ECONSTOR
Make Your Publications Visible.

A Service of

ZBW Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Kilponen, Juha; Kortelainen, Mika

Working Paper

State-dependent pricing, de-anchored expectations, and the cost of looking through inflation shocks

Bank of Finland Research Discussion Papers, No. 11/2026

Provided in Cooperation with:
Bank of Finland, Helsinki

Suggested Citation: Kilponen, Juha; Kortelainen, Mika (2026) : State-dependent pricing, de-anchored expectations, and the cost of looking through inflation shocks, Bank of Finland Research Discussion Papers, No. 11/2026, Bank of Finland, Helsinki, https://nbn-resolving.de/urn:nbn:fi-fe20261002130478

This Version is available at:
https://hdl.handle.net/10419/344113

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

Bank of Finland Research Discussion Papers
11 • 2026

Juha Kilponen – Mika Kortelainen

State-dependent pricing, de-anchored expectations, and the cost of looking through inflation shocks

Bank of Finland
Research

Bank of Finland Research Discussion Papers
Editor-in-Chief Esa Jokivuolle

Bank of Finland Research Discussion Papers 11/2026
2 October 2026

Juha Kilponen – Mika Kortelainen:
State-dependent pricing, de-anchored expectations, and the cost of looking through inflation shocks

ISSN 1456-6184, online

Bank of Finland
Research Unit

PO Box 160
FIN-00101 Helsinki

Phone: +358 9 1831

Email: research@bof.fi
Website: www.suomenpankki.fi/en/research/research-unit/

The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the Bank of Finland or the Eurosystem.

State-Dependent Pricing, De-Anchored Expectations, and the Cost of Looking Through Inflation Shocks *

Juha Kilponen[1] and Mika Kortelainen[2]

[1]Bank of Finland
[2]Bank of Finland

September 17, 2026

Abstract

We extend an estimated euro-area Smets-Wouters (2007) model with endogenous price-adjustment hazards and a credibility-weighted expectations channel and study the interaction between state-dependent pricing (SDP) and inflation-expectation de-anchoring as complementary sources of second-round inflation dynamics. We establish four results. First, the two channels are super-additive: the long-run price-level cost of a large cost-push shock exceeds the sum of individual effects when both channels are active. Second, de-anchoring is material only when routed through the forward expectation term, not through backward indexation. Third, a formal Kalman-filter state-space estimation using ECB Survey of Professional Forecasters (SPF) and Consumer Expectations Survey (CES) data finds that professional forecasters remained fully anchored while households de-anchored significantly once inflation crossed approximately 5% during 2021-2024 inflation surge, consistently with rational inattention. Fourth, looking through a large cost-push shock is costly because the supply-side and expectations-side risks compound.

JEL codes: E31, E32, E52, E58. Keywords: state-dependent pricing, inflation expectations, de-anchoring, rational inattention, second-round effects, monetary policy.

*The views expressed in this paper are of the authors do not necessarily reflect the views of the Bank of Finland or the European System of Central Banks. We thank Gene Ambrocio and Markus Haavio for invaluable advice, Erwan Gautier for providing the aggregated PRISMA data for euro area, and Chiara Osbat for pointing us to relevant additional data sources.

1. Introduction

Large supply shocks - energy surges, tariffs, global bottlenecks - do more than shift relative prices. When sufficiently large or persistent, they push the economy past a threshold beyond which pricing, wage-setting, and inflation expectations formation adjust endogenously, amplifying the initial disturbance into second-round effects. The euro-area surge of 2021-2024 is the canonical recent example: inflation that began with energy and supply-chain disruptions broadened across sectors, grew more persistent, then unwound faster and at lower output cost than standard linear models predicted. Notably, this second-round amplification did not take the textbook form many policymakers initially feared: contemporaneous assessments found little evidence of a sustained, 1970s-style wage-price spiral taking hold in the euro area (Boissay et al., 2022, Battistini et al., 2022). The mechanism we document operates instead through firms' repricing behavior and the de-anchoring of household expectations, channels largely independent of wage bargaining.

The 2021- 23 euro-area inflation episode illustrates both channels vividly.

Figure 1: Price adjustment frequency and inflation expectations

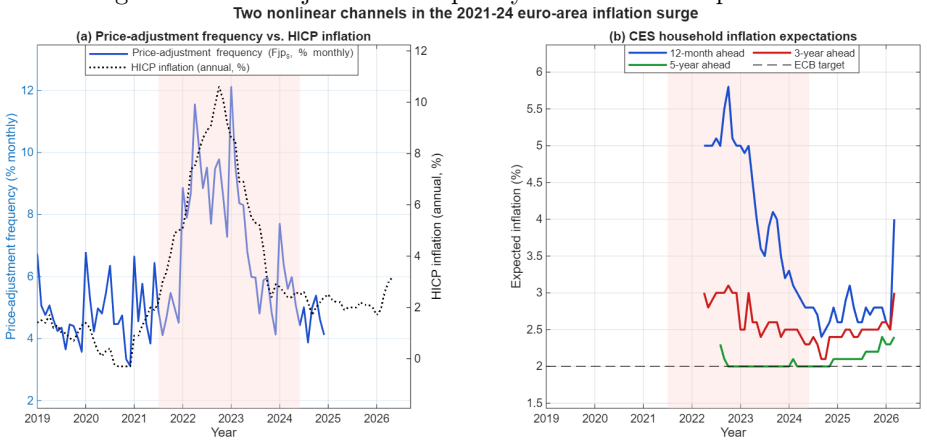

ECB/PRISMA micro-data show the monthly price-adjustment frequency rose from approximately 4.9% before the surge to a peak of 12.1% in January 2023 as headline inflation (HICP) crossed 5%. Simultaneously, the ECB Consumer Expectations Survey records household medium-term expectations rising from 2.5% to 5.1%. These are not independent phenomena. Our estimation shows that the PRISMA price-adjustment threshold (5.75%) and the CES "attention" threshold (5.15%) are statistically indistinguishable, suggesting that both nonlinear channels activate at the same inflation level.

To capture these channels in a general equilibrium setting, we extend an estimated Smets-Wouters (2007) model of the euro area (1999-2025) with (i) an endogenous

price-adjustment hazard rising with inflation following Dotsey et al. (1999) and Costain and Nakov (2011), and (ii) a credibility-weighted forward inflation expectation entering price-setting, following Alichi et al. (2009) and Carvalho et al. (2023). Both mechanisms are restricted to price-setting: wage growth does not Granger-cause inflation and shows no threshold-type nonlinearity that survives restriction to the surge episode (Appendix A.5.11), so wages follow standard Calvo pricing throughout. Our central finding and contribution is a modest super-additivity. When both SDP and de-anchoring are active, the long-run price-level response to supply shock exceeds the sum of the two individual responses by roughly one-twentieth in our setting and this specific episode. Two further results underpin and sharpen this finding. First, de-anchoring matters only via the forward-looking expectation term, not through backward indexation, and this is what makes the interaction with SDP's forward-looking pricing margin operative. Second, formal state-space estimation finds professionals maintained near-full anchoring. On the contrary, household anchoring fell sharply once inflation crossed approximately 5%, consistent with rational inattention where de-anchoring depends on the state of the economy.

A growing literature examines each channel in isolation. On the pricing side, Costain and Nakov (2011), Karadi and Reiff (2019), and Blanco et al. (2024) document nonlinear Phillips curves arising from state-dependent pricing. On the expectations side, Carvalho et al. (2023), Reis (2021), and Pfäuti (2024) model de-anchoring as a source of inflation persistence. Our contribution is to show that these two channels interact super-additively, if modestly: the whole exceeds the sum of the parts by approximately one-twentieth, because SDP steepens the Phillips curve precisely when de-anchoring shifts it outward. Neither channel alone produces the amplification observed in the data.

Our findings have practical implications for central bank's policy frameworks and optimal policy. The threshold coincidence — firms' repricing threshold (5.75%) and households' attention threshold (5.15%) activating at the same inflation level underlines the importance of monitoring current level of headline inflation. When headline inflation enters the 5–6% zone, the standard linear approaches will most likely underestimate the impact of initial shock to inflation. Our optimal policy simulations suggest that the gains from accounting for state-dependent pricing and de-anchoring depend on both the composition of the underlying shock and whether the policy response is credible and anticipated rather than a surprise. State-dependent pricing, which makes prices more flexible, improves the policy trade-off only when the inflation surge sufficiently combines demand-side pressure with the supply shock. A purely supply-driven episode leaves little additional scope for state-dependence to matter for policy. Across our optimal policy benchmarks, the welfare gains are concentrated almost entirely in regimes where the policy response is credible and anticipated rather than a surprise, indicating that the expectations channel — agents revising their inflation outlook because policy is known to be front-loaded — is the operative margin through which optimal policy delivers its benefits. The substantial welfare gains documented in Table A6.1 arise entirely from the commitment and systematic-rule-change

channels.[1]

Systematic rule changes (OSR, 20.1%) and full commitment (Ramsey, 26.2%) deliver substantial welfare gains over the estimated Taylor rule. Both are derived from linearized first-order conditions — a standard Dynare limitation on optimal-policy derivation — and the reported welfare gains are evaluated at first order for all four benchmarks to ensure an internally consistent comparison (see Table A6.1, note b). The welfare ordering confirms that under SDP and de-anchoring, the expectations channel is the key operative margin: policy gains arise from agents knowing that policy will be different, not from being surprised by it. Under the nonlinear Phillips curve, this preventive effect is especially powerful: anticipated front-loading keeps inflation below the threshold at which SDP and de-anchoring compound.

Rest of the paper is structured as follows: Section 2 summarizes euro-area empirical motivation. Section 3 presents the model and formal estimation. Section 4 analyzes the cost-push shock including the perturbation-order comparison and shock-size nonlinearity. Section 5 establishes super-additivity and placement. Section 6 evaluates optimal policy. Section 7 addresses identification and limitations. Section 8 concludes. Appendices cover model equations (A.1), Bayesian estimation (A.2), rational inattention (A.3), complete λ (anchoring of inflation expectations) estimation (A.4), joint GMM estimation of ϕ_h (price-hazard semi-elasticity with respect to inflation) and λ from PRISMA and CES data (A.5), and optimal policy derivations (A.6).

2. Two nonlinear channels in the 2021-24 Euro-Area inflation surge

2.1 State-Dependent Repricing

Euro-area CPI micro-data show a marked rise in repricing frequency during 2021–23. Excluding sales, the monthly frequency of *all* price changes increased from approximately 8.2% before the surge to 12.0% in 2022, peaking at 15.7% in January 2023 (Gautier et al., 2026). This aggregate rise was driven almost entirely by more frequent price *increases*: the ECB/PRISMA series for upward price adjustments (`ea_Fjp_s`) rose from a pre-surge mean of 4.9% per month to a peak of 12.1% in January 2023, while the frequency of price decreases barely changed. Mapped into quarterly hazards, the upward repricing frequency implies a rise from roughly 0.226 to 0.375 during the surge (Gautier et al., 2026). The spike in aggregate inflation therefore reflected primarily a surge in the *frequency* of price increases rather than their size — the signature of state-dependent pricing. Blanco et al. (2024) for UK data support a similar

[1] A linear policy diagnostic (OPP) finds a near-zero gradient under all three weight specifications (∼0.00% single-step improvement, ∼0.02–0.12% after 500 iterations, Table A6.3), confirming that the estimated Taylor rule is already locally optimal within the class of additive linear perturbations to the policy-shock sequence.

quarterly hazard semi-elasticity of approximately 0.25–0.38; a simple back-of-envelope calculation using Gautier et al.'s (2026) published frequency and regime statistics (Appendix A.4.10) corroborates this range. Consequently, the formal estimation of the hazard semi-elasticity ϕ_h in Appendix A.B uses upward price adjustments (`ea_Fjp_s`) rather than the total price-change frequency, since the model's logistic hazard governs the probability of an upward price reset in response to a cost-push shock.

Nonlinear VAR estimates show that a large rise in energy prices lifted core inflation approximately twice as much as an equivalent fall (Burriel et al. 2024). Under time dependent models a la Calvo this asymmetry cannot arise; under SDP it follows directly from the convex hazard function (see Section 3.2).

2.2 Expectation De-anchoring

Survey evidence documents another interesting fact. There is a sharp divergence between professional forecasters and household inflation expectations once euro-area inflation crossed approximately 5% in mid-2022. According to Consumer Expectations Survey a gap between 12-month and 3-year-ahead household expectations peaked at 2.70 percentage points in October 2022, then declined steadily toward zero as inflation receded. By contrast, according to Survey of Professional Forecasters gap between 1-year and 5-year-ahead professional forecasts remained below 0.5 percentage points throughout the entire episode.

This difference is consistent with state-dependent attention to inflation, a prediction shared by the rational-inattention literature (Sims, 2003) and by salience-based accounts of de-anchoring (Gennaioli et al. 2024): professionals track the rational expectation even during a surge; households allocate little attention when inflation is near the stable target but snap attention once it becomes sufficiently salient. The formal estimation in Section 3.3 confirms how these two features can be modeled and estimated jointly using familiar state-space methods.

3. Model and Estimation

3.1 Baseline: Estimated Smets-Wouters Model

The foundation is a work-horse model of Smets-Wouters (2007) estimated on quarterly euro-area data 1999Q1-2025Q4 (108 observations) via Bayesian Metropolis-Hastings. We use eight observables for estimation: real consumption, investment, GDP, and wage growth; GDP-deflator inflation; the policy rate (augmented with a shadow rate during the 2014Q3-2021Q1 ELB period following Kortela, 2016); hours worked; and the output gap. We estimate the model at first order with constant Calvo probabilities; SDP and de-anchoring is layered on for second-order simulations with pruning. Appendix A2 provides all the relevant details of the model and estimation procedure, including various diagnostic statistics.

