

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

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

November, 2024

Abstract

We use microdata on firms' prices and production costs to study inflation dynamics in high- and low-inflation environments. We document that prices are set in a state-dependent way. Both the probability of a firm adjusting its price and the magnitude of such adjustments are functions of the firm's price gap—the percentage difference between a firm's ideal and current price. We then develop a generalized state-dependent pricing model to account for Belgian PPI inflation over the period 1999:Q1 to 2023:Q4. Conditional on a path of cost shocks extracted from the data, the model explains both the low and stable inflation of the pre-pandemic period as well as the pandemic era surge. In normal times, the adjustment probabilities are approximately constant and the model resembles a framework with time-dependent pricing as in Calvo (1983). During the surge, the model captures the rise in inflation along with the change in the price adjustment frequency, which is the driver of the nonlinear dynamics.

* New York University. Email: luca.gagliardone@nyu.edu; † New York University and NBER. Email: mark.gertler@nyu.edu; ‡ New York University, Stern School of Business. Email: slenzu@stern.nyu.edu; § National Bank of Belgium and KU Leuven. Email: joris.tielens@nbb.be. We thank Al-Mahdi Ebsim for outstanding research assistance and Virgiliu Midrigan, Luminita Stevens for helpful comments and conversations. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the National Bank of Belgium, the Eurosystem, or any other institution with which the authors are affiliated.

1 Introduction

In this paper, we use microdata on firms' prices and costs to study inflation dynamics in normal times and during the recent inflation surge. The key variable in our analysis is the firm's price gap, that is the gap between its ideal price (after the realization of shocks) and the price the firm is currently charging.

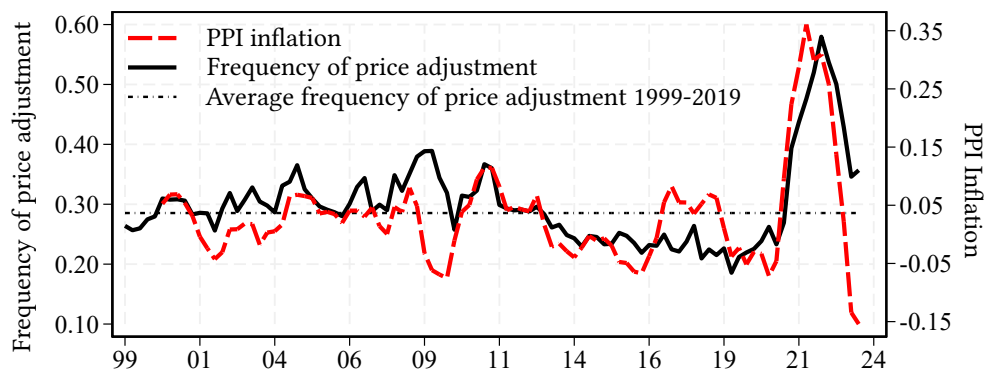
Every model with nominal price rigidity features a notion of a price gap. What distinguishes different models is how firms adjust their prices in response to their price gaps. For example, in the time-dependent pricing model of Calvo (1983), firms can adjust their prices with a fixed probability and the firm's expected price change is linear in its price gap. Conversely, in state-dependent pricing models (often known as menu cost models), the firm's expected price change is 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.

The distinction between time- and state-dependent pricing is less significant when the economy is in a low inflation environment, where shocks to desired prices are typically small on average¹. However, it becomes important in high inflation environments, which typically involve large shocks to desired prices. The recent inflation surge well illustrates these issues. Figure 1 displays the year-over-year percentage change in producer price index (PPI) for the Belgian manufacturing sector against the average frequency of price adjustment, between 1999:Q1 and 2023:Q4. Before the pandemic, both inflation and the average frequency were low and relatively stable. Starting in early 2021, inflation surged significantly and then began to collapse in mid-2022. Tracking inflation, the frequency of price adjustment displays a boom and bust, more than doubling over the course of the year, before gradually returning to its long-term trend. These dynamics are similar to what occurred in the US and various other developed economies worldwide (Blanco et al. 2024b; Cavallo et al. 2024).

We extend the micro-level data set assembled in Gagliardone et al. (2024)

¹See, for example, Dias et al. (2007), Gertler and Leahy (2008), Alvarez et al. (2017), and Auclert et al. (2024).

Figure 1: Aggregate inflation and frequency of price adjustment



Notes. This figure shows the time series of PPI manufacturing inflation along with the annual frequency of price 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.

to include the recent inflation surge. This dataset collects administrative records on product-level output quantities, sales, and production costs at the quarterly frequency for Belgian manufacturing firms. Using this data, we construct a notion of price gaps for individual firms that accounts for variation in costs, prices, and competitors prices. We analyze the data through the lens of a tractable menu cost model in tradition of the classic generalized state-dependent models proposed by Caballero and Engel (1993, 2007), Golosov and Lucas (2007), and Nakamura and Steinsson (2010). 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 accounting for the dynamics 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 probability of firm price adjustment (extensive margin) is increasing and nonlinear in the price gap. Second, when firms adjust their prices (intensive margin), they do so in order to close the price gap. Third, the non-linearity in the speed of price adjustment explains the non-linear inflation dynamics observed in the data after large shocks. During

the post-pandemic period, we show how large cost shocks shifted the entire distribution of price gaps, displacing many firms away from their optimal price and inducing large adjustments along the extensive margin. Fourth, we also show that this mechanism was not at work in the low-inflation pre-pandemic period. When shocks are small on average, price changes still depend on price gaps. However, as in a standard Calvo model, the frequency of price adjustment is roughly constant, the relationship between price changes and price is linear, and the passthrough rate in the microdata is roughly constant.

The second set of results pertains to accounting for inflation at the macro level. We leverage our micro data to construct an aggregate cost index for the Belgian manufacturing sector. Descriptive evidence shows how inflation fluctuations closely align with the variations in firm's production costs throughout the entire period. However, there is stickiness such that inflation moves less than costs do. 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 main driver of the surge and subsequent drop in inflation observed in the post-pandemic period.

Next, we formally assess the capacity of our menu cost model to explain aggregate inflation over the entire sample period. We feed into the model the marginal cost index described above and compare the model with the time series observed in the data. We find that the model tracks the high-frequency fluctuations in manufacturing inflation remarkably well, both during the moderate inflation regime characterizing the pre-pandemic period and during the post-pandemic inflation surge and bust. Specifically, it captures the stable behavior of the adjustment frequency pre-pandemic as well as the sharp jump in the adjustment frequency 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, can account for only about two-thirds of the inflation.

Earlier research provided strong evidence of the state-dependent nature of firm pricing decisions and showed how this class of models helps account for inflation surges after large aggregate shocks. These include work by Alvarez et al. (2016), Alvarez et al. (2022), and Midrigan (2011) on menu cost models; Alvarez

et al. (2019) and Karadi and Reiff (2019), presenting evidence of state-dependent pricing through case studies of hyperinflation in Argentina and major tax shocks in Hungary; and, more recently, the works of Blanco et al. (2024a, 2024b), Cavallo et al. (2024), and Morales-Jiménez and Stevens (2024), which 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 and under which conditions firms' pricing strategies conform with the theory.

Our study also bears relevance to the works of Eichenbaum et al. (2011) and Karadi et al. (2021). 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. We observe high frequency cost and price data for the entire Belgian manufacturing sector over a long sample period. 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 Belgian microdata for the pre-pandemic period, we used a time-dependent Calvo model to identify the structure parameters that enter the slope. This paper lends additional empirical support to that identification assumption 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 2 presents the theoretical framework and derives testable implications. Section 3 describes our dataset and the empirical measures of prices and price gaps. Section 4 provides empirical evidence showing

that the model predictions linking price adjustments with price gaps align with the microdata. We outline the calibration process and provide model simulations in Section 6, showing how the calibrated model explains the inflation time series and the frequency of price adjustments, including the rise during the pandemic. Section 7 offers concluding remarks.

2 Theoretical framework

Our baseline framework is a variation of a standard discrete-time menu cost model. To fit the data, we allow for both random menu costs as in Caballero and Engel (2007) and random free price adjustments as in the “CalvoPlus” model of Nakamura and Steinsson (2010).² For tractability, we follow Alvarez et al. (2023) working 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) as a special case.

2.1 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 lower case letters to denote the logarithm of the corresponding upper case 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:

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

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

²See also Dotsey et al. (1999) for a discussion of random menu costs models in a general equilibrium setting.

