggregation and optimization with state-dependent pricing. *Econometrica: Journal of the Econometric Society*, pages 601–625, 1997.

Andrew S Caplin and Daniel F Spulber. Menu costs and the neutrality of money. *The Quarterly Journal of Economics*, 102(4):703–725, 1987.

Alberto Cavallo and Roberto Rigobon. The billion prices project: Using online prices for measurement and research. *Journal of Economic Perspectives*, 30(2):151–178, 2016.

Alberto Cavallo, Francesco Lippi, and Ken Miyahara. *Inflation and misallocation in new keynesian models*. 2023.

Alberto Cavallo, Francesco Lippi, and Ken Miyahara. Large shocks travel fast. *American Economic Review: Insights*, 2024.

David Cornille and Maarten Dossche. Some evidence on the adjustment of producer prices. *The Scandinavian Journal of Economics*, 110(3):489–518, 2008. doi: 10.1111/j.1467-9442.2008.00548.x.

Luca Dedola, Erwan Gautier, Chiara Osbat, and Sergio Santoro. Price stickiness in the euro area. *Research Handbook on Inflation*, pages 265–292, 2025.

Daniel A Dias, C Robalo Marques, and JMC Santos Silva. Time-or state-dependent price setting rules? evidence from micro data. *European Economic Review*, 51(7):1589–1613, 2007.

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

Michael Dotsey, Robert G King, and Alexander L Wolman. State-dependent pricing and the general equilibrium dynamics of money and output. *The Quarterly Journal of Economics*, 114(2):655–690, 1999.

Martin Eichenbaum, Nir Jaimovich, and Sergio Rebelo. Reference prices, costs, and nominal rigidities. *American Economic Review*, 101(1):234–262, 2011.

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

Luca Gagliardone and Mark Gertler. Oil prices, monetary policy and inflation surges. Technical report, National Bureau of Economic Research, 2023.

Luca Gagliardone and Joris Tielens. Dynamic pricing under information frictions: Evidence from firm-level subjective expectations. 2025.

Luca Gagliardone, Mark Gertler, Simone Lenzu, and Joris Tielens. Anatomy of the phillips curve: micro evidence and macro implications. *American Economic Review*, 115(11): 3941–3974, 2025.

Etienne Gagnon. Price setting during low and high inflation: Evidence from mexico. *The Quarterly Journal of Economics*, 124(3):1221–1263, 2009.

Etienne Gagnon, David López-Salido, and Nicolas Vincent. Individual price adjustment along the extensive margin. *NBER Macroeconomics Annual*, 27(1):235–281, 2013.

Jordi Galí. *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press, 2015.

Erwan Gautier and Hervé Le Bihan. Shocks versus menu costs: Patterns of price rigidity in an estimated multisector menu-cost model. *Review of Economics and Statistics*, 104 (4):668–685, 2022.

Erwan Gautier and Ronan Le Saout. The dynamics of gasoline prices: Evidence from daily french micro data. *Journal of Money, Credit and Banking*, 47(6):1063–1089, 2015.

Erwan Gautier, Magali Marx, and Paul Vertier. How do gasoline prices respond to a cost shock? *Journal of Political Economy Macroeconomics*, 1(4):707–741, 2023.

Mark Gertler and John Leahy. A phillips curve with an ss foundation. *Journal of Political Economy*, 116(3):533–572, 2008.

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

Jonathon Hazell and Stephan Hobler. Do deficits cause inflation? a high frequency narrative approach. Technical report, Centre for Macroeconomics (CFM), 2024.

Colin J Hottman, Stephen J Redding, and David E Weinstein. Quantifying the sources of firm heterogeneity. *The Quarterly Journal of Economics*, 131(3):1291–1364, 2016.

Peter Karadi, Raphael Schoenle, and Jesse Wursten. Price selection in the microdata. *Journal of Political Economy Macroeconomics*, 2(2):228–271, 2024.

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

Peter J Klenow and Oleksiy Kryvtsov. State-dependent or time-dependent pricing: Does it matter for recent us inflation? *The Quarterly Journal of Economics*, 123(3):863–904, 2008.

Daniel Levy, Mark Bergen, Shantanu Dutta, and Robert Venable. The magnitude of menu costs: direct evidence from large us supermarket chains. *The Quarterly Journal of Economics*, 112(3):791–824, 1997.

Shaowen Luo and Daniel Villar. The price adjustment hazard function: Evidence from high inflation periods. *Journal of Economic Dynamics and Control*, 130:104135, 2021.

Virgiliu Midrigan. Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica*, 79(4):1139–1180, 2011.

Camilo Morales-Jiménez and Luminita Stevens. Price rigidities in us business cycles, 2024.

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.

Emi Nakamura and Jón Steinsson. Monetary non-neutrality in a multisector menu cost model. *The Quarterly Journal of Economics*, 125(3):961–1013, 2010.

Emi Nakamura, Jón Steinsson, Patrick Sun, and Daniel Villar. The elusive costs of inflation: Price dispersion during the us great inflation. *The Quarterly Journal of Economics*, 133 (4):1933–1980, 2018.

John B Taylor. Aggregate dynamics and staggered contracts. *Journal of Political Economy*, 88(1):1–23, 1980.

Joseph Vavra. Inflation dynamics and time-varying volatility: New evidence and an ss interpretation. *The Quarterly Journal of Economics*, 129(1):215–258, 2014.

Michael Woodford. Information-constrained state-dependent pricing. *Journal of Monetary Economics*, 56:S100–S124, 2009.

Mark J Zbaracki, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen. Managerial and customer costs of price adjustment: direct evidence from industrial markets. *Review of Economics and statistics*, 86(2):514–533, 2004.