3.2 State-Dependent Pricing

Following Dotsey et al. (1999) and Costain and Nakov (2011), we assume that the period-t price-adjustment hazard follows a logistic function of quarterly inflation: $h_t = 1 / [1 + exp(-(a^p + \phi_h * \pi_t))]$, where a^p is pinned so the steady-state hazard equals the Calvo probability ξ_p, and ϕ_h governs the semi-elasticity of hazard rate to inflation such that $\varepsilon_h = (1 - h_t) * \phi_h$. The same logistic structure applies to wages but with a much flatter calibration reflecting calendar-driven euro-area bargaining. The logistic specification is best understood as a reduced form approximation to the underlying menu-cost decision problem. In a fully micro-founded menu-cost economy, a firm i resets its price in period t if and only if its idiosyncratic price gap $|p_t(i) - p_t^*|$ exceeds a threshold determined by the fixed adjustment cost. The cross-sectional fraction of resetters then depends on the joint distribution of price gaps, which itself depends on the entire history of aggregate inflation. Solving this distributional fixed-point problem is computationally expensive and has been the subject of a large recent literature (see e.g. Blanco et al., 2024, and references therein). The logistic hazard collapses this distributional problem onto a single sufficient statistic — current quarterly inflation. We set baseline targets ε_h approx to 0.35 (robustness range 0.25-0.45), implying $\phi_h = 0.45$ given a steady state pricing re-adjustment probability of $\xi_p = 0.77$. Because the hazard is pinned at the Calvo steady-state value, the first-order log-linearization is observationally equivalent to standard Calvo model. Non-linearity enters only through second and higher order terms. Appendix 5 estimates ϕ_h formally from the ECB/PRISMA euro-area aggregate price-change statistics (2010–2024) via IV-GMM with Brent crude in EUR as instrument, yielding ϕ_h= 0.088 in monthly-frequency/annual-inflation units, equivalent to approximately 0.35 in the quarterly units of the model — confirming the baseline semi-elasticity target and falling inside the robustness range. As a complementary check on whether this calibration is empirically grounded, Appendix A.5 also tests directly whether firms' repricing behavior and households' attention to inflation switch on at the same point in the inflation distribution. The cross-equation threshold test in A.5 finds the price-adjustment threshold (5.75%) and the CES household attention threshold (5.15%) are statistically indistinguishable (gap = 0.60 pp, p $\approx$ 0.19), directly corroborating the rational inattention mechanism of A.3.

3.3 Formal Estimation of the Anchoring Weight λ

The anchoring weight λ governs the credibility-weighted forward expectation such that $\pi^{fwd} = (E_t\pi_{t+1})^{\lambda} * (\pi^e)^{(1-\lambda)}$, where π^e evolves via constant-gain learning. We estimate λ and constant gain parameter g jointly using a Kalman-filter state-space model. Long-horizon survey expectations ($y^{\{3,5\}}$) pin the perceived trend π^e, while λ is identified from excess movement in the short-horizon expectation (y^1) relative to the perceived trend such that $y_t^1 = \lambda * f_t^1 + (1 - \lambda) * \pi_t^e + \varepsilon_t^1$. Estimation is done by maximum likelihood with 30 random initial points.

Results are provided in Table 1. Professional forecasters maintained near-

full anchoring ($\lambda_{SPF} = 0.953$, SE = 0.000), while for households $\lambda_{CES} = 0.460$ (SE = 0.026) in the constant model. Applying the same time-varying specification to the SPF series instead finds no identifiable threshold at all over the observed inflation range (Appendix A.5.12), consistent with professional forecasters' anchoring being effectively constant throughout the surge. For households, we also estimate a time-varying specification in which the anchoring weight λ is allowed to shift with the inflation regime — switching from a "normal" value to a lower, more de-anchored value once inflation crosses an estimated threshold (full specification in Appendix A.4.4). A likelihood ratio (LR) test reveals that this time-varying model is strongly preferred over the constant-λ specification (LR = 13.51, $p = 0.004$). In the low inflation (normal) regime $\lambda_{hi} = 0.953$, while once inflation crosses approximately 5.15%, it falls sharply to $\lambda_{lo} = 0.445$. De-anchoring was strongest at the peak surge ($\lambda_{peak} = 0.369$ in April 2022-March 2023) and only partially reversed during disinflation ($\lambda_{disinfl} = 0.581$, Wald p = 0.006), confirming asymmetric re-anchoring during inflation surge and decline. Our baseline $\lambda = 0.85$ is in that sense conservative. At the estimated surge-regime anchoring weight ($\lambda_{lo} = 0.445$, Table 1), the super-additivity result in Section 5 rises from 4.8% to 15.5% of the additive benchmark. The sacrifice-ratio comparison is more sensitive still: at our baseline calibration the combined channel improves on the Calvo benchmark (5.65 vs. 5.75), whereas at λ_{lo} the ordering reverses (5.86 vs. 5.75), since the de-anchoring penalty grows much faster than SDP's efficiency gain as anchoring weakens. We adopt the more conservative $\lambda = 0.85$ as our headline calibration and treat the estimated-value result as motivating further robustness work rather than as a headline finding.

Table 1. Formal Estimation of the Anchoring Weight λ

Population	Model / Regime	lambda	SE	Interpretation
SPF professionals	Full sample	0.953	0.000	Fully anchored throughout
CES households	Constant λ (full)	0.460	0.026	Average de-anchoring
CES households	TV: normal regime	0.953	0.078	Indistinguishable from SPF
CES households	TV: surge regime	0.445	0.023	Significant de-anchoring
CES households	Peak surge (Apr 22-Mar 23)	0.369	0.014	Maximum de-anchoring
CES households	Disinflation (Apr 23+)	0.581	0.076	Partial re-anchoring

Notes: SPF estimated on quarterly data 1999Q1-2026Q2. CES estimated on monthly data April 2022-March 2026 using 12-month, 3-year, and 5-year ahead CES medians. SE = delta-method standard errors. LR test of TV vs. constant-λ: statistic 13.51, p = 0.004. Wald test of λ_{peak}not equal $\lambda_{disinfl}$: statistic 7.55, p = 0.006.

4. A Large Cost-Push Shock under State-Dependent Pricing

4.1 Calibration and Shock Size

In what follows, we discuss the results from the simulations, where the markup shock is calibrated to four times the standard deviation (0.035), as identified by the Kalman smoother for 2022-2023 period. This shock generates a peak quarterly inflation approximately 4 pp when SDP is active, consistent with the euro-area experience where quarterly GDP-deflator inflation peaked at roughly 2.5-3 pp during the surge. This peak inflation is well above the estimated household inflation attention threshold of approximately 5.15%, confirming the rational inattention mechanism discussed in A.3.

4.2 Baseline Comparison: SDP versus Calvo

Figure 2 reports impulse responses to 4σ markup shock under second-order pruned perturbation.[1] Relative to constant Calvo model, SDP generates a larger and earlier inflation peak, deeper but less persistent activity trough, sharper policy-rate response, and visible jumps in price and wage adjustment frequencies — consistent with the euro-area micro-data patterns in Section 2. Despite these timing differences, the cumulative inflation–output trade-off is almost unchanged for a pure cost-push shock: SDP steepens the effective Phillips curve but also steepens the initial contraction in real activity.

4.3 Non-linearity is a Large-Shock Phenomenon

Because adjustment hazards respond to the level of inflation, the SDP-Calvo gap grows non-linearly with shock size. A central bank using a Calvo model sees the same normalized dynamics regardless of how large the cost-push disturbance is, and the model does not provide a signal that a larger shock is qualitatively different and requires different policy responses. Under SDP this is no longer true. Figure 3 makes this directly visible via size-normalized impulse responses.

Figure 3 reports size-normalized inflation IRFs (each divided by its shock-size multiple) for multiples 1x, 2x, 3x, 5x, and 8x. Under Calvo model (left panel), all five normalized responses collapse onto a single curve representing an exact linear scaling. Under SDP (right panel), normalized responses fan out progressively: the 8x inflation response is roughly twice the 1x response per unit of shock.

Figure 4 translates this non-linearity into a summary statistic: peak quarterly inflation against shock size (left) and the SDP-Calvo peak gap (right). The gap is near-zero at 1-2 σ shock (approx 0.1-0.3 pp), confirming SDP and Calvo are observationally almost/practically indistinguishable. It then accelerates to 0.6 pp at 3σ, 2.6 pp at 5σ and 4.0 pp at 8σ sigma. The 4σ baseline produces a gap of approximately 1.5-2 pp. The 2022-23 markup shock cluster at approximately

Figure 2: State-dependent pricing versus Calvo: impulse responses to a large cost-push shock

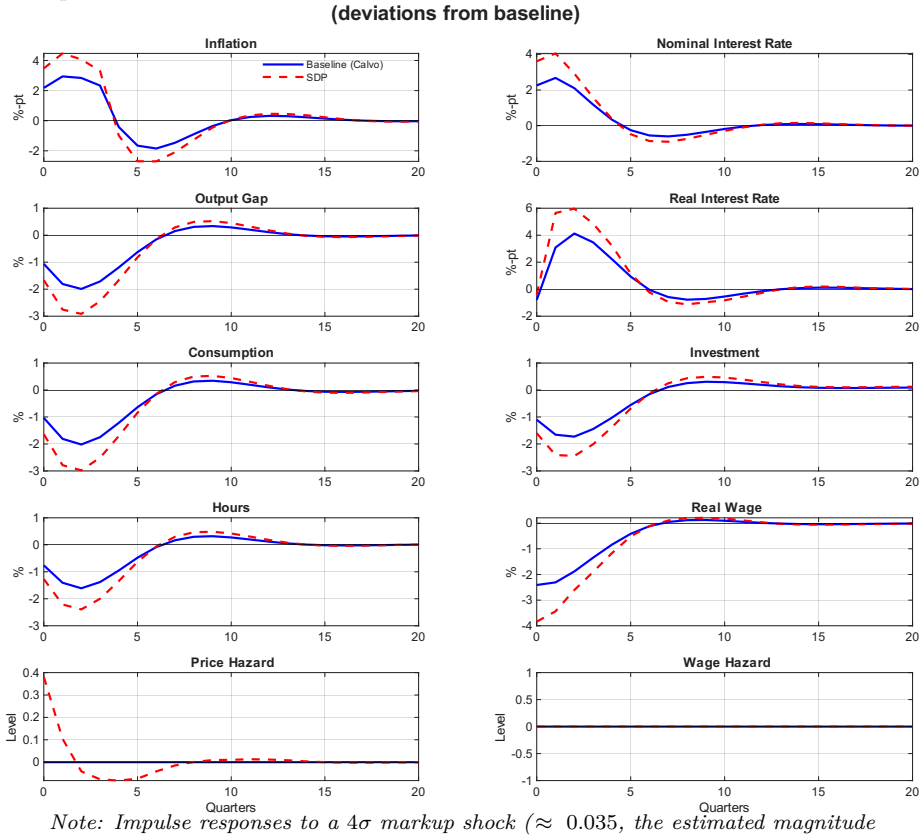

Note: Impulse responses to a 4σ markup shock ($\approx$ 0.035, the estimated magnitude of the 2022–23 euro-area innovation cluster), second-order pruned perturbation. Blue solid: constant-Calvo benchmark; red dashed: SDP model. All panels show deviations from the no-shock baseline in pp or percent. SDP generates a larger and earlier inflation peak ($\approx$ 4 pp quarterly vs. $\approx$ 2.5 pp Calvo) and deeper but shorter-lived output trough. Price hazard jumps from steady-state 0.23 to $\approx$ 0.55 at impact (empirical target: 0.226 $\rightarrow$ 0.375 during the surge, Gautier et al., 2026). Wage hazard response is much smaller ($\approx$ 0.09). Cumulative inflation–output trade-off is nearly identical across models.

Figure 3: Size-normalized inflation responses: nonlinearity grows with shock size

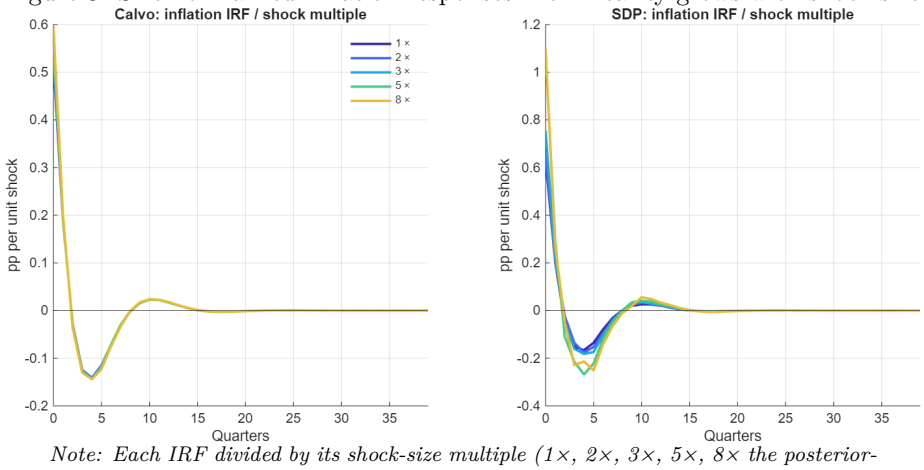

Note: Each IRF divided by its shock-size multiple (1×, 2×, 3×, 5×, 8× the posterior-mean standard deviation $\sigma(\eta^p) \approx 0.0087$), so that a linear model yields a single collapsed curve. Under Calvo (left panel), all five normalized responses overlap exactly. Under SDP (right panel), responses fan out progressively: the 8× normalized peak is roughly twice the 1× value per unit of shock. The fanning directly visualizes the state-dependent nonlinearity. Second-order pruned perturbation.

4-5σ sits in the accelerating section where SDP diverges sharply from Calvo - precisely the regime where policy design matters most.

5. Expectations De-anchoring and Super-Additivity

5.1 Four-Way Comparison

Figure 5 reports impulse responses to the 4σ markup shock disentangling the importance of two mechanisms: Calvo benchmark, SDP only, de-anchoring only (λ= 0.85), and both together. Both mechanisms enter through price-setting only; wages follow standard Calvo pricing with fully rational forward expectations throughout, consistent with the wage-side robustness checks in Appendix A.5.11, which find no Granger-causality running from wage growth to inflation and no threshold-type nonlinearity that survives restriction to the 2021–23 surge window.