Technology. Each firm operates with a constant return to scale production technology $y_t(f) = z_t(f) + l_t(f)$, which uses a composite input $l_t(f)$ and is characterized by a total factor productivity $z_t(f)$. As is standard, we assume that the latter evolves as a random walk, $z_t(f) = z_{t-1}(f) + \zeta_t(f)$, where $\zeta_t(f)$ denotes an idiosyncratic shock that is mean zero, and i.i.d. over time and firms.

Firms' nominal marginal cost is given by:

$$mc_t(f) = mc_t + z_t(f). \quad (1)$$

The term mc_t captures an aggregate nominal cost shifter. Consistent with the evidence, we assume that mc_t obeys a random walk $mc_t = mc_{t-1} + g_t$. We assume that g_t is common across firms, i.i.d. over time, drawn from a distribution with a symmetric, uni-modal, and continuously differentiable density with mean μ_g . For analytical tractability, in what follows, we assume no trend inflation ($\mu_g = 0$). Although this assumption might seem restrictive, Nakamura et al. (2018), Alvarez et al. (2019) and Alvarez et al. (2022) show that an economy with zero inflation provides an accurate approximation for economies where inflation is low, as the effects on decision rules are of second order. We relax this assumption in the quantitative exercises of Section 6, where we allow for a small trend in marginal costs to match trend inflation in the data.

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 they adjust their 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. across firms and time, and uniformly distributed over $[0, \bar{\chi}]$. 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. In Appendix A, this implies that $p_t^o(f)$

is given by the sum of the steady-state (log) markup, μ , and a convex combination of the firm's nominal marginal cost and the price index:

$$p_t^o(f) = \mu + (1 - \Omega)mc_t(f) + \Omega(p_t + \varphi_t(f)), \quad (2)$$

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.

Following Alvarez et al. (2023), we take a quadratic approximation of the per-period profit function around the static optimum $p_t(f) = p_t^o(f)$. We define the price gap:

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

Normalizing by steady-state profits then yields the following loss function that measures the cost of deviations of the price from the target:

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

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:

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

The optimal pricing policy boils down to determining a *probability of price adjustment*, denoted by $h_t(f)$, and, conditional on adjustment, an *optimal reset gap* for x_t^* :

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

which captures the difference between the static target price and p_t^* , the (dynamic) reset price set by a firm that decides to adjust its price.³

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 chosen by the firm in the previous period:

$$\begin{aligned} x'_{t-1}(f) &\equiv p_t^o(f) - p_{t-1}(f) \\ &= x_{t-1}(f) + (1 - \Omega)(g_t + \varepsilon_t(f)) + \Omega(p_t - p_{t-1}). \end{aligned} \quad (3)$$

The second line follows from replacing $p_t^o(f)$ using Equation (2), replacing $mc_t(f)$ using Equation (1), 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 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 i.i.d. idiosyncratic shock, which combines the idiosyncratic technology and taste shocks. Similarly to the aggregate shock, g_t , we assume that the composite idiosyncratic shock is drawn from a mean-zero distribution with symmetric, uni-modal, and continuously differentiable density. 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:

$$p_t^o(f) - p_t(f) = x_t(f) = \begin{cases} x_t^* & \text{w. p. } h_t(x'_{t-1}) \\ x'_{t-1}(f) & \text{w.p. } 1 - h_t(x'_{t-1}). \end{cases} \quad (4)$$

Firms adjust their price with probability $h_t(x'_{t-1})$. Upon adjustment, they set

³Note that the optimal reset gap does not have an f subscript 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.

their price to $p_t^*(f)$. If they do not adjust their price, they keep their gap at $x'_{t-1}(f)$. Thus, as in the standard “Ss” framework, the adjustment probabilities are endogenous variables and will depend on the distance between the optimal reset gap x_t^* and the price gap $x'_{t-1}(f)$. We now characterize the optimal reset probability and the optimal reset gap.

Probability of price adjustment. Let $V_t^a(f)$ be the firm’s value if it resets its price to p_t^* and $V_t(x'_{t-1}(f))$ its value if it does not. As we show below, the former depends on $x_t^*(f)$ 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. Specifically, 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:

$$\begin{aligned} h_t(x'_{t-1}) &= (1 - \theta^o) + \theta^o \cdot \Pr(V_t^a - \chi_t \geq V_t(x'_{t-1})) \\ &= (1 - \theta^o) + \theta^o \cdot \min \left\{ \frac{V_t^a - V_t(x'_{t-1})}{\bar{\chi}}, 1 \right\}, \end{aligned} \quad (5)$$

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 things, on its ex ante price gap $x'_{t-1}(f)$. Under our assumption that if p_t^* is approximately equal to p_t^o (which, as we show later, is approximately true in the data), then $h_t(0) = (1 - \theta^o)$, the probability of a free price adjustment. Also, observe that, as the upper bound for the menu costs $\bar{\chi}$ goes to infinity, the adjustment frequency becomes exogenous because it 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^* . The value of the firm in case of no adjustment 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^* :

$$V_t^a = \max_{p_t^*} V_t(p_t^o(f) - p_t^*)$$

or, equivalently, the optimal reset gap x_t^* solves the first-order condition:

$$V'_t(x_t^*) = 0.$$

Under our assumptions of no trend inflation and a quadratic profit function, $x_t^* \approx 0$ (see, e.g., Alvarez et al. 2016). The absence of trend inflation implies that the statically optimal price provides a good approximation of the dynamic optimal price ($p_t^*(f) \approx p_t^o(f)$). If there are no strategic complementarities, the approximation is exact.⁴ 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 2 and 3 suggest, which allows us to directly test the implications of the model in the microdata. In the analytical exercises that follow, we assume that $x_t^*(f) \approx 0$. In Section 5, we verify numerically that this is indeed a good approximation.

Aggregate inflation. Next, we describe the implications of firm-level price adjustment for aggregate inflation. Given the solution of the firm's problem in Equation (4) and using the formula for the price index (Equation (2.1)), we can

⁴Intuitively, 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 price $p_t^*(f)$ of the dynamic choice problem remains very close to the static optimum $p_t^o(f)$. Under strategic complementarities, the static price moves less than the dynamic one. See Alvarez et al. (2023) for a full treatment of menu cost models with strategic complementarities.

express aggregate inflation π_t as:

$$\begin{aligned}\pi_t &= \int (p_t(f) - p_{t-1}(f)) df = \int h_t(x'_{t-1}(f)) \cdot (p_t^*(f) - p_{t-1}(f)) df \\ &= \int h_t(x'_{t-1}(f)) df \cdot \int (p_t^*(f) - p_{t-1}(f)) df + \text{Cov}(h_t(x'_{t-1}), (p_t^* - p_{t-1}))\end{aligned}\quad (6)$$

$$\approx \int h_t(x'_{t-1}(f)) df \cdot \int (x'_{t-1}(f)) df + \text{Cov}(h_t(x'_{t-1}), x'_{t-1}).\quad (7)$$

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 adjustment and the average price adjustment of adjusters, and (ii) the covariance between the variables. The last line uses the assumption that $p_t^*(f) \approx p_t^o(f)$ to express aggregate inflation as a function of moments of the distribution of the ex ante price gaps (x'_{t-1}).


The first term in Equation (6) states that inflation depends on both the average price gap and the average frequency of price adjustment. With state-dependent pricing, the adjustment probability is an endogenous object that, as we will see, increases non-linearly with the absolute value of the price gap. With Calvo pricing, the adjustment probability is fixed and constant across firms. 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) and, more recently, Karadi et al. (2021), a “selection effect” increases the degree of monetary neutrality in the economy. This is captured by the covariance term in Equation (6). Firms that are more likely to adjust are also those that change their prices the most (conditional on adjustment). That is, the gap between $p_t^*(f)$ 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.

2.2 Discussion and testable implications

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

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.

In Figure 2, the horizontal axis is the density function of the ex ante price gap, denoted by $f(x'_{t-1})$, which is uni-modal 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 $p_t^o(f)$, operate with a sub-optimally low markup and therefore are more likely to increase (decrease) their price relative to the price they changed in the previous period. The greater concentration of price gaps near the optimum simply reflects that a firm is more likely to adjust the further its price gap is from the optimum.

Under our quadratic approximation of the profit function and with low trend inflation, the GHF can be approximated, up to a second order, by a quadratic

function of the price gap centered around $x'_{t-1} = 0$.⁵ Specifically, we have that:

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

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 $\phi \equiv -\frac{\theta_0}{\bar{\chi}} \frac{\partial^2 V_t(x'_{t-1})}{\partial (x'_{t-1})^2}$ controls the sensitivity of the GHF to changes in gaps (i.e., the “steepness” of the parabola).

