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

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

February, 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, pricing policies display state-dependence: 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, adjustment probabilities are 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 Dan Cao, Francesco Lippi, Virgiliu Midrigan, John Rust, Luminita Stevens, and Ludwig Straub 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.

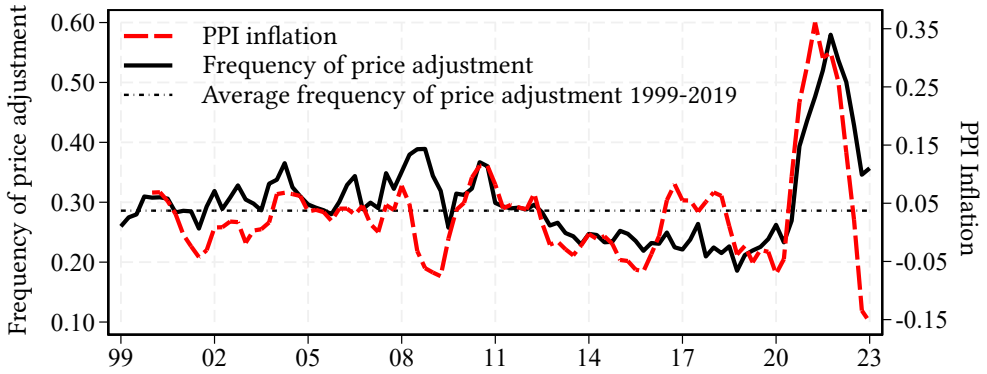
Firms adjust output prices infrequently despite continuously evolving economic conditions, leading their prices to drift temporarily from those that maximize flow profits. In all models featuring nominal rigidities, the deviation between listed and optimal prices—the *price gap*—determines firms' pricing 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 price gap. In contrast, in state-dependent pricing models (e.g., 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 points. 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, consistently with a linear mapping between expected price changes and price gaps. However, as observed globally, starting in early 2021, the frequency fluctuated sharply alongside a surge in inflation—hallmarks of state-dependent pricing and, as we will show, nonlinear cost-price dynamics.[1]

In this paper, we construct an empirical measure of price gaps using microdata on firms' prices and costs. This measure allows us to study the mapping between gaps and price changes in the cross-section of firms and derive its implications for the aggregate cost-price dynamics in the time series. The state-dependent nature of firms' pricing policies is the salient feature emerging from the cross-sectional analysis, which implies a nonlinear cost-price dynamics when large aggregate shocks induce correlated fluctuations in desired prices. In contrast, when aggregate shocks are small, the mapping between

[1] See Blanco et al. (2024b) and Cavallo et al. (2024) for evidence of similar dynamics in the U.S. and other developed economies.

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.

expected price changes and price gaps is observationally equivalent to that predicted by a time-dependent model, resulting in linear cost-price dynamics.[2] At the aggregate level, we show that a state-dependent model calibrated to match cross-sectional cost-price dynamics and fed with an aggregate cost sequence derived from microdata accurately accounts for inflation in both normal times and during inflation surges.

Our dataset collects administrative records on product-level output quantities, sales, and production costs at the *quarterly frequency* for Belgian manufacturing firms between 1999 and 2023. Using these data, we construct a notion of price gaps for individual firms that accounts for the variation in costs, prices, and competitors' prices. We analyze the data through the lens of a tractable menu-cost model in the tradition of the classic generalized state-dependent models proposed by Caballero and Engel (1993, 2007), Midrigan (2011), Nakamura and Steinsson (2010), Alvarez et al. (2016) and, more recently, Auclert et al. (2024).[3] The model nests a time-dependent Calvo (1983), as a special case. The quantitative framework serves three purposes: deriving testable predictions linking price changes to price gaps in microdata; establishing novel identification results to recover key primitive parameters from cross-sectional moments of the joint distribution of price changes and price gaps; and explaining aggregate inflation over time as a function

[2]See Dias et al. (2007), Gertler and Leahy (2008), Alvarez et al. (2017), and Auclert et al. (2024).

[3]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.

of aggregate costs.

Our analysis produces two main sets of results. At the micro-level, we document strong evidence in favor of the state-dependent nature of firms' pricing decisions. First, we show that the frequency at which firms change prices—the extensive margin—increases in the absolute value of their price gaps and hence can be well approximated by a quadratic function. Second, the mapping between expected price changes and price gaps is "S-shaped", approximately linear at small gaps with a slope that matches the average frequency of price adjustment, and highly nonlinear at large gaps with a steeper slope. Third, when firms change prices, they do so to close the gap. Fourth, focusing on the pandemic and post-pandemic period, we show how large aggregate cost shocks shifted the entire distribution of price gaps, displacing many firms away from their optimal target and inducing large and correlated price adjustments along both the intensive and extensive margins. On the other hand, this mechanism is not at work in "normal times". In the pre-pandemic period, characterized by low inflation, the price gap distribution is stable, the frequency of price adjustment is roughly constant, and therefore the relationship between price changes and gaps is linear. Altogether, these micro-facts indicate that the cost-price dynamics can be well approximated by a time-dependent model in normal times, but state dependence is needed to rationalize the effects of large aggregate shocks.

The second set of results pertains to the accounting of aggregate inflation in the time series. Leveraging our microdata, we construct an aggregate cost index for the Belgian manufacturing sector. Descriptive evidence shows how inflation and production costs align closely throughout the entire sample period. However, there is stickiness in price adjustment such that inflation moves less than costs. We also show that the sharp rise and fall in costs (and intermediates cost, in particular), rather than a change in markups, appears to be the main driver of the surge and subsequent drop in inflation observed in the post-pandemic period. We then formally assess the capacity of our menu-cost model to explain aggregate inflation. We compute a model-based inflation series after feeding into the model our marginal cost index. Comparing this sequence with the data, we show that the model tracks the high-frequency fluctuations in Belgian manufacturing inflation remarkably well, both during the moderate pre-pandemic regime and during the post-pandemic surge. Remarkably, the model captures the stable behavior of the adjustment frequency pre-pandemic as well as the sharp jump following the onset of the

pandemic, both in terms of timing and magnitude. In contrast, a standard Calvo model, fed with the same cost sequence, accounts well for inflation in normal times but can explain only about two-thirds of the inflation surge during the its recent surge.

Related literature. Earlier research provides evidence of the state-dependent nature of firm pricing decisions (Klenow and Kryvtsov 2008; Gagnon 2009; Gautier and Saout 2015) and shows how this class of models can rationalize the empirical distribution of price changes in the data (see the aforementioned references and Alvarez et al. 2022). Relatedly, Alvarez et al. (2019) and Karadi and Reiff (2019) present evidence of state-dependent pricing through case studies of hyperinflation in Argentina and major tax shocks in Hungary, respectively. More recently, Blanco et al. (2024a, 2024b), Bunn et al. (2024), Cavallo et al. (2024), Gagliardone and Tielens (2024), and Morales-Jiménez and Stevens (2024) apply state-dependent frameworks to analyze the recent inflation surge. The key distinction between these studies and ours lies in our ability to construct a high-frequency measure of price gaps at the firm level. As we have emphasized above, this is the fundamental building block of both time- and state-dependent pricing models. By analyzing how the size and frequency of price adjustment relate to price gaps in the cross-section of firms and across different economic cycles, we can directly assess the degree to which firms' pricing strategies conform with the predictions of different theories.

Our study also relates to the works of Eichenbaum et al. (2011) and Karadi et al. (2024). The former uses data on prices and costs from a large food and drugs retailer to develop a "reference price" metric. The latter employs microdata on supermarket prices to formulate a notion of reset price, derived from the average price at which the same product is offered by rivals. In our dataset, we can observe high-frequency cost and price data for the entire Belgian manufacturing sector over almost three decades. This allows us to construct an empirical measure of firm-level reset prices and price gaps that factors in both the firms' costs and the pricing of their competitors.

Finally, the results in this paper connect with our earlier work Gagliardone et al. (Forthcoming) on the estimation of the slope of the cost-based New Keynesian Phillips curve. Using microdata for the pre-pandemic period, we used a time-dependent Calvo model to identify the structure parameters that enter the slope. The findings in this paper lend additional empirical support to that identification strategy by showing that the

relationship between price gaps and price changes is approximately linear in the absence of large shocks, due to the stability of the adjustment frequency.

The paper proceeds as follows. Section 1 presents the theoretical framework and derives testable implications. Section 3 describes our dataset and the empirical measures of prices, cost, and price gaps. Section 4 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 6, showing how the calibrated model explains the inflation time series and the frequency of price adjustment, including the rise during the pandemic. Section 7 offers concluding remarks.

1 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).[4] 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.

1.1 A tractable state-dependent pricing model

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.

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

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 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.

Firms' nominal marginal cost is given by:

$$mc_t(f) = mc_t + z_t(f). \tag{2}$$

The term mc_t captures an aggregate nominal cost shifter. Consistent with the empirical evidence, we assume that mc_t obeys a random walk $mc_t = mc_{t-1} + g_t$, where g_t captures an aggregate shock, i.i.d. over time with mean μ_g, drawn from a single-peaked, symmetric, and smooth distribution. For analytical tractability, in what follows, we assume no trend inflation ($\mu_g = 0$).[5] We relax this assumption in the quantitative exercises of Section 6.

Profit maximization. Firms choose prices to maximize the present value of profits, subject to nominal rigidities. Each firm pays a fixed cost $\chi_t(f)$ when adjusting its price from the price charged in the previous period. 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 $[0, \overline{\chi}]$. As in the CalvoPlus model, we also assume that 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:

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

where the price index accounts for strategic complementarities in price setting. The scalar $\Omega \in [0, 1)$ captures the strength of such complementarities. The taste shock $\varphi_t(f)$ shows

[5] 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.

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 price from the target:

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

where σ is the steady-state price elasticity of demand and steady-state profits are equal to $1/\sigma$. Note 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:

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

The optimal pricing policy reduces to determining the *optimal probability of price adjustment*, denoted by $h_t(f)$, 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.[6] As is standard in state-dependent models, the solution of the firm problem has a "Ss flavor".

We now define the key object in our empirical analysis, the firm's (ex-ante) *price gap* in period t, $x_{t-1}(f)$. The gap serves as a state variable of the firm's problem as it captures the difference between the target price (after the realization of shocks in period t) and the price set by the firm in the previous period $(t-1)$:

$$\begin{aligned} x_{t-1}(f) &\equiv p_t^o(f) - p_{t-1}(f) \\ &= \Big(p_{t-1}^o(f) + (1-\Omega)(g_t + \varepsilon_t(f)) + \Omega(p_t - p_{t-1})\Big) - p_{t-1}(f). \end{aligned} \tag{4}$$

[6]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.

The second line follows from replacing $p_t^o(f)$ using Equation (3), replacing $mc_t(f)$ using Equation (2), and then using the expressions describing the processes for the aggregate and idiosyncratic components of $mc_t(f)$. The price gap $x_{t-1}(f)$ is measured 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 $p_t^o(f)$. Here, $\varepsilon_t(f) \equiv \zeta_t(f) + \frac{\Omega}{1-\Omega}\varphi_t(f)$ denotes a composite idiosyncratic shock with mean zero and variance denoted by σ_ε^2, which combines idiosyncratic technology and taste shocks. We assume that $\varepsilon_t(f)$ is drawn from a unimodal, symmetric, and smooth distribution. Finally, due to pricing complementarities, the inflation rate $p_t - p_{t-1}$, enters the price gap because it affects the evolution of competitors' prices.

