An Odd Approximation of Inflation for State-dependent Pricing Models

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

November, 2026

Abstract

We derive an aggregate state-space representation of state-dependent pricing models with random menu costs and strategic complementarities. The representation, derived from first principles, replaces the cross-sectional distribution of price gaps with a finite set of its moments (the first three), while preserving the key economic forces of state-dependent pricing–movements in the frequency of price adjustment and selection. Along the impulse response, inflation is an odd (cubic) polynomial in the average price gap. Embedded in a general equilibrium setting, the model can be solved by perturbation, filtered, and estimated with standard DSGE tools, and it matches well the global solution across shock sizes. Applied to the 2021–22 inflation surge in the United States and the euro area, it shows that much of the variation that linear models would attribute to cost-push shocks is in fact endogenous amplification through the frequency of price changes and selection.

Keywords: Inflation; State-dependent models; Menu-cost; Mean-field games; Strategic complementarities; Nonlinearities.

JEL: E30.

*Yale University. Email: luca.gagliardone@yale.edu; †New York University and NBER. Email: mark.gertler@nyu.edu; ‡Federal Reserve Bank of New York, University of Cambridge, and CEPR. Email: s.lenzu@jbs.cam.ac.uk; §National Bank of Belgium and KU Leuven. Email: joris.tielens@nbb.be; ★New York University. Email: ame384@nyu.edu. We thank Keshav Dogra, Francesco Lippi, Geert Mesters, Galo Nuño, Johannes Pfeifer, and Christian Wolf for helpful comments and conversations, and the participants to the 2026 NBER Summer Institute, St. Louis Fed, University of Minnesota, 2025 Northwestern Nemmers Prize Conference, 2026 CEBRA annual meeting, Hydra Workshop on Dynamic Macroeconomics, 2026 MMF conference, and 2026 RISE Conference. 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, the Federal Reserve Bank of New York, the Federal Reserve System, or any other institution with which the authors are affiliated. We acknowledge the use of AI tools for language editing and assistance with calculations. Any errors are our own.

1 Introduction

Canonical time-dependent models of price setting (Taylor 1980; Calvo 1983) imply a log-linear relationship between inflation and fundamentals. The 2021–22 inflation surge exposed the limits of this linearity: forecasts built on linear Phillips curves missed the surge, and the frequency of price adjustment, flat for two decades, rose sharply with it in the United States and in the euro area alike (Cotton and Garga 2026; Gautier et al. 2026). Models with state-dependent pricing generate exactly this pattern: a large shock displaces firms' listed prices from their targets, and more firms reprice (Blanco et al. 2026; Gagliardone et al. 2026). In menu-cost models, the firms that reprice are also those whose prices are furthest from target, so inflation accelerates further (Caballero and Engel 1993, 2007; Golosov and Lucas 2007). Despite their empirical appeal, state-dependent models remain rare in applied macroeconomic work. The deterrent is computational: their equilibrium tracks the entire cross-sectional distribution of price gaps, an infinite-dimensional state that requires global solution methods. Moreover, with strategic complementarities in price setting, the equilibrium becomes a mean-field game (Alvarez et al. 2023), difficult to solve and harder to blend into standard DSGE estimation frameworks.

This paper develops a solution to the computational challenges of the frontier class of state-dependent pricing models with random menu costs and strategic complementarities. We show that the model admits, at the aggregate level, an accurate state-space representation in terms of a finite set of moments of the distribution. In particular, denoting the average price gap by x_t, defined as the cross-sectional mean of the difference between firms' optimal reset prices and their listed prices, and by σ_t^2 and ψ_t the variance and skewness of the same distribution, we show that inflation satisfies the following equation with third-order precision in the dispersion of price gaps:

$$\pi_t \simeq \left(h(0) + 3\phi\sigma_t^2\right) x_t + \phi\left(x_t^3 + \psi_t\right),$$

where $h(0)$ and ϕ —hazard rate at the optimum and the curvature of the hazard function (Alvarez et al. 2022)— are parameters of the adjustment decision of firms that can be inferred from data on firm prices and costs.

Because inflation is an odd function of the history of the average price gap, we call this mapping the *odd approximation* of inflation. When a sufficiently large shock hits, the average price gap moves enough to activate the cubic term and to trigger adjustments in variance and skewness, which we also characterize. The variance determines the elasticity of inflation to first-order shocks, so that small shocks may have a large impact when gaps are dispersed. Incorporated in this nonlinear model are the two hallmarks of state dependence—an endogenous increase in the frequency of adjustment and selection among repricing firms. The cubic gives inflation an

S-shaped relation with the average price gap, which mirrors the relation between price gaps and price changes that Gagliardone et al. (2026) document in firm-level data.

Using our state-space representation, we further provide a novel characterization of the entire path of the impulse response function of inflation to an aggregate shock. Along the impulse response, we show that the information contained in the higher-order moments can be further collapsed into a single aggregate scalar—the average price gap. Up to a third-order approximation, inflation then takes an intuitive and interpretable form as an odd (i.e. cubic) polynomial in the average price gap:

$$\pi_t \simeq \beta_{1,t}\, x_t + \beta_{3,t}\, x_t^3.$$

From a theoretical standpoint, the approximation along the impulse response preserves and sharpens the economic forces underlying state-dependent pricing despite its parsimony and "oddness". First, its polynomial coefficients are closed-form functions of primitives measured in the model's steady state and the equilibrium persistence of the average gap. Second, we demonstrate quantitatively that the same formula with *constant* coefficients and an appropriate recalibration is well suited to capture the full nonlinear dynamics of the model—entirely abstracting from movements in the higher-order moments of the distribution. We call the latter exercise with constant coefficients "a poor man's simplification," as it provides a simple and portable block to any extension of DSGE models that aims at quantitatively accounting for nonlinearities in inflation.

The methodological innovation proposed in this paper opens up the study of the time series of inflation and frequency of price adjustments using standard methods developed for the DSGE literature: filtering, smoothing, and estimation. To illustrate, we use a textbook New Keynesian model (Galí 2015) augmented with state-dependent pricing. The model features a nonlinear feedback between the interest rate and price gaps—usually treated as exogenous by the literature on state-dependent pricing—which poses serious computational challenges when paired with strategic complementarities. Our approach instead permits the use of standard perturbation methods (e.g., Dynare) so that the model is solved in seconds. A novel insight that comes out of the analysis is a substantial compression of the importance of cost-push shocks, which previous literature based on linearized models has found—puzzlingly—to be the main driver of inflation (Smets and Wouters 2007). We interpret this finding as cost-push shocks capturing the nonlinearities that instead our model directly allows for.

Section 2 presents a canonical state-dependent pricing model in which firms face idiosyncratic shocks, random menu costs, and occasional free adjustment opportunities (Dotsey and King 2005; Caballero and Engel 2007; Nakamura and Steinsson 2010), embedded in a standard general equilibrium framework. The characterization of the general equilibrium highlights the

central obstacles to the global solution: the state space is the full cross-sectional distribution of price gaps and, when firms' desired prices depend on aggregate prices, the equilibrium is described by a nonlinear mean-field game.

Section 3 derives the odd approximation. First, we derive a quadratic approximation of the firm-level probability of price adjustment, which implies that, near its target, a firm's adjustment hazard is convex in its own price gap. Aggregating this behavior across firms allows us to express inflation in terms of moments of the price-gap distribution. Along the impulse-response function to aggregate shocks, these moments reduce aggregate inflation and frequency to cubic and quadratic polynomials in the average price gap alone, respectively. Underlying the results is a remarkably simple intuition. The third-order term arises from the product of a second-order extensive-margin response—the adjustment frequency rises with the square of the average price gap—and a first-order intensive-margin response—the size of a price change is proportional to that gap. Selection provides an additional channel of amplification. Finally, we embed the approximation in a recursive general equilibrium. The mapping from the average price gap to inflation and to the adjustment frequency closes into a system of aggregate recursions, together with the laws of motion of the distribution's variance and skewness. Standard perturbation solvers, such as Dynare, can handle this system directly.

Section 4 validates the approximation against the globally solved menu-cost model. Calibrated to the same primitive parameters as the global model, the odd approximation closely tracks the global model's inflation, adjustment-frequency, and output dynamics across small, medium, and large shocks. For small shocks, it recovers the essentially linear behavior of inflation, consistent with the local equivalence between Calvo and menu-cost models shown by Gertler and Leahy (2008) and Auclert et al. (2024). For large shocks, its cubic term activates, capturing the endogenous rise in the adjustment frequency, which Calvo misses, and the faster price-level pass-through and greater short-run monetary neutrality that come with it. We also introduce a poor man's simplification, with constant coefficients fitted directly to the global impulse responses, which—we show—can also match the global model almost exactly.

In Section 5 we take the odd approximation to US and euro-area data. We parameterize the pricing block, solve the model with third-order perturbation methods, and recover structural shocks with a nonlinear smoother. Despite the simplicity of the model, the historical shock decomposition yields a sensible decomposition of the 2021–22 surge: cost-push shocks triggered at the beginning, a large and persistent contribution of demand, amplified by monetary accommodation. In normal times, the historical shock decomposition produced by the nonlinear smoother is consistent with the one produced by a standard Kalman smoother applied to the linearized version of the model (a Calvo model). During large-shock episodes, however, the

nonlinear model attributes less of the inflation surge to exogenous cost-push shocks (*cost-push compression*) and more to demand. This is because the linear model confounds exogenous shocks with the endogenous amplification that operates through frequency and selection.

Applying our algebraic decomposition of inflation, we find that the frequency and selection channels account for about half of peak model-implied inflation during the surge in the United States and about two-fifths in the euro area. The model reproduces the observed timing of the frequency response and, especially for the euro area, its magnitude. Accordingly, the effective slope of the Phillips curve nearly doubles in the United States in 2008–09 and rises by about three-quarters in 2021–22. In the euro area, it reaches more than twice its steady-state value at the 2022 peak.

In Section 6 we leverage the tractability of the odd approximation to perform policy counterfactuals. Feeding the recovered innovations through the estimated model and a frequency-matched Calvo benchmark, we study the cost of looking through the surge. Looking through the surge costs about two to three times the Calvo prediction, because monetary accommodation leaves the frequency-selection feedback loop unchecked. The trade-off facing a central bank that tightens instead is state-dependent, as in Karadi et al. (2025): the slope-implied output cost per point of disinflation falls to about 30 percent of the constant Calvo value near the US surge peak, and to about a quarter of it in the euro area's larger surge, because state-dependent pricing adds a demand-like component of inflation that policy can lean against, most effectively when the shock is large.

Relation to the literature. Our benchmark state-dependent pricing model with random menu costs, free price adjustment, a quadratic approximation to profits, and a mean-field game structure builds on the prior contributions of Dotsey et al. (1999), Nakamura and Steinsson (2010), and Alvarez et al. (2023). Our aggregation result leverages the approximation of adjustment decisions through a quadratic hazard function of price gaps (Alvarez et al., 2022; Alvarez and Lippi, 2022; Gagliardone et al., 2026).[1]

The paper most directly related to ours is Blanco et al. (2026), which derives an aggregation result for a modification of the canonical menu cost model. In particular, they propose a variation in which multiproduct firms choose the fraction of product prices to adjust, but not which ones, and thus abstract from selection among adjusting products. Instead, we work with a conventional menu cost framework and show that, up to third order, it admits an aggregate state-space representation without imposing any non-standard behavioral assumptions. This allows us to

[1]The quadratic hazard is consistent with the convex adjustment probabilities documented in micro data (Eichenbaum et al., 2011; Gautier et al., 2023; Karadi et al., 2024a; Gagliardone et al., 2026).

assess the importance of selection, which we show significantly contributes to explaining the 2021-22 inflation surge. Also related is Adams (2026), which reduces the dimensionality of the problem by numerically approximating the dynamics of the distribution with its primary eigenfunction. Instead, we provide a closed-form characterization of the effects of the distribution on inflation dynamics that preserves and sharpens the economic intuition while delivering the same quantitative tractability.

Our results connect with Auclert et al. (2024) and Auclert et al. (2026). In Auclert et al. (2024), the authors show that, to first order in aggregate shocks, the canonical menu-cost model is equivalent to a mixture of two time-dependent models with constant adjustment frequencies. We provide a further characterization of the relation between these classes of models up to the third order: while the first-order effect of an aggregate shock depends on the dispersion in price gaps, changes in frequency contribute to inflation at a third order. Auclert et al. (2026) extend sequence-space Jacobian methods to third order, obtaining a stochastic solution to a menu-cost model by numerical differentiation of the impulse response around the steady state. The two approaches are distinct and complementary. First, our focus is on deriving an explicit state-space representation of the model. This allows us not only to solve the model efficiently, but also to produce a historical decomposition. Second, we obtain closed-form expressions for the impulse response of inflation directly from the problem's primitives, whereas their solution is numerical. At first order, our mapping recovers the small-shock equivalence results of Gertler and Leahy (2008) and Alvarez et al. (2017).

A related line of research studies the role of the extensive margin of adjustment in quantitative models and their implications for shock transmission (Costain and Nakov, 2011; Ghassibe, 2024; Ghassibe and Nakov, 2025) and monetary policy (Afrouzi et al., 2024; Karadi et al., 2025; Ascari et al., 2025). While our general equilibrium model is purposely stylized and abstracts from important quantitative aspects and optimal policies, our state-space representation can be embedded in richer models to simplify computations and perform filtering, smoothing, and parameter estimation. Moreover, our representation can be employed to study nonlinear optimal policies in a tractable yet general framework. In particular, using a textbook three-equation New Keynesian model (Galí, 2015) extended to allow for the extensive margin of adjustment, we leverage the shocks recovered from the 2021-22 inflation surge to quantify the gains from a more aggressive monetary policy.

Our paper adds to the evidence that the post-pandemic inflation surge was a state-dependent episode. The sharp rise in repricing frequencies in the United States and the euro area (Cotton and Garga, 2026; Gautier et al., 2026) is consistent with the nonlinear inflation dynamics emphasized, among others, by Harding et al. (2023), Blanco et al. (2026), Cavallo et al. (2024), and Gagliardone

et al. (2026). Our tractable higher-order approximation provides a tool to study historical episodes using aggregate data that accounts for the movements in the aggregate frequency of price changes.

Methodologically, the odd approximation confronts the obstacle of solving for equilibrium dynamics that depend on an infinite-dimensional cross-sectional distribution. The Krusell and Smith (1998) algorithm offers a numerical solution to this problem. It posits a parametric law of motion for few moments, checks it by simulation, and verifies that forecast errors are small. Within our state-dependent framework, we instead derive the mapping from moments to aggregates in closed form from the pricing problem's primitives. Our odd approximation also addresses the informational burden of full rational-expectations models emphasized by Moll (2025). Along the impulse response and in the quantitative poor man's simplification, firms need to track only the first moment of the price-gap distribution. This allows us to obtain an algebraic reduction of the (still rational-expectations) equilibrium to a recursive, low-dimensional pricing block.

2 State-dependent pricing within a canonical NK framework

We characterize firms' pricing policies and aggregate inflation in a generalized menu-cost model with smooth adjustment costs and quadratic profits. We embed the resulting pricing block in a standard three-equation NK model with consumption goods only and a nominal interest rate rule (e.g., Galí 2015).

Throughout the paper, lowercase letters denote logarithms of the corresponding uppercase variables for prices, wages, and quantities (e.g., $p_t = \ln P_t$); all other lowercase symbols are defined where they are introduced. The operators $O(\cdot)$ and $o(\cdot)$ denote the Landau orders (Big O, small o) as their arguments tend to zero. All derivations and proofs are presented in Appendices B and C. For expositional purposes, in the baseline model we assume no trend inflation, Gaussian shocks, and an invertible demand system with constant elasticity across varieties.

2.1 The pricing block

In each period t, the economy is populated by a continuum of firms indexed by $f \in [0, 1]$, each selling a single differentiated product under imperfect competition. Each firm faces a residual demand function:

$$y_t(f) = d\Big(p_t(f), \zeta_t(f), p_t\Big) + y_t,$$

defined over its own price $p_t(f)$, a taste shock $\zeta_t(f)$, and the aggregate price index p_t, defined as the average taste-adjusted listed prices, $p_t(f) - \zeta_t(f)$.

Output is produced using labor with a production function featuring short-run

returns-to-scale equal to $1/(1+\nu)$ and Hicks-neutral (log-)productivity, given by the sum of an aggregate component a_t and an idiosyncratic component $a_t(f)$:

$$y_t(f) = \Big(a_t + a_t(f)\Big) + \frac{n_t(f)}{1+\nu}.$$

We assume the firm-level component, $a_t(f)$, evolves as a random walk in levels with increments $\epsilon_t(f)$ drawn i.i.d. across firms and time from a mean-zero Gaussian distribution with variance σ^2_ϵ. These increments are realized at the beginning of the period, before firms set prices.

Profit maximization. Firms maximize intertemporal profits by choosing factor demands and pricing policies subject to nominal and real rigidities. Firms' nominal marginal cost is shifted by an aggregate cost-push disturbance u_t. Following Clarida et al. (1999), we take it as a reduced-form wedge that stands in for supply-side disturbances, such as oil or energy price increases. Given output, cost minimization implies that a firm's nominal marginal cost is given by:

$$mc^n_t(f) = -(1+\nu)a_t(f) + \xi_t + \nu\, y_t(f) \qquad\qquad \xi_t := w_t + \log(1+\nu) - (1+\nu)a_t + u_t \tag{1}$$

The first two terms capture the idiosyncratic and aggregate components of the firms' cost shifter. The idiosyncratic component, $-(1+\nu)a_t(f)$, inherits the random walk dynamics from the firm's own productivity. The aggregate component, ξ_t, combines an endogenous term, the nominal wage w_t, with two exogenous drivers: aggregate productivity, a_t, and the cost-push disturbance, u_t.

Nominal rigidities arise due to real costs of adjusting prices. Upon changing its price, a firm pays a strictly positive fixed, random cost $\chi_t(f)$, denominated in units of the final good, i.i.d. across firms and over time and drawn from an exogenous, continuous distribution $G(\cdot)$ with finite density and $G(0) = 0$. As in the Calvo-plus model of Nakamura and Steinsson (2010), we allow in addition a probability of free adjustment, carried separately from G (Equation (5)).

Given the idiosyncratic realization of the menu cost, the firm jointly chooses whether to adjust its price (optimal choice denoted with a star, $\mathbb{I}^\star_t(f) \in \{0,1\}$), and, conditional on adjustment, the (dynamic) target price $p^\star_t(f)$, to maximize firm value:

$$V_t(f) = \max_{\{\mathbb{I}_{t+i}(f)\,,\, p_{t+i}(f)\}_{i\geq 0}} \mathbb{E}_t \sum_{i=0}^{\infty} \beta^i \Big\{\Pi_{t+i}(f) - \chi_{t+i}(f)\,\mathbb{I}_{t+i}(f)\Big\}, \tag{2}$$

where $\beta \in (0,1)$ and $\Pi_t(f)$ denote the firm's discount factor and (real) flow profits (revenues minus total costs), respectively.

Definitions and assumptions. We define four firm-level objects:

(*i*) *Static target price*—the price a firm would choose in the absence of nominal rigidities.

Following Alvarez et al. (2023), we work with a quadratic approximation of the static profit function; as we show in Appendix B, profit maximization yields a static target price of:

$$p_t^o(f) = (1 - \Xi)\big(\mu + \xi_t + \nu\, y_t\big) + \Xi\, p_t + \zeta_t(f), \tag{3}$$

where μ denotes the log steady-state net markup and y_t is aggregate output. The parameter $\Xi \in [0, 1)$ is the weight the static target places on competitors' prices, which can arise from decreasing returns to scale in production or strategic complementarities in price setting. In both cases, $\Xi \neq 0$—a firm's pricing decisions depend on its own unit cost as well as on its competitors' prices—and the model is a mean-field game (Alvarez et al. 2023), in which firms need to forecast the evolution of the price distribution in order to set their price.

Taste shocks shift the demand for a firm's variety, and productivity shifts its marginal cost. Following the literature, we assume that taste shocks move one-for-one against idiosyncratic productivity in levels, $\zeta_t(f) = -a_t(f)$ (Woodford 2009; Midrigan 2011; Alvarez et al. 2016). Under this assumption, the idiosyncratic terms of the static target combine into $\zeta_t(f)$ with a unit loading (Appendix B), and the firm's pricing problem depends on a single idiosyncratic state (see Lemma 1).

(*ii*) *Price gap*—the deviation of the reset price from the price the firm carried into the period:

$$x_t(f) := p_t^\star(f) - p_{t-1}(f). \tag{4}$$

The price gap is measured before adjustment decisions are made, as it incorporates the realization of period t aggregate and idiosyncratic shocks, as we explain below. It equals the firm's price change conditional on adjustment.

We denote by $F_t(x) := \Pr\big(x_t(f) \leq x\big)$ the cross-sectional distribution of price gaps at time t, with mean x_t, variance σ_t^2, (not standardized) skewness ψ_t, and kurtosis κ_t given by:

$$x_t := \int x\, dF_t(x), \qquad \sigma_t^2 := \int (x - x_t)^2\, dF_t(x),$$
$$\psi_t := \int (x - x_t)^3\, dF_t(x), \qquad \kappa_t := \int (x - x_t)^4\, dF_t(x).$$

Throughout, we maintain regularity assumptions on the distribution that allow us to provide bounds on the higher-order moments of the cross-sectional distribution of price gaps, spelled out formally in Appendix A.[2] Finally, we denote by F_0 the distribution of price gaps in the

[2] Assumptions *4a–c* in Appendix A are three requirements on the shape of the distribution of increasing strengths. Assumption *4a* is a weak regularity requirement that permits us to bound the mean gap by the variance and is used in derivations of the bounds of Equation (17). Assumption *4b* is a moderate regularity requirement that allows us to abstract from the effects of selection on the dynamics of variance and skewness (up to, respectively, second and third order approximations) while retaining the effects of selection on inflation. It is used to derive a closed system that

stochastic steady state (when only idiosyncratic shocks are present) and by $x_0, \sigma_0^2, \psi_0, \kappa_0$ the corresponding moments.

(*iii*) The *reset gap*—the difference between the static and dynamic optima:

$$x_t^\star(f) := p_t^o(f) - p_t^\star(f).$$

As in Alvarez et al. (2016) and Gagliardone et al. (2026), the reset gap is an aggregate scalar, homogeneous across firms, $x_t^\star(f) = x_t^\star$ (see Appendix B.2).[3] Observe that the sum of the reset gap and the price gap equals the gap between the static optimal price and the price carried into the period, a feature we make use of below.

(*iv*) *Hazard rate*—the optimal probability of price adjustment between consecutive periods, conditional on the realization of the price gap but prior to observing the idiosyncratic menu cost draw:

$$h_t(x) := \mathbb{E}_t\{\mathbb{I}_t^\star(f) \mid x_t(f) = x\}.$$

Characterization of pricing policies. We now characterize firms' pricing policies—the decision to adjust (extensive margin) and the size of the adjustment (intensive margin). To do so, Lemma 1 establishes the price gap as the sufficient firm-specific state variable, which characterizes both the value function and the hazard rate.

Lemma 1 (Price gaps as state)**.** *The value function $\mathcal{V}_t : x_t(f) \mapsto \mathcal{V}_t(x_t(f)) \in \mathbb{R}$ satisfies the following conditions:*

(i) Recursive representation:

$$\mathcal{V}_t(x) = -\frac{B}{2}\left(x_t^\star + x\right)^2 + \beta\, \mathbb{E}_t\Big[h_{t+1}(x')\Big(\mathcal{V}_{t+1}(0) - \chi_{t+1}\Big) + \big(1 - h_{t+1}(x')\big)\,\mathcal{V}_{t+1}(x')\Big],$$

where $B > 0$ is the curvature of the flow profit function (Appendix B), x' next period's price gap, and χ_{t+1} the average cost paid by the firms that do adjust.

(ii) Interior optimum: $\mathcal{V}_t'(0) = 0, \quad \mathcal{V}_t''(0) < 0.$

(iii) Optimality principle: $V_t(f) = \max\Big\{\mathcal{V}_t(x_t(f)),\ \mathcal{V}_t(0) - \chi_t(f)\Big\}$

does not involve further higher-order moments in Proposition 2. Assumption *4c* is a strong regularity condition that requires that aggregate shocks and idiosyncratic heterogeneity vanish at the same rate, and is used in the closed-form characterization of the impulse response function of inflation and frequency (Theorem 2).

[3] Two features of the environment deliver this. The static target moves one for one with the firm's own shifter, and that shifter is a random walk, so a firm's best forecast of its own future targets is its current one, displaced by aggregate news that every firm shares. Measured from its own target, every firm therefore faces the same problem and picks the same distance from it.

Moreover, in a stationary environment where $x_t^\star \equiv x^\star$ and the conditional law of motion of $x_{t+1}(f)$ is time-homogeneous, $\mathcal{V}_t(x) \equiv \mathcal{V}(x)$ and $h_t(x) \equiv h(x)$.

Proof. See Appendix C.1. □

Extensive margin decision—Under Lemma 1, the hazard function is characterized by:

$$h_t(f) \equiv h_t(x_t(f)) \;=\; h_t(0) \;+\; \big(1 - h_t(0)\big)\, G\Big(\mathcal{V}_t(0) - \mathcal{V}_t(x_t(f))\Big). \tag{5}$$

The endogeneity of firms' hazard function is at the core of state-dependent pricing. Intuitively, the probability that a firm adjusts its price increases with the distance between the value of adjusting and the value of not adjusting. As in the Calvo-plus model of Nakamura and Steinsson (2010), we also allow firms to adjust their prices for free with a positive probability given by $h_t(0)$. The Calvo (1983) model is a special case nested in this framework, in which $G \to 0$ pointwise on $(0, \infty)$—menu costs grow so large that no firm pays them and adjustment occurs only through the free-adjustment probability—so that the hazard is homogeneous across firms and is typically time-invariant ($h_t(x_t(f)) = h(0)$).

