NBER WORKING PAPER SERIES

MICRO AND MACRO COST-PRICE DYNAMICS IN NORMAL TIMES AND DURING INFLATION SURGES

Luca Gagliardone
Mark Gertler
Simone Lenzu
Joris Tielens

Working Paper 33478
http://www.nber.org/papers/w33478

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2025

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 Bureau of Economic Research, the National Bank of Belgium, the Eurosystem, or any other institution with which the authors are affiliated.

Micro and Macro Cost-Price Dynamics in Normal Times and During Inflation Surges
Luca Gagliardone, Mark Gertler, Simone Lenzu, and Joris Tielens
NBER Working Paper No. 33478
February 2025
JEL No. E0

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.

Luca Gagliardone
Department of Economics
New York University
New York City, NY 11211
luca.gagliardone@nyu.edu

Simone Lenzu
New York University
44 West Fourth Street Suite 9-79
New York, NY 10012
slenzu@stern.nyu.edu

Mark Gertler
Department of Economics
New York University
269 Mercer Street, 7th Floor
New York, NY 10003
and NBER
mark.gertler@nyu.edu

Joris Tielens
National Bank of Belgium
18 Rue de Berlaimont
Brussels 10000 Belgium
joris.tielens@nbb.be

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, making expected price changes a linear function of the 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 adjustments 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 expected price changes and price gaps is observationally equivalent to that predicted by a

[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 adjustments

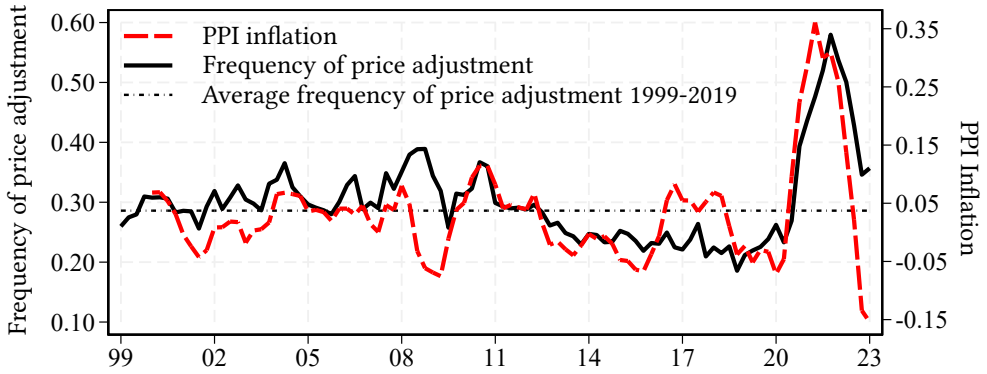

Notes. This figure shows the time series of PPI manufacturing inflation along with the annual frequency of price adjustments. The former is computed as the year-over-year percentage change in the aggregate PPI. The latter is calculated as a rolling average of the quarterly frequency of price adjustments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

time-dependent model, resulting in linear cost-price dynamics.[2] This evidence contributes to providing a unifying, microfounded accounting of inflation during both normal times and 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 time-dependent Calvo (1983), as a special case. We use this framework to derive testable predictions that relate, in the microdata, price changes to price gaps. We also use the model to explain the time series of aggregate inflation as a function 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

[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.

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 adjustments, 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 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 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 and small aggregate shocks, the price gap distribution is stable, the frequency of price adjustments 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. First, descriptive evidence shows how inflation and production costs align closely throughout the entire sample period. However, there is stickiness in price adjustments 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. Second, we formally assess the capacity of our menu-cost model to explain aggregate inflation. We feed into the model the marginal cost index described above and compare the model with data. We find that the model tracks the high-frequency fluctuations in 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 during the 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 evidence on the distribution of prices (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 adjustments relate to price gaps in the microdata, 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 drug retailer to develop a "reference price" metric. The latter employs microdata on supermarket prices to formulate a concept of reset prices, 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 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. (2024) 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 2 describes our dataset and the empirical measures

of prices, cost, and price gaps. Section 3 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 5, showing how the calibrated model explains the inflation time series and the frequency of price adjustments, including the rise during the pandemic. Section 6 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 adjustments 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. 2024), 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.

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, $z_t(f) =$

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

$z_{t-1}(f) + \zeta_t(f)$, 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 5.

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 up in the target price as noise.

We define the price gap:

$$x_t(f) \equiv p_t^o(f) - p_t(f).$$

[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.

Following Alvarez et al. (2023), we take a quadratic approximation of the per-period profit function around the static optimum $x_t(f) = 0$ and normalize it by steady-state profits. This yields the following loss function measuring the cost of deviations of the price from the target:

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

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:

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

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". As we shall discuss, it is convenient to define the firm's *"ex-ante" price gap* in period t, $x'_{t-1}(f)$, which captures the difference between the target price and the price set by the firm in the previous period:

$$\begin{aligned} x'_{t-1}(f) &\equiv p_t^o(f) - p_{t-1}(f) \\ &= x_{t-1}(f) + (1-\Omega)(g_t + \varepsilon_t(f)) + \Omega(p_t - p_{t-1}). \end{aligned} \tag{4}$$

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 ex-ante price gap $x'_{t-1}(f)$ is measured before

[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 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 single-peaked, 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 ex-ante price gap. Then the solution to the firm's problem can be expressed as a function of the ex-ante price gap:

$$p_t^o(f) - p_t(f) = x_t(f) = \begin{cases} x_t^\star & \text{w. p. } h_t(x'_{t-1}) \\ x'_{t-1}(f) & \text{w. p. } 1 - h_t(x'_{t-1}). \end{cases} \tag{5}$$

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 ex-ante 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.

Optimal reset gap. We now characterize $V_t(x'_{t-1})$, V_t^a, and therefore $x_t^\star$. As discussed, $x_t(f) = x'_{t-1}(f)$ for a firm that does not adjust its price. In this case, the value of the firm is given by:

$$V_t(x'_{t-1}) = \Pi_t(x'_{t-1}) + \beta\, \mathbb{E}_t \left\{ h_{t+1}(x'_t) \cdot V_{t+1}^a + \left(1 - h_{t+1}(x'_t)\right) \cdot V_{t+1}(x'_t) \right\}.$$

It is a function of current profits Π_t, evaluated at the ex-ante price gap $x'_{t-1}(f)$, and of the discounted expected continuation value. The latter depends on the probability of adjustment at time $t+1$, $h_t(x'_t)$. 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 ex-ante price gap $x'_{t-1}(f)$ as a simple function 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.

[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 full treatment of menu-cost models with strategic complementarities.

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:

$$\begin{aligned}\pi_t &= \int \Big(p_t(f) - p_{t-1}(f)\Big)\, df = \int h_t(x'_{t-1}(f)) \cdot \big(p_t^\star(f) - p_{t-1}(f)\big)\, df \\ &= \int h_t(x'_{t-1}(f))\, df \cdot \int \Big(p_t^\star(f) - p_{t-1}(f)\Big)\, df + \mathbb{C}\mathrm{ov}\Big(h_t(x'_{t-1})\,,\ (p_t^\star - p_{t-1})\Big) \qquad (7)\end{aligned}$$

The first line shows that, to a first-order approximation, aggregate inflation is an average of firm-level price adjustments, 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 adjustments 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 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 (7). 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 $h(x'_{t-1})$. Thus, the selection effect positively contributes to generating aggregate inflation.

1.2 Discussion and testable implications

Price gaps and the generalized hazard function. To develop some intuition on firm pricing decisions, we use the diagram originally presented in Caballero and Engel (2007) describing the adjustment process in the stationary equilibrium.

We simulate our model by calibrating the parameters of interest to the values described in Section 5.1. In Figure 2, the horizontal axis depicts the density function of the ex-ante price gaps in the steady-state of the model, denoted by $f(x'_{t-1})$, which is unimodal

Figure 2: Generalized Hazard Function and distribution of price gaps

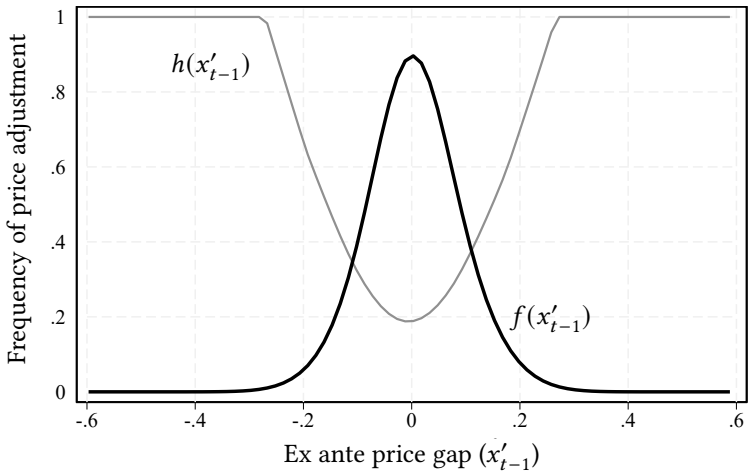

Notes. This figure plots the probability density function of price gaps, $f(x'_{t-1})$, against the Generalized Hazard Function (GHF), $h(x'_{t-1})$, evaluated at the steady state of the model.

and bell-shaped. The vertical axis reports the GHF at different points of the price gap distribution, $h_t(x'_{t-1})$. Firms in the right (left) tail of the price gap distribution are firms that, given the realization of the shocks that affect $p^o_t(f)$, operate with a suboptimally low (high) markup and therefore are more likely to increase (decrease) their price relative to the price they changed in the previous period.

With low trend inflation, assuming stationarity of the value function, the GHF can be accurately approximated by a quadratic function of the price gap centered around $x'_{t-1} = 0$.[8] Specifically, we have that:

$$h_t(x'_{t-1}(f)) \approx (1-\theta^o) + \phi \cdot \left(x'_{t-1}(f)\right)^2. \tag{8}$$

The GHF is U-shaped and symmetric around the point where the price gap is zero, which corresponds to the optimum in the stationary equilibrium. At this point, the adjustment frequency is at a minimum, corresponding to the probability of a free price adjustment $(1-\theta^o)$. As price gaps widen, the adjustment frequency monotonically increases with it. The parameter $\phi \equiv -\frac{\theta^0}{\bar{\chi}} \frac{\partial^2 V(x)}{\partial x^2}\Big|_{x=0}$ controls the sensitivity of the GHF to changes in gaps (i.e., the "steepness" of the parabola).

[8] In Appendix A.3 for the derivation of the quadratic functional form under a second-order approximation and Alvarez et al. (2022) for a general treatment.