Let $h_t(x_{t-1})$ be the probability that a firm adjusts the price at t conditional on its price gap. Then the solution to the firm's problem can be expressed as a function of the price gap:

$$p_t^o(f) - p_t(f) = \begin{cases} x_t^\star & \text{w. p. } h_t(x_{t-1}) \\ x_{t-1}(f) & \text{w. p. } 1 - h_t(x_{t-1}). \end{cases} \tag{5}$$

Firms adjust their price with probability $h_t(x_{t-1})$. Upon adjustment, they set their price to $p_t^\star(f)$. If they do not adjust their price, they keep their gap at $x_{t-1}(f)$.

We now characterize the optimal reset probability and the optimal reset gap. The derivations are provided in Appendix A.

Probability of price adjustment. As in a standard "Ss" framework, the adjustment probabilities are endogenous variables that depend on the distance between the optimal reset gap $x_t^\star$ and the price gap $x_{t-1}(f)$.

Let V_t^a be the firm's value if it resets its price to $p_t^\star(f)$ and $V_t(x_{t-1}(f))$ its value if it does not. As we show below, the former depends on $x_t^\star(f)$ while the latter is a function of $x_{t-1}(f)$. The probability that a firm adjusts its price positively depends on the gap between the two values. 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:

$$\begin{aligned} h_t(x_{t-1}) &= (1-\theta^o) + \theta^o \cdot \Pr(V_t^a - \chi_t(f) \geq V_t(x_{t-1})) \\ &= (1-\theta^o) + \theta^o \cdot \min\left\{\frac{V_t^a - V_t(x'_{t-1})}{\overline{\chi}}, 1\right\}, \end{aligned} \tag{6}$$

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:*

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

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 local 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. We now characterize $V_t(x_{t-1})$, V_t^a, and therefore $x_t^\star$. As discussed, $p_t^o(f) = x_{t-1}(f) + p_{t-1}(f)$ for a firm that does not adjust its price. In this case, the value of the firm is given by:

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

It is a function of current profits Π_t, evaluated at the price gap $x_{t-1}(f)$, and of the discounted expected continuation value. The latter depends on the probability of adjustment at time $t+1$, $h_{t+1}(x_t)$. The value of the firm conditional on adjusting is the

optimized value of V with respect to the reset price $p_t^\star$:

$$V_t^a = \max_{p_t^\star} V_t\big(p_t^o(f) - p_t^\star\big).$$

Equivalently, the optimal reset gap $x_t^\star$ solves the first-order condition $V_t(x_t^\star) = 0$.

Under our assumptions of no trend inflation and a quadratic profit function, $x_t^\star \approx 0$ (see, e.g., Alvarez et al. 2016). The absence of trend inflation implies that the static optimal price provides a good approximation of the dynamic optimal price ($p_t^\star(f) \approx p_t^o(f)$). If there are no strategic complementarities, the approximation is exact.[7] For our purposes, this result has important practical implications. Our data allow us to construct a measurable counterpart of the price gap $x_{t-1}(f)$ in terms of observables, as Equations (3) and (4) suggest, which allows us to directly test the implications of the model in the microdata. In the analytical exercises that follow, we assume that $x_t^\star \approx 0$. In Section 5, we verify numerically that this is a good approximation.

Aggregate inflation. Next, we describe the implications of firm-level price adjustment for aggregate inflation. Given the solution of the firm's problem in Equation (5) and using the formula for the price index in Equation (1), we can express aggregate inflation π_t as:

$$\begin{aligned}\pi_t &= \int \Big(p_t(f) - p_{t-1}(f)\Big)\, df = \int h_t(x_{t-1}(f)) \cdot \big(p_t^\star(f) - p_{t-1}(f)\big)\, df \\ &= \int h_t(x_{t-1}(f))\, df \cdot \int \Big(p_t^\star(f) - p_{t-1}(f)\Big)\, df + Cov\Big(h_t(x_{t-1}(f))\,,\ (p_t^\star(f) - p_{t-1}(f))\Big) \quad (8)\end{aligned}$$

The first line shows that, to a first-order approximation, aggregate inflation is an average of firm-level price adjustment, which can be expressed 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 probability is an endogenous object that, as we will see, increases nonlinearly with the absolute value of the price gap. With

[7] Intuitively, under our assumptions, the combined shocks that affect firms' pricing decisions (i.e., the sum of aggregate and idiosyncratic shocks) is a highly persistent variable that approximately evolves as a random walk. Therefore, the optimal dynamic price $p_t^\star(f)$ remains very close to the static optimum $p_t^o(f)$. Under strategic complementarities, the optimal dynamic price is a linear combination of the expected present values of future costs and future inflation rates. Hence, it is more volatile than the static price because inflation does not follow a random walk. See Alvarez et al. (2023) for a treatment of menu-cost models with strategic complementarities in continuous time.

Calvo pricing, the adjustment probability is fixed and constant between firms; 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 (8). 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.

2 Testable implications and identification results

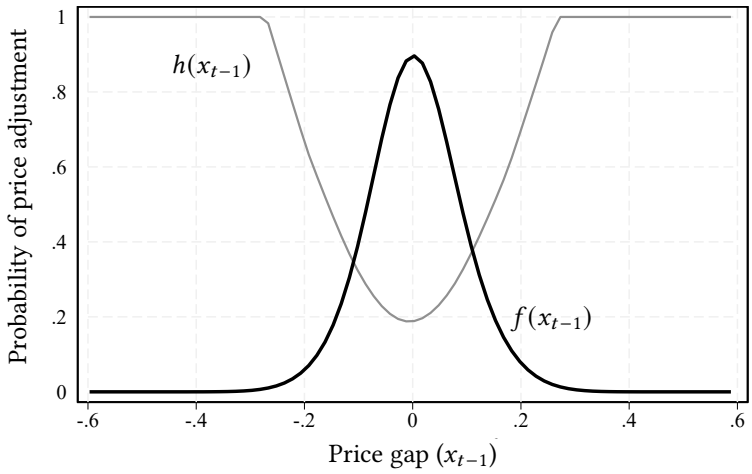
We now discuss testable implications of our model and derive identification results used for calibration purposes.

2.1 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 ($x_b \equiv$

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 $f(x_{t-1})$.)

$\int_{f\in b} x_{t-1}(f)df$) 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.

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

$$h_b \equiv \int_{f\in b} h(x_{t-1}(f))df = (1-\theta^o) + \phi\left(x_b^2 + \sigma_b^2\right) + o(x_b^2), \tag{9}$$

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 ($1-\theta^o$):

$$h_b = a_0 + a_1 \cdot x_b^2 + u_b, \tag{10}$$

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. We will rely on this identification result for model calibration, as detailed in Section 6.1.

2.2 The mapping from price gaps to price changes

Endogenous fluctuations in the probability of price adjustment are the main driver of nonlinear shock transmission in state-dependent models. To illustrate this and establish testable implications, we derive the mapping from price gaps to price changes along the price gap distribution.

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 $b \in \mathcal{B}$.[8] The following proposition characterizes inflation within a bin (i.e., the average within-bin price change) as a function of the average price gap of the bin:

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

$$\pi_b \equiv \int_{f\in b} \big(p_t(f) - p_{t-1}(f)\big) df = \phi_b^0 \cdot x_b + \phi \cdot x_b^3 + \omega_b. \tag{11}$$

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

Proof. See Appendix A.5. □

Proposition 2 states that the inflation rate within a bin can be well approximated by a third-order (odd) polynomial of the gaps. It is straightforward to derive the analog of Equation (11) in the case of a time-dependent Calvo model. Given a constant exogenous hazard rate $h^c := (1-\theta^c)$, we have:[9]

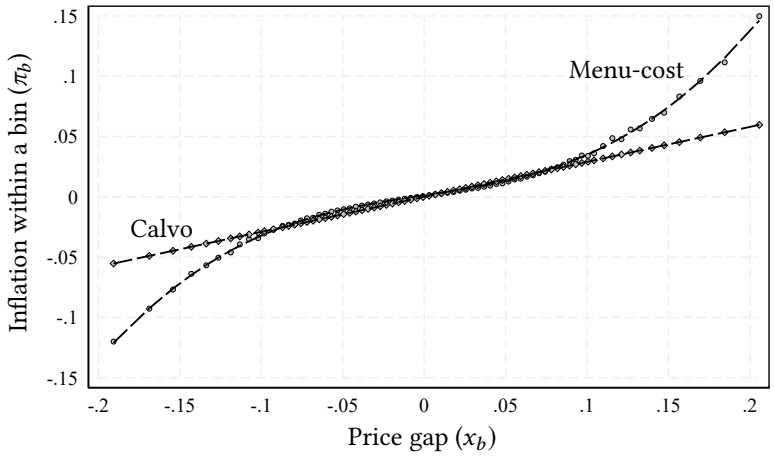
$$\pi_b = (1-\theta^c) \cdot x_b. \tag{12}$$

The binned scatterplot in Figure 3 visually represents the mapping from price gaps to price change in menu-cost model (gray circles) and in the Calvo model (gray diamonds). The dashed lines show fitted values from a regression of bin-level inflation

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

[9] In a Calvo model, $\mathbb{E}[p_t(f)|\mathcal{I}_t(f)] = \mathbb{E}[p_t(f)|p_t^\star(f), p_{t-1}(f)] = (1-\theta^c)p_t^\star(f) + \theta^c p_{t-1}(f)$, 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 3: Nonlinear price dynamics

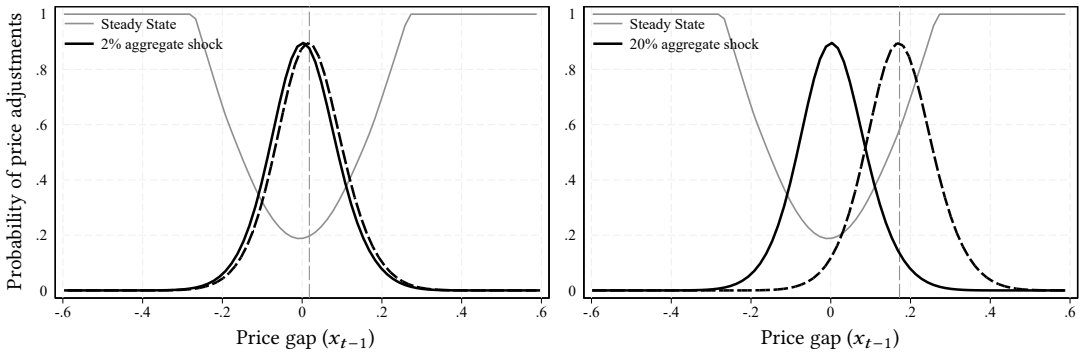


Notes. In this figure, we partition the distribution of price gaps into quantiles (b) and plot the average price gap of each bin, x_b, against the average logarithmic price change for observations in the same bin, π_b. The gray dots represent simulations of our menu-cost model, while the gray diamonds 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 (11) 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 (12) for the Calvo model.

on either a polynomial including the first and third orders of the average gap (Equation (11)) or a first-order polynomial in the average gap (Equation (12)). When price gaps are sufficiently close to zero, the third-order term becomes negligible, making inflation within a bin proportional to the corresponding price gap. As a result, firms operating near their optimal price exhibit linear pricing dynamics in both state- and time-dependent models. This result underlies the approximate equivalence of these two classes of models in response to small shocks, as illustrated in Gertler and Leahy (2008), Alvarez et al. (2017), and Auclert et al. (2024).

