Micro and macro cost-price dynamics in normal times and during inflation surges

Luca Gagliardone* Mark Gertler† Simone Lenzu‡ Joris Tielens§

July, 2025

Abstract

We study cost-price dynamics in normal times and during inflation surges. Using microdata on firms' prices and production costs we construct an empirical measure of *price gaps*—the deviation between a firm's listed and optimal price. We then examine the mapping between gaps and price changes in the cross-section of firms and derive implications for inflation dynamics in the time-series. In the microdata, we provide nonparametric evidence on the nonlinear passthrough from costs to prices, which arises from the state-dependent nature of the pricing policies: firms are more likely to adjust prices as their price gap widens, a mechanism that becomes quantitatively significant when large aggregate cost shocks occur. In normal times, the frequency of price adjustment is approximately constant and the microdata conform with the predictions of time-dependent models (e.g., Calvo 1983). Conditional on a path of aggregate cost shocks extracted from the data, we show that a generalized state-dependent pricing model accounts well for the pre-pandemic era's low and stable inflation and the nonlinear surge observed during the pandemic.

* New York University. Email: luca.gagliardone@nyu.edu; † New York University and NBER. Email: mark.gertler@nyu.edu; ‡ New York University, Stern School of Business. Email: slenzu@stern.nyu.edu; § National Bank of Belgium and KU Leuven. Email: joris.tielens@nbb.be. We thank Al-Mahdi Ebsim for outstanding research assistance and Andrey Alexandrov, Fernando Alvarez, George-Marios Angeletos, Andres Blanco, Dan Cao, Jonathon Hazell, Peter Karadi, Francesco Lippi, Virgiliu Midrigan, Luigi Paciello, John Rust, Luminita Stevens, Ludwig Straub and seminar participants at various institutions, workshops, and conferences for helpful comments and conversations. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the National Bank of Belgium, the Eurosystem, or any other institution with which the authors are affiliated.

1 Introduction

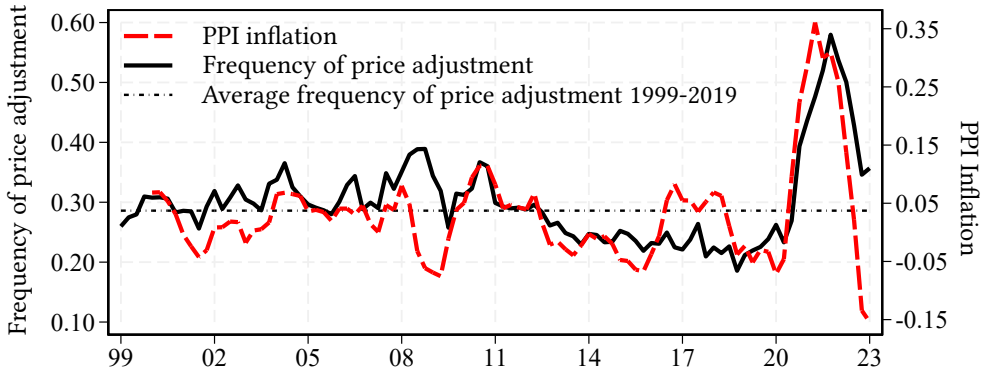
Firms adjust output prices infrequently despite continuously evolving economic conditions, leading their prices to drift temporarily from those that maximize flow profits. In all models featuring nominal rigidities, the deviation between listed and optimal prices—the *price gap*—determines firms' pricing 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. Given that firms set their price equal to the optimal one upon adjustment, in these models expected price changes are a linear function of the price gap. 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 the price gaps because both the change in prices conditional on adjustment and the adjustment frequency are endogenous functions of the gap. While the distinction between time- and state-dependent pricing is less significant when shocks to desired prices are small, it becomes important—and the nonlinearities apparent—when large aggregate shocks hit the economy.

The recent surge in inflation well illustrates these mechanisms. Figure 1 plots the year-over-year percentage change 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. However, as observed globally, starting in early 2021, the frequency increased sharply alongside a surge in inflation—hallmarks of state-dependent pricing and, as we will show, nonlinear cost-price dynamics.[1]

Recent contributions have demonstrated that state-dependent models can rationalize the inflation dynamics during the recent 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

[1] 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. This figure shows the time series of PPI manufacturing inflation along with the annual frequency of price adjustment. The former is computed as the year-over-year percentage change in the aggregate PPI. The latter is calculated as a rolling average of the quarterly frequency of price adjustment over the previous four quarters.

build (Nakamura and Steinsson 2010; Midrigan 2011; Alvarez et al. 2016), is to invert the model in order to discipline its structural parameters by fitting the observed distribution of price changes. Under some assumptions, it is also possible to recover the implied distribution of price gaps that rationalize the observed price changes (Alvarez et al. 2022). We emphasize that all these papers identify state-dependent behavior using *only* data on prices. In this paper, we follow a different approach, leveraging information on *both* prices and costs to construct an empirical measure of firm-level price gaps—the key object of our analysis—that allows us to study pricing policies in the micro data and their connection to aggregate inflation dynamics under different economic conditions.

We assemble a rich micro dataset of administrative records covering output prices, output quantities, and production costs at a quarterly frequency for Belgian manufacturing firms from 1999 to 2023. These data allow us to construct an empirical measure of firm-level price gaps that maps directly to its theoretical counterpart. We analyze the joint empirical distribution of price gaps and price changes through the lens of a tractable, state-dependent model—a menu-cost framework in the spirit of Nakamura and Steinsson (2010) and Auclert et al. (2024)—that nests a time-dependent Calvo (1983) model as a special case.[2] Using the model, we derive closed-form (algebraic) expressions characterizing price adjustments along both the extensive and intensive margins as

[2]Other seminal contributions include the works by Caplin and Spulber (1987), Caplin and Leahy (1991, 1997) and Dotsey et al. (1999) on state-dependent model and monetary neutrality.

a function of the firms' price gaps. By taking these expressions to the microdata, we can directly test the mapping from price gaps—and therefore the extent to which the predictions of different models conform with the data—under shocks of different magnitudes. In addition, the availability of microdata on production costs allows us to construct an aggregate cost index (and a corresponding sequence of aggregate cost shocks), which we feed into our model with the goal of accounting for the time series of aggregate inflation. Our analysis makes the following contributions.

Pass-through 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. This analysis allows us to produce the empirical analog of the generalized hazard function (GHF) postulated by the menu-cost model. Furthermore, by comparing the distribution of price gaps in pre-pandemic and pandemic periods, we empirically document how large aggregate shocks shift the entire price gap distribution, displacing many firms to the tails of the price gap distribution where the GHF is steeper. This shift results in a sharp rise in the average frequency of price adjustment, consistent with the evidence in Figure 1. Second, studying the intensive margin, we find that the data closely align with the model's prediction: firms that adjust their prices do so in a way that nearly closes the price gap, with an elasticity close to one. Specifically, a 10 percent price gap translates, on average, into a price change of 8 to 9 percent.

Finally, we derive an algebraic expression that combines the adjustments along the extensive and intensive margins and characterizes the pass-through for shocks of varying magnitudes. Our main proposition establishes that expected price changes can be approximated by a third-order (odd) polynomial in price gaps. Consistent with this prediction, we show that the relationship between price gaps and average price changes is S-shaped, allowing us to partition the cross-section of firms into two distinct regions. The first region, which we refer to as the "Calvo region," includes firms-quarter observations with small price gaps—those experiencing modest shocks to their desired prices. In this region, the frequency of price adjustment is approximately constant and close to its long-run (steady-state) average, and expected price changes increase linearly with the gap. Thus, consistent with the predictions of the time-dependent Calvo model, the elasticity of

average price changes with respect to price gaps is constant (roughly equal to the average frequency of adjustment), resulting in linear cost-price dynamics. The second region corresponds to the tails of the price gap distribution, where the occurrence of large shocks leads to large deviations from firms' optimal prices (approximately 20% or more). 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 price gaps is approximately twice as large as in the Calvo region.

Pass-through in the time series—The second set of results pertains to accounting for aggregate inflation in the time series. Leveraging our microdata, we construct a real marginal cost index for the Belgian manufacturing sector. Descriptive evidence reveals that aggregate inflation and production costs are strongly correlated throughout the sample period. However, due to price stickiness, inflation responds less than one-for-one to changes in costs. A simple decomposition exercise further reveals that the sharp rise and subsequent decline in the real marginal cost index (and therefore inflation) are driven primarily by movements in intermediate input costs, with labor costs playing a more limited role. Markups, by contrast, appear relatively stable throughout the inflation surge.

We use the results from our micro-level analysis and the constructed real cost index to evaluate the ability of our menu-cost model to account for aggregate inflation dynamics. We proceed in two steps. First, we develop a novel identification strategy to calibrate the key parameters governing nominal rigidities. In contrast to the standard calibration approach—which relies on model inversion and fits only the distribution of price changes—we exploit moments from the joint distribution of price changes and price gaps. Notably, the minimum of the GHF (achieved when firms have a price gap close to zero) provides a direct calibration of the free adjustment probability (the Calvo-plus parameter). We do not target the kurtosis of price adjustments, a theoretically informative moment that is, in practice, highly sensitive to issues related to unobserved heterogeneity and measurement error (Alvarez et al. 2016; Cavallo and Rigobon 2016). Nevertheless, our calibrated model reproduces a steady-state kurtosis close to the empirical value and, importantly, implies modest menu costs relative to firm revenues (about 1 to 2 percent), consistent with the empirical estimates from prior studies (Levy et al. 1997; Zbaracki et al. 2004).

Second, we compute model-based inflation series by feeding our marginal cost index into both the calibrated menu-cost model and a canonical Calvo model calibrated to match the long-run average frequency of price adjustment. We show that the menu-cost model closely tracks high-frequency fluctuations in Belgian manufacturing inflation—both during the stable pre-pandemic period and the post-pandemic surge. The model captures not only the subdued behavior of adjustment frequencies before the pandemic, but also the sharp increase in adjustment activity following its onset, both in timing and magnitude. By contrast, the 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 findings echo the conclusions of prior work highlighting that the distinction between time- and state-dependent pricing models is modest in the presence of small aggregate 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 robust evidence consistent with state-dependent pricing by studying micro-moments extracted from the empirical distribution of price changes (Klenow and Kryvtsov 2008; Gagnon 2009; Vavra 2014; Gautier and Saout 2015) and has shown that state-dependent pricing models can rationalize key features of these distributions (see the references cited earlier in this section). Direct empirical evidence of state dependence remains limited. Notable exceptions include Gagnon (2009) and Alvarez et al. (2019) who present evidence of state-dependent pricing by showing how inflation surges in Mexico and Argentina are associated with significant increases in the average frequency of price changes. The key distinction between these studies and ours lies in our ability to construct a high-frequency, firm-level measure of price gaps. This allows us to go beyond the distribution of price changes and directly test how pricing decisions relate to deviations from optimal prices. We derive algebraic expressions linking the size and frequency of price adjustment to price gaps, which we take to the data to directly test how firms' pricing strategies align with the predictions of different pricing models, conditional on the magnitude of the aggregate shocks. Our micro-based approach also yields novel identification results, allowing us to pin down key model parameters using a new moment derived from the distribution of

price changes conditional on price gaps.

Campbell and Eden (2014), Eichenbaum et al. (2011), Gagnon et al. (2013), 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) uses 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) uses data on retail and wholesale prices from a large food and drug retailer to construct a "reference price" metric. Similarly, 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 measure of price gaps captures deviations not only from competitors' prices but also from a firm's own cost structure. This yields a theoretically grounded measure that incorporates both supply-side fundamentals and strategic pricing considerations, enabling a direct micro-to-macro mapping. Second, the scope and representativeness of our dataset are substantially broader. We construct firm-level price gaps for a near-complete sample of domestic manufacturing producers in Belgium—covering 80–90 percent of the sector—over a 25-year period that includes both low-inflation regimes and inflationary episodes. This rich time-series variation enables us to study environments with stable, low adjustment frequencies as well as periods marked by sharp increases in price-changing behavior. Combined with the theoretical consistency of our price gap measure and broad cross-sectional coverage, it enables a systematic characterization of the nonlinearities in pass-through arising from shocks of varying magnitudes.

Finally, the results in this paper connect to Gagliardone et al. (Forthcoming), 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 non-parametric characterization of the nonlinear pass-through from cost to prices—which determines the slope of the NKPC—as a function of the magnitude of the cost shocks.

The paper proceeds as follows. Section 2 presents the theoretical framework and

derives testable implications. Section 4 describes our dataset and the empirical measures of prices, costs, and price gaps. Section 5 provides empirical evidence showing that the model predictions linking price adjustments with price gaps align with the microdata. We outline the calibration process and provide model simulations in Section 7, showing how the calibrated model explains the inflation time series and the frequency of price adjustment, including the rise during the pandemic. Section 8 offers concluding remarks.

2 Theoretical framework

Our baseline framework is a variation of a standard discrete-time menu-cost model. To fit the data, we allow for both random menu costs as in Caballero and Engel (2007) and random free price adjustment as in the "CalvoPlus" model of Nakamura and Steinsson (2010).[3] As is standard (Alvarez et al. 2023), we work with a quadratic approximation of the firm's profit function and permanent idiosyncratic shocks. In addition, motivated by our previous work (Gagliardone et al. Forthcoming), we allow for strategic complementarities in price setting. This framework nests a standard Calvo (1983) model as a special case.

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 (1995). Using lowercase letters to denote the logarithm of the corresponding uppercase variables, we denote by $p_t(f)$ the firm's price and by p_t the aggregate price index. Up to a first-order approximation around the symmetric steady state, the latter is given by:

$$p_t \approx \int_{[0,1]} \Big(p_t(f) - \varphi_t(f) \Big) df, \tag{1}$$

where $\varphi_t(f)$ denotes a firm-specific mean-zero log-taste shock, i.i.d. over firms and time.

Technology. Each firm operates with a constant return to scale production technology $y_t(f) = z_t(f) + l_t(f)$, which uses a composite input $l_t(f)$ and is characterized by total

[3] See also Dotsey et al. (1999) for a treatment of random menu-cost models with idiosyncratic shocks in a general equilibrium setting.

factor productivity $z_t(f)$. 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 captures an aggregate shock, i.i.d. over time with mean μ_g, drawn from a unimodal, symmetric, and smooth distribution. For analytical tractability, in what follows, we assume no trend inflation ($\mu_g = 0$).[4] We relax this assumption in the quantitative exercises of Section 7.