A Derivations and proofs

A.1 Derivation of the markup function

Assume that a perfectly competitive retailer assembles a bundle of intermediate inputs into a final product, Y_t. The bundle is a Kimball aggregator of differentiated goods produced by a continuum of producers (indexed by f):

$$\int_0^1 \Upsilon\left(\frac{Y_t(f)e^{\varphi_t(f)}}{Y_t}\right) df = 1,$$

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

Taking as given demand Y_t, each firm minimizes costs subject to the aggregate constraint:

$$\min_{Y_t(f)} \int_0^1 P_t(f)Y_t(f)df \qquad \text{s.t.} \int_0^1 \Upsilon\left(\frac{Y_t(f)e^{\varphi_t(f)}}{Y_t}\right) df = 1.$$

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

$$P_t(f) = \psi\Upsilon'\left(\frac{Y_t(f)e^{\varphi_t(f)}}{Y_t}\right)\frac{e^{\varphi_t(f)}}{Y_t}. \tag{A.1}$$

Define implicitly the industry price index P_t as:

$$\int_0^1 \phi\left(\Upsilon'(1)\frac{\widetilde{P}_t(f)}{P_t}\right) df = 1,$$

where $\widetilde{P}_t(f) \equiv P_t(f)e^{-\varphi_t(f)}$ and $\phi \equiv \Upsilon \circ (\Upsilon')^{-1}$. Evaluating the first-order condition (A.1) at symmetric prices, $\widetilde{P}_t(f) = P_t$, we get $\psi = \frac{P_tY_t}{\Upsilon'(1)}$. Replacing for ψ and denoting by

$\widetilde{Y}_t(f) \equiv Y_t(f)e^{\varphi_t(f)}$, we get the inverse demand function:

$$\frac{\widetilde{P}_t(f)}{P_t} = \frac{1}{\Upsilon'(1)}\Upsilon'\left(\frac{\widetilde{Y}_t(f)}{Y_t}\right). \tag{A.2}$$

Therefore, the demand function faced by firms when setting the (frictionless) price is:

$$\mathcal{D}_t(f) = e^{-\varphi_t(f)}(\Upsilon')^{-1}\left(\Upsilon'(1)\frac{\widetilde{P}^o_t(f)}{P_t}\right)Y_t. \tag{A.3}$$

Taking logs and differentiating, we obtain the following expression for the residual elasticity of demand:

$$\varsigma_t(f) \equiv -\frac{\partial \ln \mathcal{D}_t(f)}{\partial \ln \widetilde{P}^o_t(f)} = -\frac{\Upsilon'\left(\frac{\widetilde{Y}_t(f)}{Y_t}\right)}{\Upsilon''\left(\frac{\widetilde{Y}_t(f)}{Y_t}\right)\cdot\left(\frac{\widetilde{Y}_t(f)}{Y_t}\right)}. \tag{A.4}$$

We now use this result to derive the expression for the log-linearized desired markup. As above, we focus on the symmetric steady state. Denote the steady-state residual demand elasticity by $\varsigma = -\Upsilon'(1)/\Upsilon''(1)$. Then the derivative of the residual demand elasticity $\varsigma_t(f)$ in (A.4) with respect to $\widetilde{Y}_t(f)/Y_t$, evaluated at the steady state, is given by:

$$\varsigma' = \frac{\Upsilon'(1)\left(\Upsilon'''(1) + \Upsilon''(1)\right) - \left(\Upsilon''(1)\right)^2}{\left(\Upsilon''(1)\right)^2} \leq 0, \tag{A.5}$$

which holds with equality if the elasticity is constant (e.g., under CES preferences).

The desired markup is given by the Lerner index. Log-linearizing the Lerner index around the steady state and using Equation (A.5), we have that, up to a first-order approximation, the log-markup (in deviation from the steady state) is equal to:

$$\mu_t(f) - \mu(f) = \frac{\varsigma'}{\varsigma(\varsigma-1)}\left(y_t(f) - y_t + \varphi_t(f)\right).$$

Finally, log-linearizing the inverse demand function (A.2) and using it to replace the log difference in output, we obtain:

$$\mu_t(f) - \mu(f) = -\Gamma\left(p^o_t(f) - p_t - \varphi_t(f)\right)$$

where the sensitivity of the markup to the relative price is given by $\Gamma \equiv -\frac{\varsigma'}{\varsigma(\varsigma-1)}\frac{\Upsilon'(1)}{\Upsilon''(1)} \geq 0$. Let $\Omega \equiv \frac{\Gamma}{1+\Gamma}$. Replacing the log-linearized markup into the formula for the static optimal target price:

$$\begin{aligned} p^o_t(f) &= \mu_t(f) + mc^n_t(f) \\ &= (1-\Omega)(\mu(f) + mc^n_t(f)) + \Omega(p_t + \varphi_t(f)). \end{aligned}$$

A.2 Derivation of the optimal reset gap

The definition of intra-period nominal profits is the difference between revenues and costs:

$$\Pi_t(f) \equiv P_t(f)Y_t(f) - W_t e^{l_t(f)}$$

Define as before $\widetilde{P}_t(f) \equiv P_t(f)e^{-\varphi_t(f)}$. Replace the demand function (A.3) using market clearing ($Y_t(f) = \mathcal{D}_t(f)$):

$$\begin{aligned}\Pi_t(f) &= \left(P_t(f) - W_t e^{-z_t(f)}\right)e^{-\varphi_t(f)}(\Upsilon')^{-1}\left(\Upsilon'(1)\frac{\widetilde{P}_t(f)}{P_t}\right)Y_t \\ &= \left(\widetilde{P}_t(f) - W_t\right)(\Upsilon')^{-1}\left(\Upsilon'(1)\frac{\widetilde{P}_t(f)}{P_t}\right)Y_t.\end{aligned}$$

The second line uses the standard assumption that productivity cancels out with the taste shock, $z_t(f) = -\varphi_t(f)$ (Woodford 2009, Midrigan 2011, Alvarez et al. 2016). This assumption allows the problem to be described by a single scalar state variable. To see this, replace the definition of marginal cost:

$$\begin{aligned}p_t^o(f) &= (1-\Omega)(\mu(f) + mc_t^n(f)) + \Omega(p_t + \varphi_t(f)) \\ &= (1-\Omega)(\mu(f) + \ln(W_t)) + \Omega p_t + \varphi_t(f) \\ \Rightarrow p_t(f) - p_t^o(f) &= \widetilde{p}_t(f) - (1-\Omega)(\mu(f) + \ln(W_t)) - \Omega p_t\end{aligned}$$

This shows that the state variable of the problem can be taken to be $p_t(f) - p_t^o(f)$. As firms take aggregate prices as given, we can expand up to second order $\Pi_t(f)$ in $(p_t(f) - p_t^o(f))$ up to a term that does not depend on choices of the firm (Proposition 1 in Alvarez et al. 2023). Because no other firm-level state enters the profit function ($\mu(f) = \mu$ in the symmetric steady state), the maximization of profits pins down $x_t^\star(f) \equiv p_t^\star(f) - p_t^o(f)$ in terms of aggregate variables only. Therefore, $x_t^\star(f) = x_t^\star$ does not depend on f.