Intensive margin decision—The first-order condition of problem (2) with respect to $p_t^\star(f)$ is:

$$0 = \left\{ \frac{\partial \Pi_t(f)}{\partial p_t(f)} + \beta \mathbb{E}_t \left(\frac{\partial V_{t+1}(f)}{\partial p_t(f)} \right) \right\} \Bigg|_{p_t(f) = p_t^\star(f)}$$

As we show in Appendix B, iterating the equation above forward and approximating $\Pi_t(f)$ up to second order around the static optimum leads to the dynamic reset price being a present value of current and expected static target prices and hazards:

$$p_t^\star(f) = \frac{\mathbb{E}_t \sum_{i=0}^{\infty} \beta^i \left(\prod_{j=1}^{i} (1 - h_{t+j}(f)) \right) p_{t+i}^o(f)}{\mathbb{E}_t \sum_{i=0}^{\infty} \beta^i \left(\prod_{j=1}^{i} (1 - h_{t+j}(f)) \right)}. \tag{6}$$

Here $h_{t+j}(f)$ is the firm's adjustment probability evaluated at the price gap it would carry into period $t + j$ had it not reset since t. Thus, $\mathbb{E}_t \prod_{j=1}^{i} (1 - h_{t+j}(f))$ is the probability that the price set at t is still in effect at $t + i$. As in Calvo (1983), when choosing its reset price, a firm recognizes that today's price may remain in effect for several periods, so expected cost dynamics matter. Unlike the Calvo model, the adjustment probability is firm- and time-specific, varying with current and expected price gaps (Dotsey and King 2005).

Aggregate inflation. Let $h_t := \int h_t(x)\, dF_t(x)$ denote the average frequency of price adjustment at time t. We define inflation as the log change in average of firm-level taste-adjusted listed prices:

$$\begin{aligned}\pi_t &:= \int_0^1 \Delta p_t(f)\, df = \int h_t(x) \cdot x\, dF_t(x) \\ &= h_t \cdot x_t + \mathbb{C}\mathrm{ov}_{F_t}\Big(h_t(x), x\Big) \end{aligned} \tag{7}$$

The first equality uses the fact that only adjusters change their price and applies the law of large numbers to replace the adjustment indicator with the hazard. In the second line, we use the cross-sectional covariance identity to decompose aggregate inflation into two terms. The first is the product of the *extensive margin*—the average frequency of price adjustment h_t—and the *intensive margin*—the average price gap x_t. The second is the *selection effect*—the cross-sectional covariance between firms' adjustment probability and their price gap.[4]

2.2 Closing the model

The economic relationships outside the pricing block are standard. (Because we do not have capital or investment, for example, we do not need to include firms in the non-pricing block.) We deliberately work with standard log-linearized versions of these equations, allowing us to focus on the pricing block as a source of nonlinear inflation dynamics. Richer interactions or nonlinearities elsewhere are conceptually straightforward to add but beyond our scope.

Household. A representative household chooses consumption C_t and labor supply L_t to maximize lifetime utility:

$$\max_{C_t, L_t} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t e^{\gamma b_t} \Big(\frac{C_t^{1-\gamma}}{1-\gamma} - \frac{L_t^{1+\varphi}}{1+\varphi} \Big),$$

where $\gamma > 0$ denotes the inverse of the intertemporal elasticity of substitution, and $\varphi > 0$ denotes the inverse Frisch elasticity of labor supply. The variable $b_t = \rho_b\, b_{t-1} + \varkappa_t$ is an exogenous preference disturbance (a pure demand shock) that shifts the household's willingness to postpone consumption.[5] As is standard, in our baseline model, the aggregate consumption bundle C_t aggregates the differentiated goods produced by firms via a constant-elasticity demand system with the firm-specific taste shocks $\zeta_t(f)$ entering as demand shifters.

The household can transfer resources intertemporally by purchasing one-period riskless

[4] The accounting decomposition in Equation (7) is closely related to that in Costain and Nakov (2011), who quantify the intensive, extensive, and selection margins in a globally solved model of smoothly state-dependent pricing.

[5] Because b_t scales the entire period utility, it leaves the household's labor supply schedule, and hence marginal costs and the pricing block, untouched. Writing the shifter as $e^{\gamma b_t}$ rather than e^{b_t} normalizes the units of b_t so that it enters the consumption Euler equation with a unit coefficient; the two coincide under the log utility we calibrate.

discount bonds, B_t, which pay a nominal yield i_t. The intertemporal budget constraint is given by $P_t C_t + B_t = W_t L_t + D_t + (1 + i_{t-1}) B_{t-1}$, where W_t denotes the nominal wage and D_t the aggregate nominal dividend, equal to firms' real profits net of menu costs times the price level, accruing from the ownership of the production sector.

The household side of the model is characterized by standard labor supply and log-linearized consumption Euler equations:

$$w_t - p_t = \gamma c_t + \varphi l_t \tag{8}$$

$$c_t = \mathbb{E}_t\{c_{t+1}\} - \gamma^{-1}(r_t + \ln\beta) + b_t - \mathbb{E}_t\{b_{t+1}\}, \tag{9}$$

where the real rate satisfies the Fisher equation $r_t = i_t - \mathbb{E}_t \pi_{t+1}$. Market clearing requires:

$$N_t = L_t \tag{10}$$

$$C_t = Y_t - \mathcal{M}_t, \qquad \mathcal{M}_t := \int_0^1 \chi_t(f)\, \mathbb{I}_t^{\star}(f)\, df, \tag{11}$$

once the resources absorbed by price adjustment (menu costs) are accounted for.[6]

Monetary policy. To close the model, we take a stance on how the interest rate i_t is determined. We use a nominal-rate rule that obeys the Taylor principle, as in the textbook NK model (Galí 2015):

$$i_t = -\ln\beta + \phi_\pi\, \pi_t + \iota_t, \tag{12}$$

$$\iota_t = \rho_\iota\, \iota_{t-1} + \varsigma_t, \tag{13}$$

where ι_t denotes a monetary policy shock.

2.3 The general equilibrium

Given a price-gap distribution and exogenous shock processes at date 0, a stochastic general equilibrium of the global model is a path of cross-sectional price-gap distributions $\{F_t\}_{t\geq 0}$, firm-level pricing policies $\{h_t(\cdot), p_t^{\star}(\cdot)\}_{t\geq 0}$, and aggregate quantities $\{y_t, \pi_t, i_t, p_t\}_{t\geq 0}$ consistent with the rest of the model: the household's optimality conditions and market clearing (Equations (8)–(9), (10), and (11)); the monetary block (Equations (12)–(13) and the Fisher equation); and each firm's cost minimization and pricing decisions (Equations (1)–(6)). The price level p_t is the average taste-adjusted listed price of the resulting cross-section, and inflation $\pi_t = p_t - p_{t-1}$ is the average price change of Equation (7).

[6] The quantitative exercises below hold the cost of price adjustment at its steady-state value. Being an even function of the average price gap, this term changes the level of output but leaves the odd polynomial of Theorem 2 and its coefficients unaffected.

As we have been discussing, because the price adjustment frequency is endogenous, price and inflation dynamics are potentially nonlinear in state-dependent pricing models. Capturing these nonlinearities requires solving the model globally, which raises three computational challenges.

First, the state space is infinite-dimensional. Because each firm's hazard $h_t(x)$ and reset price $p_t^\star(f)$ depend on its own price gap, the cross-sectional distribution of gaps is itself a state variable of the equilibrium, updating firm by firm according to:

$$F_{t+1}(x') = \int \left[h_t(x)\, \Phi\big(x' - \Delta p_{t+1}^\star\big) + \big(1 - h_t(x)\big)\, \Phi\big(x' - x - \Delta p_{t+1}^\star\big) \right] dF_t(x),$$

where Φ is the c.d.f. of the idiosyncratic increment to the price gap, $-\epsilon_{t+1}(f)$, and $\Delta p_{t+1}^\star := p_{t+1}^\star - p_t^\star$ is the common component of the change in the reset price. Moreover, as the adjustment probability is a nonlinear function of the gap, the law of motion of any given moment of F_t involves higher moments. Computing aggregate inflation therefore requires tracking the entire distribution.

Second, in the presence of macroeconomic or strategic complementarities in price setting ($\Xi \neq 0$, in our setting), the model has the structure of a mean-field game (Alvarez et al. 2023). The optimal reset price $p_t^\star(f)$ depends on expected future realizations of aggregate prices, which depend on the future path of F_t, which in turn depends on firms' collective policies. This implies that policy and the distribution have to be solved for jointly, which is a nonlinear fixed point over functions with no closed-form solution.[7]

Finally, with the exception of simplified general equilibrium setups often used in the literature (e.g., Golosov and Lucas 2007, which effectively treats firms' nominal marginal cost as exogenous), the aggregate drift entering the law of motion of price gaps is itself a general equilibrium object: nominal marginal costs depend on nominal wages and output, which depend on the price level, which in turn depends on firms' pricing policies. This challenge is nested inside the previous two, adding the additional complication that the law of motion for F_t needs to be evaluated simultaneously with both the pricing and demand blocks.

These challenges make global equilibrium computation costly and, *a fortiori*, severely limit the scope for using state-dependent pricing models to conduct quantitative policy exercises—e.g., parameter estimation, filtering and historical shock decomposition, and counterfactual policy analysis—that are routine in time-dependent frameworks.

In the following section, we develop an approximation of the state-dependent pricing block. Under a mild set of assumptions, we derive a representation of inflation as a function of three moments

[7] Costain and Nakov (2011) address this problem numerically in a smoothly state-dependent pricing model, representing the full distribution of firm-level states on a grid and computing its aggregate dynamics. See also Adams (2026), who derives an analytical small-shock solution to a menu-cost mean-field game with trend inflation and proposes a low-dimensional approximation based on the primary eigenfunction of the distribution of price gaps.

of the price-gap distribution: the average gap, its variance, and its skewness. Along the impulse response to an aggregate shock, the representation collapses to a polynomial in the average gap alone. As we show, the approximate model preserves and sharpens the key economic forces of state-dependent pricing, while making it possible to solve for equilibrium and transitional dynamics using standard perturbation methods.

3 The approximate state-dependent model

We develop the pricing block of the approximate state-dependent model in four steps. First, we approximate firm-level adjustment hazards as quadratic in the price gap and aggregate them into the average frequency of adjustment. Second, we project aggregate inflation onto the moments of the price-gap distribution. Third, we provide a recursive characterization of the general equilibrium that embeds a state-dependent pricing block and solvable with standard perturbation methods; we show that inflation and frequency impulse responses are odd and even polynomials, respectively, in the average price gap. In doing so, we maintain the assumptions of zero trend inflation and symmetric Gaussian shocks.

3.1 Approximation of firm-level hazard and aggregate frequency

Approximation of firm-level hazards. Following Alvarez et al. (2022) and Gagliardone et al. (2026), we approximate firms' hazard rates as a convex function of a single state variable, their price gaps. Unlike these papers, we build on Lemma 1 and work with the (dynamic) price gap $x_t(f)$ defined in Equation (4)—rather than its static counterpart $p_t^o(f) - p_{t-1}(f)$. As we make clear below, this allows us to account for the continuation value of the pricing problem and characterize the dynamic evolution of the system. The following proposition formalizes this result.

Proposition 1 (Quadratic approximation of the hazard rate)**.** *Under Lemma 1, up to second order around the reset point, firms' hazard rates are a quadratic polynomial in the price gap:*

$$h_t(x_t(f)) = h_t(0) + \phi_t x_t^2(f) + o(x_t^2(f)), \tag{14}$$

where $h_t(0) \geq 0$ is the adjustment probability at the optimum and $\phi_t := -\frac{1}{2}(1-h_t(0))G'(0)\mathcal{V}_t''(0) \geq 0$ is the curvature of the hazard function. In a stationary environment, Equation (14) features constant coefficients $h_t(0) = h(0)$ and $\phi_t = \phi$ measured in the steady state.

Proof. See Appendix C.2. □

Theoretically, the quadratic approximation holds in the class of state-dependent pricing models whose menu-cost distribution has a continuous density at zero—the random menu costs of the

generalized-hazard framework of Caballero and Engel (1993, 2007) and Dotsey and King (2005), and Calvo-plus specifications in which the free-adjustment probability is carried alongside a smooth menu-cost distribution, as here (Equation (5)).[8] The Calvo (1983) model is the limit $\phi = 0$, $h(x) = h(0)$, the steady-state frequency. Empirically, the quadratic approximation in (14) is consistent with micro evidence of a convex relationship between firm-level adjustment probabilities and price gaps (Gagnon et al., 2013; Gautier et al., 2023; Karadi et al., 2024a; Gagliardone et al., 2026).

Approximation of average frequency of price adjustment. Using Equation (14) and assuming stationarity, we obtain the following expressions for the average frequency of price adjustment in period t and in steady state:

$$h_t = h(0) + \phi\left(\sigma_t^2 + x_t^2\right) + o\left(\sigma_t^2\right) \tag{15}$$

$$\bar{h} = h(0) + \phi\sigma_0^2 + o(\sigma_0^2), \tag{16}$$

Equation (15) illustrates how the average frequency of price adjustment depends on the first and second moments of the price gap distribution in period t. The frequency varies positively with the variance of price gaps. The intuition is the following. In normal times, when aggregate shocks are small, the distribution of price gaps is close to the steady-state distribution ($x_t \approx 0$). During these periods, most firms sit in the flat part of their hazard function (the bottom of the parabola), so the average adjustment frequency is low and stable. By contrast, a large aggregate shock—say, a surge in energy prices—raises firms' optimal prices (and therefore their price gap) in a correlated way. Since firms' hazards are convex, this effect is nonlinear and captured by x_t^2. The steepness of the generalized hazard function (GHF), ϕ, governs this effect.

3.2 The odd approximation of inflation

We are now ready to illustrate the core result of the paper: under the quadratic approximation of the hazard, aggregate inflation collapses to a projection onto low-order moments of the price-gap distribution. The projection is an odd function of the history of the average price gap. Before stating it formally, we walk through the steps that deliver this result.

Start from the definition of aggregate inflation in Equation (7), $\pi_t = \int h_t(x) \cdot x \, dF_t(x)$. Applying Proposition 1, we substitute the quadratic approximation of the firm-level hazard into

[8] The quadratic hazard requires $h_t(x)$ twice-differentiable in the gap, which Assumption 5 imposes through a continuous density for G at the origin. Models with fixed (non-random) menu-cost (Golosov and Lucas 2007; Nakamura and Steinsson 2008) and the Calvo-plus of Nakamura and Steinsson (2010) violate this assumption, as G places a point mass at the menu cost. Their hazard is flat around the reset point, so $G'(0) = 0$ and $\phi = 0$, and adjustment happens at the Ss bounds. The quadratic approximation can nevertheless be applied to them through a smooth hazard estimated on their simulated aggregate behavior (Caballero and Engel 2007).

the integral and distribute the product to obtain:

$$\begin{aligned}\pi_t &= \int \left[h(0) + \phi x^2\right] x \, dF_t(x) + o(\sigma_t^3) \\ &= h(0) \int x \, dF_t(x) \;+\; \phi \int x^3 \, dF_t(x) + o(\sigma_t^3) \\ &= h(0)\, x_t \;+\; \phi \int x^3 \, dF_t(x) + o(\sigma_t^3),\end{aligned}$$

where the last line uses the definition of the average gap, $x_t \equiv \int x \, dF_t(x)$. By the moment expansion identity, we decompose the non-central third moment into its mean, variance, and (non-standardized) skewness, $\int x^3 \, dF_t(x) = x_t^3 + 3\, x_t\, \sigma_t^2 + \psi_t$. In the next theorem, we substitute and collect terms, obtaining inflation as a projection onto the first three moments of the price-gap distribution alone.

Theorem 1 (Odd approximation of inflation)**.** *In a stationary environment, inflation satisfies:*

$$\begin{aligned}\pi_t &= \left(h(0) + 3\phi\sigma_t^2\right) x_t + \phi\left(x_t^3 + \psi_t\right) + o(\sigma_t^3) \\ &= \underbrace{\bar{h}\, x_t}_{\text{“Calvo term”}} + \underbrace{(h_t - \bar{h})\, x_t}_{\text{Extensive margin}} + \underbrace{2\phi\sigma_t^2\, x_t \;+\; \phi\, \psi_t}_{\text{Selection}} + o(\sigma_t^3)\end{aligned} \tag{17}$$

Let $\boldsymbol{x}^t := (x_0, \ldots, x_t)$ be the history of the mean gap up to time t. For every admissible $\boldsymbol{x}^t$, inflation is an odd function:

$$\pi_t\left[-\boldsymbol{x}^t\right] = -\pi_t\left[\boldsymbol{x}^t\right].$$

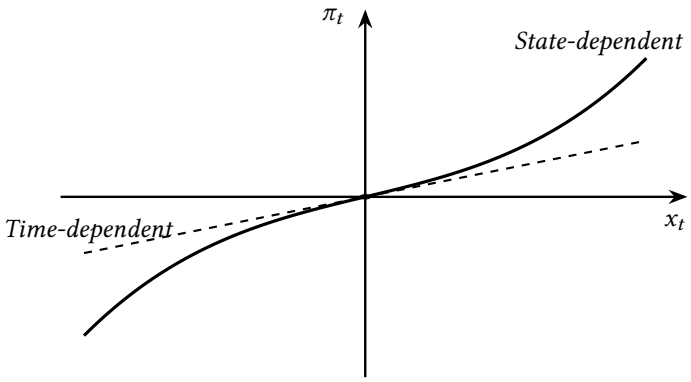
Proof. See Appendix C.3. □

The second line of Equation (17) shows that, like the general formula in Equation (7), the odd approximation decomposes inflation into a Calvo term, an extensive margin, and a selection effect capturing the cross-sectional covariance between hazards and gaps, $\mathbb{C}\mathrm{ov}_{F_t}\left(h(x), x\right) = 2\phi\sigma_t^2\, x_t + \phi\, \psi_t$. The last two are zero when the frequency does not move ($\phi = 0$).

Illustration with constant variance and skewness. To illustrate the key intuition behind Theorem 1, consider the simple case where the variance and skewness of the price gap distribution are held constant at their steady state values. Under this simplification, the aggregate frequency of price adjustment is given by a quadratic polynomial in the average (i.e., aggregate) price gap:

$$h_t \simeq \bar{h} + \alpha_2\, x_t^2,$$

Figure 1: An odd approximation of inflation



Notes. The solid line plots the odd approximation of inflation from Theorem 1, under the simplifying assumption of constant variance and skewness, $\beta_1 x_t + \beta_3 x_t^3$. The dashed line plots the time-dependent (Calvo) model, in which inflation is linear in the price gap with slope equal to the steady-state adjustment frequency, $\bar{h} x_t$.

and Equation (17) provides a characterization of aggregate inflation as a cubic (odd) polynomial in the average price gap:

$$\pi_t \simeq \beta_1 x_t + \beta_3 x_t^3,$$

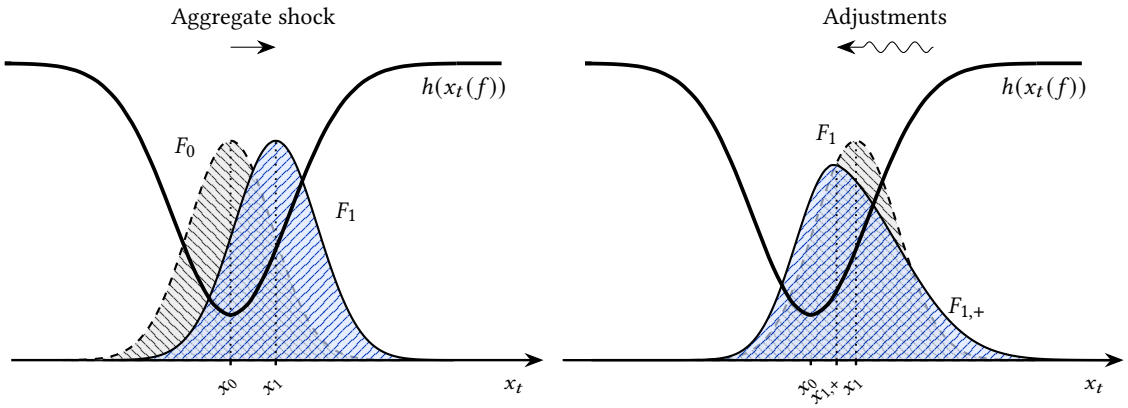
with coefficients $\beta_1 := h(0) + 3\phi\sigma_0^2$ and $\alpha_2 \equiv \beta_3 := \phi$.

The intuition behind the cubic term is straightforward: inflation is frequency times the gap and substituting the quadratic approximation for the frequency turns it into a cubic polynomial. In Figure 1, both schedules are approximately linear for small aggregate disturbances, with slopes that differ by the steady-state selection wedge, $\beta_1 - \bar{h} = 2\phi\sigma_0^2$. The odd approximation (solid line) then bends away from the time-dependent Calvo benchmark (dashed line) as $|x_t|$ grows under large shocks. This nonlinearity, driven by the endogenous response of adjustment frequency and selection, mirrors the evidence documented in the microdata by Gagliardone et al. (2026): an S-shaped relationship between price gaps and price changes, with a linear "Calvo region" at small gaps and a tail region where nonlinearities activate. It is also consistent with the near-equivalence between state-dependent and time-dependent pricing for small shocks established by Gertler and Leahy (2008), Alvarez et al. (2017), and Auclert et al. (2024).

3.3 The recursive general equilibrium of the approximate model: a finite state-space representation

Leveraging the odd approximation in Theorem 1, we now show how to characterize the general equilibrium of an economy with state-dependent pricing in recursive form and with a finite state-space representation that depends on aggregate variables only and can be recovered with

Figure 2: Dynamic effect of aggregate shocks on price gap distribution's location and shape



Notes. Dynamic effect of an aggregate shock of size $\delta > 0$ on the distribution of price gaps. The U-shaped black curve is the hazard rate $h(x_t(f))$. Left panel: on impact, the distribution shifts right by δ with its shape unchanged ($\sigma_0^2 = \sigma_1^2, \psi_0 = \psi_1$). Right panel: within the period, adjusting firms reset their price gap to zero, while non-adjusters carry the right tail forward, so the mean price gap falls ($x_{1,+} \downarrow$) and skewness rises ($\psi_{1,+} \uparrow$), as does the variance ($\sigma_{1,+}^2 \uparrow$) when the shock is large relative to the dispersion of price gaps.

standard solution methods for recursive models (e.g., Dynare).

Firm-level recursions and aggregation. We first write Equation (6) in recursive form:

$$\Lambda_t(f)\, p_t^\star(f) = p_t^o(f) + \beta\, \mathbb{E}_t \left\{ \big(1 - h_{t+1}(f)\big)\, \Lambda_{t+1}(f)\, p_{t+1}^\star(f) \right\}, \tag{18}$$

$$\Lambda_t(f) = 1 + \beta\, \mathbb{E}_t \left\{ \big(1 - h_{t+1}(f)\big)\, \Lambda_{t+1}(f) \right\}, \tag{19}$$

where $\Lambda_t(f)$ is the firm-specific effective discount factor (the denominator of Equation (6)).

We then aggregate the firm-level reset-price recursions in Equations (18)–(19) across firms, obtaining the following aggregate recursions, derived in Appendix B:

$$\Lambda_t\, p_t^\star = p_t^o + \beta\, \mathbb{E}_t \left\{ (1 - h_{t+1})\, \Lambda_{t+1}\, p_{t+1}^\star + C_{t+1}^\star \right\}, \tag{20}$$

$$\Lambda_t = 1 + \beta\, \mathbb{E}_t \left\{ (1 - h_{t+1})\, \Lambda_{t+1} \right\}. \tag{21}$$

Equations (18)–(19) are products of firm-level objects, so aggregation leaves a residual, $C_{t+1}^\star := \mathbb{C}\mathrm{ov}\left((1 - h_{t+1}(f))\Lambda_{t+1}(f), p_{t+1}^\star(f)\right)$: the covariance between a firm's effective discounting and its own reset price. It captures how the pricing behavior of the cohort of firms that has just reset—how often it will adjust again, and where it resets when it does—differs from that of the average firm.[9]

[9] Let $W(v) := \big(1 - h(v)\big)\Lambda(v)$ denote the product evaluated at price gap v, where $\Lambda(v)$ is the stationary solution of the firm-level recursion (19). The relevant gap is the one a firm that resets at t faces at $t + 1$: the change in the aggregate reset price plus its new idiosyncratic increment, $v = p_{t+1}^\star - p_t^\star + \epsilon_{t+1}(f)$. Up to third-order precision, $C_{t+1}^\star = J\,(p_{t+1}^\star - p_t^\star) + o(\eta^3)$, with $J := W''(0)\,\sigma_\epsilon^2$ (Proposition 3 in Appendix B).

Transitional dynamics of higher-order moments. In Section 3.2 we illustrated the odd approximation under the simplifying assumption that the higher-order moments of the price gap distribution—variance and skewness—remain constant under aggregate displacements. In fact, aggregate disturbances not only shift the distribution of price gaps but also affect their shape over time. As we show in Appendix B (Propositions 4 and 5), variance and skewness evolve with the aggregate displacement as:

$$\sigma^2_{t+1} - \sigma^2_0 = (1 - h(0))\big(\sigma^2_t - \sigma^2_0\big) + h(0)\big(1 - h(0)\big)\, x_t^2 \tag{22}$$

$$\psi_{t+1} - \psi_0 = (1 - h(0))\,\big(\psi_t - \psi_0\big) + 3h(0)\big(1 - h(0)\big)\, x_t\, \sigma^2_t + h(0)\big(1 - h(0)\big)\big(2h(0) - 1\big)\, x_t^3. \tag{23}$$

where $\psi_0 = 0$ under a symmetric price gap distribution. As the shock is absorbed, the mean, variance, and skewness return to their values in the stochastic steady state.

To illustrate, consider the economy in steady state, hit by a positive aggregate shock of size δ. On impact ($t = 1$), the shock shifts the price-gap distribution from F_0 to F_1 by $\delta = x_1 - x_0$. This translation preserves the distribution's shape, leaving both its variance and skewness unaffected: $\sigma^2_0 = \sigma^2_1$ and $\psi_0 = \psi_1$ (Figure 2, left panel). Within the period, firms that adjust pile up at a zero gap, which creates a mass point below the displaced mean. Firms that do not adjust carry the right tail forward. As a result, the mean price gap, after firms adjustments, falls ($x_{1,+} \downarrow$) while skewness rises ($\psi_{1,+} \uparrow$). The mass point has two opposite effects on the variance. It compresses the distribution, but it also separates adjusters from non-adjusters, who remain around the displaced mean. The second effect grows with the square of the displacement, so when the displacement is large relative to the dispersion in the stochastic steady state, $\sigma^2_{1,+}$ also increases (Figure 2, right panel).[10]

These effects persist over time. Each period starts from the variance and skewness left by the previous period's adjustment, so the distortion carries over and decays only gradually.

Recursive general equilibrium. The following proposition collects the system of equations that characterizes the general equilibrium of the approximate state-dependent model.