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 in its relative price $(p_t(f) - p_t)$, which makes the desired markup decrease in relative prices. As we show in Appendix A.2, 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.

Following Alvarez et al. (2023), we take a quadratic approximation of the per-period profit function around the static optimum $p_f(f) = p_t^o(f)$ and normalize it by steady-state profits. This yields the following loss function measuring the cost of deviations of the

[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 inflation on firms' decision rules is of second order.

price from the target:

$$\Pi_t(f) \approx -\frac{\sigma(\sigma-1)}{2(1-\Omega)}\left(p_t(f)-p_t^o(f)\right)^2,$$

where σ is the steady-state price elasticity of demand and steady-state profits are equal to $1/\sigma$. Note how the weight on the loss function is increasing in the complementarity parameter Ω. This is due to the fact that strategic complementarities increase the curvature of the profit function, and therefore raise the firm's desire to keep the price close to the target relative to the cost of adjustment.

Let $\mathbb{I}_t(f)$ be an indicator function that equals one if the firm adjusts its price and zero otherwise. Then, the value of the firm normalized by steady-state profits is given by:

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

2.2 Characterization of firms' pricing policies

As is standard in state-dependent models, the solution of the firm problem has an "Ss flavor." The optimal pricing policy reduces to determining an optimal probability of price adjustment and, conditional on adjustment, an optimal reset gap:

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

which captures the difference between the static target price and $p_t^\star$, the *(dynamic) reset price* set by a firm that decides to adjust its price.[5]

To characterize firms' 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 price gap captures the percentage difference between the static target price and the price set by the firm in the previous period. It serves as a state variable of the firm's problem because it is measurable before the firm decides whether to adjust its price (ergo, the "ex-ante"), but incorporates the realization of all time t shocks through their impact on $p_t^o(f)$. To see this, we can replace $p_t^o(f)$ using Equation (3) and $mc_t^n(f)$ using Equation (2) into Equation (5), and then use the expressions describing the processes for the aggregate

[5]Note that the optimal reset gap $x_t^\star$ varies over time due to aggregate shocks but it does not have an f subscript. This is because, to a first-order approximation, the idiosyncratic shocks, $\zeta_t(f)$ and $\varphi_t(f)$, enter both prices in an identical way and therefore cancel out once we take the difference. Important for this result is the assumption that idiosyncratic shocks evolve as a random walk.

and idiosyncratic components of $mc_t^n(f)$ to obtain:

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

where $\varepsilon_t(f) \equiv \frac{\Omega}{1-\Omega}\varphi_t(f) - \zeta_t(f)$ denotes 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.

We can characterize the solution to the firm's problem as a function of the price gap:[6]

$$p_t^o(f) - p_t(f) = \begin{cases} x_t^\star & \text{w. p.} \quad h_t(x_{t-1}) \\ x_{t-1}(f) & \text{w. p.} \quad 1 - h_t(x_{t-1}). \end{cases} \tag{6}$$

Here $h_t(x_{t-1})$ denotes the probability that a firm adjusts its price at t. Upon adjustment, they set their price to $p_t^\star(f)$. If they do not adjust their price, they keep their gap at $x_{t-1}(f)$.

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

We denote by $V_t(x_{t-1}(f))$ the value of a firm that decides not to revise its price and by V_t^a the value of a firm that decides to do so. The former is a function of current profits Π_t evaluated at the price gap $x_{t-1}(f)$ and of the discounted expected continuation value:

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

The latter is a function of the optimal reset gap, $V^a(x_t^\star)$, as we show below. The probability that a firm adjusts its price positively depends on 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:

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

[6] All derivations are provided in Appendix A.

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. Also, observe 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:*

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

where $o((x(f))^2)$ is an approximation error and $\phi \equiv -\frac{\theta^0}{2\bar{\chi}} \frac{\partial^2 V(x)}{\partial x^2}\Big|_{x=0}$.

Proof. See Appendix A.3. □

Lemma 1 states that, under a quadratic approximation, the GHF is U-shaped and symmetric around the point where the price gap is zero. At this point, the adjustment probability is at its global minimum, corresponding to the probability of a free price adjustment $(1-\theta^o)$. As price gaps widen, the adjustment probability monotonically increases. The parameter ϕ controls the sensitivity of the GHF to changes in gaps (i.e., the "steepness" of the parabola).

Optimal reset gap. The value of the firm conditional on adjusting is the optimized value of V in Equation (4) with respect to the reset price $p_t^\star$:

$$V_t^a = \max_{p_t^\star} V_t\big(x_t^\star\big).$$

Equivalently, the optimal reset gap $x_t^\star$ solves the first-order condition $V_t(x_t^\star) = 0$. Under low trend inflation, a quadratic approximation of the profit function, and assuming that the firm's marginal cost follows a random walk, the static optimal price provides a good approximation of the dynamic optimal price ($p_t^\star(f) \approx p_t^o(f)$), implying that $x_t^\star \approx 0$

(Alvarez et al. 2016). In Section 5, we provide micro-level evidence indicating that this approximation is broadly supported in the data.[7]

2.3 Aggregate inflation

Given the solution of the firm's problem in Equation (6) and using the formula for the price index in Equation (1), we can express aggregate inflation π_t as:

$$\begin{aligned}\pi_t &= \int \Big(p_t(f) - p_{t-1}(f)\Big)\, df = \int h_t(x_{t-1}(f)) \cdot \big(p_t^\star(f) - p_{t-1}(f)\big)\, df \\ &= \int h_t(x_{t-1}(f))\, df \cdot \int \Big(p_t^\star(f) - p_{t-1}(f)\Big)\, df + Cov\Big(h_t(x_{t-1}(f))\,,\ (p_t^\star(f) - p_{t-1}(f))\Big) \quad (9)\end{aligned}$$

The first line that aggregate inflation—average of firm-level price adjustments—can be expressed, to a first-order approximation, as the product of a firm's adjustment probability and its price change conditional on adjustment. The second line decomposes inflation into (i) the product of the average frequency of price adjustment and the average distance between the ideal reset price and the previous period price and (ii) the covariance between the variables.

With state-dependent pricing, the adjustment frequency is an endogenous object that, as we will see, increases nonlinearly with the absolute value of the price gap. With Calvo pricing, the adjustment frequency is orthogonal to the price gap and therefore uniform across firms and independent of the size of shocks; the price adjustment is a linear function of the price gap, and inflation is equal to the product of the constant adjustment frequency and the average price gap.

As in Caballero and Engel (2007), Golosov and Lucas (2007) and, more recently, Karadi et al. (2024), a "selection effect" increases the degree of monetary neutrality in the economy with large shocks. This force is captured by the covariance term in Equation (9). Firms that are more likely to adjust are also those that change their prices the most (conditional on adjustment). That is, the gap between $p_t^\star(f)$ and $p_{t-1}(f)$ positively co-moves with x_{t-1} and therefore with $h_t(x_{t-1})$. Thus, the selection effect positively contributes to generating aggregate inflation.

[7] Appendix B present quantitative exercises analyzing the response of static and dynamic reset prices to shocks of varying magnitudes.

3 Testable implications and identification results

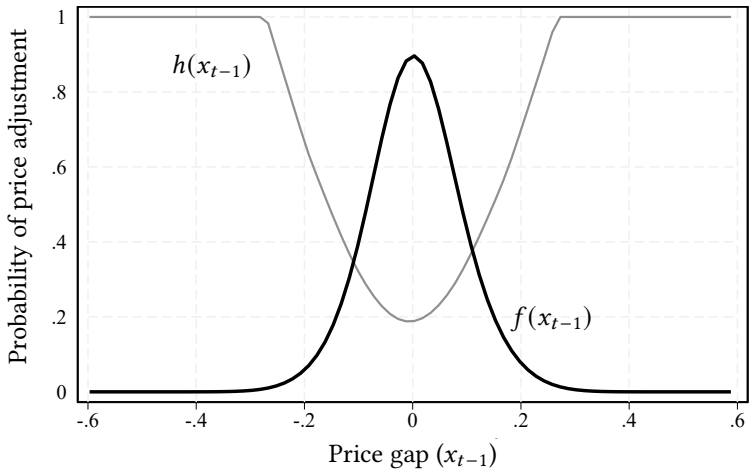
We now derive algebraic expressions that nonparametrically characterize the mapping from price gaps to firms' pricing decisions along the extensive margin, the intensive margin, and their interaction. This framework yields testable implications that we use to validate the model and its underlying assumptions (Section 5), and to derive novel identification results that inform the calibration of the structural parameters (Section 7).

3.1 The generalized hazard function

Figure 2 illustrates the relationship between price gaps and the extensive margin of price adjustments in the stationary equilibrium. The horizontal axis depicts the steady-state probability density function of price gaps, denoted by $f(x_{t-1})$, which is unimodal and bell-shaped. The vertical axis shows the GHF, $h_t(x_{t-1})$, the probability of price adjustment at different points in the price gap distribution. Under our assumptions, it is U-shaped and centered around $x_{t-1} = 0$. As discussed above, the GHF reaches its minimum precisely at this point, $h(0) = (1-\theta^0)$, matching the free adjustment probability. Given the realization of the shocks affecting $p_t^o(f)$, firms in the right (left) tail of the price gap distribution operate with a suboptimally low (high) markup and are therefore more likely to increase (decrease) their price relative to the price they previously changed.

We derive an expression relating the frequency of price adjustment—the measurable counterpart of the GHF—to a second-order (even) polynomial in the price gaps. Formally, we partition the price gap distribution into a countable set of quantiles, with the lowest quantiles containing firms with the most negative gaps. We refer to these quantiles as "bins" and denote them by $b \in \mathcal{B}$. Within each bin, we compute the average price gap, $x_b \equiv \int_{f\in b} x_{t-1}(f)df$, and the dispersion of gaps within the bin, $\sigma_b^2 \equiv \int_{f\in b}(x_{t-1}(f))^2 df - x_b^2$. 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 $b \in \mathcal{B}$). We show in Appendix A.6 that this assumption is satisfied by our definition of bins. The next proposition derives the average frequency of price adjustment for a bin and shows how to use it to identify the frequency of free price changes using microdata.

Figure 2: Generalized hazard function and distribution of price gaps



Notes. This figure shows the generalized hazard function (GHF), $h(x_{t-1})$ as a function of the underlying distribution of price gaps $f(x_{t-1})$.

Proposition 1. *The average frequency of price adjustment in bin b is given by:*

$$h_b \equiv \int_{f \in b} h(x_{t-1}(f))df = (1-\theta^o) + \phi\left(x_b^2 + \sigma_b^2\right) + o(x_b^2). \tag{10}$$

Restricting the estimation sample to bins in a sufficiently small neighborhood around $x_b = 0$, the estimator $\hat{a}_0$ of the following cross-sectional regression model recovers the probability of free price adjustment $(1-\theta^o)$:

$$h_b = a_0 + a_1 \cdot x_b^2 + u_b, \tag{11}$$

where $u_b \equiv \phi\sigma_b^2 + o(x_b^2)$ denotes the error term.

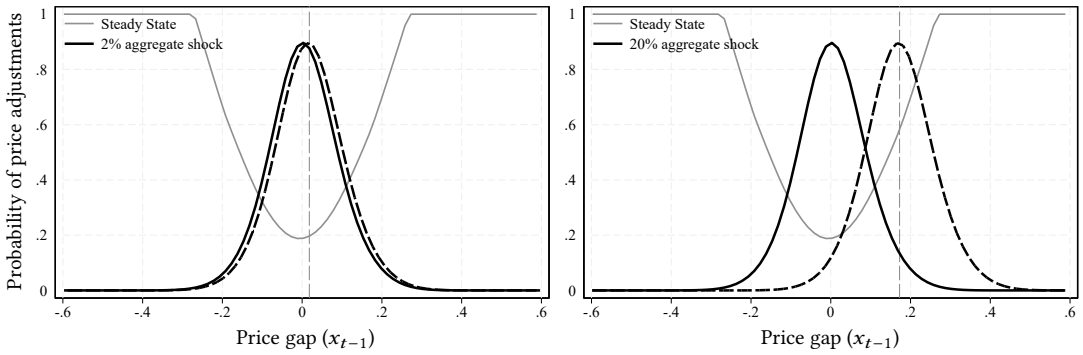
Proof. See Appendix A.4. □

Conditional on the GHF being well approximated by a U-shaped function of price gaps, Proposition 1 offers a robust method for estimating θ^o using microdata on the average frequency of adjustment and average price gaps. In Section 7.1, we will rely on this identification result for model calibration.

3.2 The impact of large aggregate shocks

Aggregate cost shocks—i.e., shocks that do not average out across firms—affect the reset prices of all firms in the economy. When such shocks are large, they shift the entire

Figure 3: Large versus small aggregate cost shocks



Notes. This figure reports the price gap distribution in steady state (black solid line) and after a small and a large aggregate cost shock (black dashed lines).

distribution of price gaps, pushing a substantial share of firms into regions where the generalized hazard function is steep. This displacement amplifies the responsiveness of price adjustment and reduces the degree of monetary non-neutrality in the economy.

Figure 3 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 cost shock of size $g_t > 0$. The left panel shows the effect of a small shock; the right panel displays the response to a large shock. In both panels, the solid lines depict the stationary generalized hazard function (GHF), while the dashed lines represent the distribution of price gaps after the shock.

A small aggregate shock induces only a modest shift in the price gap distribution, resulting in little or no change in the average frequency of price adjustment. In contrast, a large shock pushes a significant number of firms away from their target prices, increasing the average gap in the economy. This displacement leads to a substantial rise in adjustment frequency: more firms raise prices, while fewer reduce them. Under these conditions, the quadratic term in Equation (8) becomes quantitatively relevant, driving aggregate inflation beyond what would be predicted by the change in the average gap alone.

3.3 Price adjustment along the intensive margin

Under our assumptions, the dynamic reset price is approximately equal to the static target, $p_t^\star(f) \approx p_t^0(f)$ and $x_t^\star = 0$. It follows that, for firms that choose to adjust their prices, the

optimal price change is approximately equal to their current price gap:

$$\Delta p_{t-1}(f) \equiv p_t^\star(f) - p_{t-1}(f) \approx x_{t-1}(f) \qquad \text{if } \mathbb{I}_t(f) = 1. \tag{12}$$

Equation (12) provides a testable implication for the mapping from price gaps to price changes along the intensive margin.