Nonlinear price dynamics along the price gap distribution. The endogenous fluctuations in the probability of price adjustment are the key driver of nonlinear shock transmission in state-dependent models. To illustrate this, we divide the distribution of ex-ante price gaps into equally sized bins, each narrow enough to ensure that the price gap is nearly constant within them. Starting from the inflation equation in (7) and applying the quadratic approximation in Equation (8) to substitute for the hazard function, we derive the following expression, which characterizes inflation within bin b as a third-order (odd) polynomial of the price gap:[9]

$$\pi_b \approx \phi_b^0 \cdot \left(x_b'\right) + \phi \cdot \left(x_b'\right)^3, \tag{9}$$

where $\pi_b \equiv \int_{f\in b} \left(p_t(f) - p_{t-1}(f)\right)\, df$ and $x_b' \equiv \int_{f\in b} x_{t-1}'(f)\, df$ measure the average price change the average ex-ante price gap (across both adjusters and non-adjusters) that belong to a given bin b. The parameter $\phi_b^0 \equiv 1 - \theta^o + \phi\sigma_b^2$ is the sum of the probability of free adjustment $(1 - \theta^o)$, which is common between bins, and a term equal to the variance of the price gaps within the bin ($\sigma_b^2 \equiv \int_{f\in b}(x_{t-1}'(f))^2\, df - (x_b')^2$), scaled by the slope parameter of the GHF (ϕ). The latter captures the effect of deviations of the price gap from zero on the adjustment frequency. It is straightforward to derive the analog of Equation (9) in the case of a time-dependent Calvo model. Given the constant exogenous hazard rate $h^c = (1 - \theta^c)$, we have $\pi_b = (1 - \theta^c) \cdot \left(x_b'\right)$.[10]

The binned scatter plot in Figure 3 illustrates the relationship between price gaps and inflation across the price gap distribution for the menu-cost model (gray circles) and the Calvo model (gray diamonds). The dashed lines represent: (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 (9) for the menu-cost model; and (ii) the fitted values from a regression of bin-level inflation on a first-order polynomial average gap, as predicted by the Calvo model.

When price gaps are sufficiently close to zero, the third-order term is negligible and the average price adjustments are directly proportional to the average price gap. Thus, the pricing dynamics of firms that operate close to their optimum are linear in both state- and time-dependent models. This observation is at the core of the approximate equivalence

[9] See Appendix A.5 for the derivations.

[10] 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

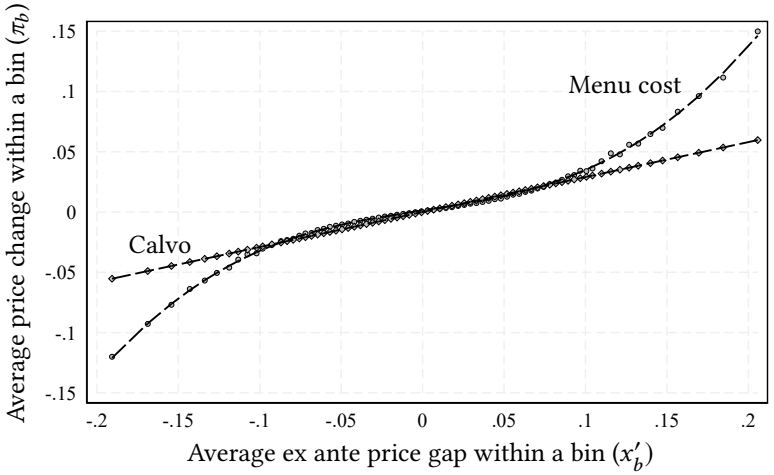

Notes. In this figure, we partition the distribution of price gaps into narrow equal-size bins (b). We plot the average ex-ante price gap of each bin, x'_b, against the average logarithmic price change for observations in the same bin, π_b. The gray dots refer to simulations of our menu-cost model. The gray diamonds are simulations from a Calvo model calibrated to the same average frequency of price adjustment. The dashed lines represent: (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 (9) for the menu-cost model; and (ii) the fitted values from a regression of bin-level inflation on a first-order polynomial average gap, as predicted by the Calvo model. The weight of each bin in the regression is proportional to the number of observations within it

result between the Calvo and menu-cost models with small shocks illustrated in Gertler and Leahy (2008), Alvarez et al. (2017), and Auclert et al. (2024).

The cubic term opens the door for nonlinear inflation dynamics. Firms in the tails of the price gap distribution are more likely to adjust their prices (the frequency effect captured by the GHF).

The impact of aggregate cost shocks. Aggregate cost shocks, that is, shocks that do not average out, affect the 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 cubic term becomes non-negligible. 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

Figure 4: Large versus small aggregate cost shocks

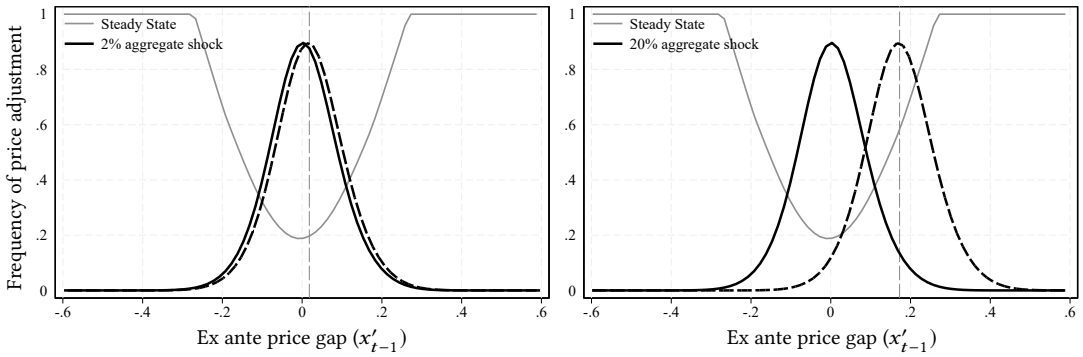

Notes. This figure reports the ex-ante price gap distribution in steady state (black solid line) and after a small and a large cost shock (black dashed line).

GHF and the distribution of ex-ante price gaps in the stationary equilibrium. The dashed lines show the post-shock distributions.

A small shock induces a small shift in the ex-ante price gap distribution, with little or no impact on the average frequency of price adjustment. This observation is in line with Gertler and Leahy (2008) and Auclert et al. (2024), which show how the price dynamics generated by a state-dependent model is observationally similar to those generated by a time-dependent model when the economy is hit by small shocks. In contrast, a large cost shock displaces a sufficient number of firms far from their target price, widening the average gap in the economy. This results in a sharp increase in the fraction of firms that want to raise their prices and in a lower fraction that want to reduce them. On average, this causes a nontrivial increase in the adjustment frequency, which amplifies aggregate inflation beyond what is accounted for by the increase in the average gap.

2 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. (2024)

in two significant ways.[11] First, the data in Gagliardone et al. (2024) 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 adjustments.

2.1 Measurement of prices, costs, and price gaps

The unit of observation in our data is a firm-industry pair, where industries are narrowly defined based on 4-digit NACE rev.2 product codes. Our final dataset tracks 5, 348 domestic firm-industry pairs, denoted by a lower-script f, distributed across 169 manufacturing industries, denoted by lower-script i, 18 sectors (3-digit codes) and 9 macro sectors (2-digit codes).

Price indices. For each domestic firm, we use PRODCOM data on product-level 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:

$$\frac{P_{ft}}{P_{ft-1}} = \prod_{p \in \mathcal{P}_{ft}} \left(\frac{P_{pt}}{P_{pt-1}} \right)^{\bar{s}_{pt}}, \tag{10}$$

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

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

$$\frac{P_{it}^{-f}}{P_{it-1}^{-f}} = \prod_{k \in \mathcal{F}_i \setminus f} \left(\frac{P_{kt}}{P_{kt-1}} \right)^{\bar{s}_{kt}^{-f}}, \tag{11}$$

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).

[11] We refer to Gagliardone et al. (2024) for details about the data sources and variable definitions.

Finally, we recover the times series of firms' prices (in levels) by concatenating the price indices in Equation (10), $P_{ft} = P_{f0} \prod_{\tau=t_f^0+1}^{t} \left(P_{f\tau}/P_{f\tau-1}\right)$, 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 (11).

Frequency of Price Adjustment. The frequency of price adjustments is a crucial variable to characterize the strength of nominal rigidities. To accurately measure this variable, we use additional micro-level records from the National Bank of Belgium Business Survey (NBB-BS). This survey interviews a representative sample of firms within each manufacturing industry on a monthly basis, asking about their pricing decisions. In a manner similar to the official Producer Price Index (PPI) data collection, the survey asks firms if they increased, decreased, or left unchanged the price of a given product in their portfolio. This allows us to define a Boolean variable that takes value one if, within a given quarter, the firm reports adjusting prices at least once relative to the previous month. Averaging the boolean variables across firms and industries in any given quarter, we compute the average frequency of price adjustments for the manufacturing sector. We denote this variable by $\bar{h}_t$.

The information on firms' price adjustments also helps us to clean up spurious price changes in the microdata. As discussed, our measure of prices is constructed from product-level unit values, which tends to overstate the frequency of small price changes due to small measurement errors, as shown by Eichenbaum et al. (2014).[12] To address this measurement problem, we combine the firm-level price changes with information on the frequency of price adjustments from the NBB-BS to define firm-specific thresholds, κ^+ and κ^-, such that a small price adjustment below these thresholds is treated as no price change:

$$\mathbb{I}_{ft}^{+} = 0 \iff \Delta p_{ft} < \kappa_h^{+} \cdot Var_f(\Delta p_{ft}) \qquad \text{if } \Delta p_{ft} > 0$$

$$\mathbb{I}_{ft}^{-} = 0 \iff \Delta p_{ft} > -\kappa_h^{-} \cdot Var_f(\Delta p_{ft}) \qquad \text{if } \Delta p_{ft} < 0$$

To account for different degrees of upward and downward nominal rigidity in the data, we

[12] For example, data from various countries reveals that the share of regular price changes that are smaller than 1 percent in absolute value is 3 to 4 percent (see, e.g., Cavallo and Rigobon 2016). This figure is 30 percent in our data, suggesting that many of the small price changes are spurious price changes.

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

Marginal cost indices. To derive a firm-level marginal cost index, we assume a cost structure in which the nominal marginal cost of a firm is proportional to its average variable costs: $MC^n_{ft} = (1 + \nu_f)AVC_{ft}$. 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:

$$mc^n_{ft} = (tvc^n_{ft} - y_{ft}) + \ln(1 + \nu_f) \tag{12}$$

In the data, we measure total variable costs as the sum of intermediate costs (materials and services purchased) and labor costs (wage bill), both of which are available at the firm-quarter level (the former are taken from tax declarations, the latter from social security declarations). We compute a quantity index by dividing a firm's domestic revenues by its domestic price index.[14] 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 provides a second rationale for removing firm-fixed effects from our measure of price gaps, together with the normalization of the price index discussed above.