Proposition 2 (Recursive general equilibrium)**.** *The approximate state-dependent model is characterized by the following system of recursive equations in the endogenous variables* $\{y_t,\, i_t,\, r_t,\, p_t,\, p^\star_t,\, p^o_t,\, x_t,\, \pi_t,\, h_t,\, \Lambda_t,\, \sigma^2_t,\, \psi_t\}$*, the exogenous states* $\{u_t,\, a_t,\, \iota_t,\, b_t\}$*, and their innovations* $\{\varrho_t,\, \omega_t,\, \varsigma_t,\, \varkappa_t\}$*:*

[10] The post-adjustment variance is $\sigma^2_{1,+} = (1 - h(0))\,\sigma^2_1 + h(0)\big(1 - h(0)\big)\, x_1^2$, which exceeds σ^2_1 if and only if $|x_1| > \sigma_1/\sqrt{1 - h(0)}$. For smaller shocks, compression dominates and the variance falls.

Pricing block

$$x_t = p_t^\star - p_{t-1} \tag{NK–1}$$

$$p_t = \pi_t + p_{t-1} \tag{NK–2}$$

$$p_t^o = (1-\Xi)\left[\, \mu + (\gamma+\varphi)\, y_t - (1+\varphi)\, a_t + u_t \,\right] + p_t \tag{NK–3}$$

$$\pi_t = \left(h(0) + 3\phi\, \sigma_t^2\right) x_t + \phi\left(x_t^3 + \psi_t\right) \tag{NK–4}$$

$$h_t = h(0) + \phi\left(\sigma_t^2 + x_t^2\right) \tag{NK–5}$$

$$\Lambda_t\, p_t^\star = p_t^o + \beta\, \mathbb{E}_t\left\{(1-h_{t+1})\, \Lambda_{t+1}\, p_{t+1}^\star + C_{t+1}^\star\right\} \tag{NK–6}$$

$$\Lambda_t = 1 + \beta\, \mathbb{E}_t\left\{(1-h_{t+1})\, \Lambda_{t+1}\right\} \tag{NK–7}$$

$$\sigma_{t+1}^2 - \sigma_0^2 = (1-h(0))\left(\sigma_t^2 - \sigma_0^2\right) + h(0)\left(1-h(0)\right) x_t^2 \tag{NK–8}$$

$$\psi_{t+1} = (1-h(0))\, \psi_t + 3h(0)\left(1-h(0)\right) x_t\, \sigma_t^2 + h(0)\left(1-h(0)\right)\left(2h(0)-1\right) x_t^3 \tag{NK–9}$$

Demand block

$$y_t = \mathbb{E}_t\{y_{t+1}\} - \gamma^{-1}\left(r_t + \ln\beta\right) + b_t - \mathbb{E}_t\{b_{t+1}\} \tag{NK–10}$$

$$i_t = -\ln\beta + \phi_\pi\, \pi_t + \iota_t \tag{NK–11}$$

$$r_t = i_t - \mathbb{E}_t\{\pi_{t+1}\} \tag{NK–12}$$

Shocks

$$u_t = \rho_u\, u_{t-1} + \varrho_t \tag{NK–13}$$

$$a_t = \rho_a\, a_{t-1} + \omega_t \tag{NK–14}$$

$$\iota_t = \rho_\iota\, \iota_{t-1} + \varsigma_t \tag{NK–15}$$

$$b_t = \rho_b\, b_{t-1} + \varkappa_t \tag{NK–16}$$

Equations (NK–1)–(NK–9) capture the pricing block of the approximate model. Equations (NK–1) and (NK–2) define the aggregate price gap (Equation (4)) and the price level (Equation (7)). Equation (NK–3) is the aggregate counterpart of the static target price (Equation (3)), derived using the household's labor-supply condition and market clearing. Equations (NK–4) and (NK–5) are the projection of inflation and of the average frequency of price adjustment onto the moments of the price-gap distribution (Equations (17) and (15)). Equations (NK–6) and (NK–7) are the aggregate reset-price and continuation-value recursions (Equations (20) and (21)). Equations (NK–8) and (NK–9) are the recursive laws of motion of the pre-adjustment variance and skewness (Equations (22) and (23)), closing the system: they pin down σ_t^2 and ψ_t, which Equations (NK–4) and (NK–5) take as states.

Equations (NK–10)–(NK–12) constitute a standard New-Keynesian demand block. Recall that, for simplicity, we log-linearize this block to focus on inflation dynamics as the sole source of nonlinearities. Equations (NK–13)–(NK–16) are the dynamics of the shock processes.

Note that the Calvo model is the zero-curvature case $\phi = 0$: Equations (NK–4) and (NK–5) then collapse to $\pi_t = h(0)\, x_t$ and $h_t = h(0)$, and σ_t^2, ψ_t drop out of the pricing block entirely even though they continue to evolve; the block then reduces to the textbook New Keynesian model with a constant frequency of price adjustment.

The odd approximation in Theorem 1 and the recursive general equilibrium representation address the three challenges to solving state-dependent pricing models globally raised in Section 2.3. Together they deliver a *finite state-space representation* of inflation, with x_t, σ_t^2 and ψ_t standing in for the whole price gap distribution F_t—a reduction that preserves all the economic forces of state-dependent pricing. This state-space reduction follows the spirit of the approach in Krusell and Smith (1998). Unlike theirs, the mapping from cross-sectional moments to aggregates is in closed form and third-order accurate—second-order for the frequency—rather than guessed and verified numerically. Conceptually, the reduction also narrows the informational burden that rational expectations imposes on agents (Moll 2025): firms need only form correct beliefs about the three moments and their law of motion rather than information on their full distributions.

The nonlinear impulse-response functions. We now study the effects of aggregate shocks (e.g., cost-push, monetary, or TFP shocks) within our framework. Building on the representation in Theorem 1 and accounting for the endogenous dynamics of the price-gap distribution's location and shape, the following theorem characterizes the impulse responses of aggregate frequency of price adjustment and inflation to an aggregate shock as, respectively, an even (quadratic) and an *odd* (cubic) polynomial in the average price gap with deterministic coefficients.

Theorem 2 (Nonlinear impulse response functions)**.** *Consider an unanticipated aggregate shock of size $x_0 = \delta$. Let the first-order response of the aggregate gap be $x_{t+1} = \lambda x_t + o(\delta)$ for some $\lambda \neq 0$.*

The impulse response functions of the average frequency of price adjustment and aggregate inflation $t \geq 0$ periods after the shock are, to second and third order in δ, given by:

$$h_t = \bar{h} + \alpha_{2,t} x_t^2 + o(\delta^2), \tag{24}$$

$$\pi_t = \beta_{1,t} x_t + \beta_{3,t} x_t^3 + o(\delta^3). \tag{25}$$

Denoting by $h \equiv h(0)$, $H \equiv 1 - h(0)$, the coefficients of the polynomials (the slopes) are:

$$\alpha_{2,t} = \phi \left[1 + \frac{hH}{H - \lambda^2} \left\{ \left(\frac{H}{\lambda^2} \right)^t - 1 \right\} \right], \qquad \beta_{1,t} = h + 3\phi\sigma_0^2 \left[1 + \frac{hH}{H - \lambda} \left\{ \left(\frac{H}{\lambda} \right)^t - 1 \right\} \right],$$

$$\beta_{3,t} = \phi \left[1 + \frac{3hH}{H - \lambda^2} \left\{ \left(\frac{H}{\lambda^2} \right)^t - 1 \right\} + \frac{3h^2 H}{(H - \lambda^2)(1 - \lambda)} \left\{ \left(\frac{H}{\lambda^3} \right)^t - \left(\frac{H}{\lambda^2} \right)^t \right\} + \frac{hH}{H - \lambda^3} \left(2h - 1 - \frac{3hH}{H - \lambda^2} \right) \left\{ \left(\frac{H}{\lambda^3} \right)^t - 1 \right\} \right].$$

provided that the coefficients are finite.

Proof. See Appendix C.4. □

Formally, Theorem 2 follows from substituting the dynamics of variance and skewness as functions of the average price gap (Equations (22) and (23)) into the projection of inflation onto these variables (Equation (17)). As in the case of constant higher-order moments illustrated in Figure 1, Theorem 2 shows that the IRF of inflation is also an odd polynomial in the average price gap induced by the aggregate shock. In this case, however, the time-variation in the coefficients accounts for the response of higher-order moments of the price gap distribution along the shock's transition path, yielding an approximation that achieves third-order precision (second-order for the frequency).[11]

Importantly, the coefficients of the polynomial approximation are closed-form functions of the steady-state pricing primitives $(\phi, h(0), \sigma_0^2)$. They also depend on the equilibrium persistence of the average price gap, λ, an endogenous object pinned down, as a fixed point, by the general equilibrium of the model.[12] Neither input depends on the size of the aggregate displacement δ, so the coefficients can be calibrated or estimated independently of the magnitude of the shock used to identify them.

4 Transitional dynamics: global vs. approximate models

We now ask whether the approximation reproduces the transitional dynamics of the state-dependent economy. We consider two variants: the odd approximation of the previous section, and a "poor man's" alternative of the odd approximation featuring constant coefficients. The first derives its coefficients from the model's primitives. The second is a simplification that ignores movements in the higher-order moments of the price gap distribution. A log-linearized Calvo model serves as a benchmark, calibrated to the same average frequency of price adjustment and flat hazard ($\phi = 0$), to compare the transitional dynamics of the global and approximate models across shocks of very different size.

[11] These orders refer to the joint limit of Assumption *4c*, in which aggregate shocks and the dispersion of price gaps shrink together. Section 4 assesses the accuracy of the approximation at a given calibration in response to aggregate shocks of various magnitudes.

[12] The fixed point can be computed numerically, as we do in our code: λ is the decay rate of the first-order impulse response of x_t, which we obtain from the solution of the log-linearized model. When that response is not geometric, Theorem 3 in Appendix C.4 gives the coefficients for any first-order path $x_t = \tilde{a}_t \delta + o(\delta)$, of which Theorem 2 is the case $\tilde{a}_t = \lambda^t$.

4.1 Calibration

The general-equilibrium block features log utility ($\gamma = 1$), a unit inverse Frisch elasticity ($\varphi = 1$), an elasticity of substitution across varieties $\theta = 6$, a quarterly discount factor $\beta = 0.96^{1/4}$, and constant returns to scale in production, which under the constant-elasticity demand system sets $\Xi = 0$. The Taylor coefficient is $\phi_\pi = 1.5$, a balanced policy stance under which the aggregate price gap moves enough for the cubic term to be quantitatively relevant. The cost-push shock is the only disturbance active in this exercise, and follows an AR(1) process with persistence $\rho_u = 0.5$.

The calibration of the parameters governing the price block varies across models:

The global model—For the global menu-cost model, we calibrate the three parameters of the price block to replicate conventional features of the micro-level distribution of price changes—frequency and size—in the US. Menu costs are drawn from a uniform distribution with upper bound $\chi^{\max} = 0.346$. We set the idiosyncratic dispersion to $\sigma_\epsilon = 0.058$ and the free-adjustment probability to $h(0) = 0.15$. This calibration delivers a quarterly frequency of price adjustment of 0.250, at the lower end of the 0.25–0.31 range in the US micro data, a standard deviation of price changes of 11.6 percent, and implies that about forty percent of adjustments are driven by the menu cost rather than by a free opportunity.

The odd approximation—We calibrate the odd approximation directly to the primitives that govern firms' pricing in the global solution: the free-adjustment probability $h(0)$, the curvature of the hazard function at the optimum ϕ, and the stochastic variance of the price gap σ_0^2. Only $h(0)$ is a parameter we set, at the same value used to calibrate the global model; $\phi = 15.580$ and $\sigma_0^2 = 0.007421$ are computed numerically from the global model's stochastic steady-state distribution. Given the three, the variance σ_t^2 and the skewness ψ_t evolve along the transition according to Equations (22) and (23), so the mapping from the average price gap to inflation and to the adjustment frequency changes as the economy moves away from the stochastic steady-state. No moment of the global model's impulse responses enters this calibration, so Figure 3 compares the global solution with a prediction of the approximation rather than with a fit to it.

A poor man's simplification—As discussed in Section 3.2, if we ignore the response of the variance and skewness to an aggregate disturbance, the frequency and inflation simplify to second- and third-order polynomials with constant coefficients, respectively. Under this simplification, the

pricing block of Proposition 2 reduces to:

$$x_t = p_t^\star - p_{t-1} \tag{NK'–1}$$

$$p_t = \pi_t + p_{t-1} \tag{NK'–2}$$

$$p_t^o = (1-\Xi)\left[\,\mu + (\gamma+\varphi)\,y_t - (1+\varphi)\,a_t + u_t\,\right] + p_t \tag{NK'–3}$$

$$\pi_t = \beta_1\,x_t + \beta_3 x_t^3 \tag{NK'–4}$$

$$h_t = \bar{h} + \alpha_2 x_t^2 \tag{NK'–5}$$

$$\Lambda_t\,p_t^\star = p_t^o + \beta\,\mathbb{E}_t\left\{(1-h_{t+1})\,\Lambda_{t+1}\,p_{t+1}^\star + C_{t+1}^\star\right\} \tag{NK'–6}$$

$$\Lambda_t = 1 + \beta\,\mathbb{E}_t\left\{(1-h_{t+1})\,\Lambda_{t+1}\right\} \tag{NK'–7}$$

Borrowing from Jarociński and Karadi (2020), we dub the constant-coefficient system the poor man's simplification. While this restriction does not capture the dynamics of the higher-order moments implied by the moment recursions, it is computationally lighter and easier to implement. It also allows $\{\bar{h}, \alpha_2, \beta_1, \beta_3\}$ to be calibrated or estimated more flexibly, without taking a stance on the primitives that discipline the pricing problem. Leveraging this flexibility, we calibrate the coefficients directly to the transitional dynamics produced by the global model: we set the steady-state frequency $\bar{h}$ to the global model's own, and the remaining coefficients are chosen to match the global model's impulse responses of inflation, output, and the frequency of price adjustment to cost-push shocks of different sizes (the same magnitudes the panel below displays).[13]

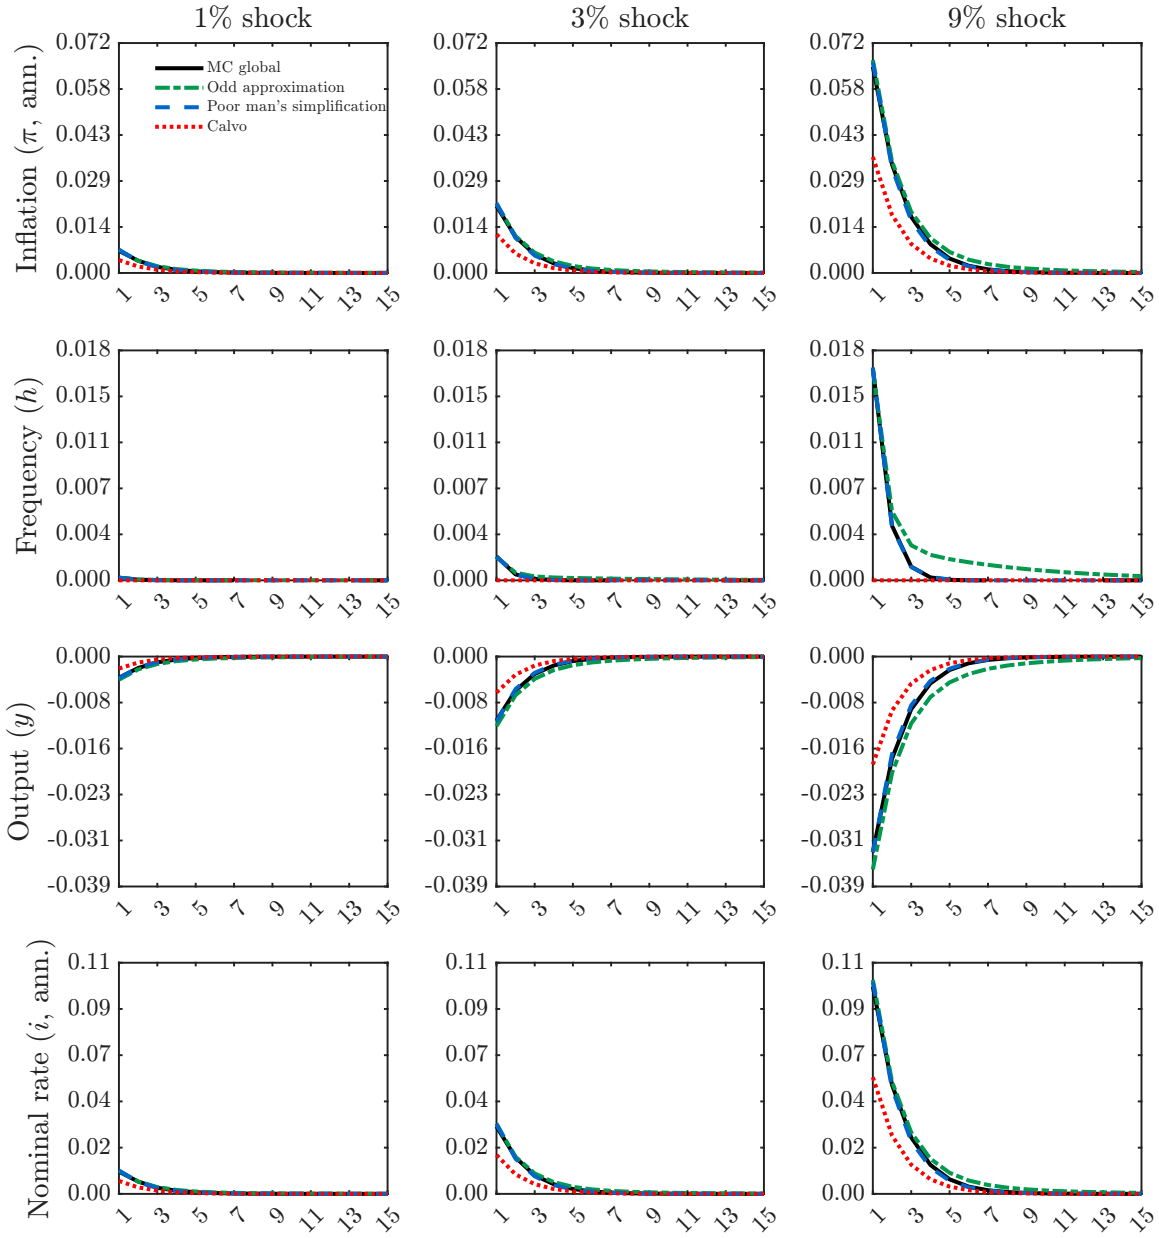
4.2 Comparison of transitional dynamics

We recover all models' impulse-response functions as deterministic perfect-foresight problems for a single unanticipated cost-push shock, so that all sets of responses are computed on an identical footing. Figure 3 reports the impulse responses of the four models to one-time, unanticipated cost-push shocks of $u_0 = \{1, 3, 9\}$ percent. As a benchmark, we compute the same transitional dynamics for a log-linearized Calvo model. In our framework the Calvo economy is the approximate model with zero hazard curvature: the cubic and the state-dependent frequency terms vanish, leaving a linear Phillips curve and a constant frequency of adjustment. We set that constant frequency equal to the menu-cost model's, so that the state- and time-dependent economies share the same frequency of price adjustment in the absence of aggregate disturbances.

As the shock grows, the menu-cost response becomes nonlinear (black solid lines), while

[13] The fit minimizes the sum, over inflation, output, and the adjustment frequency with equal weights, of squared deviations from the global model's paths over the first twelve quarters, each path normalized by its root mean square, averaged across the three shock sizes.

Figure 3: Impulse responses of the NK model to cost-push shocks of varying sizes



Notes. Impulse-response functions of inflation (π), the average frequency of price adjustment (h), real output (y), and the nominal interest rate (i) in the New Keynesian model to (quarterly) cost-push shocks of different magnitudes ($\{1, 3, 9\}$ percent) under the interest-rate rule $\phi_\pi = 1.5$. Black solid lines are the globally solved menu-cost model, green dash-dotted lines the odd approximation, blue dashed lines the poor man's simplification, and red dotted lines the Calvo benchmark.

the Calvo response remains proportional by construction (red dotted lines). The nonlinearity is clearest on the extensive margin. Each tripling of the shock, from one to three and from three to nine percent, raises the impact response of the adjustment frequency about ninefold: the frequency

responds to the square of the shock, not to its size. Large shocks push firms into regions of the price-gap distribution where the hazard function is steep, a channel absent from the Calvo model, in which the frequency does not move. Inflation and output inherit this convexity.

Figure 3 shows that both the odd approximation (green dash-dotted lines) and its poor man's simplification (blue dashed lines) track the global solution closely at all three magnitudes. The accuracy across shock sizes is the counterpart of the IRFs shock-size invariance property discussed in Section 4 (Theorem 2). For small shocks the average price gap is small and the cubic term is negligible relative to the linear one, so inflation is effectively linear in the gap, as in Calvo. For large shocks the cubic term becomes quantitatively relevant, and the odd approximation still reproduces the global solution's impact response, including the faster price-level pass-through that an endogenously higher adjustment frequency delivers. Its frequency response, however, decays more slowly than the global model's after very large aggregate shocks. This is due to the dynamic response of the variance of price gaps, which—at our calibration—displays stickier dynamics compared to the global model as it reverts to its steady-state value at a fixed rate.[14]

The poor man's simplification (blue dashed lines) matches the global model almost exactly. This results from its calibration, which chooses the three coefficients to match the global model's impulse responses, so its transition path is a fit. The fitted coefficients absorb whatever error the constant-coefficient restriction creates along the transition, including the movement in the variance and the skewness that the restriction rules out. We ask something much more demanding of the odd approximation. Its coefficients come from the global model's steady-state primitives alone. Its transition path is therefore a prediction.

5 Historical decomposition of US and EA data

We now bring the model to US and euro-area data. We estimate the pricing block using observable moments, solve the model to third order, and employ a nonlinear smoother to recover structural shocks that rationalize the data through the lens of the model. We show that the recovered shocks provide an economically plausible account of inflation dynamics in both normal times and during large aggregate disturbances. The model also reproduces the observed timing of the increase in the adjustment frequency during the 2021–22 surge, while leaving that frequency broadly stable

[14]The global model's adjustment frequency returns to its steady state within five quarters, while the odd approximation's tail persists for many more quarters, at the rate $(1 - h(0))$ implied by the variance recursion (22) once the price gap itself has decayed. The variance reverts at this fixed rate regardless of the path x_t actually took, so once the gap has died down the odd approximation's tail is pinned by a single calibrated constant. The global model's selection dynamics, by contrast, are not bound to any fixed rate and can clear faster/slower depending on the location of the distribution. Cumulated over the first fifteen quarters, the odd approximation's frequency response is more than fifty percent larger than the global model's.

in normal times. Relative to the standard order-1 Kalman smoother applied to the linearized version of the same model, the nonlinear procedure recovers shocks that are highly correlated with, and comparable in magnitude to, those obtained from the linear model. However, it attributes less of the inflation surge to exogenous cost-push shocks—a cost-push compression—because nonlinearities embedded in the pricing block endogenously amplify inflation through frequency and selection, especially after large shocks.

5.1 Data, parametrization, and model solution

Data. We require the model to match three quarterly observables for 2000Q1–2025Q1: year-over-year inflation, year-over-year real output growth, and the short-term policy rate. For the US these are CPI inflation, real GDP growth, and the effective federal funds rate, respectively. For the euro area, they are HICP inflation, real GDP growth, and the overnight interbank rate. We demean the three observables and take a perturbation around the model's zero-inflation-gap steady state. At the effective lower bound, the policy rate stays at its floor whatever the policy stance, so we treat it as missing in those quarters (Appendix D).

In addition, we use data on the frequency of price adjustment, based on regular, non-sale prices consistent with the model's adjustment hazard, to discipline the nonlinear parameters of the pricing block and validate the model-implied frequency path externally. For the United States, we use the expenditure-weighted median frequency from Cotton and Garga (2026), smoothed by a twelve-month trailing average. For the euro area, the expenditure-weighted mean frequency from Gautier et al. (2026).

The US frequency series appears to be less informative about the surge than the euro-area one. It is available only as a twelve-month trailing average, which smooths and delays the peak, and its pandemic-era level is likely influenced by sampling noise and a shift in the sample toward online retailers. At the height of the surge, the US frequency rose above 0.5, against a peak of about 0.32 in the euro area. That level exceeds even the US peak of the 1970s, when inflation was much higher. We account for these features in the estimation, by targeting a quarterly peak with an uncertainty band, and in the validation, by applying the same trailing average to the model (Figure 5, panel c). The US estimates are correspondingly less precise than those of the euro area.

Calibration and estimation. We calibrate the general-equilibrium and policy blocks to conventional values, reported in the top panel of Table 1: log utility ($\gamma = 1$), an inverse Frisch elasticity $\varphi = 2$, a quarterly discount factor $\beta = 0.99$, and a Taylor rule with $\phi_\pi = 1.50$.

The pricing block has four primitives: the free-adjustment probability $h(0)$, the curvature of the hazard function ϕ, the variance of price gaps in the stochastic steady state σ_0^2, and the loading

on competitors' prices Ξ. We calibrate $\sigma_0 = 0.026$ (Gagliardone et al. 2026), assume constant returns to scale at the firm level, and calibrate strategic complementarities as in Gagliardone et al. (2026), yielding a target loading on competitors' prices $\Xi = 0.50$. We then estimate $h(0)$ and ϕ by the simulated method of moments (SMM), matching two aggregate time-series moments: (i) the frequency of price adjustment before the pandemic (its 2010Q1–2019Q4 mean); and (ii) the peak frequency during the 2021–22 surge. Given σ_0^2, the pre-pandemic frequency ties $h(0)$ to ϕ, and the surge peak pins down ϕ. (Appendix D further describes the estimation procedure and inference.) For the euro area, we target the peak of the seasonally adjusted series, $h^* = 0.316$. Mindful of possible measurement issues described above, we adopt a conservative target for the United States, $h^* = 0.35$, which errs toward a lower state dependence.

Each evaluation of the SMM objective consists of two steps. In the first step, given a candidate for $h(0)$ and ϕ, we linearize the model and estimate the parameters that control the persistence and measurement-error variances of the cost-push, monetary, and demand shocks by Kalman maximum a posteriori (MAP). The Kalman filter delivers the likelihood, which we combine with the priors reported in Table A.1, and we use the posterior mode. The prior on the persistence of the cost-push shock is tight, and we fix the persistence of the monetary shock, for the reasons discussed below. This step is repeated at every candidate because the first-order slope of the Phillips curve, $h(0) + 3\phi\sigma_0^2$, moves with both estimated primitives. In the second step, given the candidate primitives and the estimated shock parameters, we solve the approximate state-dependent model at third order. Finally, we run the particle smoother and compute the model counterparts of the two frequency moments from the smoothed frequency path, $\mathbb{E}[h_t \mid Y_{1:T}]$. Note that, at first order, the frequency is constant at $\bar{h}$. So we compute the frequency moments from the third-order solution, which also carries the cubic terms of inflation.

