

Nominal Rigidities and Monetary Policy in the Euro Area *

Guido Ascari[†] Alexandre Carrier[‡] Emanuel Gasteiger[§]
Alex Grimaud[¶] Gauthier Vermandel^{||}

October 27, 2026

Abstract

We study monetary policy in a nonlinear New Keynesian model with state-dependent price- and wage-adjustment frequencies. The extensive margin of nominal adjustment makes shock propagation sign-, size-, and state-dependent. Inflationary shocks generally increase firms' incentives to reprice, whereas wage reoptimization responds only when shocks sufficiently affect desired real wages and labor-market conditions. Estimating the model on euro area data from 1999Q1 to 2024Q4, we find a pronounced price-wage asymmetry. During the 2021–2022 surge, price adjustment increased sharply, while wage inflation mainly reflected delayed real-wage catch-up. Consequently, monetary policy could reduce inflation at lower output cost.

Keywords: New Keynesian Phillips Curve, nonlinearity, inflation, monetary policy.

JEL codes: C51, E31, E47, E52.

***Disclaimer:** The views expressed are those of the authors and do not necessarily reflect those of the Banque de France, De Nederlandsche Bank (DNB), the European Central Bank (ECB), the Oesterreichische Nationalbank (OeNB), and the Eurosystem. This paper contains research conducted within the network "Challenges for Monetary Policy Transmission in a Changing World Network" (ChaMP). More information is provided on its [website](#). This paper previously circulated under the title "*Monetary Policy in the Euro Area, When Phillips Curves... Are Curves*".

[†]Department of Economics and Management, University of Pavia, Pavia, IT and De Nederlandsche Bank, Amsterdam, NL and CEPR.

[‡]European Central Bank, Frankfurt, DE.

[§]Institute of Statistics and Mathematical Methods in Economics, TU Wien, Vienna, AT and Business Research Unit, Instituto Universitário de Lisboa, Lisboa, PT. Financial support from the [Austrian National Bank \(OeNB, grant no. 18611\)](#) and Fundação para a Ciência e a Tecnologia (grant no. UIDB/00315/2020, DOI: [10.54499/UIDB/00315/2020](#)) is gratefully acknowledged.

[¶]Monetary Policy Section, Oesterreichische Nationalbank, Vienna, AT and Institute of Statistics and Mathematical Methods in Economics, TU Wien, Vienna, AT.

^{||}CMAP, École polytechnique, Palaiseau, FR; Universités PSL & Paris-Dauphine, Paris, FR.

‘The relation between unemployment and the rate of change of wage rates is therefore likely to be highly nonlinear.’

(Phillips, 1958, p. 283)

1. Introduction

The 2021–2022 inflation surge across advanced economies has renewed fundamental questions about the mechanisms governing nominal price and wage adjustment. Inflationary pressures were initially concentrated in prices, but were followed by a marked acceleration in wage inflation, raising concerns about the transmission from price shocks to wages and about the risk of second-round effects – i.e. the so called prices-to-wages pass-through. Yet inflation subsequently declined without a deep recession in most advanced economies. It appears that both the European Central Bank and the Federal Reserve have managed a so-called “*soft landing*”, stabilizing inflation while avoiding a sharp contraction in real activity. Understanding how price and wage adjustment interact in high-inflation environments is therefore central to assessing the propagation of large shocks and the stabilization trade-offs faced by monetary policy.

A natural interpretation is that nominal rigidities respond endogenously to economic conditions, but not necessarily in the same way for prices and wages. During inflationary episodes, firms may adjust prices more frequently, activating the extensive margin of price setting and allowing large shocks to pass through more rapidly. The wage response may instead be more delayed and depend on the nature of the shock: wage inflation can rise through catch-up dynamics even when the frequency of wage reoptimization changes only modestly. This view is supported by a growing body of microeconomic evidence documenting state-dependent price adjustment. Firms are found to adjust prices more frequently during inflation surges, consistent with a nonlinear activation of the extensive margin of price setting (see, e.g., [Gautier et al., 2026](#); [Montag and Vallenas, 2023](#)). Evidence on wage-setting frequencies is more limited, but recent work suggests that wage adjustment may also become more frequent in high-inflation environments ([Gödl and Gödl-Hanisich, 2024](#)). Taken together, this evidence points to state-dependent nominal adjustment, while leaving open whether price and wage adjustment margins are activated simultaneously or with the same intensity.

Consistent with these micro-level findings, macroeconomic data also point to sign and size dependence in inflation dynamics. Figure 1 presents evidence from local projections for the euro area, showing that inflation responses depend both on the sign *and* size of oil supply shocks. Panel A illustrates sign dependence: inflationary supply shocks generate stronger and more persistent inflationary responses than disinflationary shocks of comparable magnitude. Panel B highlights size dependence: larger shocks induce disproportionately

stronger inflation responses, indicating a departure from linearity in shock size. Related evidence from [Bobeica et al. \(2025\)](#) and [Cavallo et al. \(2024\)](#) shows that large energy shocks can trigger faster price adjustment and stronger pass-through, consistent with state-dependent pricing mechanisms. Taken together, the micro- and macro-level evidence suggests that inflation dynamics are state-dependent. In particular, large inflationary shocks appear more likely to move the economy into a region in which pass-through is stronger and nominal adjustment margins become more relevant.




Figure 1: Inflation Impulse Response Functions to a one standard deviation supply shock

Note: Monthly local projections for year-on-year HICP (Harmonised Index of Consumer Prices) inflation spanning 1997M1-2024M9. For each horizon $h = 0, 1, \dots, H$, we estimate $\pi_{t+h} = \alpha_h + \sum_{\ell=0}^{11} (g_h(z_{t-\ell}) + \delta_{\ell,h} \pi_{t-\ell}) + \varepsilon_{t+h}$, where π_t denotes year-on-year inflation and z_t is the monthly oil supply shock constructed by [Känzig \(2021\)](#). The function $g_h(\cdot)$ captures horizon-specific nonlinearities. Following [Carravella and Martínez-Bruera \(2024\)](#), Panel A (sign dependence) uses $g_h(z) = \beta_{1,h}z + \beta_{2,h}z|z|$, while Panel B (size dependence) uses $g_h(z) = \beta_{1,h}z + \beta_{2,h}z^2$. Dashed lines report the estimate from a linear benchmark $g_h(z) = \beta_{1,h}z$. Shaded area reports the 90% and 68% confidence intervals. Appendix A provides further details.

These patterns are difficult to reconcile with the standard estimated New Keynesian (NK) model, which remains the workhorse framework for monetary policy analysis. In its canonical formulation, the NK model relies on time-dependent pricing and wage-setting mechanisms, such as Calvo contracts, under which adjustment occurs with a constant probability independent of economic conditions ([Calvo, 1983](#)). These mechanisms do not typically generate the sign and size dependence documented in Figure 1, since the degree of nominal rigidity is fixed rather than responding to the state of the economy. Moreover, once estimated, standard NK models typically imply relatively flat and state-invariant price and wage Phillips curves, so that inflation dynamics are largely attributed to exogenous supply shocks, while demand shocks and monetary policy play a more limited role (see, e.g., [Fratto and Uhlig, 2020](#)). In such an environment, the inflation-output stabilization trade-off is approximately state-invariant, and disinflation tends to require sizable and persistent output losses.

Against this background, we develop a tractable New Keynesian model in which both price- and wage-adjustment frequencies can vary endogenously with economic conditions. The framework extends the state-dependent price-setting mechanism of [Gasteiger and Grimaud \(2023\)](#) to a joint price–wage setting environment, using the rational-inattention foundations of [Matějka and McKay \(2015\)](#) to model the extensive margin of nominal adjustments. This allows the model to distinguish between the activation of the price and wage extensive margins and the subsequent adjustment of wages.

An important implication of the framework is that the source of inflation determines which nominal-adjustment margin is activated. For a firm, higher inflation directly erodes its relative price, widening the gap between the inherited and optimal price and thereby raising the gains from repricing. Price-setting state dependence therefore arises for a broad range of inflationary shocks, largely irrespective of their source. By contrast, for a union, the gap between its inherited and optimal wage depends not only on inflation, but also on labor demand and the marginal rate of substitution. The wage gap therefore depends more closely on the underlying source of inflation. Consequently, similar increases in inflation can generate markedly different incentives to reoptimize wages depending on the shock driving inflation, making wage adjustment more shock-contingent.

Beyond determining which margin is activated, the framework makes shock transmission sign-, size-, and state-dependent. Size dependence arises because the extensive margin responds nonlinearly to the gains from reoptimization, so that larger inflationary shocks can generate more-than-proportional changes in adjustment frequencies. Sign dependence arises because these gains are asymmetric: inflationary shocks can activate the extensive margin more strongly than comparable disinflationary shocks. This asymmetry is consistent with recent micro-level evidence showing that price increases are more likely than equally sized price decreases (see, e.g., [Karadi, Schoenle and Wursten, 2024](#); [Luo and Villar, 2021](#)).¹ State dependence arises because the effect of a given shock also depends on inherited price and wage gaps and on the corresponding adjustment frequencies. The stabilization trade-off therefore varies with the state of the economy: in high-inflation states, greater nominal flexibility can increase the traction of monetary policy on inflation relative to real activity.

We take this mechanism to euro area data and make three contributions. First, we implement a nonlinear solution and estimation strategy, since a linearized version of the model would suppress the nonlinear propagation channels that we want to quantify. Following [Adjemian and Juillard \(2024\)](#) and [Sahuc et al. \(2025\)](#), and building on the extended path method originally proposed by [Fair and Taylor \(1983\)](#), we estimate the nonlinear model using an inversion filter that treats the model itself as the data-generating process. This approach exploits information in aggregate macroeconomic data to investigate how price and wage ad-

¹While many state-dependent pricing models generate nonlinear dynamics, they typically imply responses that are symmetric across signs in the absence of additional mechanisms (see, e.g., [Blanco et al., 2024](#); [Ghassibe and Nakov, 2025](#); [Karadi, Schoenle and Wursten, 2024](#)).

justment frequencies behaves endogenously with the business cycle. It therefore allows us to assess not only whether state-dependent adjustment is quantitatively relevant, but also which shocks activate it and how it changes historical and policy counterfactuals. Moreover, despite exploiting only aggregate macroeconomic data, the estimated share of re-optimized prices closely aligns with the share of price changes in the euro area microdata, as reported by [Gautier et al. \(2026\)](#), providing external validation to our framework (see the end of Section 4.3 and Figure G.3 in the Appendix.)

Second, our main quantitative finding is that state-dependent nominal adjustment mattered quantitatively during the euro-area inflation surge, but that price and wage adjustment margins were not activated mechanically in parallel. In our estimated macro model, the extensive-margin response during the episode is predominantly price-led: inflationary shocks strongly increase firms' incentives to reprice, whereas wage reoptimization responds more selectively and depends on the source of the shock and its labor-market effects. Wage inflation nevertheless rises persistently through delayed real-wage catch-up and price-to-wage transmission, despite only limited direct activation of the wage-reoptimization margin. Locally approximated Phillips-curve slopes confirm this asymmetry: the price Phillips curve is flat, but steepens markedly with high inflation, whereas the wage Phillips-curve slope is shock-dependent and much less closely related to either price or wage inflation. The state-dependent specification also attains a substantially higher posterior objective at the mode than the corresponding fixed-frequency model. Notably, this improvement is concentrated in the 2021–2024 inflation episode, when our model rationalizes the data with substantially less unlikely cost-push and labor-supply innovations than a standard Calvo model.

Finally, we use the estimated model to quantify how endogenous price and wage adjustment shape shock propagation and monetary-policy trade-offs over the euro area business cycle. The relevant adjustment margin depends on the shock. Cost-push shocks operate primarily through firms' marginal costs and become more inflationary in states in which price-setting frequencies are elevated. Monetary policy shocks also affect labor demand and desired reset wages and can therefore activate both price- and wage-setting margins. In the estimated high-inflation states, the endogenous response of both margins makes inflation more responsive to monetary policy relative to real activity, lowering the output cost of disinflation.

The resulting state dependence of monetary-policy transmission therefore reflects not only stronger price pass-through, but also the shock-contingent interaction between price and wage adjustment margins.

Related literature. Our paper is closely connected to the very large literature on state-dependent nominal rigidities that has flourished over the past two decades with the growing availability of micro-level data on prices – less so on wages. It has generated several seminal contributions that have enhanced our understanding of the distribution, dynamics and drivers of price changes. A comprehensive survey is beyond the scope of this paper; see [Costain and](#)

[Nakov \(2024\)](#) for a recent and thorough overview. Of course, our modeling approach, cannot – and does not aim at – reproducing the distribution of price (or wage) changes. It provides a short-cut way to take into account some insights from this recent literature, and make them operational in estimating large NK models.

Early menu-cost models imply strong selection effects, whereby price adjustments following aggregate shocks are concentrated among the most misaligned firms (e.g., [Golosov and Lucas, 2007](#)). However, a growing body of micro-level evidence suggests that this is not a dominant feature of observed price-setting behavior. [Costain and Nakov \(2024\)](#) stress this point and argue in favor of a shift toward a more general class of state-dependent pricing frameworks in which adjustment probabilities rise smoothly with firms’ incentives to adjust. In these models, state-dependence operates primarily through movements in the gross extensive margin – i.e., the increased fraction of adjusting prices – rather than through a strong selection effect – i.e., large price adjustment from firms more distant from their desired price – implying larger real effects of monetary policy. Recent evidence on the distribution of price changes ([Alvarez et al., 2016](#); [Midrigan, 2011](#)) and on adjustment hazards conditional on price gaps ([Carvalho and Kryvtsov, 2021](#); [Gagliardone et al., 2025](#); [Karadi, Schoenle and Wursten, 2024](#)) supports this view. Our modeling choice is designed to capture this smooth state-dependent behavior in adjustment frequencies and should therefore be interpreted as abstracting from selection effects, in line with the empirical literature.

Moreover, state-dependent models, which match patterns from microdata, share many of the macro implications of the Calvo model when the latter is accordingly recalibrated.² However, the key difference identified by [Costain and Nakov \(2024\)](#) between the state-dependent models and the standard Calvo model is that the former imply rapid and flexible price changes in some circumstances, as for high inflation ([Alvarez et al., 2019](#)), for large shocks ([Cavallo et al., 2024](#); [Karadi and Reiff, 2019](#)), making the Phillips curve slope state-dependent ([Karadi, Nakov, Nuno, Pasten and Thaler, 2024](#)).

While admittedly using a short-cut relative to a fully microfounded state-dependent model that features a selection effect, we provide a tractable way to incorporate state-dependent price- *and* wage- setting into the standard model. “Tractable way” has a very precise meaning for us: a way that allows us to estimate a nonlinear quantitative NK model with state-dependent price- *and* wage-setting with a nonlinear method. This is particularly important as fully microfounded state-dependent models, while able to match the micro data, are still computationally quite far away from being incorporated in a nonlinearly estimated NK model. To the best of our knowledge, this paper is the first one that includes state-dependent price- *and* wage-setting in a nonlinear NK model, nonlinearly estimated on aggregate euro

²As shown by [Auclert et al. \(2024\)](#), the first-order approximation of a wide class of state-dependent models shares many macro-implications with the Calvo framework. This motivates our focus on a fully nonlinear environment, where such first-order equivalence breaks down and state-dependence becomes quantitatively relevant.

area data, to provide policy analysis and counterfactuals.

From this same perspective, two papers are particularly close to ours, but along different dimensions. [Costain et al. \(2022\)](#) is close in allowing for state-dependent adjustment of both prices and wages. Their framework is, however, based on control costs and is calibrated to micro evidence on price and wage changes, whereas our goal is to develop a tractable extensive-margin mechanism that can be embedded in and nonlinearly estimated within a quantitative New Keynesian model on aggregate euro area data. In addition, our empirical application emphasizes a different mechanism: price and wage state-dependence need not be activated by the same shocks. Among tractable frameworks with an endogenous extensive margin estimated for aggregate policy analysis, [Blanco et al. \(2024\)](#) is also close to ours. Developed independently and in parallel, they propose a different but related mechanism where the fraction of price changes evolves endogenously, as in state-dependent models, while retaining the tractability of time-dependent models for empirical and policy analysis. At the same time, the two approaches are complementary since they abstract from state-dependent wage setting and from additional elements – such as habit persistence or price and wage indexation – that have been shown to be important for empirical analysis. Moreover, their approach estimates a more limited subset of parameters using US data. By contrast, our framework incorporates these features while being applied to euro area data.

[Reiter and Wende \(2024\)](#) is another recent paper that proposes a yet different mechanism, based on a generalization of the Rotemberg pricing model. The mechanism is based on the assumption that marginal cost follows a sigmoid function, and it is then embedded in a calibrated DSGE model to show that it has aggregate implications similar to the menu cost models in [Costain and Nakov \(2019\)](#).

Our paper also relates to recent studies on optimal monetary policy in nonlinear calibrated state-dependent models. [Karadi, Nakov, Nuno, Pasten and Thaler \(2024\)](#) find that central banks should respond to major cost-push shocks with a more aggressive anti-inflation policy, since the more favorable stabilization trade-off reduces the economic cost of controlling inflation. Additionally, they find that in the case of an aggregate productivity shock, which does not create any trade-off, the optimal policy emphasizes strict price stability. In a multisector economy with menu costs, [Caratelli and Halperin \(2025\)](#) find that optimal policy should, instead, allow inflation to move after sectoral productivity shocks, with inflation and output moving in opposite direction. The optimal policy prescription in their model entails nominal wage targeting, i.e., nominal wages, rather than inflation, should be stabilized, despite wages being flexible.

Finally, our work relates to a large empirical research dealing with nonlinear inflation dynamics and its implications for monetary policy. Some studies provide empirical evidence of nonlinearity in the aggregate response of inflation to shocks, that are consistent with state-dependent price- and wage-setting (see, e.g., [Alvarez et al., 2017](#); [Ascari and Haber, 2022](#);

Cavallo et al., 2024; de Veirman, 2023). Nonlinearity in the Phillips curve could arise for different reasons than state-dependent pricing. Harding et al. (2022, 2023) show that the nonlinear standard Calvo model with a Kimball aggregator gives rise to a nonlinear Phillips curve, and then study the implications for monetary policy under de- or accelerating inflation. Erceg et al. (2024) extend this framework by allowing for endogenous indexation of prices and wages, whereby the share of backward-looking price and wage setting by firms and households increases with the level and persistence of above-target inflation. Benigno and Eggertsson (2023) and Gitti (2025) provide evidence suggesting that the Phillips curve is strongly nonlinear in the vacancy-to-unemployment ratio, as a measure of labor market tightness. Faber and Züllig (2025) present similar evidence for a panel of 18 euro area countries, using unemployment as measure of labor market slack. However, Beaudry et al. (2024, 2025) argue that this evidence is fragile, and that nonlinearities may only reflect shifts in short-run inflation expectations. This is in line with evidence for the US in Gorodnichenko and Coibion (2025), and for Europe in Acharya et al. (2023).

2. The model

In this section, we first describe the main features of the model in Subsection 2.1 and then describe the nonlinear solution in Subsection 2.2.

2.1 Structure of the model

The model is a standard New Keynesian model with no capital, consumption habits, where wages and prices are sticky and they are both indexed to lagged inflation. The model departs from the standard Calvo (1983) setup by allowing state-dependent sticky prices and wages. The state-dependency allows for intensive *and* extensive margin adjustments in prices and wages. Given that the model is otherwise standard, its detailed description and derivations are confined in Appendix B, while here we just sketch its structure and focus on how we introduce state-dependency into the model.