Ex-ante price gaps. The availability of high-frequency firm-level price and cost data allows us to construct an empirical counterpart to ex-ante price gaps:

$$x'_{ft-1} = p^o_{ft} - p_{ft-1}.$$

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

[14] Specifically, we compute $Y_{ft} = (PY)_{ft}/\bar{P}_{ft}$, where $\bar{P}_{ft}$ denotes the firm-quarter domestic price index. For single-industry firms, $\bar{P}_{ft}$ coincides with the firm-industry price index P_{ft}. For multi-industry firms, we construct $\bar{P}_{ft}$ as an average of the different firm-industry price indices using as weights the firm-specific revenue shares of each industry. As discussed in Gagliardone et al. (2024), 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.

As discussed in Section 1, when p^o_{ft} and $p^\star_{ft}$ are sufficiently close, $x'_{t-1}(f)$ captures inefficiencies driven by nominal rigidities. A positive ex-ante 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. (2024). 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 (Gagliardone et al. 2024), they are a source measurement error weakening the connection between our empirical measures of price gaps and price changes at the firm level.

2.2 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 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

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

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]

2.3 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 ex-ante 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 changes is $\bar{h} = 0.29$. 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 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

[16] As noted above, we can identify firms' reset gaps only up to a firm-specific additive constant. This constant controls for 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 (12)), and the normalization of price levels discussed in Section 2.1. This demeaning also removes the average trend inflation rate, which is approximately 0.6 percent quarter-on-quarter before the 2021 surge. Finally, we remove seasonal patterns in nominal variable costs by including quarter intercepts, which are higher in the second and fourth quarters and vary between industries.

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 ex-ante price gaps (x'_{ft-1}) before (panel a) and after the inflation surge (panel b).

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 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.

Figure 5: Empirical distribution of ex-ante price gaps

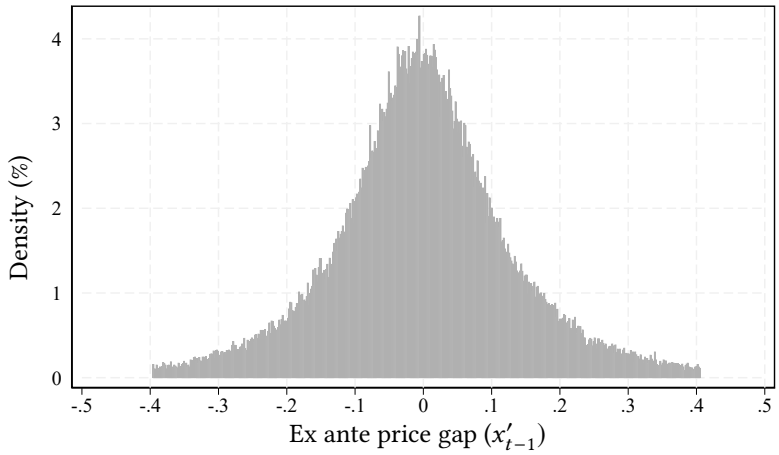

Notes. The figure presents the empirical probability density function of the price gaps, $f(x'_{ft-1})$, in the pre-pandemic period (2000-2019).

3 Micro evidence of state dependent pricing

Guided by the theoretical 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.

3.1 The empirical Generalized Hazard Function

The relationship between price gaps and the frequency of price adjustments is the watershed between state- and time-dependent pricing models. State-dependent models imply a monotonically increasing relationship between a firm's probability of price adjustment (captured by the GHF) and the absolute value of its price gap. In contrast, the two variables are independent in time-dependent models, resulting in a flat GHF.

We test these predictions using microdata on the frequency of price adjustments and price gaps, focusing on the pre-pandemic period (2009–2019). This period is treated as a steady state, with zero aggregate shocks and gap variations driven solely by idiosyncratic cost and markup shocks. In Figure 6, the black line shows the probability density function of price gaps, while the red line represents the empirical analog of the theoretical GHF, measuring the fraction of firms adjusting prices in each bin of the price gap distribution.

Figure 6: Empirical GHF and distribution of ex-ante price gaps

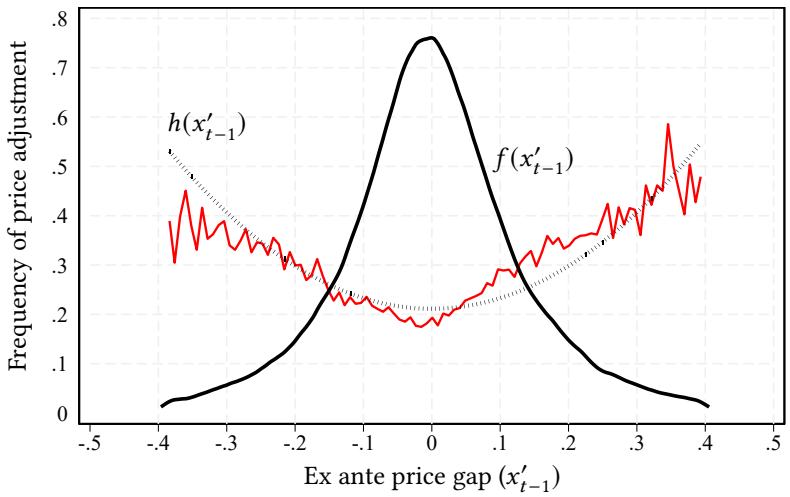

Notes. The figure plots the empirical probability density function of the ex-ante price gaps $f(x'_{t-1})$ (black line) against the empirical GHF, $h(x'_{t-1})$ (red line). The black dotted line is the fitted value obtained from a cross-sectional regression of the frequency of price adjustments of a given bin (b) on a constant and the square of the average price gap of the same bin, as dictated by Equation (8). In the regression, we weight each bin by the number of observations it counts.

The data reveal a strong connection between the size of price gaps and the frequency of price adjustments. Wider gaps (i.e., higher deviations from $x'_{t-1} = 0$) are associated with a greater likelihood of price adjustment, with a striking similarity to the theoretical GHF of state-dependent pricing models. The shape of the empirical generalized hazard function is equally noteworthy. The GHF displays a steeper slope to the right, indicating an asymmetry. This asymmetry suggests that firms' incentives to adjust prices are greater when their prices are too low (i.e., when x is positive, resulting in a realized markup that is too low) compared to when their prices are too high.

As discussed in Section 1, under some assumptions, we can approximate the GHF parametrically as a quadratic function of the price gap, as shown in Equation (8). To empirically evaluate this expression, we partition the distribution of price gaps into 500 narrowly defined and equally sized bins (b). We then estimate a cross-sectional regression of the frequency of price adjustments for each bin, h_b, on a constant term and the square of the average price gap for the same bin, x'_b:

$$h_b = a_0 + a_1 \cdot (x'_b)^2 + u_b, \tag{13}$$

where u_b is a white noise. The black dotted line in Figure 6 represents the fitted values

obtained from the model. In the regression, we weight each bin by the number of observations it counts. As we can see, a simple quadratic polynomial fits the data quite well, despite the asymmetry between upward and downward adjustment probabilities. Note that, in the vicinity of $x'_{t-1} = 0$, the intercept estimate ($\hat{a}_0$) provides an estimate of the parameter θ^o, which controls free price adjustments. This mapping will prove useful for calibration purposes.

3.2 Nonlinear price dynamics along the price gap distribution

In our next exercise, we document the nonlinear cost-price dynamics resulting from the state dependence of firms' policies. To illustrate this point, we make use of Equation (9), which links inflation to (odd) moments of the price gap distribution.

Figure 7: Nonlinear price dynamics

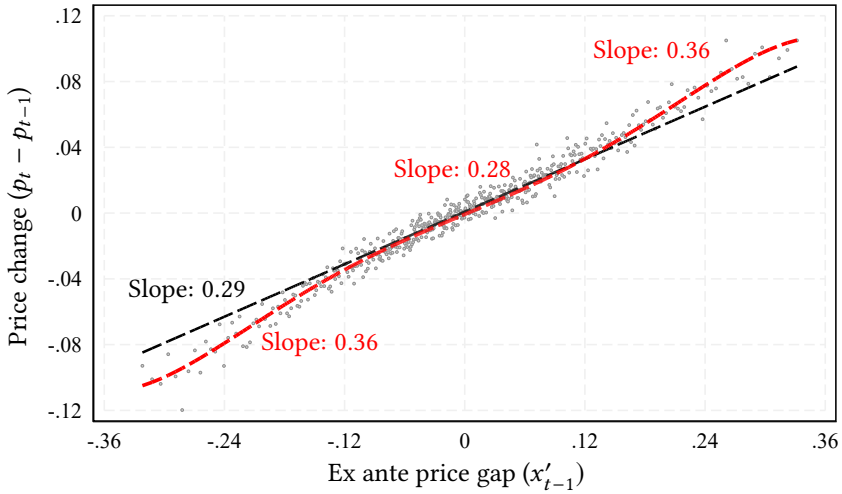

Notes. This figure presents a scatter plot of the average ex-ante price gap of a given bin of the price gap distribution, $x'(b)$, against the average logarithmic price change belonging to the same bin, $\pi(b)$. Bins are defined partitioning the price gap distribution into 500 narrowly defined intervals of width ≈ 0.002. The black dashed line depicts a linear fit of price changes on price gaps, $\hat{\pi}(b) = \hat{a}_1 \cdot x'(b)$, estimated using only bins that belong to the center of the price gap distribution (from the 25th to the 75th percentile). We report in black the estimated slope ($\hat{a}_1$). The red dashed line is the fit of a polynomial in the 1 first and 3 rd order of the gap, $\hat{\pi}_b = \hat{b}_1 \cdot (x'_b) + \hat{b}_2 \cdot (x'_b)^3$, estimated using bins throughout the entire distribution of price gaps. We report in red the slope of the polynomial fit in the center and in the tails of the distribution of price gaps. In all regressions, each bin is weighted by the number of observations it contains.

Again we sort observations into narrowly defined, equally sized bins spanning the entire price gap distribution. In Figure 7 we plot the average price gap for a given bin (x'_b,

x-axis) against the average price change for observations falling within the same bin (π_b, y-axis). We then estimate the following cross-sectional regressions, which constitute the sample analog of Equation (9):

$$\pi_b = b_1 \cdot \left(x'_b\right) + b_2 \cdot \left(x'_b\right)^3 + \eta_b$$

In Appendix A.6, we show that the coefficient b_1 consistently estimates the average frequency of price adjustments in the regression sample ($\bar{h}$) up to the third-order approximation in the gaps. The coefficient b_2 captures the importance of nonlinearities driven by state-dependent pricing as it estimates ϕ, the curvature parameter of the GHF.