We now characterize how, under the quadratic profits, the optimal $x_t^\star$ depends on aggregate conditions. The intertemporal problem of the firm at time t is:

$$\max_x \ -B(x)^2 + \beta\,\mathbb{E}_t\left\{h_{t+1}(x)\cdot V_{t+1}^a + \left(1 - h_{t+1}(x)\right)\cdot V_{t+1}(x)\right\},$$

where $B \equiv \frac{\sigma(\sigma-1)}{2(1-\Omega)}$. The first-order condition evaluated at the optimal reset gap $x_t^\star$ is:

$$Bx_t^\star = \beta\mathbb{E}_t\left\{(1 - h_{t+1}(x_t^\star))\frac{\partial V_{t+1}(x)}{\partial x}\Big|_{x=x_t^\star} + (V_{t+1}^a - V_{t+1}(x_t^\star))\frac{\partial h_{t+1}(x)}{\partial x}\Big|_{x=x_t^\star}\right\}.$$

Because the adjustment probability is minimized at $x_t^\star$, the condition simplifies to:

$$Bx_t^\star = \beta \mathbb{E}_t \left\{ (1 - h_{t+1}(x_t^\star)) \frac{\partial V_{t+1}(x)}{\partial x} \Big|_{x=x_t^\star} \right\}.$$

Using that $\partial x_t / \partial x_{t-1} = 1$, the envelope condition is:

$$\frac{\partial V_{t+1}(x_t)}{\partial x_{t-1}} = -Bx_t + \beta \mathbb{E}_t (1 - h_{t+1}(x_t)) \frac{\partial V_{t+1}}{\partial x_t}.$$

Denoting $h_{t+\tau}(x_t) := h_{t+\tau}$ to ease notation, we repeatedly replace the derivative above into the first-order condition to obtain:

$$\mathbb{E}_t \left\{ \sum_{i=0}^{\infty} \beta^i \prod_{\tau=0}^{i} (1 - h_{t+\tau}) x_{t+\tau}^\star \right\} = 0, \quad h_t \equiv 0.$$

Rearranging the condition by using the random walk dynamics of $mc_t^n(f)$ and the fact that taste shocks are i.i.d., we obtain the following expression that characterizes the optimal reset price:[17]

$$\begin{aligned} p_t^\star(f) &= (1 - \Omega)(\mu(f) + mc_t^n(f)) + \Omega p_t + \Omega \frac{\mathbb{E}_t \{\sum_{i=1}^{\infty} (p_{t+i} - p_t) \beta^i \prod_{\tau=1}^{i} (1 - h_{t+\tau})\}}{\mathbb{E}_t \{\sum_{i=0}^{\infty} \beta^i \prod_{\tau=0}^{i} (1 - h_{t+\tau})\}} \\ &= p_t^o(f) + \Omega \Psi_t. \end{aligned} \tag{A.6}$$

The equation above decomposes the optimal (dynamic) reset price into two terms. The first is the static reset price, $p_t(f)^o \equiv (1 - \Omega)(\mu(f) + mc_t^n(f)) + \Omega p_t$, which captures the effects of current cost shocks and the price index. The second term, Ψ_t, captures the expected future dynamics of aggregate prices. These influence the optimal price $p_{ft}^\star$ as the firm anticipates that the price set today may also apply to future periods due to nominal rigidities.[18] Equation (A.6) also highlights that endogenous expected probability of price adjustment effectively imply that the firm optimizes with a time-varying discount factor. This force is not present in a (time-dependent) Calvo model, where the probability of price adjustment is constant and exogenous, $(1 - h_t) = \theta^c \ \forall t$, and $\Psi_t = (1 - \beta\theta) \sum_{i=1}^{\infty} (\beta\theta^c)^k \Omega (p_{t+k} - p_t)$.

Finally, under our assumption that costs follow a random walk and i.i.d. taste shocks, the second term in Equation (A.6) (i) does not depend on the identity of the firm, (ii) is exactly zero in the absence of strategic complementarities ($\Omega = 0$), and

[17] Here, we maintain the assumption that the exogenous aggregate shocks g_t are i.i.d. over time, drawn from a distribution with a symmetric, unimodal and continuously differentiable density with mean μ_g.

[18] See Dotsey and King (2005) for a discussion of the properties of the dynamic reset price under general assumptions about cost and demand dynamics.

(iii) is approximately zero even with strategic complementarities when trend inflation is sufficiently low ($p_{t+k} - p_t \approx 0 \ \forall k$). Properties (i)–(ii) are inherited by the optimal gap.

A.3 Proofs

Proof of Lemma 1. To derive the quadratic approximation of the hazard function in Equation (8), we take a second-order approximation of the hazard function $h_t(x_{t-1})$ characterized in Equation (7) around $x_t^\star$ to obtain:

$$\begin{aligned} h_t(x_{t-1}) &= (1-\theta^o) - \frac{\theta^o}{\bar{\chi}} \frac{\partial V_t(x)}{\partial x}\Big|_{x=x_t^\star} (x_{t-1} - x_t^\star) - \frac{\theta^o}{2\bar{\chi}} \frac{\partial^2 V_t(x)}{\partial x^2}\Big|_{x=x_t^\star} (x_{t-1} - x_t^\star)^2 + o(x_{t-1} - x_t^\star)^2 \\ &= (1-\theta^o) - \frac{\theta^o}{2\bar{\chi}} \frac{\partial^2 V_t(x)}{\partial x^2}\Big|_{x=x_t^\star} (x_{t-1})^2 + o\big((x_{t-1})^2\big), \end{aligned}$$

where the second equation follows from $\frac{\partial V_t(x)}{\partial x}\big|_{x=x_t^\star} = 0$ for a firm that is resetting its price and from our assumption that $x_t^\star \approx 0$. Assuming stationarity of the value function, and defining $\phi \equiv -\frac{\theta^o}{2\bar{\chi}} \frac{\partial^2 V(x)}{\partial x^2}\big|_{x=0}$, we have that the GHF can be approximated, up to second order, by a quadratic function of the price gap as in Equation (8):

$$h(x_{t-1}(f)) = (1-\theta^o) + \phi \cdot (x_{t-1}(f))^2 + o\big((x_{t-1}(f))^2\big), \tag{A.7}$$

where the parameter ϕ controls the sensitivity of the GHF to changes in gaps (i.e., the "steepness" of the parabola). □

Proof of Proposition 1. Partition the price gap distribution into a finite number of narrowly defined quantiles ("bins") $b \in \mathcal{B}$, with the lowest quantiles containing firms with the most negative gaps. We work under the assumption bins are sufficiently narrow so that the variance within a bin is smaller than the squared mean ($\sigma_b^2 \leq x_b^2$ for all $b \in \mathcal{B}$). This assumption holds under our definition of bins (Figure A.1, panel a and b).