First, de-anchoring alone already generates second-round effects: it lifts the long-run price level from about 1.24 pp (Calvo) to about 1.31 pp, a permanent increment of roughly 0.07 pp generated purely through the expectations channel. Second, the two channels are mutually reinforcing, though only modestly so: SDP alone raises the long-run price level to about 2.60 pp; with both active it settles at about 2.73 pp. Individual increments sum to approximately 1.42 pp, while the joint effect is approximately 1.49 pp—an excess of roughly 0.07 pp, or

Figure 4: Peak inflation and the SDP–Calvo gap as functions of shock size

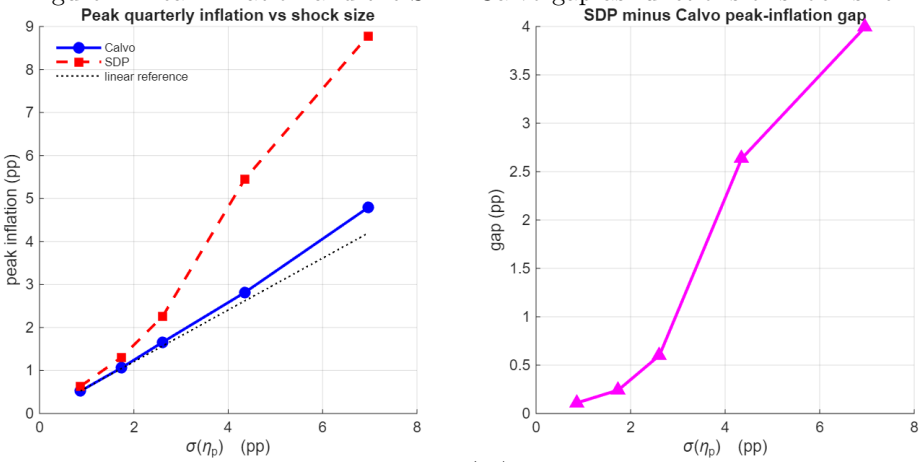

Note: Left panel: peak quarterly inflation (pp) against markup-shock size for Calvo (blue circles), SDP (red squares), and a linear reference (black dotted). Calvo scales linearly; SDP curves convexly above it. Right panel: SDP minus Calvo peak-inflation gap (pp, magenta triangles). The gap is near-zero for small shocks ($\approx$ 0.1 pp at 1σ, $\approx$ 0.3 pp at 2σ) and accelerates: 0.6 pp at 3σ, 2.6 pp at 5σ, 4.0 pp at 8σ. The baseline 4σ benchmark produces a gap of $\approx$ 1.5–2 pp. The 2022–23 markup shock cluster ($\approx$ 4–5σ) sits in the accelerating section. Second-order pruned perturbation.

about 5%, over the additive benchmark. The reason is that SDP-driven realized inflation feeds the learning rule more strongly, the perceived trend de-anchors further, and that drifted belief passes back through the forward expectation term into pricing.

The sacrifice ratio does not simply rise when both channels are active, as one might expect from the price-level result. SDP alone *lowers* the long-run sacrifice ratio, from 5.75 (Calvo) to 5.58, consistent with the classic state-dependent-pricing selection effect: faster endogenous price adjustment when inflation is high reduces the output cost of a given disinflation. De-anchoring alone raises the ratio, to 5.88. With both channels active the ratio is 5.65—higher than under SDP alone, but still below the Calvo benchmark, since SDP's efficiency gain dominates the de-anchoring penalty on net. This is not a reversal of SDP's mechanism, which continues to lower the sacrifice ratio by a roughly constant margin regardless of anchoring; rather, it reflects the imperfect-credibility logic underlying the de-anchoring mechanism itself (Alichi et al., 2009): the more expectations are governed by a backward-looking, adaptive process rather than credible forward-looking belief, the more the central bank must engineer real economic pain, rather than manage beliefs, to bring inflation down. Managing expectations is costly precisely when they are least manageable — and the finding here is that this cost grows fast enough with de-anchoring to swamp SDP's fixed efficiency gain once anchoring weakens to the estimated surge-regime level. The two mechanisms thus have genuinely distinct, partially offsetting

real-economy footprints: SDP scales up the size of the episode while improving the inflation-output tradeoff; de-anchoring adds persistence while worsening it.

Figure 5: Super-additivity of SDP and de-anchoring: four-way mechanism comparison

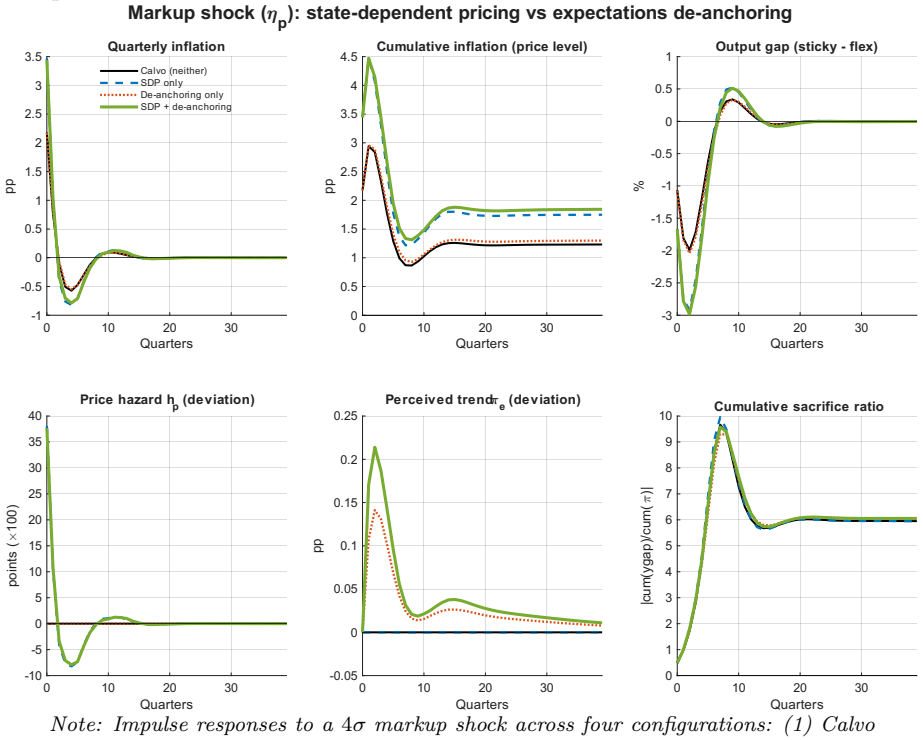

Note: Impulse responses to a 4σ markup shock across four configurations: (1) Calvo benchmark; (2) SDP only ($\lambda = 1$); (3) de-anchoring only (constant Calvo, $\lambda = 0.85$); (4) both channels active. Both mechanisms are price-only; wages are Calvo throughout (Appendix TA.B). Long-run price-level increments over Calvo: de-anchoring only $\approx +0.07$ pp; SDP only $\approx +1.35$ pp; both $\approx +1.49$ pp. Sum of individual increments ≈ 1.42 pp; joint increment exceeds the additive benchmark by ≈ 0.07 pp (about 5%). Long-run sacrifice ratio: 5.75 (Calvo), 5.58 (SDP only), 5.88 (de-anchoring only), 5.65 (both)—SDP lowers the ratio, de-anchoring raises it, and the combined effect remains below the Calvo benchmark. Third-order pruned perturbation, $\lambda = 0.85$.

5.2 Robustness to the Anchoring Weight

Figure 6 varies λ from 1.0 (Calvo benchmark, since $\lambda = 1$ collapses the forward blend to the fully rational rate) to 0.25, for the de-anchoring-only configuration (SDP off). The price-level effect grows smoothly and monotonically as λ falls across the full grid tested, $\lambda \in \{1.00, 0.90, 0.75, 0.50, 0.25\}$: long-run cumulative inflation rises from 1.24 pp at $\lambda = 1.00$ to 1.29, 1.35, 1.46, and 1.56 pp at $\lambda = 0.90$, 0.75, 0.50, and 0.25 respectively. The output-gap trough and the long-run sacrifice ratio move in step, deepening from -1.88% and 5.75 at $\lambda = 1.00$ to -2.10% and 6.34 at $\lambda = 0.25$. Perceived inflation trend paths are nearly

identical across λ values, because λ governs how much the drifted belief weights pricing, not how far beliefs drift.

Figure 6: Robustness of de-anchoring to the anchoring weight and hazard semi-elasticity

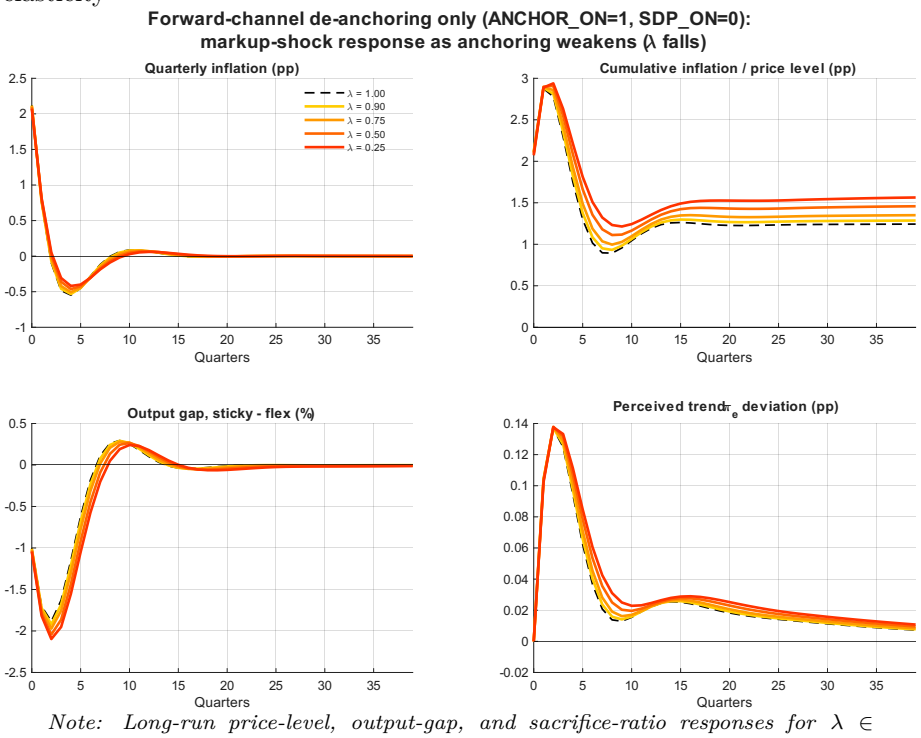

Note: Long-run price-level, output-gap, and sacrifice-ratio responses for $\lambda \in \{1.00, 0.90, 0.75, 0.50, 0.25\}$, *de-anchoring-only configuration (SDP off).* $\lambda = 1.00$ *reproduces the Calvo benchmark. All three metrics move monotonically as* λ *falls. Each* λ *point was solved as an independent run (not a parameter sweep within a single session) to rule out cross-contamination between grid points. Third-order pruned perturbation throughout.*

5.3 Placement of De-anchoring

De-anchoring could in principle enter price-setting through either of two channels: the forward-looking expectation term used in the reset-price decision, or the backward-looking indexation rule applied to prices that are not re-optimized in a given period. Non-reoptimizing firms index according to

$$P_t(i) = P_{t-1}(i) \cdot \pi_{ss}^{1-\iota_p} \, \pi_{t-1}^{\iota_p}, \tag{1}$$

where $\iota_p \in [0, 1]$ is the indexation weight and π_{t-1} is *observed, realized* lagged inflation — not an expectation, and therefore not a quantity that can de-anchor. The forward-looking channel enters instead through the reset-price recursions

of A.1.3, where the model-consistent expectation $E_t\pi_{t+1}$ is replaced by the credibility-weighted blend

$$\pi_t^{fwd} = (E_t\pi_{t+1})^{\lambda}(\pi_t^e)^{1-\lambda}, \tag{2}$$

with π_t^e the perceived inflation trend from constant-gain learning (Section 3.3). Only π_t^{fwd} is exposed to de-anchoring; the indexation rule above is left unmodified.

We tested an alternative formulation in which the perceived trend π_t^e instead entered through the indexation term — replacing π_{ss} in the equation above with a de-anchored analogue — and found this produced negligible effects at any learning gain. The reason is quantitative rather than conceptual: at the estimated price-indexation weight $\iota_p = 0.85$, only the raw indexation share re-indexes to *any* trend, anchored or not, while re-optimized prices — which carry first-order leverage on the Phillips curve through π_t^{fwd} — are unaffected by indexation altogether. De-anchoring therefore matters only insofar as it reaches price-setters through the forward expectation term, and we restrict the mechanism to that channel throughout the paper.

5.4 Rational Inattention Micro-Foundation

Appendix A.3 replaces the constant gain with a state-dependent gain rising logistically once inflation crosses the attention threshold $\pi^* = 5.15\%$ (Table 1): $g_{ri} = g_{low} + (g_{high} - g_{low}) \;/\; [1 + exp(-\kappa * (x - \pi^*))]$. This generates a mechanism where households update beliefs less than the constant-gain benchmark below the threshold (g_{low}= 0.01 vs $g_{learn} = 0.05$), explaining why their expectations were slow to de-anchor early in the surge. The gain structure also predicts asymmetric re-anchoring — fast de-anchoring when the threshold is crossed, slow re-anchoring when inflation recedes — matching the λ_{peak} vs $\lambda_{disinfl}$ asymmetry documented in Table 1. Strikingly, the PRISMA estimation in A.5 finds that firms' price-adjustment threshold (5.75%) and households' attention threshold (5.15%) are statistically indistinguishable, confirming that both nonlinear channels activate at the same inflation level. A.3 documents both findings with shock sizes up to 3σ.

6. Optimal Monetary Policy

6.1 Why Policy Matters More under Nonlinearity

Section 4 established that SDP leaves the cumulative inflation-output trade-off essentially unchanged for a pure cost-push shock under an estimated fixed monetary policy rule. This section shows that policy optimization leads to significant gains. Gains from optimal policy come from two sources. The first is commitment, or history-dependence. A planner who promises to undo the price-level effect of the shock can reallocate inflation and output over time and

stabilize the model economy at lower cost than a backward-looking inertial rule, independently of SDP. The second is specific to SDP. By steepening the effective Phillips curve in high-inflation states, SDP raises the potency of any given tightening, enlarging the gap between a well-designed rule and the estimated one. Figure 4 shows why this potency is largest precisely when inflation is highest.

6.2 Policy Benchmarks

We evaluate optimal policies under three commonly used benchmarks: Quadratic Ramsey, optimal simple rules (OSR), and the Erceg et al. (2000) welfare loss, supplemented by the model-light Optimal Policy Perturbation (OPP) of Barnichon and Mesters (2023). We evaluated them under three alternative loss-weight specifications (BM unit weights, Ramsey weights, and EHL welfare-derived weights). The Ramsey planner optimizes a second-order approximation to the welfare loss.