The cubic term arises from the state dependence of firms' pricing decisions and, as gaps widen, introduces the potential for nonlinear inflation dynamics. It captures the fact that firms in the tails of the price gap distribution are more likely to adjust prices. In this case, the data-generating process for inflation within a bin can be approximated by a cubic expression in Equation (11).

Figure 4: 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).

2.3 The impact of aggregate cost shocks

Aggregate cost shocks, that is, shocks that do not average out, impact optimal reset prices of all firms in the economy. When these shocks are large, the entire distribution of price gaps shifts, and a substantial number of firms are displaced into regions of the price gap distribution where the GHF is steep. This displacement increases the degree of monetary neutrality in the economy.

Figure 4 illustrates this point. In the spirit of the exercise in Cavallo et al. (2024), we shock the economy in its stationary equilibrium with an unexpected cost shock $g_t > 0$, which increases the marginal cost for all firms. The left panel shows the effect of a small shock, while the right panel shows the effect of a large shock. The solid lines represent the GHF before the aggregate shocks. The dashed lines show the post-shock distributions.

A small aggregate shock induces a small shift in the price gap distribution, with little or no impact on the average probability of price adjustment. In contrast, pushes a significant number of firms away from their target price, widening the average gap in the economy. This, in turn, leads to a substantial rise in adjustment frequency: more firms seek to raise prices, while fewer opt to lower them. Under these conditions, the cubic term in Equation (1) matters, driving aggregate inflation beyond what is explained by the increase in the average gap alone.

3 Data and measurement

We assemble a micro-level dataset that covers the manufacturing sector in Belgium between 1999 and 2023. The dataset is compiled from administrative sources and contains information on firms' pricing and production decisions at the business cycle (quarterly) frequency, with a detailed snapshot of firm's variable production costs (labor costs and intermediates). This dataset extends and enriches the one in Gagliardone et al. (Forthcoming) in two significant ways.[10] First, the data in Gagliardone et al. (Forthcoming) cover a period characterized by low and stable inflation (1999:Q1 to 2021:Q1). We extend the time-series dimension to include the recent inflation surge and subsequent tapering (2021:Q2 to 2023:Q4). Second, we merge new microdata that allow us to accurately measure the frequency of price adjustment.

3.1 Measurement of prices, costs, and price gaps

The unit of observation in our data is a firm-industry pair. Our final dataset tracks 5,348 domestic firm-industry pairs, denoted by a lower-script f, distributed 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:[11]

$$\frac{P_{ft}}{P_{ft-1}} = \prod_{p \in \mathcal{P}_{ft}} \left(\frac{P_{pt}}{P_{pt-1}} \right)^{\bar{s}_{pt}}, \tag{13}$$

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

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

[11] PRODCOM surveys all Belgian firms involved in manufacturing production with more than 10 employees, covering over 90% of production in each NACE 4-digit industry. We recover domestic values and quantities sold by combining information from PRODCOM with international trade data on firms' product-level exports (quantities and sales).

domestic price changes of products sold by domestic and international competitors selling in the same industry as f ($\mathcal{F}_i$):

$$\frac{P_{it}^{-f}}{P_{it-1}^{-f}} = \prod_{k \in \mathcal{F}_i \setminus f} \left(\frac{P_{kt}}{P_{kt-1}} \right)^{\bar{s}_{kt}^{-f}}, \tag{14}$$

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

Finally, we recover the times series of firms' prices (in levels) by concatenating the price indices in Equation (13), $P_{ft} = P_{f0} \prod_{\tau=t_f^0+1}^{t} \left(P_{f\tau}/P_{f\tau-1} \right)$, 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 (14).

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:

$$\mathbb{I}^{+}_{ft} = 0 \iff \Delta p_{ft} < \kappa^{+}_{h} \cdot Var_f(\Delta p_{ft}) \quad \text{if } \Delta p_{ft} > 0$$
$$\mathbb{I}^{-}_{ft} = 0 \iff \Delta p_{ft} > -\kappa^{-}_{h} \cdot Var_f(\Delta p_{ft}) \quad \text{if } \Delta p_{ft} < 0$$

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: $\sum_t \sum_f \mathbb{I}^{+}_{ft} = h^{+}$ and $\sum_t \sum_f \mathbb{I}^{-}_{ft} = h^{-}$, where $h^{+} + h^{-} = h$.[12]

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: $MC^{n}_{ft} = (1 + \nu_f)AVC_{ft}$. 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:

$$mc^{n}_{ft} = (tvc^{n}_{ft} - y_{ft}) + \ln(1 + \nu_f). \tag{15}$$

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. We compute a quantity index by dividing a firm's domestic revenues by its domestic price index.[13] Firm-specific short-run returns to scale are not directly observable in the data. Therefore, to the extent that individual firms' production technologies might deviate from constant returns to scale ($\nu_f \neq 0$), our measure of log-marginal costs would be missing an additive 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.

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

[13] 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 lion's share of the 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 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 1, 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 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 measurement error that possibly weakens the connection between our empirical measures of price gaps and price changes at the firm level.[14]

3.2 Harmonization and data cleaning

Previous literature highlighted how issues related to measurement error and unobservable cross-sectional heterogeneity can lead to a substantial bias in statistics that describe the distribution of price changes and *a fortiori*, of price gaps.[15] To address these issues, we follow the literature and apply the following data-cleaning steps and harmonization procedures.

The use of unit values often introduces spurious price changes, either incorrectly indicating small price changes where none occurred or reporting abnormally large price changes (Eichenbaum et al. 2011). To mitigate the impact of measurement error, we set

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

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

price changes smaller than 1 percent in absolute value to zero and exclude observations in the top and bottom 1.5 percent of the price change distribution. Additionally, we address (unobserved) heterogeneity among firms competing in the same industry but producing differentiated goods, which can significantly bias the measured standard deviation and Kurtosis of price changes upward (Klenow and Kryvtsov 2008; Alvarez et al. 2016). To correct for this, we standardize price changes at a disaggregated level by demeaning the price change observations within each cell. Here, a cell is defined by a firm-industry pair, the finest level of aggregation in our data. Furthermore, 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, which arise from firm-specific intercepts in the definition of target prices, trend inflation, and industry-specific seasonal patterns in nominal costs.[16]

3.3 Distribution of price changes and gaps: Summary statistics

Table 1 presents summary statistics of the distribution of firm-level log price changes, $p_{ft} - p_{ft-1}$, and price gaps, x_{ft-1}. The first four columns present moments describing the distribution of price changes. Panel a focuses on the 1999–2019 period, characterized by low inflation and, with the exception of 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 (approximately 0.5 percent quarter-on-quarter, on average).

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. Panel b presents the same statistics for the

[16] 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 (15)), and the normalization of price levels discussed in Section 3.1. This demeaning also removes the average trend inflation rate, which was approximately 0.6 percent quarter-on-quarter before the 2021 surge.

Table 1: Summary statistics of price changes and price gaps

Price change ($p_{ft} - p_{ft-1}$)				Price gap (x_{ft-1})		
Panel a: Time period 2000-2019						
Mean	Std.	Freq. Adj.	Kurt.	Mean	Std.	Kurt.
0.00	0.11	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			

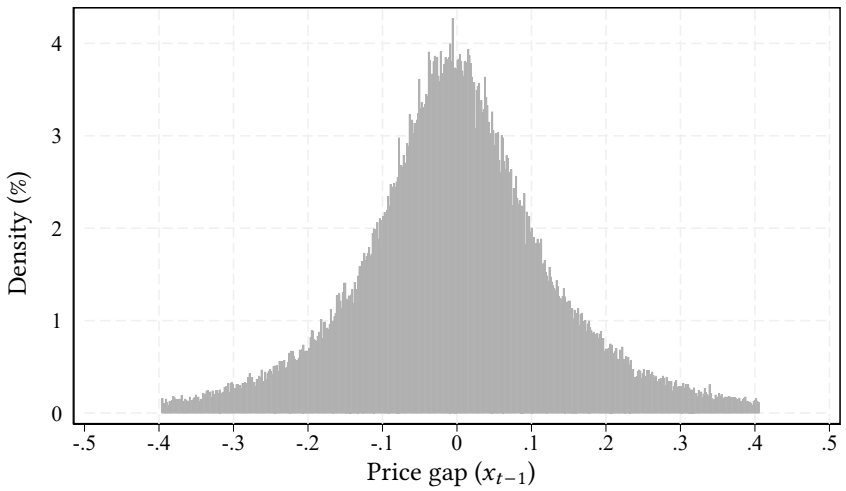
Notes. This table reports the summary statistics of the distributions of price changes ($p_{ft-1} - p_{ft}$), and price gaps (x_{ft-1}) before (panel a) and after the inflation surge (panel b).

period 2020–2023, characterized by high inflationary pressure and subsequent tapering. During this period, we observe a quarterly inflation rate that is on average 1 percentage point higher than the trend. At the same time, we observe a substantial increase in the frequency of price changes by 10 percentage points, on average.

The fourth column reports the kurtosis of price changes. We calculate this statistic following the approach in Klenow and Kryvtsov (2008), which scales the demeaned price changes by firm(-industry)-level standard deviations. The estimated kurtosis is 3.26, consistent with estimates in the literature (ranging between 3 and 5). This indicates that the distribution of price changes is more peaked (a higher frequency of small price changes) and has fatter tails (a greater number of large price changes) compared to a normal distribution (kurtosis equals 3). These features become even more pronounced during the inflation surge, as the kurtosis rises to 3.56.

The last three columns of Table 1 present summary statistics of the price gap distribution. This distribution, which is typically unobserved, is of great interest, as it contains information on inefficiencies due to the rigidities of nominal prices. Figure 5 presents the probability density function of the price gaps, $f(x_{t-1})$, in the pre-pandemic period. Consistent with our theory, the microdata reveal a price gap distribution that is

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).

unimodal, bell-shaped, and symmetric about the mean. During the inflation surge, on average, the price gap increased by 1 percentage point relative to its long-term trend. In line with theoretical predictions, this increase maps to the average average price change observed over the same period.

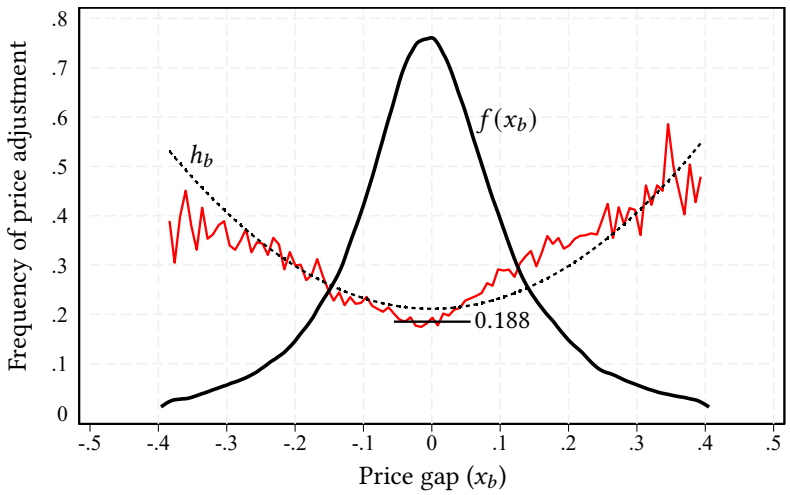
4 Micro evidence of state dependent pricing

Guided by the theoretical representations and identification results presented in Section 1, we design direct empirical tests of key model predictions that relate the micro-level pricing dynamics to the underlying price gap distribution. These exercises provide strong evidence of the state-dependent nature of firm pricing decisions, which becomes quantitatively significant when large aggregate cost shocks occur.