The lifetime utility of household features multiplicative external consumption habits and disutility from labor. The representative household owns the firms and receives dividends. Two exogenous AR(1) processes, $\varepsilon_{w,t}$ and $\varepsilon_{d,t}$, are introduced to generate variations in distortionary taxes—or subsidies—on wages and on bond income. These processes are introduced for estimation purposes. They effectively represent a labor supply shock ($\varepsilon_{w,t}$) and a demand shock ($\varepsilon_{d,t}$) creating a wedge between the central bank’s policy rate and the return on assets held by households, similar to Smets and Wouters’s (2007) risk-premium shock.

Production is divided into several layers. A representative wholesale firm hires labor from the labor packer and produces output, which is sold to a continuum of retail firms at price $P_{w,t}$, subject to exogenous AR(1) cost-push ($\varepsilon_{p,t}$). Each retail firm $i \in [0, 1]$ repackages wholesale

output, and sells it, at the price $P_t(f)$, to a competitive final goods firm which transforms retail output into a final output good according to a standard CES function, sold at the price P_t .

Retailers are price-setters. Following [Gasteiger and Grimaud \(2023\)](#), the price-setting problem is:

$$\max_{P_t(f)} \mathbb{E}_t \sum_{j=0}^{\infty} \left(\prod_{k=0}^j \phi_{p,t+k} \right) \phi_{p,t}^{-1} \Lambda_{t,t+j} \times \\ \times \left\{ P_t(f)^{1-\epsilon_p} (P_{t+j}/\Pi_{t,t+j}^{\epsilon_p})^{\epsilon_p-1} Y_{t+j} - P_{w,t+j} P_t(f)^{-\epsilon_p} (P_{t+j}/\Pi_{t,t+j}^{\epsilon_p})^{-\epsilon_p} Y_{t+j} \right\},$$

where $\Lambda_{t,t+j}$ is the stochastic discount factor of the household between period t and $t+j$, Y_t is final output good output that determines the demand for retail goods, $p_{w,t} = P_{w,t}/P_t$ represents retailers' real marginal costs, $0 \leq \epsilon_p < 1$ is a parameter governing indexation to past inflation and:

$$\Pi_{t,t+j} \equiv \begin{cases} \frac{P_{t+1}}{P_t} \times \frac{P_{t+2}}{P_{t+1}} \times \dots \times \frac{P_{t+j}}{P_{t+j-1}} & \text{for } j = 1, 2, \dots \\ 1 & \text{for } j = 0. \end{cases}$$

This price-resetting problem is the same as the one in a standard Calvo framework, with one key difference: the probability of not-being able to adjust the price in a future period—and so the implied price-resetting frequency in equilibrium—is time-varying, i.e., $\phi_{p,t+k}$ but independent to the firm (f) decision. All the price-resetting firms are symmetric and face the same problem. So we can define the optimal reset price as $P_t^\#$, because it is the same for all the firms and it does not depend on the index f . The first-order condition is recursively given by

$$x_{1,t} = p_{w,t} Y_t + \mathbb{E}_t \phi_{p,t+1} \Lambda_{t,t+1} \left(\frac{\pi_{t+1}}{\pi_t^{\epsilon_p}} \right)^{\epsilon_p} x_{1,t+1} \\ x_{2,t} = Y_t + \mathbb{E}_t \phi_{p,t+1} \Lambda_{t,t+1} \left(\frac{\pi_{t+1}}{\pi_t^{\epsilon_p}} \right)^{\epsilon_p-1} x_{2,t+1} \\ p_t^\# = \frac{\epsilon_p}{\epsilon_p - 1} \frac{x_{1,t}}{x_{2,t}}$$

with $p_t^\# = P_t^\#/P_t$. Again, these equations defined the first-order condition also in the standard model, the only difference being the Calvo probability of not-adjusting being time-varying.

Price-resetting agency. In the price-setting problem above, firms recognize that the probability of not adjusting prices is time-varying but independent of their own decisions. Consequently, this time variation cannot depend on firm-specific variables and must instead be driven solely by aggregate conditions. This assumption is key for tractability.

We formalize this mechanism by introducing a price-resetting agency that authorizes firms

to adjust prices. Firms that are not authorized must keep their existing prices (which we allow to be partially indexed automatically). All firms comply with the agency's decision, so in each period the agency effectively chooses the fraction of firms that are unable to reset prices, that is, the non-resetting probability $\phi_{p,t}$. The agency is owned by the representative household, which also owns the continuum of intermediate-goods firms. As a result, the agency internalizes firm profits, and any net payoff is rebated lump-sum to households.

We assume that it is costly for the agency to process information about individual firms, including their prices and their profit implications of price changes. As a consequence, the agency allocates price-adjustment permissions randomly across firms. In making its decision, the agency evaluates the expected continuation value of the future profits. Firms granted permission to reset prices share a common expected continuation value, $V_{p,t}^\#$, net of a menu cost of changing price, τ_p . By contrast, firms that are not allowed to reset prices have heterogeneous continuation values, reflecting differences in their inherited prices. However, because permissions are allocated randomly, the agency uses as a proxy the expected continuation value of future profits of a non-adjusting firm that keeps its price equal to the aggregate price level in the previous period. The agency then chooses the optimal fraction of firms permitted to adjust prices, taking into account a cost of processing information.

Formally, the agency would solve the following problem. It chooses $\phi_{p,t}$ to maximize

$$\max_{\phi_{p,t}} \left\{ \phi_{p,t} V_{p,t}^f + (1 - \phi_{p,t})(V_{p,t}^\# - \tau_p) - \frac{1}{\gamma_p} \Omega(\phi_{p,t}) \right\},$$

where $\Omega(\phi_{p,t})$ is an information-processing cost and $\gamma_p > 0$ is the inverse marginal cost of information processing. Under an entropy-based information cost and a symmetric prior, [Matějka and McKay \(2015, Theorem 1\)](#) shows that the first-order condition yields the logit policy

$$\phi_{p,t} = \frac{\exp(\gamma_p V_{p,t}^f)}{\exp(\gamma_p V_{p,t}^f) + \exp(\gamma_p (V_{p,t}^\# - \tau_p))}. \quad (1)$$

In each period t , the agency authorizes a fraction $(1 - \phi_{p,t})$ of firms to reset their price to $P_t^\#$, while the remaining firms have to keep their old prices (partially indexed to the average inflation). Equation (1) implies that the non-resetting share $\phi_{p,t}$ is a smooth, state-dependent function of the difference between the continuation values of non-adjustment and adjustment. The fixed burden $\tau_p > 0$ shifts the adjustment value $V_{p,t}^\# - \tau_p$ downward and therefore raises $\phi_{p,t}$, all else equal. The intensity of choice parameter $\gamma_p > 0$ governs the strength of state contingency. When γ_p is small, information is expensive and the agency responds weakly to movements in $V_{p,t}^f$ and $V_{p,t}^\#$; adjustment becomes nearly random and $\phi_{p,t}$ is close to a constant, delivering the Calvo benchmark as $\gamma_p \rightarrow 0$. As γ_p increases, information becomes cheaper, the agency conditions more sharply on the state, and $\phi_{p,t}$ varies more with the underlying present values. In the limit $\gamma_p \rightarrow \infty$, choices become nearly deterministic:

the agency puts almost all probability on the higher-value action, so price adjustment becomes effectively state-dependent. For any finite γ_p , adjustment remains probabilistic and prices are therefore rigid in the sense that a positive mass of firms does not re-optimize each period.³

Labor market. The labor market structure mirrors the one of production and consists of three layers. There is a continuum of labor unions index by $l \in [0, 1]$. They purchase labor from households at MRS_t and repackage it for sale to a representative labor packer $W_t(l)$. The labor packer combines differentiated labor inputs into final labor, $L_{d,t}$, using a CES technology with elasticity of substitution $\epsilon_w > 1$. Labor unions are wage setters, and their problem is symmetric with respect to the one just described for firms. In each period t , a wage-setting agency chooses the fraction of unions that can re-optimize their wage, but not which unions adjust. The wage-setting agency authorizes a, randomly chosen, fraction $(1 - \phi_{w,t})$ of unions to reset their wage to $W_t^\#$, while the remaining unions must keep their old wages. Resetting entails a fixed non-pecuniary cost $\tau_w > 0$, while non-reset wages are partially indexed to past inflation. The agency evaluates the expected continuation value of adjusting unions, $V_{w,t}^\#$, and that of a representative non-adjusting union that has the wage equals to the average wage from the previous period, $V_{w,t}^f$. Under an entropy-based information cost, the implied non-resetting share follows the multinomial-logit form (Matějka and McKay, 2015):

$$\phi_{w,t} = \frac{\exp(\gamma_w V_{w,t}^f)}{\exp(\gamma_w V_{w,t}^f) + \exp(\gamma_w (V_{w,t}^\# - \tau_w))},$$

where $\gamma_w > 0$ is the intensity of choice, i.e., the inverse marginal cost of information processing.

Finally, monetary policy follows a standard Taylor rule and features a monetary policy shock ε_t^r . All shocks follow stationary AR(1) processes.

³Similar devices to our assumed agency have been widely employed in the literature to circumvent heterogeneity and ensure tractability. Classic examples include the indivisible labor assumption (see Hansen, 1985; Rogerson, 1988) and the household–union (insurance) construct commonly used in New Keynesian models—as our labor packer, see below—with heterogeneity in wages or labor supplies across households (see Woodford, 2011). A closely related logic also appears in Blanco et al. (2024). They assume gigantic multiproduct firms that produce all the products in the economy and are able to choose how many prices to adjust in each period, but not which ones; price adjustments are then assigned randomly across products. In this sense, their multiproduct firm plays exactly the same role as our agency. By aggregating firms in this way, the model abstracts from the selection effects that arise when adjustment decisions depend on firm-level variables, allowing adjustment shares to be treated as probabilities without tracking the cross-sectional composition of adjusters. Our model differs from theirs in the way the problem is formulated. In particular, we follow the rational inattention approach to discrete choice developed in Matějka and McKay (2015), which results in a different first-order condition governing the evolution of the fraction of price-resetting firms.

2.2 The extended path solution method

This paper explores the nonlinearities of the Phillips curve and the interactions between the intensive and extensive margins of inflation. A first-order linearization is inadequate for approximating the model, as the state-dependent adjustment frequencies for prices and wages generate essential nonlinearities and asymmetries in inflation and wage responses to shocks, requiring a fully nonlinear estimation approach. To accurately capture these nonlinearities, we solve the model using the extended path method developed by [Fair and Taylor \(1983\)](#), following the recent implementation by [Adjemian and Juillard \(2024\)](#); [Sahuc et al. \(2025\)](#). This method provides a nonlinear solution under rational expectations, responding to MIT-type shocks without aggregate uncertainty.⁴ Specifically, agents are assumed to be surprised by contemporaneous shocks, but consistently expect future shocks to be zero, consistent with rational expectations. The numeric solution can be expressed as follows:

$$y_t = \mathbb{E}_{t,t+S} \{g_{\Theta}(y_{t-1}, \bar{y}, \eta_t)\}, \quad (2)$$

where y_t denotes the vector of endogenous variables, $g_{\Theta}(\cdot)$ is the nonlinear transition function implied by the structural equations, Θ represents the set of structural parameters, and shocks are normally distributed as $\eta_t \sim N(0, \Sigma_{\eta})$. The expectation operator $\mathbb{E}_{t,t+S}$ implies a sequence-space approximation of expected fluctuations over a horizon of S periods, beyond which the economy has fully returned to steady state $\bar{y}$. We choose a sequence-space horizon of 10 years ($S = 40$ quarters), as extending this horizon further does not significantly alter the likelihood. Importantly, this numerical solution explicitly accounts for endogenous variations in the Calvo adjustment probabilities, which are central to our model.

3. Monetary policy under nonlinear Phillips curves

We begin by calibrating the model to match a small-scale setup⁵ to build intuition about the role of endogenous versus exogenous price and wage setting frequencies. Section 3.1 contrasts the shape of both the price and wage New Keynesian Phillips curves when the economy is hit by demand and supply shocks. Section 3.2 compares state-dependent IRFs in

⁴Explicit integration over future uncertainty would offer limited additional accuracy given the CRRA specification of the utility function and does not justify the significant computational cost involved.

⁵We assume that the discount factor is set to $\beta = 0.995$, the curvature of consumption utility is set to $\sigma = 2$, and the one of labor disutility to $\chi = 1$. The Taylor rule follows standard values for the euro area, with coefficients $\{\theta_{\pi}, \theta_y, \rho\} = \{2, 0.125, 0.8\}$. The gross steady-state inflation rate is set to $\bar{\pi} = 1.005$. The elasticities of demand for individual goods and labor are set to $\varepsilon_p = \varepsilon_w = 9$. The steady-state shares of firms and unions not optimally resetting prices and wages are calibrated to $\{\phi_p, \bar{\phi}_w\} = \{0.75, 0.85\}$, consistent with standard values used in the literature on Calvo pricing. We follow [Gasteiger and Grimaud \(2023\)](#) and set the state-dependency parameters to $\gamma_p = \gamma_w = 10$. The autoregressive coefficients of the exogenous shocks are set to $\rho_d = \rho_p = \rho_w = 0.7$, while monetary policy shocks are assumed to be purely i.i.d. Finally, to simplify the analysis and focus on the role of state-dependency, we shut down most sources of endogenous persistence by setting $h = \varrho_p = \varrho_w = 0$.




Figure 2: **Asymmetry in the price-setting decision.** The blue line depicts firm f 's present value as a function of its relative price $p(f) = P(f)/P$, holding the aggregate environment at steady state. Point A is the optimal relative price. Points B and C correspond to relative prices that are respectively 10% below and 10% above the optimum. The double-headed arrows measure the present-value loss relative to the optimum. The loss from being underpriced at B is substantially larger than the loss from being equally overpriced at C.

response to different shocks, depending on the value of the underlying inflation processes and shows how the sign and the size of shocks matter for their propagation.

3.1 Implied Phillips curves and intuition

This subsection provides the intuition behind wage and price setting in the model, emphasizing the differences between fixed and endogenous adjustment frequencies. We show that endogenous price- and wage-setting frequencies can give rise to nonlinear New Keynesian Phillips curves, with stronger adjustment incentives when firms and unions are pushed toward relatively low inherited prices and wages. This result arises from an asymmetry in the price- and wage-setting decisions. For a given updating cost, the expected benefit of resetting can be substantially larger when an inherited price or wage is below its desired level than when it is comparably above it.

Asymmetry in price and wage setting decisions. Figure 2 depicts firm f 's present value as a function of its relative price $p(f) = P(f)/P$. The depicted present value is evaluated in partial equilibrium, holding all other aggregate variables at their steady-state values. Point A corresponds to the optimal relative price. At point B the firm is 10% underpriced relative to this optimum, while at point C it is 10% overpriced.

The present value exhibits a strong asymmetry, which arises from the nonlinear profit function under [Dixit and Stiglitz \(1977\)](#) monopolistic competition. Low relative prices imply rela-

tively large reductions in the present value. The opportunity cost is mainly in terms of profit per unit: an underpriced firm sells a large quantity at an insufficient markup. In contrast, high relative prices imply smaller reductions in the present value, with the opportunity cost mainly taking the form of foregone quantity. Hence, despite B and C being equally distant from the optimal relative price in percentage terms, the present-value loss at B is substantially larger than at C .

The figure is evaluated at steady state to isolate this asymmetry, but the same intuition applies conditional on any state of the economy. Aggregate conditions affect the location and scale of the present-value function through demand, marginal costs, and expectations. Given these conditions, however, a firm whose inherited relative price is pushed sufficiently below its desired reset price faces a particularly large cost of not adjusting. This asymmetry therefore weakens the incentive to cut prices when firms become relatively overpriced and strengthens the incentive to increase prices when they become relatively underpriced.

Positive trend inflation further reinforces this mechanism. When a firm does not reset its nominal price, an increase in the aggregate price level gradually erodes its relative price. If the firm is initially overpriced, trend inflation therefore moves its relative price toward the desired price and reduces the present-value cost of not adjusting. If the firm is already underpriced, by contrast, the same erosion pushes it farther down the steep left-hand side of the present-value function and increases the cost of remaining inactive. Positive trend inflation thus further amplifies the asymmetry between upward and downward price-adjustment incentives.

Moreover, because Figure 2 depicts the present value as a function of the relative price, any increase in the aggregate price level, all else equal, reduces $p(f) = P(f)/P$. An inflationary shock can therefore move a non-adjusting firm toward the underpriced region of the present-value function and sharply increase the gain from resetting its price. The opposite effect is weaker for a disinflationary shock, since moving toward the overpriced side generates a smaller present-value loss. Through Equation (1), this difference in the value of non-adjustment translates directly into asymmetric movements in the price-resetting frequency.

The wage-setting problem follows the same basic present-value logic. The corresponding present-value function for unions is reported in Appendix Figure C.1. As in the price-setting problem, being stuck with a relative wage below the optimal one is more costly than an equally sized upward deviation. In general equilibrium, however, the desired reset wage is determined by labor demand and the marginal rate of substitution. As shown below, this makes the activation of the wage extensive margin more dependent on the source of the shock than the corresponding price-setting margin.




Figure 3: The source of the fluctuation matters. Shocks with the same impact on inflation trigger different price and wage adjustment margins.

Notes: IRFs to monetary-policy, labor-supply, and cost-push shocks. All shock sizes are normalized such that on impact they generate a +0.25 percentage-point increase in quarterly inflation.

Extensive-margin activation across shocks. Figure 3 compares monetary-policy, labor-supply, and cost-push shocks scaled to raise quarterly inflation by 0.25 percentage point on impact.⁶ This normalization holds the initial inflation response fixed and isolates how its source shapes adjustment. All three shocks widen firms’ reset-price gaps and increase the price-resetting frequency. However, wage adjustment exhibits important differences: the monetary-policy

⁶The model includes four structural shocks for four observables later used in the estimation, thereby avoiding stochastic singularity. We omit the risk-premium shock from Figure 3 because it generates qualitatively similar dynamics to the monetary-policy shock for the variables displayed.

and labor-supply shocks generate sizeable reset-wage gaps and activate the wage extensive margin, whereas the cost-push shock leaves it almost unchanged on impact. Thus, similar inflation outcomes can mask sharply different price- and wage-setting dynamics.


For firms, inflation directly erodes relative prices and widens the gap with the reset price. Inflation is therefore close to a sufficient statistic for the initial activation of the price extensive margin. For unions, the relevant object is instead the gap between the inherited real wage and the reset wage. Because this gap also depends on labor demand and the marginal rate of substitution, equal inflation responses need not generate equal wage-resetting incentives.

The two supply-side shocks make this distinction transparent. The labor-supply shock raises the desired reset wage through the marginal rate of substitution, generating a large wage gap and a strong increase in wage resetting. The cost-push shock instead acts primarily through firms: it sharply activates price resetting, while weaker labor demand keeps the wage gap small and leaves the wage extensive margin largely inactive on impact. More generally, higher inflation does not automatically trigger state-dependent wage adjustments. The wage extensive margin is a contingent amplifier: it responds strongly only when shocks generate a sufficiently large gap between inherited and desired reset wages. When inflation operates mainly through firms' margins, or when weaker labor demand offsets the erosion of real wages, wage resetting can remain close to steady state. A muted response of the wage extensive margin does not imply that wage dynamics are irrelevant. As labor demand recovers and real wages catch up, nominal wage inflation can rise gradually through the intensive margin, even with only limited subsequent activation of wage resetting.