All four benchmarks point to similar optimal policy patterns of earlier tightening and a smoother rate path relative to the estimated rule. The inertial estimated rule delays the response to the cost-push shock, allowing inflation and markups to propagate before monetary policy tightens more aggressively. Optimal commitment pre-empts this propagation. The lower peak and smoother path of the optimal nominal rate in Figure 7a should not be read as a more accommodative stance: because the tightening is anticipated and credible, it changes agents' expectations before it changes rates, so less subsequent rate movement is needed to achieve the same disinflation — the mechanism is earlier and smaller, not weaker. Importantly, the OPP confirms that the estimated Taylor rule is already locally optimal within the class of additive linear perturbations: the gradient is negligible under all three weight specifications ($\sim$0.00% single-step improvement), and even after 500 iterations the improvement remains no larger than 0.12% (Table A6.3). The substantial welfare gains documented in Table A6.1 (Ramsey: 26.2%; OSR: 20.1%) therefore arise entirely from the commitment and systematic-rule-change channels — agents revising their inflation outlook because the *rule itself* is known to be front-loaded — rather than from any local sub-optimality of the estimated rule's shock-by-shock response.

Figure 7b complements the impulse responses above by showing the cumulative price level implied by each policy benchmark, illustrating the difference between commitment versus non-commitment distinction. Estimated Taylor rule settles at $\approx$ P = 1.020 (permanent +2 pp price-level increment). OSR converges to the same long-run level ($\approx$ 1.020) and hence does not deliver price-level stationarity. Quadratic Ramsey returns the price level to $\approx$ 0.996, reflecting history-dependent commitment to undo the price-level effect of the shock, while EHL achieves the most aggressive correction ($\approx$ 0.996–0.998). The commitment benchmarks reduce the permanent price-level cost by $\approx$ 2 pp relative to the estimated rule. These price-level paths are generated under the order = 2 pruned approximation.[2]

[2]See Table A6.1, note b for a caveat on the numerical reliability of the Ramsey-augmented

Figure 7a: Optimal monetary policy: four-benchmark comparison

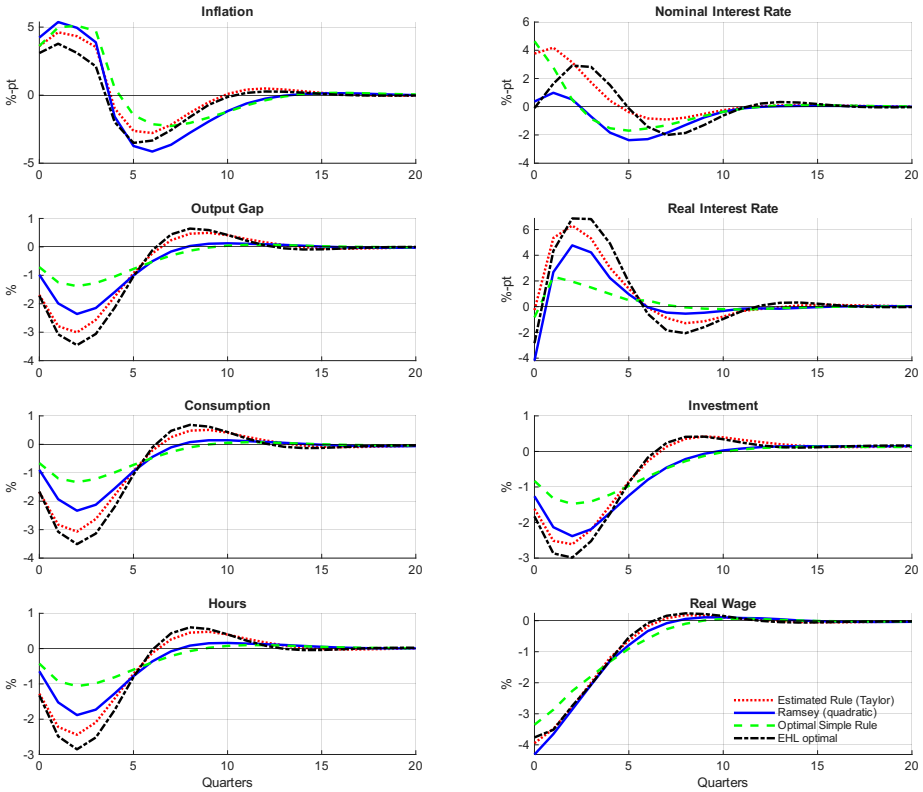

Note: Impulse responses to a 4σ markup shock under four optimal policy benchmarks. Red dotted line shows estimated Taylor rule (inertial, delays tightening), blue solid line is quadratic Ramsey under commitment), green dashed line is optimal simple rule (OSR) over grid $r_\pi \in [1.1, 4.5]$, $\rho_R \in [0.30, 0.97]$), Black dash-dot shows Erceg–Henderson–Levin (EHL) welfare-optimal policy path. In all simulations, $\lambda = 0.75$.

Figure 7b: Cumulative price level under alternative policy benchmarks

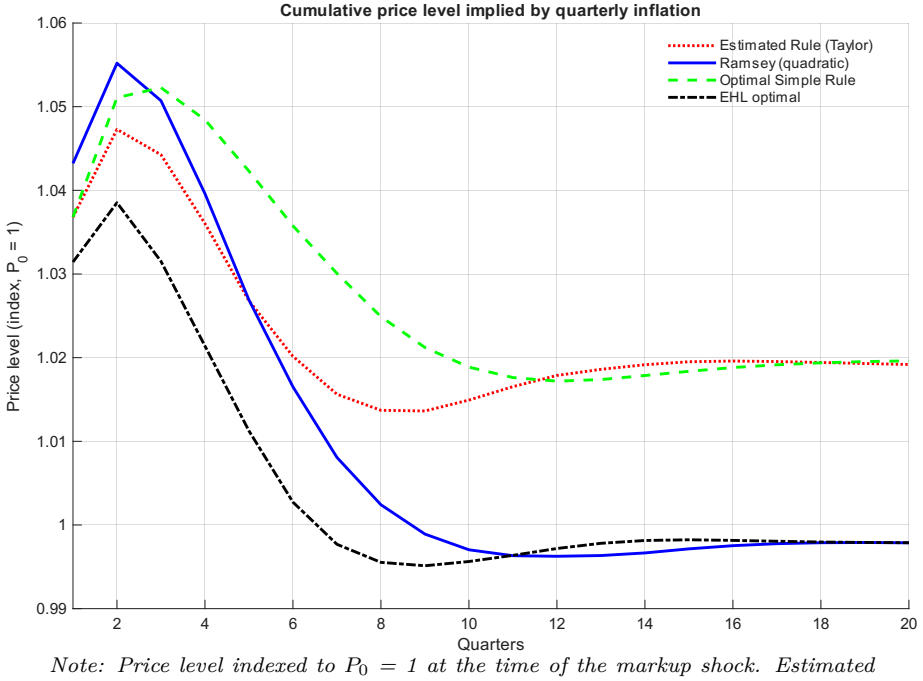

Note: Price level indexed to $P_0 = 1$ at the time of the markup shock. Estimated Taylor rule settles at $\approx P = 1.020$ (permanent +2 pp price-level increment). OSR converges to the same long-run level (≈ 1.020)—it does not deliver price-level stationarity. Quadratic Ramsey and EHL return the price level to ≈ 0.998, reflecting history-dependent commitment to undo the price-level effect of the shock). The commitment benchmarks reduce the permanent price-level cost by $\approx$ 2 pp relative to the estimated rule. Second-order pruned perturbation, 4σ markup shock, $\lambda = 0.75$.

The smoother optimal paths should not be read as a relaxed stance. They reflect a history-dependent commitment to undo the price-level effect of the shock, internalized by forward-looking agents. The central message is that looking through a large supply shock is costly precisely because it risks activating the SDP and de-anchoring channels simultaneously, and those channels compound rather than add to each other.

7. Discussion on identification and limitations

In this section we discuss possible limitations of our approach related to estimation, calibration and data. The core structural parameters are estimated from the Bayesian likelihood via four chains of 400,000 Metropolis-Hastings draws see A.2). The anchoring weight λ and learning gain g are estimated via Kalman-filter maximum likelihood from SPF and CES data (see section 3.3 and A.4).

The price-hazard semi-elasticity ϕ_h= 0.45 is calibrated rather than estimated from the aggregate likelihood. Full nonlinear estimation was attempted (auxiliary particle filter, Gaussian mixture filter, conditional particle filter) but we encountered prohibitive weight degeneracy on COVID-period observations at all feasible particle counts which is typical challenge in medium-scale DSGE models, documented in A.2.7. The perturbation order choice (A.2.9) reflects the same local-approximation constraint. However, all qualitative results are robust across the full micro-data-supported range 0.25-0.45. Formal estimation from the ECB/PRISMA euro-area aggregate price-change statistics (Appendix B) yields a quarterly-equivalent ϕ_h of approximately 0.35, inside the robustness range and consistent with the baseline semi-elasticity target. The PRISMA estimate is, if anything, closer to the lower boundary of the range, confirming that the paper's SDP amplification results are conservative rather than overstated.

A natural concern is whether this exercise is circular: ϕ_h is estimated from a sample (2010–2024) that includes the very episode it is used to explain. Appendix B.8 addresses this directly with an out-of-sample design. We re-estimate ϕ_h using only pre-surge data (2010M1–2021M6). It turns out that it is descriptively consistent with the full-sample estimate. The resulting model explains 77% of the out-of-sample variance in surge-period repricing frequency, far outperforming both a Calvo benchmark and a simple AR(1) persistence model. The one point where this pre-surge-only model falls short is the single month of peak realized inflation (January 2023), but even there roughly half of the apparent miss reflects a calendar effect which refers to well-documented concentration of annual price-list and contract resets in January.

When it comes to data, firm-level PRISMA microdata for the 2021–24 surge period is restricted to the ChaMP research network and was not accessible for this paper. We therefore identify ϕ_h from the published PRISMA aggregate series (Appendix B) and corroborate it with two independent reduced-form checks

solution at this order.

on euro-area aggregate data — inflation swaps and a single-equation logistic smooth-transition regression (A.5.6, A.5.7) — together with the out-of-sample test above (A.5.8), rather than relying on a single microdata-based identification strategy.

As an independent, non-survey corroboration of the threshold-coincidence result above, we examined whether EUR HICP zero-coupon inflation swap rates (Bloomberg, 2005–2026, tenors 1Y–10Y) show a comparable nonlinear de-anchoring response to realized inflation. The exercise was inconclusive: the long-run anchor (5-year-5-year forward) shows no reliably identified threshold at all, while a medium-tenor (1-year-2-year forward) measure does show a well-identified curve in the full sample, with an implied threshold near 1.8%. However, this estimate shifts to roughly 3%, with sharply worse precision, once the 2014–21 below-target "lowflation" period is excluded, indicating it largely reflects cross-regime co-movement rather than a stable within-regime threshold. We attribute this to a data-richness constraint — roughly two decades of quarterly swap data span only two or three independent macro episodes — rather than a substantive disagreement with the PRISMA/CES estimates, and draw no threshold estimate from it (Appendix 5.6).

A second, complementary check points to the same conclusion. Using a single-equation logistic smooth-transition regression (LSTR) on euro-area aggregate data, where core inflation responds to energy-price pass-through and the transition is triggered by lagged headline inflation. It turns out that Luukkonen-Saikkonen-Teräsvirta linearity test strongly rejects a linear specification. However, the transition location itself is not well pinned down by the full sample. The point estimate converges to the edge of the admissible search range, pointing to a weak identification in smooth-transition models when the regime-switching episode is short (six of 107 quarters above 5% inflation). A block bootstrap places the location with substantial uncertainty instead (90% CI, [0.3%, 9.4%]) — comfortably containing both the PRISMA and CES thresholds without identifying specifically neither. As with the swap-rate exercise above, we read this as compatible with, rather than confirming, the threshold-coincidence result, and draw no separate threshold estimate from it (Appendix 5.7).

Finally, we abstract from production-network propagation, financial-structure filters (loan-rate fixation, NBFI financing), and fiscal policy are documented in the empirical record as quantitatively important and are left to future work. Current results characterize what the aggregate pricing-and-expectations block alone implies, providing a clean lower bound on the full set of second-round amplification mechanisms.

A further limitation concerns the empirical counterpart to the model's expectations mechanism. The credibility-weighted forward term in Section 3.3 enters price-setting through firms' expectations, but the anchoring estimates in Section 3.3 and Appendix A.4 are drawn from professional forecasters (SPF) and households (CES) — neither of which is a direct measure of firm expectations. This is a genuine gap rather than an innocuous simplification: existing survey

evidence (Savignac et al., 2024) documents that firms' expectations follow their own dynamics, distinct from both professionals and households. Unfortunately, a direct, contemporaneous test is not possible for the episode this paper studies: the ECB's Survey on the Access to Finance of Enterprises (SAFE) began collecting firm-level inflation expectations only in June 2023 (Ferrando et al., 2025), after the 2021–23 surge and the de-anchoring episode documented in Section 3.3 had already largely unwound. No quantitative firm-expectations series — to our knowledge — covers the window in which the paper's threshold-coincidence and de-anchoring results are identified. We therefore use SPF and CES as bracketing proxies rather than as direct measurements: firms plausibly have better information access than households, arguing for anchoring closer to the professional estimate, but remain boundedly rational agents facing attention costs of their own, arguing against assuming full SPF-level anchoring. The paper's baseline calibration ($\lambda = 0.85$) sits close to the professional end of this range and, as discussed in Section 3.3, is conservative relative to the household-implied surge-regime estimate ($\lambda_{lo} = 0.445$). Firm-level SAFE data, now available from mid-2023 onward, offers a natural avenue for validating this bracketing approach in future work, once a comparably large post-surge sample accumulates.

8. Conclusion

This paper establishes that state-dependent pricing and inflation-expectation de-anchoring reinforce one another, if only modestly during the 2021-24 episode: their combined long-run price-level cost exceeds the sum of individual effects by roughly one-twentieth, a modest but robust super-additivity result with a transparent mechanism: SDP-driven inflation feeds the household belief-updating rule more strongly, drifted beliefs re-enter via the forward expectation term, and the two margins amplify one another. A formal state-space estimation confirms the empirical relevance. Professionals remained anchored while household anchoring fell sharply once inflation crossed approximately 5%, consistent with rational inattention.