Results. Table 1 summarizes the calibrated and estimated parameters for both economies. The estimated measurement-error standard deviations are reported in Table A.1. The persistence of the monetary shock requires separate treatment. In the simple three-equation New Keynesian model, a persistent monetary innovation generates an anticipated path of future interest rates that has implausibly large effects on current activity—the well-known forward-guidance puzzle (Del Negro et al. 2023; McKay et al. 2016). As a result, the likelihood pushes the estimated persistence toward the upper bound of its admissible range. We therefore fix $\rho_\iota = 0.20$ for the United States and 0.30 for the euro area, values that generate economically plausible responses to monetary shocks.[15] We discipline the persistence of the cost-push shock, ρ_u, with a tight prior

[15]Raising the admissible upper bound for ρ_ι simply pulls the free optimum toward the new edge, which is a manifestation of the same non-identification pattern documented by Carrillo et al. (2007).

Table 1: Model parameters

		US	Euro area
General-equilibrium block			
Risk aversion	γ	1.00	1.00
Inverse Frisch elasticity	φ	2.00	2.00
Discount factor	β	0.99	0.99
Taylor rule, inflation coefficient	ϕ_π	1.50	1.50
Shock processes			
Cost-push shock SD	σ_ϱ	0.024	0.028
Monetary shock SD	σ_ς	0.007	0.006
Demand shock SD	$\sigma_\varkappa$	0.014	0.019
Cost-push persistence	ρ_u	0.820	0.779
Monetary persistence	ρ_ι	0.200	0.300
Demand persistence	ρ_b	0.457	0.407
ME SD, inflation obs. (pp)	σ_π^{me}	0.24	0.21
ME SD, output obs. (pp)	$\sigma_{\Delta y}^{me}$	0.54	0.67
ME SD, rate obs. (pp)	σ_i^{me}	0.32	0.31
Pricing primitives			
Strategic complementarity	Ξ	0.50	0.50
SD of the price gap	σ_0	0.026	0.026
Free-adjustment probability	$h(0)$	0.202 (0.046)	0.199 (0.007)
Curvature of the hazard function	ϕ	81.0 (53.5)	21.4 (4.6)

Notes. Model parameters. The shock-process parameters are order-1 posterior modes at the estimated pricing primitives (Table A.1). The pricing primitives $h(0)$ and ϕ are estimated by SMM, and Ξ and σ_0 are calibrated, as described in Section 5.1. Numbers in parentheses are standard errors: the sandwich covariance of the exactly identified SMM estimator, plus simulation error, mapped into $h(0)$ and ϕ by the delta method. The covariance of the data moments uses a Newey–West (HAC) estimator for the pre-pandemic mean frequency and, for the surge peak, the uncertainty of the target and of the frequency in the peak quarter. Standard errors condition on σ_0 and on the observed macro history (Appendix D).

centered at 0.70 (Appendix D).

Moving to the pricing block, the euro area has a lower steady-state adjustment frequency than the United States ($\bar{h} = 0.213$ against 0.257, with standard errors of 0.005 and 0.012). The difference comes from the curvature of the hazard function, while the free-adjustment probabilities are almost the same. In both economies, the surge peak points to a sizable curvature, which triggers nonlinearities during large-displacement episodes. The steady-state frequency is estimated precisely in both economies. The curvature is estimated precisely in the euro area and less so in the United States. As explained above, the uncertainty about the US frequency adds sampling uncertainty and contributes to generating larger standard errors for ϕ.

5.2 Historical shock decomposition

Model solution and shock recovery. At the estimated parameters, we perform historical shock decomposition exercises that recover two sets of structural innovations: one through the

lens of our odd approximation of the state-dependent pricing model and one through the lens of a time-dependent Calvo model. Both solutions are perturbations around the deterministic steady state, condition on the same sample, and use the same initial-state distribution. They differ in that the former filters the third-order solution of the model, while the latter filters its first-order approximation. Their comparison therefore isolates the role of the state-dependent pricing block in shock recovery and propagation.

At third order, frequency and selection enter the decision rules through quadratic and cubic terms, and the filtering distribution is no longer Gaussian. We use a tempered particle filter with $N = 50{,}000$ particles, average the results across $n = 8$ independent seeds (Fernández-Villaverde and Rubio-Ramírez 2007; Herbst and Schorfheide 2019), and apply a fixed-lag particle smoother in the spirit of Olsson et al. (2008).[16] The smoothed innovations are then propagated through the third-order decision rules. At first order, the model is linear and Gaussian, so we use a standard Kalman filter and smoother with diffuse initialization for the unit root in the price level (An and Schorfheide 2007; Durbin and Koopman 2012).

Linear and nonlinear shock histories. Figure 4 presents the time-domain decomposition of inflation, output growth, and the policy rate (third-order smoother's fitted paths) and the contributions of different shocks. Despite its simplicity, the nonlinear model recovers a sensible historical decomposition and reading of the 2021–22 inflation surge (Gagliardone and Gertler 2026; Hazell and Hobler 2025). In the United States, demand shocks, amplified by monetary accommodation, explain most of the surge; in the euro area, cost-push shocks also contribute substantially. Demand shocks, partly offset by reverting cost-push shocks, explain its persistence.

The comparison with the order-1 solution and its recovered shocks provides a useful benchmark for evaluating where the nonlinear model changes the economic interpretation of the data, and a sanity check for the odd approximation.[17] The first-order terms of the decision rule are the same in the first-order and third-order model (Schmitt-Grohé and Uribe 2004), so for most of the sample, when shocks and price gaps are small, the two filters should recover almost the same shocks. Indeed, the first-order Kalman smoother and the nonlinear particle filter recover a broadly similar history of shocks: the cost-push, monetary, and demand series are strongly correlated with those recovered by the order-3 smoother (0.79, 0.88, 0.89 in the United States; 0.74, 0.79, 0.89 in the euro area) and have comparable dispersion, with standard-deviation ratios to the order-3 series between 0.84 and 1.07.

[16] The particle filter infers states and shocks; it is not used to evaluate a likelihood for the primitives, which are estimated by matching moments. As a diagnostic, its seed-averaged log-likelihood is −467.1 in the United States and −466.0 in the euro area.

[17] Figure A.1 in Appendix D reports the historical decomposition implied by the first-order Kalman smoother.

Table 2: Shocks' contribution to inflation surge: order-1 vs. order-3 smoother

	US (CPI), 2021Q1–2023Q1			Euro area (HICP), 2021Q3–2023Q1		
	Kalman sm. (order 1)	PF sm. (order 3)	Δ	Kalman sm. (order 1)	PF sm. (order 3)	Δ
Cost-push	48.2%	6.6%	−41.6 ±12.6	60.9%	22.0%	−38.8 ±4.4
Monetary	20.9%	21.5%	+0.5 ±10.5	8.0%	9.9%	+1.9 ±3.1
Demand	29.6%	58.9%	+29.3 ±16.2	30.5%	48.5%	+18.0 ±4.0
Joint interaction	0.0%	11.5%	+11.5 ±5.0	0.0%	17.7%	+17.7 ±1.6
Data minus joint	1.2%	1.5%	+0.3 ±1.1	0.7%	1.9%	+1.2 ±1.8

Notes. Share of the surge in observed year-over-year inflation attributed to each component over the surge window W shaded in Figure 4: $S(k) = \sum_{t\in W} \text{contrib}_k(t) / \sum_{t\in W} \pi_t^{obs}$, where π_t^{obs} is inflation in deviation from its 2000–25 mean and contrib_k is the path of inflation when only shock k is fed through the decision rule of the model that recovered it: order 1 for the Kalman-smoothed innovations and order 3 for the particle-smoothed ones. Particle-smoother entries are means across seeds; ± values are cross-seed standard deviations. The Δ column, their difference, measures the move from the linear to the nonlinear model. The row "Joint interaction" measures the part of the path generated by all three shocks together that no single shock accounts for (zero under the order-1 rule). The row "Data minus joint" measures the gap between observed inflation and that path. The five shares sum to 100%.

Cost-push compression. The two smoothers, however, distribute inflation differently across shocks when large aggregate displacements occur. Fed through the order-3 decision rule, the shocks recovered by the linear smoother overshoot the recent 2021–22 inflation surge, by 39% in the United States and 70% in the euro area. This is because the nonlinear model amplifies large impulses, so it needs a smaller cost-push shock and attributes more of the surge to demand.

Table 2 reports the share of surge inflation attributed to different shocks by the first-order Kalman smoother and the third-order particle smoother, with each smoother's shocks fed through the decision rule of its own model.[18] The difference measures how the relative importance of different structural shocks moves from the linear to the nonlinear model. The cost-push share falls from 48.2 to 6.6 percent (−41.6 percentage points) in the United States and from 60.9 to 22.0 (−38.8 percentage points) in the euro area. While the monetary shares are essentially unchanged, demand absorbs about 70 percent of the cost-push decline in the United States and almost half in the euro area. The joint interaction of the shocks absorbs most of the rest.

When the economy is hit by large shocks, adjustment frequency and selection steepen the response of inflation to a given cost-push shock; the model consequently needs a smaller exogenous cost-push disturbance to match the observed surge. In other words, the linear model needs to resort to exogenous shocks to capture what is instead endogenous amplification due to

[18] We attribute surge inflation to the shocks one shock at a time. Under a first-order solution, the single-shock paths add up to the fitted series, so the residual is measurement error alone. At third order, shocks propagate through nonlinear decision rules. This implies that the residual captures both interactions between shocks, which cannot be assigned to any single disturbance, and the gap between the data and the path of inflation the model produces when all three smoothed shocks are simultaneously fed through the order-3 decision rules (Appendix D).

state-dependent pricing. The same mechanism explains the larger demand share. A given demand shock raises inflation by more in a high-inflation environment, because more firms reprice and those that do make larger price changes. Moreover, shocks that hit together widen price gaps further, which is why their joint effect exceeds the sum of their separate effects. In the linear model, by contrast, the effects of different shocks simply add up.

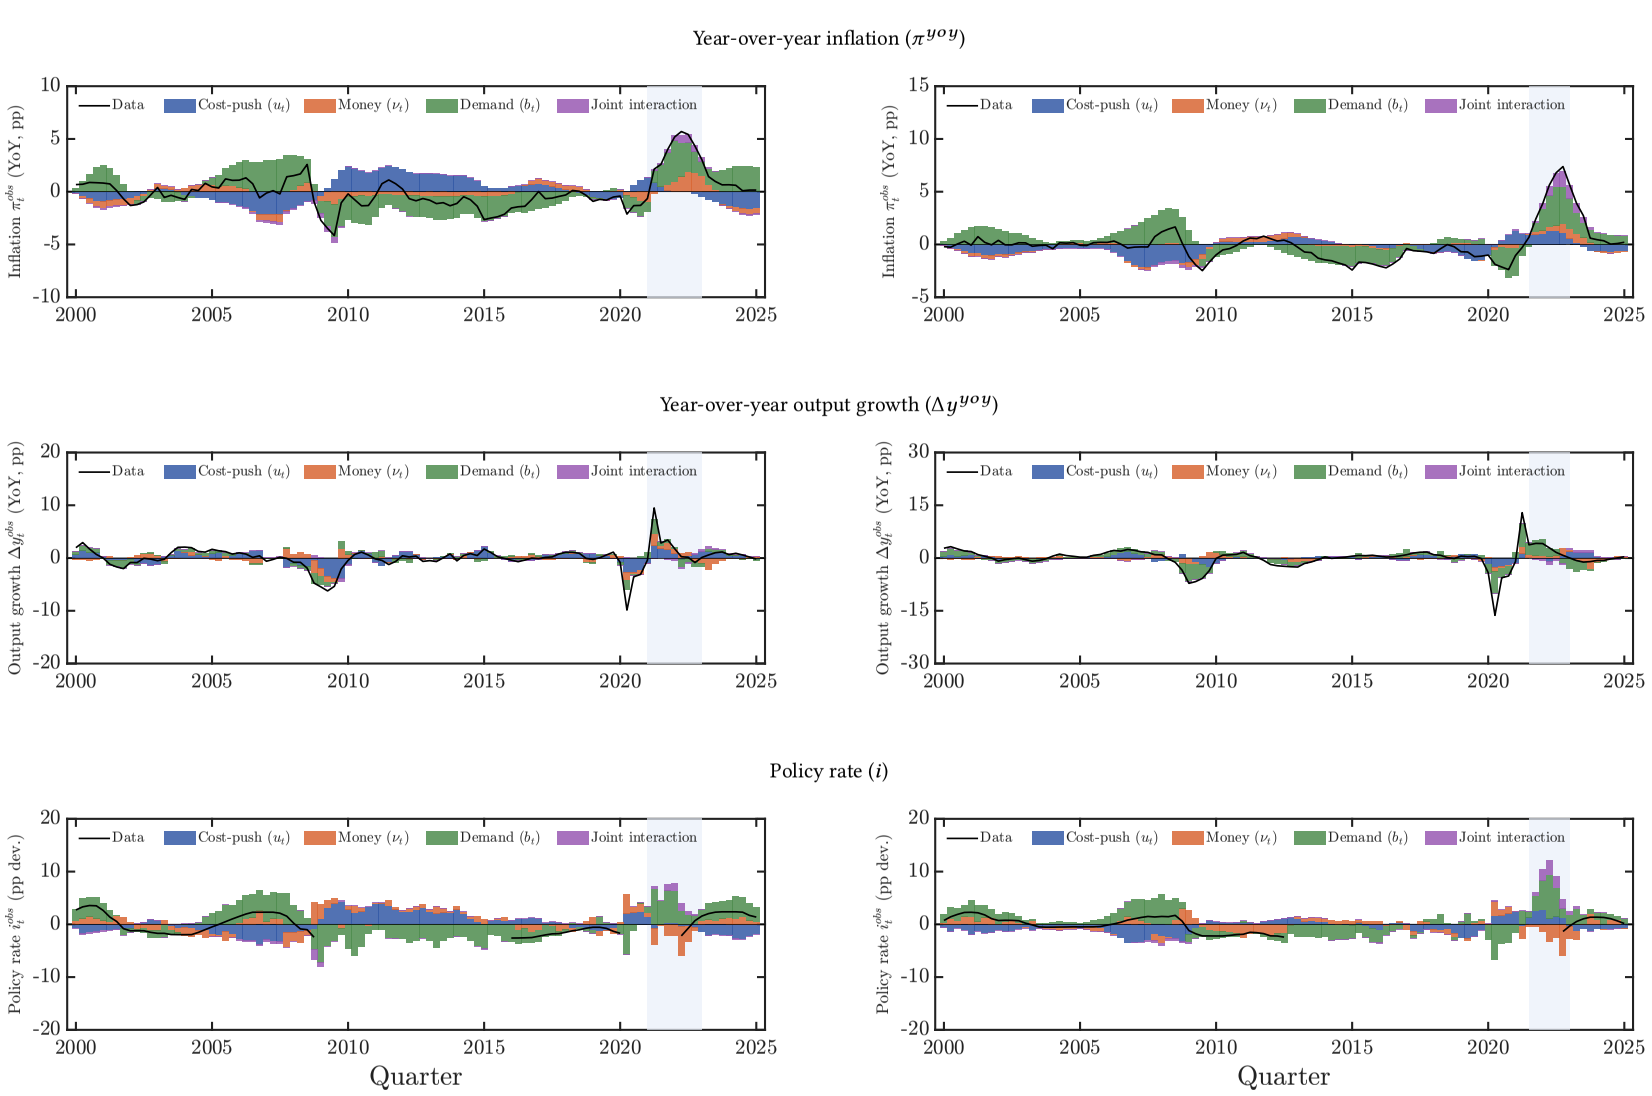
Inflation decomposition and the slope of the Phillips curve. Unlike the global model, the odd approximation provides us with a closed-form decomposition of inflation into its forces. It also gives us a sense of how the slope of the Phillips curve changes over the business cycle.

Applying Equation (17) to the smoothed moments of the price gap distribution, Panel a of Figure 5 splits model-implied inflation into linear (Calvo), frequency, and selection terms. This exercise highlights how, in normal times, the frequency term is small in both economies and inflation is accounted for by the Calvo and selection terms. Selection matters even then, particularly in the United States, through two terms that are linear in price gaps: the steady-state selection wedge, $2\phi\sigma_0^2 x_t$, and the skewness $\phi\psi_t$ that past gaps leave in the distribution. During the 2021–22 surge, the frequency and selection components widen sharply and together account for about half of peak model-implied inflation in the US and about two-fifths in the euro area.

Panel b shows this mechanism through the lens of the (nonlinear) Phillips curve. It plots the expected slope of the Phillips curve, $h(0) + 3\phi(\sigma_t^2 + x_t^2)$, averaged over the smoothed cross-section and across seeds (solid line), against its steady-state value, $h(0) + 3\phi\sigma_0^2$ (dashed line).[19] A linear model with either economy's steady-state slope captures inflation dynamics in normal times, when cost-price dynamics are quasi-linear (Gagliardone et al. 2025), but misses the surge and attributes too much of it to exogenous shocks, when state-dependence activates the cubic mechanism (Gagliardone et al. 2026). Our estimates suggest that the slope of the Phillips curve rose by about three-quarters in the United States and more than doubled in the euro area during the post-pandemic surge, and nearly doubled in the United States after the global financial crisis.

[19] Note that there is a wedge between the models' Phillips curves which is present even at small shock scales. This is because steady-state selection raises the menu-cost model's linear Phillips-curve slope above the slope of a Calvo model, even if the two are matched on the same average adjustment frequency (Theorem 1).

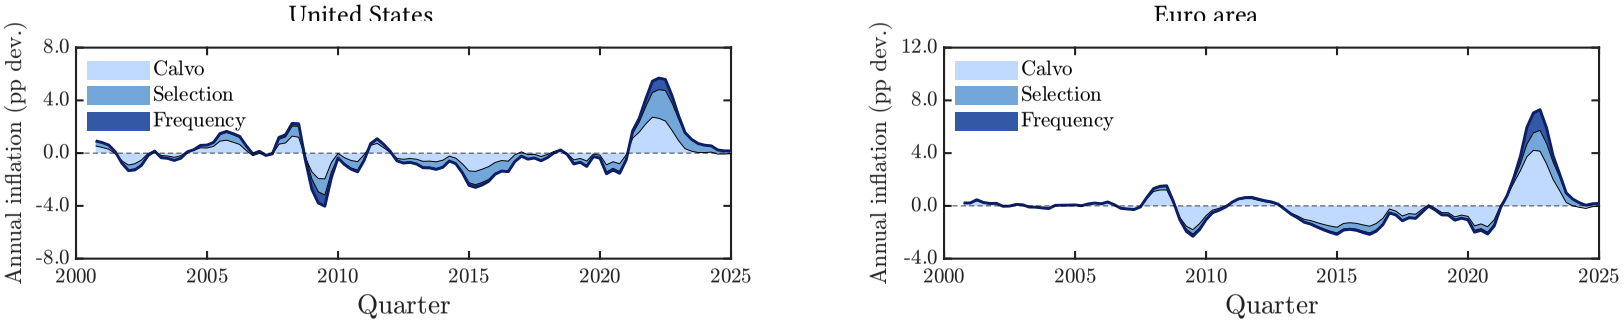
Figure 4: Historical shock decomposition of inflation, output, and the policy rate



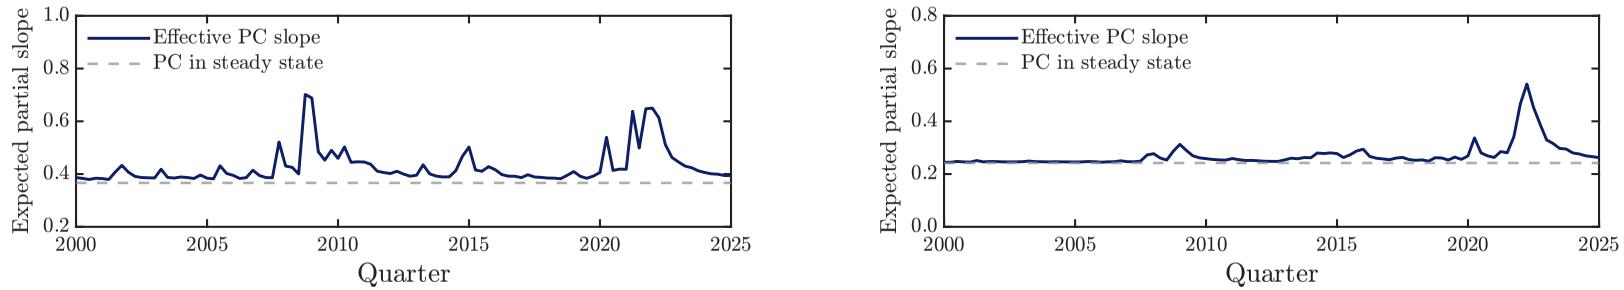
Notes. Historical decomposition of the nonlinear smoother's fitted series for year-over-year inflation, year-over-year output growth, and the policy rate (top to bottom), for the United States (left column) and the euro area (right column). Stacked bars are the contributions of the cost-push (u_t, blue), monetary (ι_t, orange), and demand (b_t, green) shocks; the black line is the observed series; the shaded band marks the 2021–22 surge. Purple bars are the joint interaction of the three shocks, and light-grey bars the gap between the data and the fitted series. The shocks are the particle-filter smoother innovations of each seed, propagated through the order-3 decision rule; the bars average the resulting contributions across the eight seeds. Each series is plotted in deviation from its sample mean.

Figure 5: Model-implied inflation dynamics: decomposition, effective Phillips-curve slope, and adjustment frequency

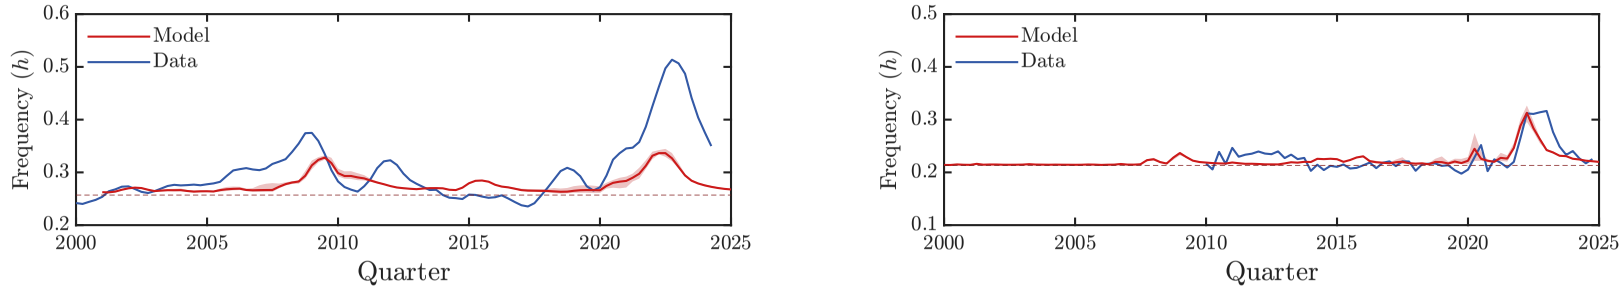
Panel a: Decomposition of model-implied inflation during the surge



Panel b: The Phillips curve's effective slope, linear vs. nonlinear



Panel c: Model-implied frequency of price adjustment vs. micro data



Notes. Panel a: decomposition of model-implied annual inflation, computed from the order-3 smoothed cross-section of price gaps for one seed, into the Calvo, frequency, and selection terms of Equation (17). Inflation is plotted in deviation from its sample mean. Panel b: the slope of the Phillips curve (solid), against its steady-state value (dashed). Panel c: quarterly frequency of price adjustment, smoother-implied (red, mean across seeds, 5–95% band across seeds) against the micro series (blue); dashed line is the estimated steady-state frequency. The euro-area micro series is seasonally adjusted (Appendix D). The US micro series is the monthly expenditure-weighted median frequency from Cotton and Garga (2026), published as a twelve-month trailing average and compounded to quarters. We apply the same construction to the smoothed US frequency, so the US model series starts in 2001Q1.

The frequency of price adjustment. The model reproduces the broad timing of the frequency response (Panel c of Figure 5)—rising during the surge and remaining flat in normal times. (Recall that we did not target the time-series path of adjustment frequency, and only used the pre-surge level and surge peak to estimate the pricing primitives.) The mean correlation across seeds between the model-implied and observed series is 0.69 in the United States and 0.62 in the euro area. In the euro area, the model also reproduces the short-lived rise in the frequency in 2020, peaking one quarter before the data. For the United States, the model matches the rise during the 2008–09 recession. It also captures the rise during the 2021–22 surge, but not its full swing: the model peaks at 0.34 in mid-2022, while the observed series keeps rising to 0.51 in 2022Q4, a level we attribute partly to sample composition (Section 5.1). As noted above, that spike is an outlier relative both to the US experience of the 1970s and to the one observed in the euro area during the same surge. It has been a challenge in the literature to explain this movement in the frequency with a standard menu-cost model (see Blanco et al. 2024, and Kaplan and Miyahara 2026).

6 Policy trade-offs under state-dependent pricing

We now illustrate how our odd approximation of a state-dependent pricing framework can be easily leveraged to produce policy counterfactuals, adding insight into the role of nonlinear inflation dynamics. In our specific application, we turn to a classic policy question: How aggressively should a central bank respond to a large but transitory inflationary shock that poses a trade-off between output and inflation stabilization?

A common answer is to "look through" it. Provided expectations remain anchored, the shock will reverse on its own, and the output losses from tightening would outweigh the gains from lower inflation (Bernanke and Mishkin 1997; Mishkin 2007; Bernanke 2007). By contrast, the textbook New Keynesian model—with a quadratic loss function and a linear Phillips curve—calls for a weighted stabilization of inflation and the output gap following cost-push shocks (Clarida et al. 1999; Woodford 2011). Crucially, the resulting policy response scales *linearly* with shock size, and the inflation–output trade-off is essentially invariant to it. Recent work shows that state-dependent pricing changes this prescription: larger cost-push shocks warrant a more forceful policy response because state dependence makes the trade-off more favorable (Karadi et al. 2025).

In this section, we use our approximate state-dependent model to take these prescriptions to the data. First, we show that the cost of looking through large aggregate shocks is greater than implied by the linear model, and increasingly so as the shock grows. Second, consistent with the normative results of Karadi et al. (2025), we show that the sacrifice ratio associated with a tighter policy stance would have been lower during the recent inflation surge.