Nonlinear price dynamics along the price gap distribution. The endogenous nature of the probability of price adjustment is the key driver of the nonlinear transmission of shocks into prices in state-dependent models is. To illustrate this point, we divide the distribution of ex ante price gaps into equally sized bins, each defined narrowly enough to ensure that the price gap remains nearly constant within each bin. Starting from the formula for inflation in Equation (7) and applying the quadratic approximation from Equation (8) to substitute the hazard function, we derive the following expression, which characterizes inflation within a bin b as a function of the (odd) moments of the price gap distribution:⁶

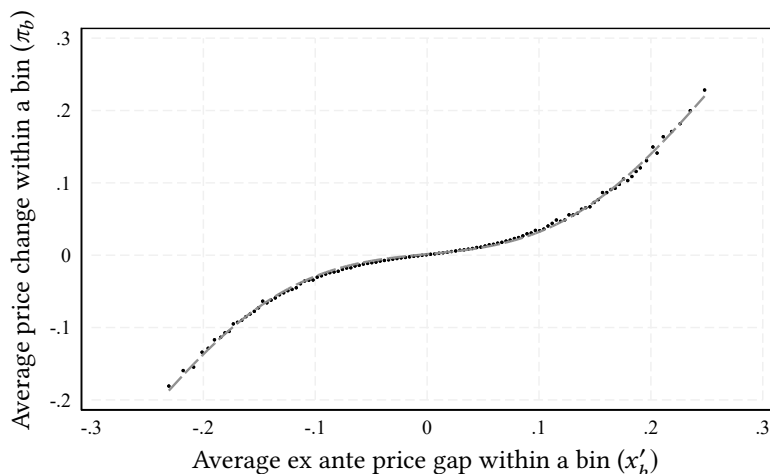
$$\pi_b \approx \phi_b^0 \cdot (x'_b) + \phi \cdot (x'_b)^3, \quad (9)$$

where $\pi_b := \int_{f \in b} (p_t(f) - p_{t-1}(f)) df$ and $x'_b := \int_{f \in b} x'_{t-1}(f) df$ 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 free-adjustment probability ($1 - \theta^o$), which is common across bins, and a term equal to the variance of the price gaps within bin b scaled by the GHF slope parameter ($\phi \sigma_b^2$). The latter captures the effect of deviations of the price gap from zero on the adjustment frequency. It is straightforward to derive the analog

⁵In Appendix A for the derivation of this result and Alvarez et al. (2022) for a generalization to asymmetric GHFs.

⁶See Appendix A for the analytical derivations.

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 dashed line represents the fitted values from a regression of bin-level inflation on a polynomial in the first and third orders of the average gap, as specified by Equation (9). The weight of each bin in the regression is proportional to the number of observations within it.

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 $\pi_b = (1 - \theta^c) \cdot (x'_b)$.⁷

The binned scatter plot in Figure 3 illustrates the relationship between price gaps and inflation across the distribution of gaps. The gray dotted line represents the fitted values from a regression of bin-level inflation on a polynomial in the first and third orders of the average gap, as specified by Equation (9).

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 result between the Calvo and Menu cost models with small shocks illustrated in Gertler and Leahy (2008), Alvarez et al.


⁷In a Calvo model, $\mathbb{E}[p_t(f)|\mathcal{I}_t(f)] = \mathbb{E}[p_t(f)|p_t^*(f), p_{t-1}(f)] = (1 - \theta^c)p_t^*(f) + \theta^c p_{t-1}(f)$, where $\mathcal{I}_t(f)$ denotes the information set of a firm entering period t . Using the approximation $p_t^*(f) \approx p_t^o(f)$ and re-arranging we obtain the equation in the text.

(2017), and Auclert et al. (2024).

The cubic term opens the door for nonlinear inflation dynamics. Firms at the tails of the distribution of price gaps are more likely to adjust their prices (the frequency effect captured by the GHF), and tend to do so more aggressively (the selection effect discussed in Golosov and Lucas (2007)).

The impact of aggregate cost shocks. Aggregate cost shocks—i.e., shocks that do not average out—affect the optimal reset prices of all firms in the economy, shifting the entire distribution of price gaps. When these shocks are large, 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: Large versus small aggregate cost shocks



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

Figure 4 illustrates this point. In the spirit of the exercise in Cavallo et al. (2024), we shock the economy in its stationary equilibrium with an unexpected cost shock $g_t > 0$, which increases the marginal cost for all firms. The left panel shows the effect of a small shock, while the right panel shows the effect of a large shock. The solid lines represent the GHF 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 which has little or no impact on the average frequency of price adjustment. This observations is in line with Gertler and Leahy (2008) and Auclert et al. (2024), showing how the price dynamics generated by a state-dependent model are observationally similarly to those generated by a time-dependent model when the economy is hit by small shocks. By 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 a lower fraction that wants to reduce them. On average, this causes a non-trivial increase in the adjustment frequency, which amplifies aggregate inflation beyond what is accounted by the increase in the average gap.

3 Data and measurement

The backbone of our empirical analysis is the dataset assembled by Gagliardone et al. (2024). Constructed by integrating different administrative micro-datasets, this dataset encompasses market interactions between domestic and international competitors across manufacturing industries in Belgium. It contains information about firms' production decisions and a detailed snapshot of firms' variable production costs (labor costs and intermediates) at a business cycle (quarterly) frequency.⁸

We enhance this data set in two significant ways. 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.

⁸We refer to Gagliardone et al. (2024) for details about the data sources and variable definitions.

3.1 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 includes 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 sectors (2-digit codes).

Price indexes. 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 the changes in the domestic prices across the different products of firm f in industry i :

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

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 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):

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

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).

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

Frequency of Price Adjustment. The frequency of price adjustments is a crucial variable to characterize the nature of nominal rigidity in the data. 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 adjustment for the manufacturing sector. Mirroring the notation in the model, we denote this variable by \bar{h}_t .

This variable also helps us to clean for spurious price changes in the micro data. Our measure of prices is based on product-level unit values. Due to small measurement error, this measure tends to overstate the frequency of small price changes, as shown by Eichenbaum et al. (2014).⁹ To address this measurement problem, we combine the firm-level price changes with information on the frequency of price adjustment 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:

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

⁹For 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.

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

To account for different degrees of upward and downward nominal rigidity in the data, we set the thresholds $\kappa^+ = 0.75$ and $\kappa^- = 0.87$ to separately match the average frequency of upward and downward price changes measured using the NBB-BS micro data: $\sum_t \sum_f \mathbb{I}_{ft}^+ = \bar{h}^+$ and $\sum_t \sum_f \mathbb{I}_{ft}^- = \bar{h}^-$, where $\bar{h}^+ + \bar{h}^- = \bar{h}$.¹⁰

Marginal cost indexes. 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: $MC_{ft}^n = (1 + v_f)AVC_{ft}$. The coefficient v_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 ($v_f \equiv 1/RS_f - 1$). Using the definition of average variable costs (total variable costs over output, TVC_{ft}^n/Y_{ft}) and applying a logarithmic transformation, we have that firm-level log-nominal marginal cost is given by:

$$mc_{ft}^n = (tvc_{ft}^n - y_{ft}) + \ln(1 + v_f) \quad (12)$$

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. We compute a quantity index by dividing a firm's domestic revenues by its domestic price index.¹¹ Firm-specific short-run returns to scale are not directly observable in the data. Therefore, to the extent that individual firms' production technologies deviate from constant returns to scale ($v_f \neq 0$), 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.

¹⁰See Karadi et al. (2021) and Luo and Villar (2021) for evidence of asymmetric upward and downward rigidity.

¹¹Specifically, 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 indexes 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.

Ex-ante price gaps. The availability of high-frequency data on firm-level prices, marginal costs, and competitors' price indexes enables us to construct an empirical counterpart of the firm-level ex ante price gaps defined in Equation (3):

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

Guided by our theoretical framework, we construct a proxy of firms' target prices as a convex combination of the firm's own marginal cost and its competitors' price index: $p^o_{ft} = (1 - \Omega)mc^n_{ft} + \Omega p_t^{-f}$. We calibrate Ω to 0.5 to match the micro estimate in Gagliardone et al. (2024). As we discussed in Section 2, when p^o_{ft} and p^*_{ft} are sufficiently close to each other, $x'_{t-1}(f)$ provides information on inefficiencies driven by nominal rigidities. A positive ex ante price gap indicates that a firm is operating with a markup below the profit-maximizing one and, therefore, absent nominal rigidities, would adjust its price upward.