The shock-size variations in Section 4.3 provides direct empirical grounding for why 2021-23 episode was different. The SDP-Calvo peak inflation gap is near-zero below 2σ sized markup shock and accelerates sharply above 3σ, reaching 4 pp when shocks is at 8σ. The 2022-23 markup shock cluster at approximately 4-5σ sits squarely in the accelerating regime. A commonly used Calvo model does not capture this non-linearity and the SDP model illustrates how the episode was predictably different, given only the shock's size.

Three methodological findings have broader relevance. First, de-anchoring is material only via the forward expectation term - models locating credibility risk in indexation will systematically understate de-anchoring risk. Second, the PRISMA price-adjustment threshold (5.75%) and the CES household attention threshold (5.15%) are statistically indistinguishable, meaning both nonlinear channels activate around the same inflation level, providing the empirical foundation of

the super-additivity result. Third, the welfare gains documented in Section 6 come from three distinct sources. The OPP diagnostic's nonlinear re-solve confirms that shock-by-shock deviations from the estimated rule offer essentially no scope for improvement

This suggests that a policymaker relying on this local, linear diagnostic alone would correctly infer that the estimated rule is already well-designed and locally optimal reaction to shock-by-shock. OSR's gain (20.1%) is a purely linear-quadratic result, provably invariant to whether SDP or de-anchoring is present (Table A6.1), reflecting the standard value of better-tuned rule coefficients rather than any nonlinear channel. Ramsey's gain (26.2%) is larger still: its optimal rule is derived from linearized first-order conditions, reflecting the the SDP channel. This favors early, decisive, and well-communicated tightening over wait-and-see gradualism when inflation approaches the threshold zone, on grounds that combine standard commitment-value logic (OSR, Ramsey) with the paper's nonlinear threshold-avoidance result (OPP) — consistent with the ECB's own review of its 2021–22 response.[3]

For monetary policy an environment with more frequent repricing is also one where a given inflation surprise de-anchors beliefs more. The two risks compound above a shock-size threshold. Early, credible policy prevents the economy from entering the self-reinforcing regime. The smoother optimal rate paths reflect commitment to undo the price-level effect of the shock, and not accommodation.

For policymakers operating within the ECB's three-pillar analytical framework, the results offer three operational lessons, each subject to the model uncertainty inherent in any single structural analysis. First, when headline inflation enters the high regime (5–6% range in our case) where both the repricing and attention thresholds are estimated to activate, standard linear Phillips -curve models will systematically underestimate inflation dynamics (persistence or volatility), and the analytical framework should shift accordingly. Second, the headline-versus-core decomposition is informative about whether SDP improves the policy trade-off — it does so only when a demand component is present alongside the supply shock (A.6.6). Third, the welfare gains documented in Sections 5 and 6 arise from agents believing that policy will be systematically front-loaded, not from surprise hawkishness. A linear, model-light policy diagnostic (OPP) finds negligible scope for improving on the estimated rule through shock-by-shock deviations alone ($\sim 0.00\%$ single-step; up to 0.12% after 500 iterations, Table A6.3) — confirming that the substantial gains documented above operate through the systematic-rule and commitment channels rather than through better shock responses. Systematic rule changes (OSR, 20.1%) and full commitment (Ramsey, 26.2%) deliver genuine welfare gains because they operate through the expectations channel: anticipated front-loading keeps inflation below the threshold where SDP and de-anchoring compound. This favors early, decisive, and well-communicated tightening over wait-and-see gradualism when inflation

[3]The reported welfare figure is evaluated at second order and so may partially, though likely incompletely, reflect the SDP channel.

approaches the threshold zone — consistent with the ECB's own review of its 2021–22 response.

References

Alichi, A., Chen, H., Clinton, K., Freedman, C., Johnson, M., Kamenik, O., Kişinbay, T., and D. Laxton et al. (2009). Inflation Targeting under Imperfect Policy Credibility. IMF Working Paper 09/94.

Andreasen, M. M., J. Fernandez-Villaverde, and J. F. Rubio-Ramirez (2018). The Pruned State-Space System for Non-Linear DSGE Models. *Review of Economic Studies 85, 1-49.*

Barnichon, R. and G. Mesters (2023). A Sufficient Statistics Approach for Macro Policy Evaluation. *American Economic Review 113(11), 2852-2898.*

Battistini, N., H. Grapow, E. Hahn, and M. Soudan (2022). Wage Share Dynamics and Second-Round Effects on Inflation after Energy Price Surges in the 1970s and Today. ECB Economic Bulletin, Issue 5/2022.

Blanco, A., C. Boar, C. Jones and V. Midrigan (2024). Non-Linear Inflation Dynamics in Menu Cost Economies. Working Paper 32094, National Bureau of Economic Research.

Boissay, F., F. De Fiore, D. Igan, A. Pierres-Tejada, and D. Rees (2022). Are Major Advanced Economies on the Verge of a Wage-Price Spiral? BIS Bulletin No. 53.

Burriel Llombart, P., Odendahl, F., and S. Párraga Rodríguez (2024), "Asymmetries in the transmission of energy price increases and decreases to underlying inflation in the euro area and Spain," Economic Bulletin, Banco de España, Issue 2024/Q1.

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2023). Anchored Inflation Expectations. *American Economic Journal: Macroeconomics 15(1), 1-47.*

Costain, J. and A. Nakov (2011). Price Adjustments in a General Model of State-Dependent Pricing. *Journal of Money, Credit and Banking 43(2), 385-406.*

De Rezende, R. and A. Ristiniemi (2023). Monetary Policy Strategies for the Euro Area: Optimal Rules in the Presence of the ELB, ECB Working Paper Series No. 2797.

Dotsey, M., R. G. King, and A. L. Wolman (1999). State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output. *Quarterly Journal of Economics 114(2), 655-690.*

Erceg, C. J., D. W. Henderson, and A. T. Levin (2000). Optimal Monetary Policy with Staggered Wage and Price Contracts. *Journal of Monetary Economics 46(2), 281-313.*

Ferrando, A., S. Lamboglia, and J. Rariga (2025). Determinants of inflation expectations of firms in the SAFE. ECB Economic Bulletin, Issue 4/2025.

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 E. Wieland (2026), "Consumer price stickiness in the euro area during an inflation surge," ECB WP No.3181.

Gennaioli, N., M. Leva, R. Schoenle and A. Shleifer (2024). "How Inflation Expectations De-Anchor: The Role of Selective Memory Cues." Working Paper 32633, National Bureau of Economic Research.

Karadi, P. and A. Reiff (2019). Menu Costs, Aggregate Fluctuations, and Large Shocks. *American Economic Journal: Macroeconomics 2019, 11(3): 111–146.*

Koester, G. and H. Grapow (2021). The Prevalence of Private Sector Wage Indexation in the Euro Area and Its Potential Role for the Impact of Inflation on Wages. ECB Economic Bulletin, Issue 7/2021.

Kortela, T. (2016). A Shadow Rate Model with Time-Varying Lower Bound. Bank of Finland Research Discussion Paper 19/2016.

Mackowiak, B. and M. Wiederholt (2009). Optimal Sticky Prices under Rational Inattention. *American Economic Review 99(3), 769-803.*

Pfäuti, O. (2024). Inflation — Who Cares? Monetary Policy in Times of Low Attention. *Journal of Money, Credit and Banking Vol. 57, No. 5 (August 2025), 1211-1239.*

Reis, R. (2021). Losing the Inflation Anchor. *Brookings Papers on Economic Activity, Fall 2021, 307–361.*

Savignac, F., E. Gautier, Y. Gorodnichenko, and O. Coibion (2024). Firms' Inflation Expectations: New Evidence from France. *Journal of the European Economic Association 22(6), 2748–2781.*

Sims, C. A. (2003). Implications of Rational Inattention. *Journal of Monetary Economics 50(3), 665-690.*

Smets, F. and R. Wouters (2007). Shocks and Frictions in US Business Cycles. *American Economic Review 97(3), 586-606.*

APPENDIX

This appendix provides technical details of the model and simulations as well as various robustness checks for the main results.

A.1 Full Model Equations

This appendix presents the equilibrium conditions of the nonlinear Smets and Wouters (2007) model used in the main text, augmented with state-dependent price- and wage-adjustment hazards. The block structure follows the original Smets and Wouters (2007) specification: a sticky-price/sticky-wage block whose variables carry no subscript, and a flex-price/flex-wage block whose variables carry the subscript f and pin down the natural rate of output y_t^f used to define the output gap. All real variables are detrended by the deterministic balanced-growth factor γ^t, where γ is the gross quarterly productivity growth rate; nominal variables are detrended by the steady-state price level. The detrended marginal utility of consumption is denoted ξ_t, gross quarterly inflation is $\pi_t = P_t/P_{t-1}$, and the gross nominal interest rate is R_t.

For brevity we present the sticky block in detail; the flex block is identical with $\xi_p = \xi_w = 0$ (i.e. price- and wage-setting frictions removed) and serves only to compute y_t^f.

Complete nonlinear system consists of following elements: (A.1.1) household optimality conditions, capital accumulation, and Euler equation; (A.1.2) capital-rental and investment-adjustment-cost blocks; (A.1.3) price-setting with endogenous logistic hazards, Kimball curvature, price-dispersion recursions, and Kimball wedge; (A.1.4) wage-setting block; (A.1.5) state-dependent logistic hazards, steady-state calibration, and semi-elasticity formula; (A.1.6) resource constraint; (A.1.7) monetary policy rule and optimal-policy variants; (A.1.8) exogenous shock processes and solution method; (A.1.9) deterministic steady state; (A.1.10) de-anchoring modification of forward expectation in price- and wage-setting recursions.

A1.1 Households

There is a continuum of households indexed by $j \in [0, 1]$ who supply differentiated labor services and pool consumption risk through state-contingent claims so that consumption C_t and asset holdings are common across households. Household j maximizes

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \, \varepsilon_t^b \, \frac{1}{1-\sigma_c} \left(C_t - \frac{\kappa}{\gamma} \, C_{t-1} \right)^{1-\sigma_c} \exp\left(\frac{\sigma_c - 1}{1+\sigma_l} \, L_t(j)^{1+\sigma_l} \right), \qquad \text{(A1.1)}$$

subject to a sequence of nominal flow budget constraints (stated below). Here $\beta \in (0, 1)$ is the household discount factor, σ_c governs the intertemporal elasticity

of substitution, κ controls external consumption habits[4], σ_l is the inverse Frisch elasticity, and ε_t^b is a stochastic risk-premium / discount-factor wedge that follows

$$\log \varepsilon_t^b = \rho_b \log \varepsilon_{t-1}^b + \eta_t^b, \qquad \eta_t^b \sim \mathcal{N}\left(0, \sigma_b^2\right). \tag{A1.2}$$

A positive realization of η_t^b raises ε_t^b and acts as a contractionary aggregate-demand wedge in the bond Euler equation.

A1.1.1 Budget constraint

Household j accumulates physical capital K_t, supplies capital services $K_t^s = U_t K_{t-1}$ at utilization rate U_t, and rents these services to intermediate-goods firms at the real rental rate r_t^k. Its detrended period-t flow budget constraint[5], expressed in real terms, reads

$$\begin{aligned} C_t + I_t + \frac{B_t}{P_t\, \varepsilon_t^b R_t} = \frac{B_{t-1}}{P_t\, \gamma} + \frac{W_t(j)}{P_t} L_t(j) + r_t^k U_t \frac{K_{t-1}}{\gamma} \\ - a\left(U_t\right) \frac{K_{t-1}}{\gamma} + Div_t - T_t, \end{aligned} \tag{A1.3}$$

where B_t are one-period nominal bonds, $W_t(j)$ is the nominal wage of household j, Div_t are dividends from intermediate-goods firms, T_t are lump-sum taxes, and $a\left(U_t\right) K_{t-1}/\gamma$ is the resource cost of varying capital utilization. The factor $1/\gamma$ reflects the detrending of beginning-of-period capital. The capital-utilization cost satisfies $a(1) = 0$ and is parameterized as

$$a\left(U_t\right) = \frac{r_{ss}^k}{\widetilde{z}_1} \left(\exp\left(\widetilde{z}_1 \left(U_t - 1\right)\right) - 1\right), \qquad \widetilde{z}_1 \equiv \frac{\psi}{\left(1 - \psi\right) r_{ss}^k}, \tag{A1.4}$$

so that $a'(1) = r_{ss}^k$ and the curvature parameter $\widetilde{z}_1$ is pinned by $\psi \in (0, 1)$, the elasticity of capital-utilization cost.

Capital accumulates according to

[4]Because C_t is detrended by γ^t while C_{t-1} in the habit term corresponds to the previous period's detrended value, habit persistence enters the detrended utility function as $(\kappa/\gamma)\, C_{t-1}$ rather than $\kappa\, C_{t-1}$; this follows the standard treatment of habit persistence in balanced-growth models (cf. Smets and Wouters, 2007). The first-order condition in A1.6 is stated in these already-detrended units.

[5]The risk-premium wedge ε_t^b enters multiplicatively on bond holdings for algebraic convenience, following the standard convention in Smets and Wouters (2007) and the broader New Keynesian DSGE literature; it should be understood as a reduced-form representation of a financial-intermediation or preference-driven wedge between the policy rate and the household's effective return on saving, rather than as a literal component of the household's resource constraint. Equivalent formulations introduce the same wedge directly into the bond-pricing Euler equation (A1.7) with an identical effect on equilibrium dynamics.

$$K_t = (1-\delta)\frac{K_{t-1}}{\gamma} + \varepsilon_t^i \left[1 - S\left(\frac{I_t \gamma}{I_{t-1}}\right)\right] I_t, \qquad S(x) = \frac{\varphi}{2}(x-\gamma)^2, \quad \text{(A1.5)}$$

with $\delta \in (0,1)$ the depreciation rate, ε_t^i an investment-specific technology shock, and $S(\cdot)$ the standard Smets and Wouters (2007) investment adjustment-cost function with $S(\gamma) = S'(\gamma) = 0$ and $S''(\gamma) = \varphi$.