Two observations follow. First, in our setting, the approximation is exact in the absence of strategic complementarities in price setting. Under our assumptions, the combined shock—comprising both aggregate and idiosyncratic components—that drives firms' pricing decisions follows a random walk. In turn, the expected present value of future costs coincides with the current realization. When strategic complementarities influence firm pricing, the optimal dynamic price is more volatile than the static one. This is because it also reflects variation in the expected present values of future inflation, which need not follow a random walk. However, with small trend inflation, the additional volatility is negligible, and the approximation remains accurate.[8]

Second, regardless of the process for marginal cost, Equation (12) becomes more accurate after large shocks when firms adjust their prices in a state-dependent manner. As discussed by Dotsey and King (2005), when the probability of a price adjustment in the near future is high, profit flows beyond that horizon have a low weight on today's pricing decisions. Consequently, during periods of large aggregate shocks—when price adjustments become more likely—the wedge between the static and dynamic targets tends to shrink.

3.4 The mapping from price gaps to price changes in the cross-section of firms

Finally, we derive a proposition that combines the effects of price adjustments along both the extensive and intensive margins to characterize the nonlinear pass-through as a function of the shock magnitude.

We again partition the distribution of price gaps into equal-frequency bins (quantiles). Denote by $\gamma_b \equiv \int_{f\in b}((x(f) - x_b)/\sigma_b)^3 df$ the skewness of price gaps within a bin b and define bins sufficiently small so that $|\gamma_b| \leq 1$ for all $b \in \mathcal{B}$.[9]

[8] See Alvarez et al. (2023) for a treatment of menu-cost models with strategic complementarities in continuous time.

[9] We show that this condition is met given our definition of bins in Appendix A.6.

Proposition 2. *The inflation rate within a bin is given by:*

$$\pi_b \equiv \int_{f \in b} \big(p_t(f) - p_{t-1}(f)\big) df = \phi_b^0 \cdot x_b + \phi \cdot x_b^3 + \omega_b. \tag{13}$$

where $\phi_b^0 \equiv 1 - \theta^o + 3\phi\sigma_b^2$ *is a bin-specific coefficient and* $\omega_b \equiv \gamma_b \sigma_b^3 + o(x_b^3)$ *is small when the bins are sufficiently narrowly defined.*

Proof. See Appendix A.5. □

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 (odd) polynomial in the average price gap of 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). 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. The product of these two components gives rise to a cubic approximation of the mapping from price gaps to average price changes.

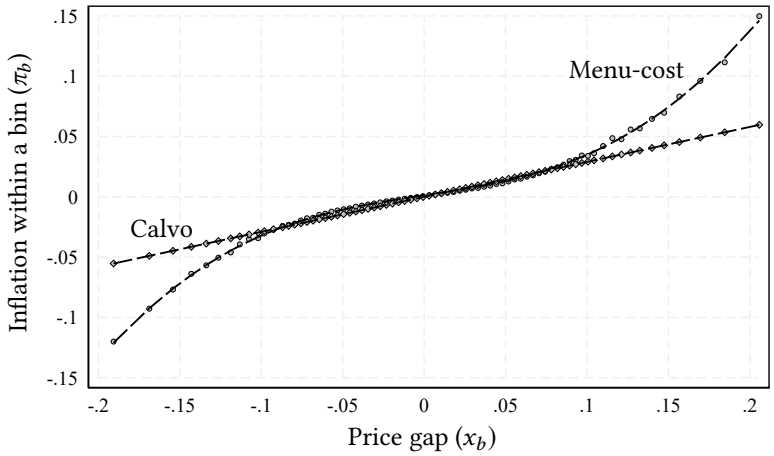
It is straightforward to derive the analog of Equation (13) for a time-dependent Calvo model. Under the assumption of a constant exogenous hazard rate $h^c := (1 - \theta^c)$, we obtain:[10]

$$\pi_b = (1 - \theta^c) \cdot x_b. \tag{14}$$

The binned scatterplot in Figure 4 visually represents the mapping from price gaps to price changes in the menu-cost model (gray circles) and the Calvo model (gray diamonds). The dashed lines show fitted values from regressions of bin-level inflation on either a third-order polynomial in the average gap (Equation 13) or a first-order polynomial (Equation 14). The graph shows that when price gaps are close to zero, the third-order term becomes negligible, and inflation within a bin is approximately 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. This feature underlies the approximate equivalence of the two modeling frameworks in response to small shocks (Gertler and Leahy 2008; Alvarez et al. 2017; Auclert et al. 2024).

[10] 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 the equation in the text.

Figure 4: Nonlinear price dynamics



Notes. In this figure, we partition the distribution of price gaps into quantiles (b) and plot the average price gap of each bin, x_b, against the average logarithmic price change for observations in the same bin, π_b. The gray dots represent simulations of our menu-cost model, while the gray diamonds correspond to simulations from a Calvo model calibrated to the same average frequency of price adjustment. The dashed lines depict: (i) the fitted values of a regression of bin-level inflation on a polynomial in the first and third orders of the average gap, as specified by Equation (13) for the menu-cost model; and (ii) the fitted values from a regression of bin-level inflation on a first-order polynomial in the average gap, as specified by Equation (14) for the Calvo model.

The mapping between price gaps and average price changes becomes nonlinear at the tails of the price gap distribution. Firms in these regions experience large shocks that move them substantially away from their optimal prices. As a result, the frequency of price adjustment increases more than proportionally with the size of the gap (Proposition 1). Consequently, the data-generating process for inflation can be approximated by the cubic expression in Equation (13), which captures the nonlinearities inherent in state-dependent pricing policies when shocks are large.

4 Data and measurement

We assemble a micro-level dataset that covers the Belgian manufacturing sector between 1999 and 2023. Compiled from administrative sources, the dataset contains information on firms' pricing and production decisions at the business cycle (quarterly) frequency, including detailed information on variable production costs, such as labor and intermediate inputs. This dataset extends and enriches the one in our prior work

(Gagliardone et al. Forthcoming) in two key dimensions.[11] First, the data in Gagliardone et al. (Forthcoming) cover a period characterized by low and stable inflation (1999:Q1 to 2021:Q1). Here we extend the time series to include the recent inflation surge and its subsequent decline (2021:Q2 to 2023:Q4). Second, we incorporate new microdata that enable precise measurement of the frequency of price adjustment.

4.1 Measurement of prices, costs, and price gaps

The unit of observation in our analysis is a firm-industry pair. Our final dataset tracks 5,348 domestic firm-industry pairs, denoted by a subscript f, spread across 169 narrowly defined manufacturing industries (4-digit NACE rev.2 product codes), denoted by lower-script i.

Price indices. For each domestic firm, we use PRODCOM data on product-level domestic unit values (sales over quantity sold) to construct a firm-industry price index that aggregates domestic price changes across the products sold by firm f in industry i:[12]

$$\frac{P_{ft}}{P_{ft-1}} = \prod_{p \in \mathcal{P}_{ft}} \left(\frac{P_{pt}}{P_{pt-1}} \right)^{\bar{s}_{pt}}, \tag{15}$$

where $\mathcal{P}_{ft}$ represents the set of 8-digit products manufactured by the firm, P_{pt} is the unit value of product p in $\mathcal{P}_{ft}$, and $\bar{s}_{pt}$ is a Törnqvist weight given by the average within-firm sales share of the product between t and $t-1$, $\bar{s}_{pt} \equiv \frac{s_{pt}+s_{pt-1}}{2}$.

Using data from PRODCOM and from the customs declarations filed by foreign firms exporting to Belgium, we construct firm f competitors' price index by aggregating the domestic price changes of products sold by domestic and international competitors selling in the same industry as f ($\mathcal{F}_i$):

$$\frac{P_{it}^{-f}}{P_{it-1}^{-f}} = \prod_{k \in \mathcal{F}_i \setminus f} \left(\frac{P_{kt}}{P_{kt-1}} \right)^{\bar{s}_{kt}^{-f}}, \tag{16}$$

where $\bar{s}_{kt}^{-f} \equiv \frac{1}{2}\left(\frac{s_{kt}}{1-s_{ft}} + \frac{s_{kt-1}}{1-s_{ft-1}} \right)$ represents the Törnqvist weight assigned to competitor k

[11] We refer to Gagliardone et al. (Forthcoming) for details about the data sources and variable definitions.

[12] PRODCOM's sampling scheme is designed to survey all Belgian manufacturing firms with more than 10 employees with the goal of accounting for at least 90% of output in each NACE 4-digit industry. We recover domestic values and quantities sold by merging the product-level information from PRODCOM with firms' product-level export quantities and sales from the Belgian customs declarations.

given by the average residual revenue share of competitor k in the industry (excluding firm f's revenues).

Finally, we recover the time series of firms' prices (in levels) by concatenating the price indices in Equation (15), $P_{ft} = P_{f0} \prod_{\tau=t_f^0+1}^{t} \left(P_{f\tau}/P_{f\tau-1}\right)$, where t_f^0 denotes the first quarter when f appears in our data. We set the base period P_{f0} to one for all firms. As discussed in the following section, this normalization is one rationale for removing firm-fixed effects from our empirical measures of price gaps. The series of competitors' prices, P_{it}^{-f}, is constructed similarly, concatenating the price indices in Equation (16).

Frequency of price adjustment. We use micro-level records from the National Bank of Belgium Business Survey (NBB-BS) to accurately measure the frequency of price adjustment. This survey regularly interviews a representative sample of firms across manufacturing industries about their pricing decisions. Similar to the official Producer Price Index (PPI) data collection, the NBB-BS asks firms whether they increased, decreased, or maintained the price of a specific product in their portfolio. Using this information, we define a Boolean variable that equals one if a firm reports adjusting prices at least once within a given quarter relative to the previous month. Averaging this variable across firms and industries in each quarter, we compute the manufacturing sector's average frequency of price adjustment (h_t).

Information on firms' price adjustments from the NBB-BS also helps filter out spurious price changes in the microdata. As discussed, our price measure is derived from product-level unit values, which tend to overstate the frequency of small price changes due to minor measurement errors, as shown by Eichenbaum et al. (2014) and Cavallo and Rigobon (2016). To address this issue, we combine firm-level price changes with information on price adjustment frequency from the NBB-BS to define firm-specific thresholds, κ^+ and κ^-, such that small price variations within these bounds are treated as no price update:

$$\begin{aligned} \mathbb{I}_{ft}^{+} = 0 &\iff \Delta p_{ft} < \kappa^{+} \cdot Var_f(\Delta p_{ft}) \qquad \text{if } \Delta p_{ft} > 0 \\ \mathbb{I}_{ft}^{-} = 0 &\iff \Delta p_{ft} > -\kappa^{-} \cdot Var_f(\Delta p_{ft}) \qquad \text{if } \Delta p_{ft} < 0 \end{aligned}$$

To account for different degrees of upward and downward nominal rigidity in the data, we set the thresholds $\kappa^+ = 0.75$ and $\kappa^- = 0.87$ to separately match the average frequency of upward and downward price changes measured using the NBB-BS

microdata: $\sum_t \frac{\sum_f \mathbb{I}^+_{ft}}{\sum_f \mathbb{I}^+_{ft}+\sum_f 1-\mathbb{I}^+_{ft}} = h^+$ and $\sum_t \frac{\sum_f \mathbb{I}^-_{ft}}{\sum_f \mathbb{I}^-_{ft}+\sum_f 1-\mathbb{I}^-_{ft}} = h^-$, where $h^+ + h^- = h$.[13]

Marginal cost indices. To derive a firm-level marginal cost index, we assume a cost structure in which the nominal marginal cost of a firm is proportional to its average variable costs: $MC^n_{ft} = (1 + \nu_f)AVC_{ft}$. Relaxing the assumption in our theoretical framework, the coefficient ν_f captures the curvature of the short-run cost function, and it is inversely related to the firm's short-run returns to scale in production ($\nu_f \equiv 1/RS_f - 1$). Using the definition of average variable costs (total variable costs over output, TVC^n_{ft}/Y_{ft}) and applying a logarithmic transformation, we have that firm-level log-nominal marginal cost is given by:

$$mc^n_{ft} = (tvc^n_{ft} - y_{ft}) + \ln(1 + \nu_f). \tag{17}$$

We measure total variable costs as the sum of intermediate costs (materials and services purchased) and labor costs (wage bill), sourced from firms' quarterly VAT and social security declarations. To construct a firm-level quantity index, we divide domestic revenues by the firm-specific domestic price index.[14] Firm-specific short-run returns to scale are not directly observable in the data. Thus, if production technologies deviate from constant returns to scale ($\nu_f \neq 0$), our measure of log marginal costs will omit an additive firm-specific constant. This is the second rationale for removing firm-fixed effects from our measure of price gaps, together with the normalization of the price index discussed above.

Price gaps. The availability of high-frequency firm-level price and cost data allows us to construct an empirical counterpart of the (ex-ante) price gaps defined in our model:

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

As discussed in Section 2, when p^o_{ft} and $p^\star_{ft}$ are sufficiently close, x_{ft-1} captures inefficiencies driven by nominal rigidities and influences firms' pricing policies. A positive

[13] See Karadi et al. (2024) and Luo and Villar (2021) for evidence of asymmetric upward and downward rigidity.

[14] Specifically, we compute $Y_{ft} = (PY)_{ft}/\bar{P}_{ft}$, where $\bar{P}_{ft}$ denotes the firm-quarter domestic price index. For single-industry firms, $\bar{P}_{ft}$ coincides with the firm-industry price index P_{ft}. For multi-industry firms, we construct $\bar{P}_{ft}$ as an average of the different firm-industry price indices using as weights the firm-specific revenue shares of each industry. As discussed in Gagliardone et al. (Forthcoming), the vast majority of firms in our sample operate in only one industry, and the main industry accounts for the lion's share of sales of multi-industry firms.

price gap indicates that a firm operates with a markup below the profit-maximizing level and, absent nominal rigidities, would adjust its price upward.