By comparing the patterns in Figure 3 with their theoretical counterpart in Figure 7 we can see just how closely the microdata align with the predictions of the model. Consider first the bins located at the center of the distribution (bins covering the 25th to 75th percentiles). Observations in this range are characterized by relatively low gaps, which means relatively small deviations in their prices from the target price. We can therefore think of them as representing the distribution of the economy in "normal times," with low inflation and small aggregate shocks. For these observations, the cubic term is small and can be ignored. Thus, the relationship between inflation and price gaps is approximately linear, as in the Calvo model. This result echoes those of Gertler and Leahy (2008), Alvarez et al. (2017), and Auclert et al. (2024), which highlight how, in "normal times," the price dynamics generated by a state-dependent model resembles those generated by a time-dependent model. To this point, the estimate of the slope coefficient is $\hat{b}_1 = 0.28$ (or 0.29, if we ignore the cubic term), which matches well the frequency of price adjustments observed in our sample in low inflation environment (Table 1, panel a) and the estimates of the degree of nominal rigidity in Gagliardone et al. (2024).[17]

Now consider the relationship between price gaps and price changes across the entire price gap distribution, including its tails. The red dashed line in Figure 7 represents the projection of inflation on the price gap and the cubic of the price gap of each bin. The microdata offer an opportunity to appreciate the nonlinearities generated by the state-dependent nature of price adjustments. When price gaps are wide (for example, when a large aggregate shock hits the economy), the sensitivity of price changes to price

[17] As shown in Appendix A.6, omitting higher-order terms from the regression would result in an upward bias in the estimate of $\hat{b}_1$. However, this bias in the linear specification turns out to be small (as it increases the estimated frequency by only one percentage point). For our purposes, this indicates that a third-order approximation is enough to correctly estimate the linear coefficient.

gaps increases substantially because of the presence of the cubic term. In fact, the gradient between price gaps and price changes steepens by 25 percent (from 0.28 to 0.36) as we move from the center to the tails of the price gap distribution.

3.3 Price gaps and price changes conditional on adjustment

So far, we have studied the micro-level relationship between gaps and pricing, averaging across both firms that adjust and those that do not. We now focus our attention on the former group of firms. The theory predicts that, conditional on adjusting, the firms set $p_{ft} = p^{\star}_{ft}$. This implies that $p_{ft} - p_{ft-1} = x^{\star}_{ft}$, when $p_{ft} \neq p_{ft-1}$. Although we cannot measure $x^{\star}_{ft}$ in the data, to the extent that p^{o}_{ft} provides a reasonable approximation for $p^{\star}_{ft}$, we should still observe an elasticity of price changes with respect to ex-ante price gaps, $(p_{ft} - p_{ft-1})/x'_{ft-1}$, that is approximately one. Figure 8 shows that this is the case.

It presents two binned scatter plots that report, on the x-axis, the average price gap of a given percentile of the price gap distribution of adjusters (that is, firms for which $p_{ft} \neq p_{ft-1}$) and the corresponding average percentage change in the prices of firms in the same percentile (y-axis). The left panel focuses on the pre-pandemic period (1999–2019), and the right panel on the pandemic and post-pandemic period (2020–2023). In each panel, the black dashed line depicts the linear fit of price changes on price gaps across the percentiles of the distribution of price gaps. The data indicate a gradient that is not only positive but also very close to one, as the theory suggests. 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, we find that the gradient is particularly steep during the post-pandemic period. This could be due to firms being more attentive and reactive to cost movements when inflation is high, as shown by Gagliardone and Tielens (2024) using a model with state-dependent information frictions.

3.4 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

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

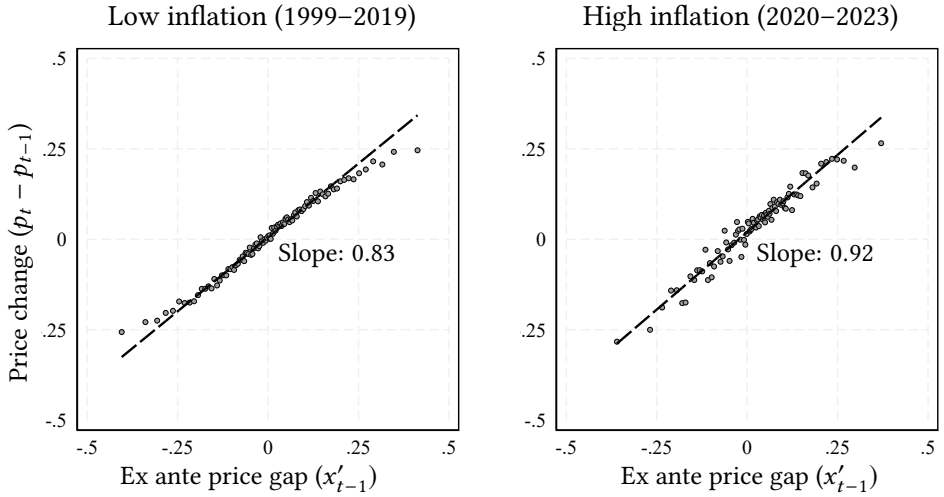

Notes. This figure presents a binned scatter plot of the log price change for adjusters (i.e., firms for which $p_{ft} \neq p_{ft-1}$) against the ex-ante price gap. Each dot marks the average price gap of a given percentile of the price gap distribution (x-axis) and the corresponding average percentage 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.

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 the price gaps before the pandemic. In panel a, the red dashed line represents the distribution in 2022:Q2. During this quarter, on average, firms' marginal costs increased by 6.2 percent relative to the previous quarter. Accordingly, and consistent with the theoretical predictions, a cost shock of this magnitude shifts the entire price gap distribution to the right, so that a significant number of firms' prices are now further away from their desired levels, resulting in a shift of the distribution and fatter tails. Because the shock reduces firms' profit margins, more firms move to regions of the price gap distribution where the GHF is high and the cost of not adjusting is large. Consequently, over a single quarter, the average probability of price adjustment almost doubles relative to the frequency observed in normal times. Once again, a time-dependent model would not produce any of these results, as its GHF is flat and orthogonal to the gap distribution.

In panel b, we repeat the same exercise, but now the red line represents the distribution of price gaps in 2023:Q3. During this quarter, on average, firms' marginal costs decreased by 3.8 percent relative to the previous quarter, as energy prices and

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

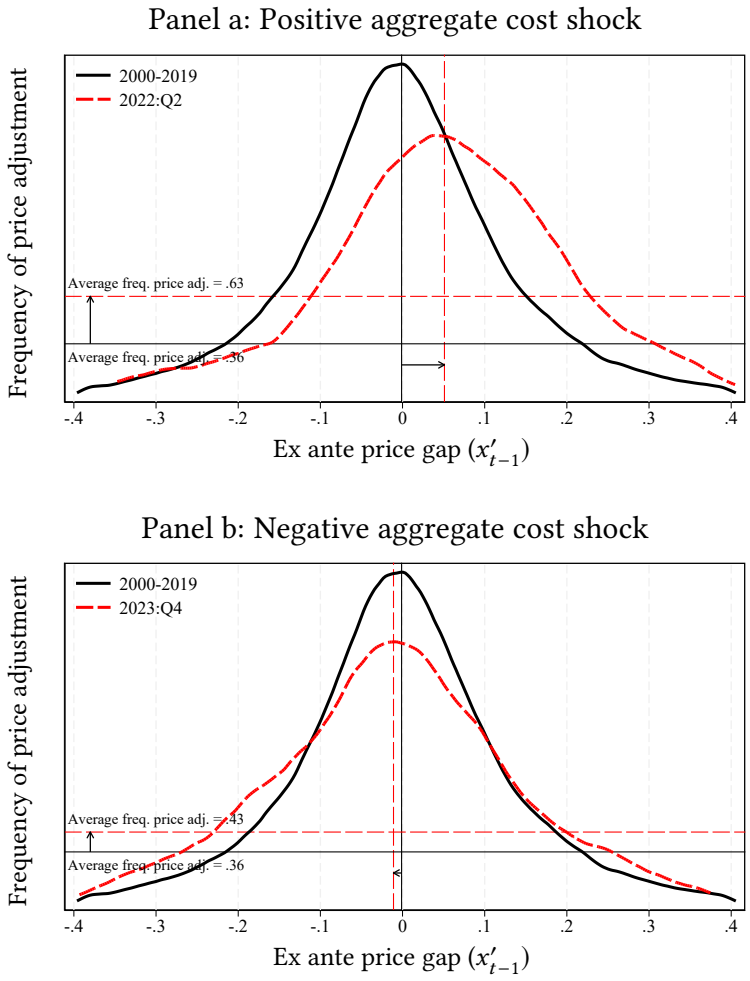

Notes. This figure presents the empirical probability density function of the ex-ante price gaps in the pre-pandemic period, 1999–2019, (black solid line) and in two snapshots of the post-pandemic period, in 2022:Q2 (red dashed line, panel a) and 2023:Q4 (red dashed line, panel b). The solid and dashed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

international supply chains began to normalize. This (negative) cost shock shifted the price gap distribution to the left, which led to an increase in the frequency of price adjustments as firms began to lower their prices.

4 Aggregate cost-price dynamics

We now turn our focus to the aggregate level. We show that the micro-level dynamics—and their state-dependent nature–translate into nonlinear inflation dynamics, with the degree passthrough of aggregate cost shocks into prices varying with the magnitude of the shock.

4.1 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):

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

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

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 results. 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,

Figure 10: Inflation, cost, and markup dynamics

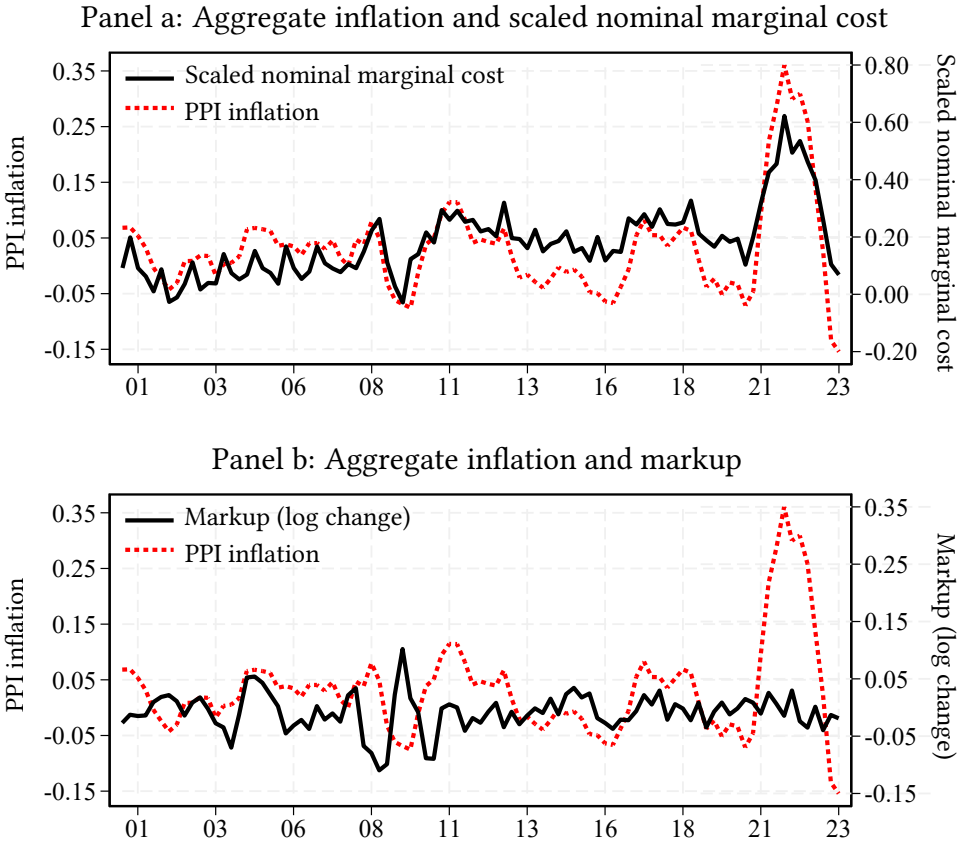

Notes. This figure shows the time series of year-over-year manufacturing PPI inflation ($p_t - p_{t-4}$) alongside the times series of the scaled nominal marginal cost index ($mc_t^n - p_{t-4}$, panel a) and the log change in average realized markups ($\Delta \ln(\text{Markup}_t)$, panel b) for the Belgian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

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 recover the latter as the difference between the former and the change in our nominal aggregate nominal marginal cost measure: $\Delta \ln(\text{Markup}_t) \equiv \pi_t - \Delta mc_t^n$. This exercise illustrates that, at least in our sample, the hypothesis that a rise in markups can explain the recent inflation surge seems to have no bite in the data.[18]

[18]Studying the price and cost data for a large global manufacturer, Alvarez et al. (2024) also find that markups remained stable over time, including during the inflation surge.