A1.1.2 First-order conditions

Letting $\Xi_t = \varepsilon_t^b \xi_t$ denote the Lagrange multiplier on the period-t budget constraint scaled by detrending, the household's first-order conditions deliver:

(A1.1.a) Marginal utility of consumption. Differentiating with respect to C_t,

$$\xi_t = \varepsilon_t^b (C_t - \frac{\kappa}{\gamma} C_{t-1})^{-\sigma_c} \exp\left(\frac{\sigma_c - 1}{1+\sigma_l}(L_t^h)^{1+\sigma_l}\right), \quad \text{(A1.6)}$$

where $L_t^h = L_t \frac{\epsilon_w + s_t^w}{1+\epsilon_w}$ is an "effective" labor index that accounts for cross-sectional wage dispersion under Kimball aggregation (see A.4); ϵ_w is the Kimball curvature parameter for wages and s_t^w is the wage-dispersion auxiliary defined in A.4.

(A1.1.b) Bond (pricing) Euler equation. Differentiating with respect to B_t,

$$\frac{1}{R_t} = \beta \gamma^{-\sigma_c} \mathbb{E}_t \left[\varepsilon_t^b \frac{\xi_{t+1}}{\xi_t} \frac{1}{\pi_{t+1}}\right]. \quad \text{(A1.7)}$$

The risk-premium wedge ε_t^b enters multiplicatively. A persistent rise in ε_t^b raises the effective discount factor on future consumption and depresses current demand — this is the "discount-factor / risk-premium".

(A1.1.c) Investment Euler equation. Letting Q_t denote Tobin's Q (the shadow price of installed capital in consumption units), the FOC for I_t is

$$1 = Q_t \varepsilon_t^i [1 - S(x_t) - S'(x_t)\, x_t] + \beta \gamma^{-\sigma_c} \mathbb{E}_t \left[\frac{\xi_{t+1}}{\xi_t} Q_{t+1} \varepsilon_{t+1}^i S'(x_{t+1})\, x_{t+1}^2\right], \quad \text{(A1.8)}$$

with $x_t \equiv I_t \gamma / I_{t-1}$.

(A1.1.d) Capital (pricing) Euler equation. Differentiating with respect to K_t,

$$Q_t = \beta\gamma^{-\sigma_c}\,\mathbb{E}_t\left[\frac{\xi_{t+1}}{\xi_t}((1-\delta)Q_{t+1} + r^k_{t+1}U_{t+1} - a\,(U_{t+1}))\right]. \tag{A1.9}$$

(A1.1.e) Capital utilization. Differentiating with respect to U_t,

$$r^k_t = a'\,(U_t) = r^k_{ss}\,\exp\left(\widetilde{z}_1\,(U_t - 1)\right). \tag{A1.10}$$

Households also choose their wage $W_t(j)$ subject to a state-dependent renegotiation friction; the FOCs from this problem are derived in A.4.

A1.2 Firms

A1.2.1 Final-goods aggregator (Kimball)

A representative competitive final goods firm aggregates a continuum of intermediate varieties $Y_t(i)$ into the final good Y_t via the Kimball (1995) aggregator

$$1 = \int_0^1 G\left(\frac{Y_t(i)}{Y_t};\,\phi_p\right)di, \tag{A1.11}$$

where G is strictly increasing and concave with $G\,(1;\phi_p) = 1$, the steady-state markup is ϕ_p, and the demand-curve curvature $\epsilon_p \geq 0$ controls strategic complementarities in pricing. The Dixit–Stiglitz CES case is recovered as $\epsilon_p \to 0$.

Cost minimization yields the following demand for variety i:

$$Y_t(i) = Y_t{\cdot}\mathcal{D}\left(\frac{P_t(i)}{P_t};\,\phi_p,\epsilon_p\right), \tag{A1.12}$$

where the demand schedule $\mathcal{D}$ is more elastic far from $P_t(i)/P_t = 1$ than the constant-elasticity case, generating realistic "kinked" demand.

A1.2.2 Intermediate-goods producers

Intermediate firm i produces variety i using capital services $K^s_t(i)$ and labor $L_t(i)$ via Cobb–Douglas technology with a fixed cost:

$$Y_t(i) = \varepsilon^a_t\,K^s_t(i)^{\alpha}L_t(i)^{1-\alpha} - (\phi_p - 1)\;Y_{ss}, \tag{A1.13}$$

where ε^a_t is total factor productivity and the fixed cost is calibrated so that profits are zero in steady state. Cost minimization gives the standard optimal capital/labor ratio

$$\frac{K^s_t}{L_t} = \frac{\alpha}{1-\alpha}\,\frac{w_t}{r^k_t}, \tag{A1.14}$$

and real marginal cost

$$mc_t = \left(\frac{\alpha}{1-\alpha}\,\frac{w_t}{r_t^k}\right)^{-\alpha}\frac{w_t}{(1-\alpha)\,\varepsilon_t^a} \;=\; \frac{1}{\varepsilon_t^a}\,\alpha^{-\alpha}(1-\alpha)^{-(1-\alpha)}\,w_t^{1-\alpha}\,\left(r_t^k\right)^{\alpha}. \tag{A1.15}$$

A1.2.3 Aggregate production

Aggregating across varieties and accounting for cross-sectional price dispersion through the Kimball aggregator delivers the aggregate production function

$$Y_t \cdot \frac{\epsilon_p + s_t^p}{1+\epsilon_p} = \varepsilon_t^a\, K_t^{s\,\alpha} L_t^{1-\alpha} - (\phi_p - 1)\; Y_{ss}, \tag{A1.16}$$

where s_t^p is the price-dispersion auxiliary defined in A.3.

A1.3 Price setting with state-dependent adjustment

A fraction h_t^p of intermediate firms re-optimize their price in period t; the remaining fraction $1 - h_t^p$ partially index to lagged inflation according to

$$P_t(i) = P_{t-1}(i) \cdot \pi_{ss}^{1-\iota_p}\,\pi_{t-1}^{\iota_p}, \tag{A1.17}$$

where $\iota_p \in [0,1]$ is the degree of price indexation. The novel feature of this model is that h_t^p is not a constant Calvo parameter but a state-dependent hazard.

A1.3.1 Optimal reset price

A re-optimizing firm chooses its reset price P_t^* to maximize the present discounted value of profits, knowing it will not re-optimize in any future period $t+s$ with probability $\prod_{k=1}^{s}\left(1-h_{t+k}^p\right)$ but will instead index. After detrending, the FOC can be written compactly as a system of three forward-recursive auxiliary variables $\gamma_{1,t}^p$, $\gamma_{2,t}^p$, $\gamma_{3,t}^p$, mirroring the Schmitt-Grohé and Uribe (2007) representation:

$$\frac{1+\phi_p\epsilon_p}{1+\epsilon_p}\,p_t^*\,\gamma_{1,t}^p \;=\; \phi_p\,\gamma_{2,t}^p \;+\; \frac{\epsilon_p\,(\phi_p-1)}{1+\epsilon_p}\,(p_t^*)^{\frac{\phi_p(1+\epsilon_p)}{\phi_p-1}+1}\,\gamma_{3,t}^p, \tag{A1.18}$$

where $p_t^* = P_t^*/P_t$ is the reset relative price and $\gamma_{1,t}^p, \gamma_{2,t}^p, \gamma_{3,t}^p$ satisfy the recursions

$$\gamma_{1,t}^p = \delta_t^{p\,\frac{\phi_p(1+\epsilon_p)}{\phi_p-1}}\,Y_t + \beta\gamma^{1-\sigma_c}\,\mathbb{E}_t\left[\left(1-h_{t+1}^p\right)\frac{\xi_{p,t+1}}{\xi_{p,t}}\left(\frac{\pi_{ss}^{1-\iota_p}\pi_t^{\iota_p}}{\pi_{t+1}}\right)^{-\frac{1+\phi_p\epsilon_p}{\phi_p-1}}\gamma_{1,t+1}^p\right], \tag{A1.19}$$

$$\gamma_{2,t}^{p} = \delta_t^{p\,\frac{\phi_p(1+\epsilon_p)}{\phi_p-1}}\, mc_t\, \varepsilon_t^{p} Y_t + \beta\gamma^{1-\sigma_c}\, \mathbb{E}_t \left[\left(1-h_{t+1}^{p}\right) \frac{\xi_{p,t+1}}{\xi_{p,t}} \left(\frac{\pi_{ss}^{1-\iota_p}\pi_t^{\iota_p}}{\pi_{t+1}} \right)^{-\frac{\phi_p(1+\epsilon_p)}{\phi_p-1}} \gamma_{2,t+1}^{p} \right], \tag{A1.20}$$

$$\gamma_{3,t}^{p} = Y_t + \beta\gamma^{1-\sigma_c}\, \mathbb{E}_t \left[\left(1-h_{t+1}^{p}\right) \frac{\xi_{p,t+1}}{\xi_{p,t}} \left(\frac{\pi_{ss}^{1-\iota_p}\pi_t^{\iota_p}}{\pi_{t+1}} \right) \gamma_{3,t+1}^{p} \right], \tag{A1.21}$$

and δ_t^p is the wedge between the aggregate price and the Kimball aggregator (see below). Here ε_t^p is the price-markup shock.

Key feature of the SDP formulation. The non-adjustment continuation factor $\left(1-h_{t+1}^{p}\right)$ enters the recursions as a *time-varying* discount-like term. Under standard Calvo it is the constant ξ_p; under SDP it falls when expected future inflation is high (because more firms will reset), shortening the effective horizon of fixed-price firms and steepening the Phillips curve.

A1.3.2 Price index, dispersion, and Kimball wedge

Combining the reset-price decision with the indexation rule and aggregating over varieties yields the price-index identity

$$1 = h_t^{p} \left(\frac{p_t^{*}}{\delta_t^{p}} \right)^{-\frac{1+\phi_p\epsilon_p}{\phi_p-1}} + \left(1-h_t^{p}\right) \left(\frac{\pi_{ss}^{1-\iota_p}\pi_{t-1}^{\iota_p}}{\pi_t} \frac{\delta_{t-1}^{p}}{\delta_t^{p}} \right)^{-\frac{1+\phi_p\epsilon_p}{\phi_p-1}}, \tag{A1.22}$$

the price-dispersion recursion

$$s_t^{p} = h_t^{p} \left(\frac{p_t^{*}}{\delta_t^{p}} \right)^{-\frac{\phi_p(1+\epsilon_p)}{\phi_p-1}} + \left(1-h_t^{p}\right) \left(\frac{\pi_{ss}^{1-\iota_p}\pi_{t-1}^{\iota_p}}{\pi_t} \frac{\delta_{t-1}^{p}}{\delta_t^{p}} \right)^{-\frac{\phi_p(1+\epsilon_p)}{\phi_p-1}} s_{t-1}^{p}, \tag{A1.23}$$

and the Kimball-wedge identity

$$1+\epsilon_p = \left(1+\epsilon_p\, s_t^{p,l}\right) \delta_t^{p}, \qquad s_t^{p,l} = h_t^{p}\, \frac{p_t^{*}}{\delta_t^{p}} + \left(1-h_t^{p}\right) \frac{\pi_{ss}^{1-\iota_p}\pi_{t-1}^{\iota_p}}{\pi_t} \frac{\delta_{t-1}^{p}}{\delta_t^{p}}\, s_{t-1}^{p,l}. \tag{A1.24}$$

The auxiliary variable $s_t^{p,l}$ is a "level" version of the dispersion recursion that closes the system in the presence of Kimball curvature; under Dixit–Stiglitz ($\epsilon_p = 0$) it drops out and $\delta_t^p = 1$.

A1.4 Wage setting with state-dependent renegotiation

The wage block mirrors the price block. A fraction h_t^w of households re-optimize their wage in period t, while the remainder index according to $W_t(j) = W_{t-1}(j)\,\pi_{ss}^{1-\iota_w}\pi_{t-1}^{\iota_w}\gamma$. A re-optimizing household chooses its reset real wage $w_t^* = W_t^*/P_t$ to maximize expected lifetime utility from labor income, subject to its labor-demand schedule from the Kimball aggregator over differentiated labor varieties (with markup ϕ_w and curvature ϵ_w).

The FOC has the same structure as in A1.3:

$$\frac{1+\phi_w\epsilon_w}{1+\epsilon_w}\,w_t^*\,\gamma_{1,t}^w \;=\; \phi_w\,\gamma_{2,t}^w \;+\; \frac{\epsilon_w\,(\phi_w-1)}{1+\epsilon_w}\,(w_t^*)^{\frac{\phi_w(1+\epsilon_w)}{\phi_w-1}+1}\,\gamma_{3,t}^w, \tag{A1.25}$$

with

$$\gamma_{1,t}^w = \delta_t^{w\,\frac{\phi_w(1+\epsilon_w)}{\phi_w-1}}\,L_t + \beta\gamma^{1-\sigma_c}\,\mathbb{E}_t\left[\left(1-h_{t+1}^w\right)\frac{\xi_{w,t+1}}{\xi_{w,t}}\left(\frac{\pi_{ss}^{1-\iota_w}\pi_t^{\iota_w}}{\pi_{t+1}}\right)^{-\frac{1+\phi_w\epsilon_w}{\phi_w-1}}\gamma_{1,t+1}^w\right], \tag{A1.26}$$

$$\gamma_{2,t}^w = \delta_t^{w\,\frac{\phi_w(1+\epsilon_w)}{\phi_w-1}}\,MRS_t\,\varepsilon_t^w\,L_t + \beta\gamma^{1-\sigma_c}\,\mathbb{E}_t\left[\left(1-h_{t+1}^w\right)\frac{\xi_{w,t+1}}{\xi_{w,t}}\left(\frac{\pi_{ss}^{1-\iota_w}\pi_t^{\iota_w}}{\pi_{t+1}}\right)^{-\frac{\phi_w(1+\epsilon_w)}{\phi_w-1}}\gamma_{2,t+1}^w\right], \tag{A1.27}$$

$$\gamma_{3,t}^w = L_t + \beta\gamma^{1-\sigma_c}\,\mathbb{E}_t\left[\left(1-h_{t+1}^w\right)\frac{\xi_{w,t+1}}{\xi_{w,t}}\left(\frac{\pi_{ss}^{1-\iota_w}\pi_t^{\iota_w}}{\pi_{t+1}}\right)\gamma_{3,t+1}^w\right], \tag{A1.28}$$

where $MRS_t = \left(C_t - \frac{\kappa}{\gamma}C_{t-1}\right)(L_t^h)^{\sigma_l}$ is the marginal rate of substitution between consumption and (effective) labor.