Guided by our theoretical framework, we construct a measurable proxy for firms' target prices as a convex combination of the firm's own marginal cost and its competitors' price index, $p^o_{ft} = (1 - \Omega)mc^n_{ft} + \Omega p_t^{-f}$, calibrating Ω to 0.5 to match the micro estimate in Gagliardone et al. (Forthcoming). This empirical measure of firms' target prices differs from the theoretical one (Equation (3)) in three dimensions. First, it does not capture variation in firms' steady-state markups. Second, it does not directly account for potential curvature in firms' short-run cost functions. The harmonization of the data discussed below helps address these two limitations. Third, we cannot directly measure the realized idiosyncratic taste shocks (φ_{ft}). While such shocks are likely transitory and average out in the cross-section, they are a source of measurement error that possibly weakens the connection between our empirical measures of price gaps and price changes at the firm level.[15]

4.2 Harmonization and data cleaning

Following prior literature, we apply the standard data-cleaning steps and harmonization procedures to address issues related to measurement error and unobservable cross-sectional heterogeneity, which tend to bias the statistics describing the distribution of price changes and, *a fortiori*, of price gaps.[16]

First, the use of unit values tends to introduce in the measurement of price changes, either by recording small changes where none occurred or by overstating the magnitude of genuine adjustments (Eichenbaum et al. 2011). To mitigate the impact of such measurement error, we recode all price changes smaller than 1 percent in absolute value as zero and exclude observations in the top and bottom 1.5 percent of the price change distribution. Second, we standardize price changes at a granular level by demeaning observations within each cell, where a cell is defined as a firm-industry pair—the most disaggregated level in our dataset. This correction helps address issues related to unobserved heterogeneity among firms, which can inflate the measured standard deviation and kurtosis of price changes (Klenow and Kryvtsov 2008; Alvarez et al.

[15] See the evidence in Gagliardone et al. (Forthcoming).

[16] See, Klenow and Kryvtsov (2008) and Alvarez et al. (2016) for discussion of these issues and proposed solutions.

2016) and also removes the average trend inflation rate, which was approximately 0.6 percent quarter-on-quarter before the 2021 surge. Finally, we remove industry-specific calendar-quarter averages to account for seasonality in price-setting behavior.

We apply the same trimming and harmonization procedure to the distribution of price gaps. Demeaning helps reconcile differences between our empirical measure of price gaps and its theoretical counterpart. These differences stem from firm-specific intercepts in target price definitions, trend inflation, and industry-specific seasonal patterns in nominal costs.[17]

4.3 The empirical distribution of price changes and price gaps

Table 1 presents summary statistics of the distribution of firm-level log price changes and price gaps. The first four columns report moments describing the distribution of price changes. 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. Drawing an analogy with our model, we view this period as representing the economy in its steady-state distribution. During this period, the (harmonized) average price change is close to zero, which implies that inflation is generally aligned with the long-term industry trend. During this period, the standard deviation of price changes is 0.11 and the average frequency of price adjustment is $h = 0.29$. The latter implies that, in a low inflation environment, firms adjust their prices every 3 to 4 quarters, on average (Gagliardone et al. Forthcoming).

Panel b reports the same statistics for the pandemic and post-pandemic period (2020–2023), which was marked by high inflationary pressures followed by a gradual tapering. During this period, the PPI quarterly inflation rate was, on average, 1 percentage point above the long-run trend. At the same time, the frequency of price changes increased by 10 percentage points, on average.

The fourth column reports the kurtosis of price changes, calculated following the approach in Klenow and Kryvtsov (2008). The measured kurtosis is 3.26—consistent with values commonly found in the literature (typically ranging between 3 and 5)—indicating

[17] As noted above, we can identify firms' static reset prices, p^o_{ft}, up to a firm-specific additive constant. This constant reflects a combination of unobserved steady-state markups (the term $\mu(f)$ in Equation (3)), unobserved deviations from constant short-run returns to scale affecting marginal costs (the term $\ln(1+\nu_f)$ in Equation (17)), and the normalization of price levels discussed in Section 4.1.

Table 1: Summary statistics of price changes and price gaps

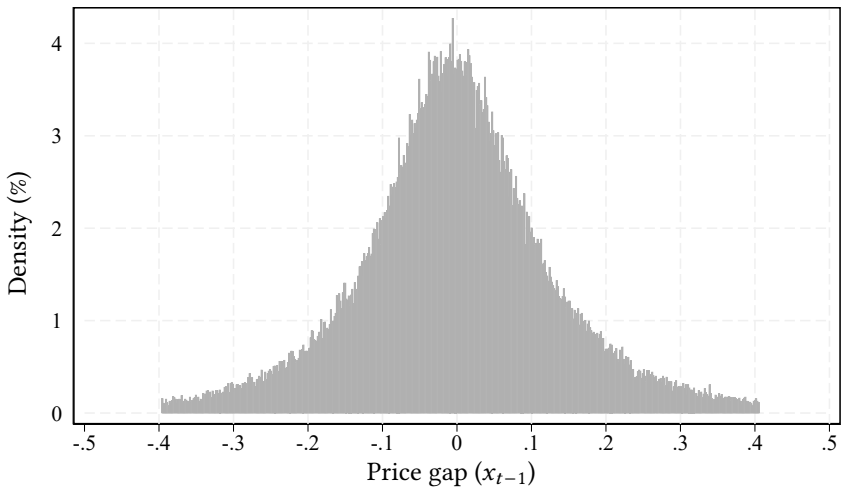
Price changes ($p_{ft} - p_{ft-1}$)				Price gaps (x_{ft-1})		
Panel a: Time period 2000-2019						
Mean	Std.	Freq. Adj.	Kurt.	Mean	Std.	Kurt.
0.00	0.12	0.29	3.26	-0.00	0.13	2.86
Panel b: Time period 2020-2023						
Mean	Std.	Freq. Adj.	Kurt.	Mean	Std.	Kurt.
0.02	0.11	0.44	3.56	0.01	0.14	2.83
Number of observations:			118,308			
Number of firm-industry pairs:			4,974			
Number of firms:			4,488			

Notes. This table reports the summary statistics of the distributions of price changes ($p_{ft-1} - p_{ft}$), and price gaps (x_{ft-1}) before (panel a) and after the inflation surge (panel b).

that the distribution of price changes is more peaked (a higher frequency of small adjustments) and exhibits fatter tails (a greater frequency of large adjustments) than a normal distribution (kurtosis = 3). These features are consistent with empirical patterns associated with state-dependent price-setting behavior and become even more pronounced during the inflation surge, with kurtosis rising to 3.56 over the 2000–2023 period.

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 from nominal rigidities. The histogram in Figure 5 plots the probability density function of the price gaps, $f(x_{t-1})$, for the pre-pandemic period. Consistent with our theory, the empirical distribution of price gaps is unimodal, bell-shaped, and symmetric around the mean. Table 1 reports that the average price gap rose by approximately 1 percentage point relative to its long-term trend. In line with model predictions, this increase corresponds closely to the average price change observed over the same period.

Figure 5: Empirical distribution of price gaps



Notes. The figure presents the empirical probability density function of the price gaps, $f(x_{ft-1})$, in the pre-pandemic period (2000-2019).

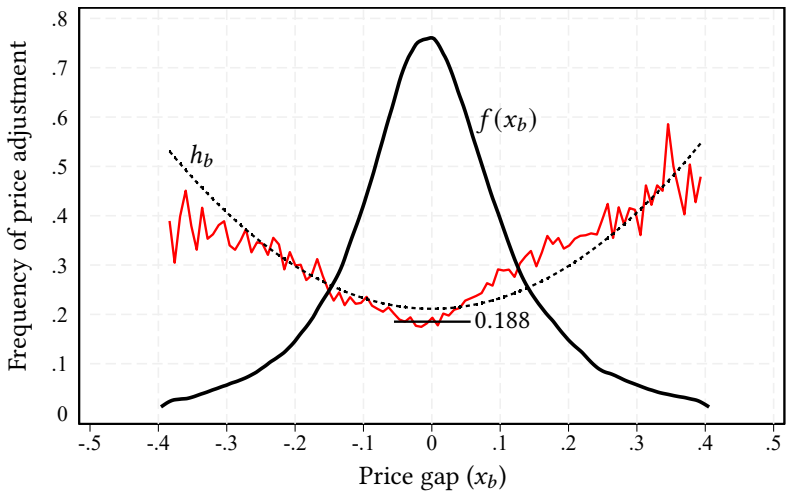
5 Micro evidence of state dependent pricing

Guided by the theoretical framework presented in Section 2, we design empirical tests of key model predictions that non-parametrically 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 novel identification results that inform the calibration of structural parameters.

5.1 The empirical generalized hazard function

We begin by analyzing the mapping between price gaps and the frequency of price adjustment—*the extensive margin*. As discussed above, in time-dependent models, the two variables are independent, yielding a flat generalized hazard function (GHF). By contrast, state-dependent models predict that the frequency of price adjustment increases monotonically with the absolute value of the price gap. In our framework, the GHF can be approximated by a convex function—specifically, a second-order polynomial in the price gap—as formalized in Lemma 1. Proposition 1 provides the sample analog of this relationship, linking the average frequency of price adjustment at different points in the

Figure 6: Empirical GHF and distribution of price gaps



Notes. The figure plots the empirical probability density function of the price gaps $f(x_b)$ (black line) against the empirical GHF, h_b (red line). The black dotted line is the fitted value obtained from a cross-sectional regression of the frequency of price adjustment of a given bin (b) on a constant and the square of the average price gap of the same bin, as dictated by Equation (8). In the regression, we weight each bin by the number of observations it counts.

gap distribution to the average squared gaps.

We take these predictions to the data microdata by looking at the relationship between the frequency of price adjustment and price gaps. In Figure 6, the black line depicts the kernel estimate of the empirical probability density function of price gaps. The red line plots the empirical GHF—defined as the fraction of firms adjusting prices within each quantile (bin) of the price gap distribution. The black dotted line presents fitted values from a cross-sectional regression of bin-level adjustment frequencies on the square of the average price gap, $\hat{h}_b = \hat{a}_0 + \hat{a}_1 \cdot x_b^2$.

The data reveal a strong, convex relationship between the size of price gaps and the frequency of price adjustment. In striking similarity to the theoretical GHF implied by state-dependent pricing models, wider gaps (i.e., greater deviations from $x_{ft-1} = 0$) are associated with a higher likelihood of adjustment. Moreover, consistent with our model's prediction, a simple quadratic polynomial in price gaps fits the frequency data remarkably well. We also note that the empirical GHF exhibits a steeper slope for positive gaps, suggesting asymmetry in nominal rigidity: firms are more likely to adjust prices when they are too low (when the realized markup is below the desired one) than when they are too high.

We then leverage Proposition 1 to estimate the probability of free price adjustment directly from the empirical GHF. Specifically, fitting the quadratic polynomial over an interval centered around the bottom of the parabola ($x_b \in (-0.04, 0.04)$) yields an estimate of $\hat{a}_0 = 0.188$, implying a Calvo-type parameter of $\widehat{\theta^o} = 0.812$. We use this estimate in our model calibration, as detailed in Section 7.1.

5.2 The impact of large aggregate shocks

In Section 2, we discussed how small idiosyncratic shocks generate dispersion in the price gap distribution, while large aggregate shocks shift the entire distribution, substantially increasing the fraction of firms with incentives to adjust prices (see Figure 3). The pronounced increase and subsequent normalization of production costs in the post-pandemic period offer a compelling empirical setting to test this prediction.

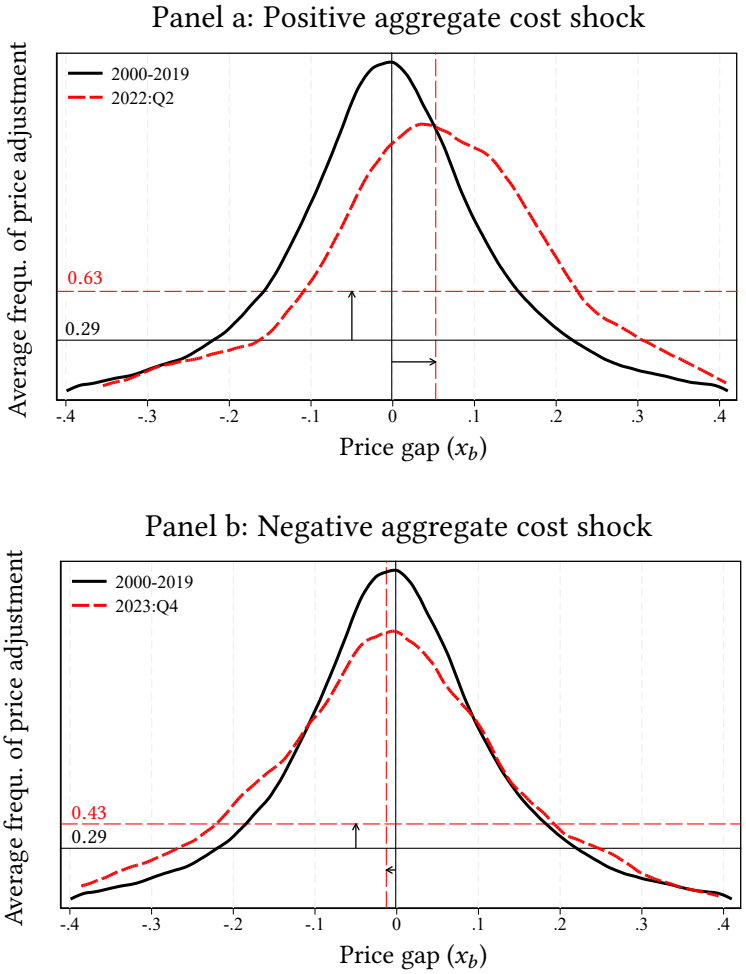
In Figure 7, the solid black line shows the distribution of price gaps in the pre-pandemic period. In panel (a), we overlay the distribution in 2022:Q2 (red dashed line), a quarter in which firms' marginal costs *increased* by an average of 6.2 percent relative to the previous quarter. As predicted by the model, a shock of this magnitude shifts the entire distribution to the right, displacing many firms from their optimal price levels. The resulting compression in markups raises the cost of inaction, pushing a greater share of firms into regions where the GHF is steep. As a consequence, the average frequency of price adjustment nearly doubles within a single quarter relative to its long-run average.

Panel (b) repeats the exercise for 2023:Q3, a quarter in which firms' marginal costs *declined* by an average of 3.8 percent, driven by falling energy prices and easing supply chain pressures. Consistent with model predictions, the distribution of price gaps shifts to the left. The adjustment frequency responds to this shift, as more firms reduce prices than increase them.