3.2 Harmonization and data cleaning

We apply the following data cleaning steps and harmonization procedures to the empirical distributions of price changes and price gaps. To mitigate the noise in price changes due to the use of unit values, we set to zero the price changes that are less than 1 percent in absolute value. Note that this adjustment does not affect our measure of the average frequency of price adjustments, which is computed precisely using the NBB-BS data on price adjustments, as discussed above. To remove outliers, we trim observations at the top and bottom 2 percent of the price changes and the price gaps distribution.

Next, we need to address some differences between our empirical measure of price gaps and its theoretical counterpart in Equation (3). First, our measure of p^o_{ft} does not account for the realization of unobservable idiosyncratic taste shocks (φ_{ft}). Although such shocks average out, they also introduce measurement error that weakens the connection between price gaps and price changes, at the individual firm's level. Second, given our measures of prices and marginal costs, we can only identify firm's reset gaps up to an additive constant, which captures

a combination of unobserved steady-state markups (the constant μ in Equation (3)), unobserved deviations from constant short-run returns to scale affecting marginal costs (the term $\ln(1+v_f)$ in Equation (12)), and the normalization of price levels. Third, for analytical tractability, we developed a one-sector model assuming zero trend inflation. In reality, inflation exhibits a small trend (approximately 0.6 percent quarter-on-quarter, before the 2021 surge), which varies between industries. We also observe strong seasonal patterns in nominal variable costs, which are higher in the second and fourth quarters, on average. To address these differences and align the empirical and theoretical price gap measures, we harmonize the price gaps by removing firm-specific and industry-specific calendar quarter averages. For consistency, we apply the same harmonization to the distribution of price changes. By doing so, we account for firm-specific intercepts, industry-specific seasonal patterns in nominal costs, and trend inflation. It also mitigates measurement errors in price changes and reset gaps caused by the use of unit values for price measurements and inaccuracies in measuring marginal costs.

3.3 The joint distribution of price changes and price 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 frequency of price changes by 10 percentage points, on average.

Table 1: Summary statistics of price changes and price gaps

Price change ($p_{ft} - p_{ft-1}$)				Ex ante price gap (x'_{ft-1})		
<i>Panel a: Time period 2000-2019</i>						
Mean	Std.	Freq. Adj.	Kurt.	Mean	Std.	Kurt.
-0.00	0.11	0.29	5.79	-0.00	0.14	2.96
<i>Panel b: Time period 2020-2023</i>						
Mean	Std.	Freq. Adj.	Kurt.	Mean	Std.	Kurt.
0.01	0.12	0.38	5.00	0.01	0.15	2.81
Number of observations:			133,401			
Number of firm-industry pairs:			5,348			
Number of firms:			4,811			


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).

The fourth column reports the kurtosis of price changes. We calculate this statistic following the approach in Klenow and Kryvtsov (2008), which involves standardizing the distribution of price changes by removing firm-level means and scaling by firm-level standard deviations. The estimated kurtosis is on the high end of the distribution of estimates found in the literature. This is likely due to the high sensitivity of this statistic to measurement error and to unobserved heterogeneity, which we are not able to fully account for.¹²

The last three columns of Table 1 present summary statistics of the price

¹²To this point, Alvarez et al. (2016) shows that the empirical estimates of the kurtosis can be biased upward by 30 percent or more if measurement error and heterogeneity are not appropriately filtered.

Figure 5: Empirical distribution of ex ante price gaps



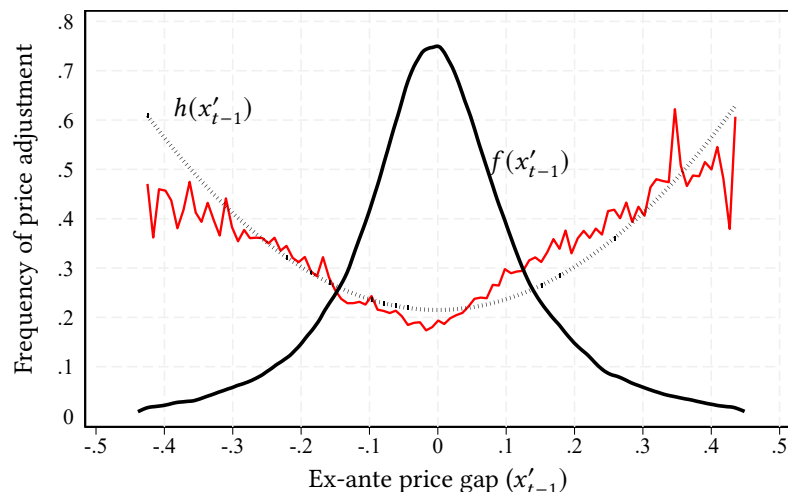
Notes. The figure presents the empirical probability density function of the price gaps, $f(x'_{t-1})$, in the pre-pandemic period (2000-2019).

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. The data 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.

4 Micro evidence of state dependent pricing

Guided by the theoretical results presented in Section 2, we design direct empirical tests of key model predictions that relate the micro-level pricing dynamics to the underlying price gaps distribution in both low- and high-inflation regimes. These exercises provide strong evidence of the state-dependent nature of firm pricing decisions.

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 adjustment of a given bin (b) on a constant and the square of the average price gap of the same bin, as dictated by Equation (8). In the regression, we weight each bin by the number of observations it counts.

4.1 The empirical Generalized Hazard Function

The relationship between price gaps and the frequency of price adjustments is the watershed between models featuring state- and time-dependent pricing. State-dependent models imply a monotonically increasing relationship between a firm's probability of price adjustments (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 in the microdata using information on the frequency of price adjustment and price gaps. We focus on the pre-pandemic period (2009–2019). Throughout the paper, we treat this period as a representation of the economy in steady state, when aggregate shocks are zero and the only variation in gaps is driven by the realization of idiosyncratic cost and markup shocks. In Figure 6, the black line represents the probability density function of price gaps. The red line is the empirical analog of the theoretical GHF, measuring

the fraction of firms that adjust their prices for each bin of the distribution of price gaps.

The data reveal a strong connection between the size and magnitude of price gaps and the frequency of price adjustment. A greater deviation from $x'_{t-1} = 0$ induces a larger fraction of firms to adjust their prices, with a striking resemblance 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 2, 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 adjustment for each bin, h_b , on a constant term and the square of the average price gap for the same bin, x'_b :


$$h_b = a_0 + a_1 \cdot (x'_b)^2 + v_b,$$

where v_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.

4.2 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, $\hat{\pi}(b) = \hat{a}_0 + \hat{a}_1 \cdot x'(b)$, 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 1st and 3rd 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 counts.

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):

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

The coefficient b_1 captures the average price change associate with a small increase in the price gap. In Appendix A we show that \widehat{b}_1 converges in probability to the average frequency of price adjustment in the regression sample and that \widehat{b}_2 converges in probability to ϕ , the curvature parameter of the GHF.

By comparing the patterns in Figure 3 to 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 of 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, which implies that the relationship between inflation and price gaps is approximately linear, as in Calvo model. This result echoes those in 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 resemble those generated by a time-dependent model, up to first-order. The linearity of the relationship between price gaps and price changes is also at the core of the identification argument in Gagliardone et al. (2024), which used microdata to identify the slope of the cost-based New Keynesian Philips curve in a low inflation environment. To this point, the estimate of the slope coefficient is $\widehat{b}_0 = 0.28$ (or 0.29, if we ignore the cubic term), which almost exactly matches the frequency of price adjustments observed in our sample in low inflation environment (Table 1, panel a).

Now consider the relationship between price gaps and price changes across the entire price gap distribution, including its tails. The red dashed red line in Figure 7 represents the projection of inflation on a the price gap and the cubic of the price gap of each bin. The micro data offer the opportunity to appreciate the non-linearities 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 connection between movements in gaps and movements in prices 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.

4.3 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 don't. We now focus our attention on the former group of firms. The theory predicts that, conditional on adjusting, the firms set $p_{ft} = p_{ft}^*$. This implies that $p_{ft} - p_{ft-1} = x_{ft}^*$, when $p_{ft} \neq p_{ft-1}$. Although we can't measure x_{ft}^* in the data, to the extent that p_{ft}^o provides a reasonable approximation for p_{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 indeed 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_{ft}^* with p_{ft}^o 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 movements in costs when inflation is high, as shown by Gagliardone and Tielens (2024) using a model with state-dependent information frictions.