Step 1. First, we derive the expression for the average frequency of price adjustment in a bin. Denote by $x_b \equiv \frac{1}{n_b} \int_{f\in b} x_{t-1}(f) df$ and $\sigma_b^2 \equiv \frac{1}{n_b} \int_{f\in b} (x_{t-1}(f))^2 df - x_b^2$ the within-bin average price gap and within-bin dispersion of gaps. The quadratic approximation (8) means that every $\epsilon/2 > 0$ there exists a $\delta_f > 0$ such that for $|x_f| \leq \delta_f$:

$$|h_f - (1-\theta^o) - \phi x_f^2| \leq \frac{\epsilon}{2} x_f^2.$$

Taking the average of both sides of the equation between observations falling within a given bin b (n_b is the mass of the bin b):

$$\begin{aligned} n_b \cdot |h_b - (1-\theta^o) - \phi(x_b^2 + \sigma_b^2)| &\leq \int_{f \in b} |h_f - (1-\theta^o) - \phi x_f^2| df \\ &\leq \frac{\epsilon}{2} \int_{f \in b} x_f^2 df \\ &= \frac{\epsilon}{2} n_b (x_b^2 + \sigma_b^2) \leq \epsilon n_b x_b^2, \end{aligned}$$

where the last inequality uses the fact that, when bins are arbitrarily small, $\sigma_b^2 \leq x_b^2$. Figure A.1 shows that this inequality holds in the data given our choice of quantiles. It follows that, letting $\delta_b \equiv \sup_{f \in b} \delta_f$, we get that:

$$h_b = (1-\theta^o) + \phi(x_b^2 + \sigma_b^2) + o(x_b^2).$$

Step 2. Next, we show how knowledge of the empirical frequency of price adjustment within a bin allows us to recover the probability of free price adjustment, θ^o. We denote by $\tilde{b}$ the bin such that $x_{\tilde{b}} = 0$. Label $\tilde{b} = 0$ and let $b' = -b'' \iff x_{b'} = -x_{b''}$. Let $h(b) \equiv h_b$ for all bs. As we have shown above, the frequency $h(b)$ is a convex function of the bins. Therefore, for any open interval of gaps around $\tilde{b}$ it holds that:

$$\int_{(-b,b)} h(b) db \geq (1-\theta^o) \int_{(-b,b)} db.$$

We want to show that the integral on the LHS converges to the RHS as the interval shrinks.

Let the bins take values on $b \in \{-\frac{1}{N}, -\frac{1}{N+1}, \ldots, 0, \ldots, \frac{1}{N+1}, \frac{1}{N}\}$ for some finite $N \in \mathbb{N}_+$. Consider a sequence of decreasing bounds indexed by $n \in \mathbb{N}_+$. Let

$$\bar{h}_n \equiv \frac{1}{|\mathcal{B}_n|} \sum_{b \in \mathcal{B}_n} h_b, \qquad \mathcal{B}_n \equiv \{b : |x_b| \leq \delta_n\}.$$

Then, the sequence $(1-\bar{h}_n)$ is non-decreasing (as h convex and integral is monotone in the support) and bounded above by θ^o. Therefore, by the monotone convergence theorem, it converges to its supremum which is given by:

$$\lim_{n \to \infty} \bar{h}_n = h(0) = 1 - \theta^o.$$

For a sufficiently small interval around the central interval, the mean of the frequencies of price adjustment over that interval is a consistent estimator of the probability of free adjustment. □

Proof of Proposition 2. Starting from the expression for aggregate inflation in Equation (2.3), we want to derive the cubic expression for inflation within a bin in Equation (12), under the assumption that $p_t^\star(f) = p_t^o(f)$ and associated cross-sectional regression model.

We again partition the distribution of price gaps into equal frequency bins (quantiles) denoted by b and adopt the same labeling convention of bins described in the proof of Proposition 1. Denote by $\gamma_b \equiv \frac{1}{n_b}\int_{f\in b}((x(f) - x_b)/\sigma_b)^3 df$ the skewness of price gaps within a bin b and define bins sufficiently small so that $|\gamma_b| \leq 1$ for all $b \in \mathcal{B}$. This assumption holds under our definition of bins (Figure A.1, panel c).

Consider a bin b in the positive range ($x_f > 0$ for all $f \in b$). Then for every $\epsilon/5 > 0$, there exists a $\delta_f > 0$ such that for $x_f \leq \delta_f$:

$$\begin{aligned}\frac{1}{n_b}\Big|\int_{f\in b}(h_f - (1-\theta^o) - \phi x_f^2)\cdot x_f df\Big| &\leq \frac{1}{n_b}\int_{f\in b}\big|h_f - (1-\theta^o) - \phi x_f^2\big|\cdot x_f df \\ &\leq \frac{\epsilon}{5}\frac{1}{n_b}\int_{f\in b} x_f^3 df \\ &= \frac{\epsilon}{5}\left(x_b^3 + 3\sigma_b^2 x_b + \gamma_b\sigma_b^3\right) \leq \epsilon x_b^3\end{aligned}$$

where the last step uses that bins can be chosen arbitrarily small (and the distribution of gaps is smooth) so that $\sigma_b^2 \leq x_b^2$ and $|\gamma_b| \leq 1$ and that $\gamma_b \geq 0 \iff x_b \geq 0$ because the distribution of gaps is single-peaked at zero. We note that the same argument applies for the negative range by switching signs of x_f (x_b) and γ_b and reversing the inequalities. Setting $\delta_b \equiv \sup_{f\in b}\delta_f$, for every $\epsilon/5 > 0$ it then holds that for $|x_f| \leq \delta_b$ for all $f \in b$:

$$\left|\frac{\mathbb{E}_b(h_f\cdot x_f) - (1-\theta^o)x_b - \phi\,\mathbb{E}_b(x_f^3)}{x_b^3}\right| \leq \epsilon.$$

Hence, $\mathbb{E}_b(h_f\cdot x_f) = (1-\theta^o)x_b + \phi\,\mathbb{E}_b(x_f^3) + o(x_b^3)$.
Using this approximation, the covariance within a bin satisfies:

$$\begin{aligned}Cov_b(h_f, x_f) &= \phi\,\mathbb{E}_b(x_f^3) + o(x_b^3) - \phi\,\mathbb{E}_b(x_f^2)x_b - o(x_b^2)x_b \\ &= \phi(2x_b\sigma_b^2 + \gamma_b\sigma_b^3) + o(x_b^3).\end{aligned}$$