Phillips Curves are curves. Figure 4 visualizes the implications of the asymmetry in price and wage decisions for the model's response to demand and supply disturbances. Each panel reports 10,000 simulated observations conditional on monetary-policy or labor-supply shocks in the state-dependent model and the standard Calvo model.⁷

Figure 4a reports results for monetary policy shocks, one of the model's demand-side shocks. Demand shocks move both output and inflation in the same direction. As demand shocks, monetary policy shocks in Figure 4a trace out the nonlinear aggregate supply curve. The blue circles reveal the nonlinear dynamics of the model under state-dependent price and wage-setting. The nonlinear behavior of the Phillips curve in our model is driven by the interaction of two distinct channels: (i) on the intensive margin, higher demand increases the optimal resetting price for firms, directly raising inflation (as in a Calvo model), (ii) on the extensive margin, higher demand boosts the present discounted value of resetting prices, further amplifying inflation by increasing the frequency of price adjustments. The same mechanisms apply to labor unions' wage-setting decisions. Note that, in general equilibrium, the non-linearity of wage and price setting decisions reinforce each other. Higher price level reduces

⁷In the absence of habit and indexation, the degree of endogenous persistence is relatively modest, allowing us to isolate the conditional Phillips curves. These curves can be interpreted as proxies for the underlying aggregate demand and supply relationships (see [Bergholt et al., 2024](#)).



(a) Monetary policy shocks



(b) Labor supply shocks

Figure 4: Phillips curves are curves. The dots show the results of stochastic simulated runs of the output gap and of (i) price inflation, (ii) wage inflation, (iii) the share of not-adjusting firms; (iv) the share of not-adjusting unions, respectively, in the case of monetary policy shocks (Panel (a)) and labor supply shocks (Panel (b)).

Notes: For each panel, we simulate 10,000 observations of the model while allowing only the indicated structural shock to vary. Shock innovations are drawn from a Gaussian distribution, with their scale calibrated to reproduce the relevant volatility of the corresponding disturbance. Each dot represents one simulated period. Blue circles refer to the benchmark model with state-dependent adjustment frequencies, while red crosses refer to the corresponding model with fixed adjustment frequencies. For brevity, we refer to the shares of non-reoptimized prices and wages as the Calvo shares. Results for risk-premium and cost-push shocks are reported in Appendix Figure G.5.

real wages, increasing the real wage gap and thus the incentive for unions to reset nominal wages upward. This adjustment pushes up marginal costs and increases the price gap, reinforcing inflationary pressures through a feedback loop. Consequently, both price *and* wage Phillips curves exhibit convexity, linking inflation dynamics across the goods and labor markets. Without state-dependency (i.e., setting $\gamma_p, \gamma_w \rightarrow 0$), the model reverts to the standard NK model, which produces a quasi-linear relationship between output and price or wage inflation in response to demand-side shocks, as shown by the red crosses in Figure 4a.

Figure 4b presents similar simulations conditional on only labor supply shocks, as an example of supply disturbances. The shocks move price and wage inflation in the same direction and output in opposite directions. As supply shocks, labor supply shocks in Figure 4b trace out the nonlinear aggregate demand relationship. By reducing labor supply, the shocks drive up real wages and the real wage gap, increasing marginal costs and, ultimately, prices. In the standard NK model (red crosses), these shocks generate a quasi-linear relationship between marginal costs, inflation, and output. In contrast, in the state-dependent model, the extensive margin creates nonlinear effects due to the asymmetry analyzed above, which is evident when comparing the effects of inflationary and deflationary supply shocks. Inflationary supply shocks, which raise marginal costs, lead to more frequent price and wage resets (see Panel (iii) and (iv) in Figure 4b), thus generating strong upward pressures on prices and wages. On the contrary, deflationary shocks have weaker effects, as adjustment frequencies are less responsive to falling prices. As a result, the blue circles in Panel (i) and (ii) in Figure 4b stretch and fan out relative to the red crosses when inflation is on the upside, while they lie close to the red crosses when inflation is on the downside. The extensive margin amplifies the inflationary pressure of adverse supply shocks, but this effect is not symmetric: disinflationary pressure by lowering the nominal reset gaps of benign supply shocks is not amplified with respect to the standard Calvo model.

Appendix Figure G.5 reports the corresponding simulations for the remaining shocks. The risk-premium shock reproduces the demand-side curvature generated by monetary policy, whereas the cost-push shock generates pronounced price-side nonlinearity but slightly weaker activation of the wage margin than the labor supply shock. Cost-push shocks differ because they operate more directly through firms' marginal costs. As noted above, therefore, they generate strong price-side nonlinearity, while their adverse effect on labor demand mutes the contemporaneous activation of the wage extensive margin because they lower labor demand.

In sum, Figure 4 visualizes the strong interactions between inflation and the frequency of price and wage resets that our model generates. During periods of high inflation, price and wage resets can increase, reflecting the state-dependent nature of the adjustment process. In the fixed adjustment probability model of Calvo (1983), these channels are deactivated, leading to constant adjustment frequencies, and quasi-linear price and wage Phillips curves un-

der both demand and supply side shocks. In such a setting, inflationary responses to shocks are proportional to marginal cost fluctuations, and the rich interplay between the resetting frequency and inflation vanishes. This underscores the key role of state-dependency in shaping nonlinear inflation-output trade-offs observed in the price and wage Phillips curves.

The role of trend inflation. The discussion above highlights why positive trend inflation reinforces the asymmetry in adjustment incentives. Appendix D studies this mechanism explicitly, and Figures D.3 and D.4 show how Figure 4 changes with trend inflation. With incomplete indexation, positive inflation continuously erodes inherited relative prices and wages. At zero trend inflation this channel disappears, making adjustment substantially more symmetric across inflationary and disinflationary states. Once inflation turns negative, the direction of the mechanism reverses, as non-adjustment instead raises inherited relative prices and wages.

3.2 State-, sign-, and size-dependent effects of exogenous shocks

In our nonlinear benchmark, the effects of the exogenous shocks are state-, sign- and size-dependent.

State-dependence. Unlike the quasi-linear standard NK Calvo model, where responses are uniform regardless of initial conditions, our model implies that the responses to shocks are state-dependent, meaning that they differ depending on the condition of the economy. To isolate state dependence from the sign and size dependence investigated below, we hit the economy with inflationary labor supply shocks – which, as shown in Panels (iii) and (iv) of Figure 4 activates both the price- and wage-resetting margins – of different size. On top of each initial state, we add either the same contractionary monetary-policy shock or the same adverse cost-push shock. Figure E.1 displays the state-dependent impulse response functions. The effect of the same shocks depends on the macroeconomic states. The larger the initial level of inflation, the more prices and wages are flexible, and the more the effects of both the monetary policy and the cost-push shocks on the responses of nominal variables are quantitatively amplified. While for cost-push shocks the larger response of inflation translate in an even larger negative response of output, for monetary policy shocks lower nominal rigidities translate in a lower (negative) response of output.

State-dependency is therefore particularly relevant for the stabilization trade-off that monetary policy is facing. The increased flexibility due to a higher initial level of inflation steepens the Phillips curves, leading to larger variations in inflation and smaller ones in output in response to a given monetary policy shock. The efficacy of monetary policy in controlling inflation therefore depends on macroeconomic conditions. To underscore those dynamics, Figure 5 shows the output-inflation stabilization trade-off, computed from simulations conditioned on different starting inflation levels (the solid blue line is the case in which we condition on labor supply shocks only, the red dashed line is with risk-premium shocks only).




Figure 5: **The output-inflation stabilization trade-off for monetary policy depends on the level of inflation.** The lines decreases in the level of initial inflation, meaning that the higher the level of inflation, the more monetary policy affects inflation relative to output.

Notes: The stabilization trade-off is measured throughout this paper as the average ratio of cumulative output gap to inflation deviations over 20 quarters in response to +0.25% monetary policy shock.

The stabilization trade-off is defined as the average ratio of cumulative output gap deviations to cumulative inflation deviations over 20 quarters in response to a 0.25% monetary policy shock. The results are striking: at an annualized inflation rate of 0.5%, the stabilization trade-off is 7 times larger than at 5%. This indicates that monetary policy is more effective at controlling inflation when starting from high inflation levels, as the steeper Phillips curves amplify the response of inflation to monetary policy shocks. Notably, this effect holds regardless of whether underlying inflation is driven by labor supply or risk-premium shocks.

Sign- and size-dependence. Given the model's inherent nonlinearity, shock propagation is both sign- and size-dependent. This can have important implications for the question of whether central banks should look through supply shocks. Figure 6 reports rescaled responses to positive and negative cost-push shocks of different sizes, highlighting that larger inflationary shocks amplify the price response.⁸

Regarding size dependence, larger inflationary cost-push shocks generate a stronger price-inflation response per unit of shock, but lower persistence, as price adjustment accelerates in response to stronger inflationary pressures. A cost-push shock directly affects firms' marginal costs and hence the gap between inherited and desired prices. As a result, the price non-resetting Calvo share responds immediately on impact and increasingly with the size of

⁸Because the responses are normalized by the size of the shock, differences across lines measure departures from proportional scaling.




Figure 6: **Size and sign dependency in the response to a cost-push shock.** The different lines correspond to the scaled responses to different sizes of (positive and negative) cost-push shocks

Notes: The responses to the shocks are rescaled such that shocks are normalized to a 1% cost-push shock. We refer to the share of non-reoptimized wages or prices displayed in the chart as the Calvo shares.

the shock (see Panel e), generating a correspondingly stronger impact response of price inflation. This contrasts sharply with wage adjustment. As discussed above, the decline in labor demand also lowers the desired reset wage mitigating the gap between inherited and desired wages. Hence, a cost-push shock has a much weaker direct effect on the wage non-resetting Calvo share (see Panel f), which responds only modestly, even for relatively large shocks. Nominal wage inflation consequently displays a more sluggish and hump-shaped response. Its dynamics primarily reflect the gradual catch-up of real wages following the initial inflationary erosion, rather than an immediate increase in state-dependent wage adjustment.

As for cost-push shocks the nonlinearity operates primarily through the price extensive margin, the asymmetry in the present value of firms' profit functions generates sign dependency. While inflationary shocks increase price-resetting frequency, disinflationary shocks exhibit weaker size effects, as disinflation partly closes the price gap and therefore generates weaker incentives to reset prices downward. This is evident in the muted response of price and wage resetting frequencies following disinflationary shocks. As a result, large inflationary

shocks make the economy more flexible by increasing price adjustment rates, while disinflationary shocks fail to generate a symmetric response in price-setting activity, implying approximately linear scaling on the downside.⁹

4. A tale of two quantitative models

Building on the intuition of the small-scale model in Section 3, this section examines how time-varying updating frequencies affect the estimation of quantitative NK models, particularly focusing on the recent inflation surge period. To make the model more operational, we extend the framework to allow for consumption habits, as well as price and wage indexation to past inflation. We then estimate two specifications of this nonlinear model. In the first specification, referred to as the benchmark model or the model with time-varying updating frequencies, the parameters γ_p and γ_w are estimated. In the second specification, called the model with fixed updating frequencies, we effectively remove state-dependency by setting $\gamma_p \rightarrow 0$ and $\gamma_w \rightarrow 0$. Both estimations rely on the same dataset and follow an identical estimation procedure. Additionally, the priors are consistently specified across both versions.

4.1 Estimation strategy

Filtering method. The state-dependent adjustment frequencies for prices and wages require a fully nonlinear estimation approach. To bring this nonlinear model to the data, we use the inversion filter developed by [Cuba-Borda et al. \(2019\)](#); [Guerrieri and Iacoviello \(2017\)](#) and adapted to the extended path context by [Sahuc et al. \(2025\)](#). Given the observed sample $\{\mathcal{Y}_t\}_{t=1}^T$, the inversion filter recursively recovers structural shocks $\{\eta_t\}_{t=1}^T$ from the model observation equations and nonlinear solution paths derived from equation (2). The observation equation is given by:

$$\mathcal{Y}_t = h_\Theta(y_t), \quad (3)$$

where $h_\Theta(y_t)$ is the vector of measurement equations in (2) and Θ is the set of structural parameters. For a candidate vector of estimated parameters $\theta \in \Theta$, the likelihood function of the model $L(\theta, \mathcal{Y}_{1:T})$ is evaluated from the filtered sequence of shocks $\{\eta_t\}_{t=1}^T$. Following standard Bayesian practice in the DSGE literature ([An and Schorfheide, 2007](#)), we combine likelihood with prior distributions. Parametric uncertainty is assessed numerically via Markov Chain Monte Carlo (MCMC) sampling using the Metropolis-Hastings algorithm. Specifically, we generate 400,000 posterior draws across eight parallel chains, each running 50,000 iterations after a burn-in period of 50,000 draws. The scale parameter of the proposal distribution is adjusted to achieve an acceptance rate close to 33%.

⁹Figure G.1 in Appendix G compares local projection impulse responses to supply shocks with the impulse responses to a cost-push shock generated by the estimated model described in Section 4. The comparison shows that the model replicates both the sign and size asymmetries present in the data.

Observation equations and Data. Our dataset consists of four quarterly macroeconomic time series spanning the period from 1999Q1 to 2024Q4. The mapping between these observed variables and the model-implied quantities is summarized in the system of observation equations below, corresponding to an expanded version of (3):

$$\begin{bmatrix} \text{Real output growth per cap., lin. detrended} \\ \text{Inflation rate quarterly (HICP)} \\ \text{Quarterly policy rate (Krippner SR) adj. lvl} \\ \text{Real hourly wage growth, lin. detrended} \end{bmatrix} = \begin{bmatrix} 100 \times (\frac{Y_t}{Y_{t-1}} - 1) + \mathbb{1}_{t=2020Q2} \alpha_{y,1} + \mathbb{1}_{t=2020Q3} \alpha_{y,2} \\ 100 \times (\pi_t - 1) \\ 100 \times (R_t - 1) + \bar{r} \\ 100 \times (\frac{w_t}{w_{t-1}} - 1) + \mathbb{1}_{t=2020Q2} \alpha_{w,1} + \mathbb{1}_{t=2020Q3} \alpha_{w,2} \end{bmatrix}.$$

We next discuss the data transformation. All data is based on EA changing composition. We first start with GDP transformation. We de-trend the real output growth series with linear trend to account for the fact that our model abstracts from exogenous long-run GDP growth. For prices and wages, we use HICP and real hourly wages – detrended nominal compensation of employees divided by total employment deflated by HICP –, respectively, and compute their quarterly growth rates. As for the interest rate, we adopt the shadow rate constructed by Krippner (2013), using the quarterly average of monthly values, to account for the constraints imposed by the Effective Lower Bound (ELB).

In the measurement equation, we introduce two types of measurement errors related to two events in our sample. First, our sample includes the COVID-19 outbreak, a major macroeconomic shock that violates the "all fluctuations are alike" principle of Lucas (1977).¹⁰ Without specific treatment of this episode, the slope of the Phillips Curve would appear artificially flat, leading to a biased and downwardly distorted estimate of nominal rigidities. We therefore introduce $\alpha_{y,1}$ and $\alpha_{y,2}$ and $\alpha_{w,1}$ and $\alpha_{w,2}$, calibrated to average out the extreme recession and expansion caused by lockdown policies over those two quarters. Since we exclude 2020 from likelihood calculation, these two additional terms are automatically excluded from the likelihood computation and thus do not directly interfere with the identification of structural parameters.

The second event concerns the period at the Zero Lower Bound (ZLB), which requires the use of a shadow rate that pushes the average interest rate to very low levels. Without treatment, the estimated interest rate would be consistent with a near-one discount factor. We therefore introduce a correction term, $\bar{r}$, in the measurement equation for the interest rate. Given the low nominal rate environment and the presence of trend inflation, this adjustment ensures that the steady-state condition $\frac{\bar{\pi}}{\beta} - 1 + \bar{r}$ aligns with the observed average interest rate, while preserving the restriction $\beta < 1$.

Priors and posteriors estimate in the benchmark model. Table 1 reports the prior and posterior distributions for the estimated parameters. We first focus on the priors and benchmark

¹⁰In particular, the COVID recession was marked by a sharp contraction in output without a corresponding decline in real wages, largely due to mitigation policies implemented by governments. This disconnect temporarily breaks the empirical relationship underpinning the New Keynesian Phillips Curve.

		PRIOR DISTRIBUTION			BENCHMARK MODEL			FIXED UPDATING FREQ. MODEL		
		Shape	Mean	Std	$\phi_{p,t}, \phi_{w,t}$ - POSTERIOR DISTRIBUTION			ϕ_p, ϕ_w - POSTERIOR DISTRIBUTION		
					Mode	Mean	[5%:95%]	Mode	Mean	[5%:95%]
Panel A: Shock processes										
Monetary policy shock std	σ_r	$\mathcal{IG}_2$	0.001	1	0.0011	0.0012	[0.001:0.0013]	0.0011	0.0012	[0.001:0.0013]
Risk-premium shock std	σ_d	$\mathcal{IG}_2$	0.001	1	0.0029	0.003	[0.0024:0.0036]	0.0029	0.003	[0.0025:0.0036]
Cost-push shock std	σ_p	$\mathcal{IG}_2$	0.001	1	0.018	0.0183	[0.015:0.0219]	0.0252	0.0271	[0.0221:0.033]
Labor supply shock std	σ_w	$\mathcal{IG}_2$	0.001	1	0.0298	0.0305	[0.0246:0.0374]	0.0384	0.0413	[0.0326:0.0514]
Half-life, risk-premium shock	$\frac{\log(0.5)}{\log(\rho_d)}$	$\mathcal{G}$	4	1	13.0084	13.2201	[10.785:15.9057]	14.5451	14.2904	[11.757:17.0361]
Half-life, cost-push shock	$\frac{\log(0.5)}{\log(\rho_p)}$	$\mathcal{G}$	4	1	12.2363	12.4602	[9.682:15.6644]	11.4573	12.2717	[9.266:15.7839]
Half-life, labor supply shock	$\frac{\log(0.5)}{\log(\rho_w)}$	$\mathcal{G}$	4	1	9.2332	9.265	[7.0103:11.7264]	8.4392	8.9047	[6.8619:11.1949]
Panel B: Structural parameters										
Curvature consumption utility	σ	$\mathcal{G}$	2	1	4.0832	4.0926	[3.2369:5.0245]	4.7637	4.8712	[3.8234:6.0353]
Curvature disutility of labor	χ	$\mathcal{G}$	2	1	0.079	0.1034	[0.0317:0.2092]	0.1414	0.2006	[0.0625:0.4205]
Consumption habit	h	$\mathcal{B}$	0.5	0.1	0.7503	0.754	[0.6612:0.8602]	0.6813	0.6993	[0.6172:0.7756]
Inflation stance	θ_π	$\mathcal{N}$	2	0.05	1.9408	1.9448	[1.8614:2.028]	1.9469	1.9429	[1.8593:2.027]
Output stance	θ_y	$\mathcal{N}$	0.125	0.01	0.1197	0.1189	[0.1031:0.1348]	0.118	0.1172	[0.1007:0.1336]
Interest rate smoothing	ρ	$\mathcal{B}$	0.5	0.1	0.9244	0.9239	[0.9084:0.9383]	0.9294	0.928	[0.9135:0.9416]
Panel C: Price and wage setting										
Price Calvo share	$\phi_{p,ss}$	$\mathcal{B}$	0.7	0.02	0.8317	0.8316	[0.8086:0.8531]	0.7896	0.7954	[0.7766:0.8121]
Intensity of choice price	γ_p	$\mathcal{N}$	10	1	7.9831	7.7925	[5.7736:9.7402]		–	
Price indexation	ϱ_p	$\mathcal{B}$	0.5	0.1	0.1893	0.1908	[0.1235:0.2615]	0.2458	0.263	[0.1779:0.3528]
Wage Calvo share	$\phi_{w,ss}$	$\mathcal{B}$	0.7	0.02	0.8395	0.8363	[0.8124:0.8575]	0.8113	0.8168	[0.797:0.8342]
Intensity of choice wage	γ_w	$\mathcal{N}$	10	1	9.5346	9.443	[7.776:11.0997]		–	
Wage indexation	ϱ_w	$\mathcal{B}$	0.5	0.1	0.3632	0.3653	[0.2675:0.4705]	0.2466	0.2563	[0.1622:0.3633]

Table 1: Priors and posteriors of the quantitative models. Sample: 1999Q1–2024Q4.

model's posterior. Priors are chosen based on standard conventions in the literature under the constraint that the model dynamics are more sensitive to parameter change than their usual linearized counterpart. In all the estimations in the paper, we assume that the discount factor is $\beta = 0.995$, the yearly inflation trend is 2% or $\bar{\pi} = 1.005$, and that the elasticity of good and labor demand are equal to $\varepsilon_p = \varepsilon_w = 9$. The autocorrelation of the monetary policy shock is usually weakly identified suggesting we could assume an i.i.d. shock and $\rho_r = 0$. Results are not sensitive to change in those parameters but their identification is beyond the scope of this paper.