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



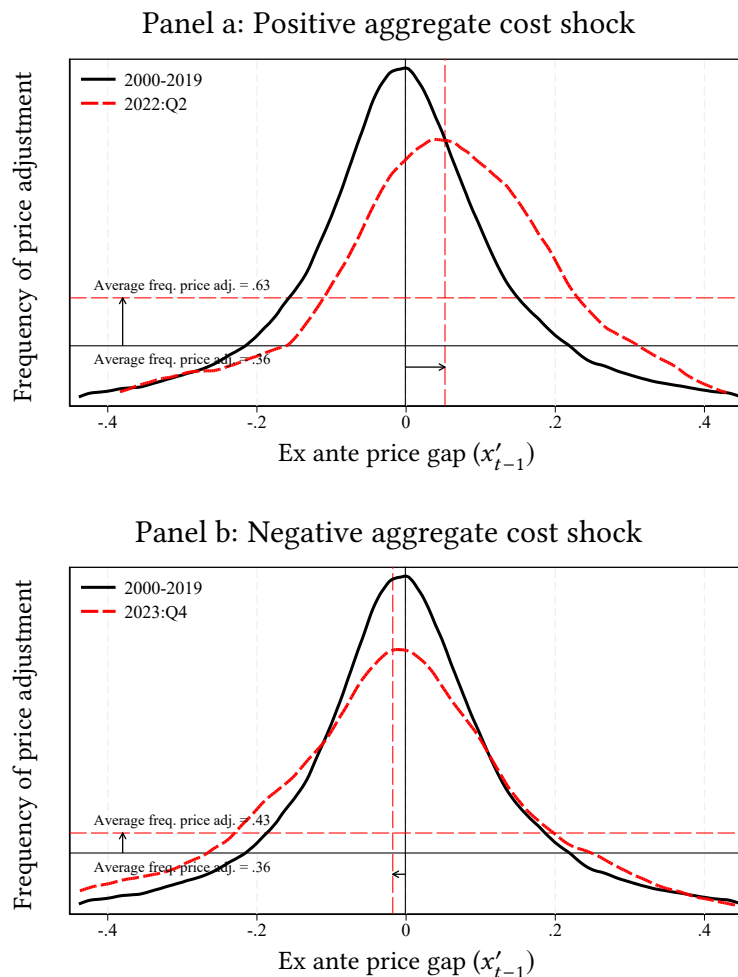
Notes. This figure presents a binned scatter plot of the log price change for adjusters (i.e., firms for which $p_{ft} \neq p_{ft-1}$) against the ex ante price gap. Each dot marks the average price gap of a given percentile of the price gap distribution (x-axis) and the corresponding average percentage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

4.4 Large cost shocks and shifts in the price gaps distribution

We discussed in Section 2 how small, idiosyncratic shocks generate dispersion in the price gap distribution, while large aggregate shocks shift the entire distribution of price gaps, significantly increasing the fraction of firms who want to adjust their prices (Figure 4). The drastic surge and subsequent normalization of production costs observed in the post-pandemic period allows us to directly test this model's prediction in the microdata.

In Figure 9, the black solid 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

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



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 adjustment in the pre-pandemic period (black solid line) and in 2022:Q2 and 2023:Q4 (red dashed lines).

now further away from their desired levels, resulting in a shift of the distribution and fatter tails. Because the shock reduces firms' profit margins, the cost of not adjusting is larger. Firms adjust their prices upward as more firms move to regions of the price gap distribution where the GHF is high. 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 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 adjustment as firms began to lower their prices.

This exercise also reveals that the shocks not only shifted the distribution of price gaps but also altered its shape, thickening the tails of the distribution. This observation suggests that, in reality, the aggregate shock due to the hikes of intermediate costs prices affected firms in a heterogeneous way. This heterogeneity is accounted for by our theoretical framework and it is an interesting topic for future research.

5 Aggregate cost-price dynamics

The analysis of the microdata highlighted the state-dependent nature of firms' price adjustments. We now shift our focus to the aggregate level. We show that the micro-level dynamics translate into nonlinear inflation dynamics, with the degree passthrough of aggregate cost shocks into prices varying with the magnitude of the shock.

5.1 Aggregate inflation and aggregate costs

We use our micro data 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 $\bar{s}_{ft} \equiv \frac{s_{ft} + s_{ft-1}}{2}$:

$$\pi_t = \sum_{f \in \mathcal{F}} \bar{s}_{ft} \cdot \Delta p_{ft} - 1.$$

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 we normalized to zero the value of the index the first quarter of our data ($mc_{1999:Q1}^n = 0$).

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) and of scaled nominal marginal costs (black line) throughout our sample period. Keep in mind that the scale of the two axes differs for the two variables.

As the theory predicts, inflation closely tracks the fluctuations of scaled marginal cost over the whole sample. But, also consistent with theory, there is stickiness such that inflation moves less than costs do. At the same time, underlying the significant surge and subsequent normalization of inflation in the post-pandemic period was 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,

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.

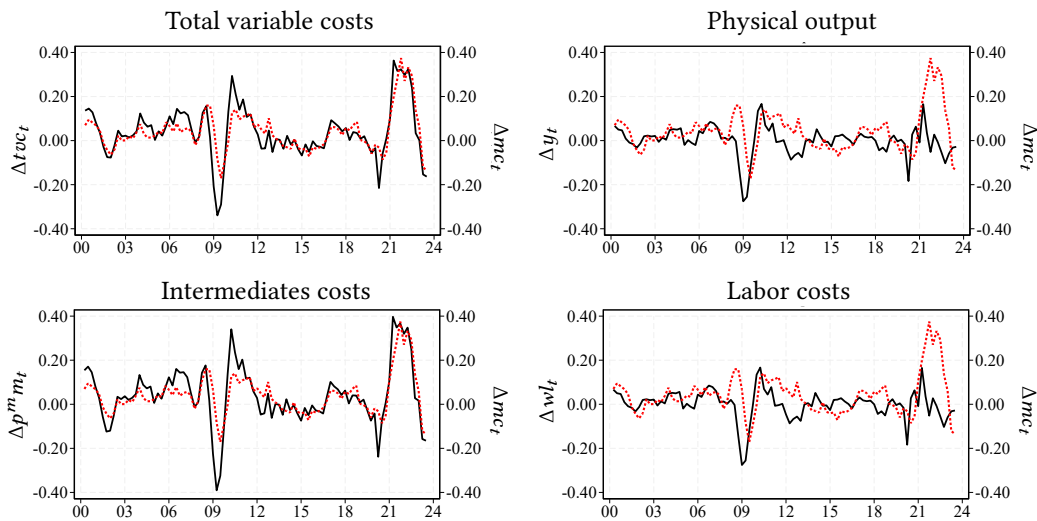
the hypothesis that a rise in markups can explain the recent inflation surge seems to have no bite in the data.¹³

Finally, to get a sense of what may drive 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

¹³Studying the price and cost data for a large global manufacturer, Alvarez et al. (2024b) also find that markups remained stable markups over time over time, including during the inflation surge.

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 the time-series evolution of nominal marginal cost.

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 (Δwl_t , bottom left panel).

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—is the main driver of the inflation surge between 2021 and 2023, at least in our sample.

5.2 Nonlinear cost-price passthrough in the macro data

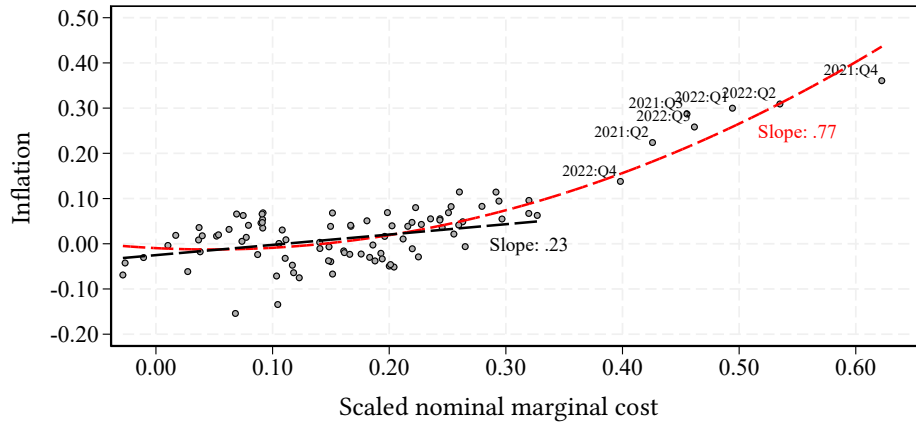
In panel a of Figure 12 we sort quarters by their measured scaled marginal cost index and plot this index against PPI inflation. The black dashed line represents the linear fit between the two variables during periods of low inflation (below 10% year-over-year), while the red dashed line represents a quadratic fit across both high- and low-inflation periods. The slope of these curves provide descriptive evidence on the aggregate passthrough of cost shocks into prices as a function of the size of the shock to marginal cost.