Figure 11: Decomposition of aggregate nominal marginal cost index

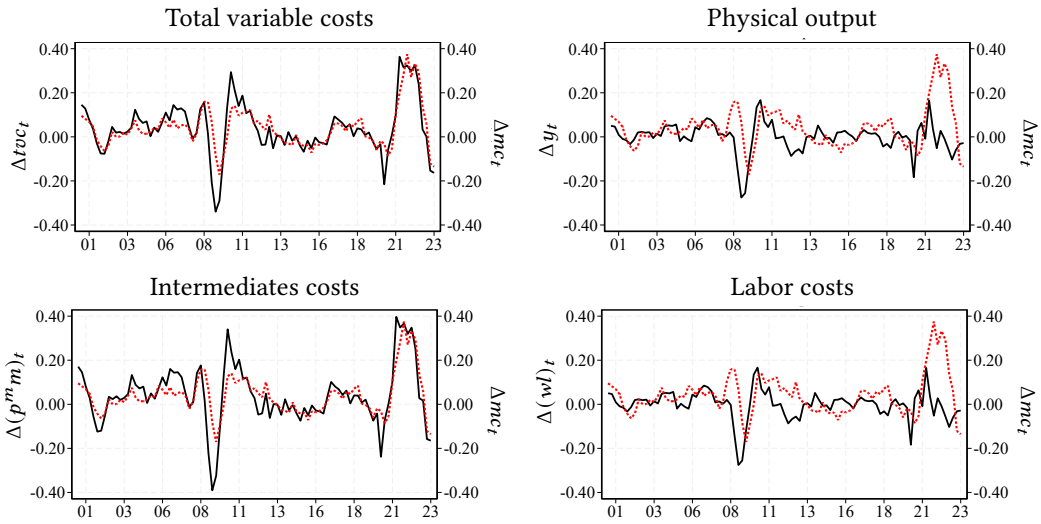

Notes. This figure decomposes the log change in our nominal aggregate marginal cost index into the log change in total variable costs (Δtvc_t, top left panel), real output (Δy_t, top right panel), intermediates costs ($\Delta(p^m m)_t$, bottom left panel), and labor costs ($\Delta(wl)_t$, bottom left panel).

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. Recall that we measure marginal cost as the ratio of total variable cost to real output (Equation (12)). 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 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.

4.2 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 3, we find a linear relationship between aggregate inflation and nominal costs during normal times. 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. (2024) in a low inflation environment.

Importantly, the linear relationship between the two variables breaks down when the economy is hit by large aggregate shocks. In fact, 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 adjustments. As before, the black dashed line shows the linear fit during periods of low inflation, and the red dashed line represents the fit across all inflation regimes. In low-inflation environments, we observe essentially no relationship between the average frequency of price adjustments and inflation. Once again, this observation aligns with the assumption of a Calvo model and suggests that a time-dependent pricing model provides a good framework to capture nominal rigidities in a low-inflation environment. However, inflation and the frequency of price adjustments are highly correlated in high inflation environments, as shown by Alvarez et al. (2019), Cavallo et al. (2023), and Blanco et al. (2024a).

Figure 12: Passthrough of costs into inflation

Panel a: Nonlinear passthrough of costs into inflation

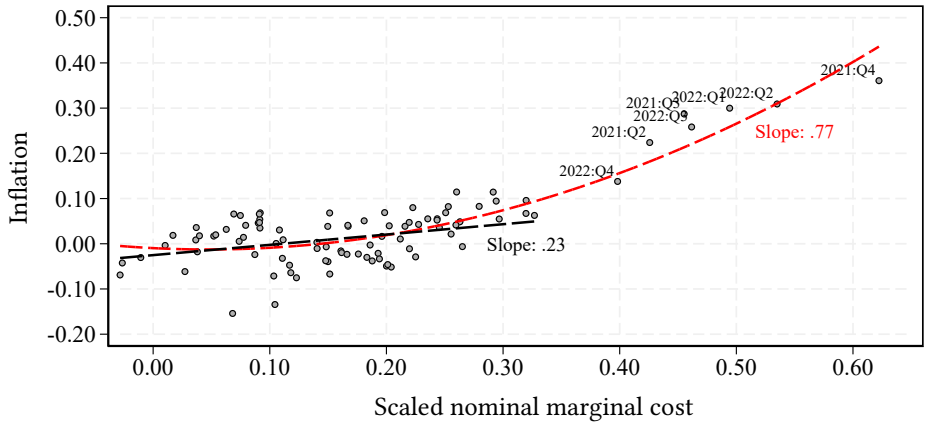

Panel b: Frequency of price adjustment and inflation

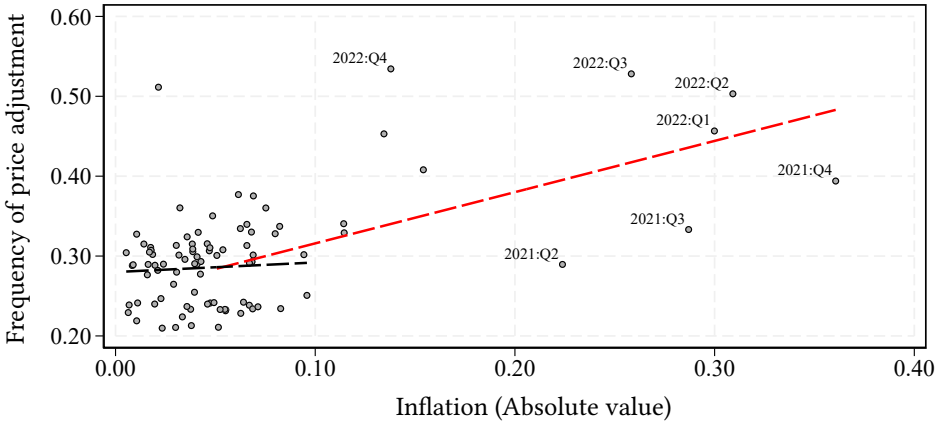

Notes. In panel a, we sort the different quarters in our data according to their realized aggregate scaled nominal marginal cost ($mc_t^n - p_t$) and plot this variable against year-over-year manufacturing PPI inflation (π_t) in the same quarter. In panel b, we sort the data according to realized year-over-year manufacturing PPI inflation and plot this variable against the average frequency of price adjustments ($\bar{h}_t$) in the same quarter. The average frequency of price adjustments is a rolling average of the quarterly frequency of price adjustments 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.

5 Quantitative implications

Having established the close connection between theory and data, we now use moments derived from microdata to calibrate and simulate the quantitative model presented in

Section 1. We use the calibrated model to perform two types of quantitative exercises. In a first, more standard, set of exercises, we compare the dynamics of our state-dependent model to those of a standard time-dependent Calvo model in response to small and large shocks. The second set feeds the model a sequence of aggregate marginal costs extracted from the data and compares the model-generated aggregate inflation series to the one observed in the data.

5.1 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. (2024). 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 adjustments, mediating the relationship between price gaps and price adjustments. The standard approach to calibrate these parameters leverages the theoretical mapping between the unobservable price gap distribution and measurable moments of the price change distribution: standard deviation, kurtosis, and average frequency of price changes.[19] Producing unbiased empirical measures of these moments is challenging in our context. The measurement of kurtosis is particularly problematic. As we discussed earlier, unit values tend to be more prone to measurement error. Furthermore, controlling for all relevant sources of unobserved heterogeneity can be difficult without detailed product-level information on goods and their points of sale.[20] To address these issues, we developed an alternative calibration procedure that does not rely on targeting the kurtosis of price adjustments. Instead, we leverage information subsumed in the joint distribution of price changes and

[19] See, e.g., the approach in Alvarez et al. (2022) and Blanco et al. (2024a).

[20] See Alvarez et al. 2016, Alvarez et al. (2022) and Cavallo and Rigobon 2016 for a discussion on biases in the measured kurtosis and how to address them when barcode data are available.

price gaps and in the empirical GHF during the pre-pandemic period.[21]

First, we have shown that the relationship between price gaps and the frequency of price adjustments observed in the microdata is well approximated by a quadratic GHF (see Figure 6), $h_{ft} = (1 - \theta^o) + \phi \cdot \left(x'_{ft-1}\right)^2$. Thus, we calibrate the free price adjustment parameter θ^o to match the frequency of price adjustments in a neighborhood of $x'_{ft-1} \approx 0$. Specifically, we estimate Equation (13) restricting the estimation sample to bins capturing observations in the 25 to 75 percentiles of the gaps distribution and calibrate $(1 - \theta^o) = \hat{a}_0 = 0.188$. In Appendix A.4, we provide a formal argument for the consistency of this estimator. Note that by averaging across observations in the zero-gap neighborhood, this calibration is robust to small measurement errors due to spurious changes in unit values.

Second, when trend inflation is low (as is the case during the pre-pandemic period) and idiosyncratic shocks are drawn from a Gaussian distribution, Alvarez et al. (2016) show that the following identity links the average frequency of price adjustments ($\bar{h}$), the variance of the price changes, and the variance of idiosyncratic shocks (σ_ϵ^2) in steady-state:

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

We thus simulate the model assuming that the idiosyncratic shocks $\varepsilon_t(f)$ are i.i.d draws from a Gaussian distribution $\mathcal{N}(0, \sigma_\epsilon^2)$ and calibrate σ_ϵ^2 to 0.0036 to match the product of the average frequency of price adjustments 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

[21] In Appendix C, we show that the results of our quantitative exercises are robust to an alternative calibration that targets the kurtosis of price adjustment.