The standard deviations of the shocks $\sigma_r, \sigma_d, \sigma_p, \sigma_w$ are assigned inverse gamma priors with a mean of 0.001. Posterior means remain close to these priors for the monetary policy shock $\sigma_r = 0.0012$ and the risk-premium shock $\sigma_d = 0.003$, with, in the absence of the scaling generated by the linearization, higher values for the cost-push shock ($\sigma_p = 0.0183$) and labor supply shock ($\sigma_w = 0.0305$). The persistence of shocks is estimated in terms of their half-life, defined as $\log(0.5)/\log(\rho_j)$, rather than directly through their AR(1) coefficients.¹¹ The posterior mean half-lives are 13.2201 quarters for the risk-premium shock, 12.4602 quarters for the price cost-push shock, and 9.265 quarters for the labor supply shock, corresponding to AR(1) coefficients of $\rho_d = 0.9489$, $\rho_p = 0.9459$, and $\rho_w = 0.9279$, respectively.

The curvature of consumption utility σ and the disutility of labor χ follow gamma priors with means of 2 and 2, respectively. Posterior mode estimates indicate large increases for σ , with $\sigma = 4.0832$ while the estimation for $\chi = 0.079$ implies a nearly linear curvature, as the exponent $1 + \chi$ in the disutility of labor approaches 1. The consumption habit prior follows a beta distribution with a mean at 0.5 and exhibits posterior distribution with a mode at $h = 0.7503$.

In terms of the monetary policy rule, the inflation response parameter we found a posterior mean of $\theta_\pi = 1.9448$, reflecting a strong monetary policy response to inflation deviations, while the mode of the output response parameter is $\theta_y = 0.1197$ and the interest rate smoothing is $\rho = 0.9244$, consistent with findings in the literature on interest rate inertia.

Parameters governing steady-state price and wage stickiness, $\phi_{p,ss}$ and $\phi_{w,ss}$, are assigned relatively tight priors centered at 0.7, reflecting standard estimates for the euro area and helping to anchor the nonlinear model to empirically plausible degrees of nominal rigidity. Their posterior modes, $\phi_{p,ss} = 0.8317$ and $\phi_{w,ss} = 0.8395$, imply substantial price and wage rigidity around the steady state. The state-dependence parameters, γ_p and γ_w , determine instead how strongly adjustment probabilities respond to changes in the gains from reoptimization. Their posterior estimates, $\gamma_p = 7.7925$ and $\gamma_w = 9.443$, remain relatively close to their priors, indicating that aggregate data are less informative about the precise strength

¹¹This parameterization penalizes processes close to the unit-root boundary, which otherwise arise in both models due to the brevity of the sample and the relative sparsity of the models. The problem is even more pronounced in the fixed updating frequency model.

of this margin than about steady-state nominal rigidity ¹². We therefore focus below on the model-implied variation in price and wage adjustment frequencies rather than on the γ coefficients in isolation. Finally, the posterior estimates of price and wage indexation lie between 0.2 and 0.4, despite priors centered at 0.5, implying moderate indexation to lagged inflation.

The comparison with the fixed-frequency model is informative because most conventional structural and policy parameters have broadly overlapping posterior distributions across the two specifications. The main differences are instead concentrated in the shocks and in the parameters directly governing nominal adjustment. Hence, the comparison largely contrasts two different ways of accounting for the same inflation and wage dynamics.

First, the fixed-frequency model relies more heavily on exogenous supply disturbances. The posterior mean standard deviation of the cost-push shock rises from $\sigma_p = 0.0183$ in the benchmark model to 0.0271, an increase of almost 50%, with non-overlapping 90% posterior intervals. The labor-supply shock standard deviation also rises, from $\sigma_w = 0.0305$ to 0.0413, although the corresponding posterior intervals overlap. By contrast, the monetary-policy and risk-premium shock estimates change little. Thus, once adjustment frequencies are prevented from responding to economic conditions, a larger share of price and wage fluctuations is accounted for by exogenous supply disturbances.

Second, the fixed-frequency model is estimated with somewhat greater nominal flexibility in normal times. The distinction is therefore not simply that the benchmark is more flexible than the alternative. The fixed-frequency specification builds this flexibility permanently into the economy, whereas the benchmark allows substantial rigidity around the steady state and generates additional flexibility endogenously when the gains from reoptimization become large. Put differently, the fixed-frequency model treats nominal flexibility as a level; the benchmark treats it partly as a state. This distinction is particularly relevant for the inflation surge: the benchmark can generate a substantial increase in adjustment frequencies when shocks become large without requiring the economy to be equally flexible throughout the historical sample.

Model fit. We next compare how well the two estimated specifications account for the observed data. Since estimation uses the nonlinear inversion filter described above, it is useful to make explicit the object being compared. For a candidate parameter vector θ , the filter recursively recovers the structural innovations $\{\eta_t(\theta)\}_{t=1}^T$ that rationalize the observations. Following [Cuba-Borda et al. \(2019\)](#), the log posterior objective used in the estimation can be

¹²Their magnitudes are nevertheless economically plausible: in particular, the estimate of γ_p is very close to that obtained by [Gasteiger and Grimaud \(2023\)](#); [Gautier et al. \(2026\)](#) when the mechanism is disciplined using euro-area and U.S micro price data.

written, up to an additive constant C , as

$$\begin{aligned} \log \tilde{p}(\theta \mid \mathcal{Y}_{1:T}) = C + & \underbrace{\log p(\theta)}_{\text{Log prior density}} \\ & \underbrace{-\frac{Tn_y}{2} \log(2\pi) - \frac{T}{2} \log |\Sigma_\eta|}_{\text{Gaussian normalization}} \\ & + \underbrace{\sum_{t=1}^T \left[\underbrace{-\frac{1}{2} \eta_t(\theta)' \Sigma_\eta^{-1} \eta_t(\theta)}_{\text{Structural-shock term}} + \underbrace{\log \left| \det \left(\frac{\partial \eta_t(\theta)}{\partial \mathcal{Y}_t} \right) \right|}_{\text{Jacobian term}} \right]}_{\text{Shock-and-Jacobian contribution}}. \end{aligned} \quad (4)$$

Here n_y is the number of observables, Σ_η is the covariance matrix of structural innovations.

We refer to $\log \tilde{p}(\theta \mid \mathcal{Y}_{1:T})$ as the posterior objective: it is the sum of the log prior and log likelihood used to locate the posterior mode, rather than a posterior model probability.

The likelihood has a useful economic interpretation. The Gaussian normalization and the structural-shock term jointly assess how plausible the innovations required by the model are under their estimated distribution. A tighter Σ_η raises the density around zero through the normalization term, but also penalizes large standardized innovations more strongly through $-\frac{1}{2} \eta_t' \Sigma_\eta^{-1} \eta_t$. A model therefore benefits from a more concentrated shock distribution only if the innovations needed to rationalize the data remain plausible under that distribution.

The Jacobian captures a different margin. Because the inversion filter maps observed variables into structural innovations, the likelihood must correct for how strongly a given innovation moves the observables. The Jacobian is this change-of-variables correction: it prevents a model from obtaining a mechanical fit advantage simply because very small shocks generate very large observable responses. The shock and Jacobian terms together therefore capture the time-varying part of the inversion-filter likelihood, whereas the Gaussian normalization is constant over time within each estimated specification.

Table 2 shows that the benchmark attains a substantially higher posterior objective at its mode: 433.65, compared with 409.87 for the fixed-frequency model. The difference reflects a much higher log likelihood, 554.68 against 499.144, partly offset by a lower prior density. The 55.54 point likelihood gain is large relative to the two additional parameters.¹³

More informative is where this likelihood gain arises. Before 2020, the Shock-and-Jacobian contribution is slightly less favorable in the benchmark. During 2021–2024, the ordering re-

¹³As a robustness check on model complexity, applying the conventional Schwarz/Bayesian Information Criterion penalty to the likelihood evaluated at the posterior mode gives -1021.11 for the benchmark and -919.33 for the fixed-frequency model. Since standard BIC is formally evaluated at the maximum-likelihood estimate, we use this only as a penalized-fit diagnostic and do not report it as a model-selection statistic.

verses. The reversal is driven by the structural-shock term: during the inflation surge, the benchmark rationalizes the observations with substantially less unlikely innovations than the fixed-frequency model. Almost the entire improvement comes from the cost-push and labor-supply shocks. The Jacobian moves in the opposite direction, so the gain is not a mechanical consequence of stronger nonlinear propagation.

	BENCHMARK	FIXED UPDATING FREQ.
Panel A: Posterior-mode comparison		
Log posterior objective	433.65	409.87
Log likelihood	554.68	499.14
Log prior	-121.03	-89.27
Panel B: Average contribution per quarter, 1999Q1–2019Q4		
Gaussian normalization	16.495	15.923
Structural-shock term	-2.189	-1.862
Jacobian term	-9.279	-9.501
Total log-likelihood contribution	5.027	4.560
Panel C: Average contribution per quarter, 2021Q1–2024Q4		
Gaussian normalization	16.495	15.923
Structural-shock term	-2.030	-3.810
Jacobian term	-10.315	-8.837
Total log-likelihood contribution	4.151	3.276


Notes: Both specifications are evaluated at their respective posterior modes. The likelihood decomposition corresponds to equation (4). The treatment of 2020 follows the estimation; these observations are excluded from the pre/post comparison.

Table 2: Comparison with fixed updating frequencies


Because the two specifications are re-estimated separately, this decomposition should not be read as holding all common parameters and shocks fixed. Rather, it describes how the two estimated models rationalize the same observed data. With fixed updating frequencies, a larger part of the inflation surge is accounted for by unusually large exogenous supply disturbances. In the benchmark, large shocks also change the incentives to reoptimize prices and wages and thereby alter their own propagation. The empirical value of state-dependent adjustment therefore becomes largest precisely during the high-inflation episode.

4.2 Drivers of inflation in both models

Historical decompositions are more challenging in a nonlinear setting than in a linear model. In a linear model, the effects of individual shocks are additive, so that their separate contributions exactly reconstruct the historical path of the variable. This is no longer the case in our model: when several shocks act simultaneously, their joint effect may differ from the sum of their effects when simulated separately. We therefore first simulate the demand- and supply-side shock bundles separately and then compare their sum with the path obtained when both bundles act jointly. The resulting demand–supply interaction captures this non-



(a) Benchmark



(b) Fixed updating frequencies

Figure 7: Historical decomposition of price inflation in both model specifications. Panel (a): benchmark model. Panel (b): model with fixed updating frequencies.

Notes: The black line reports the historical path generated when all shocks act jointly. Demand-side shocks collect the joint contribution of the monetary-policy and risk-premium shocks, while supply-side shocks collect the joint contribution of the labor-supply and cost-push shocks. Because the model is nonlinear, the joint effect of the demand- and supply-side shock bundles need not equal the sum of their separate effects. The demand-supply interaction captures this non-additive component, so that the bars reconstruct the historical path by construction. The corresponding decomposition for nominal wage inflation is reported in Figure G.6 of Appendix G.

additive component. It should not be interpreted as an additional structural shock, but as the contribution arising from the nonlinear propagation and interaction of the identified shocks. The more granular decompositions in Appendix G further separate the supply-side bundle into labor-supply shocks, cost-push shocks, and their within-supply interaction.

Figure 7a highlights striking differences between the two versions of the model. In the benchmark model with endogenous updating frequencies, the recent inflation surge is explained in roughly equal proportions by supply-side shocks and demand-side shocks, and partly also by their nonlinear interaction. Supply-side shocks are predominant at the onset of the surge, but subsequently demand-side shocks mostly contribute to its amplification and persistence. The mechanism is that the initial supply-side shocks raise inflation, increase the frequency of price adjustment, and steepen the price Phillips curve, thereby amplifying the effects of subsequent shocks. As inflation persists and labor demand recovers, wage resetting may also increase more gradually, reinforcing the price-to-wage feedback and helping sustain inflation. Section 4.3 examines the role of endogenous price- and wage-setting rigidities in generating these dynamics.

In comparison, in Figure 7b, the historical decomposition in the model with fixed updating frequencies attributes roughly two-thirds of the recent inflation surge to supply-side shocks and one-third to demand-side shocks, with virtually no contribution from demand–supply interactions. Without endogenous changes in price- and wage-setting frequencies, the effects of the shocks remain approximately additive during the inflationary episode.


4.3 Inflation channels in the benchmark model

We now turn to counterfactuals that demonstrate the role of time-varying price- *and* wage-updating frequencies in explaining euro area data from 2021 onwards. Figure 8 reports the benchmark simulation with time-varying price- and wage-adjustment frequencies (blue solid line, coinciding with the data by construction) next to the counterfactuals. The red dotted line is the counterfactual with both adjustment frequencies fixed. The yellow dashed line is the counterfactual with time-varying wage-adjustment only. Finally, the purple dash-dotted line is the counterfactual with time-varying price-adjustment frequency only.¹⁴

Panel 8a exhibits the simulated paths for all scenarios over the entire sample. Up to 2021, the counterfactual predicted paths deviate only slightly from the data. This suggests that time-varying price- and wage-adjustment frequencies play only a minor role in explaining the data during the pre-2021 period. The adjustment frequencies show only little variation around their posterior mode in the benchmark as well in the counterfactual simulations.

However, Panel 8b makes clear that time-varying adjustment frequencies play a crucial role

¹⁴All counterfactuals use the benchmark parameters and filtered shock sequence. They differ only in whether the price and/or wage Calvo shares are allowed to vary endogenously or are fixed at their benchmark steady-state values. They are therefore distinct from the separately estimated fixed-frequency model considered in Section 4.1.




(a) Full sample

Figure 8: Contribution of state-dependent price and wage setting to the dynamics of selected variables. Counterfactual exercise switching off the state-dependency in wage and/or price setting. Dotted red lines: fixed updating frequencies, dashed yellow lines: only *wage* updating frequency is time-varying, dash-dotted purple lines: only *price* updating frequency is time-varying. We refer to the share of non-reoptimized wages or prices displayed in the chart as the Calvo shares.

in explaining the data from 2021 onwards. In particular, the model relies predominantly on the time-varying price-adjustment frequency to generate the surge in inflation and the associated dynamics of the output gap, the real wage, and nominal wage inflation. The counterfactual with only time-varying price adjustment remains close to the benchmark for these variables. Without state-dependent price setting, the model fails to reproduce the magnitude of the nominal wage inflation surge and the associated real-wage dynamics, as illustrated by the red dotted and yellow dashed lines.

By contrast, the counterfactual with only time-varying wage adjustment does not reproduce the data well. The relatively limited quantitative contribution of state-dependent wage set-



(b) Recent sample

Figure 8: (Cont.) **Contribution of state-dependent price and wage setting to the dynamics of selected variables.** Counterfactual exercise switching off the state-dependency in wage and/or price setting. Dotted red lines: fixed updating frequencies, dashed yellow lines: only *wage* updating frequency is time-varying, dash-dotted purple lines: only *price* updating frequency is time-varying. We refer to the share of non-reoptimized wages or prices displayed in the chart as the Calvo shares.

ting reflects the composition of the shocks driving this episode. Cost-push shocks directly affect firms' marginal costs and reduce labor demand. Consequently, the price Calvo share falls sharply during the inflation surge, whereas changes in the wage-updating frequency have a smaller effect on aggregate dynamics. Much of the movement in nominal wage inflation and the real wage is therefore generated by the price-side response and the subsequent adjustment of wages, rather than by a strong activation of the extensive margin of wage setting. This does not imply that endogenous wage setting is irrelevant in general. Rather, under the estimated sequence of shocks, allowing the wage-resetting frequency to vary adds relatively little beyond state-dependent price setting.

Finally, we benchmark our model-implied extensive margin of price adjustment against euro area micro evidence. Appendix Figure G.3 compares the latent share of re-optimized prices in the quantitative model, $1 - \phi_{p,t}$, with the share of changed prices reported by [Gautier et al. \(2026\)](#) over their common sample (2012–2024). The two series comove strongly, with a correlation of 0.6146 (90% confidence interval: [0.4609, 0.7325]). The model-generated share of re-optimized prices closely aligns with the share of price changes in the data. Before the 2021–2022 inflation surge, these shares are around 0.15 (model) and 0.23 (data). Both series display a sharp acceleration in price adjustment during the inflation episode, with the increase occurring about one quarter earlier in the model. A natural interpretation is composition: our model-based series pertains to headline HICP and therefore includes energy items, whose adjustment frequency rose abruptly already in late 2021 as energy prices began to surge, while [Gautier et al. \(2026\)](#) exclude energy from their micro-data coverage. Accordingly, the data series picks up the broad-based increase in price changes with a short delay, once non-energy components adjust more strongly. This consistency suggests that our model captures key features of the extensive margin of price adjustment during the recent inflation episode.

5. Policy implications

Next, we analyze the policy implications of our model. We proceed in three steps. First, we measure directly how the local price and wage Phillips-curve slopes vary along the filtered euro area history. Second, we examine how this variation maps into the state-dependent transmission of monetary-policy and cost-push shocks. Third, we study counterfactual outcomes under alternative monetary-policy stances.

5.1 Time-varying Phillips-curve slopes

We first measure directly how the slopes of the price and wage Phillips curves vary along the estimated euro area history. In a standard linearized New Keynesian model, these slopes are constant because the Phillips curves are linearized once around the steady state. In our nonlinear model, instead, the relevant first-order slopes depend on the state of the economy. We therefore linearize the nonlinear equilibrium conditions locally around each filtered historical state. Let $\mathcal{S}_t$ denote the endogenous state implied by the filtered shock history through quarter t . The resulting state-specific price and wage NKPC slopes are

$$\kappa_{p,t} \equiv \left. \frac{\partial \log \pi_t}{\partial \log p_{w,t}} \right|_{\mathcal{S}_t}, \quad \kappa_{w,t} \equiv \left. \frac{\partial \pi_t^w}{\partial \log(mrs_t/w_t)} \right|_{\mathcal{S}_t},$$

where $p_{w,t}$ denotes firms' real marginal cost and $\pi_t^w = \log(w_t) - \log(w_{t-1}) + \log(\pi_t)$ is nominal wage inflation. Hence, $\kappa_{p,t}$ and $\kappa_{w,t}$ have the same interpretation as the κ coefficients in the

“textbook” linearized price and wage NKPCs, except that they are evaluated at the sequence of states actually visited by the estimated economy rather than at a single steady state.¹⁵




Figure 9: Time-varying local price and wage Phillips-curve slopes. The figure reports the price and wage NKPC slopes, $\kappa_{p,t}$ and $\kappa_{w,t}$, obtained by locally linearizing the nonlinear model around each filtered historical state.

Notes: The local slopes measure the first-order response of price inflation to real marginal cost and of nominal wage inflation to the wage-setting gap, respectively. They are computed using small symmetric contemporaneous perturbations around each filtered state; see the footnote in the text.