For both exercises, we calibrate the US and euro-area model economies as in Section 5.1 (Table 1) and feed each model the order-3 particle-smoothed innovations recovered in Section 5, starting from the steady state in 2000Q1. Until 2020Q4, both exercises feature our baseline Taylor rule parametrization ($\phi_\pi = 1.5$) and receive all three recovered innovations, so the menu-cost economy tracks the history estimated in Section 5, including the policy response to the pandemic. In 2021Q1, we impose an unexpected and permanent switch in the policy stance from the baseline to a counterfactual rule: monetary accommodation ($\phi_\pi = 1.05$) in the first exercise and monetary tightening ($\phi_\pi = 2$) in the second. We date the switch in 2021Q1, when inflation started rising but before central banks started raising their policy rates (the Federal Reserve in March 2022 and the ECB in July 2022). These counterfactual policies draw a stark comparison between looking through the surge and tightening against it. Both are rule counterfactuals rather than the realized policy paths.[20] From the switch on, we shut down monetary-policy shocks because the exercises compare hypothetical policy rules: the relevant question is how the recovered supply-and-demand disturbances propagate under each rule.[21] We simulate the third-order solution with pruning (Andreasen et al., 2018) and display the 2018–2025 window.

Looking through large shocks. The first exercise considers an accommodative central bank ($\phi_\pi = 1.05$) facing the large post-pandemic shocks (Figure 6, panel a). We treat the menu-cost model as the data-generating process and compare it with a frequency-matched Calvo benchmark. Under this stance, the recovered supply-and-demand disturbances generate a menu-cost inflation surge that peaks at 6.5 percentage points above its sample mean in 2022, about 3.2 times the response of the Calvo model to the same shocks (10.7 percentage points and 2.4 times in the euro area). These effects are monotone in the policy parameter. With $\phi_\pi = 1.1$, 1.2, and 1.3, the ratio falls to 2.8, 2.4, and 2.2 in the United States and to 2.2, 1.9, and 1.7 in the euro area.

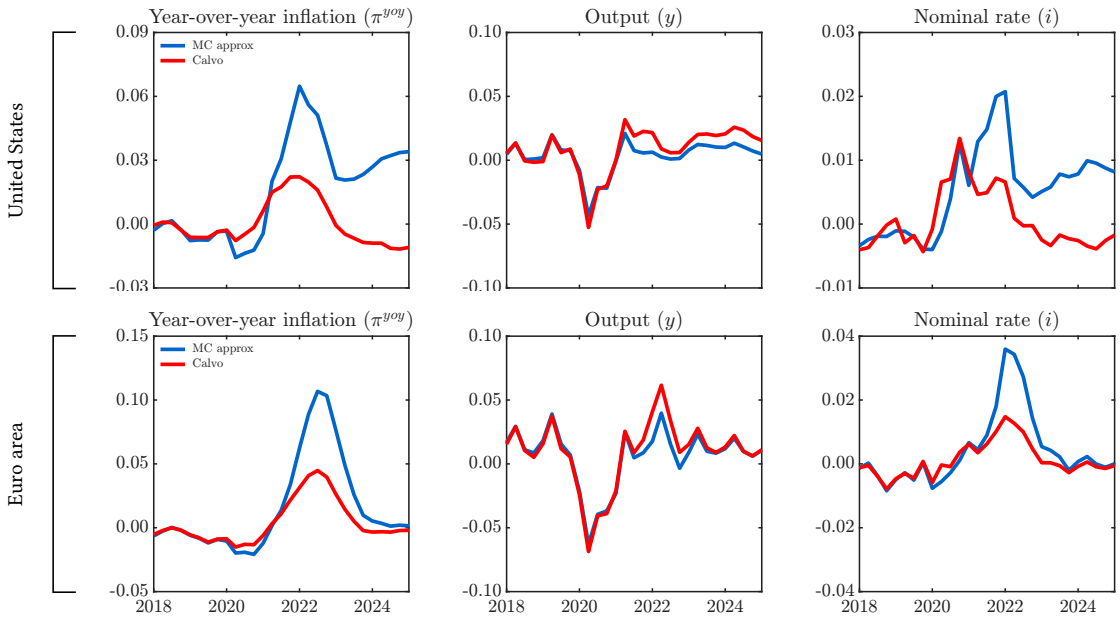
The difference between the two models is, again, the feedback loop made explicit by Equation (17). The shocks move the distribution of price gaps onto the steep flanks of the adjustment hazard, raising the adjustment frequency and strengthening selection. Repricing then raises the aggregate price level, moving other firms' real prices toward their adjustment thresholds and inducing further repricing. The frequency part of this loop is the inflation-accelerator mechanism of Blanco et al. (2026), who abstract from selection, and is close to the pricing cascades of Ghassibe and Nakov (2025). Selection adds a second meaningful feedback in our model. After

[20] The comparison between the two models does not hinge on the date. The message is unchanged if we move the policy switch to 2020Q1 or to 2021Q3.

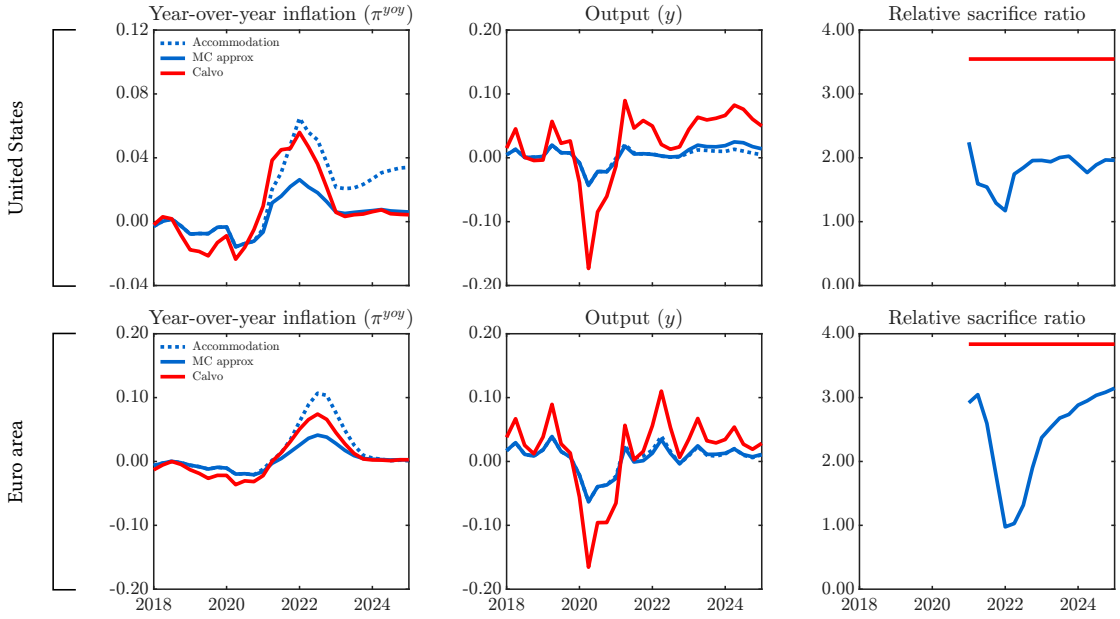
[21] Including monetary-policy innovations would answer a different question—how the full recovered shock history would have propagated—rather than isolate the consequences of alternative policy rules. Because monetary innovations were accommodative in 2021–22, including them would further activate the cubic and widen the look-through gap.

Figure 6: Looking through large inflationary shocks and sacrifice ratios under different monetary policy regimes

Panel a: Nominal and real effects of the recovered shocks under accommodative monetary policy



Panel b: Tightening monetary policy and relative sacrifice ratios



Notes. Panel a presents the responses of aggregate year-over-year inflation (π^{yoy}), real aggregate output (y), and the nominal rate (i) to the order-3 smoothed innovations recovered in Section 5 (monetary innovations excluded after the switch), 2018–2025, under the baseline Taylor rule ($\phi_\pi = 1.5$) until 2020Q4 and monetary accommodation ($\phi_\pi = 1.05$) from 2021Q1, for the United States (top row) and the euro area (bottom row, Table 1 estimates); red lines are the log-linearized (Calvo) model, blue lines the approximate menu-cost model. Panel b repeats this with tight policy ($\phi_\pi = 2$) from 2021Q1; dotted lines reproduce the accommodative responses for reference. Recall that inflation and the nominal rate are plotted in deviation from their sample means, and output in deviation from its trend. The right panel in each row shows, from 2021Q1, the slope-implied relative sacrifice ratio of the menu-cost model (blue) and the Calvo sacrifice ratio (red). Plotted paths are the pointwise median across seeds.

the switch, output moves similarly in both models: common demand innovations dominate the output path, while the nonlinearity is concentrated in inflation.

The sacrifice ratio. Should the central bank then respond more aggressively, and at what cost? Figure 6, panel b, repeats the experiment with a switch to a tight rule ($\phi_\pi = 2$) in 2021Q1 and compares it with the accommodative case discussed above ($\phi_\pi = 1.05$). The two paths coincide until the switch. Because policymakers observe inflation rather than the shocks behind it, the relevant cross-model comparison equates peak inflation. We therefore rescale the shocks fed to the Calvo model so that its peak inflation under accommodation matches that of the menu-cost economy, and compare the output cost of the resulting disinflation.[22]

At the estimated values, tightening lowers peak year-over-year inflation by 3.9 percentage points (from 6.5 to 2.6 percentage points above its sample mean) in the state-dependent model but by only about 0.9 points in the matched Calvo economy (from 6.5 to 5.6). We summarize this trade-off with the sacrifice ratio, the output forgone per point of inflation reduction when policy switches from accommodation to tightening. For the Calvo economy, we compute it over the whole surge, as total output forgone divided by total inflation reduction:

$$\mathrm{SR}^{\mathrm{Calvo}} = \frac{\sum_{t\in\mathcal{S}} \left(y_t^{accom} - y_t^{tight}\right)}{\sum_{t\in\mathcal{S}} \left(\pi_t^{accom} - \pi_t^{tight}\right)},$$

where $\mathcal{S}$ collects the quarters of the inflation surge: 2021Q1–2023Q1 in the United States and 2021Q3–2023Q1 in the euro area. Recall that the demand block is log-linear in both the state-dependent and time-dependent economies so that they differ only in the pricing block. Therefore, holding the output cost of a given fall in the price gap equal across them, the output cost per point of disinflation is inversely proportional to the slope of the Phillips curve. Using this property and the fact that the slope of the Phillips curve does not vary in the Calvo model, we obtain the sacrifice ratio for the state-dependent economy by scaling the Calvo ratio by the ratio of the two slopes. The *relative sacrifice ratio*, evaluated along the menu-cost path under accommodation, is therefore given by:

$$\mathrm{RSR}_t = \mathrm{SR}^{\mathrm{Calvo}} \frac{\bar{h}}{h(0) + 3\phi\left(\sigma_t^2 + x_t^2\right)}.$$

Note that, even though the two models share the same average adjustment frequency, the relative sacrifice ratio is already below one in the steady state because selection makes the menu-cost slope

[22] Because the Calvo model is linear, a single scale factor applied to the full shock vector matches its peak exactly. Away from the peak, its rescaled inflation path is flatter than the menu-cost surge. Neither result depends on this match: the look-through comparison feeds both models the common recovered innovations, and the Calvo sacrifice ratio is invariant to shock scale.

$h(0) + 3\phi\sigma_0^2$ exceed $\bar{h}$.

Figure 6 shows how the cost of bringing down inflation is state-dependent in the menu-cost model. Its sacrifice ratio starts below the Calvo value (due to the selection wedge) but then further falls as the slope steepens, and rises again as inflation normalizes. In the United States, the relative sacrifice ratio reaches about 30 percent of the constant Calvo value at the peak of the surge. In the euro area's larger surge, it falls to about a quarter of the Calvo value.

Taken together, these exercises show that state dependence raises the cost of looking through large shocks and lowers the cost of leaning against them. A linear Calvo benchmark understates the first and overstates the second, most of all when price gaps are large, as in the 2021–22 surge. This is because a tightening compresses both the linear pass-through and the nonlinear amplification generated by state-dependent pricing. These results provide a data-contingent positive counterpart to the normative analysis of Karadi et al. (2025), who show in a Ramsey menu-cost model that the planner should respond more aggressively to large inflationary shocks. Our policy counterfactuals also complement Ascari et al. (2025), who find, in an estimated nonlinear model with state-dependent price and wage adjustment for the euro area, that earlier tightening during the recent surge would have reduced inflation at relatively low output cost.

7 Concluding remarks

State-dependent pricing models are widely seen as more persuasive descriptions of firms' price-setting than time-dependent ones—especially after large aggregate shocks. Yet the latter class is the dominant choice of researchers and practitioners because it is significantly easier to implement. This paper narrows that gap with a tractable state-dependent pricing model built from an algebraic approximation of a standard menu-cost framework. In our setting, inflation admits a third-order representation in three moments of the price-gap distribution: the average price gap, its variance, and its skewness. Along the impulse response, under the poor man's simplification, it is an odd (cubic) polynomial in the average price gap alone, the difference between the economy's desired reset price and the price level it enters the period with. The cubic term captures the nonlinearities that state-dependent pricing introduces.

From a practical standpoint, the approximation allows us to apply conventional perturbation methods to solve the model, putting the analysis on par with standard practice in the (nonlinear) DSGE literature. We then illustrate its applicability with standard historical decomposition exercises using both US and euro-area data, and use the identified shocks to run policy counterfactuals. For this exercise, we chose a simple, textbook macro framework. Because the pricing block is portable, however, it can replace the Calvo block of richer models, from

medium-scale models such as Smets and Wouters (2007) to multi-sector economies, in which cross-sector differences in the frequency of adjustment drive monetary non-neutrality and the propagation of shocks (Pastén et al. 2020; Ghassibe and Nakov 2025). It also makes policy relevant questions tractable that so far required global methods: optimal monetary policy when the slope of the Phillips curve depends on the size of shocks, the pass-through of large cost shocks such as tariffs, energy prices, or devaluations, and the interaction of state-dependent pricing with household heterogeneity or with shocks to idiosyncratic volatility. We leave these for future work.

References

Jonathan J. Adams. Sticky prices for inflationary economies: A tractable linear approximation to menu cost models with trend inflation. Working paper, Federal Reserve Bank of Kansas City, July 2026.

Hassan Afrouzi, Saroj Bhattarai, and Edson Wu. The welfare cost of inflation in production networks. Technical report, Working Paper. https://afrouzi. com/# workingpapers, 2024.

Fernando Alvarez and Francesco Lippi. The analytic theory of a monetary shock. *Econometrica*, 90(4):1655–1680, 2022.

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, 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.

Sungbae An and Frank Schorfheide. Bayesian analysis of DSGE models. *Econometric Reviews*, 26 (2–4):113–172, 2007.

Martin M. Andreasen, Jesús Fernández-Villaverde, and Juan F. Rubio-Ramírez. The pruned state-space system for non-linear DSGE models: Theory and empirical applications. *Review of Economic Studies*, 85(1):1–49, 2018.

Guido Ascari, Alexandre Carrier, Emanuel Gasteiger, Alex Grimaud, and Gauthier Vermandel. Monetary policy in the Euro Area, when Phillips curves... are curves. *Available at SSRN 5383859*, 2025.

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.

Adrien Auclert, Rodolfo Rigato, Matthew Rognlie, and Ludwig Straub. Higher-order perturbation in sequence space: The certainty correspondence. Working paper, July 2026.

Ben S. Bernanke. Inflation expectations and inflation forecasting, July 2007. Speech at the Monetary Economics Workshop of the NBER Summer Institute, Cambridge, MA, July 10, 2007.

Ben S. Bernanke and Frederic S. Mishkin. Inflation targeting: A new framework for monetary policy? *Journal of Economic Perspectives*, 11(2):97–116, 1997.

Andres Blanco, Corina Boar, Callum Jones, and Virgiliu Midrigan. Nonlinear inflation dynamics in menu cost economies. *Working paper*, 2024.

Andrés Blanco, Corina Boar, Callum J Jones, and Virgiliu Midrigan. The inflation accelerator. Technical report, American Economic Review, forthcoming, 2026.

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.

Julio Carrillo, Patrick Fève, and Julien Matheron. Monetary policy inertia or persistent shocks: A DSGE analysis. *International Journal of Central Banking*, 3(2):1–38, 2007.

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

Richard Clarida, Jordi Galí, and Mark Gertler. The science of monetary policy: a new Keynesian perspective. *Journal of economic literature*, 37(4):1661–1707, 1999.

James Costain and Anton Nakov. Distributional dynamics under smoothly state-dependent pricing. *Journal of Monetary Economics*, 58(6–8):646–665, 2011.

Christopher D. Cotton and Vaishali Garga. The price-change statistics we've weighted for. Working Paper 26-6, Federal Reserve Bank of Boston, 2026.

Marco Del Negro, Marc P. Giannoni, and Christina Patterson. The forward guidance puzzle. *Journal of Political Economy Macroeconomics*, 1(1):43–79, 2023. doi: 10.1086/724214.

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

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

James Durbin and Siem Jan Koopman. *Time Series Analysis by State Space Methods*. Oxford University Press, Oxford, 2nd edition, 2012.

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

Jesús Fernández-Villaverde and Juan F. Rubio-Ramírez. Estimating macroeconomic models: A likelihood approach. *Review of Economic Studies*, 74(4):1059–1087, 2007.

Luca Gagliardone and Mark Gertler. Oil prices, monetary policy, and inflation surges. *American Economic Journal: Macroeconomics*, 18(4):468–501, 2026.

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

Luca Gagliardone, Mark Gertler, Simone Lenzu, and Joris Tielens. Micro and macro cost-price dynamics in normal times and during inflation surges. Technical report, National Bureau of Economic Research, 2026.

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

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

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

Erwan Gautier, Cristina Conflitti, Daniel Enderle, Ludmila Fadejeva, Alex Grimaud, Eduardo Gutiérrez, Valentin Jouvanceau, Jan-Oliver Menz, Alari Paulus, Pavlos Petroulas, Pau Roldan-Blanco, and Elisabeth Wieland. Consumer price stickiness in the euro area during an inflation surge, 2026.

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

Mishel Ghassibe. Endogenous production networks and non-linear monetary transmission. Working Paper 1449, Barcelona School of Economics, 2024.

Mishel Ghassibe and Anton Nakov. Business cycles with pricing cascades. 2025.

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

Martin Harding, Jesper Lindé, and Mathias Trabandt. Understanding post-covid inflation dynamics. *Journal of Monetary Economics*, 140:S101–S118, 2023.

Jonathon Hazell and Stephan Hobler. Do deficits cause inflation? A high frequency narrative approach. Technical report, London School of Economics, 2025.

Edward Herbst and Frank Schorfheide. Tempered particle filtering. *Journal of Econometrics*, 210 (1):26–44, 2019.

Marek Jarociński and Peter Karadi. Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43, 2020. doi: 10.1257/mac.20180090.

Greg Kaplan and Ken Miyahara. How does monetary and fiscal policy affect the economy in the face of large shocks? Technical report, National Bureau of Economic Research, 2026.

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

Peter Karadi, Anton Nakov, Galo Nuño, Ernesto Pastén, and Dominik Thaler. Strike while the iron is hot–optimal monetary policy under state-dependent pricing. 2025.

Per Krusell and Anthony A Smith, Jr. Income and wealth heterogeneity in the macroeconomy. *Journal of political Economy*, 106(5):867–896, 1998.

Alisdair McKay, Emi Nakamura, and Jón Steinsson. The power of forward guidance revisited. *American Economic Review*, 106(10):3133–3158, 2016. doi: 10.1257/aer.20150063.

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

Frederic S. Mishkin. Headline versus core inflation in the conduct of monetary policy. In *Speech delivered at the Business Cycles, International Transmission and Macroeconomic Policies Conference, HEC Montreal*. October 2007. Federal Reserve Board speech, October 20, 2007.

Benjamin Moll. The trouble with rational expectations in heterogeneous agent models: A challenge for macroeconomics. *The Economic Journal*, pages 1173–1216, 2025.

Emi Nakamura and Jón Steinsson. Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics*, 123(4):1415–1464, 2008.

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

Jimmy Olsson, Olivier Cappé, Randal Douc, and Éric Moulines. Sequential monte carlo smoothing with application to parameter estimation in nonlinear state space models. *Bernoulli*, 14(1): 155–179, 2008. doi: 10.3150/07-BEJ6150.

Ernesto Pastén, Raphael Schoenle, and Michael Weber. The propagation of monetary policy shocks in a heterogeneous production economy. *Journal of Monetary Economics*, 116:1–22, 2020. doi: 10.1016/j.jmoneco.2019.10.001.

Stephanie Schmitt-Grohé and Martín Uribe. Solving dynamic general equilibrium models using a second-order approximation to the policy function. *Journal of Economic Dynamics and Control*, 28(4):755–775, 2004.

Frank Smets and Rafael Wouters. Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review*, 97(3):586–606, 2007.

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

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

Michael Woodford. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, 2011.

An Odd Approximation of Inflation for State-depdendent Pricing Models

L. Gagliardone M. Gertler S. Lenzu

J. Tielens M. Ebsim

Appendix

A Assumptions for the main text and appendices B–C

For clarity of exposition, before developing the proofs we restate the maintained assumptions.

Assumptions. The following assumptions are maintained in the text and Appendices B–C:

1. Zero trend inflation—Trend inflation is zero. This assumption, together with the idiosyncratic process (below), implies that the steady-state price-gap distribution F_0 is even; this delivers the odd-polynomial structure of inflation (Theorems 1–2).

2. Idiosyncratic process—Firm-level productivity follows a random walk in levels, with new increments $\epsilon_t(f)$ that are i.i.d. across firms and over time, mean zero, Gaussian, and independent of the gap the firm carries into the period.

3. Taste–productivity offset—The demand-side taste shock offsets idiosyncratic productivity at its *level*, $\zeta_t(f) := -a_t(f)$. This assumption makes the price gap the only idiosyncratic state of the value function.

4. Regularity conditions for the higher-order moments—We impose the following bounds on the moments of the price gap distribution to formally discuss the approximation remainders. We split this assumption into three parts (*a*, *b*, and *c*), which are progressively stronger so that *4b* implies *4a*, and *4c* implies *4b*.

4a. Weak regularity—We assume the following two conditions: (i) the mean gap stays at the scale of the dispersion, $|x_t| \leq K\,\sigma_t$ for some $K < \infty$; and (ii) the standardized gap has a uniformly bounded $(3+\varepsilon)$ moment, $\sup_t \mathbb{E}\left|(x - x_t)/\sigma_t\right|^{3+\varepsilon} < \infty$ for some $\varepsilon > 0$. This assumption is used in deriving the bounds on inflation in Theorem 1.

4b. Moderate regularity—Consider a family of economies indexed by the scale $\eta \downarrow 0$ of

idiosyncratic heterogeneity. Normalize the idiosyncratic increment as $\epsilon_t^\eta(f) = \eta\,\bar{\epsilon}_t(f), \mathbb{E}[\bar{\epsilon}_t(f)] = 0, \mathbb{V}\mathrm{ar}(\bar{\epsilon}_t(f)) = 1$. Define the normalized pre-adjustment price gap $Z_t^\eta(f) := x_t(f)/\eta$. We assume that, for some $\eta_0 > 0$ and $T < \infty$,

$$\sup_{0<\eta\leq\eta_0} \sup_{0<t\leq T} \mathbb{E}\left|Z_t^\eta(f)\right|^6 < \infty.$$

The assumption is used in the derivations of the dynamics of variance and skewness (Propositions 4–5) as it allows us to derive a closed system in terms of the first three moments only, by eliminating the dependence of variance and skewness on the higher-order moments (up to, respectively, second and third orders). Under this assumption, we abstract from the effects of selection on the dynamics of variance and skewness while retaining the effects of selection on the first moment.

4c. Strong regularity— For the impulse-response results (Theorems 2 and 3), augment the family of economies in Assumption *4b* by a normalized aggregate-shock magnitude $\bar{\delta} \in \mathcal{D}$, where $\mathcal{D}$ is compact and $0 < \underline{d} \leq |\bar{\delta}| \leq \overline{d} < \infty$. In the economy indexed by η, the actual impact displacement is $\delta = \eta\bar{\delta}$, so that the mean gap on impact is $x_0 = \delta$ (the mean gap is zero in the stochastic steady state). Thus η indexes the scale of idiosyncratic heterogeneity, while $\bar{\delta}$ indexes the magnitude and sign of the aggregate shock within a given economy. Define the normalized pre-adjustment price gap by $Z_t^{\eta,\bar{\delta}}(f) := x_t(f)/\eta$. We assume that, for some $\eta_0 > 0$ and $T < \infty$,

$$\sup_{0<\eta\leq\eta_0} \sup_{\bar{\delta}\in\mathcal{D}} \sup_{0\leq t\leq T} \mathbb{E}\left|Z_t^{\eta,\bar{\delta}}(f)\right|^6 < \infty.$$

This condition requires that the aggregate shock and idiosyncratic heterogeneity vanish at the same rate, and allows us to bound the remainders of the dynamics of variance and skewness by those of the aggregate shock. All local remainder statements are understood as $\eta \downarrow 0$, holding $\bar{\delta}$ fixed, uniformly over $\bar{\delta} \in \mathcal{D}$. Since $\eta = |\delta|/|\bar{\delta}|$, and $|\bar{\delta}|$ is bounded above and away from zero, $o(\eta^k) = o(\delta^k)$ and $O(\eta^k) = O(\delta^k)$ uniformly over $\mathcal{D}$.

5. Regularity conditions for the menu-cost distribution—The menu-cost distribution G satisfies $G(0) = 0$ and has a continuous density on $(0, \infty)$ that extends continuously to the origin, so that $G'(0) := \lim_{y\downarrow 0} G(y)/y$ exists and is finite; without it the hazard carries no quadratic term. The hazard is four times differentiable at zero. The value function is strictly concave at the reset point, $\mathcal{V}_t''(0) < 0$; otherwise $\phi_t = 0$ in Proposition 1.

B Derivations

B.1 Static target price

We derive the static target price of Equation (3). The firm's static optimality condition implies that firms price at a markup over nominal marginal cost. Moreover, under our baseline demand system, the elasticity of demand is constant and the same at every relative price, which implies that the demand curve is log-linear and the desired markup is constant.