It follows that inflation within a bin simplifies to:

$$\begin{aligned}\pi_b &= \frac{1}{n_b}\int_{f\in b} h(x(f))\,df\cdot\frac{1}{n_b}\int_{f\in b} x(f)\,df + Cov_b(h(f), x(f)) \\ &= \left((1-\theta^o) + 3\phi\sigma_b^2\right)x_b + \phi x_b^3 + \phi\gamma_b\sigma_b^3 + o(x_b^3).\end{aligned}$$

□

Figure A.1: Within bin variance, square of the mean, and skewness

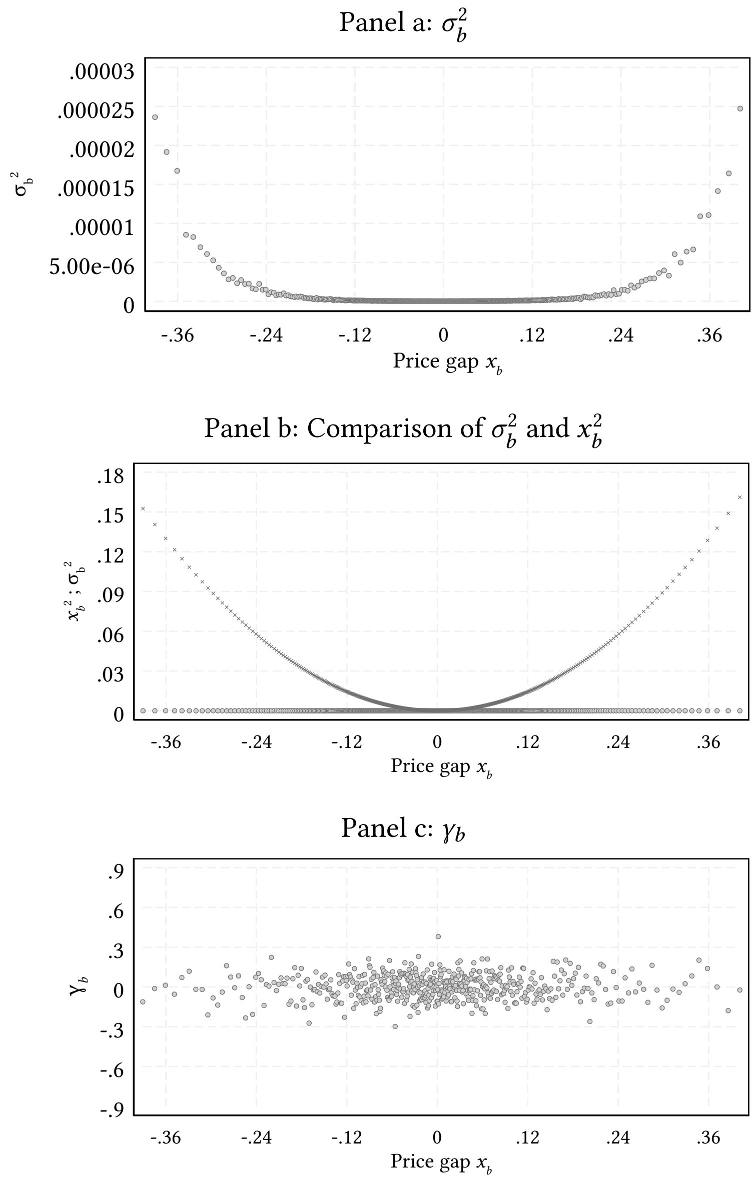

Notes. Three moments of the price gap distribution by bin (quantile): the within-bin variance of price gaps (σ_b^2, panel a), the square of the within-bin average gap (x_b^2, panel b), and the within-bin skewness of price gaps (γ_b, panel c). In panel b, x_b^2 (crosses) is plotted alongside σ_b^2 (circles) on the same scale to facilitate comparison.

Figure A.2: Impulse responses: Static vs. dynamic price targets

Panel a: State-dependent pricing (Menu costs)

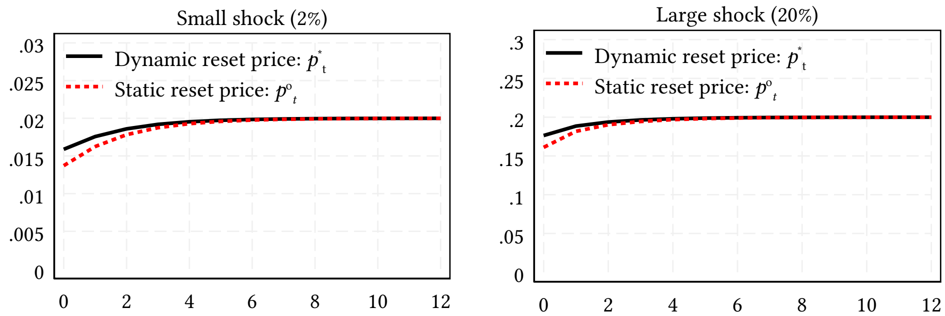

Panel b: Time-dependent pricing (Calvo)

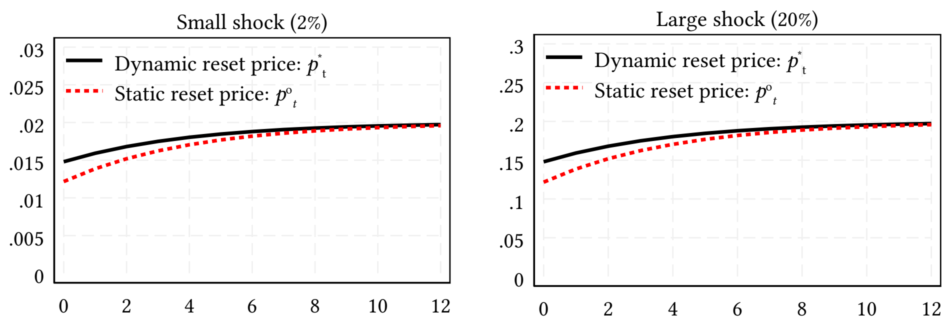

Notes. Impulse responses of the static reset price (p^o_t) and dynamic reset price ($p^\star_t$) to aggregate cost shocks of different magnitudes for the menu-cost and Calvo models. The x-axis denotes quarters since the shock.

B Additional quantitative exercises

Approximation of $p^\star_{ft}$ with p^o_{ft}. As discussed in Section 2, the two prices coincide in a steady state with zero trend inflation and constant markups. We also argued that the two prices remain sufficiently close to each other as long as trend inflation is not too large, even in the presence of strategic complementarities in pricing. We therefore assumed $p^o_{ft} \approx p^\star_{ft}$, which implies that $x^\star_{ft} \approx 0$, and derived expressions for aggregate inflation and within-bin inflation as a function of price gaps (Equations (2.3) and (12), respectively). The question is how well p^o_{ft} approximates $p^\star_{ft}$ away from the steady state.