The wage-index identity, dispersion recursion, level auxiliary $s_t^{w,l}$ and Kimball wedge δ_t^w are exactly analogous to the price block with w^* in place of p^* and $(\phi_w, \epsilon_w, \iota_w)$ in place of $(\phi_p, \epsilon_p, \iota_p)$. The aggregate-wage-versus-aggregator wedge is

$$(1+\epsilon_w)\;w_t = \left(1+\epsilon_w\,s_t^{w,l}\right)\,\delta_t^w. \tag{A1.29}$$

A1.5 The state-dependent hazards rates

The central modeling departure from Smets and Wouters (2007) is that the price- and wage-adjustment hazards h_t^p and h_t^w are functions of current quarterly inflation rather than constants. We adopt a logistic specification:

$$h_t^p = \frac{1}{1 + exp\left(-\left(a^p + \phi_h^p\, \widehat{\pi}_t\right)\right)}, \qquad h_t^w = \frac{1}{1 + exp\left(-\left(a^w + \phi_h^w\, \widehat{\pi}_t\right)\right)}, \quad \text{(A1.30)}$$

where $\widehat{\pi}_t = 100 \cdot (\pi_t - 1)$ is net quarterly inflation expressed in percentage points, ϕ_h^p and ϕ_h^w are sensitivity parameters (semi-elasticity of the hazard with respect to inflation), and a^p, a^w are intercepts pinned down by the requirement that the steady-state hazards equal the calibrated Smets and Wouters (2007) Calvo probabilities:

$$a^p = log\left(\frac{1-\xi_p}{\xi_p}\right) - \phi_h^p\, \widehat{\pi}_{ss}, \qquad a^w = log\left(\frac{1-\xi_w}{\xi_w}\right) - \phi_h^w\, \widehat{\pi}_{ss}. \quad \text{(A1.31)}$$

This formulation ensures that the deterministic steady state of the SDP model coincides with that of the standard Smets and Wouters (2007) Calvo model regardless of ϕ_h^p, ϕ_h^w, meaning that any difference between the two models in simulation is driven purely by the time variation in the hazards along the equilibrium path.

The semi-elasticity of the hazard with respect to net quarterly inflation, evaluated at any state, is

$$\varepsilon_t^h \equiv \frac{\partial log h_t^p}{\partial \widehat{\pi}_t} = (1 - h_t^p)\ \phi_h^p, \quad \text{(A1.32)}$$

and similarly for wages. At the steady state this gives $\varepsilon_{ss}^h = (1 - h_{ss}^p)\ \phi_h^p = \xi_p\, \phi_h^p$. The baseline calibration $\phi_h^p = 0.45$ together with $\xi_p = 0.77$ yields $\varepsilon_{ss}^h \approx 0.35$, in line with the empirical target of $\varepsilon^h \approx 0.32$–$0.38$ documented in Section 4 of the main text. For wages, the much smaller calibration $\phi_h^w = 0.075$ delivers a near-flat wage hazard, reflecting the predominantly calendar-driven nature of euro-area collective wage bargaining.

A1.5.1 Interpretation as an approximation

The logistic specification is best understood as a *reduced form approximation* to the underlying menu-cost decision problem. In a fully micro-founded menu-cost economy, a firm i resets its price in period t if and only if its idiosyncratic price gap $|p_t(i) - p_t^*|$ exceeds a threshold determined by the fixed adjustment cost. The cross-sectional fraction of resetters then depends on the joint distribution of price gaps, which itself depends on the entire history of aggregate inflation. Solving this distributional fixed-point problem is computationally expensive and has been the subject of a large recent literature (see e.g. Blanco et al., 2024, and references therein).

The logistic hazard collapses this distributional problem onto a single sufficient statistic — current quarterly inflation $\widehat{\pi}_t$. The economic justification for this

approximation is that during inflation surges, accumulated real-wage and real-price misalignment correlate strongly with realized inflation. Consequently, $\widehat{\pi}_t$ acts as a sufficient statistic for "how stretched" the cross-sectional distribution is. The semi-elasticity ϕ_h^p can then be disciplined by the empirical relationship between adjustment frequencies and inflation reported in Gautier et al. (2026) for the euro area and Blanco et al. (2024) for the United Kingdom, both of which imply $\varepsilon^h \approx 0.3$ at surge regimes, or estimated as done in this paper.

In sum, the approximation has three attractive properties:

1. *Tractability.* It preserves the recursive structure of the Smets and Wouters (2007) price- and wage-setting blocks; the only modification is that h_t^p, h_t^w become endogenous variables rather than parameters.
2. *First-order equivalence to Calvo model.* Because the hazards are pinned at the Calvo SS values, the first-order log-linearization of the SDP model is observationally equivalent to standard Calvo. Nonlinearity matters only through second- and higher-order terms.
3. *Empirically disciplined extensive margin.* The single parameter ϕ_h^p controls how much of the inflation surge is attributable to the extensive margin (more firms resetting) versus the intensive margin (resetters choosing larger price changes). Calibrating to $\varepsilon^h \approx 0.33$ closely matches the menu-cost evidence.

This said, the reduced-form nature of the hazard means that the model does not endogenize the cross-sectional distribution of price gaps. In particular, the hazard depends only on current $\widehat{\pi}_t$, not on the history of past inflation. During very prolonged inflation surges, when price gaps are accumulated over many quarters, the true menu-cost hazard would be higher than the logistic specification predicts at any given level of contemporary inflation. Conversely, immediately after a surge ends, the logistic hazard would over-predict adjustment because the cross-sectional distribution is still concentrated near the optimum. The model is therefore best suited to studying the impulse response to a single inflationary disturbance, not very protracted inflation regimes.

A1.6 Resource constraint and aggregate accounting

Real output is allocated between private consumption, investment, government spending, and capital-utilization costs:

$$C_t + I_t + g_t Y_t + a\left(U_t\right) \frac{K_{t-1}}{\gamma} = Y_t, \tag{A1.33}$$

where $g_t Y_t$ is government spending and g_t is its share of output. The government share follows

$$\log g_t = (1-\rho_g)\log g_{ss} + \rho_g \log g_{t-1} - \rho_{ga}\log \varepsilon_{t-1}^a + \eta_t^g, \tag{A1.34}$$

allowing for feedback from lagged TFP. Capital services satisfy $K_t^s = U_t K_{t-1}/\gamma$.

1.7 Monetary policy rule

The policy rate follows a Taylor-type rule with smoothing, response to inflation and to the output gap, and an i.i.d. policy shock:

$$\log R_t = \rho_R \log R_{t-1} + (1-\rho_R)\left[\log\left(\frac{\pi_{ss}}{\bar{\beta}}\right) + r_\pi \log\left(\frac{\pi_t}{\pi_{ss}}\right) + r_y \log\left(\frac{Y_t}{Y_t^f}\right)\right] + r_{\Delta y}\,\Delta \log\left(Y_t/Y_t^f\right) + \log \varepsilon_t^r, \tag{A1.35}$$

equivalently in multiplicative form

$$R_t = R_{t-1}^{\rho_R}\left[\frac{\pi_{ss}}{\bar{\beta}}\right]^{1-\rho_R}\left(\frac{\pi_t}{\pi_{ss}}\right)^{r_\pi(1-\rho_R)}\left(\frac{Y_t}{Y_t^f}\right)^{r_y(1-\rho_R)}\left(\frac{Y_t/Y_t^f}{Y_{t-1}/Y_{t-1}^f}\right)^{r_{\Delta y}}\varepsilon_t^r, \tag{A1.36}$$

where $\bar{\beta} = \beta\,\gamma^{-\sigma_c}$ is the steady-state real discount factor and $\pi_{ss}/\bar{\beta}$ is the steady-state nominal rate (Fisher equation). The policy shock follows $\log \varepsilon_t^r = \rho_r \log \varepsilon_{t-1}^r + \eta_t^r$.

In the optimal-policy variants of Section 6 of the main text, this rule is replaced by either the Ramsey first-order condition associated with a quadratic loss, the OSR-optimized parameters (ρ_R, r_π, r_y) in the same rule, or the Erceg–Henderson–Levin (2000) loss function over price inflation, wage inflation, and the output gap.

For completeness, when the effective lower bound (ELB) is active, the rule is replaced by

$$R_t = max\left\{1+\bar{b},\ R_t^{\text{Taylor}}\right\}, \tag{A1.37}$$

where $\bar{b}$ is the floor on the net nominal rate. This reformulation is used only in the perfect-foresight variant of the model; for stochastic simulations the ELB constraint is suppressed because $\max(\cdot)$ is non-differentiable.

A1.8 Exogenous processes and solution method

The seven structural shocks follow

$$\log \varepsilon_t^a = \rho_a \log \varepsilon_{t-1}^a + \eta_t^a, \tag{A1.38}$$
$$\log \varepsilon_t^b = \rho_b \log \varepsilon_{t-1}^b + \eta_t^b, \tag{A1.39}$$
$$\log \varepsilon_t^i = \rho_i \log \varepsilon_{t-1}^i + \eta_t^i, \tag{A1.40}$$
$$\log \varepsilon_t^p = \rho_p \log \varepsilon_{t-1}^p + \left(\eta_t^p - \mu_p \eta_{t-1}^p\right) / \pi_{2,p}, \tag{A1.41}$$
$$\log \varepsilon_t^w = \rho_w \log \varepsilon_{t-1}^w + \left(\eta_t^w - \mu_w \eta_{t-1}^w\right) / \pi_{2,w}, \tag{A1.42}$$
$$\log \varepsilon_t^r = \rho_r \log \varepsilon_{t-1}^r + \eta_t^r. \tag{A1.43}$$

where the markup shocks are ARMA(1,1) with the rescaling factors

$$\pi_{2,p} = \frac{\left(1 - \xi_p \bar{\beta}\right)(1 - \xi_p)}{\xi_p \left(1 + (\phi_p - 1)\, \epsilon_p\right)}, \qquad \pi_{2,w} = \frac{\left(1 - \xi_w \bar{\beta}\right)(1 - \xi_w)}{\xi_w \left(1 + (\phi_w - 1)\, \epsilon_w\right)}, \tag{A1.44}$$

calibrated so that the markup innovations enter the Phillips curves with unit coefficient (this is the standard Smets and Wouters (2007) normalization).

Government purchases relative to trend output, $g_t = G_t/(\gamma^t Y)$, follow an exogenous AR(1) process with feedback from lagged technology, as in Smets and Wouters (2007):

$$log\ g_t = (1 - \rho_g)\, log\ g + \rho_g\, (log\ g_{t-1} - \rho_{ga}\, log\ \varepsilon_{t-1}^a) + \eta_t^g, \quad \eta_t^g \sim N(0,\, \sigma_g), \tag{A1.45}$$

where ρ_g is the autoregressive coefficient, ρ_{ga} captures the response of government purchases to lagged technology, and σ_g is the standard deviation of the innovation. Government purchases enter the aggregate resource constraint (equation A.1 of the nonlinear system) and do not affect the marginal utility of private consumption or production, with Ricardian equivalence imposed through lump-sum financing.

The model is solved numerically with Dynare 7.0 using third-order perturbation around the deterministic steady state, with the pruning algorithm of Andreasen, Fernández-Villaverde and Rubio-Ramírez (2018). The first-order solution is observationally equivalent to the standard Calvo Smets and Wouters (2007) model because the hazards have zero first-order derivatives at the steady state when evaluated as Calvo constants; non-trivial state dependence shows up only at second and higher orders of the policy function. Impulse responses reported in the main text are computed at order two, where possible, starting from the ergodic mean of the stochastic steady state, and are averaged over a sufficiently large number of Monte Carlo simulations.

A1.9 Steady state

In the deterministic steady state all shock processes are at their unconditional means: $\varepsilon^a_{ss} = \varepsilon^b_{ss} = \varepsilon^i_{ss} = \varepsilon^p_{ss} = \varepsilon^w_{ss} = \varepsilon^r_{ss} = 1$, g_{ss} equals the calibrated share, capital utilization $U_{ss} = 1$, and reset prices and wages equal aggregates so that $p^*_{ss} = 1$ and $w^*_{ss} = w_{ss}$. Inflation equals its target π_{ss}; the gross nominal rate equals the Fisher level $R_{ss} = \pi_{ss}/\bar{\beta}$; price and wage dispersion are unity ($s^p_{ss} = s^w_{ss} = 1$) and Kimball wedges are unity ($\delta^p_{ss} = \delta^w_{ss} = 1$). The rental rate of capital is $r^k_{ss} = \gamma^{\sigma_c}/\beta - (1-\delta)$ from the capital Euler equation. The marginal cost in steady state equals $1/\phi_p$, pinning the real wage from the cost-minimization condition. Capital, investment, output, and consumption are then pinned recursively by the production function, capital-accumulation identity, and the resource constraint with $g_{ss}Y_{ss}$ as the residual government component. Hours L_{ss} are determined by the labor-supply equation with $w_{ss} = (1/\phi_w)\ MRS_{ss}$, which combined with the production FOC fully closes the steady state.

A1.10 Expectations de-anchoring (forward channel)

The de-anchoring extension used in Section 5 modifies the price- and wage-setting blocks in two steps, leaving the rest of the model unchanged. First, a perceived inflation trend π^e evolves through constant-gain adaptive learning, see Carvalho et al. (2023), Alichi et al. (2009), and Orphanides and Williams (2005), updating toward lagged realized inflation: $log\ \pi^e_t = log\ \pi^e_{t-1} + g*\ (log\ \pi_{t-1} - log\ \pi^e_{t-1})$, with constant gain g (baseline 0.05). The lag keeps π^e predetermined, so it enters as a state and the system remains solvable by perturbation at any order; setting g = 0 (or anchoring fully) pins π^e to the central bank inflation target $\overline{\pi} = 0$.

Second, in the forward-looking term of the price recursions only, the model-consistent expectation of next-period inflation $E_t\pi_{t+1}$ is replaced by the geometric average $\pi^{fwd} = (E_t\pi_{t+1})^{\lambda}\ (\pi^e)^{1-\lambda}$. The corresponding forward-looking term in the wage recursions retains the fully rational $E_t\pi_{t+1}$ unchanged, so wages carry neither the price hazard nor the de-anchoring mechanism; both channels are restricted to price-setting, consistent with the wage-side Granger-causality and threshold results in Appendix TA.B. The anchoring weight $\lambda \in [0,1]$ is the only new structural parameter. Because at the deterministic steady state $\pi = E_t\pi_{t+1} = \pi^e = \overline{\pi}$, meaning steady state π^{fwd} equals $\overline{\pi}$ for any λ, so the steady state and the calibration of Appendix A2 are left intact. We set $\lambda = 0.85$ in the baseline. The indexation terms in the price and wage equations remain anchored at $\overline{\pi}$, so de-anchoring enters exclusively through the forward expectation; an alternative routing through the indexation terms was found to be quantitatively negligible under the estimated indexation weight and is not used.