Figure 9 reveals substantial time variation in both Phillips-curve slopes, but also a pronounced asymmetry between prices and wages. This provides a direct quantitative counterpart to the “Phillips curves are curves” result developed above: the curvature of the nonlinear model appears along the estimated history as variation in the slope of its local tangent Phillips curves.

The pattern is particularly clear on the price side. The price NKPC is flat but steepens sharply during the recent inflation surge and closely tracks the extensive margin of price adjustment. The correlation between $\kappa_{p,t}$ and the price-reset frequency is 0.964, while its correlation with HICP inflation is 0.850.

The wage side is markedly different. The wage NKPC slope is also tightly linked to its own extensive margin: the correlation between $\kappa_{w,t}$ and the wage-reset frequency is 0.929. Yet it

¹⁵We recover these local first-order effects numerically. For each quarter t , we retain the estimated filtered shock history through t . Around this local baseline, we introduce a small iid wedge of 0.01%, separately in the price marginal-cost condition and in the wage-setting condition. The wedges are applied only in period t and are zero otherwise. Using equal positive and negative perturbations, we approximate the derivatives by symmetric finite differences: $\kappa_{p,t} \simeq [\log \pi_t^{+p} - \log \pi_t^{-p}] / [\log p_{w,t}^{+p} - \log p_{w,t}^{-p}]$ and $\kappa_{w,t} \simeq [\pi_t^{w,+w} - \pi_t^{w,-w}] / [\log(mrs_t^{+w} / w_t^{+w}) - \log(mrs_t^{-w} / w_t^{-w})]$. The diagnostic wedges are used only to approximate the local derivative and do not enter the estimated shock processes. Their small size ensures that the exercise isolates the first-order response of the nonlinear equilibrium mapping around each historical state.

is essentially unrelated to either HICP inflation or nominal wage inflation, with correlations of only 0.099 and 0.152, respectively, neither statistically different from zero.¹⁶ Neither high price inflation nor high wage inflation is therefore, by itself, a sufficient statistic for wage-setting state dependence.

This distinction is particularly relevant for the recent inflation episode. Price inflation directly activated the price-resetting margin and steepened the local price NKPC. Wage inflation, by contrast, could rise through delayed real-wage catch-up without a comparable activation of the wage-resetting margin. The local slopes therefore make precise the price–wage asymmetry emphasized above: state dependence was strongly activated on the price side, whereas on the wage side it remained more contingent on the source of the shock.

We next examine how these changes in local Phillips-curve slopes translate into the state-dependent transmission of monetary-policy and cost-push shocks.

5.2 Shock transmission over the business cycle

Having documented direct time variation in the local price and wage Phillips-curve slopes, we now examine how this variation maps into the transmission of shocks along the euro area business cycle. We first consider monetary-policy shocks and then cost-push shocks. To compute IRFs conditional on the business cycle, we simulate the path of the endogenous variables in our model, in response to the filtered shocks, which provides the observed data. We then run a second simulation, adding a monetary policy (or a cost-push) shock of fixed size to the filtered shocks for each quarter in the sample. The difference between these two simulations yields the IRFs to the monetary policy (or cost-push) shock conditional on the business cycle. Hence, for each quarter in our sample, we obtain a different IRF to the same shock, allowing us to analyze how the transmission of these shocks varied over time along the euro area business cycle. Note that in a linearized version of the model (e.g., [Smets and Wouters, 2007](#)), these IRFs would all be identical, while in the nonlinear version IRFs are state-dependent.

Figure 10 shows the responses of inflation and output to a contractionary monetary policy shock for each quarter, highlighting the IRFs corresponding to periods associated with the peak and trough of quarterly inflation series. This figure nicely visualizes how the effects of monetary policy shocks in the euro area differed in our sample according to our estimates. In 2008Q4 – orange dashed line – a monetary policy would have had a strong contractionary impact on output with little effect on inflation. In contrast, during the peak of the recent inflation surge, in 2022Q1, a contractionary monetary policy shock would have largely affected inflation and much less output.

¹⁶The corresponding 90% confidence intervals are [0.951,0.974] for $\text{corr}(\kappa_{p,t}, 1 - \phi_{p,t})$, [0.798,0.890] for $\text{corr}(\kappa_{p,t}, \pi_t)$, [0.902,0.948] for $\text{corr}(\kappa_{w,t}, 1 - \phi_{w,t})$, [-0.064,0.257] for $\text{corr}(\kappa_{w,t}, \pi_t)$, and [-0.011,0.306] for $\text{corr}(\kappa_{w,t}, \pi_t^w)$. Confidence intervals are obtained using the Fisher transformation of the sample correlations.




Figure 10: **IRFs to a + 0.25% monetary policy shock as function of the state.** State-dependent IRFs are computed as local counterfactual responses around each filtered state. For each impact date, we first simulate a baseline path that keeps the filtered shock history up to the impact period and sets all subsequent innovations to zero. We then simulate a counterfactual path with the same shock history, but add a contemporaneous monetary policy shock at impact. The IRF is the difference between the counterfactual and baseline paths.

This pattern is consistent with the local slopes in Figure 9. In high-inflation states, the price Phillips curve is steeper because the price-resetting margin is more active. Monetary-policy shocks also affect labor demand and the desired reset wage, so their propagation can involve the wage-setting margin as well. The resulting state dependence is therefore a general-equilibrium consequence of the time-varying price and wage slopes rather than a separate property imposed on monetary-policy transmission.

Next, we turn to the transmission of cost-push shocks. Cost-push shocks are particularly informative for our purposes because they operate directly through firms' marginal costs. They therefore provide a case in which price-setting state dependence can be strongly activated, while wage-setting state dependence need not move in parallel. For this analysis, we repeat the exercise above, replacing the monetary policy shock with a fixed-size cost-push shock. For each date in the sample, we compare the path generated by the filtered historical shock sequence with a counterfactual path in which an additional cost-push shock is added on impact. The difference between the two paths gives the marginal response to the cost-push shock conditional on the state of the euro area business cycle.

Figure 11 reports the full set of state-dependent responses to the cost-push shock. The shock raises inflation and reduces output, as expected. More importantly, the response varies substantially with the initial state. In high-inflation states, the inflation response is larger because the shock activates the price-setting margin more strongly. Since the Calvo share is the share of firms or unions that do not re-optimize, a fall in the Calvo share corresponds to




Figure 11: **IRFs to a + 0.5% cost-push shock as function of the state.** State-dependent IRFs are computed as local counterfactual responses around each filtered state. For each impact date, we first simulate a baseline path that keeps the filtered shock history up to the impact period and sets all subsequent innovations to zero. We then simulate a counterfactual path with the same shock history, but add a contemporaneous cost-push shock at impact. The IRF is the difference between the counterfactual and baseline paths.

an increase in the corresponding reset frequency. The price Calvo share falls sharply in high-inflation states, indicating a strong rise in the price-resetting frequency. The wage Calvo

share, in contrast, moves much less on impact. Thus, cost-push shocks generate strong state dependence in price inflation without requiring a contemporaneous surge in wage reoptimization.

This distinction is central for the wage response. A cost-push shock directly raises firms' marginal costs and triggers a rapid price response. At the same time, it weakens labor demand, which limits the increase in the desired reset wage and therefore dampens the activation of the wage extensive margin on impact. Nominal wages nevertheless respond with a delay and very persistently, so that real wages fall initially and recover only gradually. Hence, wage inflation can be persistent even when the extensive margin of wage setting is only weakly activated: the persistence comes from delayed real-wage catch-up and price-to-wage transmission, rather than from an immediate wage-resetting spiral.

This is consistent with the local-slope evidence in Figure 9. Price-side steepening is closely associated with the activation of price resetting and therefore with the inflationary state of the economy. Wage-side steepening instead follows the wage-resetting margin and need not move contemporaneously with HICP inflation. The recent euro area episode is therefore price-led in the model: cost-push shocks activate price-setting state dependence first, while nominal wages adjust subsequently through real-wage catch-up and price-to-wage transmission.

These results imply that the case for looking through a supply shock is itself state-dependent. When a shock moves the economy into a high-inflation state with elevated price resetting, the shock becomes more inflationary, but monetary policy also gains traction. This raises the question of whether different systematic policy responses could have materially changed the outcome during the recent euro-area episode. We examine this next through illustrative counterfactuals that vary the inflation response coefficient in the estimated Taylor rule.

5.3 Policy stance and stabilization trade-offs

This section examines how the shape of the price and wage Phillips curves influences the trade-off between inflation and output, and how this can lead to different macroeconomic outcomes under alternative monetary policy stances. Counterfactual simulations are used to illustrate how recent dynamics might have evolved under different policy settings.

Specifically, in this subsection, we investigate how policy counterfactuals change between models with endogenous and fixed price and wage adjustment frequencies. Figure 12 compares the effects of different monetary policy stances toward inflation for the output gap, inflation, nominal wage inflation, and the policy rate. Solid lines represent the paths of variables when all parameters are at the posterior mode, i.e., the baseline scenarios. By construction, the simulated paths coincide with the data in both cases.

The counterfactual simulations during the 2021Q1–2024Q4 period use the same model-specific filtered shocks as in the baseline scenarios. However, the counterfactuals consider inflation




Figure 12: Impact of the monetary policy stance on output and inflation in the euro area recent inflation surge. Counterfactual exercise varying the monetary policy stance in the benchmark and fixed updating frequencies models. Red lines: fixed updating frequencies, blue lines: benchmark model. Dotted lines: decrease of the inflation reaction coefficient in the Taylor rule parameter by 20%, Dashed lines: increase of the same parameter by 20%.

Notes: Counterfactual simulations in the benchmark model (blue lines) and the model with fixed adjustment frequencies (red lines) during the 2021Q1-2024Q4 period, with different monetary policy stance. The dotted lines correspond to cases with a change -20% to the estimated reaction of the deviation of inflation to the target in the Taylor rule, the dashed lines are cases with a change $+20\%$ to the same parameter.

reaction coefficients (θ_π) adjusted by $\pm 20\%$.¹⁷

Under a tighter monetary policy stance ($\theta_\pi + 20\%$, dashed lines), the two models exhibit markedly different outcomes. In this counterfactual, both models predict comparable price inflation, as expected, below the baseline.¹⁸ However, in the model with fixed adjustment frequencies (red lines), the slight decrease in inflation comes at the cost of generating a longer-lasting recession with a more persistent negative output loss relative to the benchmark model (blue lines), whose output gap behavior, instead, almost coincides with the baseline. In sum, the model with fixed adjustment frequencies predicts a less favorable output-inflation trade-off, i.e., flatter Phillips curves relative to the benchmark model. Therefore this model suggests that a more aggressive monetary policy would have led to a deeper recession during this episode of high inflation. In contrast, the benchmark model predicts that a more aggressive monetary policy stance could have reduced price and wage inflation

¹⁷To save space, we omit the pre-2021 period, where no significant deviations are observed.

¹⁸At times, the model with fixed adjustment frequencies predicts lower wage inflation than the benchmark.

by more during this episode of high inflation with only limited additional output cost. The latter is true, because in the benchmark model the Phillips curves are steeper when inflation is high, which improves the output-inflation trade-offs.

Finally, notice from Figure 12 that the counterfactual with a more passive monetary policy stance ($\theta_\pi - 20\%$, dotted lines), suggests opposite conclusions compared to a more aggressive monetary policy stance. The model with fixed adjustment frequencies predicts a counterfactual expansion at the cost of moderately higher price and wage inflation. In contrast, the benchmark predicts hardly any difference in output relative to baseline, but much higher price and wage inflation.

The takeaway from these counterfactuals is that, through the lens of the benchmark model, a more proactive monetary policy during 2021Q1–2024Q4 could have helped moderate inflation with relatively minimal impact on output. In contrast, a more dovish monetary policy would have led to higher inflation.¹⁹

The difference in monetary policy effectiveness between the two models, described above, also affects the resulting policy prescriptions. Appendix F extends the analysis by considering a recursive inflation-targeting counterfactual, in the spirit of Sims and Zha (2006), implemented through contemporaneous monetary policy shocks. The exercise reveals important differences across the two models. In the fixed-updating-frequency model, the inflation surge is accounted for by larger exogenous supply disturbances and transmitted through a relatively flat, fixed Phillips curve. Since nominal flexibility cannot rise endogenously in high-inflation states, inflation is weakly responsive to economic slack. The recursive inflation-targeting rule therefore calls for a strong and persistent tightening, which generates substantial output losses and only gradually produces disinflation.

In the benchmark model, the policy problem is different. The model does not mainly rely on shorter-lived structural shocks: instead, inflation dynamics are shaped by state-dependent propagation. In high-inflation states, monetary-policy shocks are transmitted through steeper price and wage Phillips curves, because firms' pricing incentives and unions' wage-setting incentives both respond to the induced changes in demand, inflation, and labor-market conditions. As a result, monetary policy has more traction on inflation for a given movement in real activity. The inflation-targeting counterfactual therefore calls for a tightening path that is somewhat more aggressive than, but remains close to, the observed policy path, and delivers lower inflation with a smaller output loss than in the fixed-updating-frequency model.

¹⁹A caveat is that the analysis primarily captures nonlinearities in price and wage setting, while abstracting from other potentially relevant sources of nonlinearity — such as financial frictions — which could influence the inflation-output trade-off and alter the effects of alternative policy stances, possibly in the opposite direction. More broadly, the analysis does not account for the elevated real-time uncertainty that characterized the period, including the economic effects of the COVID-19 pandemic and Russia's invasion of Ukraine.

6. Conclusion

The findings of this paper highlight the crucial role of state-dependent nominal adjustments, and the distinct behavior of its price and wage extensive margins, in shaping inflation dynamics and monetary policy trade-offs during high inflation episodes. Unlike standard New Keynesian models with quasi-linear Phillips curves, our framework captures nonlinear inflation responses when the economy is hit by large shocks. Inflationary shocks broadly increase firms' incentives to reset prices, whereas the wage extensive margin responds only when shocks generate a sufficiently large gap between inherited and desired reset wages. Consistent with this distinction, the estimated local price Phillips-curve slope steepened sharply during the recent euro area inflation surge, while the wage slope did not move mechanically with HICP inflation. During this period, the extensive-margin response was therefore predominantly price-led: price-resetting frequencies rose sharply, while wage inflation primarily reflected delayed real-wage catch-up and price-to-wage transmission, with only limited activation of the wage-reoptimization margin. The episode thus featured price-to-wage transmission, but limited amplification through a generalized increase in wage-resetting frequency.

State dependence changes the transmission of both structural shocks and monetary policy to inflation and real activity. Cost-push shocks become more inflationary when the price-setting margin is highly active. Monetary-policy shocks also affect labor demand and desired reset wages and can therefore activate both price- and wage-setting margins. In high-inflation states, monetary policy consequently has greater traction on inflation relative to real activity, lowering the output cost of disinflation.

Beyond its direct implications for monetary policy, our methodology offers a flexible and scalable framework that can be embedded in larger DSGE models. By capturing the endogenous interaction between price- and wage-setting frequencies in a reduced-form, state-dependent setup, our approach provides a computationally efficient alternative to a fully microfounded menu cost model. This enables the integration of state-dependent inflation dynamics into more complex macroeconomic environments without significantly increasing model complexity.

Overall, our findings do not imply that central banks should react aggressively to every supply shock. Rather, the case for looking through such shocks depends on their size, persistence, and capacity to activate state-dependent nominal adjustment. In the estimated euro-area episode, the sizeable and persistent inflation overshoot moved the economy into a state of strongly elevated price resetting, making monetary policy more effective in curbing inflation than in a fixed-frequency model would imply. In this context, our model suggests that a somewhat stronger policy response could have reduced inflation with limited additional output losses.

References

- Acharya, V. V., Crosignani, M., Eisert, T. and Eufinger, C. (2023), How do supply shocks to inflation generalize? evidence from the pandemic era in europe, Working Paper 31790, National Bureau of Economic Research.
- Adjemian, S. and Juillard, M. (2024), ‘Stochastic extended path’, SSRN Working Paper. Available at SSRN: <https://ssrn.com/abstract=5189260> or <http://dx.doi.org/10.2139/ssrn.5189260>.
- Alvarez, F., Beraja, M., Gonzalez-Rozada, M. and Neumeyer, P. A. (2019), ‘From hyperinflation to stable prices: Argentina’s evidence on menu cost models’, *The Quarterly Journal of Economics* **134**(1), 451–505.
- Alvarez, F., Le Bihan, H. and Lippi, F. (2016), ‘The real effects of monetary shocks in sticky price models: a sufficient statistic approach’, *American Economic Review* **106**(10), 2817–2851.
- Alvarez, F., Lippi, F. and Passadore, J. (2017), ‘Are state-and time-dependent models really different?’, *NBER Macroeconomics Annual* **31**(1), 379–457.
- An, S. and Schorfheide, F. (2007), ‘Bayesian analysis of DSGE models’, *Econometric reviews* **26**(2-4), 113–172.
- Ascari, G. and Haber, T. (2022), ‘Non-linearities, state-dependent prices and the transmission mechanism of monetary policy’, *The Economic Journal* **132**(641), 37–57.
- Ascari, G. and Sbordone, A. M. (2014), ‘The macroeconomics of trend inflation’, *Journal of Economic Literature* **52**(3), 679–739.
- Auclert, A., Rigato, R., Rognlie, M. and Straub, L. (2024), ‘New pricing models, same old phillips curves?’, *The Quarterly Journal of Economics* **139**(1), 121–186.
- Beaudry, P., Hou, C. and Portier, F. (2024), The Dominant Role of Expectations and Broad-Based Supply Shocks in Driving Inflation, in M. Eichenbaum, J. Leahy and V. Ramey, eds, ‘NBER Macroeconomics Annual 2025’, The University of Chicago Press. forthcoming.
- Beaudry, P., Hou, C. and Portier, F. (2025), On the fragility of the nonlinear phillips curve view of recent inflation, Technical report, National Bureau of Economic Research.
- Benigno, P. and Eggertsson, G. B. (2023), It’s baaack: The surge in inflation in the 2020s and the return of the non-linear phillips curve, Technical report, National Bureau of Economic Research.
- Bergholt, D., Furlanetto, F. and Vaccaro-Grange, E. (2024), ‘Did monetary policy kill the phillips curve? some simple arithmetics’, *The Review of Economics and Statistics* pp. 1–45.
- Blanco, A., Boar, C., Jones, C. J. and Midrigan, V. (2024), The inflation accelerator, Technical report, National Bureau of Economic Research.
- Bobeica, E., Holton, S., Huber, F. and Hernández, C. M. (2025), ‘Beware of large shocks! a non-parametric structural inflation model’.