The impulse response functions shown in Figure A.2 indicate that, as expected, the dynamic reset price responds more than the static one to cost shocks, since $p^\star_t$ accounts for marginal cost being a persistent process, though not a pure random walk, due to strategic pricing motives. However, this exercise also shows that the gap between the

Figure A.3: Quarter-over-quarter inflation: Static vs. dynamic price targets

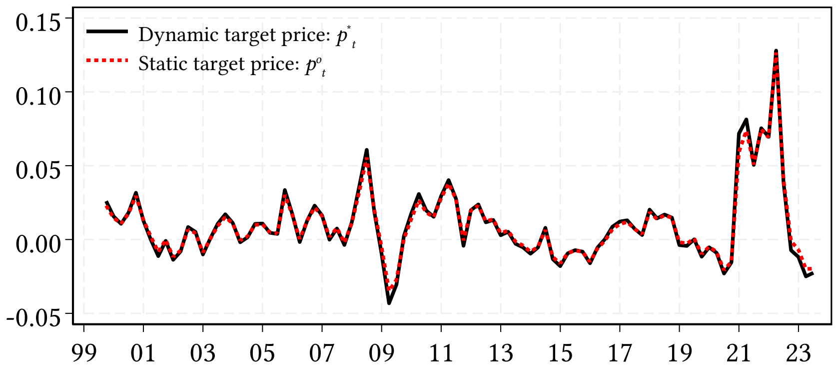

Notes. Inflation dynamics generated by the menu-cost model using $p^\star$ (the exact, dynamic reset price) and using p^o (the static approximation of $p^\star$) when solving the model. As in Figure 13, we solve the model feeding it a sequence of aggregate nominal marginal cost shocks that matched the one observed in the data.

two prices is negligible if the shock is small, as expected, and remains small even when the shock is large. Thus, the assumption that $p^o_{ft} \approx p^\star_{ft}$ is sensible. Additionally, this exercise demonstrates how the dynamics of the two prices are particularly close in the context of the menu-cost model relative to the Calvo model.

Next, we verify that using p^o_{ft} as an approximation for $p^\star_{ft}$ has a small impact on the aggregate inflation dynamics once we feed the model a sequence of aggregate nominal marginal cost shocks that matched the one observed in the data. Figure A.3 repeats the same quantitative exercise presented in Figure 13. The black line displays the time series of model-based quarterly inflation using $p^\star_{ft}$ as a measure of target price; the red dashed line displays the time series of model-based inflation, solving the model with p^o_{ft} as a proxy for $p^\star_{ft}$.

The role of strategic complementarities. Strategic complementarities in price setting help explain the differing dynamics of static and dynamic reset prices in time- and state-dependent models. Figure A.4 compares inflation dynamics after high- and low-cost shocks, without strategic complementarities ($\Omega = 0$) and with strategic complementarities ($\Omega = 0.5$). As before, panels a and b report the impulse response functions for the menu-cost model and the Calvo model, respectively. As expected, strategic complementarities generate additional discounting, which reduces cost passthrough in both models. However, we can see how the difference between the impulse-response

Figure A.4: The role of strategic complementarities

Panel a: State-dependent pricing (Menu costs)

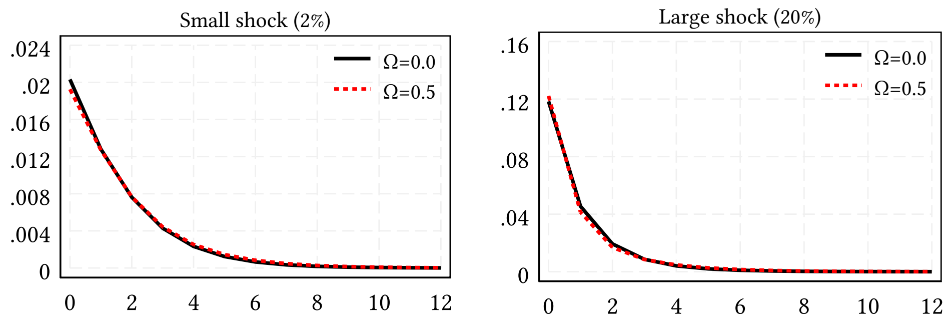

Panel b: Time-dependent pricing (Calvo)

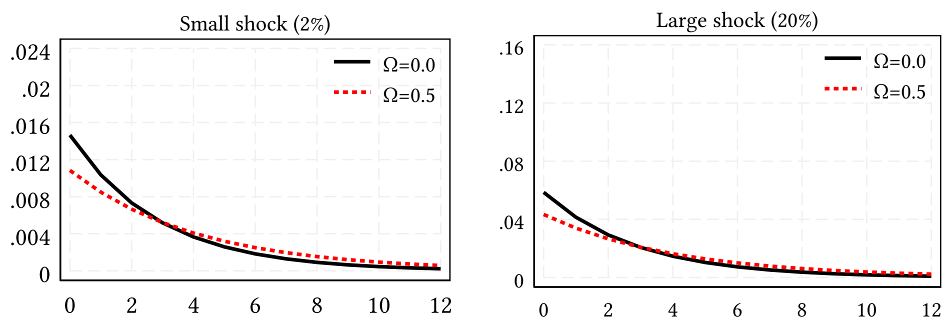

Notes. Inflation impulse responses to aggregate cost shocks under two pricing assumptions: no strategic complementarities ($\Omega = 0$, black line) and with complementarities ($\Omega = 0.5$, red dotted line) the menu-cost and Calvo models. The x-axis reports quarters since the shock.

functions with and without complementarities is narrower in the menu-cost model. This is due to the greater curvature of the value function under state-dependent pricing.

C Benchmarking the frequency of price adjustment

This appendix presents different robustness exercises using official PPI data. We first document the consistency between our baseline measure of the frequency of price adjustment and (i) a PPI-based quarterly frequency and (ii) monthly estimates from existing work on manufacturing.