Consistent with the micro-level dynamics presented in Section 4, we find a linear 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. Note that this figure aligns with aggregate passthrough coefficient in Gagliardone et al. (2024) in a low inflation environment. The linear relationship between the two variables breaks down when shocks are large. In fact, the passthrough coefficient more than tripled during the recent inflation surge, revealing highly nonlinear cost-price dynamics.

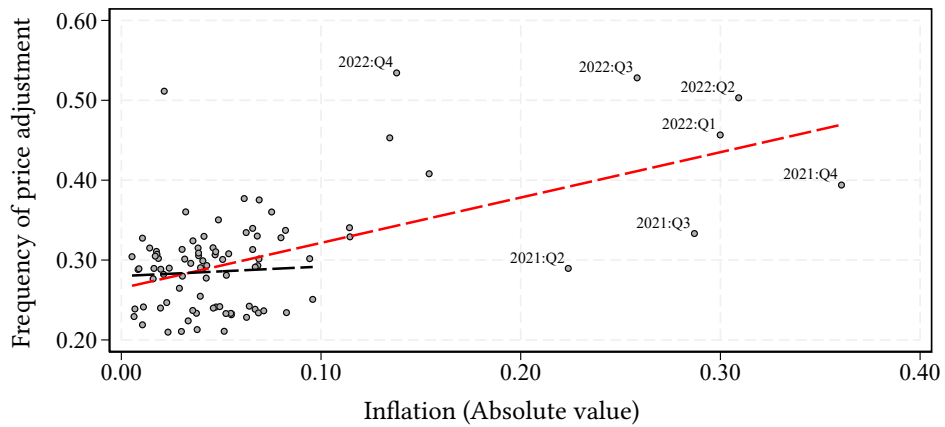
At the core of this result is the endogenous nature of the frequency of price adjustment. In Panel b, we sort quarters by their annual inflation rates and plot aggregate inflation against the average frequency of price adjustments. As before, the black dashed line shows the linear fit during periods 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 assumption of a Calvo model and it suggests that a time-dependent pricing model provides a good framework to capture nominal rigidities in a low-inflation environment. However, inflation and frequency of price adjustment 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 adjustment (\bar{h}_t) in the same quarter. The average frequency of price adjustment 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 between 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.

6 Quantitative implications

Having established the close connection between theory and data, we now use moments from the microdata to calibrate and simulate the quantitative model presented in Section 2. We use the calibrated model to perform two types of quantitative exercises. In a first, more standard, set of exercises, we contrast 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.

6.1 Calibration

We have a total of seven parameters to calibrate. We calibrate four of them to standard values in the literature. We calibrate the elasticity of substitution between goods σ , to 6, which implies a markup of 20 percent in the symmetric steady state equilibrium. We set β , the firm's risk neutral discount factor, at 0.99. As in our empirical analysis, we calibrate $\Omega = 0.5$ to reflect the importance of strategic complementarities estimated in Gagliardone et al. (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 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-dependency of price adjustments. The standard approach to calibrate these parameters leverages the theoretical mapping between the unobservable distribution of price gaps and the observable distribution of price changes, targeting standard deviation, kurtosis, and frequency of price changes (see, e.g., Alvarez et al. (2022), and Blanco et al. (2024a)). In theory, we could use microdata on price changes to recover these moments. In practice, producing unbiased empirical measures of these moments can be challenging, especially in our context. The measurement of kurtosis is particularly problematic, as this moment tends to be sensitive to small measurement error (Alvarez et al. 2016) and

unobserved heterogeneity (Alvarez et al. 2022), which can mechanically generate upward bias in the measured statistic (see, e.g., Cavallo and Rigobon 2016).

To circumvent 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 price gaps and in the empirical GHF during the pre-pandemic period.

First, we have shown strong evidence in favor of a V-shaped GHF (see Figure 6), as captured by Equation (8). We therefore calibrate the free price adjustment parameter θ^o to match the frequency of price adjustment in a neighborhood of $x'_{ft-1} \approx 0$.¹⁴ This yields an estimate of the free adjustment probability of $(1 - \theta^o) = 0.20$. Note that by averaging between thousands of observations in a neighborhood of $x'_{ft-1} \approx 0$, this calibration tends to be 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 adjustment (\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 $\epsilon_t(f)$ 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 adjustment and the variance of price changes reported in panel a of Table 1.

Finally, given σ_ϵ^2 and θ^o , we calibrate $\bar{\chi}$ (the upper limit of the uniform distribution from which the random menu costs are drawn) to 0.6, to allow the model to match the frequency of price changes in the pre-pandemic period.

¹⁴We estimate the empirical analog of Equation (8), $h_b = \widehat{1 - \theta^o} + \widehat{\phi} \cdot (x'_b)^2 + \epsilon_b$, where h_b and x'_b denote the within-bin average frequency of price adjustment and the average price gap. To obtain more precise estimates of this parameter, we restrict the estimation sample to bins capturing observations in the 25 to 75 percentiles of the gaps distribution and assign each bin a regression weight equal to the share of in each observations.

We conclude with two observations lending empirical support to our calibration procedure. First, in a recent paper, Blanco et al. (2024a) show how a standard menu costs 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 the 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. Caveat the measurement issues discussed above, we calibrated model displays a kurtosis of price changes that is about 25 percent lower than the one computed in the data. This figure is consistent with the results in Alvarez et al. (2024a), which show that the estimated kurtosis of price adjustment of French CPI data shrinks by 30 percent once unobserved heterogeneity is controlled for.

Table 2: Calibration: Data vs model

	Price change ($p_{f_t} - p_{f_{t-1}}$)				Ex ante price gap ($x'_{f_{t-1}}$)			Share MC Mean (%)
	Mean	Std	Kurt	Freq. Adj.	Mean	Std	Kurt.	
Data	-0.00	0.11	3.23	0.29	-0.00	0.14	4.14	1.22
Menu cost	0.00	0.12	2.62	0.29	0.00	0.09	3.30	1.70
Calvo	0.00	0.12	5.21	0.29	0.00	0.12	5.21	

Notes. This table reports moments of the distribution of price changes and price gaps computed during the period 2000–2020 (panel a) and the corresponding moments for the menu cost model and Calvo model, in steady-state, under our baseline calibration. As we do in the data, the model-based moments are computed after de-meaning the distribution of price changes and ex ante price gaps. The last column reports the average share of menu costs paid by firms as a fraction of firms' revenues. The estimate of this moment in the data comes from Zbaracki et al. (2004).

Table 2 compares the empirical moments of the price changes (panel a) and 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. Notably, although we did not directly target moments of the distribution ex ante price gaps, our simulated model displays a dispersion and a degree of leptokurtosis that closely aligns with the empirical one.


We also consider a standard Calvo model, calibrated to match the steady-state frequency of price adjustment observed in the data. As explained in Section 2, our menu costs model nests the Calvo as a special case when the maximum menu costs, $\bar{\chi}$, go to infinity and the probability of free price adjustment, $1 - \theta^o$, is re-calibrated to match the steady-state frequency of price adjustment.

6.2 Impulse-responses to small and large aggregate shocks

We use the calibrated menu costs 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 steady, we shock the system with unanticipated, permanent shocks to aggregate marginal cost of different magnitudes, $g_t = \{2\%, 10\%, 20\%\}$.