4.1 The empirical generalized hazard function

We begin by analyzing the relationship between price gaps and the frequency of price adjustment—the defining distinction between state- and time-dependent pricing models. In time-dependent models, the two variables are independent, yielding a flat probability of price adjustment (captured by the GHF). In contrast, state-dependent models predict

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 (7). In the regression, we weight each bin by the number of observations it counts.

that the probability of adjustment increases monotonically with the absolute value of the price gap. Specifically, in our framework, the GHF can be approximated by 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 gap distribution to the average squared gaps.

We use microdata on the frequency of price adjustment and price gaps from the pre-pandemic period (2009–2019) to test these predictions. In Figure 6, the black line depicts the probability density function of price gaps described in the previous section. The red line represents the empirical counterpart of the theoretical GHF, showing the fraction of firms adjusting prices in each quantile b (bin) of the price gap distribution. The black dotted line in Figure 6 displays the fitted values from a cross-sectional quadratic regression:

$$h_b = a_0 + a_1 \cdot x_b^2 + u_b.$$

The data reveal a strong relationship between price gap size and the frequency of price adjustment. With striking similarity to the theoretical GHF of state-dependent pricing models, wider gaps (i.e., greater deviations from $x = 0$) correspond to a higher likelihood of price adjustment. Moreover, consistent with model predictions, a quadratic

polynomial in price gaps fits the frequency data well. Figure 6 also shows that the empirical GHF exhibits a steeper slope on the right, indicating asymmetry in upward and downward rigidity. This suggests that firms are more likely to adjust prices when they are too low (i.e., when x is positive and the realized markup is too low) than when they are too high.

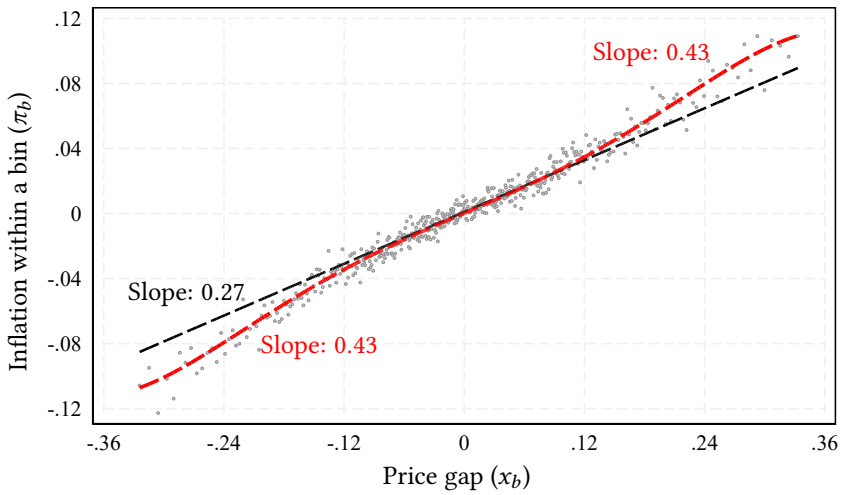
Given the strong fit of the quadratic approximation, we can apply Proposition 1 to obtain an estimate of the probability of free price adjustment by fitting the quadratic polynomial model within the interval centered around the bottom of the parabola, $x_b \in (-0.04, 0.04)$. This procedure yields $\hat{a}_0 = 0.188$, implying $\widehat{\theta^o} = 0.812$. We rely on this estimate for model calibration, as detailed in Section 6.1.

4.2 The empirical mapping from price gaps to price changes

The next set of results use cross-sectional variation in the joint distribution of price changes and price gaps to document the nonlinear cost-price dynamics resulting from firms' state-dependent policies. As before, we sort observations into quantiles (bins) spanning the entire price gap distribution. Figure 7 plots the average price change within the same bin (π_b, y-axis) as a function of the average price gap for each bin (x_b, x-axis). Comparing this empirical pattern with its theoretical counterpart in Figure 3 reveals a strong alignment between the microdata and model predictions linking inflation to the odd moments of the price gap distribution, as described by Equation (11).

Consider first the central quantiles of the distribution (covering the 25th to 75th percentiles). Observations in this range are characterized by relatively small price gaps, meaning moderate deviations of their prices from the target prices. We can think of this set as representing the mapping between gaps and price changes in "normal times," with low inflation and small aggregate shocks. As we can see, over this range, the mapping between inflation and price gaps is essentially linear, as in the Calvo model (Equation (12)). This result echoes findings from Gertler and Leahy (2008), Alvarez et al. (2017), and Auclert et al. (2024), which show that in "normal times," the price dynamics in state-dependent models resemble those of time-dependent models. Interestingly, we find that the gradient between price changes and price gaps (0.27) is approximately equal to the average frequency of price adjustment in the pre-pandemic period. This observation supports the identification strategy in Gagliardone et al. (Forthcoming), which relies

Figure 7: 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.

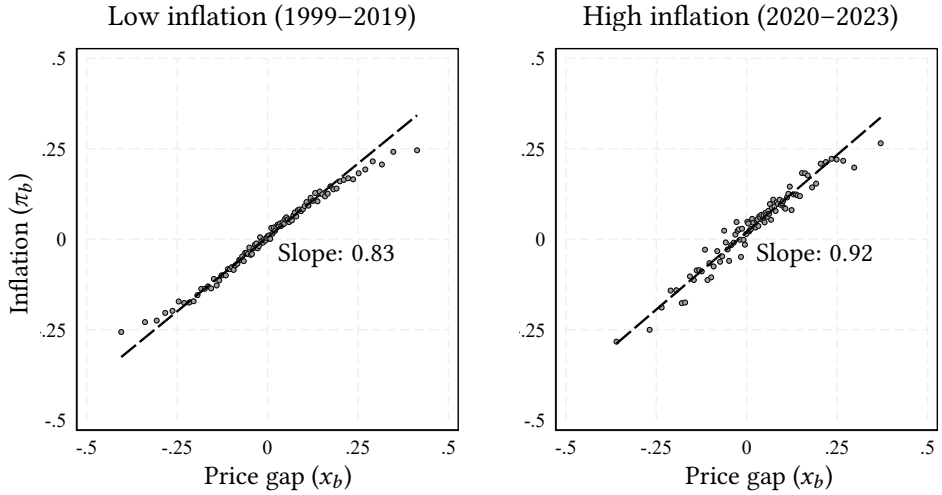
on the linear mapping between price changes and price gaps to estimate parameters underlying the slope of the Phillips curve in a low inflation environment.

To appreciate the nonlinearities induced by the state-dependent nature of price adjustments, consider the relationship between price changes and price gaps across the entire price gap distribution, including its tails. The red dashed line in Figure 7 represents the fit of the cross-sectional regression model estimated over the full support of the price gap distribution:

$$\pi_b = b_1 \cdot x_b + b_2 \cdot x_b^3 + \eta_b.$$

When price gaps widen—such as after a large aggregate shock—the sensitivity of price changes to price gaps increases significantly. The gradient between price gaps and price changes at the tails of the distribution steepens by nearly 60 percent compared to the gradient at the center (rising from 0.27 to 0.43). As shown in Proposition (2), the cubic term captures these nonlinearities at the tails, reflecting adjustments along the extensive margin (increases in the adjustment frequency).

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.

4.3 Price gaps and price changes conditional on adjustment

So far, we have examined the mapping between price gaps and price changed averaging across both firms that adjust and those that do not. We study the mapping focusing specifically on adjusters. According to theory, conditional on adjustment, firms set $p_{ft} = p^{\star}_{ft}$, implying that $p_{ft} - p_{ft-1} = x^{\star}_{ft}$. Although we cannot observe $x^{\star}_{ft}$ in the data, if p^{o}_{ft} provides a reasonable approximation for $p^{\star}_{ft}$, the elasticity of price changes with respect to price gaps, $(p_{ft} - p_{ft-1})/x_{ft-1}$, should be approximately one. Figure 8 confirms this prediction.

The figure presents two binned scatterplots, where the x-axis shows the average price gap for each percentile of the price gap distribution among adjusters (firms with $p_{ft} \neq p_{ft-1}$), and the y-axis reports the corresponding average percentage price change. The left panel focuses on the pre-pandemic period (1999–2019), while the right panel covers the pandemic and post-pandemic period (2020–2023). In both panels, the black dashed line represents the linear fit of price changes on price gaps across quantiles of the price gap distribution.

Consistent with the theoretical predictions, the data reveal a gradient that is not only positive but also close to one. Measurement error and the approximation of $p^{\star}_{ft}$ with $p^{o}_{ft}$are likely the two main factors explaining why the gradient is not exactly one. Interestingly, the gradient is particularly steep in the post-pandemic period, suggesting that firms may have become more attentive and responsive to cost changes when inflation is high. This is consistent with findings in Gagliardone and Tielens (2024), which use a model with state-dependent information frictions to show greater firm responsiveness in high-inflation environments.

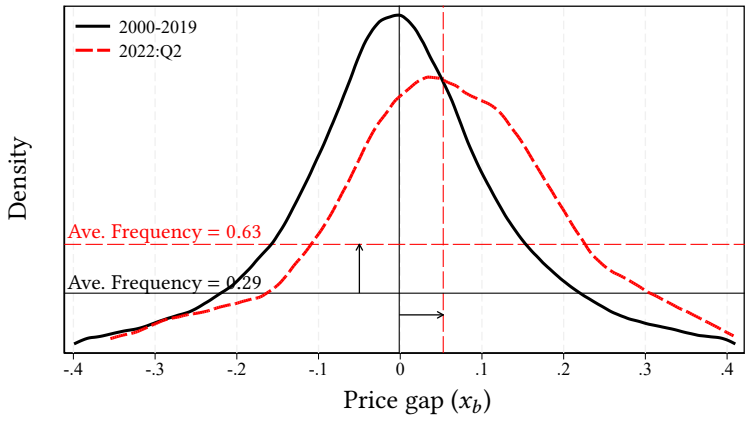
4.4 Large cost shocks and shifts in the price gaps distribution

In Section 1, we discussed how small idiosyncratic shocks generate dispersion in the price gap distribution, while large aggregate shocks shift the entire distribution of price gaps, significantly increasing the fraction of firms who want to adjust their prices (Figure 4). The drastic surge and subsequent normalization of production costs observed in the post-pandemic period allows us to directly test this model's prediction in the microdata.