Sample differences. The PPI sampling procedure is designed to achieve representative coverage of domestic manufacturing output for the purpose of constructing aggregate price indices. By contrast, PRODCOM is designed to track the near universe of Belgian

Figure A.5: Quarterly frequency of price adjustment in the PPI and in the NBB-BS data

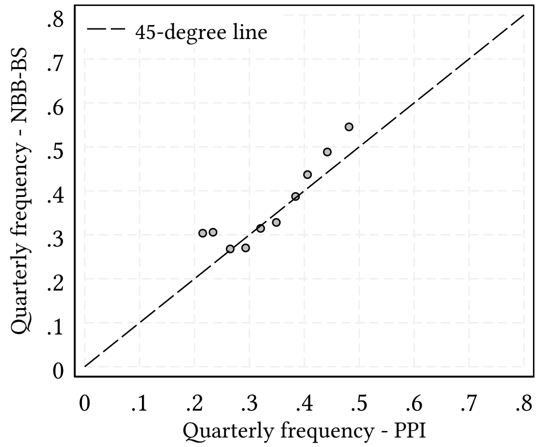

Notes. Binned scatter plot of quarterly frequency of price adjustment from the National Bank of Belgium Business Survey (y-axis) against the quarterly frequency from PPI microdata (x-axis).

manufacturing production: the survey covers at least 90 percent of domestic production value within each 4-digit NACE industry by including all firms with at least 20 employees or annual revenues above 4.5 million Euros.

As a result of the PPI limited cross-sectional coverage, the PPI-PRODCOM matched sample spans only 728 firm-industry pairs from 2003 onward, that is 15 percent of our baseline sample. Within each matched firm, we observe PPI price changes for a subset of products—typically those accounting for the largest share of sales. On average, PPI-covered products represent roughly 80 percent of firm-industry sales.

PPI-based quarterly frequency. During the period 2003–2019 the quarterly frequency of price adjustment in the manufacturing PPI—the share of products with a recorded price change in a given quarter relative to the previous quarter—is 0.275, within one-and-half percentage points of the NBB-BS baseline. The residual gap reflects differences in product and firm coverage across the two samples. Within the overlapping matched sample, the two series track each other closely in the time series (Figure A.5). The Pearson correlation coefficient is 67 percent for the full sample. During the inflation surge, the correlation coefficient raises to 90 percent: large common shocks dominate idiosyncratic noise, so both series move together strongly.

Reconciling monthly and quarterly rates. The monthly PPI frequency—the fraction of products with a price change relative to the previous month—is 19.5 percent over 1999–2019. If monthly adjustment decisions were independent across months, this rate would imply a quarterly frequency of $1 - (1 - 0.195)^3 = 47.8$ percent. The observed quarterly frequency is 27.5 percent, roughly twenty percentage points lower. The difference reflects within-quarter price reversals: a firm that adjusts its price in one month is less likely to adjust again within the same quarter, generating negative autocorrelation in month-to-month adjustment timing. Approximately 29 percent of firm-quarter observations that record no net quarterly price change contain at least one month-to-month reversal. The quarterly frequency of 0.275–0.290 thus implies, under the independence assumption, a monthly frequency of roughly 10–11 percent. As a reference point, Nakamura and Steinsson (2008) report a 10.8 percent monthly adjustment frequency for finished goods in US PPI data over 1998–2005, closely matching our implied monthly estimate.[19] The within-quarter reversal mechanism documented above reconciles the observed monthly frequency of 19.5 percent with the quarterly-implied rate of 10–11 percent: once reversals are accounted for, the two are consistent.

The empirical GHF based on PPI data. We reproduce the empirical GHF in Figure 4 using PPI data to measure both the frequency of price adjustment (y-axis) and price gaps (x-axis). The first one is directly observable in the PPI—whether a firm revised its price relative to the previous quarter. The latter requires us to construct firm(-industry) based price indices, p_{ft}, based on PPI. To do that, we follow the same steps described in Section 3, concatenating the (weighted) average price changes across different products of the same firm(-industry) surveyed in the PPI.

The resulting empirical GHF is shown in Figure A.6. Its shape is qualitatively similar to the baseline: the curve is U-shaped with an intercept of $(1 - \hat{\theta}^o) = 0.196$ and a slope of $\hat{\phi} = 2.41$, compared to $(1 - \hat{\theta}^o) = 0.188$ and $\hat{\phi} = 2.56$ obtained by fitting the GHF produced by the combined PRODCOM-NBB-BS baseline. As further discussed in Section D, the slope of the GHF based on PPI data possibly suffers from the same downward bias as the GHF in Figure A.6 due to measurement error in the price gaps and heterogeneity

[19] In earlier work, Cornille and Dossche (2008) report a higher average monthly frequency of approximately 24 percent for Belgian manufacturing over 2001–2005; the series shows a visible decline in adjustment frequency after 2005, which likely accounts for much of the difference.

Figure A.6: Empirical GHF and price gaps distribution

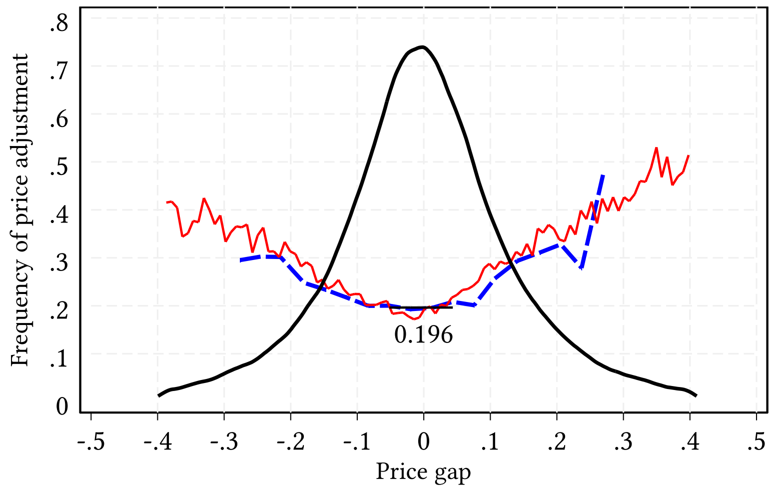

Notes. Empirical probability density function of price gaps, $f(x_b)$ (solid black line), alongside the empirical GHF, h_b (blue dashed line) using information on of price changes and price gaps constructed from the PPI microdata. As a reference, we also plot the empirical GHF computed using the NBB-PRODCOM as Figure 4 (red solid line). The black dotted line shows the fitted values from a cross-sectional regression of the adjustment frequency in each bin (b) on a constant and the squared average price gap in the same bin, as specified by Equation (8). Each bin is weighted by its number of observations.

across firms and industries.

A few measurement limitations are worth noting. First, the PPI provides microdata on product-level price changes but not on sales, which serve as aggregation weights. We address this by merging with PRODCOM to obtain product-level sales shares. When a product is observed in PRODCOM but lacks a corresponding PPI price, we rescale the remaining products' weights within the firm so that aggregation weights sum to one. As documented above, PPI-covered products account for roughly 80 percent of within-firm sales on average, so the rescaling adjusts weights over products representing only about 20 percent of each firm's revenue; any bias from non-random product selection within the PPI is unlikely to distort the shape of the GHF. Second, as discussed above, the PPI microdata covers roughly 15 percent of firm-industry pairs in our baseline sample, so we grouped firms in the PPI subsample into wider bins than in Figure 4 to reduce noise; the U-shape is not sensitive to the specific bin cutoffs used.