Figure 13 displays the impulse response function of the frequency of price adjustment (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 triggering an increase in the number of firms adjusting their prices and, therefore, an inflation. However, as discussed in Section 2 and empirically shown in Section 4, 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 adjustment and, consequently, a rapid and substantial surge in inflation. These graphs highlight the non-linearities 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 adjustment and inflation is about three times larger than the effect of the medium shock, although the former is only twice as large as the latter (10% vs. 20%). To highlight these features, it is useful to contrast the IRFs of the menu costs model to those obtained from the Calvo model (Figure 13, panel b). By construction, in the Calvo the number of firms adjusting their prices is not affected by the magnitude of the shock (the GHF is a 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

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

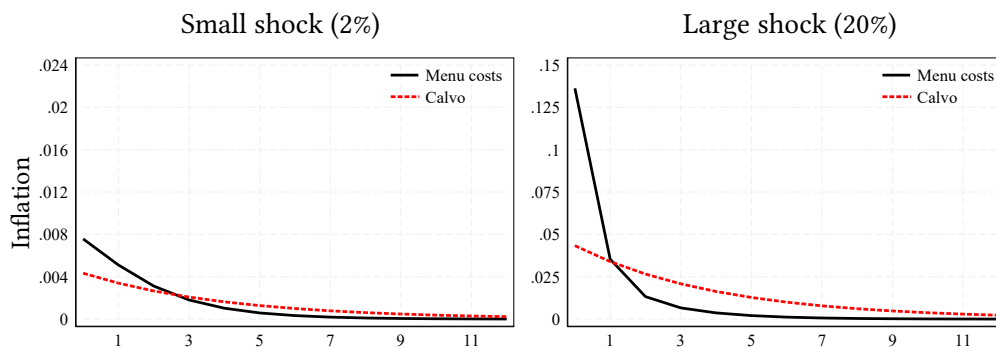


Notes. This figure presents the impulse responses of inflation and frequency to aggregate cost shocks of different magnitudes. Panel a reports the impulse response for our state-dependent pricing model (menu costs model). Panel b reports the impulse responses for a time-dependent model (Calvo model).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

shock, but in a proportional way.

The second observation regards the speed at which the permanent cost shocks are fully incorporated into prices. Figure 13 highlights how large shocks induce firms to react faster than small shocks do. The interplay of the endogenous change in the frequency of adjustment and the selection effect (the observed price changes come from firms who need it the most) translates into a faster passthrough from costs into prices. The passthrough is notably slower in the Calvo model, especially in response to large shocks. Figure 14 helps visualize this result, overlaying the IRFs of inflation in the menu costs and Calvo models in response to the same size 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.

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_{ft}^o and the dynamic optimal price p_{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 particularly close in the context of the menu cost model relative to the Calvo model. These results are important because they suggest that the assumption that $p_{ft}^o \approx p_{ft}^*$ needed to derive the expressions for aggregate inflation and within-bin inflation as a function of ex ante gaps (Equations (7) and (9), respectively) is sensible.

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, the strategic complementarities lead to a reduction in the cost pass-through in both the menu costs and the Calvo model. The greater curvature of the value function under state-dependent pricing implies that the difference between the impulse-response functions with and without complementarities is narrower in the menu cost model, especially in response to a large shock.

6.3 Explaining the time series of inflation


Having characterized the price dynamics in response to shocks of different magnitude, 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 frequency of price adjustment. We conduct the following quantitative exercise for the state-dependent menu costs 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 price gap as the new model's equilibrium, we repeat this feeding exercise for all subsequent quarters until 2023Q4, 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. Panel a and b show that the menu costs 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 of only about

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



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

two-thirds of that generated by the menu cost model. This exercise also highlights the more sluggish behavior of inflation produced by the Calvo model relative to that generated by the menu cost model. This is consistent with the faster cost pass-through generated by the state-dependent pricing policies documented in the impulse response function of Section 6.2.

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

7 Concluding Remarks

We have developed a state-dependent pricing model designed to provide an accounting of aggregate price dynamics across both high and low inflation regimes. The model explains both the low stable inflation of the pre-pandemic period and the pandemic era surge. It also captures the associated changes in the price adjustment frequency.

Unlike previous studies, we leverage detailed information on prices and costs to construct a direct measure of firm's price gaps. Studying the joint variation of prices and price gaps, we show how firms' behavior is consistent with the state-dependent framework. At the micro level, variation in price gaps determines both the likelihood that a firm adjusts its price and how much its price changes conditional on adjustment.

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, cost-price dynamics were highly nonlinear. The sharp increase in marginal cost led to not only a jump in price gaps but also a significant increase in adjustment probabilities. This extensive margin of price adjustment is the hallmark of state-dependent pricing models, but it 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 main variation in marginal cost during the inflation surge was due to sharp increases in the cost of intermediate inputs. Extending a state-dependent version of the New Keynesian model to allow for 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.
- Fernando Alvarez, Andrea Ferrara, Erwan Gautier, Hervé Le Bihan, and Francesco Lippi. Empirical investigation of a sufficient statistic for monetary shocks. *Review of Economic Studies*, page rdae082, 2024a.
- Santiago Alvarez, Alberto Cavallo, Alexander MacKay, and Paolo Mengano. Markups and cost pass-through along the supply chain. *Unpublished manuscript, Harvard Business School*, 2024b.
- 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.
- 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.
- 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, 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.
- Jordi Galí. *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press, 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. Measuring price selection in microdata: It’s not there. 2021.
- 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.
- 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 during inflation surges versus normal times

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

Appendix

A Derivations

A.1 Derivation of markup function

Assume that a perfectly competitive retailer assembles a bundle of intermediate inputs into a final product, Y_t . The bundle is 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:

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

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

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

Define implicitly the industry price index P_t as:

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

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

function:

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

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

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

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

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

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:

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

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:

$$\mu_t(f) - \mu_f = \frac{\epsilon'}{\epsilon(\epsilon - 1)} (y_t(f) - y_t)$$

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

$$\mu_t(f) - \mu_f = -\Gamma (\tilde{p}_t^o(f) - p_t)$$

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

minimization):

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

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

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 $h_t(x'_{t-1})$ in characterized in Equation (5) around x_t^* to obtain:

$$\begin{aligned} h_t(x'_{t-1}) &\approx (1 - \theta^0) - \frac{\theta_0}{\bar{\chi}} \frac{\partial V_t(x_t^*)}{\partial(x'_{t-1})} (x'_{t-1} - x_t^*) - \frac{\theta_0}{\bar{\chi}} \frac{\partial^2 V_t(x_t^*)}{\partial(x'_{t-1})^2} (x'_{t-1} - x_t^*)^2 \\ &= (1 - \theta^0) - \frac{\theta_0}{\bar{\chi}} \frac{\partial^2 V_t(x_t^*)}{\partial(x'_{t-1})^2} (x'_{t-1})^2, \end{aligned}$$

where the second equation follows from $\frac{\partial V_t(x_t^*)}{\partial(x'_{t-1})} = 0$ for a firm that is resetting its price and from our assumption that $x_t^* \approx 0$. Defining $\phi \equiv -\frac{\theta_0}{\bar{\chi}} \frac{\partial^2 V_t(x_t^*)}{\partial(x'_{t-1})^2}$ 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):

$$h_t(x'_{t-1}(f)) \approx (1 - \theta^0) + \phi \cdot \left(x'_{t-1}(f)\right)^2.$$

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

Averaging across firms we have that the average frequency of price adjustment is given by:

$$\bar{h}_t(x'_{t-1}) \equiv \int_{[0,1]} h_t(x'_{t-1}(f)) \approx (1 - \theta^0) + \phi \cdot \int_{[0,1]} \left(x'_{t-1}(f)\right)^2 df \quad (\text{A.5})$$

To take Equation (A.5) to the data, we partition the support of the distribution ex ante price gap into equally spaced bins. Assume that the first and second moments of the distribution of $x'_t(f)$ within each bin exist and denote them by

$x'_b \equiv \int_{f \in b} x'_{bt-1}(f) df$ and $\sigma_b^2 \equiv \int_{f \in b} (x'_{t-1}(f) - x'_b)^2 df$.

Adding a white noise disturbance, v_b , we obtain the cross-sectional regression model:

$$h_b(x'_b) = a_1 + a_2 \cdot (x'_b)^2 + v_b. \quad (\text{A.6})$$

Estimating model (A.6) via weighted least squares (weighting observations by the number of observations within each bin) allows us to calibrate the free-adjustment parameter, $\theta^0 = \hat{a}_1 - 1$, and the steepness parameter, $\phi = \hat{a}_2$. Plugging these estimates into Equation (A.5) and using the information on the first and second moments of the price gaps across bins, we can characterize the empirical GHF, as shown by the dotted line in Equation (6).

Cubic approximation of inflation within a bin

We partition the distribution of price gaps into bins denoted by b . As before, we denote by x'_b and σ_b^2 the first and second moments of the ex ante price gap distribution within a bin.