Final goods are assembled competitively from the differentiated varieties using the constant-elasticity technology with elasticity of demand $\theta > 1$:

$$Y_t = \left[\int_0^1 \left(Y_t(f)\, e^{\zeta_t(f)} \right)^{\frac{\theta-1}{\theta}} df \right]^{\frac{\theta}{\theta-1}}, \tag{A.1}$$

which yields, in logarithmic terms, for each firm f a residual demand function of the form:

$$y_t(f) = -\theta \left(p_t(f) - p_t \right) + (\theta - 1)\, \zeta_t(f) + y_t, \tag{A.2}$$

where p_t is the CES price index. Given Equation (A.2), the elasticity of demand is the same at every relative price, and the price that maximizes static profit is a constant markup $\mu := \ln \frac{\theta}{\theta-1}$ over nominal marginal cost:

$$p_t^o(f) = \mu + mc_t^n(f). \tag{A.3}$$

Nominal marginal cost is given by Equation (1), $mc_t^n(f) = \xi_t - (1+\nu) a_t(f) + \nu\, y_t(f)$, which is a function of a common price shifter, an idiosyncratic shifter, and the firm's own physical output. Evaluating the demand curve (A.2) at the static optimum and substituting removes the firm's own output:

$$mc_t^n(f) = \xi_t - (1+\nu) a_t(f) + \nu \Big[y_t - \theta \big(p_t^o(f) - p_t \big) + (\theta - 1)\, \zeta_t(f) \Big]. \tag{A.4}$$

The expression in Equation (A.4) holds for any returns to scale. Relative to the standard constant returns case ($\nu = 0$), three terms appear. A relative-price term $-\nu\theta \big(p_t^o(f) - p_t \big)$: a firm that prices above the market sells less and, with decreasing returns, produces at a lower marginal cost, which pulls its target back toward competitors' prices. A common level term νy_t. And a taste term $\nu(\theta - 1)\, \zeta_t(f)$, which operates for the same reason in reverse—a more desired variety is produced in greater quantity, and at $\nu > 0$ greater quantity means higher marginal cost.

Substituting (A.4) into (A.3) and collecting the terms in $p_t^o(f)$ yields:

$$p_t^o(f) = (1 - \Xi) \big(\mu + \xi_t + \nu y_t \big) + \Xi\, p_t - (1 - \Xi)(1+\nu) a_t(f) + (1 - \Xi)\, \nu(\theta - 1)\, \zeta_t(f),$$

where the weight the target places on competitors' prices is given by $\Xi := \frac{\nu\theta}{1+\nu\theta} \in [0, 1)$, and $\nu\theta$ is

the rate at which a higher relative price lowers the firm's own marginal cost, so Ξ rises with ν and is zero at constant returns.

Finally, assuming offsetting taste shocks, $\zeta_t(f) = -a_t(f)$, the two idiosyncratic terms in the equation above collapse into one, with a loading on $\zeta_t(f)$ equal to one for any ν and any θ, yielding Equation (3) in the paper.

B.2 Dynamic reset price

We now derive the dynamic target price in Equation (6). Throughout, $x_t^o(f) := p_t^o(f) - p_{t-1}(f)$ denotes the static gap a firm would carry into period t absent adjustment, and $\widetilde{\mathcal{V}}_t(\cdot)$ denotes the value of not adjusting before the menu-cost draw, related to the indirect value function (Lemma 1) by $\mathcal{V}_t(x) = \widetilde{\mathcal{V}}_t(x_t^\star + x)$. We abstract (without loss of generality for this subsection) from the free-price adjustments to simplify the notation (i.e., set $h_t(0) = 0$ in what follows).

Following Alvarez et al. (2023), we use a quadratic approximation to characterize firms' (real) flow profits around the flexible (static) optimal price $p_t^o(f)$ of Equation (3)—this is the second-order (symmetric-loss) profit approximation, which makes profit an even function of the price gap.

Let $\Pi(x)$ denote flow profit, expressed as a share of the firm's steady-state profit Ψ_0, as a function of a generic gap x between the static optimum and the firm's price. The menu cost is measured in the same units: from here on $\chi_t(f)$ denotes the adjustment cost as a share of Ψ_0, and G its distribution. Its quadratic form is:

$$\Pi_t(f) \equiv \Pi\big(p_t^o(f) - p_t(f)\big) := -\frac{B}{2}\left(p_t^o(f) - p_t(f)\right)^2, \tag{A.5}$$

where $B > 0$ is, by this normalization, directly the curvature $-\Pi''(0)$ of the profit share. Since $p_t^o(f)$ maximizes static profits, Π attains its maximum at a zero gap, so that $\Pi'(0) = 0$ and $\Pi''(0) = -B < 0$. Define, for any gap $x \in \mathbb{R}$, the value of not changing the price at date t as

$$\widetilde{\mathcal{V}}_t(x) \;:=\; \Pi(x) \;+\; \beta\, \mathbb{E}_t\big[\, V_{t+1}(f) \mid x \,\big], \qquad \beta \in (0,1). \tag{A.6}$$

Hence, the optimal reset gap

$$x_t^\star(f) \;\in\; \arg\max_x \widetilde{\mathcal{V}}_t(x)$$

is independent of $\chi_t(f)$, and the corresponding optimal reset price is

$$p_t^\star(f) \;=\; p_t^o(f) - x_t^\star(f).$$

Under the regularity conditions on the primitives—a differentiable value function and a smooth menu-cost distribution—$\widetilde{\mathcal{V}}_t(\cdot)$ is differentiable and the maximizer is interior, so that the first-order

condition holds:

$$\widetilde{\mathcal{V}}_t'\big(x_t^\star(f)\big) = 0. \tag{A.7}$$

In terms of the value of adjusting and not adjusting, the value function is

$$V_t(f) = \max\Big\{\widetilde{\mathcal{V}}_t\big(x_t^o(f)\big),\ \widetilde{\mathcal{V}}_t\big(x_t^\star(f)\big) - \chi_t(f)\Big\}. \tag{A.8}$$

Write $x' := p_{t+1}^o(f) - p$, with $p := p_t^o(f) - x$, for the static gap the firm would carry into $t+1$ having held its price at the counterfactual level p; at $x = x_t^\star(f)$ this coincides with the firm's actual decision. It exceeds the price gap of Equation (4) by the reset gap, $x' = x_{t+1}^\star + x_{t+1}(f)$, and the hazard below is written at that argument. Because $\widetilde{\mathcal{V}}_t\big(x_t^\star(f)\big) - \chi_t(f)$ does not depend on x, and $\chi_{t+1}(f)$ is drawn independently of x', Equation (A.8) one period ahead can be written with the menu cost integrated out,

$$\mathbb{E}\left[V_{t+1}(f) \mid x'\right] = \widetilde{\mathcal{V}}_{t+1}(x') + \mathbb{E}\left[\max\left\{0,\ \Delta_{t+1}(x') - \chi_{t+1}(f)\right\}\right], \tag{A.9}$$

with

$$\Delta_{t+1}(x') := \widetilde{\mathcal{V}}_{t+1}\big(x_{t+1}^\star\big) - \widetilde{\mathcal{V}}_{t+1}(x')\ \geq 0.$$

The second term is a differentiable function of the value gain Δ_{t+1}, with derivative equal to the adjustment probability $\Pr\{\chi_{t+1}(f) \leq \Delta_{t+1}\}$, while $\Delta_{t+1}'(x') = -\widetilde{\mathcal{V}}_{t+1}'(x')$. The right-hand side of (A.9) therefore has derivative $\big(1-\mathbb{I}_{t+1}^\star(f)\big)\widetilde{\mathcal{V}}_{t+1}'(x')$ in x' (by the envelope theorem) and $\partial x'/\partial x = 1$,

$$\frac{d}{dx}\ \mathbb{E}_t\left(V_{t+1}(f) \mid x\right) = \mathbb{E}_t\left[\left(1 - \mathbb{I}_{t+1}^\star(f)\right)\widetilde{\mathcal{V}}_{t+1}'\big(x'\big)\ \middle|\ x\right]. \tag{A.10}$$

Differentiating (A.6) with respect to x and using that $\partial(p_{t+1}^o(f) - p)/\partial x = 1$ yields

$$\widetilde{\mathcal{V}}_t'(x) = \Pi'(x) + \beta\ \mathbb{E}_t\Big[\big(1 - \mathbb{I}_{t+1}^\star(f)\big)\ \widetilde{\mathcal{V}}_{t+1}'\big(p_{t+1}^o(f) - p\big)\Big], \tag{A.11}$$

with $p = p_t^o(f) - x$. Evaluate (A.11) at $x = x_t^\star(f)$, so that $p = p_t^\star(f)$, and use (A.7) to obtain

$$0 = \Pi'\big(p_t^o(f) - p_t^\star(f)\big) + \beta\ \mathbb{E}_t\Big[\big(1 - \mathbb{I}_{t+1}^\star(f)\big)\ \widetilde{\mathcal{V}}_{t+1}'\big(p_{t+1}^o(f) - p_t^\star(f)\big)\Big]. \tag{A.12}$$

Next, apply (A.11) one period ahead to the derivative $\widetilde{\mathcal{V}}_{t+1}'\big(p_{t+1}^o(f) - p_t^\star(f)\big)$, substitute into (A.12), and use the law of iterated expectations. Repeating this substitution N times yields, for any $N \geq 0$,

$$\begin{aligned} 0 = \mathbb{E}_t\Bigg[\sum_{i=0}^{N}\beta^i\left(\prod_{j=1}^{i}\left(1 - \mathbb{I}_{t+j}^\star(f)\right)\right)\Pi'\big(p_{t+i}^o(f) - p_t^\star(f)\big)\Bigg] \\ + \beta^{N+1}\ \mathbb{E}_t\Bigg[\left(\prod_{j=1}^{N+1}\left(1 - \mathbb{I}_{t+j}^\star(f)\right)\right)\widetilde{\mathcal{V}}_{t+N+1}'\big(p_{t+N+1}^o(f) - p_t^\star(f)\big)\Bigg], \end{aligned} \tag{A.13}$$

where by convention $\prod_{j=1}^{0}(\cdot) = 1$. The second line of (A.13) is β^{N+1} times a term of order $\mathbb{E}_t\left|p^o_{t+N+1}(f) - p^\star_t(f)\right|$ (since $\widetilde{\mathcal{V}}'$ satisfies the transversality condition and the survival product is bounded by one); it therefore vanishes as $N \to \infty$ whenever $\beta^N\ \mathbb{E}_t\left|p^o_{t+N}(f) - p^\star_t(f)\right| \to 0$, which the random-walk target of Assumption 2 satisfies. Therefore, taking limits in (A.13) gives the infinite-horizon first-order condition:

$$0 = \mathbb{E}_t\left[\sum_{i=0}^{\infty}\beta^i\left(\prod_{j=1}^{i}\left(1-\mathbb{I}^\star_{t+j}(f)\right)\right)\Pi'\left(p^o_{t+i}(f) - p^\star_t(f)\right)\right]. \tag{A.14}$$

Using the quadratic profit approximation (A.5), $\Pi'(x) = -Bx$, and (A.14) becomes:

$$0 = -B\ \mathbb{E}_t\left[\sum_{i=0}^{\infty}\beta^i\left(\prod_{j=1}^{i}\left(1-\mathbb{I}^\star_{t+j}(f)\right)\right)\left(p^o_{t+i}(f) - p^\star_t(f)\right)\right].$$

Divide by $-B \neq 0$ and rearrange to isolate $p^\star_t(f)$:

$$p^\star_t(f)\ \mathbb{E}_t\left[\sum_{i=0}^{\infty}\beta^i\left(\prod_{j=1}^{i}\left(1-\mathbb{I}^\star_{t+j}(f)\right)\right)\right] = \mathbb{E}_t\left[\sum_{i=0}^{\infty}\beta^i\left(\prod_{j=1}^{i}\left(1-\mathbb{I}^\star_{t+j}(f)\right)\right)p^o_{t+i}(f)\right].$$

Hence,

$$p^\star_t(f) = \frac{\mathbb{E}_t\left[\sum_{i=0}^{\infty}\beta^i\left(\prod_{j=1}^{i}\left(1-\mathbb{I}^\star_{t+j}(f)\right)\right)p^o_{t+i}(f)\right]}{\mathbb{E}_t\left[\sum_{i=0}^{\infty}\beta^i\left(\prod_{j=1}^{i}\left(1-\mathbb{I}^\star_{t+j}(f)\right)\right)\right]}. \tag{A.15}$$

Equation (A.15) expresses the reset price through the realized survival indicators $\prod_{j=1}^{i}\left(1-\mathbb{I}^\star_{t+j}(f)\right)$. Defining the firm's effective discount factor $\Lambda_t(f) \coloneqq \mathbb{E}_t\sum_{i=0}^{\infty}\beta^i\prod_{j=1}^{i}\left(1-\mathbb{I}^\star_{t+j}(f)\right)$, the indicator products factor as follows. Taking expectations,

$$\mathbb{E}_t\left[\prod_{j=1}^{i}\left(1-\mathbb{I}^\star_{t+j}(f)\right)p^o_{t+i}(f)\right] = \mathbb{E}_t\left[\prod_{j=1}^{i}\left(1-h_{t+j}(f)\right)p^o_{t+i}(f)\right], \tag{A.16}$$

and likewise for the denominator, which delivers the form of Equation (6) used in the main text.

Finally, because the idiosyncratic component of $p^o_t(f)$ follows a random walk, the reset gap $x^\star_t(f) \coloneqq p^o_t(f) - p^\star_t(f)$ does not depend on the identity of the firm. That is, $x^\star_t(f) = x^\star_t$.

B.3 Firm-level recursions and aggregation

Define the firm-level object $W_t(f) \coloneqq \left(1-h_t(f)\right)\Lambda_t(f)$ and denote the covariance as:

$$C^\star_{t+1} \coloneqq \mathrm{Cov}_f\left(W_{t+1}(f),\ p^\star_{t+1}(f)\right). \tag{A.17}$$

In the stationary benchmark below, $W_{t+1}(f) = W\left(d_{t+1} + \epsilon_{t+1}(f)\right)$, where $W(v) \coloneqq \left(1-h(v)\right)\Lambda(v)$ is a time-invariant function of the gap v.

Proposition 3 (Third-order covariance correction)**.** *Consider a stationary symmetric benchmark with* $\Lambda_t(v) = \Lambda$*. The covariance is given by:*

$$C^\star_{t+1} = J\,(p^\star_{t+1} - p^\star_t) + o(\eta^3), \qquad J := W''(0)\sigma^2_\epsilon. \tag{A.18}$$

Proof. Integrating the firm recursion (18) over all firms gives

$$\begin{aligned} \Lambda_t p^\star_t =& p^o_t + \beta\,\mathbb{E}_t\left[\int_0^1 W_{t+1}(f)p^\star_{t+1}(f)df\right] \\ =& p^o_t + \beta\,\mathbb{E}_t\left[\int_0^1 W_{t+1}(f)df\,p^\star_{t+1} + C^\star_{t+1}\right] \end{aligned} \tag{A.19}$$

Let $\overline{W}_{t+1}$ be the average of $W_{t+1}(f)$. Denoting by $d_{t+1} := p^\star_{t+1} - p^\star_t$, under stationarity, $W_{t+1}(f) = W(d_{t+1} + \epsilon_{t+1}(f))$. Since the common component is additive, the covariance can be written as

$$C^\star_{t+1} = \int_0^1 W(d_{t+1} + \epsilon_{t+1}(f))\,[\zeta_t(f) + \epsilon_{t+1}(f)]\,df. \tag{A.20}$$

Independence of $\zeta_t(f)$ and $\epsilon_{t+1}(f)$, together with $\int_0^1 \zeta_t df = 0$, gives

$$C^\star_{t+1} = \int_0^1 \epsilon_{t+1}(f)W(d_{t+1} + \epsilon_{t+1}(f))\,df. \tag{A.21}$$

Since the covariance already contains one factor $\epsilon = O(\eta)$, and W is even in the stationary symmetric benchmark, so that $W'''(0) = 0$, a quadratic Taylor expansion of W is sufficient to obtain the covariance through third order:

$$W(d+\epsilon) = W(0) + W'(0)(d+\epsilon) + \frac{1}{2}W''(0)(d+\epsilon)^2 + O(\eta^4) \tag{A.22}$$

Multiplying (A.22) by ϵ and averaging leads to:

$$C^\star_{t+1} = W'(0)\sigma^2_\epsilon + W''(0)d_{t+1}\sigma^2_\epsilon + \frac{1}{2}W''(0)\int_0^1 \epsilon^3_{t+1}df + O(\eta^4). \tag{A.23}$$

Moreover, stationarity and evenness of h imply $W'(0) = 0$, and, since $\epsilon_{t+1}(f)$ is Gaussian with mean zero, its cross-sectional third moment vanishes, $\int_0^1 \epsilon^3_{t+1}\,df = 0$. Substituting these two results into (A.23) yields:

$$\begin{aligned} C^\star_{t+1} &= W''(0)\sigma^2_\epsilon d_{t+1} + O(\eta^4) \\ &= J\,d_{t+1} + o(\eta^3), \qquad J := W''(0)\sigma^2_\epsilon, \end{aligned} \tag{A.24}$$

where the last step uses $O(\eta^4) = o(\eta^3)$. □

B.4 Dynamics of higher-order moments

For notational convenience, write $h \equiv h(0)$ and $H \equiv 1 - h$. In this section, we work under the moderate regularity Assumption *4b* (Appendix A).

We distinguish between two cross-sectional distributions of price gaps, measured at two distinct points within the period. The *pre-adjustment* distribution, introduced in Section 2.1, captures the cross-section of the price gaps $x_t(f) = p_t^\star(f) - p_{t-1}(f)$ of Equation (4), measured after the realization of the period-t aggregate and idiosyncratic shocks but *before* firms take their adjustment decisions. We denote the distribution of price gaps by F_t and its four moments—mean, variance, skewness, and kurtosis—by $x_t, \sigma_t^2, \psi_t, \kappa_t$, as defined in Section 2.1. These are the objects that enter the contemporaneous frequency and inflation approximations: aggregate inflation (7) accumulates the price changes of the firms that adjust, and the adjustment decision is taken on the pre-adjustment gap. When only idiosyncratic shocks are drawn, F_t collapses to the steady-state distribution F_0, with moments $x_0, \sigma_0^2, \psi_0, \kappa_0$.

The *post-adjustment* distribution captures the cross-section of price gaps once the adjustment decisions $\mathbb{I}_t^\star(f) \in \{0, 1\}$ are realized in period t. The cross-section is truncated because a firm that adjusts ($\mathbb{I}_t^\star(f) = 1$) resets its price to the target and carries a zero gap into the next period, while a firm that does not adjust ($\mathbb{I}_t^\star(f) = 0$) keeps its gap. The gap a firm carries forward is therefore the *post-adjustment gap*

$$x_{t,+}(f) := \left(1 - \mathbb{I}_t^\star(f)\right) x_t(f), \tag{A.25}$$

and we denote its cross-sectional distribution by $F_{t,+}$, with central moments

$$\begin{aligned} x_{t,+} &:= \int x \, dF_{t,+}(x), & \sigma_{t,+}^2 &:= \int (x - x_{t,+})^2 \, dF_{t,+}(x), \\ \psi_{t,+} &:= \int (x - x_{t,+})^3 \, dF_{t,+}(x), & \kappa_{t,+} &:= \int (x - x_{t,+})^4 \, dF_{t,+}(x). \end{aligned} \tag{A.26}$$

The truncation at $\mathbb{I}_t^\star(f) = 1$ is the selection effect: because adjustment is more likely at large gaps (the hazard is convex), it removes mass from the tails of F_t and reshapes the distribution that firms carry into the following period. Absent aggregate shocks these post-adjustment moments take their steady-state values, which we denote $x_{0,+}, \sigma_{0,+}^2, \psi_{0,+}, \kappa_{0,+}$.

The two cross-sections are linked across time by the idiosyncratic innovation. Between $t-1$ and t every carried gap $x_{t-1,+}(f)$ receives the new idiosyncratic increment $\epsilon_t(f)$—drawn i.i.d. across firms and over time, mean zero, symmetric, with variance σ_ϵ^2, and independent of the carried gap. The pre-adjustment gap at t is thus the carried gap displaced by the new draw,

$$x_t(f) = x_{t-1,+}(f) + \epsilon_t(f),$$

up to the common aggregate displacement that shifts every gap and is absorbed in the mean.

Under our assumption of independence of $\epsilon_t(f)$ from the carried gap, we have the following variance timing identity:

$$\sigma_t^2 = \sigma_{t-1,+}^2 + \sigma_\epsilon^2, \tag{A.27}$$

while the assumption of symmetric idiosyncratic shocks implies no third cumulant, which carries the skewness forward unchanged:

$$\psi_t = \psi_{t-1,+}. \tag{A.28}$$

In the stochastic steady state Equation (A.27) reads:

$$\sigma_0^2 = \sigma_{0,+}^2 + \sigma_\epsilon^2. \tag{A.29}$$

That is, the steady-state pre-adjustment variance exceeds the carried variance by exactly the variance of one idiosyncratic draw. Equations (A.27) and (A.28) link the two cross-sections: they map the post-adjustment variance and skewness carried from $t-1$ into the pre-adjustment variance and skewness at t, which are the moments that enter the contemporaneous formulas.

Let X_t denote the pre-adjustment gap $x_t(f)$ of a randomly drawn firm, $M_{k,t} := \mathbb{E}[X_t^k]$ and $M_{k,t}^+ := \mathbb{E}[X_{t,+}^k]$. For every integer $k \geq 1$, the carried gap satisfies

$$X_{t,+}^k = (1 - \mathbb{I}_t^\star) X_t^k,$$

because $(1 - \mathbb{I}_t^\star)^k = 1 - \mathbb{I}_t^\star$. Conditioning on the pre-adjustment gap therefore gives

$$\begin{aligned} M_{k,t}^+ &= \mathbb{E}\left[\mathbb{E}\left[1 - \mathbb{I}_t^\star \mid X_t\right] X_t^k\right] \\ &= \mathbb{E}\left[(1 - h(X_t)) X_t^k\right]. \end{aligned} \tag{A.30}$$

Write the local hazard approximation as

$$h(z) = h + \phi z^2 + r(z), \qquad r(z) = o(z^2).$$

Assumption *4b* implies, for $j = 0, 1, 2, 3$,

$$\mathbb{E}\left[|X_t|^j |r(X_t)|\right] = o(\eta^{j+2}). \tag{A.31}$$

Indeed, write $r(z) = z^2 \omega(z)$, with $\omega(z) \to 0$ as $z \to 0$, and split the expectation over $|X_t| \leq \bar{z}$ and $|X_t| > \bar{z}$. On the first set, the expectation is bounded by an arbitrarily small constant times $\mathbb{E}\,|X_t|^{j+2} = O(\eta^{j+2})$. On the second set, the uniform sixth-moment bound gives uniform integrability for every power required in (A.31). Hence

$$M_{k,t}^+ = H M_{k,t} - \phi M_{k+2,t} + o(\eta^{k+2}), \qquad k = 1, 2, 3. \tag{A.32}$$

The raw moments satisfy

$$M_{1,t} = x_t, \tag{A.33}$$

$$M_{2,t} = \sigma_t^2 + x_t^2, \tag{A.34}$$

$$M_{3,t} = \psi_t + 3x_t\sigma_t^2 + x_t^3, \tag{A.35}$$

$$M_{4,t} = \kappa_t + 4x_t\psi_t + 6x_t^2\sigma_t^2 + x_t^4, \tag{A.36}$$

$$M_{5,t} = m_{5,t} + 5x_t\kappa_t + 10x_t^2\psi_t + 10x_t^3\sigma_t^2 + x_t^5. \tag{A.37}$$

Here $m_{5,t}$ denotes the fifth central moment of F_t. Absent the aggregate shock, the pre- and post-adjustment means vanish. Using the variance timing identity (A.29) and (A.32) at $k = 2$ gives

$$\begin{aligned}\sigma_0^2 &= \sigma_{0,+}^2 + \sigma_\epsilon^2 \\ &= M_{2,0}^+ + \sigma_\epsilon^2 \\ &= H\sigma_0^2 - \phi\kappa_0 + \sigma_\epsilon^2 + o(\eta^4),\end{aligned} \tag{A.38}$$

and therefore

$$\sigma_\epsilon^2 = h\sigma_0^2 + \phi\kappa_0 + o(\eta^4). \tag{A.39}$$

Proposition 4 (Dynamics of the variance)**.** *The pre-adjustment variance satisfies*

$$\sigma_{t+1}^2 - \sigma_0^2 = H\left(\sigma_t^2 - \sigma_0^2\right) + hHx_t^2 + o(\eta^2). \tag{A.40}$$

Proof. By the variance timing identity,

$$\sigma_{t+1}^2 = M_{2,t}^+ - \left(M_{1,t}^+\right)^2 + \sigma_\epsilon^2.$$

Using (A.32) at $k = 1, 2$,

$$\begin{aligned}\sigma_{t+1}^2 =& HM_{2,t} - \phi M_{4,t} - \left(HM_{1,t} - \phi M_{3,t}\right)^2 + \sigma_\epsilon^2 + o(\eta^4) \\ =& H\sigma_t^2 + hHx_t^2 - \phi M_{4,t} + 2H\phi x_t M_{3,t} - \phi^2 M_{3,t}^2 + \sigma_\epsilon^2 + o(\eta^4).\end{aligned}$$

Subtracting σ_0^2 and using (A.39) yields

$$\sigma_{t+1}^2 - \sigma_0^2 = H\left(\sigma_t^2 - \sigma_0^2\right) + hHx_t^2 - \phi\left(M_{4,t} - \kappa_0\right) + 2H\phi x_t M_{3,t} - \phi^2 M_{3,t}^2 + o(\eta^4). \tag{A.41}$$

Assumption *4b* in Appendix A gives

$$M_{4,t} - \kappa_0 = O(\eta^4), \qquad x_t M_{3,t} = O(\eta^4), \qquad M_{3,t}^2 = O(\eta^6).$$

Hence the last two terms of (A.41), together with $-\phi(M_{4,t} - \kappa_0)$, are $O(\eta^4) = o(\eta^2)$. This proves (A.40). □

Proposition 5 (Dynamics of skewness)**.** *The pre-adjustment third central moment satisfies*

$$\psi_{t+1} = H\psi_t + 3hHx_t\sigma_t^2 + hH(2h-1)x_t^3 + o(\eta^3). \tag{A.42}$$

Proof. The idiosyncratic increment is independent of the carried gap and has zero third central moment. Hence the third central moment crosses the timing step unchanged:

$$\psi_{t+1} = M_{3,t}^+ - 3M_{1,t}^+M_{2,t}^+ + 2\big(M_{1,t}^+\big)^3.$$

By (A.32) and the raw-moment relations (A.33)–(A.35),

$$M_{1,t}^+ = HM_{1,t} + O(\eta^3), \qquad M_{2,t}^+ = HM_{2,t} + O(\eta^4), \qquad M_{3,t}^+ = HM_{3,t} + O(\eta^5).$$

Consequently,

$$\psi_{t+1} = HM_{3,t} - 3H^2M_{1,t}M_{2,t} + 2H^3M_{1,t}^3 + O(\eta^5).$$

Using (A.33)–(A.35),

$$\begin{aligned}\psi_{t+1} &= H\big(\psi_t + 3x_t\sigma_t^2 + x_t^3\big) - 3H^2x_t\big(\sigma_t^2 + x_t^2\big) + 2H^3x_t^3 + O(\eta^5)\\ &= H\psi_t + 3H(1-H)x_t\sigma_t^2 + H(1-H)(1-2H)x_t^3 + O(\eta^5).\end{aligned}$$

Since $1 - H = h$ and $1 - 2H = 2h - 1$, this is (A.42), because $O(\eta^5) = o(\eta^3)$. □

C Proofs

In this section we provide the proofs and derivations of the results presented in Sections 2 and 3. The road map is the following:

- Section C.1 proves Lemma 1, reducing each firm's pricing problem to the single scalar price gap and establishing the recursive representation, interior optimum, and adjustment threshold of the resulting indirect value function.
- Section C.2 proves Proposition 1, the quadratic approximation of the hazard rate.
- Section C.3 formalizes the projection of aggregate inflation onto the moments of the price-gap distribution in two steps: the general representation (Theorem 1), and its specialization to the stationary, symmetric environment maintained here, which recovers the odd cubic of the main text.