- Calvo, G. A. (1983), ‘Staggered prices in a utility-maximizing framework’, *Journal of Monetary Economics* **12**(3), 383–398.
- Caratelli, D. and Halperin, B. (2025), ‘Optimal monetary policy under menu costs’.
- Caravello, T. and Martínez-Bruera, P. (2024), ‘Disentangling sign and size non-linearities’, SSRN Working Paper. Available at SSRN: <https://ssrn.com/abstract=4704050> or <http://dx.doi.org/10.2139/ssrn.4704050>.
- Carvalho, C. and Kryvtsov, O. (2021), ‘Price selection’, *Journal of Monetary Economics* **122**, 56–75.
- Cavallo, A., Lippi, F. and Miyahara, K. (2024), ‘Large shocks travel fast’, *American Economic Review: Insights* **6**(4), 558–574.
- Costain, J. and Nakov, A. (2019), ‘Logit price dynamics’, *Journal of Money, Credit and Banking* **51**(1), 43–78.
- Costain, J. and Nakov, A. (2024), Models of price setting and inflation dynamics, in G. Ascari and R. Trezzi, eds, ‘Research Handbook of Inflation’, Edward Elgar Publishing, UK.
- Costain, J., Nakov, A. and Petit, B. (2022), ‘Flattening of the phillips curve with state-dependent prices and wages’, *The Economic Journal* **132**(642), 546–581.
- Cuba-Borda, P., Guerrieri, L., Iacoviello, M. and Zhong, M. (2019), ‘Likelihood evaluation of models with occasionally binding constraints’, *Journal of Applied Econometrics* **34**(7), 1073–1085.
- de Veirman, E. (2023), How Does the Phillips Curve Slope Vary with Repricing Rates?, ECB Working Paper, 2023/2804.
- Dixit, A. K. and Stiglitz, J. E. (1977), ‘Monopolistic competition and optimum product diversity’, *The American Economic Review* **67**(3), 297–308.
- Erceg, C. J., Lindé, J. and Trabandt, M. (2024), Monetary policy and inflation scares, IMF Working Paper No. 2024/260.
- Faber, M. and Züllig, G. (2025), ‘Kinky europe: Evidence from the regional phillips curve in the euro area’.
- Fair, R. C. and Taylor, J. B. (1983), ‘Solution and maximum likelihood estimation of dynamic nonlinear rational expectations models’, *Econometrica* **51**(4), 1169–1185.
- Fratto, C. and Uhlig, H. (2020), ‘Accounting for post-crisis inflation: A retro analysis’, *Review of Economic Dynamics* **35**, 133–153.
- Gagliardone, L., Gertler, M., Lenzu, S. and Tielens, J. (2025), Micro and macro cost-price dynamics in normal times and during inflation surges, Technical report, National Bureau of Economic Research.
- Gasteiger, E. and Grimaud, A. (2023), ‘Price setting frequency and the Phillips curve’, *European Economic Review* **158**, 104535.

- Gautier, E., Conflitti, C., Enderle, D., Fadejeva, L., Grimaud, A., Gutiérrez, E., Jouvanceau, V., Menz, J.-O., Paulus, A., Petroulas, P., Roldan-Blanco, P. and Wieland, E. (2026), 'Price stickiness in the euro area in times of high inflation', *Unpublished Manuscript*.
- Ghassibe, M. and Nakov, A. (2025), 'Business cycles with pricing cascades'.
- Gitti, G. (2025), Nonlinearities in the regional phillips curve with labor market tightness, mimeo.
- Gödl, M. and Gödl-Hanisch, I. (2024), Wage Setting in Times of High and Low Inflation, Working Paper.
- Golosov, M. and Lucas, Jr., R. E. (2007), 'Menu costs and phillips curves', *Journal of Political Economy* **115**(2), 171–199.
- Gorodnichenko, Y. and Coibion, O. (2025), Inflation, Expectations And Monetary Policy: What Have We Learned And To What End?, mimeo.
- Guerrieri, L. and Iacoviello, M. (2017), 'Collateral constraints and macroeconomic asymmetries', *Journal of Monetary Economics* **90**, 28–49.
- Hansen, G. D. (1985), 'Indivisible labor and the business cycle', *Journal of Monetary Economics* **16**(3), 309–327.
- Harding, M., Lindé, J. and Trabandt, M. (2022), 'Resolving the missing deflation puzzle', *Journal of Monetary Economics* **126**, 15–34.
- Harding, M., Lindé, J. and Trabandt, M. (2023), 'Understanding post-covid inflation dynamics', *Journal of Monetary Economics* **140**, S101–S118.
- Känzig, D. R. (2021), 'The macroeconomic effects of oil supply news: Evidence from opec announcements', *American Economic Review* **111**(4), 1092–1125.
- Karadi, P., Nakov, A., Nuno, G., Pasten, E. and Thaler, D. (2024), Strike while the iron is hot: optimal monetary policy with a nonlinear phillips curve, CEPR Discussion Paper No. 19339, CEPR Press, Paris & London.
- Karadi, P. and Reiff, A. (2019), 'Menu costs, aggregate fluctuations, and large shocks', *American Economic Journal: Macroeconomics* **11**(3), 111–46.
- Karadi, P., Schoenle, R. and Wursten, J. (2024), 'Price selection in the microdata', *Journal of Political Economy Macroeconomics* **2**(2), 228–271.
- Krippner, L. (2013), 'Measuring the stance of monetary policy in zero lower bound environments', *Economics Letters* **118**(1), 135–138.
- L'Huillier, J.-P. and Schoenle, R. (2024), 'Raising the inflation target: What are the effective gains in policy room?'
- Lucas, Jr., R. E. (1977), Understanding business cycles, in 'Carnegie-Rochester Conference Series on Public Policy', Vol. 5, Elsevier, pp. 7–29.
- Luo, S. and Villar, D. (2021), 'The price adjustment hazard function: Evidence from high inflation periods', *Journal of Economic Dynamics and Control* **130**, 104135.

- Matějka, F. and McKay, A. (2015), ‘Rational inattention to discrete choices: A new foundation for the multinomial logit model’, *American Economic Review* **105**(1), 272–298.
- Midrigan, V. (2011), ‘Menu costs, multiproduct firms, and aggregate fluctuations’, *Econometrica* **79**(4), 1139–1180.
- Montag, H. and Vallenas, D. V. (2023), Price-setting during the covid era, FEDS Notes. Washington: Board of Governors of the Federal Reserve System, August 29, 2023.
- Nakamura, E., Steinsson, J., Sun, P. and Villar, D. (2018), ‘The elusive costs of inflation: Price dispersion during the U.S. Great Inflation’, *Quarterly Journal of Economics* **133**(4), 1933–1980.
- Phillips, A. W. (1958), ‘The relation between unemployment and the rate of change of money wage rates in the united kingdom, 1861-1957’, *Economica* **25**(100), 283–299.
- Reiter, M. and Wende, A. (2024), Generalized rotemberg price-setting, CESifo Working papers 11297.
- Rogerson, R. (1988), ‘Indivisible labor, lotteries and equilibrium’, *Journal of Monetary Economics* **21**(1), 3–16.
- Sahuc, J.-G., Smets, F. and Vermandel, G. (2025), ‘The new Keynesian climate model’.
- Sims, C. A. and Zha, T. (2006), ‘Does monetary policy generate recessions?’, *Macroeconomic Dynamics* **10**(2), 231–272.
- Sims, E. and Wu, J. C. (2021), ‘Evaluating central banks’ tool kit: Past, present, and future’, *Journal of Monetary Economics* **118**, 135–160.
- Smets, F. and Wouters, R. (2007), ‘Shocks and frictions in US business cycles: A Bayesian DSGE approach’, *American Economic Review* **97**(3), 586–606.
- Woodford, M. (2011), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.

Online Appendix

A. Macroeconomic evidence

This appendix describes the local projection framework used to estimate the impulse responses of inflation to oil supply shocks. We work with monthly data and measure inflation as year-on-year HICP inflation. Let P_t denote the HICP at month t . Inflation is defined as $\pi_t = 100 \times (\log P_t - \log P_{t-12})$, so that π_t corresponds to the annual inflation rate expressed in percentage points. Using year-on-year inflation mitigates seasonality and ensures comparability across horizons in the local projection setting.

Impulse response functions are estimated using the local projection methodology. For each horizon $h = 0, 1, \dots, H$, we estimate

$$\pi_{t+h} = \alpha_h + \sum_{\ell=0}^{11} g_h(z_{t-\ell}) + \mathbf{X}'_{t-1} \gamma_h + \varepsilon_{t+h},$$

where z_t denotes the monthly oil supply shock constructed by [Känzig \(2021\)](#), standardized to unit variance. The vector $\mathbf{X}_{t-1}$ includes lagged controls; in the baseline specification, it consists of 11 lags of year-on-year inflation. Standard errors are computed using Newey–West corrections. The function $g_h(\cdot)$ governs the transmission of the shock and allows for different forms of nonlinearity.

Following [Caravello and Martínez-Bruera \(2024\)](#), we consider three specifications.

Linear specification.

$$g_h(z) = \beta_{1,h}z, \quad \text{IRF}_h(s) = \beta_{1,h}s,$$

which implies responses that are linear and symmetric in both sign and size.

Sign-dependent specification.

$$g_h(z) = \beta_{1,h}z + \beta_{2,h}z|z|, \quad \text{IRF}_h(s) = \beta_{1,h}s + \beta_{2,h}s|s|,$$

which allows responses to differ between positive and negative shocks of equal magnitude.

Size-dependent specification.

$$g_h(z) = \beta_{1,h}z + \beta_{2,h}z^2, \quad \text{IRF}_h(s) = \beta_{1,h}s + \beta_{2,h}s^2,$$

which allows responses to vary with the magnitude of the shock independently of its sign.

B. The model

Time t is discrete. We borrow the notation from [Sims and Wu \(2021\)](#) and develop a New Keynesian model with consumption habits, indexation to lagged inflation, and state-dependent sticky prices and wages following [Gasteiger and Grimaud \(2023\)](#). The state-dependency allows for intensive *and* extensive margin adjustments in prices and wages.

Households. The lifetime utility of household $k \in [0, 1]$ is given by

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \left\{ \frac{1}{1-\sigma} \left(\frac{C_{t+j}(k)}{C_{t-1+j}^h} \right)^{1-\sigma} C_{t-1+j}^h - \frac{\psi}{1+\chi} L_{t+j}(k)^{1+\chi} \right\},$$

with relative risk aversion $\sigma > 1$, discount factor $0 < \beta < 1$, $\psi > 0$ is a scaling parameter for labor disutility, and, χ is the inverse Frisch elasticity. We incorporate multiplicative external consumption habits into the utility function via parameter $0 \leq h < 1$. This creates complementarities between individual agent consumption at time t and aggregate consumption at $t - 1$. The utility function is increasing and concave in both individual and aggregate consumption. When the habit parameter $h = 0$, the utility function collapses to standard CRRA preferences. This formulation allows for the introduction of consumption smoothing in Section 4 while avoiding negative marginal utility, which could destabilize the model under nonlinear simulations in response to large shocks.

Household k , faces the nominal budget constraint

$$P_t C_t(k) + B_t(k) \leq (1 - \varepsilon_{w,t}) MRS_t L_t(k) + (1 - \varepsilon_{d,t}) R_{t-1} B_{t-1}(k) + DIV_t + T_t,$$

where P_t represents the price level of goods, and B_t denotes a one-period bond that pays the short-term gross interest rate R_t . The term MRS_t corresponds to the nominal remuneration households earn by supplying labor to the union. DIV_t captures the dividends households receive as owners of the private sector, while T_t represents lump-sum taxes net of subsidies, ensuring the government budget constraint is balanced.

Two exogenous AR(1) processes, $\varepsilon_{w,t}$ and $\varepsilon_{d,t}$, are introduced to generate variations in distortionary taxes or subsidies on wages and on bond income. These processes are introduced for estimation purposes. They effectively represent a labor supply shock ($\varepsilon_{w,t}$) moving the desired real wage and a shock ($\varepsilon_{d,t}$) creating a wedge between the central bank's policy rate and the return on assets held by households, similar to [Smets and Wouters's \(2007\)](#) risk-premium shock.

Aggregating over homogeneous households and dropping the k index, the first-order conditions for a representative agent in real terms are

$$\psi L_t^\chi (C_t C_{t-1}^{-h})^\sigma = (1 - \varepsilon_{w,t}) mrs_t$$

$$1 = (1 - \varepsilon_{d,t}) R_t \mathbb{E}_t \Lambda_{t,t+1} \pi_{t+1}^{-1}$$

$$\Lambda_{t,t+1} = \beta \left(\frac{C_{t+1} C_t^{-h}}{C_t C_{t-1}^{-h}} \right)^{-\sigma}$$

with $\pi_t = P_t/P_{t-1}$ the gross inflation rate, $mrs_t = MRS_t/P_t$ the real remuneration for labor supply and $\Lambda_{t,t+1}$ the real stochastic discount factor. Clearly, multiplicative consumption habits affect both the labor supply and the inter-temporal consumption decision.

Production. Production is divided into several layers. A representative wholesale firm hires labor from the labor packer and produces output. The latter is sold to a continuum of retail firms at price $P_{w,t}$. Each retail firm $i \in [0, 1]$ repackages wholesale output, and sells it to a competitive final goods firm at $P_t(f)$.

A final goods firm transforms retail output into a final output good according to $Y_t = \left[\int_0^1 Y_t(f)^{\frac{\epsilon_p - 1}{\epsilon_p}} df \right]^{\frac{\epsilon_p}{\epsilon_p - 1}}$. Profit maximization by the final goods firm gives a demand for each retail output as $Y_t(f) = \left(\frac{P_t(f)}{P_t} \right)^{-\epsilon_p} Y_t$, and, the aggregate price index as $P_t^{1-\epsilon_p} = \int_0^1 P_t(f)^{1-\epsilon_p} df$.

Retailers are price-setters. Following [Gasteiger and Grimaud \(2023\)](#), the retailers discount future real dividends with the stochastic discount factor of the household, as well as the time-varying price-setting frequency. The price-setting problem is therefore:

$$\begin{aligned} \max_{P_t(f)} \mathbb{E}_t \sum_{j=0}^{\infty} \left(\prod_{k=0}^j \phi_{p,t+k} \right) \phi_{p,t}^{-1} \Lambda_{t,t+j} \times \\ \times \left\{ P_t(f)^{1-\epsilon_p} (P_{t+j}/\Pi_{t,t+j}^{\varrho_p})^{\epsilon_p - 1} Y_{t+j} - P_{w,t+j} P_t(f)^{-\epsilon_p} (P_{t+j}/\Pi_{t,t+j}^{\varrho_p})^{-\epsilon_p} Y_{t+j} \right\}, \end{aligned}$$

where $0 \leq \varrho_p < 1$ is a parameter governing indexation to past inflation. We define the optimal reset price as $P_t^\#$, since it does not depend on the index f . The first-order condition is recursively given by

$$\begin{aligned} P_t^\# &= \frac{\epsilon_p}{\epsilon_p - 1} \frac{X_{1,t}}{X_{2,t}} \\ X_{1,t} &= p_{w,t} (P_t/\Pi_{t,t}^{\varrho_p})^{\epsilon_p} Y_t + \mathbb{E}_t \phi_{p,t+1} \Lambda_{t,t+1} X_{1,t+1} \\ X_{2,t} &= (P_t/\Pi_{t,t}^{\varrho_p})^{\epsilon_p - 1} Y_t + \mathbb{E}_t \phi_{p,t+1} \Lambda_{t,t+1} X_{2,t+1} \end{aligned}$$

With $p_{w,t} = P_{w,t}/P_t$ interpreted as real marginal costs. We define $x_{1,t} = X_{1,t}/P_t^{\epsilon_p}$ and $x_{2,t} =$

$X_{2,t}/P_t^{\epsilon_p-1}$ so as to have

$$\begin{aligned} x_{1,t} &= p_{w,t} Y_t + \mathbb{E}_t \phi_{p,t+1} \Lambda_{t,t+1} \left(\frac{\pi_{t+1}}{\pi_t^{\epsilon_p}} \right)^{\epsilon_p} x_{1,t+1} \\ x_{2,t} &= Y_t + \mathbb{E}_t \phi_{p,t+1} \Lambda_{t,t+1} \left(\frac{\pi_{t+1}}{\pi_t^{\epsilon_p}} \right)^{\epsilon_p-1} x_{2,t+1} \\ p_t^\# &= \frac{\epsilon_p}{\epsilon_p - 1} \frac{x_{1,t}}{x_{2,t}} \end{aligned}$$

with $p_t^\# = P_t^\# / P_t$. The wholesale firm produces output following

$$Y_{W,t} = L_{d,t}.$$

Production is linear so firms marginal cost is equal to real wage subject to a shock

$$(1 - \varepsilon_{p,t}) w_t = p_{w,t},$$

in which $\varepsilon_{p,t}$ represents an exogenous tax or subsidy on marginal costs, introduced to generate supply-side cost-push shocks. Unlike labor supply shocks, these shocks move real wage and inflation in opposite directions.

Price-resetting agency. We microfound the endogenous non-resetting share $\phi_{p,t}$ through a centralized price-setting agency. The agency is owned by the representative household and owns the continuum of intermediate-goods firms, so it internalizes firm profits; any net payoff is rebated lump-sum to households. Its role is limited to the timing of re-optimization: when adjustment is authorized, firms reset to the common optimal price $P_t^\#$. In each period t , the agency chooses $d_{p,t} \in \{0, 1\}$, where $d_{p,t} = 1$ authorizes a reset to $P_t^\#$ and $d_{p,t} = 0$ implies continuation at legacy prices. Let $\mathcal{V}_{p,i,t,d}$ denote firm i 's continuation value conditional on action d . Absent informational frictions, the agency would choose $d_{p,t}$ to maximize $\int_0^1 \mathcal{V}_{p,i,t,d} di$. We assume that for the agency, conditioning sharply on the state is costly in terms of information, therefore it chooses adjustment probabilities $\mathcal{P}_{p,0,t}$ and $\mathcal{P}_{p,1,t}$. Under symmetric implementation, these probabilities coincide with equilibrium shares,²⁰

$$\phi_{p,t} \equiv \mathcal{P}_{p,0,t}, \quad 1 - \phi_{p,t} \equiv \mathcal{P}_{p,1,t}.$$

Accordingly, aggregate continuation value takes the Calvo-style mixture form

$$\int_0^1 \mathcal{V}_{p,i,t} di = \mathcal{P}_{p,1,t} \int_0^1 \mathcal{V}_{p,i,t,1} di + \mathcal{P}_{p,0,t} \int_0^1 \mathcal{V}_{p,i,t,0} di.$$

²⁰As in Blanco et al. (2024), our setup abstracts from the selection effect that arises when adjustment decisions depend on firm-level states: the agency's symmetric implementation and the insured non-adjustment payoff imply that adjustment shares can be treated as probabilities without tracking the cross-sectional composition of adjusters.