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)

empirical evidence of small menu costs documented in Levy et al. (1997) and Zbaracki et al. (2004). Second, as discussed above, we did not target the kurtosis of price changes in our calibration. Proviso the measurement issues discussed above, the calibrated model displays a kurtosis of price changes that is close to the ones computed in our data.

We also consider a standard Calvo model calibrated to match the steady-state frequency of price adjustments 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, $1 - \theta^o$, is re-calibrated to match the steady-state frequency of price adjustment. 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 inconsistent with the data.

5.2 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 adjustments (panel a, left) and aggregate inflation (panel a, right). All shocks increase the optimal reset price, shifting the distribution of price gaps to the right, thereby

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

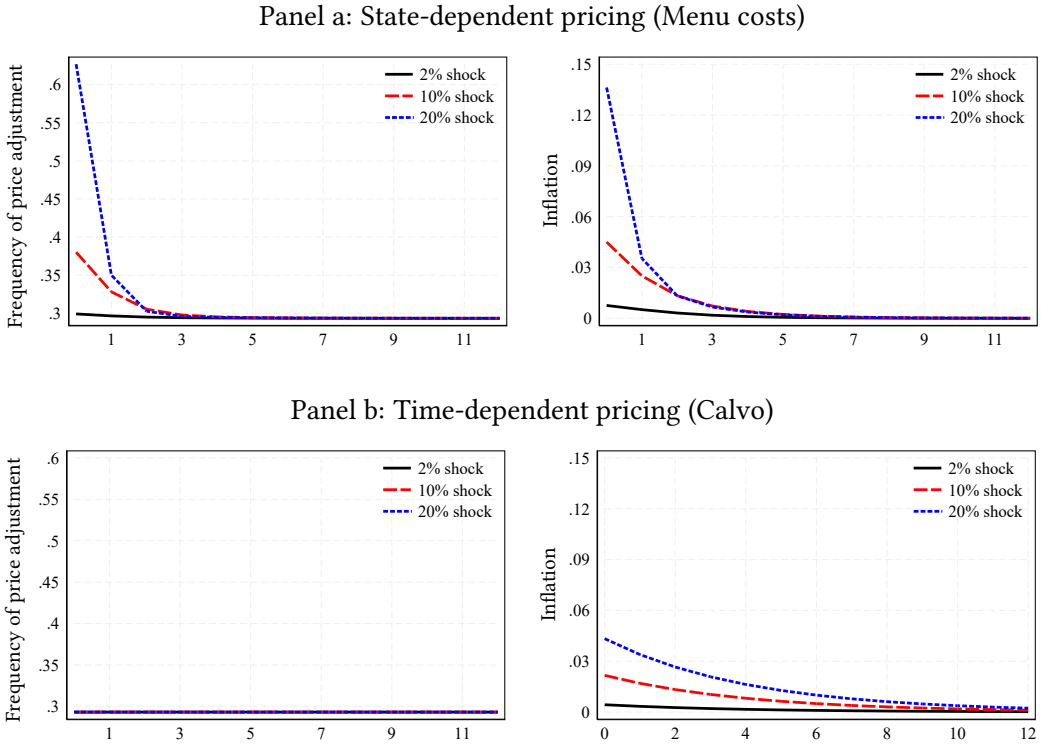

Notes. This figure presents the impulse responses of inflation and frequency to aggregate cost shocks of different magnitudes. Panel a reports the impulse response for our state-dependent pricing model (menu-cost model). Panel b reports the impulse responses for a time-dependent model (Calvo model).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

triggering an increase in the number of firms adjusting their prices and, therefore, inflation. However, as discussed in Section 1 and empirically shown in Section 3, large shocks lead to a significant shift in the price gap distribution, displacing many firms in a region where the GHF is higher, generating a spike in the frequency of price adjustments and, consequently, a rapid and substantial surge in inflation.

These exercises highlight the nonlinearities of state-dependent pricing as shocks grow in magnitude. For example, on impact, the effect of the large shock on both the frequency of price adjustments 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%). To highlight these features, it is useful to compare the IRFs of the menu-cost model with those obtained from the Calvo model (Figure 13, panel b). By construction, in Calvo, the number of firms adjusting their prices is not affected by the magnitude of the shock

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

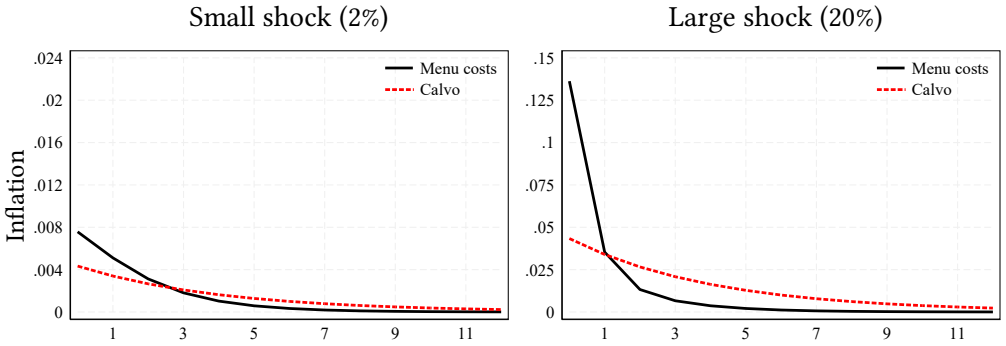

Notes. This figure presents the impulse responses of aggregate inflation to marginal cost shocks of different sizes in the menu-cost model and in the Calvo model. The x-axis reports quarters since the shock.

(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.

The 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 the same size (small or large). Although 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 adjustments 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 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.

5.3 Explaining the time series of inflation

Having characterized the price dynamics in response to shocks of different magnitudes, we now turn to evaluating the model's ability 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 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

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

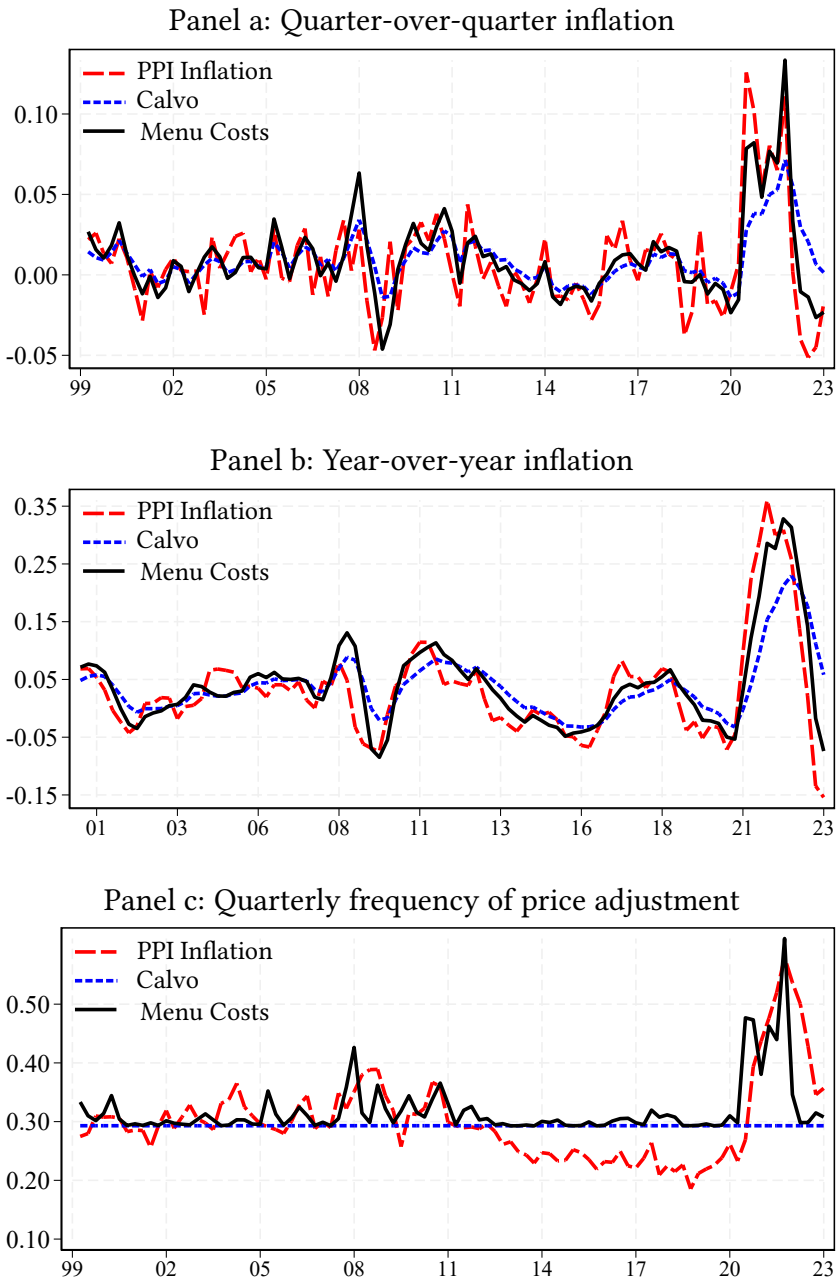

Notes. This figure contrasts the dynamics of PPI manufacturing inflation in the data to the inflation dynamics generated by the Calvo and menu-cost models, after feeding the model a sequence of aggregate nominal marginal cost shocks that matched the one observed in the data.

the faster cost passthrough generated by the state-dependent pricing policies documented in the impulse response function of Section 5.2.

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.

6 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 these data 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. 2024.

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.[22]

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

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

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

A.3 Quadratic approximation of Generalized Hazard Function

We now derive the expression for the quadratic approximation of the hazard function in Equation (8) and describe how we take this equation to the data. We take a second-order approximation of the hazard function $h_t(x'_{t-1})$ characterized in Equation (6) around $x_t^\star$ to obtain:

$$\begin{aligned} h_t(x'_{t-1}) &= (1-\theta^0) - \frac{\theta^0}{\bar{\chi}} \frac{\partial V_t(x)}{\partial x}\Big|_{x=x_t^\star}(x'_{t-1} - x_t^\star) - \frac{\theta^0}{\bar{\chi}} \frac{\partial^2 V_t(x)}{\partial x^2}\Big|_{x=x_t^\star}(x'_{t-1} - x_t^\star)^2 + o(x'_{t-1} - x_t^\star)^2 \\ &= (1-\theta^0) - \frac{\theta^0}{\bar{\chi}} \frac{\partial^2 V_t(x)}{\partial x^2}\Big|_{x=x_t^\star}(x'_{t-1})^2 + o(x'_{t-1})^2, \end{aligned}$$

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

$$h_t(x'_{t-1}(f)) = (1-\theta^o) + \phi \cdot ((x'_{t-1}(f))^2 + o(x'_{t-1}(f))^2, \tag{A.6}$$

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

To take Equation (A.6) to the data, we partition the support of the distribution of the ex-ante price gap into equally spaced bins. Denote the first and second moments of the distribution of $x'_t(f)$ within each bin by $x'_b \equiv \int_{f\in b} x'_{t-1}(f)\, df$ and $\sigma_b^2 \equiv \int_{f\in b}(x'_{t-1}(f))^2\, df - (x'_b)^2$. The frequency in bin b is given by:

$$h_b(x'_b) \equiv \int_{f\in b} h_t(x'_{t-1}(f))\, df = (1-\theta^o) + \phi\Big((x'_b)^2 + \sigma_b^2\Big) + o(x'_b)^2 \tag{A.7}$$

with $o(x'_b)^2 \equiv \int_{f\in b} o(x'_{t-1}(f))^2 df$. In the data, we plot the cross-sectional regression:

$$h_b(x'_b) = a_1 + a_2 \cdot \left(x'_b\right)^2 + \nu_b. \tag{A.8}$$

Matching the frequency of price changes ($\bar{h} = \int h_b(x'_b)db$) as explained in section 2.1, and estimating model (A.8) allows us to robustly identify the free-adjustment parameter, θ^0, by the average of the adjustment frequencies for firms with small gaps.