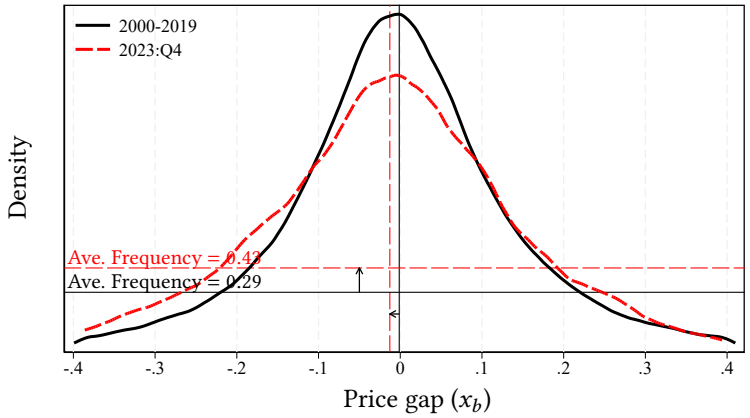
In Figure 9, the solid black line represents the distribution of price gaps before the pandemic. In panel (a), the red dashed line represents the distribution in 2022:Q2. In this quarter, firms' marginal costs increased by an average of 6.2 percent relative to the previous quarter. As predicted by the theory, a cost shock of this magnitude shifts the entire price gap distribution to the right, pushing many firms' prices away from their desired price levels. Since the shock compresses profit margins, the cost of inaction is large, and more firms move into regions where the GHF is high. As a result, within a single quarter, the average probability of price adjustment nearly doubles compared to its frequency in normal times. In panel (b), we repeat the exercise, but now the red line represents the distribution of price gaps in 2023:Q3. During this quarter, firms' marginal costs declined by an average of 3.8 percent relative to the previous quarter, as energy prices and international supply chains began to normalize. This negative cost shock shifted the price gap distribution to the left, increasing the frequency of price adjustments as more firms lowered their prices than raised them.

Figure 9: Impact of aggregate cost shocks on the price gap distribution and average frequency of price adjustment

Panel a: Positive aggregate cost shock



Panel b: Negative aggregate cost shock



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).

5 Aggregate cost-price dynamics

In this section, we shift our focus to macro-level cost-price dynamics. We show that micro-level dynamics—and their state-dependent nature—give rise to nonlinear inflation dynamics, where the pass-through of aggregate cost shocks to prices varies with the shock's magnitude.

5.1 Aggregate inflation and aggregate costs

We use our microdata to compute domestic producer price inflation and an index capturing changes in production 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 the quarterly changes in domestic firms' prices and weighting them by the Törnqvist weights $\bar{s}_{ft} \equiv \frac{s_{ft}+s_{ft-1}}{2}$:

$$\pi_t = \sum_{f \in \mathcal{F}} \bar{s}_{ft} \cdot \Delta p_{ft}.$$

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

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

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

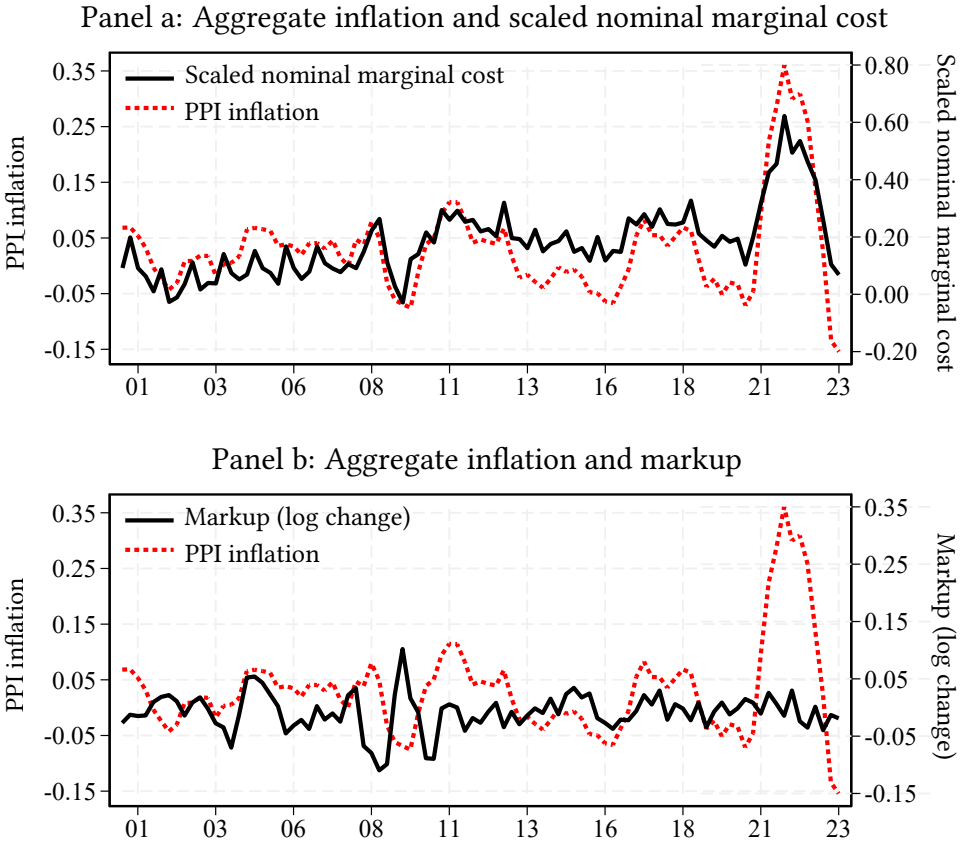
where the value of the index in the first quarter of our data is normalized to zero.

According to our theory, firms price on the basis of current and expected marginal costs. Therefore, the inflation rate between t and $t-4$ (the year-over-year rate, $p_t - p_{t-4}$) should depend on the nominal marginal cost at t, relative to the price level at $t-4$. We refer to the logarithmic difference between these variables, $mc_t^n - p_{t-4}$, as the "scaled nominal marginal cost." Figure 10 (panel a) shows the evolution of manufacturing inflation (red dashed line, left axis) and of scaled nominal marginal costs (black line, right axis) throughout our sample period. Note that the scales of the two axes differ for the variables.

Figure 10 highlights two key empirical patters. First, as predicted by the theory, inflation closely follows the fluctuations of scaled marginal cost throughout the entire sample period. However, consistent with the theory, the passthrough is imperfect, meaning that inflation responds less strongly than costs do. Second, the significant surge in inflation during the post-pandemic period, followed by its subsequent normalization, was driven by a dramatic rise and fall in scaled marginal costs.

To further stress the contribution of cost passthrough to movements in inflation, Figure 10, panel b, plots aggregate inflation against the log-change of average realized markups. We compute the latter by applying 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

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 scaled nominal 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.

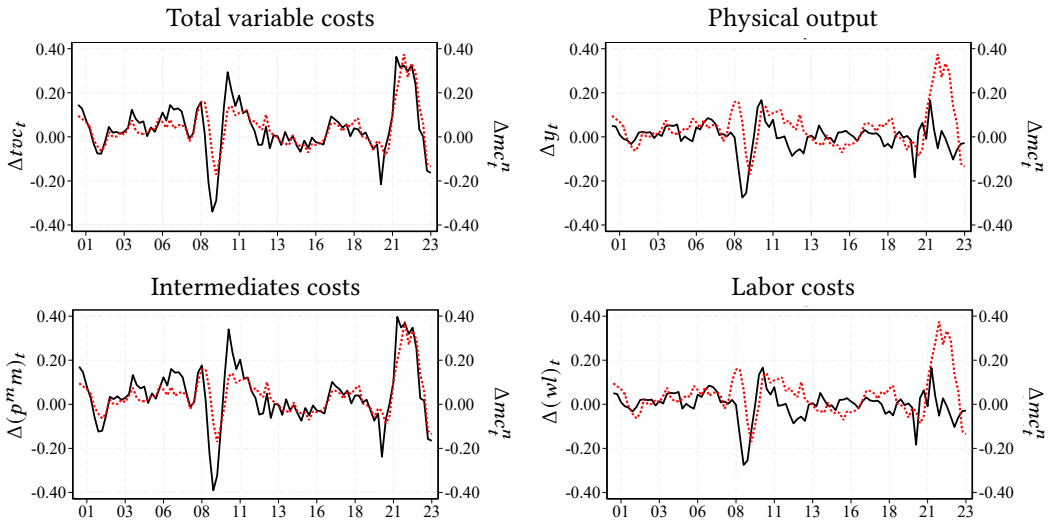
can explain the recent inflation surge seems to have no bite in the data.[17]

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 (15)). The top left panel shows the growth rate in total variable cost and real output (black lines) relative to the growth rate of the nominal marginal cost index (red dashed line). The two panels make clear that throughout the sample, and in particular during the recent inflation surge, fluctuations in total variable costs are the main drives of the time series evolution of nominal marginal cost.

The two panels at the bottom of Figure 11 further decompose total variable costs

[17] Analyzing price and cost data for a large global manufacturer, Alvarez et al. (2024) also finds that markups remained stable during the inflation surge.

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).

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 the intermediates was four times greater. This cost component alone accounts for approximately 70% of the revenues of manufacturing firms, on average. In addition, more than 80% of intermediate input costs come from importing from abroad. 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.

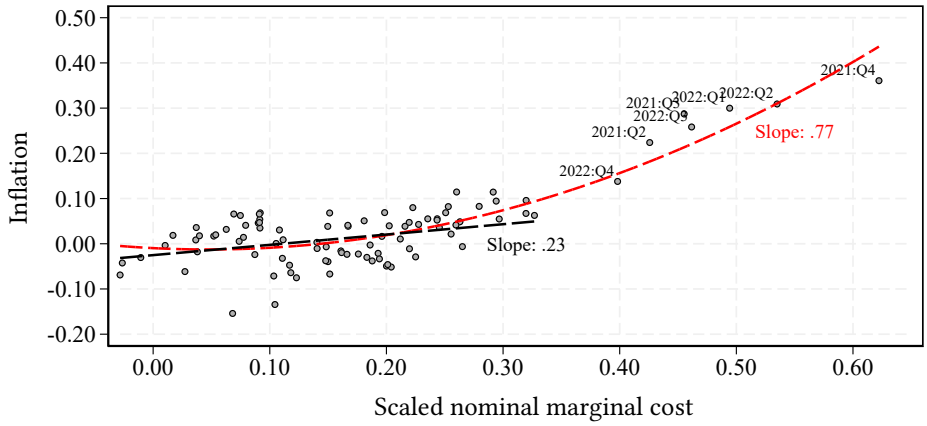
5.2 Nonlinear aggregate cost-price passthrough

In panel a of Figure 12 we sort quarters by their measured scaled 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 represents a quadratic fit across both high- and low-inflation periods. The slope 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.

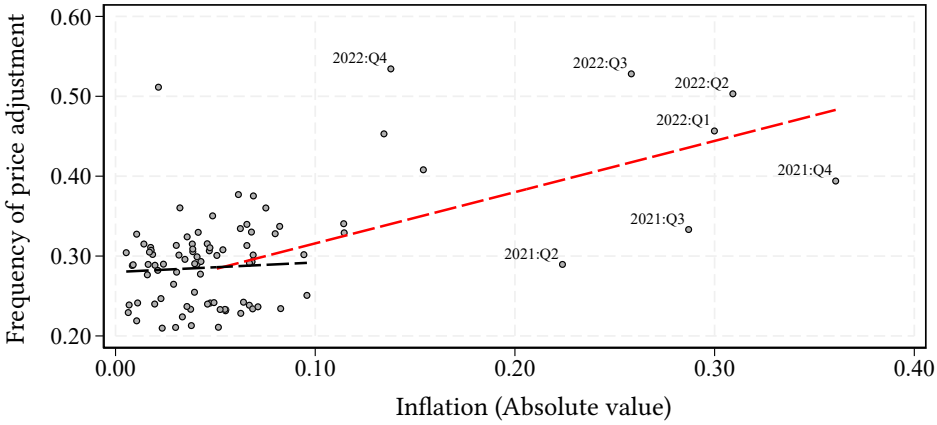
Consistent with the micro-level dynamics presented in Section 4, we find a linear

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 scaled nominal 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.

relationship between aggregate inflation and nominal costs during normal times. As discussed, a linear passthrough is consistent with the predictions of a Calvo model and with the predictions of a menu-cost model when aggregate shocks are small. The estimated reduced-form slope is 0.23, which is in close alignment with the aggregate

passthrough coefficient estimated by Gagliardone et al. (Forthcoming) in a low inflation environment.