5.3 Price gaps and price changes conditional on adjustment

We now study the mapping from price gaps to price changes, focusing on firms that adjust their prices—the *intensive margin*. Under our assumptions, p^o_{ft} provides a reasonable approximation for $p^\star_{ft}$, so that $\Delta p_{ft-1} \approx x_{ft-1}$ if $\Delta p_{ft-1} \neq 0$, as formalized in Equation (12). Accordingly, our model predicts that the elasticity of price changes with respect to

Figure 7: Impact of aggregate cost shocks on the price gap distribution and average frequency of price adjustment

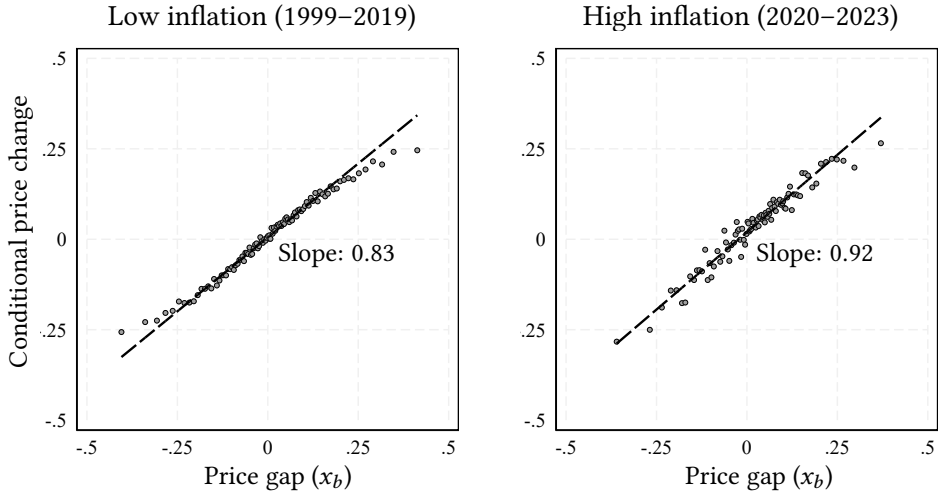


Notes. This figure presents the empirical probability density function of the price gaps in the pre-pandemic period, 1999–2019, (black solid line) and in two snapshots of the post-pandemic period, in 2022:Q2 (red dashed line, panel a) and 2023:Q4 (red dashed line, panel b). The solid and dashed vertical 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:Q4 (red dashed lines).

price gaps should be approximately one. We test this prediction in Figure 8.

The figure presents binned scatterplots for firms that adjust prices. The x-axis shows the average price gap in each percentile of the distribution, while the y-axis reports the corresponding average price change. 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

Figure 8: Price changes and price gaps, conditional on adjusting



Notes. This figure presents a binned scatterplot of the log price change for adjusters (i.e., firms for which $p_{ft} \neq p_{ft-1}$) against the price gap. 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.

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 assumptions underlying our theoretical framework–particularly the firm-level nominal marginal costs display random-walk dynamics—are empirically plausible.

That said, theory predicts a pass-through of exactly one. Three factors may explain the observed deviation from this benchmark. First, classical measurement error in price gaps likely attenuates the estimated elasticity. Second, the approximation of $p_t^\star(f)$ with $p_t^0(f)$ is an additional source of measurement error. Notably, the elasticity is closer to one in the high-inflation period (panel b), suggesting that the approximation performs better in such environments. As discussed in Section 3, this finding is consistent with the model's prediction: large aggregate shocks trigger a state-dependent response along the extensive margin, leading firms to discount future conditions more heavily. This behavior makes the static approximation more accurate even if marginal cost deviates from the random walk dynamics. Third, the model assumes that firms have complete information. In practice, firms may underreact to fundamentals due to inattention or incomplete information,

as shown in Morales-Jiménez and Stevens (2024) and Gagliardone and Tielens (2024). Importantly, the frictions in these models are themselves state-dependent, implying that pass-through becomes stronger when shocks are large—consistent with what we observe in the data.

5.4 The empirical mapping from price gaps to price changes

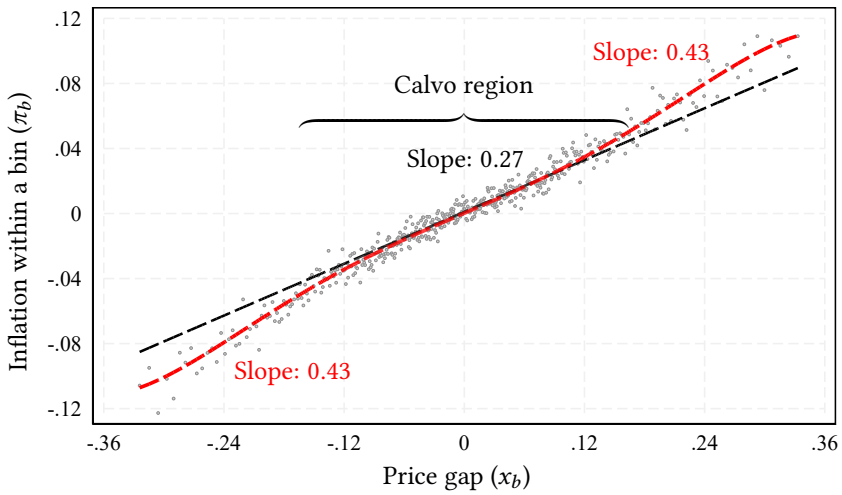
In our final analysis, we examine the joint effect of the extensive and intensive margins of price adjustment and characterize the nonlinear pass-through as a function of shocks' magnitudes. According to Proposition 2, the mapping between price gaps and inflation (i.e., expected price changes) can be approximated by a cubic polynomial in the price gaps. This functional form captures both the linear pass-through observed in response to small shocks as well as the nonlinearities introduced by state dependence when shocks are large.

The data provides strong support for the predictions summarized by Proposition 2. Figure 9 takes this prediction to the data. As before, we sort observations into quantiles (bins) spanning the price gap distribution. We then plot the average price change within the same bin (π_b, y-axis) against the corresponding average price gap (x_b, x-axis). The red dashed line in Figure 9 represents the fitted values from a cross-sectional regression 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 relationship between price gaps and average price changes is S-shaped, indicating that the cross-section of firm-quarter observations can be partitioned into two distinct regions. The first comprises observations 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 as the "*Calvo region*." In this range, the mapping between price gaps and expected price changes is approximately linear, consistent with the predictions of the Calvo model (Equation 14) and characteristic of "normal times," when shocks are modest and inflation remains low and stable. This result echoes findings from Gertler and Leahy (2008), Alvarez et al. (2017), and Auclert et al. (2024), which show that with small shocks the price dynamics in state-dependent models resemble those of time-dependent models. Notably, the estimated gradient in this region is 0.27—nearly identical to the

Figure 9: Nonlinear price dynamics



Notes. This figure presents a scatterplot of the average price gap for each bin of the price gap distribution, x_b, against the corresponding average inflation (the logarithmic price change) with the bin, π_b. Bins are constructed by partitioning the price gap distribution into 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 (odd) 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.

average frequency of price adjustment in the pre-pandemic period.

The second region comprises observations in the tails of the price gap distribution, where large shocks displace firms substantially from their optimal prices (by 20 percent or more). As shown above, in this part of the distribution, the frequency of price adjustment increases more than proportionally with the size of the gap. As a result, 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 joint effect of adjustments along both the extensive and intensive margins in response to large shocks.

6 Aggregate cost-price dynamics

In this section, we shift focus from micro-level dynamics to aggregate cost-price dynamics. We show that firms' state-dependent pricing behavior gives rise to nonlinear aggregate

inflation dynamics, whereby the pass-through of aggregate cost shocks varies with the magnitude of the shock.

6.1 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 adopted by national statistical agencies, we calculate domestic PPI inflation as a Törnqvist price index, averaging firms' quarterly price changes weighted by their 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, mc_t^n, by aggregating and concatenating the average changes in firm-level nominal marginal costs across producers (Δmc_t^n):

$$mc_t^n = \sum_{t=1999:Q2}^{2023:Q4} \Delta mc_t^n$$

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

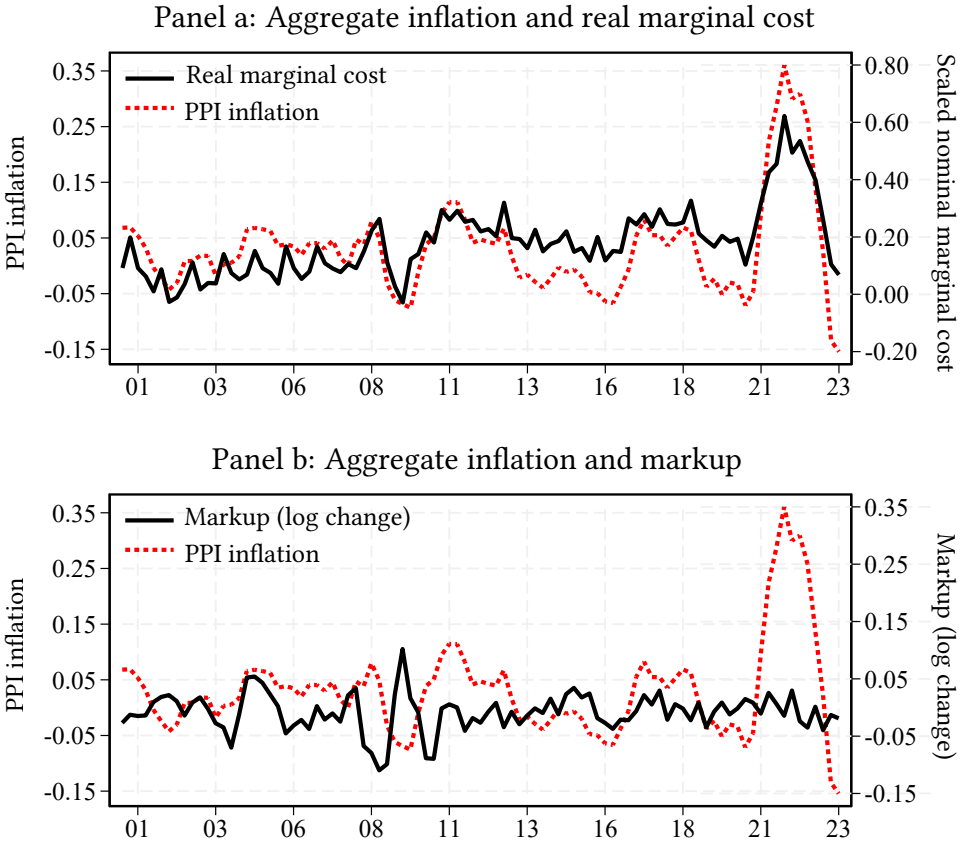
where the value of the index in the first quarter of our data is normalized to zero.

According to our theory, firms set prices based on movements in 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 $t-4$. This logarithmic difference between these variables, $mc_t^n - p_{t-4}$ provides a measure of real marginal cost. Panel a of Figure 10 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 key patterns emerge. First, consistent with the theory, inflation closely tracks the fluctuations of real marginal cost throughout the sample, albeit with incomplete pass-through: price changes respond less than one-for-one to cost movements. Second, the sharp rise and subsequent decline in inflation during the pandemic and post-pandemic period mirrors a similarly large swing in real marginal costs.

To further stress the contribution of cost passthrough to movements in inflation,

Figure 10: Inflation, cost, and markup dynamics

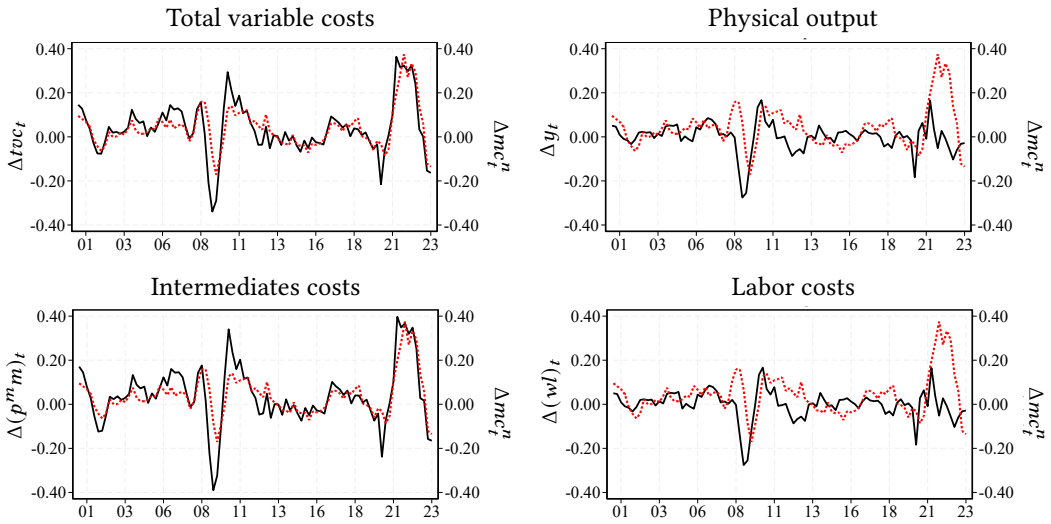


Notes. This figure shows the time series of year-over-year manufacturing PPI inflation ($p_t - p_{t-4}$) alongside the times series of the real marginal cost index ($mc_t^n - p_{t-4}$, panel a) and the log change in average realized markups ($\Delta \ln(\text{Markup}_t)$, panel b) for the Belgian manufacturing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

panel b of Figure 10 contrasts inflation with changes in average realized markups. We compute the latter using the identity $\Delta \ln(\text{Markup}_t) \equiv \pi_t - \Delta mc_t^n$. This exercise illustrates that, at least in our sample, the hypothesis that a rise in markups can explain the recent inflation surge seems to have no bite in the data. This finding echoes findings in Alvarez et al. (2024), who use firm-level data from a large multinational manufacturer.

Finally, to get a sense of what drove the fluctuations in nominal costs, Figure 11 presents a decomposition of our aggregate cost index into its different components (Equation (17)). 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

Figure 11: Decomposition of aggregate nominal marginal cost index



Notes. This figure decomposes 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), intermediates costs ($\Delta(p^m m)_t$, bottom left panel), and labor costs ($\Delta(wl)_t$, bottom left panel).

during the recent inflation surge, fluctuations in total variable costs are the main drivers of the time series evolution of nominal marginal cost.