A.4 Estimator of the frequency of free adjustments

We now show how knowledge of the empirical GHF allows us to recover the frequency of free price adjustments, θ^0. We partition the distribution of price gaps into equally spaced bins denoted by b. 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 $h(b) \equiv h_b(x'_b)$ 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:

$$\int_{(-b',b')} h(b)db \geq 1 - \theta^0.$$

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^0. Therefore, by the monotone convergence theorem, it converges to its supremum which is given by:

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

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

A.5 Cubic approximation of inflation within a bin

Starting from the expression for aggregate inflation in Equation (7), we want to derive the cubic expression for inflation within a bin in Equation (9), under the assumption that $p_t^\star(f) \approx p_t^o(f)$:

$$\pi_b \approx \phi_b^0 \cdot \left(x'_b\right) + \phi \cdot \left(x'_b\right)^3,$$

and associated cross-sectional regression model.

We partition the distribution of price gaps into equally spaced bins denoted by b and adopt the same labeling convention of bins described in Appendix A.4.

Using the convexity of h, sufficiently close to the reset gap ($x^\star = 0$), the frequency of price adjustments is approximately constant. Hence, the frequency of all firms within the same (small) bin is approximately the same and therefore the following covariance is small:

$$\begin{aligned} Cov_b(h_t(f), x'_{t-1}(f)) &\equiv \int_{f\in b} \left(h_t(f)\cdot x'_{t-1}(f)\right) df - h_b(x'_b)x'_b \\ &= \int_{f\in b} \left\{\left(h_b(x'_b) + o(x'_{t-1}(f))^2\right)\cdot x'_{t-1}(f)\right\} df - h_b(x'_b)x'_b \\ &= \int_{f\in b} \left\{o(x'_{t-1}(f))^2 \cdot x'_{t-1}(f)\right\} df = o(x'_b)^3. \end{aligned}$$

It follows that inflation within a bin simplifies to:

$$\pi_b = \int_{f\in b} h_t(x'_{t-1}(f))\, df \cdot \int_{f\in b} \left(x'_{t-1}(f)\right) df + o(x'_b)^3 = h_b(x'_b)\cdot x'_b + o(x'_b)^3.$$

We then use the expression in (A.7) to substitute for $h_b(x'_b)$ in the equation above and define the bin-specific coefficient $\phi_b^0 \equiv (1-\theta^o) + \phi\sigma_b^2$ to obtain Equation (9) in the paper:

$$\begin{aligned} \pi_b &= (1-\theta^o) + \phi\left((x'_b)^2 + \sigma_b^2\right)\cdot x'_b + o(x'_b)^3 \\ &= \phi_b^0 x'_b + \phi(x'_b)^3 + o(x'_b)^3. \end{aligned} \tag{A.9}$$

The equation above represents the data-generating process behind the binned scatter plot in Figure (7). To take this equation to the data, we estimate the following cross-sectional regression model:

$$\pi_b = b_1 x'_b + b_2 (x'_b)^3 + \eta_b. \tag{A.10}$$

A.6 Estimator of the frequency of price changes

We partition the distribution of price gaps into equally spaced bins denoted by b and adopt the same labeling convention of bins described in Appendix A.4. Denote the variance across all firms with:

$$\bar{\sigma}^2 \equiv \int_{[0,1]} (x'(f))^2 df - \left(\int_{[0,1]} x'(f) df\right)^2$$

Our normalization of gaps implies zero cross-sectional average:

$$\int_{[0,1]} x'(f) df = \int x'(b) db = 0.$$

Averaging equation (A.6), the (average) frequency of price changes is:

$$\begin{aligned}\bar{h} \equiv \int_{[0,1]} h_t(f)df &= (1-\theta^o) + \phi\bar{\sigma}^2 + \mathbb{E}(o(x'_b)^2) \\ &\approx (1-\theta^o) + \phi\int_{[0,1]}(x'(f))^2 df\end{aligned}$$

Adding and subtracting $\phi\bar{\sigma}^2$ to Equation (A.9) we obtain the formula for inflation in a bin:

$$\pi_b = \bar{h}\cdot x'_b + \phi(x'_b)^3 + \underbrace{\phi(\sigma_b^2 - \bar{\sigma}^2)\cdot x'_b + o(x'_b)^3 + \mathbb{E}(o(x'_b)^2)x'_b}_{\equiv \eta_b} \tag{A.11}$$

We want to prove that the error term is mean zero, $\mathbb{E}(\eta_b) = \mathbb{E}(o(x'_b)^3)$, and the orthogonality condition $Cov(\eta_b, x'_b) = o(x'_b)^3$ holds.

Mean-zero error. First, we show that the error term in regression (A.10) has mean zero. Using $\mathbb{E}(x'_b) = 0$, we have that:

$$\begin{aligned}\mathbb{E}(\eta_b) &= \mathbb{E}\left(\phi(\sigma_b^2 - \bar{\sigma}^2)\cdot x'_b\right) + \mathbb{E}(o(x'_b)^3) + \mathbb{E}(o(x'_b)^2)\,\mathbb{E}(x'_b) \\ &= \phi\,\mathbb{E}\left(\sigma_b^2 - \bar{\sigma}^2\right)\cdot\mathbb{E}(x'_b) + \phi Cov(\sigma_b^2, x'_b) + \mathbb{E}(o(x'_b)^3) \\ &= \phi Cov(\sigma_b^2, x'_b) + \mathbb{E}(o(x'_b)^3)\end{aligned}$$

We want to show that the covariance $Cov(\sigma_b^2, x'_b)$ is equal to zero. First, due to the relabelling, for all bins b', it holds that $x'(-b') = -x'(b')$. Denote by $B(x_b) \equiv \{x(f)|f \in b\}$ the bin around x_b and by $\mathbb{I}_f\{B(x'_b)\} \equiv \mathbb{I}\left\{x'(f) \in B(x'_b)\right\}$ the indicator function for firm f belonging in bin b. Under the assumptions of symmetry on the distribution of gaps, firms are identically distributed in specular bins b' and $-b'$. Formally, for any measurable even function g with symmetry point at zero, it holds that:

$$\int_{[0,1]}\left(g(x'(f))\cdot\mathbb{I}_f\{B(x'_{b'})\}\right)df = \int_{[0,1]}\left(g(-x'(f))\cdot\mathbb{I}_f\{B(x'_{-b'})\}\right)df.$$

It follows that $\sigma^2(b)$ coincides for every two specular bins b' and $-b'$:

$$\begin{aligned}\sigma^2(b') &= \int_{[0,1]}\left((x'(f))^2 - (x'(b'))^2\right)\cdot\mathbb{I}_f\{B(x'(b'))\}\,df \\ &= \int_{[0,1]}\left((-x'(f))^2 - (-x'(b'))^2\right)\cdot\mathbb{I}_f\{B(-x'(b'))\}\,df \\ &= \int_{[0,1]}\left((x'(f))^2 - (x'(-b'))^2\right)\cdot\mathbb{I}_f\{B(x'(-b'))\}\,df = \sigma^2(-b')\end{aligned}$$

Therefore $\sigma^2(b)$ is an even function with symmetry point at $\tilde{b} = 0$. Next, note that the

product of an even function with symmetry point at $\tilde{b}$ and odd function with the same symmetry point is an odd function with symmetry at $\tilde{b}$. Therefore the integral is zero:

$$Cov(\sigma_b^2, x_b') = \int \sigma^2(b) \cdot x'(b) db = 0.$$

Orthogonality. We now show that estimator $\hat{b}_1$ from model (A.10) consistently recovers $\bar{h}$ up to the third-order approximation in the gap. To do this, we need to show that the following exclusionary restriction across bins holds in a neighborhood of $x_b' = 0$:

$$Cov(x_b', \phi(\sigma_b^2 - \bar{\sigma}^2) \cdot x_b') \approx 0.$$

Using the formula of the covariance across bins and the results above:

$$\begin{aligned}
Cov(x_b', (\sigma_b^2 - \bar{\sigma}^2) \cdot x_b') &= \mathbb{E}\left\{(\sigma_b^2 - \bar{\sigma}^2) \cdot (x_b')^2\right\} \\
&= Cov(\sigma_b^2, (x_b')^2) + \mathbb{E}(\sigma_b^2 - \bar{\sigma}^2)\, \mathbb{E}((x_b')^2) \\
&= Cov(\sigma_b^2, (x_b')^2) - \mathbb{E}((x_b')^2) \cdot \mathbb{E}((x_b')^2) \\
&= Cov(\sigma_b^2, (x_b')^2) - \left(\mathbb{E}((x_b')^2)\right)^2 \\
&= Cov\left(\int_{f \in b} (x(f))^2 df, (x_b')^2\right) - Var((x_b')^2) - \left(\mathbb{E}((x_b')^2)\right)^2 \\
&= Cov\left(\int_{f \in b} (x(f))^2 df, (x_b')^2\right) - \mathbb{E}((x_b')^4),
\end{aligned}$$

where the term in the last equation is $o(x_b')^3$. Therefore, $\hat{b}_1$ consistently estimates the average frequency of price adjustments up to the third-order approximation.

Finally, we also prove that the estimator $\hat{b}_2$ consistently estimates ϕ up to the third-order approximation. This is immediate from the fact that:

$$Cov((x_b')^3, (\sigma_b^2 - \bar{\sigma}^2) \cdot x_b') = o(x_b')^3,$$

which follows from the same steps above.

B Additional quantitative exercises

In this section, we present additional quantitative exercises using the simulation of the calibrated menu-cost and Calvo models. The first set of exercises study how cost shocks affect both the static target price p_{ft}^o and the dynamic optimal price $p_{ft}^\star$. The second exercise studies the role of strategic complementarities.

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 ex-ante price gaps (Equations (7) and (9), respectively). The question is how well p^{o}_{ft} approximates $p^{\star}_{ft}$ away from the steady state.