D Estimates of the empirical GHF with measurement error and heterogeneity

In this section, we discuss the robustness of the estimate of the intercept of the GHF (free-price adjustment) and a potential theoretical bias in the estimate of the slope. The discussion motivates our choice of calibrating the model to match the intercept but not the slope (Section 6.1). In particular, we consider the following generalizations of the baseline analysis of the model: (i) classical measurement error in price gaps; (ii) heterogeneity in firms' (and industries') primitives.

We work with the following data-generating process for the GHF:

$$h(x^*_{ft-1}) = (1 - \theta^o_f) + \phi_f (x^*_{ft-1})^2,$$

where x^*_{ft-1} is the latent firm-level gap. The latent gap is possibly different from the measured one, x_{ft-1}, because costs, output, and prices are measured with noise; $(1 - \theta^o_f)$ and ϕ_f are the free-adjustment probability and steepness of the GHF, which are possibly heterogeneous across firms.

Measurement error. Assume homogeneous coefficients and suppose that the "true" gap X^* is measured with error:

$$X^* \sim N(0, \sigma^2_X), \qquad \rho \sim N(0, \sigma^2_\rho), \qquad X = X^* + \rho, \qquad X^* \perp \rho$$

Recall that the probability of adjusting is:

$$\Pr(\mathbb{I} = 1 \mid X^* = x) = h(x) = (1 - \theta^o) + \phi x^2$$

We compute the estimator:

$$(1 - \hat{\theta}^o) = \lim_{\delta \to 0} \mathbb{E}[\mathbb{I} \mid |X| \leq \delta]$$

By the law of iterated expectations:

$$\begin{aligned}
\mathbb{E}[\mathbb{I} \mid |X| \leq \delta] &= \mathbb{E}\Big[\ \mathbb{E}\big[\ \mathbb{I} \mid X\big]\ \Big|\ |X| \leq \delta\Big] \\
&= \mathbb{E}\Big\{\ \mathbb{E}\Big[\ \mathbb{E}\big[\mathbb{I} \mid X^*\big]\ \Big|\ X = \bar{x}\Big] \mid |\bar{x}| \leq \delta\Big\} \\
&= \mathbb{E}\Big\{\ \mathbb{E}\Big[\ h(X^*)\ \mid\ X = \bar{x}\Big]\ \Big|\ |\bar{x}| \leq \delta\Big\} \\
&= (1 - \theta^o) + \phi\mathbb{E}\Big\{\ \mathbb{E}\Big[\ (X^*)^2\ \mid\ X = \bar{x}\Big]\ \Big|\ |\bar{x}| \leq \delta\Big\} \\
&= (1 - \theta^o) + \phi\,\mathbb{E}\left(\frac{\sigma_X^2\sigma_\rho^2}{\sigma_X^2 + \sigma_\rho^2} + \left(\frac{\sigma_X^2}{\sigma_X^2 + \sigma_\rho^2}\right)^2 \bar{x}^2\ \Big|\ |\bar{x}| \leq \delta\right)
\end{aligned}$$

Taking the limit for $\delta \longrightarrow 0$, we get:

$$(1 - \hat{\theta}^o) = (1 - \theta^o) + \phi\frac{\sigma_X^2\sigma_\rho^2}{\sigma_X^2 + \sigma_\rho^2}.$$

Therefore, the bias in the intercept from measurement error is upward and the calibration of the model is conservative (less state dependence —closer to the Calvo model). Second, using the same steps as above, the OLS estimate of the slope is:

$$\hat{\phi} = \left(\frac{\sigma_X^2}{\sigma_X^2 + \sigma_\rho^2}\right)^2 \phi.$$

Finally, the bias on the intercept is smaller than the bias of the slope if and only if:

$$\sigma_X^4 + \sigma_X^2\sigma_\rho^2 \leq 2\sigma_X^2 + \sigma_\rho^2,$$

which always holds because $\sigma_X^2 < 1$.

Cross-firm and cross-industry heterogeneity. Assume now that the intercepts are heterogeneous:

$$h(x_{ft-1}) = (1 - \theta_f^o) + \phi(x_{ft-1})^2.$$

In this case, the intercept recovers a weighted average of the heterogeneous coefficients, where the weights depend on the frequency at which each firm is at the optimum. Define $q_f(\bar{\epsilon}) \equiv \Pr(|x_{ft}| \leq \bar{\epsilon} \mid f)$, then:

$$1 - \hat{\theta}^o = \frac{\int (1 - \theta_f^o) q_f(\bar{\epsilon})\, dF(f)}{\int q_f(\bar{\epsilon})\, dF(f)}..$$

As firms with more free price adjustments are sampled more often, our estimator is always (weakly) larger than the cross-sectional mean of the heterogeneous coefficients. Note that this might change the interpretation of the parameter in our one-sector/homogeneous-firm model, but not its validity for the calibration.

Assume now that the slopes are heterogeneous:

$$h(x_{ft-1}) = (1 - \theta^o) + \phi_f (x_{ft-1})^2.$$

Heterogeneity in firms' primitives may flatten the empirical GHF through a selection channel. A high-ϕ_f firm adjusts aggressively whenever its gap drifts away from zero, so its ergodic gap distribution is tightly concentrated; a low-ϕ_f firm tolerates larger deviations before paying the menu cost, and its gap distribution is more diffuse. The central bins are therefore populated disproportionately by high-curvature firms whose own GHF is steep; the tail bins, by low-curvature firms whose own GHF is shallow.

Formally, let X_f^2 denote firm f's average squared gap. The bin-level slope estimate is $\text{Cov}(h_f, X_f^2)/\text{Var}(X_f^2)$. Expanding gives that the estimate of the slope of the GHF is $\hat{\phi} = \mathbb{E}[\phi_f] + \bar{X}^2\ \text{Cov}(\phi_f, X_f^2)/\text{Var}(X_f^2)$. By the argument above, $\text{Cov}(\phi_f, X_f^2) \leq 0$, and the slope of the empirical GHF lies below the average firm-level curvature, $\hat{\phi} \leq \mathbb{E}[\phi_f]$. The same selection logic applies across industries: pooling firms whose industries differ in the dispersion of equilibrium gaps adds a further compositional contribution to the same downward bias.