The linear relationship between the two variables breaks down when the economy is hit by large aggregate shocks. The passthrough coefficient more than tripled during the recent inflation surge, revealing highly nonlinear cost-price dynamics. At the core of this result 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 Calvo, in low-inflation environments, we observe essentially no relationship between the average frequency of price adjustment and inflation. Once again, this finding highlights why time-dependent models provides a good framework to capture nominal rigidities in a low-inflation environment. 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 for Europe and the US.

6 Quantitative implications

Having established the close connection between theory and data, we now rely on moments from microdata to calibrate the quantitative model introduced in Section 1 and perform to types of quantitative exercises. First, we simulate the model and compare the dynamics of our state-dependent model to those of a standard time-dependent Calvo model in response to small and large shocks. Second, we feed the model a sequence of aggregate marginal costs shocks extracted from the data and evaluate how the model-generated aggregate inflation series compares to inflation in the data.

6.1 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 in Gagliardone et al. (Forthcoming). To align the model and the 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. The remaining three parameters, θ^o, σ_ϵ^2, and $\bar{\chi}$, control the degree of nominal rigidity and state dependence of price adjustment, mediating the relationship between price gaps and price adjustment.

The standard approach to calibrating these parameters leverages the theoretical mapping between the unobservable price gap distribution and measurable moments of the price change distribution: standard deviation, kurtosis, and the average frequency of price changes.[18] In theory, microdata on price changes could be used to recover these moments. In practice, however, previous research has shown that obtaining a reliable measure of the kurtosis of price adjustments is challenging due to measurement error and the inability to control for relevant sources of heterogeneity (Alvarez et al. 2016; Cavallo and Rigobon 2016). These issues may be particularly pronounced in our context, given our reliance on unit values. For this reason, we develop an alternative calibration procedure that avoids targeting the kurtosis of price adjustments and instead leverages the information embedded in the empirical GHF.

The calibration proceeds as follows. First, since the empirical GHF is well approximated by a quadratic polynomial in price gaps (see Figure 6), we use Proposition 1 to calibrate the probability of free adjustment, $(1 - \theta^o)$, to match the frequency of price adjustment in the vicinity of $x_{ft-1} \approx 0$. 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 4, this exercise yields $(1 - \hat{\theta}^o) = 0.188$.

Second, when trend inflation is low (as is the case during the pre-pandemic period) and assuming idiosyncratic shocks are i.i.d. draws from a 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:

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

[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)

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

Finally, 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 the empirical moments of the price change distribution (panel a) and the price gap distribution (panel b) to the corresponding moments of the menu-cost model, in steady state, under our baseline calibration. The model is able to capture the data quite well. Two observations lend additional empirical support to our calibration procedure. First, in a recent paper, Blanco et al. (2024a) show how a standard menu-cost model with single-product firms calibrated to match the kurtosis of price changes may need unreasonably high menu costs to rationalize the data. In our model, in steady state, menu costs amount to 1.7 percent of firm revenues, on average. This is consistent with empirical evidence of small menu costs documented in Levy et al. (1997) and Zbaracki et al. (2004). Second, we did not target the kurtosis of price changes in our calibration. However, despite the potential measurement issues discussed above, the calibrated model exhibits a kurtosis of price changes that is broadly in line with the one observed in our data. In Appendix C, we show that the results of the quantitative exercises presented below are robust to an alternative calibration that targets the kurtosis of price adjustment.

We also consider a standard Calvo model calibrated to match the steady-state frequency of price adjustment observed in the data. As explained in Section 1, our menu-cost model nests the Calvo model as a special case when the maximum menu cost, $\bar{\chi}$,

approaches infinity and the probability of free price adjustment is re-calibrated to match the steady-state frequency of price adjustment ($1 - \theta^c = h$). The theoretical moments derived from the steady state of the Calvo model are presented in the third row of Table 2. As expected, while the time-dependent model successfully matches the first and second moments of the price change and price gap distributions, it generates a pronounced degree of leptokurtosis, which is largely inconsistent with the kurtosis measured in the data.

6.2 Impulse-responses to small and large aggregate shocks

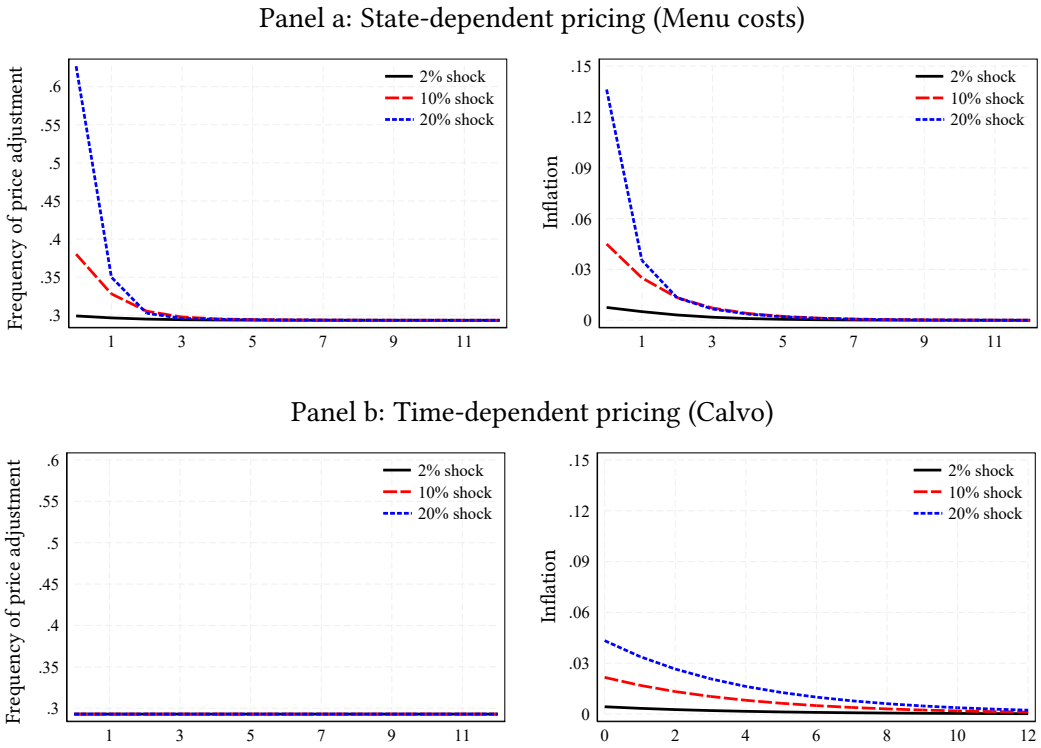
We use our calibrated menu-cost and Calvo models to study price dynamics in response to large and small shocks, under state- and time-dependent pricing. Starting from an economy in steady state, we shock the system with permanent and unanticipated aggregate cost shocks of different magnitudes, g_t=$\{2\%,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

Figure 13 displays the impulse response function of the frequency of price adjustment (left panels) and aggregate inflation (right panels) in the menu-cost model and in the Calvo model. All shocks increase the optimal reset price, shifting the distribution of price gaps to the right, thereby triggering inflation. However, as discussed in Section 1 and empirically shown in Section 4, in the menu-cost model large shocks displace many firms in a region where the GHF is higher, generating a spike in the frequency of price adjustment and, consequently, a more rapid and substantial surge in inflation compared to the Calvo model.

Comparing the IRFs for shocks of varying magnitude reveals the nonlinearities of state-dependent pricing, which become more pronounced as shocks intensify. On impact, the effect of the large shock on both the frequency of price adjustment and inflation is about three times larger than the effect of the medium shock, although the former is only twice as large as the latter (10% vs. 20%). By construction, in Calvo, the number of firms adjusting their prices is not affected by the magnitude of the shock (the GHF is flat across the price gap distribution), and adjusters are a random sample of the population (aka, there is no selection effect). As a result, inflation increases with the magnitude of the shock, but in a proportional way.

A second observation concerns the differing speeds at which permanent cost shocks of varying magnitudes are fully incorporated into prices in state- and time-dependent models. Figure 14 overlays the IRFs of inflation in both models in response to shocks of

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



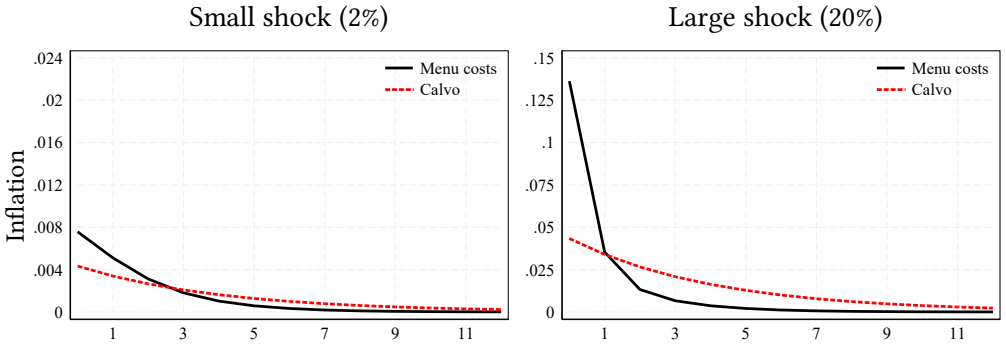
Notes. This figure presents the impulse responses of inflation and frequency to aggregate cost shocks of different magnitudes. Panel a reports the impulse response for our state-dependent pricing model (menu-cost model). Panel b reports the impulse responses for a time-dependent model (Calvo model).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

the same size (small or large). Although the passthrough is similar for small shocks, it is notably faster in the menu-cost model for large shocks. This difference arises from the endogenous change in the frequency of price adjustment and the selection effect present in the menu-cost model.

In Appendix B we present 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 dynamics of the two prices are closer in the context of the menu-cost model than in the Calvo model, consistent with our assumptions for the measurement of $p^\star_{ft}$.

The second exercise studies the role of strategic complementarities in both state- and time-dependent models. We compare inflation dynamics after high- and low-cost

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.

shocks, without strategic complementarities ($\Omega = 0$) and with strategic complementarities ($\Omega = 0.5$). As expected, strategic complementarities lead to a reduction in the cost passthrough in both the menu cost and the Calvo model. The greater curvature of the value function under state-dependent pricing implies that the difference between the impulse-response functions with and without complementarities is narrower in the menu-cost model, especially in response to a large shock.