Figure A.1: Impulse responses: Static vs. dynamic price targets

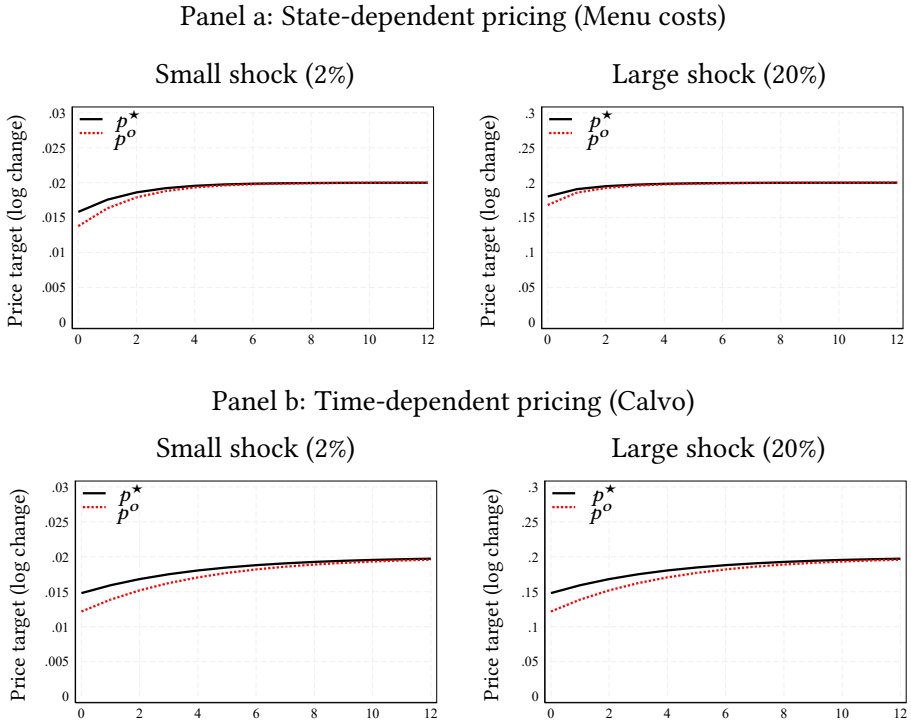

Notes. This figure presents impulse responses of the static target price (p^{o}) and the optimal reset price ($p^{\star}$) to aggregate cost shock of different sizes. The x-axis reports quarters since the shock.

The impulse response functions shown in Figure A.1 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 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.2 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 $p^\star_{ft}$.

Figure A.2: Quarter-over-quarter inflation: Static vs. dynamic price targets

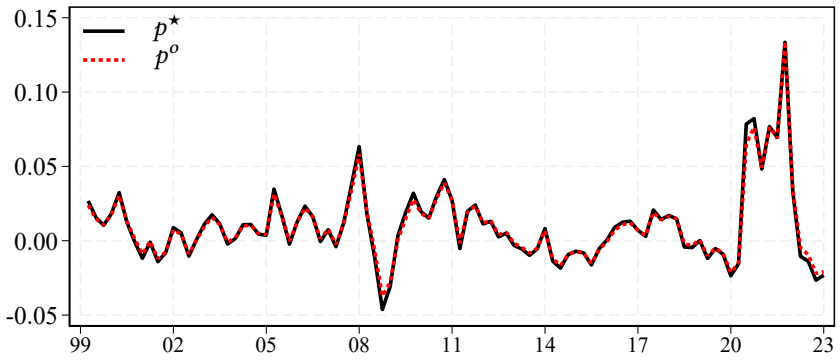

Notes. This figure contrasts the inflation dynamics generated by the menu-cost model using $p^\star$ (the exact, dynamic reset price) and using p^o (the static approximation of $p^\star$) when solving the model. As in Figure 15, we solve the model feeding it a sequence of aggregate nominal marginal cost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

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

Figure A.3: The role of strategic complementarities

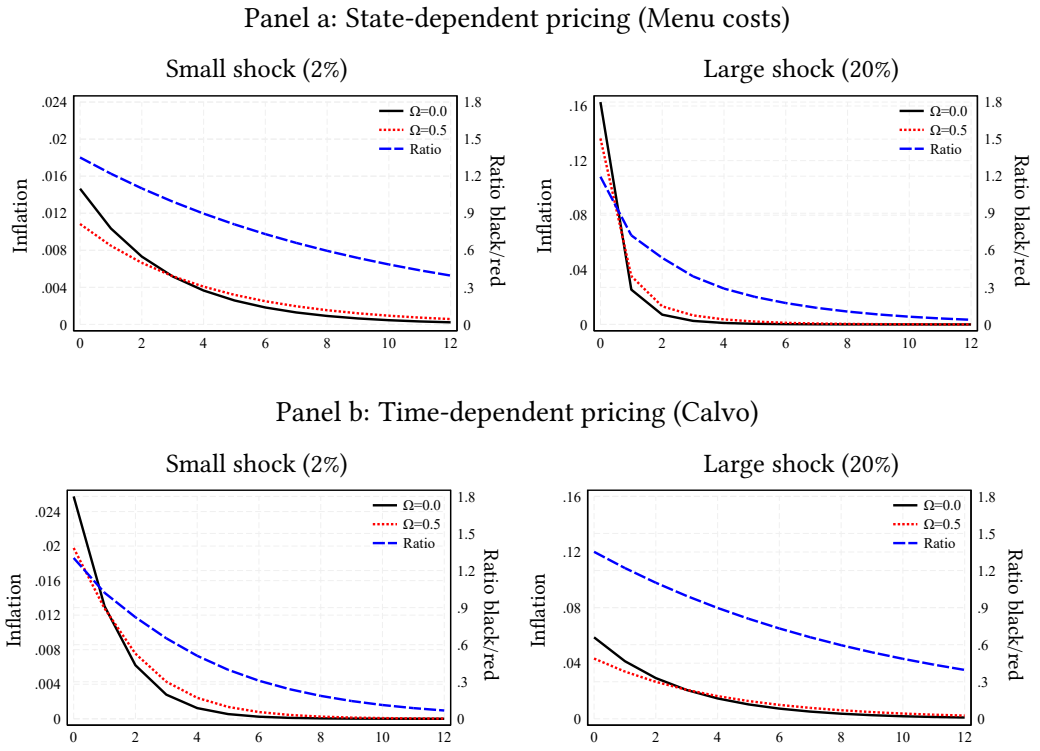

Notes. This figure presents the impulse responses of inflation to aggregate cost shocks of different sizes, without strategic price complementarities ($\Omega = 0$, black line) and with strategic complementarities ($\Omega = 0.5$, red dotted line). The blue dashed line represents the ratio of the impulse 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 adjustments as the time-dependent model. The x-axis reports quarters since the shock.

C Alternative calibration

In our baseline exercise, we used information from the empirical GHF to directly calibrate the frequency of price adjustments. As explained in the paper, this procedure tends to be more robust to small measurement error which can affect the measured kurtosis of price adjustments. As a robustness check, we perform an alternative calibration, this time targeting the kurtosis of price adjustments. We then repeat the feeding exercise discussed in Section 5.3 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

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	3.269	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 adjustments as well as the standard deviation and kurtosis of price adjustments.

average frequency of price adjustments, (ii) the standard deviation of price adjustments, and (iii) the kurtosis of price adjustments 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 adjustments 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 frequency 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 adjustments.

Figure A.4 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 adjustments well. However, the model fit is better under the 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.4: Inflation and frequency of price adjustment: Model versus data

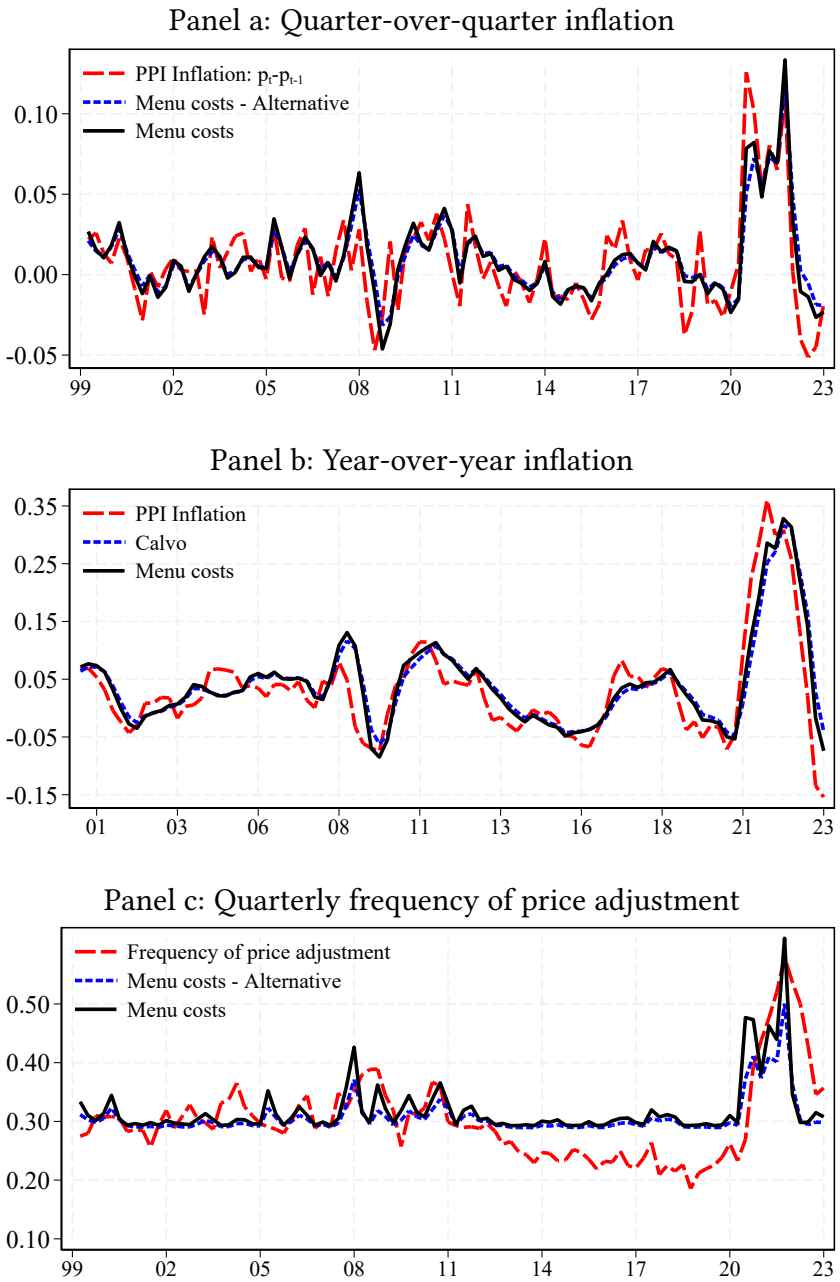

Notes. This figure contrasts the dynamics of PPI manufacturing inflation in the data to the inflation dynamics generated by our menu-cost models under our baseline (solid black line) and alternative calibration (dashed blue line).