• Steps 1 and 2 of Section C.4 derive the pre-adjustment variance's and skewness's response to the aggregate shock, each first as a polynomial in the mean gap and then in closed form as a function of the steady-state primitives $(h(0), \phi, \sigma_0^2)$ and of the equilibrium persistence of the average gap, from a shared selection integral.

• Step 3 of Section C.4 proves Theorem 2, deriving the frequency and inflation impulse-response coefficients $\bar{h}, \alpha_2, \beta_1, \beta_3$ from the moment responses of Steps 1 and 2.

C.1 Proof of Lemma 1

Proof of Lemma 1 (Price gaps as forcing variable). Fix a firm f and a date t. We work conditional on the aggregate state (suppressed in the notation) and on the firm's current draw of the adjustment cost $\chi_t(f)$, which is observed by the firm before choosing $\mathbb{I}_t(f)$ and $p_t(f)$. Because adjusting price when receiving a free adjustment opportunity is always—weakly—optimal, we characterize (without loss of generality) in what follows the value of not adjusting and associated decisions conditional on not receiving a free adjustment opportunity at t; we then extend to the case of free adjustment when deriving the hazard at the end of the proof.

Define the static gap between the frictionless optimal price and the price at the beginning of the period:

$$x_t^o(f) := p_t^o(f) - p_{t-1}(f). \tag{A.43}$$

Then,

$$\begin{aligned} x_t^o(f) &= p_t^o(f) - p_{t-1}(f) \\ &= \left(p_t^o(f) - p_t^\star(f)\right) + \left(p_t^\star(f) - p_{t-1}(f)\right) \\ &= x_t^\star(f) + x_t(f). \end{aligned} \tag{A.44}$$

Under the random walk dynamics of the idiosyncratic shocks, which renders the reset gap common across firms, $x_t^\star(f) = x_t^\star$ for all f, (A.44) becomes

$$x_t^o(f) = x_t^\star + x_t(f). \tag{A.45}$$

Define the value of not changing price as (mirroring the construction in Appendix B):

$$\widetilde{\mathcal{V}}_t(x) := \Pi(x) + \beta\, \mathbb{E}_t\left[V_{t+1}(f) \,\middle|\, x\right]. \tag{A.46}$$

Because profits depend on $p_t(f)$ only through $p_t^o(f) - p_t(f)$ (by (A.5)), and because the law of motion for future primitives is Markov, $\widetilde{\mathcal{V}}_t(\cdot)$ is a well-defined function of the scalar gap x (given

the aggregate state at t). The value of adjusting (conditional on no free adjustment) is

$$\max_{x \in \mathbb{R}} \widetilde{\mathcal{V}}_t(x) \; - \; \chi_t(f). \tag{A.47}$$

Under differentiability, an interior optimum, and Assumption 5, it satisfies

$$\widetilde{\mathcal{V}}_t'\big(x_t^\star(f)\big) = 0, \qquad \widetilde{\mathcal{V}}_t''\big(x_t^\star(f)\big) < 0. \tag{A.48}$$

In terms of the value of adjusting and not adjusting, the value function is

$$V_t(f) = \max\Big\{\widetilde{\mathcal{V}}_t\big(x_t^o(f)\big),\; \widetilde{\mathcal{V}}_t\big(x_t^\star(f)\big) - \chi_t(f)\Big\}. \tag{A.49}$$

Hence, the optimal decisions satisfy

$$p_t^\star(f) = p_t^o(f) - x_t^\star(f). \tag{A.50}$$

$$\mathbb{I}_t^\star(f) = 1 \quad \Longleftrightarrow \quad \widetilde{\mathcal{V}}_t\big(x_t^\star(f)\big) - \chi_t(f) \;\geq\; \widetilde{\mathcal{V}}_t\big(x_t^o(f)\big). \tag{A.51}$$

Define the indirect value function in terms of the dynamic gap as the re-centered version of $\widetilde{\mathcal{V}}_t$:

$$\mathcal{V}_t(x) := \widetilde{\mathcal{V}}_t\big(x_t^\star + x\big), \qquad x \in \mathbb{R}. \tag{A.52}$$

The continuation expectations in (A.46) and below are conditioned on the firm's payoff-relevant dynamic gap $x_t(f)$, which—together with the aggregate state—pins down the distribution of next period's gap x'. Using (A.45) and (A.52),

$$\widetilde{\mathcal{V}}_t\big(x_t^o(f)\big) = \widetilde{\mathcal{V}}_t\big(x_t^\star + x_t(f)\big) = \mathcal{V}_t\big(x_t(f)\big), \tag{A.53}$$

and also

$$\mathcal{V}_t(0) = \widetilde{\mathcal{V}}_t(x_t^\star). \tag{A.54}$$

Substituting (A.53)–(A.54) into (A.49) yields:

$$V_t(f) = \max\Big\{\mathcal{V}_t\big(x_t(f)\big),\; \mathcal{V}_t(0) - \chi_t(f)\Big\}. \tag{A.55}$$

Differentiate (A.52):

$$\mathcal{V}_t'(x) = \widetilde{\mathcal{V}}_t'\big(x_t^\star + x\big), \qquad \mathcal{V}_t''(x) = \widetilde{\mathcal{V}}_t''\big(x_t^\star + x\big).$$

Evaluating at $x = 0$ and using (A.48) with $x_t^\star(f) = x_t^\star$ gives

$$\mathcal{V}_t'(0) = 0, \qquad \mathcal{V}_t''(0) < 0,$$

which is condition (ii) (the interior optimum) of Lemma 1. Fix $x \in \mathbb{R}$ and interpret it as the current

dynamic gap $x_t(f) = x$. If the firm does not adjust at t, the flow payoff is $\Pi(x_t^\star + x)$, and by definition of $\widetilde{\mathcal{V}}_t$,

$$\widetilde{\mathcal{V}}_t(x_t^\star + x) = \Pi(x_t^\star + x) + \beta\, \mathbb{E}_t\left[V_{t+1}(f) \,\middle|\, x_t(f) = x\right]. \tag{A.56}$$

The left-hand side equals $\mathcal{V}_t(x)$ by (A.52). Let

$$x' \equiv x_{t+1}(f) \tag{A.57}$$

denote next period's dynamic gap implied by inaction at t (i.e. given that the price carried into $t+1$ is $p_t(f) = p_{t-1}(f)$). Its conditional law is pinned down by the Markov evolution of $p^o_{t+1}(f)$ (and thus of $p^\star_{t+1}(f)$) together with the fixed inherited price. At date $t+1$, the firm again chooses whether to adjust. Using (A.55),

$$V_{t+1}(f) = \max\left\{\mathcal{V}_{t+1}(x'),\ \mathcal{V}_{t+1}(0) - \chi_{t+1}(f)\right\}.$$

Write this maximum using the optimal decision $\mathbb{I}^\star_{t+1}(f)$:

$$V_{t+1}(f) = \mathbb{I}^\star_{t+1}(f)\left(\mathcal{V}_{t+1}(0) - \chi_{t+1}(f)\right) + \left(1 - \mathbb{I}^\star_{t+1}(f)\right)\mathcal{V}_{t+1}(x').$$

Take conditional expectations given x', and use the definition of the hazard:

$$\mathbb{E}_{t+1}\left[V_{t+1}(f) \mid x'\right] = h_{t+1}(x')\left(\mathcal{V}_{t+1}(0) - \chi_{t+1}\right) + \left(1 - h_{t+1}(x')\right)\mathcal{V}_{t+1}(x'),$$

obtained by taking the expectation of the indicator form term by term: $\mathbb{E}_{t+1}[\mathbb{I}^\star_{t+1}(f) \mid x'] = h_{t+1}(x')$; the expected menu cost paid upon adjustment is $\mathbb{E}_{t+1}\left[\mathbb{I}^\star_{t+1}(f)\,\chi_{t+1}(f) \mid x'\right] = h_{t+1}(x')\,\chi_{t+1}$, with χ_{t+1} denoting the mean cost conditional on adjustment; and $\mathbb{E}_{t+1}\left[\left(1 - \mathbb{I}^\star_{t+1}(f)\right)\mathcal{V}_{t+1}(x') \mid x'\right] = \left(1 - h_{t+1}(x')\right)\mathcal{V}_{t+1}(x')$ because $\mathcal{V}_{t+1}(x')$ is x'-measurable. Substituting this into (A.56) (and using the law of iterated expectations) yields the recursive representation in condition (i) of Lemma 1.

Finally, consider a stationary environment, in which $x_t^\star \equiv x^\star$ and the conditional law of motion of $x_{t+1}(f)$ given $x_t(f)$ is time-homogeneous. Conditions (i) and (iii) then define a per-period map sending $\mathcal{V}_{t+1}$ to $\mathcal{V}_t$ whose primitives—the profit function $\Pi(x^\star + \cdot)$, the discount factor β, the menu-cost distribution G, the free-adjustment probability $h(0)$, and the transition kernel of x'—are all time-invariant.

Call that map T. It is monotone—$\max\{\cdot,\cdot\}$ and conditional expectation both preserve the pointwise order—and it discounts: adding a constant c to $\mathcal{V}_{t+1}$ raises both branches of the maximum by c, and hence raises $T\mathcal{V}_{t+1}$ by exactly βc. These are Blackwell's two sufficient conditions.[1] In the weighted supremum norm $\|\cdot\|_w$ (with the usual treatment of the unbounded

[1] Blackwell's sufficient conditions deliver a contraction on a space of *bounded* functions, which is not available here: the quadratic flow loss evaluated along an unbounded random walk leaves $\mathcal{V}$ unbounded below—this does not conflict

support), its fixed point is therefore unique and time-independent, $\mathcal{V}_t(x) \equiv \mathcal{V}(x)$. Absent aggregate shocks, the reset gap and the common displacement vanish. Since $\Pi(x)$ is even and the idiosyncratic innovation is symmetric, the continuation problem is invariant to replacing every gap by its negative. Hence, if $\mathcal{V}(x)$ is a fixed point, so is $\mathcal{V}(-x)$. Uniqueness of the fixed point therefore implies $\mathcal{V}(-x) = \mathcal{V}(x)$.

Substituting into the hazard (5) (with $h_t(0) \equiv h(0)$ time-invariant) gives $h_t(x) = h(0) + \big(1 - h(0)\big)\, G\big(\mathcal{V}(0) - \mathcal{V}(x)\big) \equiv h(x)$, which establishes the final claim of the Lemma. Additionally, Equation (5) then implies $\mathcal{V}(-x) = \mathcal{V}(x) \Rightarrow h(-x) = h(x)$. □

C.2 Proof of Proposition 1

Proof of Proposition 1 (Quadratic approximation of the hazard rate). Fix t and condition on the event $\{x_t(f) = x\}$. From condition (iii) of Lemma 1 (the optimality principle),

$$V_t(f) = \max\big\{\mathcal{V}_t(x_t(f)),\ \mathcal{V}_t(0) - \chi_t(f)\big\},$$

so a firm that does not receive a free adjustment pays the menu cost and resets whenever the value gain from adjusting weakly exceeds the cost,

$$\chi_t(f) \le \Delta_t\big(x_t(f)\big), \qquad \Delta_t(x) := \mathcal{V}_t(0) - \mathcal{V}_t(x). \tag{A.58}$$

The firm may instead adjust "for free" with the steady probability $h_t(0) \ge 0$—zero in a pure menu-cost model, or a positive mass point as in the Nakamura and Steinsson (2010) Calvo-plus model. Conditioning on whether a free adjustment arrives, and using that the menu cost $\chi_t(f) \sim G$ is i.i.d. and independent of $x_t(f)$, the law of total probability gives the hazard:

$$h_t(x) := \mathbb{E}_t\big[\mathbb{I}_t^{\star}(f) \mid x_t(f) = x\big] = h_t(0) + \big(1 - h_t(0)\big)\, G\big(\Delta_t(x)\big), \tag{A.59}$$

which is Equation (5). Since $\Delta_t(0) = \mathcal{V}_t(0) - \mathcal{V}_t(0) = 0$, Equation (A.59) confirms that $h_t(0)$ is the adjustment probability at the optimum.

Since $\mathcal{V}_t$ is twice differentiable at 0 with the interior optimum $\mathcal{V}_t'(0) = 0$ (condition (ii) of Lemma 1), a Taylor expansion at 0 yields

$$\mathcal{V}_t(x) = \mathcal{V}_t(0) + \mathcal{V}_t'(0)\, x + \tfrac{1}{2}\mathcal{V}_t''(0)\, x^2 + o(x^2) = \mathcal{V}_t(0) + \tfrac{1}{2}\mathcal{V}_t''(0)\, x^2 + o(x^2), \tag{A.60}$$

with the low-heterogeneity condition maintained on the ergodic cross-section F_t, which bounds the dispersion of gaps across firms in the population, not the range of a single firm's value function evaluated off that ergodic path along its own unbounded random walk. We therefore work in the weighted supremum norm $\|\mathcal{V}\|_w := \sup_x |\mathcal{V}(x)|/w(x)$ with weight $w(x) := M + x^2$, for a constant $M > 0$ chosen below. Because $x' = x + d + \epsilon$ off the adjustment event, $\mathbb{E}\big[w(x') \mid x\big] = w(x) + 2xd + d^2 + \sigma_\epsilon^2 \le (1+\eta)\, w(x) + C_\eta$ for every $\eta > 0$, with $C_\eta := d^2/\eta + d^2 + \sigma_\epsilon^2$; since $w \ge M$, this gives $\mathbb{E}\big[w(x') \mid x\big]/w(x) \le 1 + \eta + C_\eta/M$, and choosing first η and then M so that $\beta\big(1 + \eta + C_\eta/M\big) < 1$ makes T map the space of w-bounded functions into itself and contract on it with that modulus.

so that the value gain from adjusting is

$$\Delta_t(x) = \mathcal{V}_t(0) - \mathcal{V}_t(x) = -\tfrac{1}{2}\mathcal{V}_t''(0)\,x^2 + o(x^2). \tag{A.61}$$

Because $\mathcal{V}_t(0)$ is the maximized value, $\Delta_t(x) \geq 0$ for all x, so G is only ever evaluated at nonnegative arguments and the one-sided expansion at zero suffices. With $G(0) = 0$ and G admitting a right derivative $G'(0) := \lim_{y\downarrow 0} G(y)/y \geq 0$ (Assumption 5),

$$G(y) = G'(0)\,y + o(y) \qquad (y \downarrow 0). \tag{A.62}$$

By (A.61), $\Delta_t(x) = O(x^2) \to 0$ as $x \to 0$, hence $o\big(\Delta_t(x)\big) = o(x^2)$. Substituting (A.61) into (A.62) and the result into (A.59) gives

$$h_t(x) = h_t(0) + \big(1 - h_t(0)\big)\Big(-\tfrac{1}{2}\,G'(0)\,\mathcal{V}_t''(0)\,x^2\Big) + o(x^2) = h_t(0) + \phi_t\,x^2 + o(x^2),$$

with the curvature coefficient

$$\phi_t := -\tfrac{1}{2}\big(1 - h_t(0)\big)\,G'(0)\,\mathcal{V}_t''(0),$$

exactly as in Proposition 1. The coefficient is non-negative: G is a distribution function, so $G'(0) \geq 0$; $\mathcal{V}_t''(0) < 0$ by condition (ii) of Lemma 1; and $0 \leq 1 - h_t(0) \leq 1$. Hence $\phi_t \geq 0$.

Finally, in a stationary environment $\mathcal{V}_t''(0)$, G, and $h_t(0)$ are time-invariant (last statement of Lemma 1), so $h_t(0) = h(0)$ and $\phi_t = \phi = -\tfrac{1}{2}\big(1 - h(0)\big)G'(0)\,\mathcal{V}''(0)$, which are the constant coefficients of Equation (14). □

C.3 Proof of Theorem 1

Proof. We first prove the projection and control its remainder. By the definition of the hazard and the law of iterated expectations,

$$\pi_t = \mathbb{E}_t\big[\mathbb{I}_t^{\star}(f)x_t(f)\big] = \int z\,h(z)\,dF_t(z).$$

Write

$$h(z) = h(0) + \phi z^2 + r(z), \qquad r(z) = z^2\omega(z), \qquad \omega(z) \longrightarrow 0 \quad \text{as } z \longrightarrow 0,$$

and let

$$R_t := \int z\,r(z)\,dF_t(z).$$

It follows that

$$\begin{aligned}\pi_t &= h(0)\int z\,dF_t(z) + \phi\int z^3\,dF_t(z) + R_t\\ &= h(0)x_t + \phi\left(x_t^3 + 3x_t\sigma_t^2 + \psi_t\right) + R_t.\end{aligned}$$

It remains to show that $R_t = o(\sigma_t^3)$. Let $X_t \sim F_t$ and define

$$Z_t := \frac{X_t - x_t}{\sigma_t}, \qquad m_t := \frac{x_t}{\sigma_t}, \qquad Y_t := \frac{X_t}{\sigma_t} = Z_t + m_t.$$

Assumption *4a* implies that $|m_t| \le K$ and that Z_t has a uniformly bounded $(3+\varepsilon)$ moment. Therefore, Y_t also has a uniformly bounded $(3+\varepsilon)$ moment, and the family $\{|Y_t|^3\}$ is uniformly integrable. Moreover,

$$\frac{R_t}{\sigma_t^3} = \mathbb{E}\left[Y_t^3\,\omega(\sigma_t Y_t)\right].$$

Since h is a probability, $h(z) \in [0,1]$ for every z, so $\omega(z) = \left(h(z) - h(0) - \phi z^2\right)/z^2 \to -\phi$ as $|z| \to \infty$; ω is continuous, so $\sup_z |\omega(z)| < \infty$. For every $A > 0$,

$$\begin{aligned}\frac{|R_t|}{\sigma_t^3} \le &\sup_{|z|\le\sigma_t A} |\omega(z)|\;\mathbb{E}\,|Y_t|^3\\ &+ \sup_z |\omega(z)|\,\mathbb{E}\left[|Y_t|^3\mathbb{I}\{|Y_t| > A\}\right].\end{aligned}$$

For fixed A, the first term converges to zero as $\sigma_t \downarrow 0$. Uniform integrability makes the second term arbitrarily small, uniformly in t, by letting $A \to \infty$. Hence, $R_t = o(\sigma_t^3)$, which proves Equation (17).

We next prove that inflation is an odd function of the history of mean gaps. Lemma 1 implies that the stationary hazard is even, $h(-z) = h(z)$, and the idiosyncratic innovation is symmetric. We want to show by induction that, for every integrable function g,

$$\int g(z)\,dF_t[-\boldsymbol{x}^t](z) = \int g(-z)\,dF_t[\boldsymbol{x}^t](z). \tag{A.63}$$

At the initial symmetric point, this follows from the symmetry of F_0 and from the fact that the initial aggregate displacement changes sign under $\boldsymbol{x}^t \mapsto -\boldsymbol{x}^t$.

For the induction step, suppress time subscripts and use primes to denote one-period-ahead objects. Let $F(z) \equiv F[\boldsymbol{x}^t](z)$ be the current distribution under $\boldsymbol{x}^t$ and $F^-(z) \equiv F[-\boldsymbol{x}^t](z)$ the current distribution under $-\boldsymbol{x}^t$. Similarly, denote with a minus superscript the variables along the reflected history. Suppose that

$$\int g(z)\,dF^-(z) = \int g(-z)\,dF(z)$$

for every integrable g by the induction hypothesis. The one-period transition of the gap is

$$X' = (1 - \mathbb{I}^\star)X + \epsilon' + d', \tag{A.64}$$

where d' is the common aggregate displacement of the mean gap. Since the mean of the next cross-section is x',

$$d' = x' - \int (1 - h(z))z \, dF(z). \tag{A.65}$$

Under the reflected history, the next mean is $x'^{,-} = -x'$. Moreover, the induction hypothesis and the evenness of h imply

$$\begin{aligned}\int (1 - h(z))z \, dF^-(z) &= \int (1 - h(-z))(-z) \, dF(z) \\ &= -\int (1 - h(z))z \, dF(z).\end{aligned}$$

Hence the common displacement in the reflected economy is

$$\begin{aligned}d'^{,-} &= -x' - \int (1 - h(z))z \, dF^-(z) \\ &= -x' + \int (1 - h(z))z \, dF(z) = -d'.\end{aligned}$$

Now, couple firms with current gaps X and $-X$. Since $h(-X) = h(X)$, they can be assigned the same adjustment realization $\mathbb{I}^\star$. Since ϵ' is symmetric, the innovation in the reflected economy can be represented by $-\epsilon'$. Therefore,

$$\begin{aligned}X'^{,-} &= (1 - \mathbb{I}^\star)(-X) - \epsilon' - d' \\ &= -\left[(1 - \mathbb{I}^\star)X + \epsilon' + d'\right] = -X'.\end{aligned}$$

Thus the next-period distributions are also reflections of one another. Induction proves Equation (A.63) for every t and for any evolution of the mean gap.

Finally, apply Equation (A.63) to $g(z) = zh(z)$. This gives the required step:

$$\begin{aligned}\pi_t[-\boldsymbol{x}^t] &= \int zh(z) \, dF_t[-\boldsymbol{x}^t](z) \\ &= \int (-z)h(-z) \, dF_t[\boldsymbol{x}^t](z) \\ &= -\int zh(z) \, dF_t[\boldsymbol{x}^t](z) = -\pi_t[\boldsymbol{x}^t],\end{aligned}$$

where the third equality uses $h(-z) = h(z)$. Therefore, inflation is odd. □

C.4 Proof of Theorem 2

For notational convenience, write $h \equiv h(0)$ and $H \equiv 1 - h$.

In this section, we work under the strong regularity Assumption *4c*. The proof proceeds in three steps. Step 1 (Lemma 2) obtains the dynamics of the variance under the uniform bound on the higher-order moments. Step 2 (Lemma 3) obtains the dynamics of the skewness under the uniform bound on the higher-order moments. Step 3 (Theorem 3) derives a more general version of Theorem 2 with time-varying coefficients for the first-order dynamics of the mean gap. Theorem 2 is the special case $\tilde{a}_t = \lambda^t$.

Step 1: IRF of the variance. The following lemma derives the evolution of the variance along the IRF. Derivations are obtained under the more general first-order dynamics for the mean gap with a time-varying coefficient, $x_t = \tilde{a}_t \delta + o(\delta)$, $\tilde{a}_t \neq 0$.

Lemma 2 (IRF of the variance)**.** *The pre-adjustment variance satisfies along the IRF:*

$$\sigma_t^2 = \sigma_0^2 + \tilde{b}_t \delta^2 + o(\delta^2), \tag{A.66}$$

where

$$\tilde{b}_0 = 0, \tag{A.67}$$

$$\tilde{b}_{t+1} = H\tilde{b}_t + hH\tilde{a}_t^2, \tag{A.68}$$

$$\tilde{b}_t = hH \sum_{j=0}^{t-1} H^{t-1-j} \tilde{a}_j^2, \tag{A.69}$$

and $\tilde{a}_t$ is the first-order impulse-response coefficient of x_t.

Proof. The impact aggregate displacement is common across firms and therefore leaves the variance unchanged, so $\tilde{b}_0 = 0$. Substituting $x_t = \tilde{a}_t \delta + o(\delta)$ and (A.66) into (A.40) (Proposition 4) and using Assumption *4c* gives

$$\tilde{b}_{t+1}\delta^2 = \left(H\tilde{b}_t + hH\tilde{a}_t^2\right)\delta^2 + o(\delta^2),$$

which proves (A.68). Iterating from $\tilde{b}_0 = 0$ gives (A.69). □

Step 2: IRF of the skewness. The following lemma derives the evolution of the skewness along the IRF. Derivations are obtained under the more general first-order dynamics for the mean gap with a time-varying coefficient, $x_t = \tilde{a}_t \delta + o(\delta)$, $\tilde{a}_t \neq 0$.

Lemma 3 (IRF of the skewness)**.** *The pre-adjustment third central moment satisfies along the IRF:*

$$\psi_t = \sigma_0^2 \tilde{q}_t \delta + \tilde{r}_t \delta^3 + o(\delta^3), \tag{A.70}$$

where

$$\tilde{q}_0 = 0, \qquad \tilde{r}_0 = 0, \tag{A.71}$$

$$\tilde{q}_{t+1} = H\tilde{q}_t + 3hH\tilde{a}_t, \tag{A.72}$$

$$\tilde{r}_{t+1} = H\tilde{r}_t + 3hH\tilde{a}_t\tilde{b}_t + hH(2h-1)\tilde{a}_t^3, \tag{A.73}$$

$$\tilde{q}_t = 3hH \sum_{j=0}^{t-1} H^{t-1-j}\tilde{a}_j, \tag{A.74}$$

$$\tilde{r}_t = hH \sum_{j=0}^{t-1} H^{t-1-j}\left[3\tilde{a}_j\tilde{b}_j + (2h-1)\tilde{a}_j^3\right], \tag{A.75}$$

and $\tilde{a}_t$ is the first-order impulse-response coefficient of x_t.