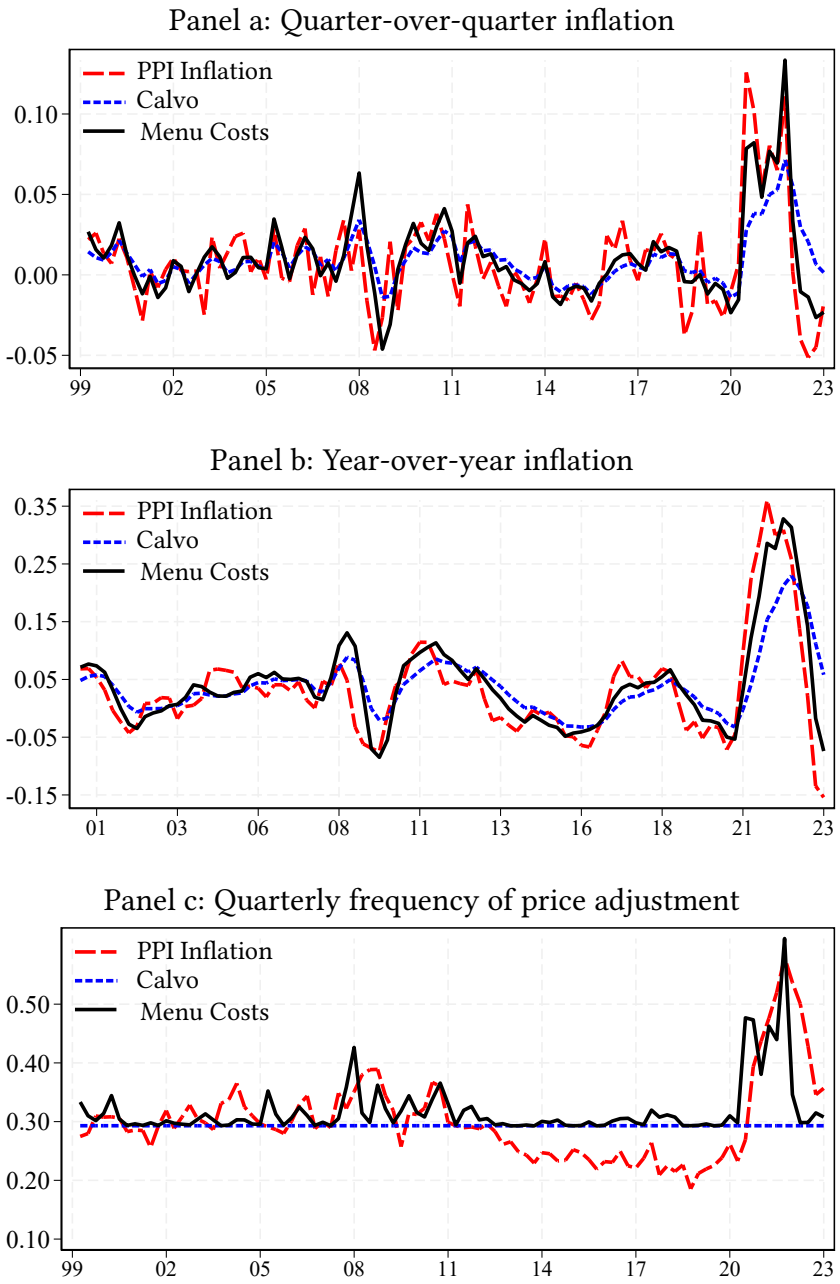
6.3 Explaining the time series of inflation

We now turn to evaluating the ability of time- and state-dependent models to explain the time series of aggregate inflation observed in the data. We feed into our 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 for the state-dependent menu-cost model and for its time-dependent Calvo counterpart, calibrated to hit the same steady-state frequency of price adjustment.

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_t^n, 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

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.

gaps and the response of inflation to the frequency of price adjustment, assuming that all future aggregate shocks are unanticipated, as in an impulse response function. 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 characterizing the pre-pandemic period and during the post-pandemic inflation surge and bust.

Note also that, 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. The Calvo model also exhibits an inflation surge during the pandemic era, but only about two-thirds of that is generated by the menu-cost model. This exercise also highlights the more sluggish behavior of inflation produced by the Calvo model relative to that generated by the menu-cost model. This is consistent with the faster cost passthrough generated by the state-dependent pricing policies documented in the impulse response function of Section 6.2.

Finally, panel c plots the quarterly frequency of price adjustment. The model captures the stable behavior of the adjustment frequency pre-pandemic, though it misses the smooth trend decline between 2012 and 2019. However, the model captures well the sharp jump in the adjustment frequency following the onset of the pandemic, both in terms of timing and magnitude. As inflation drops, the model frequency recedes faster than in the data. It is possible that firms anticipated the mean reversion in nominal marginal costs better than our random walk model would suggest.

7 Concluding remarks

In this paper, we study cost-price dynamics in normal times and during the recent inflation surge. We leverage detailed information on prices and costs to construct a direct measure of firms' price gaps and analyze the microdata through the lens of a tractable menu-cost model. Variation in price gaps determines both the likelihood that a firm adjusts its price and how much its price changes conditional on adjustment, providing strong empirical support to the predictions of our state-depending pricing model.

At the macro-level, we document linear cost-price dynamics in "normal" times,

when aggregate inflation is low. That is, aggregate inflation is well approximated by the product of a fixed price adjustment probability and the average price gap. In contrast, during the inflation surge, the cost-price dynamics is highly nonlinear. The sharp increase in the marginal cost led not only to a jump in price gaps but also to a significant increase in adjustment probabilities. This extensive margin of price adjustment is the hallmark of state-dependent pricing models but is absent in time-dependent models, such as the workhorse Calvo (1983) model.

Overall, we find that conditional on the path of marginal cost, the state-dependent pricing model does a good job of capturing price dynamics both at 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 model (for example, Galí 2015) typically includes labor as the only variable input, implying that the marginal cost is measured by the labor share. However, our analysis suggests that the sharp increase in the cost of intermediate inputs was the main driver of variation in marginal cost observed during the inflation surge. Extending a state-dependent version of the New Keynesian model to allow for wage determination, intermediate inputs, primary commodities and energy, and supply chains is on the agenda 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.

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

Andres Blanco, Corina Boar, Callum Jones, and Virgiliu Midrigan. Nonlinear inflation dynamics in menu cost economies. *Working paper*, 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.

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 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.

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.

Peter Karadi and Adam Reiff. Menu costs, aggregate fluctuations, and large shocks. *American Economic Journal: Macroeconomics*, 11(3):111–146, 2019.

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

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:

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

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 $\epsilon = -\frac{\Upsilon'(1)}{\Upsilon''(1)}$. Then the derivative of the residual demand elasticity $\epsilon_t(f)$ in (A.3) with respect to $\frac{Y_t(f)}{Y_t}$, evaluated at the steady state, is given by:

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

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:

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

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 $\Gamma \equiv \frac{\epsilon'}{\epsilon(\epsilon-1)}\frac{1}{\Upsilon''(1)}$. Finally, replacing the log-linearized markup into the formula for the static optimal target price (obtained from cost minimization):

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

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 $mc_t(f)$ and the fact that taste shocks are i.i.d., we obtain the following expression that characterizes the optimal reset price:

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

The equation above decomposes the optimal (dynamic) reset price into two terms. The first is the static reset price, $p^0_{ft} \equiv (1-\Omega)(\mu(f) + mc_t(f)) + \Omega p_t$, 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 (7) 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 (6) 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 (7):

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

where the parameter ϕ controls the sensitivity of the GHS 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 (7) 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 (8), we want to derive the cubic expression for inflation within a bin in Equation (11), 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:

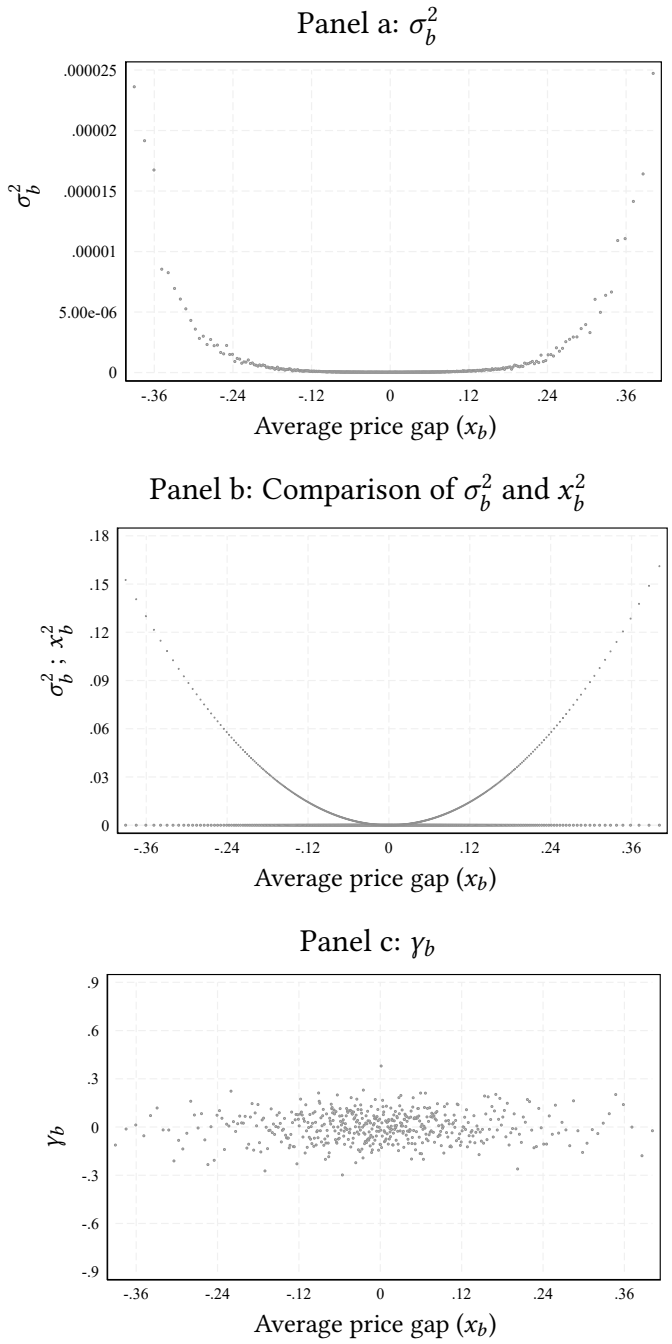
$$\begin{aligned}Cov_b(h_f, x_f) &= \phi\, \mathbb{E}_b(x_f^3) + o(x_b^3) - \phi\, \mathbb{E}_b(x_f^2) x_b - o(x_b^2) x_b \\ &= \phi(2x_b\sigma_b^2 + \gamma_b \sigma_b^3) + o(x_b^3)\end{aligned}$$

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

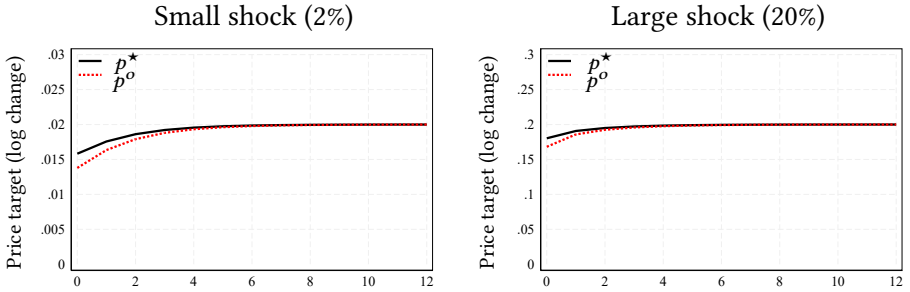
Figure A.1: Within bin variance, square of the mean, and skewness



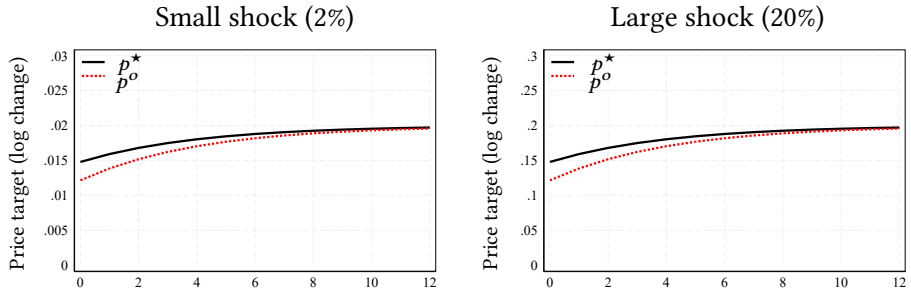
Notes. Panel a figure plots the within-bin variance of price gaps (σ_b^2), the square of the within-bin square of the within-bin average gap (x_b^2), and the within-bin skewness of price gaps (γ_b) for different bins (quantiles) along the price gap distribution. Panel b plots σ_b^2 (circles) and x_b^2 (crosses) on the same scale.