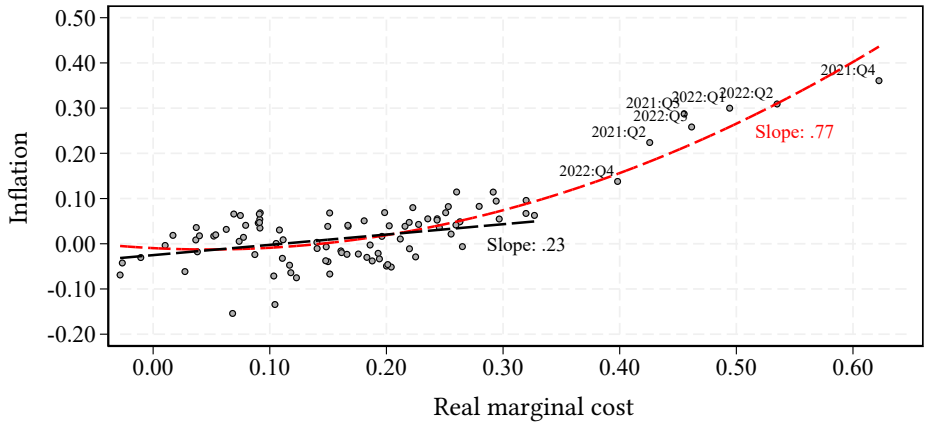
The two panels at the bottom of Figure 11 further decompose total variable costs into the cost of intermediate inputs (purchases of materials, services, and energy) and the cost of labor. As we can see, both cost components rose during the post-pandemic period. However, the increase in the cost of intermediates was four times greater. This component alone accounts for approximately 70% of the revenues of manufacturing firms, on average. In addition, more than 80% of intermediate input costs are paid to foreign suppliers. These figures make clear how the shock to the cost of (foreign-supplied) intermediates—rather than a surge in labor cost—was the main driver of the inflation surge between 2021 and 2023, at least in our sample.

6.2 Nonlinear aggregate pass-through

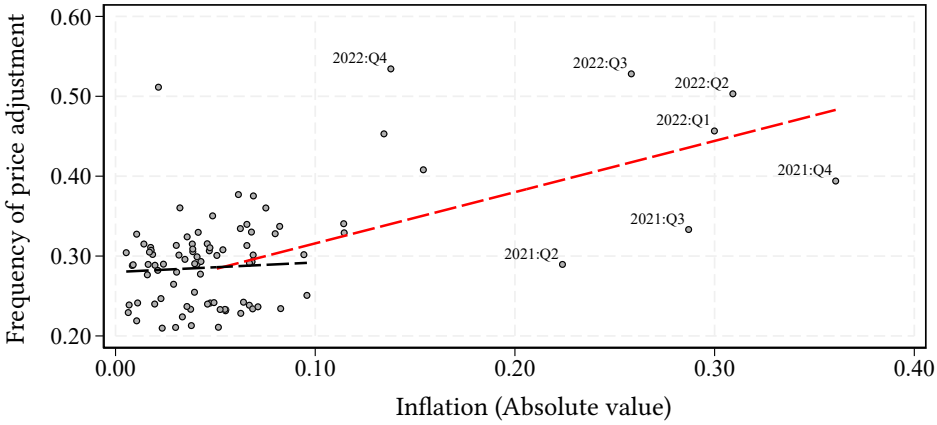
In panel a of Figure 12 we sort quarters by their measured real marginal cost index and plot this index against PPI inflation. The black dashed line represents the linear fit between the two variables during periods of low inflation (below 10% year-over-year), while the red dashed line shows a quadratic fit across high- and low-inflation periods. The slopes

Figure 12: Passthrough of costs into inflation

Panel a: Nonlinear passthrough of costs into inflation



Panel b: Frequency of price adjustment and inflation



Notes. In panel a, we sort the different quarters in our data according to their realized aggregate real marginal cost ($mc_t^n - p_t$) and plot this variable against year-over-year manufacturing PPI inflation (π_t) in the same quarter. In panel b, we sort the data according to realized year-over-year manufacturing PPI inflation and plot this variable against the average frequency of price adjustment (h_t) in the same quarter. 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% year-over-year); the red dashed line represents a quadratic fit across both high- and low-inflation periods.

of these curves provide descriptive evidence on the aggregate passthrough of cost shocks into prices as a function of the size of the shock to marginal cost.

Consistent with the micro-level dynamics documented in Section 5, we find a linear relationship between aggregate inflation and nominal costs during normal times. As

discussed, a linear passthrough is consistent with the predictions of both the Calvo model and the menu-cost model when aggregate shocks are small. The estimated reduced-form slope is 0.23, closely aligning with the aggregate passthrough coefficient estimated by Gagliardone et al. (Forthcoming) in a low inflation environment.

However, this linear relationship breaks down when the economy is hit by large aggregate shocks. The pass-through coefficient more than tripled during the recent inflation surge, revealing highly nonlinear cost-price dynamics. At the core of this nonlinearity is the endogenous nature of the frequency of price adjustment. In Panel b, we sort the quarters by their annual inflation rates and plot aggregate inflation against the average frequency of price adjustment. As before, the black dashed line shows the (linear) fit during periods of low inflation, and the red dashed line represents the (nonlinear) fit once high-inflation periods are included in the sample. As in the Calvo model, in low-inflation environments, we observe essentially no relationship between the average frequency of price adjustment and inflation. This finding reinforces the view that time-dependent models provide a good framework to capture nominal rigidities with low inflation. However, inflation and the frequency of price adjustment are highly correlated in high-inflation environments, as shown by Alvarez et al. (2019) in the case of Argentina and by Cavallo et al. (2023) and Blanco et al. (2024a) in Europe and the US.

7 Quantitative implications

Having established the close connection between theory and data, we now calibrate the menu-cost model presented in Section 2 and conduct two types of quantitative exercises. First, we simulate cost shocks of various magnitudes and compare the dynamics of our state-dependent specification to those of a standard time-dependent Calvo model. Second, we feed the model a sequence of aggregate marginal cost shocks extracted from the data and evaluate how the model-generated inflation series compares to inflation in the data.

7.1 Calibration

We have a total of seven parameters to calibrate. We calibrate four of them to standard values in the literature. We calibrate the elasticity of substitution between goods σ, to 6, which implies a markup of 20 percent in the symmetric steady state equilibrium. We

set β, the firm's risk-neutral discount factor, at 0.99. As in our empirical analysis, we calibrate $\Omega = 0.5$ to reflect the importance of strategic complementarities estimated using microdata in our prior work (Gagliardone et al. Forthcoming) in Amiti et al. (2019). To align model and data, we allow for a drift in the aggregate component of nominal marginal cost ($\mu_g = 0.5\%$), which corresponds to a trend inflation rate of 1.6% year over year.

We develop a new approach 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. In the absence of data on price gaps, the standard calibration strategy relies on model inversion to recover structural parameters that jointly match moments from the observed distribution of price changes—typically, the standard deviation of price changes, the average frequency of adjustment, and the kurtosis of price changes.[18] In theory, this approach identifies the parameters of interest through the lens of a specific model. In practice, obtaining a reliable measure of the kurtosis of price adjustments is challenging due to measurement error and the inability to control for unobserved heterogeneity (Alvarez et al. 2016; Cavallo and Rigobon 2016; Berger and Vavra 2018). These issues may be particularly pronounced in our context, given the reliance on unit values. Thus, our calibration approach avoids targeting the kurtosis and instead proceeds as follows:

Step 1. As discussed in Section 5.1, we leverage information embedded in the empirical GHF to calibrate the free-price adjustment probability $(1 - \theta^o)$. To do so, we match the frequency of price adjustment in the vicinity of the bottom of the GHF, around $x_{ft-1} \approx 0$. By averaging across observations in the zero-gap neighborhood, this calibration is robust to small measurement errors due to spurious changes in unit values.[19] As discussed in Section 5, this exercise yields $(1 - \hat{\theta}^o) = 0.188$.

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

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

Using this identity we calibrate σ_ϵ^2 to 0.0036 to match the product of the average frequency

[18] See, e.g., Alvarez et al. (2022) and Blanco et al. (2024a).

[19] See Appendix A.4, for a formal discussion on the consistency of this estimator.

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	3.26	-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. This table reports moments of the distribution of price changes and price gaps computed during the period 2000–2019 and the corresponding moments for the menu-cost model and Calvo model, in steady-state, under our baseline calibration. The last column shows the average share of menu costs as a fraction of firms' revenues, with the data estimate sourced from Zbaracki et al. (2004)

of price adjustment 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 limit of the uniform distribution from which the random menu costs are drawn) to 0.61, to allow the model to match the frequency of price changes in the pre-pandemic period.

Table 2 compares empirical moments of the price change distribution (panel a) and the price gap distribution (panel b) to those generated by the menu-cost model, in steady state, under our baseline calibration. The model fits the data closely. Notably, the calibrated model produces menu costs that amount to 1.7 percent of firm revenues in steady state. While we did not target this moment directly, the implied magnitude is reasonable and broadly consistent with empirical estimates from Levy et al. (1997) and Zbaracki et al. (2004), which suggest that menu costs are small. In addition, the model reproduces a kurtosis of price changes that is broadly in line with the one measured in the data, despite the potential measurement issues discussed earlier. To this point, Blanco et al. (2024a) shows that in a standard menu-cost model with single-product firms, calibrating to match the kurtosis of price changes may require implausibly high menu costs to generate sufficient volatility in the frequency of price adjustment. We confirm this result in our setting: a calibration that targets the kurtosis observed in our data yields sufficient variation in adjustment frequency only at the cost of implausibly high (more than double) menu costs.

Calibration of the Calvo model. In our quantitative exercises, we also consider a standard time-dependent Calvo model. As explained in Section 2, our menu-cost model nests the Calvo model as a special case. To calibrate this model, we set the maximum menu cost, $\bar{\chi}$, to an arbitrarily large number and re-calibrate the probability of free price adjustment to match the steady-state adjustment frequency ($1 - \theta^c = h$). The theoretical moments derived from the calibrated Calvo model are reported in the third row of Table 2. As expected, the Calvo model produces a pronounced degree of leptokurtosis, which is substantially at odds with the kurtosis observed in the data.

7.2 Impulse-responses to small and large aggregate shocks

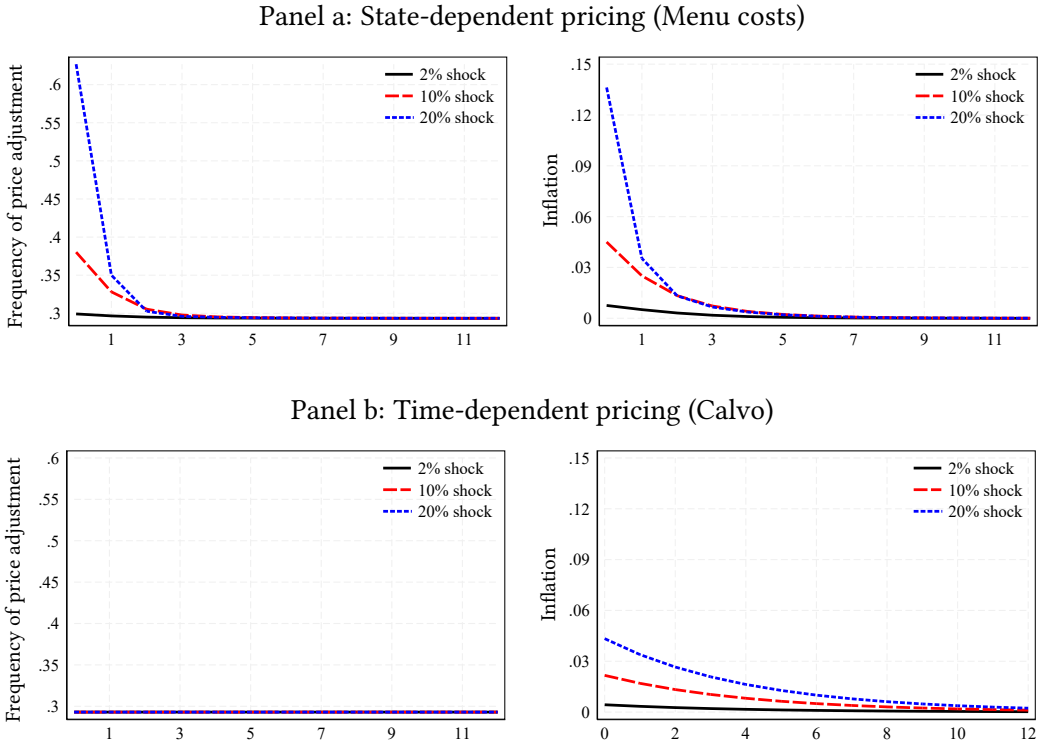
We use our calibrated menu-cost and Calvo models to study price dynamics in response to large and small shocks, under state- and time-dependent pricing. Starting from an economy in steady state, we shock the system with permanent and unanticipated aggregate cost shocks of different magnitudes, g_t=$\{2\%, 10\%, 20\%\}$.

Figure 13 shows 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 increase the optimal reset price, shifting the price gap distribution to the right and triggering an inflation response. However, as discussed in Section 2 and shown empirically in Section 5, large shocks in the menu-cost model push firms into regions of the price gap distribution where the GHF is steep. This generates a spike in the frequency and, in turn, a sharper and faster rise in inflation relative to the Calvo benchmark.

Comparing IRFs across shocks of different magnitudes highlights the nonlinearities inherent in state-dependent pricing, which become more pronounced as shock size increases. On impact, a large shock produces an inflation response and adjustment frequency nearly three times that of a medium shock, even though the shock itself is only twice as large (10% vs. 20%). In contrast, in the Calvo model, the frequency of price adjustment and, as a result, inflation rises proportionally with the size of the shock.