Using Equation (A.5) and the formula of the variance, we have that the average frequency of price adjustment for firms in bin b is given by:

$$h_b(x'_b) = \int_{f \in b} h_t(x'_{t-1}(f)) df \approx (1 - \theta^0) + \phi \left((x'_b)^2 + \sigma_b^2 \right). \quad (\text{A.7})$$

Next, consider the expression for aggregate inflation under the assumption that $p_t^*(f) \approx p_t^o(f)$ in Equation (7). We choose these bins to be sufficiently narrow such that within each bin the covariance between the GHF is approximately zero: $\int_{f \in b} h_t(x'_{t-1}(f)) \cdot x'_{t-1}(f) df \approx 0$. When this condition is satisfied, we have that inflation within a bin is given by:

$$\pi_b \approx \int_{f \in b} h_t(x'_{t-1}(f)) df \cdot \int_{f \in b} (x'_{t-1}(f)) df = h_b(x'_b) \cdot x'_b.$$

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

in the paper:

$$\begin{aligned}\pi_b &\approx (1 - \theta^o) + \phi \left((x'_b)^2 + \sigma_b^2 \right) \cdot x'_b \\ &= \phi_b^0 x'_b + \phi (x'_b)^3,\end{aligned}\tag{A.8}$$

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:

$$\pi_b = b_1 x'_b + b_2 (x'_b)^3 + \eta_b.\tag{A.9}$$

where the error term $\eta_b \equiv (\phi_b^0 - b_1)x'_b + v_b$, with v_b representing a white noise disturbance.

We want to show that (i) the coefficient in front of the linear term, b_1 , is a constant that equals the average frequency of price adjustment between observations that belong to the bins in the regression sample; (ii) the estimate of b_1 is unbiased. (iii) b_2 converges in probability to ϕ .

Denote by $\bar{\sigma}_b^2$ the average of the variances between all bins b in the regression sample. Adding and subtracting $\phi \bar{\sigma}_b^2$ to Equation (A.8) we obtain:

$$\pi_b \approx \left((1 - \theta^o) + \phi \sigma_{ss}^2 \right) \cdot x'_b + \phi (x'_b)^3 + \left(\phi (\sigma_b^2 - \sigma_{ss}^2) \cdot x'_b \right)$$

where the the $\left(\phi (\sigma_b^2 - \bar{\sigma}_b^2) \cdot x'_b \right)$ is equal to the error term, v_b , in regression (A.9). The term $\left((1 - \theta^o) + \phi \bar{\sigma}_b^2 \right)$ is equal to the coefficient b_1 . Given that the average price gap is approximately zero, $\int (x'_{ft-1})^2 df \approx 0$ and the coefficient in front of the linear term captured the average frequency of price adjustment across the bins in the regression sample: $\left((1 - \theta^o) + \phi \bar{\sigma}_b^2 \right) \approx \bar{h}_b$.

Finally, we can show that estimator \widehat{b}_1 from model (A.9) converges in probability to \bar{h}_b . To do this, we need to show that the following exclusionary restriction holds:

$$Cov(x'_b, \phi (\sigma_b^2 - \bar{\sigma}_b^2) \cdot x'_b) = 0$$

Define an indicator for a gap being positive:

$$\mathbb{I}^+ \equiv \begin{cases} 1 & \text{if } x'_b > 0 \\ 0 & \text{otherwise} \end{cases}$$

and similarly for the negative gaps (\mathbb{I}^-). Then, we have that:

$$\begin{aligned} & Cov(x'_b, \phi(\sigma_b^2 - \bar{\sigma}_b^2) \cdot x'_b) = \\ & Cov(\mathbb{I}^+ \cdot x'_b, \phi(\sigma_b^2 - \bar{\sigma}_b^2) \cdot x'_b) + Cov(\mathbb{I}^- \cdot x'_b, \phi(\sigma_b^2 - \bar{\sigma}_b^2) \cdot x'_b) = \\ & Cov(\mathbb{I}^+ \cdot x'_b, \phi(\sigma_b^2 - \bar{\sigma}_b^2) \cdot x'_b) - Cov(\mathbb{I}^- \cdot x'_b, \phi(\bar{\sigma}_b^2 - \sigma_b^2) \cdot x'_b) = \\ & Cov(\mathbb{I}^+ \cdot x'_b, \phi(\sigma_b^2 - \bar{\sigma}_b^2) \cdot x'_b) - Cov(\mathbb{I}^+ \cdot x'_b, \phi(\sigma_b^2 - \bar{\sigma}_b^2) \cdot x'_b) = 0 \end{aligned}$$

where the last line follows from the fact that both x'_b and $(\sigma_b^2 - \bar{\sigma}_b^2)$ are symmetric around zero. Finally, the same argument applies to $(x'_b)^3$, so the estimator is consistent.

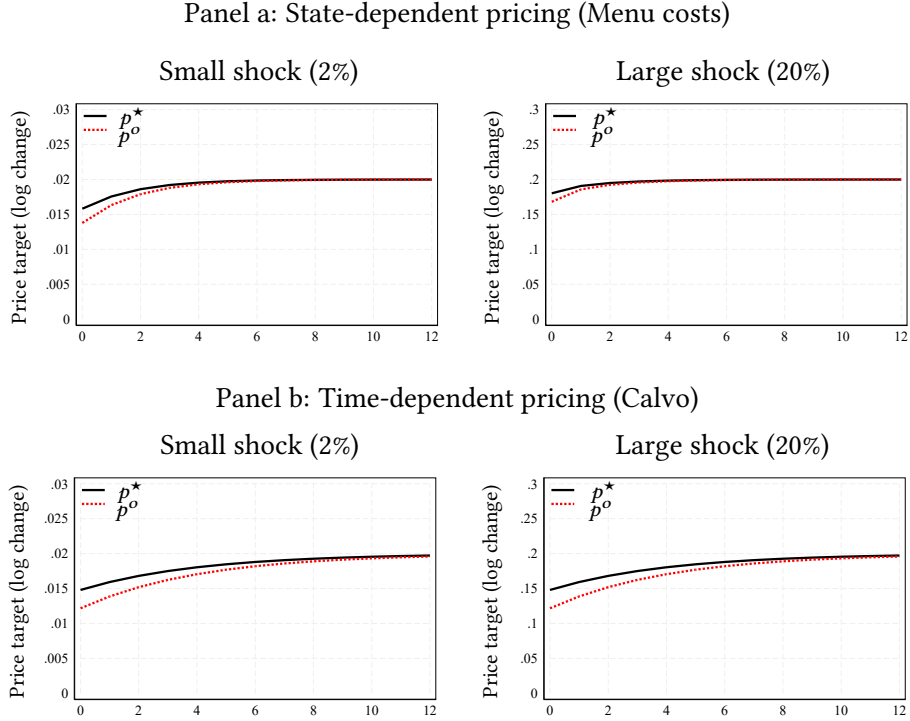
B Additional quantitative exercises

In this section, we present additional quantitative exercises using the simulation of the calibrated menu costs 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}^* . The second exercise studies the role of strategic complementarities in price setting.

Approximation of p_{ft}^* with p_{ft}^o . As discussed in Section 2, the two prices coincide in a steady state with zero trend inflation and constant markups. We also argued that the two prices remain sufficiently close to each other as long as trend inflation is not too large, even in the presence of strategic complementarities in pricing. We therefore assumed $p_{ft}^o \approx p_{ft}^*$, which implies that $x_{ft}^* \approx 0$, 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_{ft}^o approximates p_{ft}^* away from the steady state.

The impulse response functions shown in Figure A.1 indicate that, as

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^*) to aggregate cost shock of different sizes. The x-axis reports quarters since the shock.

expected, the static reset price responds more than the static one to cost shocks, since the dynamic optimum p_{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 $p_{ft}^o \approx p_{ft}^*$ 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_{ft}^o as an approximation for p_{ft}^* has a small impact on 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_{ft}^* as a measure of target price; the red dashed line displays the time-series of model-based inflation, solving the model with p_{ft}^o as a proxy for p_{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^* (the exact, dynamic reset price) and using p^o (the static approximation of p^*) 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 one 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, the strategic complementarities generate additional discounting, which reduces cost pass-through 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

Panel a: State-dependent pricing (Menu costs)



Panel b: Time-dependent pricing (Calvo)



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 without strategic complementarities ($\Omega = 0.5$, red dashed line). The blue dotted 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 costs model). Panel b reports the impulse responses for a time-dependent model (Calvo model), calibrated to display the same steady-state frequency of price adjustment as the time-dependent model. The x-axis reports quarters since the shock.

References