Derivations of the present values of the firms. We assume that each firm i in every period either updates the price optimally, $P_{i,t}^\#$, or believe to update the price to the average legacy price, $P_{i,t}^f$. By symmetry, we have $P_{i,t}^x = P_t^x$ for $x \in \{\#, f\}$. From the FOC conditions, the value of a firm of type x is

$$\begin{aligned} V_{p,t|t}^x &= \mathbb{E}_t \left\{ \sum_{j=0}^{\infty} \left(\prod_{k=0}^j \phi_{p,t+k} \right) \phi_{p,t}^{-1} \Lambda_{t,t+j} \left[\left(\frac{P_t^x}{P_{t+j} \Pi_{t,t+j}^{-\varrho_p}} \right)^{1-\epsilon_p} - p_{w,t+j} \left(\frac{P_t^x}{P_{t+j} \Pi_{t,t+j}^{-\varrho_p}} \right)^{-\epsilon_p} \right] Y_{t+j} \right\} \\ &= \mathbb{E}_t \left\{ \sum_{j=0}^{\infty} \left(\prod_{k=0}^j \phi_{p,t+k} \right) \phi_{p,t}^{-1} \Lambda_{t,t+j} \left[\left(\frac{P_t^x}{P_t} \frac{P_t}{P_{t+j} \Pi_{t,t+j}^{-\varrho_p}} \right)^{1-\epsilon_p} - p_{w,t+j} \left(\frac{P_t^x}{P_t} \frac{P_t}{P_{t+j} \Pi_{t,t+j}^{-\varrho_p}} \right)^{-\epsilon_p} \right] Y_{t+j} \right\} \\ &= \mathbb{E}_t \left\{ \sum_{j=0}^{\infty} \left(\prod_{k=0}^j \phi_{p,t+k} \right) \phi_{p,t}^{-1} \Lambda_{t,t+j} \left[\left(\frac{p_t^x}{\Pi_{t,t+j} \Pi_{t,t+j}^{-\varrho_p}} \right)^{1-\epsilon_p} - p_{w,t+j} \left(\frac{p_t^x}{\Pi_{t,t+j} \Pi_{t,t+j}^{-\varrho_p}} \right)^{-\epsilon_p} \right] Y_{t+j} \right\}, \end{aligned}$$

where we define the relative price $p_t^x \equiv P_t^x / P_t$. Next, we split $V_{p,t|t}^x$ into

$$V_{p,t|t}^x = V_{p,x_2,t|t}^x - V_{p,x_1,t|t}^x,$$

where

$$\begin{aligned} V_{p,x_1,t|t}^x &\equiv \mathbb{E}_t \left\{ \sum_{j=0}^{\infty} \left(\prod_{k=0}^j \phi_{p,t+k} \right) \phi_{p,t}^{-1} \Lambda_{t,t+j} \left(\frac{p_t^x}{\Pi_{t,t+j} \Pi_{t,t+j}^{-\varrho_p}} \right)^{1-\epsilon_p} Y_{t+j} \right\} \quad \text{and} \\ V_{p,x_2,t|t}^x &\equiv \mathbb{E}_t \left\{ \sum_{j=0}^{\infty} \left(\prod_{k=0}^j \phi_{p,t+k} \right) \phi_{p,t}^{-1} \Lambda_{t,t+j} p_{w,t+j} \left(\frac{p_t^x}{\Pi_{t,t+j} \Pi_{t,t+j}^{-\varrho_p}} \right)^{-\epsilon_p} Y_{t+j} \right\} \end{aligned}$$

Next, recall the definitions of the auxiliary variables $x_{1,t}$ and $x_{2,t}$. Thus, we obtain

$$V_{p,t|t}^x = p_t^{x,1-\epsilon_p} x_{2,t} - p_t^{x,-\epsilon_p} x_{1,t}.$$

The problem of the agency. Resetting entails a fixed non-pecuniary burden $\tau_p > 0$. Let $V_{p,t}^\#$ denote the continuation value associated with resetting to $P_t^\#$. Then, with symmetric firms,

$$\mathcal{V}_{p,i,t,1} \equiv V_{p,t}^\# - \tau_p, \quad V_{p,t}^\# \equiv V_p(p_{p,t}^\#).$$

When prices are not reset ($d_{p,t} = 0$), firms operate with heterogeneous legacy prices and continuation values $V_{p,i,t,0}^f$. We assume the payoff relevant for the agency's decision is

$$V_{p,t}^f \equiv V_p(p_{p,t}^f), \quad p_{p,t}^f \equiv \frac{P_{t-1}}{P_t} = \frac{1}{\pi_t}.$$

Hence, the agency compares

$$\mathcal{V}_{p,0,t} \equiv V_{p,t}^f, \quad \mathcal{V}_{p,1,t} \equiv V_{p,t}^\# - \tau_p,$$

and chooses $(\mathcal{P}_{p,0,t}, \mathcal{P}_{p,1,t})$ to maximize the continuation value against an information-processing cost $\Omega(\mathcal{P}_{p,0,t}, \mathcal{P}_{p,1,t})$

$$\begin{aligned} \max_{\mathcal{P}_{p,0,t}, \mathcal{P}_{p,1,t}} & \left\{ \mathcal{P}_{p,0,t} V_{p,t}^f + \mathcal{P}_{p,1,t} (V_{p,t}^\# - \tau_p) - \frac{1}{\gamma_p} \Omega(\mathcal{P}_{p,0,t}, \mathcal{P}_{p,1,t}) \right\}, \\ & 1 = \mathcal{P}_{p,0,t} + \mathcal{P}_{p,1,t}. \end{aligned}$$

where $\gamma_p > 0$ is the inverse marginal cost of information processing. Under an entropy-based information cost and a symmetric prior, [Matějka and McKay \(2015, Theorem 1\)](#) shows that the FOC condition boils down to the logit policy

$$\mathcal{P}_{p,0,t} \equiv \phi_{p,t} = \frac{\exp(\gamma_p V_{p,t}^f)}{\exp(\gamma_p V_{p,t}^f) + \exp(\gamma_p (V_{p,t}^\# - \tau_p))}. \quad (\text{B.1})$$

Equation (B.1) implies that the non-resetting share $\phi_{p,t}$ is a smooth, state-dependent function of the difference between the continuation values of non-adjustment and adjustment. The fixed burden $\tau_p > 0$ shifts the adjustment value $V_{p,t}^\# - \tau_p$ downward and therefore raises $\phi_{p,t}$, all else equal. The intensity of choice parameter $\gamma_p > 0$ governs the strength of state contingency. When γ_p is small, information is expensive and the agency responds weakly to movements in $V_{p,t}^f$ and $V_{p,t}^\#$; adjustment becomes nearly random and $\phi_{p,t}$ is close to a constant, delivering the Calvo benchmark as $\gamma_p \rightarrow 0$. As γ_p increases, information becomes cheaper, the agency conditions more sharply on the state, and $\phi_{p,t}$ varies more with the underlying present values. In the limit $\gamma_p \rightarrow \infty$, choices become nearly deterministic: the agency puts almost all probability on the higher-value action, so price adjustment becomes effectively state-dependent. For any finite γ_p , adjustment remains probabilistic and prices are therefore rigid in the sense that a positive mass of firms does not re-optimize each period.

Labor market. The labor market consists of three layers. There is a continuum of labor unions index by $l \in [0, 1]$. They purchase labor from households at MRS_t and repackage it for sale to a representative labor packer $W_t(l)$. The labor packer combines differentiated labor inputs into final labor, $L_{d,t}$, via a CES technology with elasticity of substitution $\epsilon_w > 1$. Therefore, the labor demand curve faced by each union l is

$$L_t(l) = \left(\frac{W_t(l)}{W_t} \right)^{-\epsilon_w} L_{d,t},$$

and, the aggregate wage index is

$$W_t^{1-\epsilon_w} = \int_0^1 W_t(l)^{1-\epsilon_w} dl.$$

Labor unions are wage setters. Following [Gasteiger and Grimaud \(2023\)](#), the problem for union l that updates optimally is to choose a wage that maximizes the discounted value of future real profits. Unions apply the household's stochastic discount factor to get the net present value of profits, and take the time-varying wage-setting frequency as given, i.e.,

$$\begin{aligned} \max_{W_t(l)} \mathbb{E}_t \sum_{j=0}^{\infty} \left(\prod_{k=0}^j \phi_{w,t+k} \right) \phi_{w,t}^{-1} \Lambda_{t,t+j} \\ \times \left\{ (W_t(l) \Pi_{t,t+j}^{\varrho_w})^{1-\epsilon_w} W_{t+j}^{\epsilon_w} P_{t+j}^{-1} L_{d,t+j} - mrs_{t+j} (W_t(l) \Pi_{t,t+j}^{\varrho_w})^{-\epsilon_w} W_{t+j}^{\epsilon_w} L_{d,t+j} \right\} \end{aligned}$$

with

$$\Pi_{t,t+j} \equiv \begin{cases} \frac{P_{t+1}}{P_t} \times \frac{P_{t+2}}{P_{t+1}} \times \dots \times \frac{P_{t+j}}{P_{t+j-1}} & \text{for } j = 1, 2, \dots \\ 1 & \text{for } j = 0. \end{cases}$$

$0 \leq \varrho_w < 1$ captures partial indexation of wages to past inflation. Assuming that the reset wage is the same across all labor unions, the index l can be dropped. The first-order condition defining optimal reset wage $W_t^\#$ can be written in real terms as

$$w_t^\# = \frac{\epsilon_w}{\epsilon_w - 1} \frac{F_{1,t}}{F_{2,t}},$$

where

$$\begin{aligned} F_{1,t} &= mrs_t w_t^{\epsilon_w} (P_t / \Pi_{t,t+j}^{\varrho_w})^{\epsilon_w} L_{d,t} + \mathbb{E}_t \phi_{w,t+1} \Lambda_{t,t+1} F_{1,t+1} \\ F_{2,t} &= w_t^{\epsilon_w} (P_t / \Pi_{t,t+j}^{\varrho_w})^{\epsilon_w - 1} L_{d,t} + \mathbb{E}_t \phi_{w,t+1} \Lambda_{t,t+1} F_{2,t+1} \end{aligned}$$

Defining $f_{1,t} = F_{1,t} / P_t^{\epsilon_w}$ and $f_{2,t} = F_{2,t} / P_t^{\epsilon_w - 1}$ we get

$$w_t^\# = \frac{\epsilon_w}{\epsilon_w - 1} \frac{f_{1,t}}{f_{2,t}},$$

where

$$\begin{aligned} f_{1,t} &= mrs_t w_t^{\epsilon_w} L_{d,t} + \mathbb{E}_t \phi_{w,t+1} \Lambda_{t,t+1} \left(\frac{\pi_{t+1}}{\pi_t^{\varrho_w}} \right)^{\epsilon_w} f_{1,t+1} \\ f_{2,t} &= w_t^{\epsilon_w} L_{d,t} + \mathbb{E}_t \phi_{w,t+1} \Lambda_{t,t+1} \left(\frac{\pi_{t+1}}{\pi_t^{\varrho_w}} \right)^{\epsilon_w - 1} f_{2,t+1}. \end{aligned}$$

Wage-resetting agency. Wage re-optimization is microfounded using the same delegated-agency setup as for prices. A wage-setting agency chooses whether unions re-optimize their wage in period t . If a reset is authorized, unions set the common optimal real wage $w_t^\# \equiv W_t^\# / P_t$, delivering continuation value $V_{w,t}^\#$. If no reset is authorized, unions continue at heterogeneous legacy wages. As in the price block, we impose the restriction that the relevant non-adjustment payoff depends only on the average legacy real wage $w_t^f \equiv W_{t-1} / P_t = w_{t-1} / \pi_t$, implying a common continuation value $V_{w,t}^f$ for the agency's decision. Resetting entails a fixed non-pecuniary burden $\tau_w > 0$.

Derivations of the present values of the unions. We assume that each union l in every period either updates the wage optimally, $W_{i,t}^\#$, or, assume to update the wage to the average old wage $W_{i,t}^f$. By symmetry, we have $W_{i,t}^x = W_t^x$ for $x \in \{\#, f\}$. Thus the value of a union of type x is

$$V_{w,t|t}^x = \mathbb{E}_t \left\{ \sum_{j=0}^{\infty} \left(\prod_{k=0}^j \phi_{w,t+k} \right) \phi_{w,t}^{-1} \Lambda_{t,t+j} (W_t^x - MRS_{t+j}) L_{x,t+j|t} P_{t+j}^{-1} \Pi_{t,t+j}^{\varrho_w} \right\}.$$

with demand for the firm's good given by

$$l_{x,t+j|t} = \left(\frac{W_t^x}{W_{t+j}} \right)^{-\epsilon_w} L_{d,t+j}.$$

Eliminating $L_{x,t+j|t}$ yields

$$\begin{aligned} V_{w,t|t}^x &= \mathbb{E}_t \left\{ \sum_{j=0}^{\infty} \left(\prod_{k=0}^j \phi_{w,t+k} \right) \phi_{w,t}^{-1} \Lambda_{t,t+j} [W_t^x - MRS_{t+j}] \left(\frac{W_t^x}{W_{t+j}} \right)^{-\epsilon_w} L_{d,t+j} P_{t+j}^{-1} \Pi_{t,t+j}^{\varrho_w} \right\} \\ &= \mathbb{E}_t \left\{ \sum_{j=0}^{\infty} \left(\prod_{k=0}^j \phi_{w,t+k} \right) \phi_{w,t}^{-1} \Lambda_{t,t+j} \left[W_t^{x,1-\epsilon_w} W_{t+j}^{\epsilon_w} P_{t+j}^{-1} L_{d,t+j} \Pi_{t,t+j}^{\varrho_w} - mrs_{t+j} W_t^{x,-\epsilon_w} W_{t+j}^{\epsilon_w} L_{d,t+j} \Pi_{t,t+j}^{\varrho_w} \right] \right\} \end{aligned}$$

Next, we split $V_{w,t|t}^x$ into

$$V_{w,t|t}^x = V_{w,f_2,t|t}^x - V_{w,f_1,t|t}^x,$$

where

$$\begin{aligned} V_{w,f_1,t|t}^x &\equiv \mathbb{E}_t \left\{ \sum_{j=0}^{\infty} \left(\prod_{k=0}^j \phi_{w,t+k} \right) \phi_{w,t}^{-1} \Lambda_{t,t+j} mrs_{t+j} W_t^{x,-\epsilon_w} W_{t+j}^{\epsilon_w} L_{d,t+j} \Pi_{t,t+j}^{\varrho_w} \right\} \quad \text{and} \\ V_{w,f_2,t|t}^x &\equiv \mathbb{E}_t \left\{ \sum_{j=0}^{\infty} \left(\prod_{k=0}^j \phi_{w,t+k} \right) \phi_{w,t}^{-1} \Lambda_{t,t+j} W_t^{x,1-\epsilon_w} W_{t+j}^{\epsilon_w} P_{t+j}^{-1} L_{d,t+j} \Pi_{t,t+j}^{\varrho_w} \right\} \end{aligned}$$

recursively we have

$$\begin{aligned} V_{w,f_1,t|t}^x &\equiv W_t^{x,-\epsilon_w} \times F_{1,t} & \text{and} \\ V_{w,f_2,t|t}^x &\equiv W_t^{x,1-\epsilon_w} \times F_{2,t} \end{aligned}$$

multiplying by $\frac{p_t^{\epsilon_w}}{p_t^{\epsilon_w}}$ and $\frac{p_t^{1-\epsilon_w}}{p_t^{1-\epsilon_w}}$ we can express all parts in real terms

$$\begin{aligned} V_{w,f_1,t|t}^x &\equiv w_t^{x,-\epsilon_w} \times f_{1,t} & \text{and} \\ V_{w,f_2,t|t}^x &\equiv w_t^{x,1-\epsilon_w} \times f_{2,t} \end{aligned}$$

so we obtain

$$V_{w,t|t}^x = w_t^{x,1-\epsilon_w} f_{2,t} - w_t^{x,-\epsilon_w} f_{1,t}.$$

In the same way as for prices, the wage-resetting agency faces an information-processing constraint and chooses a probabilistic authorization rule. Under an entropy-based information cost, the implied non-resetting share follows the multinomial-logit form ([Matějka and McKay, 2015](#)):

$$\phi_{w,t} = \frac{\exp(\gamma_w V_{w,t}^f)}{\exp(\gamma_w V_{w,t}^f) + \exp(\gamma_w (V_{w,t}^\# - \tau_w))},$$

where $\gamma_w > 0$ is the intensity of choice (inverse marginal cost of information), $\phi_{w,t} \in [0, 1]$ is the share of unions not re-optimising, and $1 - \phi_{w,t}$ is the reset share.

Monetary Policy. The gross nominal rate R_t is set according to a linear Taylor rule

$$R_t = \rho R_{t-1} + (1 - \rho) \left(\bar{R} + \theta_\pi (\pi_t - \bar{\pi}) + \theta_y \left(\frac{Y_t}{\bar{Y}} - 1 \right) \right) + \varepsilon_{r,t}$$

with θ_π the reaction to inflation deviation, θ_y the reaction to the output gap, ρ the smoothing parameter, and $\varepsilon_{r,t}$ an AR(1) monetary policy shock.

Aggregation. Using the properties of [Gasteiger and Grimaud \(2023\)](#) derived from the [Calvo \(1983\)](#) model, we obtain that the aggregate inflation rate and aggregate real wages laws of motion

$$\begin{aligned} 1 &= (1 - \phi_{p,t}) (\pi_t^\#)^{1-\epsilon_p} + \phi_{p,t} \left(\frac{\pi_t}{\pi_{t-1}^{\varrho_p}} \right)^{\epsilon_p-1} \\ w_t^{1-\epsilon_w} &= (1 - \phi_{w,t}) (w_t^\#)^{1-\epsilon_w} + \phi_{w,t} \left(\frac{\pi_t}{\pi_{t-1}^{\varrho_w}} \right)^{\epsilon_w-1} w_{t-1}^{1-\epsilon_w} \end{aligned}$$

Aggregate production function across the continuum of retailers is

$$Y_{W,t} = Y_t v_t^p$$

with $v_t^p = \int_0^1 \left(\frac{P_t(f)}{P_t} \right)^{-\epsilon_p} df$ being a measure of price dispersion, that we can write recursively, i.e.,

$$v_t^p = (1 - \phi_{p,t}) (\pi_t^\#)^{-\epsilon_p} + \phi_{p,t} \left(\frac{\pi_t}{\pi_{t-1}^{\epsilon_p}} \right)^{\epsilon_p} v_{t-1}^p.$$

For labor market clearing, we integrate the demand for union labors across unions and obtain

$$L_t = L_{d,t} v_t^w$$

with $v_t^w = \int_0^1 \left(\frac{W_t(l)}{W_t} \right)^{-\epsilon_w} dl$ being a measure of wage dispersion, which can also be written recursively, i.e.,

$$v_t^w = (1 - \phi_{w,t}) \left(\frac{w_t^\#}{w_t} \right)^{-\epsilon_w} + \phi_{w,t} \left(\frac{\pi_t}{\pi_{t-1}^{\epsilon_w}} \right)^{\epsilon_w} \left(\frac{w_t}{w_{t-1}} \right)^{\epsilon_w} v_{t-1}^w.$$

Bonds are in zero net supply. Therefore the aggregate resource constraint is

$$Y_t = C_t.$$

Exogenous process. We have 4 exogenous variables, two supply shocks – a cost-push shock ε_t^p and a labor supply shock ε_t^w – and two demand shocks – a risk-premium shock ε_t^d and a monetary policy shock ε_t^r . We assume that they follow a stationary AR(1) processes.

$$\varepsilon_t^z = \rho_z \varepsilon_{t-1}^z + \eta_{z,t},$$

with $\eta_{z,t} \sim \text{i.i.d. } \mathcal{N}(0, \sigma_z^2)$, and $0 \leq \rho_z < 1$ for $z = \{d, p, w, r\}$.




Figure C.1: **Asymmetry in the wage-setting decision.** The blue line depicts union l 's present value as a function of its relative wage $\tilde{w}(l) = W(l)/W$, holding the aggregate environment at steady state. Point A is the optimal relative wage. Points B and C correspond to relative wages that are respectively 10% below and 10% above the optimum. The double-headed arrows measure the present-value loss relative to the optimum.

C. Asymmetry in wage setting decision

In analogy to Figure 2, Figure C.1 depicts union l 's present value as a function of its relative wage $\tilde{w}(l) = W(l)/W$. The union's present value exhibits the same basic asymmetry as the firm's present value. A relative wage below the optimum implies a substantially larger present-value loss than an equally sized relative wage above the optimum. The opportunity cost of a low relative wage is mainly the low remuneration received per unit of labor supplied, whereas a high relative wage primarily reduces labor demand.

Positive trend inflation reinforces this asymmetry by gradually eroding the relative wage of a union that does not reset. For an initially high relative wage, this erosion moves the union closer to its desired wage, whereas for an already low relative wage it increases the cost of non-adjustment. Hence, as for firms, incentives to increase wages can be substantially stronger than incentives to reduce them.

The dynamic response of wage setting nevertheless depends on more than the aggregate price or wage level. The desired reset wage also varies with labor demand and the marginal rate of substitution. Consequently, although the underlying present-value asymmetry is common to firms and unions, the shocks that activate the wage-resetting extensive margin need not be the same as those that activate price resetting.