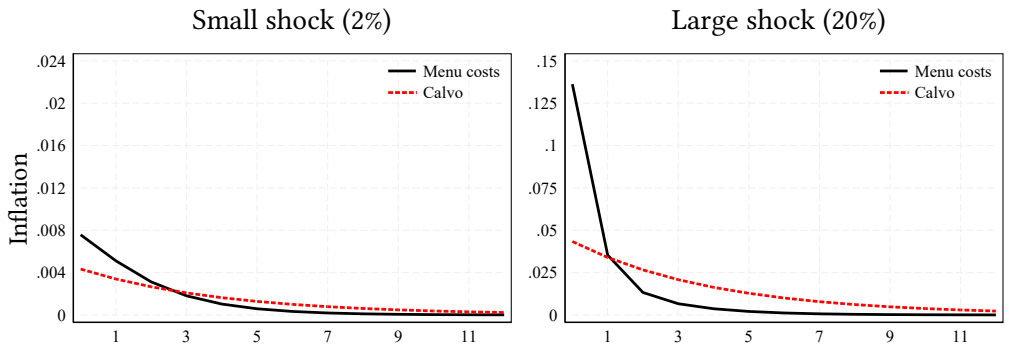
A second key difference concerns the speed at which permanent cost shocks are fully incorporated into prices. Figure 14 overlays inflation IRFs from both models under small and large shocks. While pass-through dynamics are nearly indistinguishable for small shocks, the menu-cost model exhibits significantly faster pass-through for large shocks. Again, the different response reflects the endogenous increase in the frequency

Figure 13: Impact of aggregate cost shocks in state- and price-dependent models



Notes. This figure presents the impulse responses of inflation and frequency to aggregate cost shocks of different magnitudes. Panel a reports the impulse response for our state-dependent pricing model (menu-cost model). Panel b reports the impulse responses for a time-dependent model (Calvo model). The x-axis reports quarters since the shock.

Figure 14: Persistence of inflation in state- and time-dependent models



Notes. This figure presents the impulse responses of aggregate inflation to marginal cost shocks of different sizes in the menu-cost model and in the Calvo model. The x-axis reports quarters since the shock.

of adjustment and the selection effect unique to the state-dependent setting.

Appendix B presents two additional quantitative exercises. In the first exercise, we study how cost shocks of different magnitudes affect both the static target price p^o_{ft} and the dynamic optimal price $p^\star_{ft}$. We show that the gap between the two prices is negligible if the cost shock is small, as expected, and remains small even when the shock is larger.

The static and dynamic reset prices are more closely aligned in the menu-cost model than in the Calvo model, consistent with our model's assumption and in line with empirical findings presented in Section 5, 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—it fluctuates based on movements in the current and expected probability of price adjustment. In contrast, the Calvo model features a fixed discounting structure, where the discount factor is simply a power function of the constant adjustment probability.

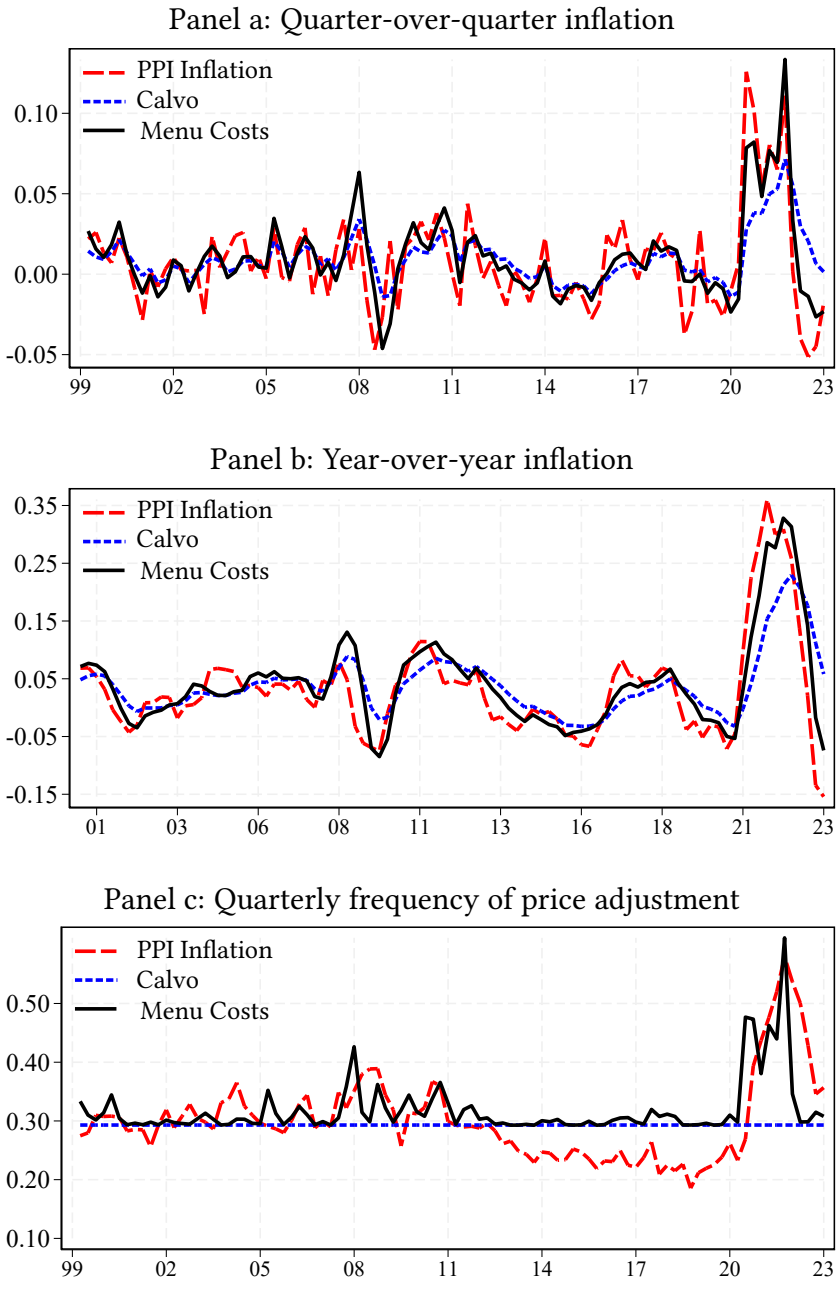
The second exercise explores the role of strategic complementarities. We compare inflation dynamics under high and low cost shocks in models with and without complementarities ($\Omega = 0$ vs. $\Omega = 0.5$). As expected, strategic complementarities dampen cost pass-through in both models. However, 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, especially with large shocks.

7.3 Explaining the time series of inflation

We now evaluate the ability of time- and state-dependent models to explain the time series of aggregate inflation observed in the data. We feed the model a sequence of aggregate cost shocks recovered from the data and simulate the model to produce a time series of aggregate inflation and the frequency of price adjustment. We perform the following quantitative exercise:

Starting in 1999:Q1, we assume that the economy is in steady state. We then feed the model a shock to the aggregate component of marginal cost, equal to the logarithmic change in our aggregate nominal marginal cost index, Δmc^n_t, between 1999:Q1 and 1999:Q2. In doing so, we maintain the model's assumption that the logarithm of the aggregate component of firms' marginal costs follows a random walk with drift. Given this shock, we solve the model and compute the new distribution of price gaps along with the resulting inflation response and frequency of price adjustment, assuming that all future

Figure 15: Inflation and frequency of price adjustment: Model versus data



Notes. This figure contrasts the dynamics of PPI manufacturing inflation in the data to the inflation dynamics generated by the Calvo and menu-cost models, after feeding the model a sequence of aggregate nominal marginal cost shocks that matched the one observed in the data.

aggregate shocks are unanticipated, as in an impulse-response function exercise. Using the updated distribution of the price gap as the new model's equilibrium, we repeat this feeding exercise for all subsequent quarters until 2023:Q4, the last period in our sample.

Figure 15 compares model simulations and 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 (black line) can capture fluctuations in manufacturing inflation well, both during the moderate inflation regime that characterized the pre-pandemic period and during the pandemic and post-pandemic inflation surge and bust. During the pre-pandemic period, the menu-cost model is nearly indistinguishable from the Calvo model, consistent with the price adjustment frequency being relatively stable over this period. In contrast, the Calvo model is capable of capturing only a modest fraction of the inflation surge in the data (about two-thirds)even though it is fed the same cost sequence as the menu-cost model. Equally important is the sluggish inflation response in the Calvo model, which, unlike the menu-cost model, fails to reproduce the rapid decline in inflation seen by the end of 2023.

Finally, panel c compares the quarterly frequency of price adjustment in the model and 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 declined 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.

8 Concluding remarks

In this paper, we study cost-price dynamics during both normal times and the recent inflation surge. Using detailed microdata on prices and costs, we construct a direct measure of firms' price gaps and provide novel nonparametric evidence on the pass-through from cost into prices for different magnitudes of shocks. Variation in price gaps explains both the likelihood that a firm adjusts its price and the magnitude of adjustment, providing strong empirical support for the predictions of our state-dependent

pricing framework.

At the macro level, the cost-price dynamics are approximately linear in normal times, when aggregate inflation is low: aggregate inflation is well approximated by the product of a fixed frequency of price adjustment and the average price gap. In contrast, during the inflation surge, the dynamics become highly nonlinear. The sharp rise in marginal costs not only increased price gaps but also triggered a significant rise in the frequency of price adjustment. This extensive margin of adjustment—a key feature of state-dependent pricing models—is absent in time-dependent frameworks such as the canonical Calvo (1983) model.

Overall, our analysis indicates that, conditional on the path of marginal cost, the state-dependent pricing model does a good job of capturing price dynamics at both the firm and aggregate levels. A natural next step is to improve the modeling of marginal cost and its connection to real activity. The conventional New Keynesian framework (e.g., Galí 2015) typically models labor as the sole variable input, implying that marginal cost is captured by the labor share. However, our analysis highlights that the sharp increase in intermediate input costs was the main driver of marginal cost variation during the inflation surge. We view extending a state-dependent New Keynesian model to incorporate intermediate inputs (e.g., primary commodities and energy) and supply chain disruptions as a promising direction for future research.

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.

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

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.

David Berger and Joseph Vavra. Dynamics of the us price distribution. *European Economic Review*, 103:60–82, 2018.

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.

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. 2024.

Luca Gagliardone, Mark Gertler, Simone Lenzu, and Joris Tielens. Anatomy of the phillips curve: Micro evidence and macro implications. *American Economic Review*, Forthcoming.

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 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.

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.

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. 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.

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.

Micro and macro cost-price dynamics in normal times and during inflation surges

L. Gagliardone M. Gertler S. Lenzu J. Tielens

Appendix

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):

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

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:

$$\min_{Y_t(f)} \int_0^1 \widetilde{P}_t(f) Y_t(f) df \qquad \text{s.t.} \int_0^1 \Upsilon\left(\frac{Y_t(f)}{Y_t}\right) df = 1.$$

where $\widetilde{P}_t(f) \equiv \frac{P_t(f)}{e^{\varphi_t(f)}}$ is the quality-adjusted price. Denoting by ψ the Lagrange multiplier of the constraint, the first-order condition of the problem is:

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

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 $\phi \equiv \Upsilon \circ (\Upsilon')^{-1}$. Evaluating the first-order condition (A.1) at symmetric prices, $\widetilde{P}_t(f) = P_t$, we get $\psi = \frac{P_t Y_t}{\Upsilon'(1)}$. Replacing for ψ, we recover the demand function:

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

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

$$\mathcal{D}_t(f) = (\Upsilon')^{-1}\left(\Upsilon'(1)\frac{\widetilde{P}_t^o(f)}{P_t}\right)Y_t$$

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

$$\varsigma_t(f) \equiv -\frac{\partial \ln \mathcal{D}_t(f)}{\partial \ln \widetilde{P}_t^o(f)} = -\frac{\Upsilon'\left(\frac{Y_t(f)}{Y_t}\right)}{\Upsilon''\left(\frac{Y_t(f)}{Y_t}\right)\cdot\left(\frac{Y_t(f)}{Y_t}\right)}. \tag{A.3}$$

We now use this result to derive the expression for the log-linearized desired markup. As above, for ease of exposition, we focus on the symmetric steady state. Denote the steady-state residual demand elasticity by $\varsigma = -\frac{\Upsilon'(1)}{\Upsilon''(1)}$. Then the derivative of the residual demand elasticity $\varsigma_t(f)$ in (A.3) with respect to $\frac{Y_t(f)}{Y_t}$, evaluated at the steady state, is given by:

$$\varsigma' = \frac{\Upsilon'(1)\left(\Upsilon'''(1)+\Upsilon''(1)\right)-\left(\Upsilon''(1)\right)^2}{\left(\Upsilon''(1)\right)^2} \leq 0, \tag{A.4}$$

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.4), we have that, up to a first-order approximation, the log-markup (in deviation from the steady state) is equal to:

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

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

$$\mu_t(f) - \mu(f) = -\Gamma\left(\widetilde{p}_t^o(f) - p_t\right)$$

where, in the case of Kimball preferences, the sensitivity of the markup to the relative price is given by $\Gamma \equiv \frac{\varsigma'}{\varsigma(\varsigma-1)}\frac{1}{\Upsilon''(1)}$. Finally, replacing the log-linearized markup into the formula for the static optimal target price (obtained from cost minimization):

$$\begin{aligned} p_t^o(f) &= \mu_t(f) + mc_t^n(f) \\ &= (1-\Omega)(\mu(f) + mc_t^n(f)) + \Omega(p_t + \varphi_t(f)) \end{aligned}$$

where $\Omega \equiv \frac{\Gamma}{1+\Gamma}$ is the degree of strategic complementarities.