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

Panel a: State-dependent pricing (Menu costs)



Panel b: Time-dependent pricing (Calvo)



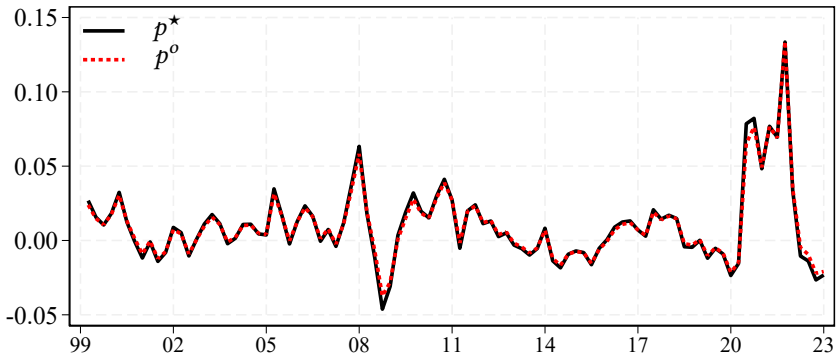
Notes. This figure presents impulse responses of the static target price (p^o) and the optimal reset price ($p^\star$) to aggregate cost shock of different sizes. The x-axis reports quarters since the shock.

B Additional quantitative exercises

Approximation of $p^\star_{ft}$ with p^o_{ft}. As discussed in Section 1, 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 (8) and (11), 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



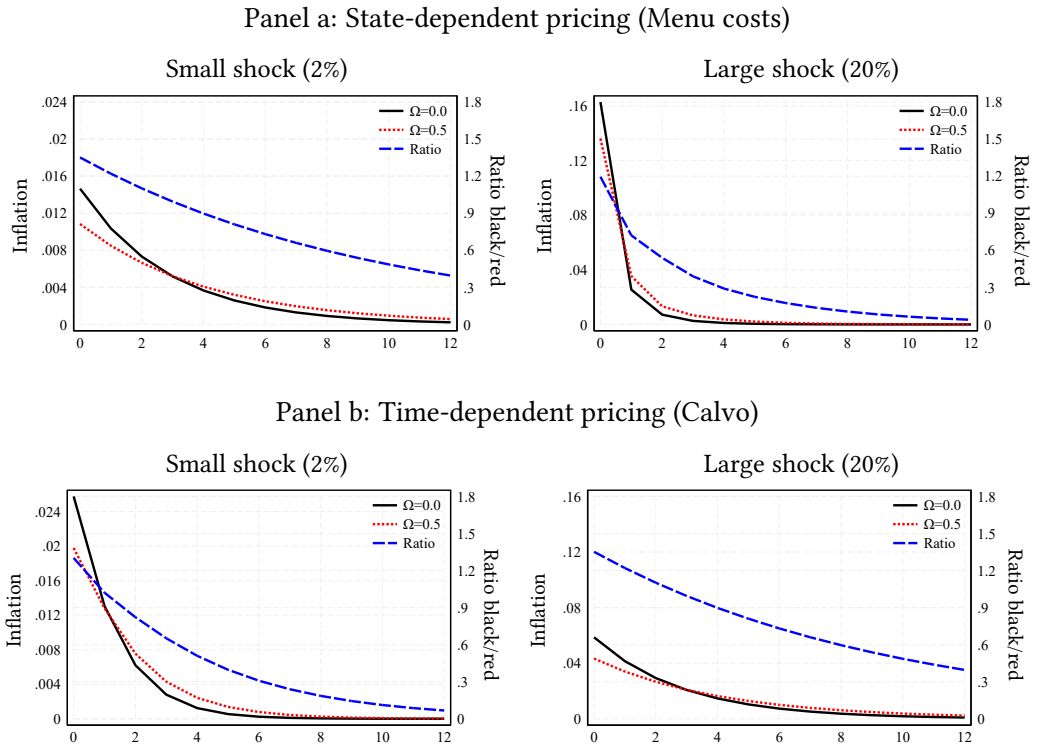
Notes. This figure contrasts the inflation dynamics generated by the menu-cost model using $p^\star$ (the exact, dynamic reset price) and using p^o (the static approximation of $p^\star$) when solving the model. As in Figure 15, we solve the model feeding it a sequence of aggregate nominal marginal cost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

C Alternative calibration

In our baseline exercise, we used information from the empirical GHF to directly calibrate the frequency of price adjustment. As explained in the paper, this procedure tends to be more robust to small measurement error which can affect the measured kurtosis of price adjustment. As a robustness check, we perform an alternative calibration, this time targeting the kurtosis of price adjustment. We then repeat the feeding exercise discussed

Table A.1: Calibration: Data vs. model under alternative calibration

	Price change ($p_{ft} - p_{ft-1}$)			Price gap (x_{ft-1})	
	Std	Freq. Adj.	Kurt.	Std	Kurt.
Data	0.12	0.29	3.26	0.13	2.86
Menu cost	0.12	3.27	2.62	0.01	3.61

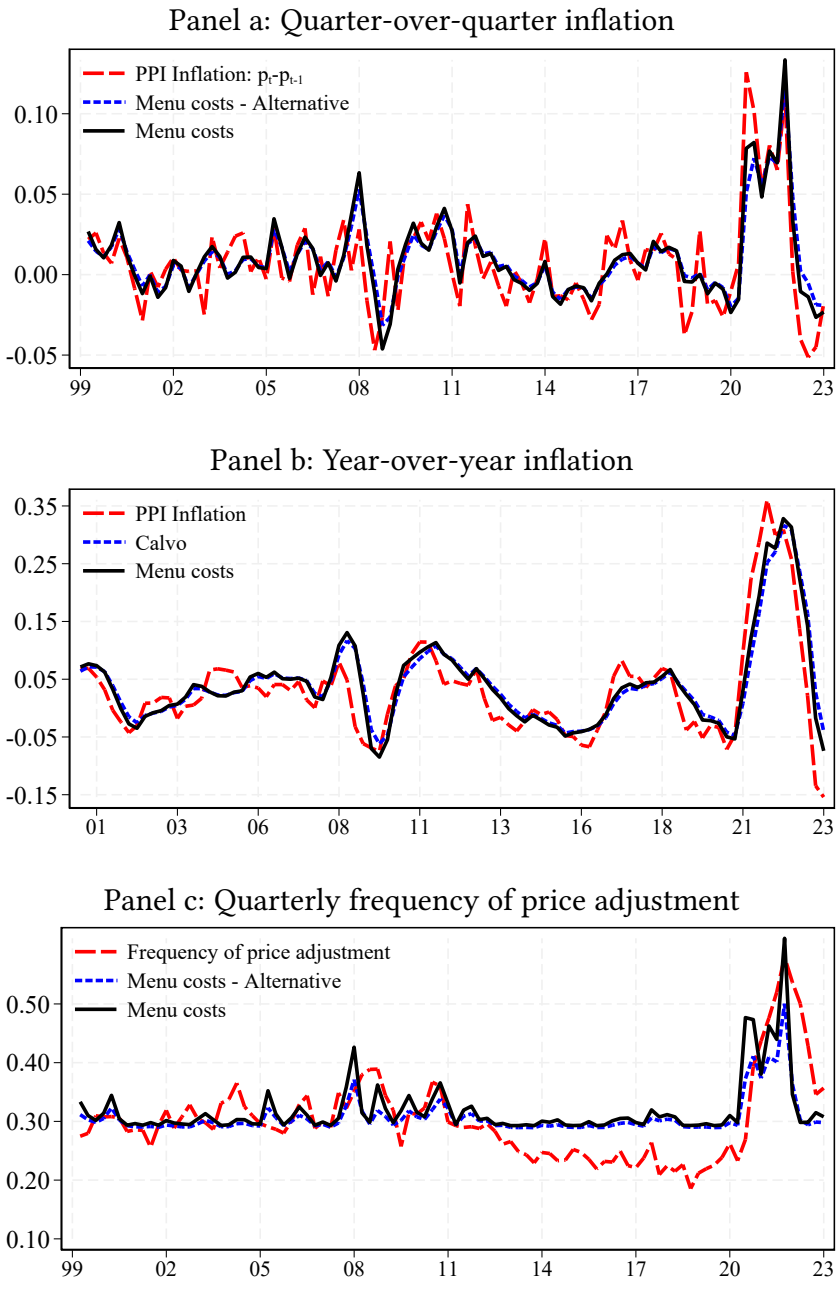
Notes. This table reports moments of the distribution of price changes and price gaps computed during the period 2000–2019 and the corresponding moments for the menu-cost model, in steady-state, under an alternative calibration that targets the average frequency of price adjustment as well as the standard deviation and kurtosis of price adjustment.

in Section 6.3 to derive an alternative model-based time series of aggregate inflation and the average frequency of price adjustment.

We choose the parameters $\sigma, \lambda, \bar{\chi}$ to jointly minimize the distance between (i) the average frequency of price adjustment, (ii) the standard deviation of price adjustment, and (iii) the kurtosis of price adjustment observed in the data and their model-based counterparts. Table A.1 reports the targeted moments in both the data and the calibrated model. The calibration yields the following parameter values $\{\sigma, \lambda, \bar{\chi}\} = \{0.063, 0.765, 1.441\}$. To match the kurtosis of price adjustment in the data, this calibration requires a higher value for the upper bound of the menu cost parameter (1.441 compared to 0.61 in our baseline calibration). Given the higher incidence of menu costs, the probability of free price adjustment must increase relative to our baseline ($1 - \theta^0 = 0.235$ compared to 0.188) to match the average frequency of price adjustment.

Figure A.5 compares aggregate inflation in the data with the model-based sequence under our baseline (solid black line) and alternative calibration (dashed blue line). Both sequences capture fluctuations in aggregate inflation and the frequency of price adjustment well. However, the model fit is better under our baseline calibration, as lower menu costs allow the model to better capture high-frequency movements in the frequency, particularly in response to large aggregate cost shocks.

Figure A.5: 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 our menu-cost models under our baseline (solid black line) and alternative calibration (dashed blue line).