D. Effect of trend inflation






Figure D.1: **Trend inflation and steady state.** The panels show the effect of trend inflation on the steady-state values of variables related to price (blue lines) and wage (red lines) setting, price and wage dispersion, and output. Dashed lines in Panels (e) and (f) refer to the present value of a non-resetting firm (union) that keeps the price at the average price (wage) level.

From the steady-state price and wage Calvo shares, $\bar{\phi}_p$ and $\bar{\phi}_w$, we can derive the associated resetting costs, τ_p and τ_w . Assuming a constant steady-state labor supply of $\bar{L} = 1$ and setting τ_p and τ_w to their posterior modes, we compute the model's deterministic steady state as a function of the trend inflation rate, $\bar{\pi}$. Figure D.1 presents the effect of changes in trend inflation on selected steady-state variables.

The model predicts that higher steady-state inflation increases the frequency of both wage and price adjustment—reducing $\phi_{p,t}$ and $\phi_{w,t}$, see Panel (a)—and therefore makes the economy more flexible. The mechanism is intuitive: the higher trend inflation is, the faster inflation erodes the relative price of non-resetting firms. As a result, a price-resetting firm in the standard Calvo model sets a higher price (see the discussion in [Ascari and Sbordone, 2014](#)) to account for the trend increase in the aggregate price level. In our model, firms also choose to adjust prices more frequently. The same holds for wages and unions. Higher trend inflation widens the gap between agents that reset their prices or wages and those that do not (see Panels (e) and (f)). The resulting increase in the opportunity cost of non-adjustment raises



(a) Responses to a +0.25% monetary policy shock



(b) Responses to a -0.5% cost-push shock

Notes: The parameters used are the posterior mode of the baseline model (see Table 1).

Figure D.2: IRFs as a function of the steady-state trend inflation level

the fraction of firms and unions that reset.


Positive trend inflation also amplifies the model’s sign asymmetry. Following a disinflationary shock, as long as inflation remains positive, not adjusting continues to lower an inherited relative price. Inaction therefore moves the relative price in the same direction as the lower desired reset price and partly substitutes for active adjustment. Following an inflationary shock, instead, the desired reset price rises while inaction further erodes the inherited relative price. The reset-price gap therefore widens and the incentive to adjust increases sharply. Once inflation turns negative, this logic changes: under deflation, inaction raises the inherited relative price and no longer delivers the desired downward adjustment. The same passive-adjustment mechanism applies to inherited relative wages, although the desired reset wage also depends on labor demand and the marginal rate of substitution.

Figures D.3 and D.4 isolate the role of trend inflation in the sign asymmetry. At zero trend inflation, state dependence remains, but sufficiently large shocks move inflation to either side of zero. Adjustment therefore increases in both inflationary and deflationary states, producing a more symmetric relationship between inflation and the Calvo shares. With positive trend inflation, moderate disinflationary shocks generally leave inflation above zero. Inaction then continues to lower inherited relative prices in the direction of the desired adjustment, and the extensive margin reacts only weakly. Inflationary shocks instead widen the reset-price gap and strongly increase the adjustment frequency. The asymmetry consequently becomes more pronounced as trend inflation rises. This pattern is clearest for price adjustment. Wage adjustment is also affected by this mechanism, but remains more dependent on the source of the shock because the desired reset wage responds to labor demand and the marginal rate of substitution.


This main feature of our model is both quantitatively and qualitatively consistent with results from the literature. Following L’Huillier and Schoenle (2024), we calculate the best-fitting linear steady-state relationship $f_t = \beta_0 + \beta_1 \bar{\pi} + \varepsilon_t$, where f_t is the steady-state average monthly frequency of price changes in percent²¹, and $\bar{\pi}$ is net annualized trend inflation, also in percent, implied by Panel (a) of Figure D.1. Using the parameter set calibrated at the posterior mode from the estimation of the quantitative model in Section 4, we obtain $\beta_0 = 3.36$ and $\beta_1 = 0.97599$. These figures are close to the L’Huillier and Schoenle’s (2024) estimates, $\beta_0 \in [4.61, 7.42]$ and $\beta_1 \in [0.98, 2.26]$, obtained by regressing the price-setting frequencies of Nakamura et al. (2018) on alternative measures of US trend inflation.

Figure D.2 shows how the IRFs to monetary-policy and cost-push shocks change with trend inflation. The higher the level of trend inflation, the more flexible are wages and prices. Hence, consistent with the literature, monetary-policy shocks have larger effects on inflation, while cost-push shocks have larger effects on output.

²¹When making this comparison, we disregard model-specific features such as price indexation, which implies that all prices mechanically change each period. The model-generated f_t therefore does not correspond to the overall frequency of price changes, but rather to the frequency of optimally re-optimized prices.




(a) Monetary policy shocks




(b) Labor supply shocks

Figure D.3: **Phillips curves are curves under alternative levels of trend inflation.** Stochastic simulations conditional on monetary-policy shocks (Panel (a)) and labor-supply shocks (Panel (b)). The four columns plot the output gap against, respectively, (i) price inflation, (ii) wage inflation, (iii) the share of firms not re-optimizing prices, and (iv) the share of unions not re-optimizing wages, under alternative levels of trend inflation.

Notes: For each level of trend inflation, we simulate 10,000 observations of the model while allowing only the indicated structural shock to vary. Shock innovations are drawn from a Gaussian distribution, with their scale recalibrated separately for each trend-inflation specification to reproduce the relevant volatility of the corresponding disturbance. Each dot represents one simulated period. The different series correspond to the alternative levels of trend inflation indicated in the figure. For brevity, we refer to the shares of non-reoptimized prices and wages as the *Calvo shares*. Results for risk-premium and cost-push shocks are reported in Figure D.4.



(a) Risk-premium shocks



(b) Cost-push shocks

Figure D.4: Phillips curves are curves under alternative levels of trend inflation: additional shocks. Stochastic simulations conditional on risk-premium shocks (Panel (a)) and cost-push shocks (Panel (b)). The four columns plot the output gap against, respectively, (i) price inflation, (ii) wage inflation, (iii) the share of firms not re-optimizing prices, and (iv) the share of unions not re-optimizing wages, under alternative levels of trend inflation.

Notes: For each level of trend inflation, we simulate 10,000 observations of the model while allowing only the indicated structural shock to vary. Shock innovations are drawn from a Gaussian distribution, with their scale recalibrated separately for each trend-inflation specification to reproduce the relevant volatility of the corresponding disturbance. Each dot represents one simulated period. The different series correspond to the alternative levels of trend inflation indicated in the figure. For brevity, we refer to the shares of non-reoptimized prices and wages as the *Calvo shares*.

E. State-dependent IRFs


Figure E.1 displays state-dependent impulse response functions (IRFs), computed by simulating a vector of labor supply shocks that generate inflation paths ranging from an initial annualized inflation rate of 0.5% to 5.5%. On top of each initial labor supply shock, we add either an additional $\varepsilon_{r,1} = 0.25\%$ monetary policy shock, as shown in Figure E.1a, or a $\varepsilon_{p,1} = -0.5\%$ cost-push shock, as shown in Figure E.1b. We report all the variables in deviation from the baseline underlying IRFs (i.e. minus the response to the initial labor supply shock).²² Darker lines indicate higher initial inflation levels, corresponding to more negative underlying cost-push shocks. In a quasi-linear model, such as the standard Calvo model, all shocks would generate quasi-identical dynamics regardless of the initial inflation level, and all lines would overlap. Consequently, there would be no state-dependency in the responses.²³ In this model, instead, we observe clear state-dependency. For monetary policy, the inflation response is stronger when initial inflation is high. The underlying mechanism can be explained as follows: at higher levels of inflation, the shares of firms and unions not optimally updating the price and wage are lower, meaning that prices and wages are more flexible. This increased flexibility steepens the Phillips curves, amplifying the effects of monetary policy shocks. Thus, the higher the initial level of inflation, the larger the variations in inflation and the smaller the ones in output induced by a given contractionary monetary policy shock.

Moreover, a second-order effect due to the forward-lookingness of firms and unions further affects this state-dependency. Firms and unions would anticipate the change in the frequency of price and wage resetting. In other words, they will anticipate if in the future the economy will be on a flatter or a steeper part of the Phillips curves. These effects tend to dampen the inflation response to contractionary monetary policy shocks, because the decline in inflation makes price and wage adjustments less appealing in expectation, and nominal adjustment more rigid downward. However, the opposite occurs for expansionary monetary policy shocks. Since firms (union) anticipate future increases in price (wage) resetting frequencies effectively, they would choose a higher reset price, amplifying the response of inflation. Hence, expansionary monetary policy when initial inflation is high can ignite strong inflationary dynamics.


Finally, the relative changes in the steepness of the NKPCs lead to pronounced differences in real wage dynamics. Under conditions of high inflation, both the wage and price NKPCs become substantially steeper. This steepening amplifies the responsiveness of real wages to demand shocks, resulting in stronger and more immediate adjustments. For example, when

²²We use labor supply and not a cost-push shock because on impact the former moves wage and price inflation in the same direction.

²³In the Calvo model, most of the difference will be generated by the nonlinearity in labor supply FOC and the effect of wages and prices dispersion on marginal cost.



(a) Responses to a +0.25% monetary policy shock



(b) Responses to -0.5% cost-push shock

Figure E.1: IRFs are state-dependent. The lines represent IRFs to a monetary policy (Panel (a)) and a cost-push (Panel (b)) shock, for different level of starting inflation. The darker the line, the higher the initial inflation level.

Notes: The initial inflation level are generated by a vector of labor supply shocks of different sizes.

inflation rises, firms and unions adjust nominal wages and prices more frequently and by larger magnitudes, facilitating rapid realignment of real wages to changing economic conditions. However, despite the magnitude of these adjustments, the increased flexibility in the economy – driven by the endogeneity of price and wage setting – ensures that these real wage deviations are less persistent. In other words, while the immediate response of real wages to shocks is stronger in a high-inflation environment, the economy's enhanced adaptability allows it to absorb and dissipate these deviations more quickly.

The mechanism underlying the cost-push shock is similar to that of monetary policy shocks: the economy becomes more flexible when inflation is high, causing supply shocks to have a larger impact both on inflation and on output. A subtle yet important observation is that while these shocks propagate more strongly, they do not exhibit greater persistence. This is because, in a high-inflation environment, shocks travel faster through the economy due to its increased flexibility.

Moreover, a cost-push shock induces an opposite response of price and wage inflation. Firms increase prices, triggering a contraction in demand that reduces both nominal wage inflation and the desired real wage. This reduction in wage inflation dampens the incentives for unions to reset wages, leading to fewer adjustments on the extensive margin of wage setting. Over time, however, this initial effect gives way to a standard price-to-wage feedback loop, where wages, as inflation remains elevated, become more flexible to catch up with the higher price levels. This adjustment reflects the endogenous response of wage-setting mechanisms to prolonged inflationary pressures, highlighting the interaction between the extensive and intensive margins of price and wage adjustments. It is worth noting that this dynamic interplay in our model makes the required adjustment to a negative supply shock, through lower both real wages and output, very persistent, no matter what the starting inflation level.

F. Real-time optimal inflation targeting




Figure F.1: **Recursive inflation targeting through monetary-policy shocks implies different policy prescriptions and macroeconomic effects.** Counterfactual paths are obtained by recursively selecting the contemporaneous monetary-policy shock that minimizes the central bank's conditional expected loss in the benchmark and fixed-updating-frequency models. Blue lines: benchmark model; red lines: fixed-updating-frequency model. Solid lines: baseline historical paths; dashed blue and dotted red lines: recursive inflation-targeting counterfactuals.

Notes: Counterfactual simulations for the benchmark model (blue lines) and the model with fixed adjustment frequencies (red lines) over 2021Q1–2024Q4. Starting in 2021Q1, the central bank recursively chooses the contemporaneous monetary-policy shock that minimizes its conditional expected loss. Future non-policy innovations are set to zero, while private agents expect future monetary policy to follow the estimated Taylor rule. The loss function is

$$\mathcal{L}_t = (100(\pi_t - \bar{\pi}))^2 + (100(R_t^4 - R_{t-1}^4))^2 + \beta \mathbb{E}_t \mathcal{L}_{t+1}.$$

In Figure F.1, we extend our analysis by considering a recursive inflation-targeting counterfactual implemented through monetary-policy shocks. Unlike the exercise in Section 5, which changes the coefficients of the Taylor rule, we hold the estimated rule fixed and optimize the sequence of deviations from that rule.

Our approach follows the policy-counterfactual logic of Sims and Zha (2006): starting in 2021Q1, at each date t we condition on the state of the economy inferred from the historical observations through that date and construct a model-consistent forecast, setting future non-policy innovations to their conditional means. The central bank then chooses the con-

temporaneous monetary-policy shock, and hence the current deviation of the policy rate from the estimated Taylor rule, that minimizes the expected present value of its loss. From period $t + 1$ onward, private agents expect monetary policy to follow the estimated Taylor rule; future policy deviations are neither announced nor incorporated into their expectations. In the following quarter, newly realized shocks are incorporated and the problem is solved again. The resulting counterfactual is therefore a sequence of period-by-period policy interventions, rather than a commitment to an announced future interest-rate path.

A key question is why we rely on an ad hoc loss function instead of microfounded conditional welfare measures. Given the current parameterization at the posteriors' modes, households exhibit an extreme aversion to inflation, which leads to counterintuitive results – agents would prefer an unrealistic output contraction rather than tolerate a small overshoot in inflation. This strong aversion likely stems from the nonlinearities in price and wage dispersion, combined with the quasi-linear disutility of labor in both models alongside very concave utility gain from consumption. Additionally, conditional welfare measures vary significantly across the two different parameterizations of the models, leading to preference structures that are highly model-dependent. Finally, in the absence of financial frictions, the central bank does not take into consideration the policy rate volatility. Therefore, using a consistent ad hoc inflation targeting loss function approach allows for greater comparability across exercises and mitigates model-specific biases in welfare measurement.

Therefore, we set up the Central Bank loss function in such way

$$\mathcal{L}_t = (100 \times (\pi_t - \bar{\pi}))^2 + (100 \times (R_t^4 - R_{t-1}^4))^2 + \beta \mathbb{E}_t \mathcal{L}_{t+1}$$

where the central bank minimizes the inflation deviation from the target and policy rate volatility as a proxy for financial frictions. To be clear, the purpose of this section is not to be a fully-fledged optimal policy exercise, which would also strongly depend on the implied output gap, but to assess the monetary policy stance and maximal output loss given strict inflation-targeting objective.

A final caution is warranted when comparing the outcomes of the real-time optimal policy exercises. While both the benchmark and the fixed updating frequency models align with the data in their respective baseline cases, they rely on different structural parameters and shock dynamics. Therefore, the objective of this section is to examine robust optimal policies given model uncertainty. Accordingly, we focus on highlighting the differences in optimal inflation targeting prescriptions across the two models, rather than evaluating the intrinsic performance of these policy recommendations.

Our results highlight key differences in policy responses under alternative price and wage-setting mechanisms. In the fixed-updating frequency model, inflation is initially perceived as slightly above the target and highly persistent, primarily driven by highly persistent cost-

push shocks. The stabilization trade-off is significant due to the flat Phillips curve, which limits the sensitivity of inflation to economic slack. As a result, in real-time, the optimal inflation targeting initially implies hard tightening (here depicted by the alternative evolution of the shadow rate in the bottom right panel). This leads to a sharp increase in the real interest rate, generating a substantial decrease in the output gap. Given the muted pass-through from economic activity to inflation, even this harsh tightening leads to a strong trade-off in economic stabilization and initially little effect on inflation. However, as the level and persistence of the policy stance increase, the increase in slack leads to strong disinflationary dynamics starting in 2023.

By contrast, in our benchmark model, inflation surprises are perceived as shorter-lived but occur repeatedly. This results in a less aggressive tightening path than in the fixed-updating Calvo model, and slightly more aggressive than, but fairly close to, the realized one. Moreover, the steeper Phillips curve enhances the transmission of monetary policy, leading to far lower inflation and a less pronounced output loss than in the fixed-updating frequency model.

G. Additional figures






Figure G.1: **Year-over-year inflation responses to a cost-push shock.** Panel (a): Local projection responses to oil shocks. Panel (b): Benchmark model responses to cost-push shocks. Horizon is quarterly.

Notes: Local projection is monthly frequency for year-on-year HICP inflation spanning 1997M1-2024M9. For each horizon $h = 0, 1, \dots, H$, we estimate $\pi_{t+h} = \alpha_h + \sum_{\ell=0}^{11} (g_h(z_{t-\ell}) + \delta_{\ell,h} \pi_{t-1-\ell}) + \varepsilon_{t+h}$, where π_t denotes year-on-year inflation and z_t is the monthly oil supply shock constructed by [Känzig \(2021\)](#). The function $g_h(\cdot)$ captures horizon-specific nonlinearities. In this specification, $g_h(z) = \beta_{1,h}z + \beta_{2,h}|z|$.




(a) Responses to a +0.25% monetary policy shock



(b) Responses to -0.5% cost-push shock


Notes: The parameter used are the posterior mode of the baseline model (see Table 1).

Figure G.2: IRFs with and without state-dependency in price and wage setting




Notes: This chart compares the share of re-optimized prices in the quantitative model with the share of price changes reported in [Gautier et al. \(2026\)](#) over their common sample (2012–2024). The correlation between the two series is 0.6146, with a 90% confidence interval of [0.4609, 0.7325]. The two objects are not perfectly comparable: in the model, non-reoptimized prices are indexed each period, so the model-based series moves even when firms do not re-optimize. Moreover, [Gautier et al. \(2026\)](#) cover only part of the HICP basket and, in particular, exclude energy items.


Figure G.3: Share of re-optimized prices in the quantitative model versus the share of changed prices in ([Gautier et al., 2026](#)).




(a) Inflation responses




(b) Output responses




(c) Nominal wage inflation responses



(d) Real wage responses



(e) Calvo share prices responses



(f) Calvo share wages responses

Figure G.4: IRFs to a + 0.5% labor-supply shock as function of the state. State-dependent IRFs are calculated by comparing the IRFs from all filtered shocks, with the IRFs from those filtered shocks plus an additional contemporaneous labor-supply shock. This isolates the effect of the labor-supply shock conditional on the initial state.







Figure G.5: **Phillips curves are curves.** The dots show the results of stochastic simulated runs of the output gap and of (i) price inflation, (ii) wage inflation, (iii) the share of not-adjusting firms; (iv) the share of not-adjusting unions, respectively, in the case of monetary policy shocks (Panel (a)) and labor supply shocks (Panel (b)).

Notes: For each panel, we simulate 10,000 observations of the model while allowing only the indicated structural shock to vary. Shock innovations are drawn from a Gaussian distribution, with their scale calibrated to reproduce the relevant volatility of the corresponding disturbance. Each dot represents one simulated period: the contemporaneous output gap is reported on the horizontal axis and the corresponding price inflation, wage inflation, price non-reoptimization share, or wage non-reoptimization share on the vertical axis. Blue circles refer to the benchmark model with state-dependent adjustment frequencies, while red crosses refer to the corresponding model with fixed adjustment frequencies. For brevity, we refer to the shares of non-reoptimized prices and wages as the *Calvo shares*.



(a) Benchmark



(b) Fixed updating frequencies

Notes: The black line reports the historical path generated when all shocks act jointly. Demand-side shocks collect the joint contribution of the monetary-policy and risk-premium shocks, while supply-side shocks collect the joint contribution of the labor-supply and cost-push shocks. The demand–supply interaction captures the non-additive effect of the two shock bundles. See also the notes to Figure 7.

Figure G.6: Historical decomposition of nominal wage inflation in both model specifications. Panel (a): benchmark model. Panel (b): model with fixed updating frequencies.