A.2 Derivation of the optimal reset gap

Under the quadratic profits, the problem of the firm at time t is:

$$\max_{x} \quad -B(x)^2 + \beta\, \mathbb{E}_t \left\{ h_{t+1}(x) \cdot V^a_{t+1} + \big(1 - h_{t+1}(x)\big) \cdot V_{t+1}(x) \right\},$$

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

$$Bx^\star_t = \beta \mathbb{E}_t \left\{ (1 - h_{t+1}(x^\star_t)) \frac{\partial V_{t+1}(x)}{\partial x}\Big|_{x=x^\star_t} + (V^a_{t+1} - V_{t+1}(x^\star_t)) \frac{\partial h_{t+1}(x)}{\partial x}\Big|_{x=x^\star_t} \right\}.$$

Because the adjustment probability is minimized at $x^\star_t$, the condition simplifies to:

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

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

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

Repeatedly replacing into the first-order condition, we obtain:

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

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

$$\begin{aligned} p^\star_t(f) &= (1-\Omega)(\mu(f) + mc^n_t(f)) + \Omega p_t + \Omega \frac{\mathbb{E}_t\{\sum_{i=1}^{\infty} (p_{t+i} - p_t)\beta^i \prod_{\tau=1}^{i}(1 - h_{t+\tau})\}}{\mathbb{E}_t\{\sum_{i=0}^{\infty} \beta^i \prod_{\tau=0}^{i}(1 - h_{t+\tau})\}} \\ &= p^o_t(f) + \Omega \Psi_t. \end{aligned} \tag{A.5}$$

The equation above decomposes the optimal (dynamic) reset price into two terms. The first is the static reset price, $p^0_{ft} \equiv (1-\Omega)(\mu(f) + mc^n_t(f)) + \Omega p_t$, 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 $p^\star_{ft}$ as the firm anticipates that the price set today may also apply to future periods due to nominal rigidities.[20]

Finally, under our assumption that costs follow a random walk and i.i.d. taste shocks, the second term in Equation (A.5) (i) does not depend on the identity of the

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

firm, (ii) is exactly zero in the absence of strategic complementarities ($\Omega = 0$), and (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 Quadratic approximation of generalized hazard function

We now derive the expression for the quadratic approximation of the hazard function in Equation (8) and describe how we take this equation to the data. 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:

$$\begin{aligned} h_t(x_{t-1}) &= (1-\theta^0) - \frac{\theta^0}{\bar{\chi}} \frac{\partial V_t(x)}{\partial x}\Big|_{x=x_t^\star} (x_{t-1} - x_t^\star) - \frac{\theta^0}{2\bar{\chi}} \frac{\partial^2 V_t(x)}{\partial x^2}\Big|_{x=x_t^\star} (x_{t-1} - x_t^\star)^2 + o(x_{t-1} - x_t^\star)^2 \\ &= (1-\theta^0) - \frac{\theta^0}{2\bar{\chi}} \frac{\partial^2 V_t(x)}{\partial x^2}\Big|_{x=x_t^\star} (x_{t-1})^2 + o(x_{t-1})^2, \end{aligned}$$

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 $\phi \equiv -\frac{\theta^0}{2\bar{\chi}} \frac{\partial^2 V(x)}{\partial x^2}\Big|_{x=0}$, we have that the GHF can be approximated, up to second order, by a quadratic function of the price gap as in Equation (8):

$$h(x_{t-1}(f)) = (1-\theta^o) + \phi \cdot ((x_{t-1}(f))^2 + o\big((x_{t-1}(f))^2\big), \tag{A.6}$$

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

A.4 Estimator of the probability of free adjustment

First, we derive the expression for the average frequency of price adjustment in a bin. 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:

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

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:

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

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^0. 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:

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

We want to show that the integral on the LHS converges to the RHS as the interval $(-b', b')$ 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 $1/n$, for $n = K, K+1, \ldots$, with $N < K \in \mathbb{N}_+$. Then the sequence:

$$1 - \int_{\{-\frac{1}{n},\ldots,0,\ldots,\frac{1}{n}\}} h(b)db$$

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:

$$1 - \lim_{n\to\infty} \int_{\{-\frac{1}{n},\ldots,0,\ldots,\frac{1}{n}\}} h(b)db = 1 - h(\tilde{b}) = \theta^o.$$

For a sufficiently small interval around $\tilde{b}$, the mean of the frequencies of price adjustment over that interval recovers the probability of free adjustment.

A.5 Cubic approximation of inflation within a bin

Starting from the expression for aggregate inflation in Equation (9), we want to derive the cubic expression for inflation within a bin in Equation (13), under the assumption that $p_t^\star(f) = p_t^o(f)$ and associated cross-sectional regression model. Again we partition the distribution of price gaps into equal frequency bins (quantiles) denoted by b and adopt the same labeling convention of bins described in Appendix A.4. Denote by γ_b the skewness within a bin. 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 $x_f \leq \delta_f$:

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

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 $f \in b$:

$$\left| \frac{\mathbb{E}_b(h_f \cdot x_f) - (1-\theta^0) x_b - \phi\, \mathbb{E}_b(x_f^3)}{x_b^3} \right| \leq \epsilon.$$

Hence:

$$\mathbb{E}_b(h_f \cdot x_f) = (1-\theta^0) x_b + \phi\, \mathbb{E}_b(x_f^3) + o(x_b^3)$$

Using this approximation, the covariance within a bin satisfies:

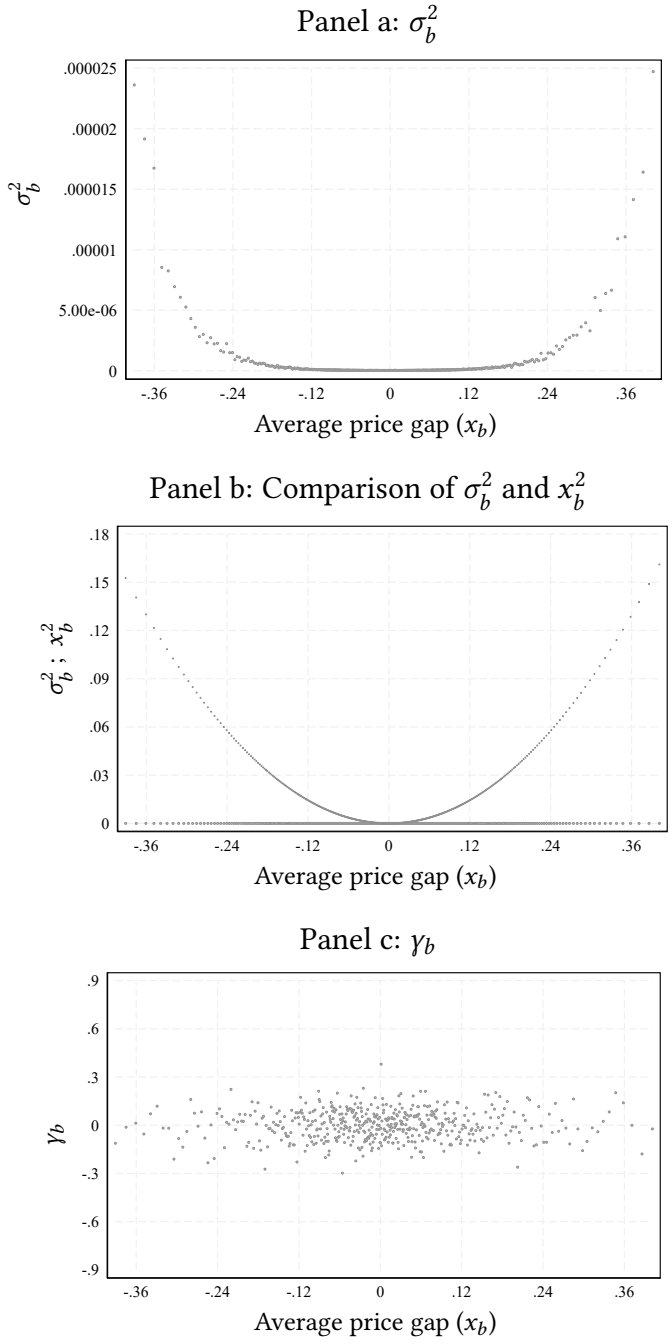
$$\begin{aligned}Cov_b(h_f, x_f) &= \phi\, \mathbb{E}_b(x_f^3) + o(x_b^3) - \phi\, \mathbb{E}_b(x_f^2) x_b - o(x_b^2) x_b \\ &= \phi(2x_b\sigma_b^2 + \gamma_b \sigma_b^3) + o(x_b^3)\end{aligned}$$

It follows that inflation within a bin simplifies to:

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

A.6 Variance and skewness across the price gap distribution

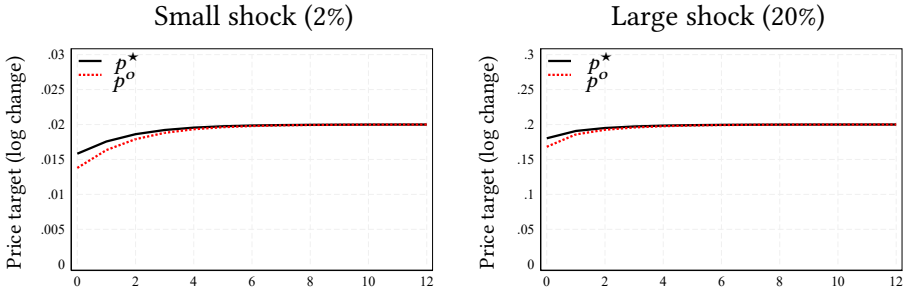
Figure A.1: Within bin variance, square of the mean, and skewness



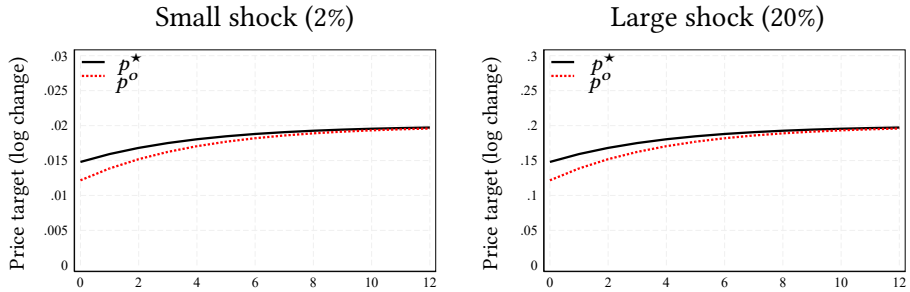
Notes. This figure plots three moments of the price gap distribution by bin (quantile): the within-bin variance of price gaps (σ_b^2, panel a), the square of the within-bin average gap (x_b^2, panel b), and the within-bin skewness of price gaps (γ_b, panel c). In panel b, x_b^2 (crosses) is plotted alongside 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)



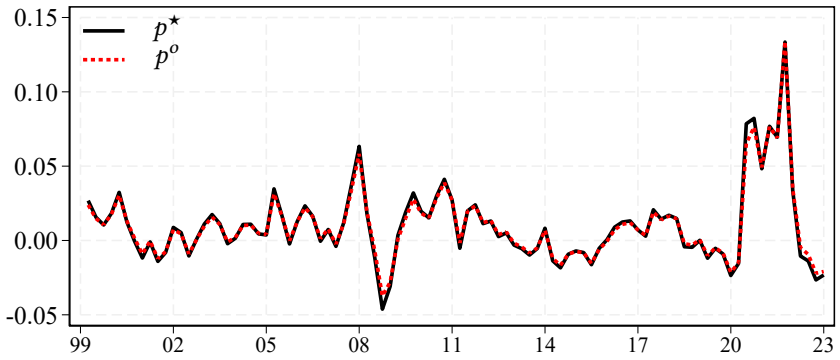
Notes. This figure presents impulse responses of the static target price (p^o) and the optimal reset price ($p^\star$) to aggregate cost shock of different sizes. The x-axis reports quarters since the shock.

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 (9) and (13), 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 static reset price responds more than the static one to cost shocks, since the dynamic optimum $p^\star_{ft}$ accounts for the 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 two prices is negligible if the shock is small, as expected, and

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



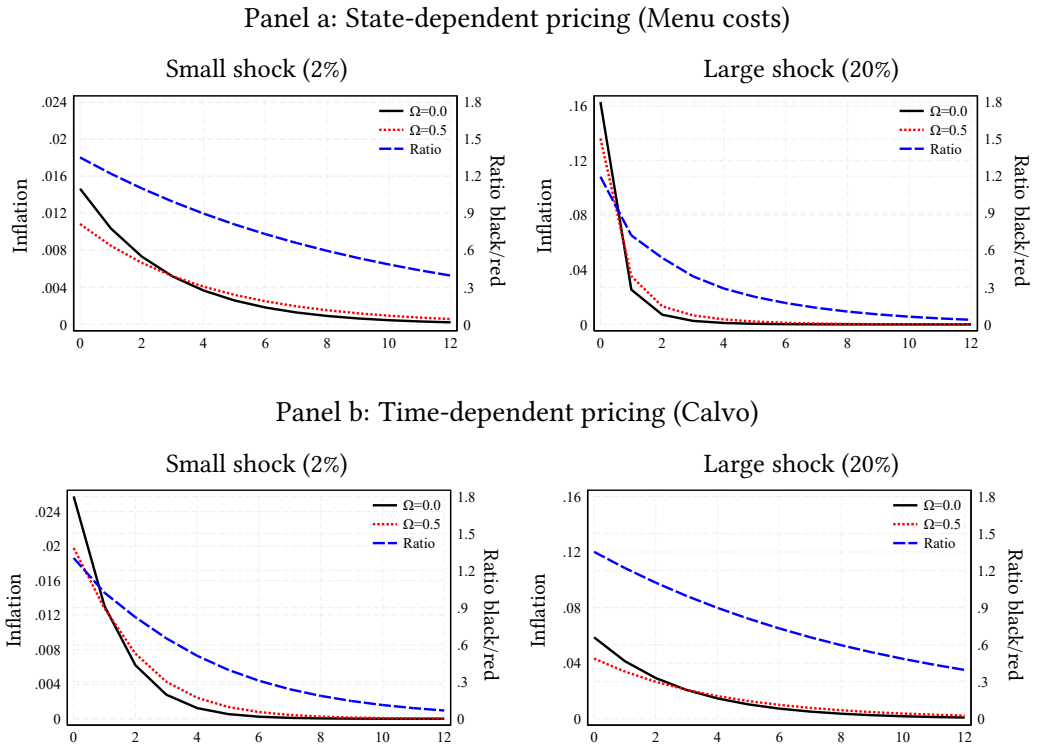
Notes. This figure contrasts the inflation dynamics generated by the menu-cost model using $p^\star$ (the exact, dynamic reset price) and using p^o (the static approximation of $p^\star$) when solving the model. As in Figure 15, we solve the model feeding it a sequence of aggregate nominal marginal cost shocks that matched the one observed in the data.

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 15. 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 the setting of prices are a factor that contributes to explaining the differential 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 with and without complementarities is narrower in the menu-cost

Figure A.4: The role of strategic complementarities



Notes. This figure presents the impulse responses of inflation to aggregate cost shocks of different sizes, without strategic complementarities ($\Omega = 0$, black line) and with strategic complementarities ($\Omega = 0.5$, red dotted line). The blue dashed line represents the ratio of the impulse response under $\Omega = 0$ over the impulse response under $\Omega = 0.5$. Panel a reports the impulse response for our state-dependent pricing model (menu-cost model). Panel b reports the impulse responses for a time-dependent model (Calvo model), calibrated to display the same steady-state frequency of price adjustment as the state-dependent model. The x-axis reports quarters since the shock.

model, especially in response to a large shock. This is due to the greater curvature of the value function under state-dependent pricing.