Proof. The impact aggregate displacement leaves the third central moment unchanged, and the steady-state distribution has zero skewness, so $\tilde{q}_0 = \tilde{r}_0 = 0$. Substituting

$$x_t = \tilde{a}_t\delta + o(\delta), \qquad \sigma_t^2 = \sigma_0^2 + \tilde{b}_t\delta^2 + o(\delta^2),$$

and (A.70) into (A.42) (Proposition 5) and using Assumption *4c* gives:

$$\begin{aligned}\psi_{t+1} = {} & \sigma_0^2\left(H\tilde{q}_t + 3hH\tilde{a}_t\right)\delta \\ & + \left[H\tilde{r}_t + 3hH\tilde{a}_t\tilde{b}_t + hH(2h-1)\tilde{a}_t^3\right]\delta^3 + o(\delta^3).\end{aligned}$$

Although $\sigma_0^2\delta$ and δ^3 have the same joint low-heterogeneity order, the two components are separated by linearity of the recursion in ψ_t: the first is the response to the forcing $3hH\sigma_0^2 x_t$, and the second is the response to the remaining cubic forcing. This gives (A.72) and (A.73). Iterating from $\tilde{q}_0 = \tilde{r}_0 = 0$ gives (A.74) and (A.75). □

Step 3: Derivation of the IRF—general case. We now derive the theorem stated in the body of the paper (Theorem 2), under the general first-order dynamics for the mean gap with a time-varying coefficient, $x_t = \tilde{a}_t\delta + o(\delta)$, $\tilde{a}_t \neq 0$.

Theorem 3 (Nonlinear impulse response functions with general coefficient)**.** *At every fixed horizon for which* $\tilde{a}_t \neq 0$*,*

$$h_t = \bar{h} + \alpha_{2,t} x_t^2 + o(\delta^2), \tag{A.76}$$

$$\pi_t = \pi_0 + \beta_{1,t} x_t + \beta_{3,t} x_t^3 + o(\delta^3), \tag{A.77}$$

where

$$\bar{h} = h + \phi\sigma_0^2 + o(\sigma_0^2), \tag{A.78}$$

$$\alpha_{2,t} = \phi\left(1 + \frac{\tilde{b}_t}{\tilde{a}_t^2}\right), \tag{A.79}$$

$$\beta_{1,t} = h + \phi\sigma_0^2\left(3 + \frac{\tilde{q}_t}{\tilde{a}_t}\right), \tag{A.80}$$

$$\beta_{3,t} = \phi\left(1 + 3\frac{\tilde{b}_t}{\tilde{a}_t^2} + \frac{\tilde{r}_t}{\tilde{a}_t^3}\right), \tag{A.81}$$

and $\tilde{b}_t, \tilde{q}_t, \tilde{r}_t$ *are given by* (A.69), (A.74)*, and* (A.75)*. Moreover,* $\pi_0 = o(\sigma_0^3)$*.*

Proof. At a fixed horizon with $\tilde{a}_t \neq 0$,

$$x_t = \tilde{a}_t\delta + o(\delta) \qquad \Longrightarrow \qquad \delta = \frac{x_t}{\tilde{a}_t} + o(\delta).$$

Lemma 2 therefore gives

$$\sigma_t^2 = \sigma_0^2 + \frac{\tilde{b}_t}{\tilde{a}_t^2} x_t^2 + o(\delta^2). \tag{A.82}$$

Similarly, Lemma 3 gives

$$\psi_t = \sigma_0^2 \frac{\tilde{q}_t}{\tilde{a}_t} x_t + \frac{\tilde{r}_t}{\tilde{a}_t^3} x_t^3 + o(\delta^3). \tag{A.83}$$

The substitution error in the linear skewness component is $o(\sigma_0^2\delta) = o(\delta^3)$ because $\sigma_0^2 = O(\delta^2)$. Substituting (A.82) into (15),

$$h_t = h + \phi\sigma_0^2 + \phi\left(1 + \frac{\tilde{b}_t}{\tilde{a}_t^2}\right) x_t^2 + o(\delta^2),$$

which proves (A.76) and (A.79).

Substituting (A.82) and (A.83) into (17),

$$\begin{aligned}\pi_t &= hx_t + \phi x_t^3 + 3\phi x_t\left(\sigma_0^2 + \frac{\tilde{b}_t}{\tilde{a}_t^2}x_t^2\right) \\ &\quad + \phi\left(\sigma_0^2\frac{\tilde{q}_t}{\tilde{a}_t}x_t + \frac{\tilde{r}_t}{\tilde{a}_t^3}x_t^3\right) + o(\delta^3) \\ &= \left[h + \phi\sigma_0^2\left(3 + \frac{\tilde{q}_t}{\tilde{a}_t}\right)\right]x_t \\ &\quad + \phi\left[1 + 3\frac{\tilde{b}_t}{\tilde{a}_t^2} + \frac{\tilde{r}_t}{\tilde{a}_t^3}\right]x_t^3 + o(\delta^3).\end{aligned}$$

This proves (A.77)–(A.81). At the symmetric zero-trend operating point the steady-state inflation rate is zero; under the local hazard formulation the weaker statement $\pi_0 = o(\sigma_0^3)$ is sufficient. □

Note that at $\lambda = H$, $\lambda^2 = H$, $\lambda^3 = H$ and $\lambda = 1$, where a denominator of the closed forms in Theorem 2 vanishes, the coefficients are defined by continuity. Specifically, the limits are finite. For example, $\lim_{\lambda\to H}\frac{(H/\lambda)^t-1}{H-\lambda} = \frac{t}{H}$, so that $\beta_{1,t} \to h + 3\phi\sigma_0^2(1+ht)$; $\lim_{\lambda^2\to H}\frac{(H/\lambda^2)^t-1}{H-\lambda^2} = \frac{t}{H}$, so that $\alpha_{2,t} \to \phi(1+ht)$; and $\lim_{\lambda\to 1}\frac{(H/\lambda^3)^t-(H/\lambda^2)^t}{1-\lambda} = tH^t$. The case $\lambda^3 \to H$ is analogous. The limits coincide with the finite sums of Theorem 3.

D Estimation and historical shock decomposition

This appendix discusses how the model is solved, how the structural innovations are recovered from the data, how the pricing primitives are estimated, and how the historical shock decomposition of Section 5 is constructed.

Solution and state space. We solve the model by third-order perturbation around its deterministic steady state, with the k-order solver of Dynare (Schmitt-Grohé and Uribe 2004; Adjemian et al. 2024). Collect the model's predetermined variables in s_t, the three structural innovations in $\varepsilon_t = (\varrho_t, \varsigma_t, \varkappa_t)'$, and the three observables in y_t. The solution is a pair of polynomials of degree three,

$$s_t = g(s_{t-1}, \varepsilon_t), \qquad y_t = Z(s_{t-1}, \varepsilon_t) + v_t, \qquad \varepsilon_t \sim \mathcal{N}(0, Q), \quad v_t \sim \mathcal{N}(0, R),$$

with Q diagonal and R the diagonal covariance matrix of the measurement errors. Demeaning the observables sets the constant of the measurement equation to zero.[2]

[2]The policy rate is the effective federal funds rate for the United States and, for the euro area, the overnight interbank rate: EONIA until 2019Q3 and the euro short-term rate thereafter, both as quarterly averages. We treat it as missing

Table A.1: Estimated parameters: priors and order-1 posteriors (US and euro area)

		US (CPI)		Euro area (HICP)	
Parameter		Prior	Posterior	Prior	Posterior
Shock innovation standard deviations					
Cost-push	σ_{ϱ}	Inv.Gamma(0.02, 0.02)	0.024	Inv.Gamma(0.02, 0.02)	0.028
Monetary	σ_{ς}	Inv.Gamma(0.02, 0.02)	0.007	Inv.Gamma(0.02, 0.02)	0.006
Demand	$\sigma_{\varkappa}$	Inv.Gamma(0.02, 0.02)	0.014	Inv.Gamma(0.02, 0.02)	0.019
Measurement-error standard deviations					
Inflation	σ_{π}^{me}	Inv.Gamma(0.5, 0.3)	0.245	Inv.Gamma(0.5, 0.3)	0.210
Output growth	$\sigma_{\Delta y}^{me}$	Inv.Gamma(1, 0.5)	0.540	Inv.Gamma(1, 0.5)	0.675
Policy rate	σ_{i}^{me}	Inv.Gamma(0.5, 0.3)	0.319	Inv.Gamma(0.5, 0.3)	0.315
Shock persistences					
Cost-push	ρ_{u}	Beta(0.7, 0.05)	0.820	Beta(0.7, 0.05)	0.779
Monetary	ρ_{ι}		0.200		0.300
Demand	ρ_{b}	Beta(0.1, 0.05)	0.457	Beta(0.1, 0.05)	0.407

Notes. First-order Bayesian estimation (Kalman filter) on the three demeaned observables (the policy rate is missing at the effective lower bound), at the SMM estimates of the pricing primitives. The prior columns give the distribution with its (mean, standard deviation); the posterior columns give the mode (point estimate). The monetary persistence ρ_{ι} is fixed at 0.20 for the US and 0.30 for the euro area, as discussed in Section 5. Because ρ_{ι} is fixed, the modes are not an unconstrained posterior maximum, so we report them without standard errors.

The third-order step: tempered particle filter. At third order neither the filtering distribution nor the likelihood is Gaussian (Fernández-Villaverde and Rubio-Ramírez 2007; Aruoba et al. 2017). A single measurement update at this sample length concentrates nearly all of the weight on a handful of particles, so we temper it into stages, following Herbst and Schorfheide (2019). Each period runs in four steps.

1. *Propagation.* For every particle we draw $\varepsilon_t^{(i)} \sim \mathcal{N}(0, Q)$ and push $(s_{t-1}^{(i)}, \varepsilon_t^{(i)})$ through g and Z.

2. *Tempered update.* The measurement density is introduced in $K = 10$ equal increments, so that at each stage every particle's weight is multiplied by $\exp\{-\frac{1}{2K}\, v^{(i)\prime} R^{-1} v^{(i)}\}$, where $v^{(i)} = y_t - Z(s_{t-1}^{(i)}, \varepsilon_t^{(i)})$ is the particle's prediction error, recomputed at each stage from the particle's current state and innovation. Whenever the effective sample size of the accumulated weights falls below $N/2$, the cloud is resampled systematically and rejuvenated by three random-walk Metropolis steps on $\varepsilon_t^{(i)}$, whose acceptance ratio carries the tempered

in the quarters at the effective lower bound, in which the federal funds rate averages at most 0.25 percent or the deposit facility rate of the ECB is at or below zero: 2009Q1–2015Q4 and 2020Q2–2022Q1 in the United States, and 2012Q4–2022Q3 in the euro area. In those quarters, both filters update on inflation and output growth alone. We demean the policy rate over the quarters in which it is observed.

likelihood and the innovation's own density.

3. *Likelihood.* The period's contribution is the sum across stages of the log of the average incremental weight, each stage averaged under the weights entering it, which telescopes to the log of the period's average full weight. We add the Gaussian normalizing constant, $-\frac{1}{2}\left(n_y \log 2\pi + \log\det R\right)$ with $n_y = 3$ observables, so that the reported log-likelihood is normalized.

4. *Resampling.* The cloud is resampled systematically at the end of the period, so that every period starts from equally weighted particles.

The filter is initialized at the covariance the first-order solution of the model implies, which solves the discrete Lyapunov equation in the stationary directions of the state. In the direction the unit root in the price level leaves undetermined, it receives a diffuse block instead, capped at a finite value so that the initial weights do not degenerate.

Smoothing. Two-sided estimates come from a fixed-lag particle smoother (Kitagawa 1996; Doucet and Johansen 2011). The filter carries the last $L+1$ periods of particle states and innovations and reindexes them through the ancestor lineage of every resampling step, so the slice stored for date τ and read at date t is an equally weighted sample from $p(s_\tau, \varepsilon_\tau \mid Y_{1:t})$. The slice is frozen once $t = \tau + L$, and the final L periods are read from the last buffers, which condition on the whole sample. We set $L = 16$ quarters.

We read the smoother for four stationary objects: the price gap, its variance, the frequency of adjustment, and the innovations themselves. At the estimated persistences, dates more than four years ahead carry negligible information about them, so the frozen slice is their smoothed distribution up to that memory.

The construction runs no backward pass and therefore needs no transition density, the smoothed moments being particle averages of forward draws. That matters here, because the alternative, backward simulation (Godsill et al. 2004), requires the density of $s_{\tau+1}$ given s_τ, which is supplied only on a three-dimensional subspace by three innovations driving a fourteen-dimensional state.

Estimation of the pricing primitives. We calibrate σ_0 and estimate the steady-state frequency $\bar{h} = h(0) + \phi\sigma_0^2$ and the curvature share $\alpha = \phi\sigma_0^2/\bar{h}$, which map into $h(0) = (1-\alpha)\bar{h}$ and $\phi = \alpha\bar{h}/\sigma_0^2$. The model moments are computed from the smoothed frequency $\mathbb{E}[h_t \mid Y_{1:T}]$, averaged across the eight seeds of the particle smoother. The first moment, m_1, is its mean over 2010Q1–2019Q4.

The second, m_2, is its maximum over the surge window, 2021Q1–2023Q1 in the United States and 2021Q3–2023Q1 in the euro area.

On the data side, the first target is the mean of the micro series over the same quarters. For the United States, the series is the weighted median frequency of Cotton and Garga (2026), a 12-month trailing average. The peak target $h^* = 0.35$ is a judgment on the quarterly peak net of the change in sample composition (Section 5.1), so we compare it with the model's quarterly peak. For the euro area, we remove month-of-year effects, estimated over 2010–2019, from the monthly frequency of Gautier et al. (2026) before compounding it to quarters, so that January repricing does not set the peak. The peak target is the maximum of the adjusted series over the surge window.

Each evaluation of the moments re-estimates the order-1 posterior mode and runs the eight particle smoothers, so the simulated moments are costly to compute and carry simulation noise. We therefore evaluate them on a design of points around the solution and fit, by weighted least squares, a local response surface for each moment in $\theta = (\bar{h}, \alpha)$, quadratic in α. The estimate $\hat{\theta}$ solves the two fitted moment conditions. We add points near the current root and refit until the moments simulated at the root match their targets within simulation error.

With as many moments as parameters, the weighting matrix drops out, and the variance of the estimate is given by:

$$V(\hat{\theta}) = G^{-1}\,\Omega\,(G^{-1})' + V_{\text{sim}},$$

where G is the slope of the fitted response surface at $\hat{\theta}$, Ω the variance of the data moments, and V_{sim} the simulation variance of the root. The two data moments use different windows of the micro series, so Ω is diagonal. Its first element is the Newey–West long-run variance, with eight lags, of the quarterly gap between the micro series and the smoothed frequency over 2010–2019, divided by the number of quarters. Its second element adds two terms. The first is the uncertainty of the peak target: for the United States, the variance of a uniform distribution on $[0.343, 0.371]$, a band around the judgmental target; for the euro area, the quarterly noise of the series. The second is the posterior variance of the frequency in the peak quarter t^*, $\phi^2\,\mathrm{Var}(\sigma^2_{t^*} + x^2_{t^*} \mid Y_{1:T})$, because the observed peak is one draw from the posterior distribution of h_{t^*}, whereas the model moment is its mean. Standard errors of $h(0)$, ϕ, the steady-state frequency $\bar{h}$, and the first-order slope of the Phillips curve follow by the delta method. All of them condition on σ_0 and on the observed macro history.

The delta method treats the moments as linear in θ, whereas the peak moment flattens as α rises. We therefore also compute a 95% confidence set for ϕ, obtained by test inversion. For each α, the statistic

$$Q(\alpha) = \min_{\bar{h}} \sum_{k=1}^{2} \frac{\left(\hat{m}_k(\bar{h}, \alpha) - m_k^{\text{data}}\right)^2}{\Omega_{kk} + v_k(\bar{h}, \alpha)}$$

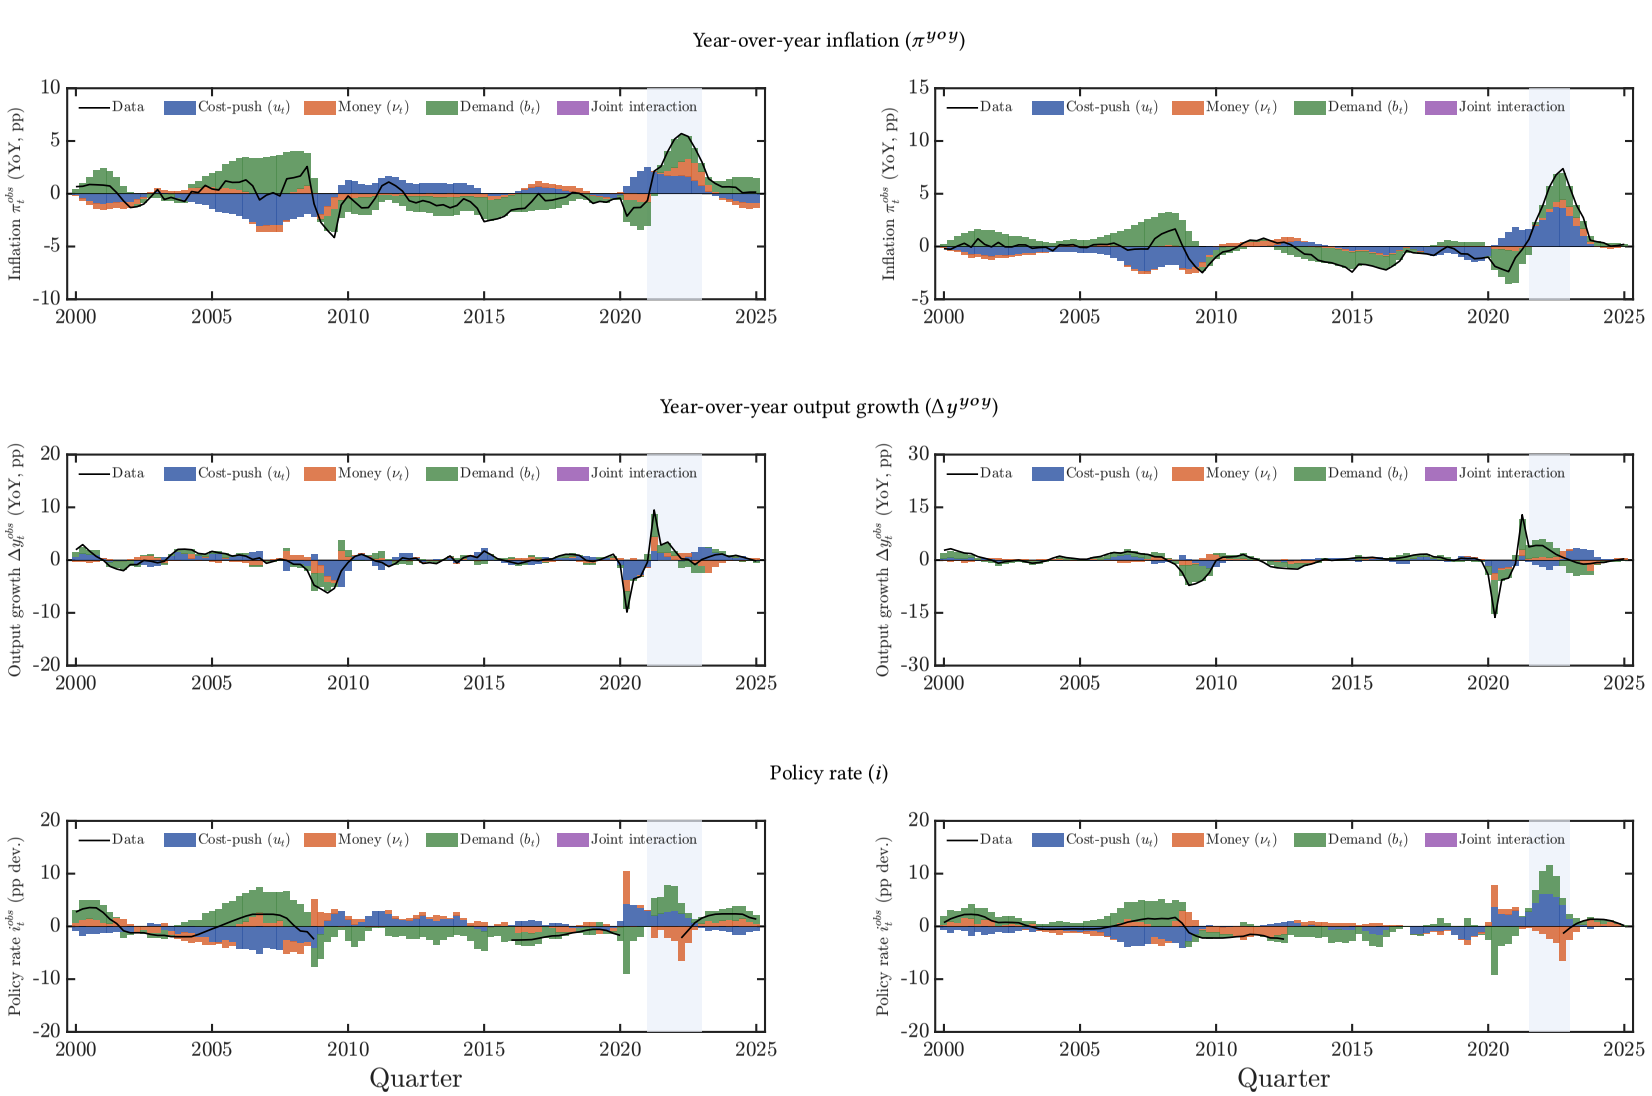
profiles out $\bar{h}$, where $\hat{m}_k$ is the fitted response surface, v_k its fit variance, and Ω_{22} is evaluated at the ϕ implied by $(\bar{h}, \alpha)$. The set collects the values of α at which $Q(\alpha)$ does not exceed the 95% critical value of a $\chi^2(1)$ distribution, mapped into ϕ along the minimizing $\bar{h}$. The set is $[13.8,\ 34.5]$ in the euro area and $[8.1,\ 165.1]$ in the United States. Both US ends lie at the bounds of the search range for α, beyond the evaluated points, where the response surface is extrapolated. In the euro area, a separate range of α, outside the set, also passes the test.

Structural shocks and historical shock decomposition. The historical decompositions of Figure 4 and Table 2 use the same construction at both first and third order. For each shock k, we set k's innovations—Kalman-smoothed or particle-smoothed—to their estimated values, hold every other shock at zero, and propagate the resulting single-shock path through the decision rule of the model that recovered the innovations: the order-1 rule for the Kalman-smoothed innovations and the order-3 rule for the particle-smoothed ones. The comparison therefore measures the move from the linear to the nonlinear model. This gives a counterfactual path $\text{contrib}_k(t)$—what the observable would have done if shock k alone had hit the economy—and the share reported in Table 2 is $S(k) = \sum_{t \in W} \text{contrib}_k(t) / \sum_{t \in W} \pi_t^{obs}$, the component's share of inflation over the surge window W. The single-shock paths and the fitted series use the smoothed means of the innovations. Because the order-3 decision rule is nonlinear, a path generated by the smoothed means need not equal the posterior mean of that path, so the decomposition should be interpreted as a "plug-in exercise," applied identically to both columns.

This construction need not sum to 100% across shocks, in either column, for two reasons. The order-3 decision rule does not satisfy superposition, so the three single-shock paths no longer sum to the fitted series even before measurement error is considered. The gap is the model's own cross-shock interaction terms, active only when multiple shocks move together, which a one-shock-at-a-time construction cannot assign to any single disturbance. The decomposition is exact for the order-1 rule, so the interaction term is zero in the Kalman column. And the fitted series differs from the data, by measurement error and because it is built from point estimates of the innovations. Because the shares are linear in the paths, naming the wedges closes the accounting exactly: writing $y_t^{obs} = \sum_k \text{contrib}_k(t) + I(t) + FG(t)$, with I the cross-shock interaction in the fitted series and FG the gap between the data and the fitted series, gives $\sum_k S(k) + S_I + S_{FG} = 100\%$ as an identity. Table 2 reports S_I and S_{FG} as the joint interaction and data minus joint.

Figure A.1 reports the historical decomposition implied by the first-order model, with the Kalman-smoothed innovations propagated through the order-1 decision rule. It is the counterpart of Figure 4 and of the Kalman column of Table 2. As we discuss in the paper, cost-push shocks play a much larger role in rationalizing inflation during the recent surge.

Figure A.1: Historical shock decomposition in the first-order model



Notes. As Figure 4, for the first-order model: the shocks are the Kalman-smoothed innovations at the estimated parameters, propagated through the order-1 decision rule. The joint interaction is zero at first order. The omitted series are the same as in Figure 4. Each series is plotted in deviation from its sample mean.

References

Stéphane Adjemian, Michel Juillard, Frédéric Karamé, Willi Mutschler, Johannes Pfeifer, Marco Ratto, Normann Rion, and Sébastien Villemot. Dynare: Reference manual, version 6. Dynare Working Papers 80, CEPREMAP, 2024.

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

S. Borağan Aruoba, Luigi Bocola, and Frank Schorfheide. Assessing DSGE model nonlinearities. *Journal of Economic Dynamics and Control*, 83:34–54, 2017.

Christopher D. Cotton and Vaishali Garga. The price-change statistics we've weighted for. Working Paper 26-6, Federal Reserve Bank of Boston, 2026.

Arnaud Doucet and Adam M. Johansen. A tutorial on particle filtering and smoothing: Fifteen years later. In Dan Crisan and Boris Rozovskii, editors, *The Oxford Handbook of Nonlinear Filtering*, pages 656–704. Oxford University Press, Oxford, 2011.

Jesús Fernández-Villaverde and Juan F. Rubio-Ramírez. Estimating macroeconomic models: A likelihood approach. *Review of Economic Studies*, 74(4):1059–1087, 2007.

Erwan Gautier, Cristina Conflitti, Daniel Enderle, Ludmila Fadejeva, Alex Grimaud, Eduardo Gutiérrez, Valentin Jouvanceau, Jan-Oliver Menz, Alari Paulus, Pavlos Petroulas, Pau Roldan-Blanco, and Elisabeth Wieland. Consumer price stickiness in the euro area during an inflation surge, 2026.

Simon J. Godsill, Arnaud Doucet, and Mike West. Monte Carlo smoothing for nonlinear time series. *Journal of the American Statistical Association*, 99(465):156–168, 2004.

Edward Herbst and Frank Schorfheide. Tempered particle filtering. *Journal of Econometrics*, 210 (1):26–44, 2019.

Genshiro Kitagawa. Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics*, 5(1):1–25, 1996.

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

Stephanie Schmitt-Grohé and Martín Uribe. Solving dynamic general equilibrium models using a second-order approximation to the policy function. *Journal of Economic Dynamics and Control*, 28(4):755–775, 2004.