A.2 Bayesian Estimation Details

(A.2.1) Data sources and measurement equations, including shadow-rate adjustment for ELB period 2014Q3-2021Q1 (Kortela, 2016); (A.2.2) prior distributions; (A.2.3) fixed/calibrated parameters; (A.2.4) identification analysis - all 21 pa-

rameters pass full-rank tests (Iskrev, 2010; Komunjer-Ng, 2011; Qu-Tkachenko, 2012); (A.2.5) MCMC and convergence diagnostics (four chains of 400,000 draws; PSRF <= 1.002; Table A.2.2); (A.2.6) smoothed structural shocks and 2022-23 markup shock cluster.

A.2.7 Perturbation order. Main-text simulations use second-order pruned perturbation throughout. Third-order was investigated and produces qualitatively identical but quantitatively overshooting results: price hazard change of approx 0.65 vs. empirical target of approx 0.15 at 4σ; hazard exceeds 1.0 at slightly larger shock sizes; Ramsey produces implausible initial rate responses of approx -3 pp. The order=2 choice is conservative: it understates SDP amplification, so the paper's quantitative results provide a lower bound on true nonlinear effects. Full comparison documented in Figure A2.9. Particle-filter nonlinear estimation was also attempted but encountered prohibitive weight degeneracy on COVID-period observations.

A2.1 Data and measurement equations

The estimation uses quarterly euro-area data over 1999Q1–2025Q4, with eight observables entering the measurement system. Real consumption (C), real investment (I), real GDP (Y), and the real wage (W) are expressed in per-capita terms and enter as quarter-on-quarter log differences. GDP-deflator inflation (π) is the quarter-on-quarter log difference of the GDP deflator. The policy rate (R) is the annual $EONIA/£STR$ equivalent, expressed in net quarterly terms. Nevertheless, we adjust this with the shadow interest rate during recent liquidity trap period 2014Q3-2021Q1, see Kortela (2016). Hours worked (L) enter as a demeaned log level. The output gap (ygap) as well as other variables taken from the BMPE database except for quarterly euro area population statistics which are from the Eurostat.

We note that the output gap observable is constructed by the BMPE using a distinct (semi-structural) methodology and therefore need not correspond exactly to the flexible-price natural-rate concept y_t^{nat} that our model equations imply. We include it, following common practice in this literature, because it disciplines the model's estimate of potential output using information beyond what the other seven observables alone would imply, at the cost of this measurement-concept mismatch.

The mapping between observables and model variables is standard. Letting hatted variables denote model log-deviations from steady state:

$dc_t^{obs} = \Delta\hat{c}_t$, $di_t^{obs} = \Delta\hat{\imath}_t$, $dy_t^{obs} = \Delta\hat{y}_t$, $dw_t^{obs} = \Delta\widehat{w}_t$, $pi_t^{obs} = \hat{\pi}_t$, $R_t^{obs} = \widehat{R}_t$, $L_t^{obs} = \widehat{l}_t$, $ygap_t^{obs} = \widehat{y}_t - \widehat{y}_t^{nat}$.

Small measurement-error variances are added on each observable's state equation to absorb residual data idiosyncrasies. For the growth observables (dc^{obs}, di^{obs}, dy^{obs}, dw^{obs}, pi^{obs}, R^{obs}, L^{obs}) the measurement-error standard deviations are set to three times approximately ten percent of the in-sample data standard deviation; for the output gap the measurement-error standard deviation is set

Figure A2.1: Observables

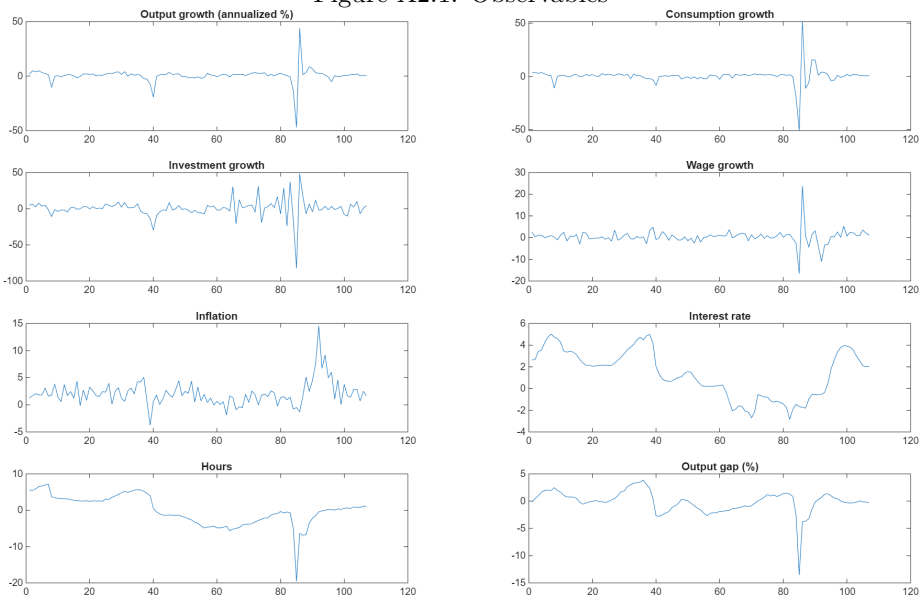

to 0.005. These values are large enough to prevent the Kalman filter from over-fitting outliers but small enough to leave most information in the data for structural inference.

A2.2 Priors

Prior distributions are centered on the posterior modes of De Rezende and Ristiniemi (2023) for the same set of structural parameters, adapted to our model specification. Prior families follow the Smets and Wouters (2007) conventions: beta distributions for parameters constrained to [0,1] (Calvo probabilities, indexation, persistence, habit, capital share, smoothing coefficient), normal distributions for unconstrained parameters (utility curvature, investment adjustment cost, Taylor-rule response coefficients), and inverse-gamma distributions for shock standard deviations. Prior means and standard deviations are listed in column 4–5 of Table A2.1.

The inverse-gamma priors on shock standard deviations are intentionally loose (prior mean 0.05, prior s.d. 0.50). This corresponds to a very weakly informative prior consistent with Smets and Wouters (2007). The objective is to let the data determine shock magnitudes; the smoothed shocks reported in Appendix A2.6 confirm that the data is highly informative about most shock variances. Figures A2.3 and A2.4 plot the prior and posterior densities for the structural parameters and the shock standard deviations, respectively; in nearly all cases the posterior is far tighter than the prior, confirming that the sample is informative.

A2.2a Estimation approach and posterior results

The model is estimated on quarterly euro-area data over the sample 1999Q1–2025Q4 (108 observations). We employ a Bayesian approach using the random-walk Metropolis–Hastings algorithm with priors centered on the De Rezende and Ristiniemi (2023) posterior modes for the euro area. The model is solved at first order using the Kalman filter to evaluate the likelihood. Posterior sampling is based on four parallel chains of 400,000 draws each, with the first 20% discarded as burn-in; convergence is verified by Brooks and Gelman (1998) diagnostics, Geweke (1992) chain-stationarity tests, and inspection of the running posterior moments.

The estimation is performed on the model with constant Calvo probabilities (i.e., the state-dependent pricing and wage hazards described in A.1 are switched off for estimation). This is the standard approach in the New Keynesian estimation literature and is supported by the practical observation that, even during the inflation surge, the empirical price-adjustment frequency deviates from its low-inflation average for only a few quarters. The estimated structural parameters are therefore best interpreted as governing the average, low-inflation pricing regime; state dependence is then layered on top of these estimates in the subsequent simulations.

Eight observables enter the measurement equations: real consumption growth, real investment growth, real GDP growth, real wage growth, GDP-deflator inflation, the policy rate (augmented with the De Rezende and Ristiniemi (2023) shadow rate during the ELB period 2014Q3–2021Q1 following Kortela (2016)), hours worked, and the output gap. Small measurement-error variances are included on each observable to absorb residual data idiosyncrasies.

Twenty-one parameters are estimated jointly: fifteen structural parameters governing preferences, technology, price- and wage-setting, and the monetary-policy rule, plus six standard deviations of structural shocks (productivity, risk-premium, government spending, investment-specific, price-markup, and wage-markup). The posterior estimates lie close to conventional values from euro-area DSGE estimations. The Calvo probabilities of non-adjustment are 0.66 for prices and 0.75 for wages, implying average price and wage spells of approximately three and four quarters respectively. Habit formation ($\kappa = 0.58$) and the inverse Frisch elasticity ($\sigma_l = 0.89$) are in line with values reported by Smets and Wouters (2007) and De Rezende and Ristiniemi (2023). The estimated monetary-policy rule features a strong inflation response ($r_\pi = 2.06$) and high interest-rate inertia ($\rho_R = 0.87$). The capital share ($\alpha = 0.15$) is somewhat lower than the calibrated values often imposed in literature.

A noteworthy finding concerns the price-markup shock. Its posterior mean standard deviation is 0.0087 with a tight 90% HPD interval of [0.0073, 0.0101], roughly five times smaller than the conventional calibrated value of 0.05 used in earlier studies. However, the smoothed time series of the price-markup shock reveals substantial heterogeneity across the sample: typical-quarter innovations are

Table A2.1: Posterior summary of parameters, euro-area estimation, 1999Q1–2025Q3

		Prior			Posterior	
Parameter		Dist.	Mean	S.d.	Mean	90% HPD
A. Structural parameters						
ξ_p	Price Calvo probability	Beta	0.50	0.10	0.66	[0.57, 0.75]
ξ_w	Wage Calvo probability	Beta	0.75	0.05	0.75	[0.71, 0.80]
ι_p	Price indexation	Beta	0.50	0.20	0.85	[0.73, 0.98]
κ	Habit persistence	Beta	0.70	0.05	0.58	[0.52, 0.64]
σ_c	Inverse IES	Normal	1.50	0.10	1.21	[1.06, 1.36]
σ_l	Inverse Frisch elasticity	Normal	2.00	0.25	0.89	[0.61, 1.17]
φ	Investment adjustment cost	Normal	6.00	1.50	5.98	[4.44, 7.43]
α	Capital share	Normal	0.20	0.025	0.15	[0.12, 0.18]
B. Monetary policy rule						
ρ_R	Interest-rate smoothing	Beta	0.85	0.03	0.87	[0.84, 0.90]
r_π	Inflation response	Normal	1.85	0.25	2.06	[1.70, 2.42]
r_y	Output-gap response	Normal	0.10	0.025	0.11	[0.07, 0.16]
C. Shock persistence						
ρ_a	Technology	Beta	0.50	0.20	0.93	[0.88, 0.99]
ρ_b	Risk premium	Beta	0.80	0.10	0.42	[0.31, 0.53]
ρ_{ga}	Technology in spending	Normal	0.50	0.20	0.82	[0.56, 1.07]
ρ_w	Wage markup	Beta	0.50	0.20	0.75	[0.51, 0.97]
D. Shock standard deviations ($\times 100$)						
σ_a	Technology	Inv. Gamma	5.00	50.00	0.68	[0.58, 0.79]
σ_b	Risk premium	Inv. Gamma	5.00	50.00	7.41	[4.14, 10.48]
σ_g	Exogenous spending	Inv. Gamma	5.00	50.00	1.48	[1.02, 1.93]
σ_i	Monetary policy	Inv. Gamma	5.00	50.00	14.02	[9.75, 18.12]
σ_p	Price markup	Inv. Gamma	5.00	50.00	0.87	[0.73, 1.01]
σ_w	Wage markup	Inv. Gamma	5.00	50.00	0.78	[0.66, 0.90]

an order of magnitude smaller than the value implied by the standard deviation, whereas crisis-period innovations during 2022–2023 reach approximately 0.035, about four standard deviations above the unconditional mean. This empirical pattern disciplines the shock size used in the subsequent simulations. In the impulse-response analyses in Sections 4 and 6 of the main text we set the markup innovation to approximately four times the estimated posterior-mean standard deviation ($4\sigma \approx 0.035$), the magnitude of the 2022–2023 innovation cluster identified by the Kalman smoother (Figures A2.6–A2.7). This brings the experiment into correspondence with what the data attribute to the markup shock during the surge, rather than the conventional 0.05 value that the estimation rejects.

A2.3 Parameters not estimated

A subset of parameters is fixed at calibrated values because either the data is uninformative about them or identification is empirically problematic in our sample. We document these choices for transparency.

Parameter	**Value**	**Rationale**
ψ (capacity util. elast.)	0.5	Standard SW calibration; weakly identified.
ι_w (wage indexation)	0.5	Identified jointly with ξ_w; calibrated to median SW07 value.
ρ_p (price-markup AR(1))	0.05	Preliminary estimation indicated near-iid behavior; consistent with De Rezende and Ristiniemi (2023).
ρ_i (inv.-specific AR(1))	0.05	As above.
ρ_g (govt. spending AR(1))	0.97	Highly persistent; calibrated to long-run fiscal expenditure dynamics.
$\sigma(\eta^r)$ monetary shock	0.005	Identified poorly given $EONIA$/€STR data and ZLB period; conservative calibration.
μ_p, μ_w (MA components)	0	Identification of MA components on markup shocks is weak; following SW07 baseline.
Kimball curvatures ε_p, ε_w	0	Set to zero to isolate state-dependent pricing as the source of nonlinearity (see main text).
Steady-state values	data-implied	Trend growth, inflation target, depreciation, etc., calibrated to sample means.

The structural parameters used in the simulations are the posterior means reported in Table A2.1 jointly with the fixed values listed in the table above. The remaining state-dependent pricing parameters not covered by the Bayesian

estimation are the price-hazard sensitivity $\phi_h^p = 0.45$, calibrated to the surge-regime semi-elasticitity documented in Gautier et al. (2026); the Kimball curvatures $\epsilon_p = \epsilon_w = 0$, set to zero to isolate state-dependent pricing as the source of nonlinearity (as noted in the table above); and the impulse-response shock magnitude $\sigma(\eta^p) \approx 0.035$, set to roughly four times the estimated posterior-mean standard deviation (4σ) and used in the simulations of Sections 4 and 6 to match the magnitude of the smoothed 2022–2023 markup innovation cluster identified by the Kalman smoother (Figures 2.6–2.7), rather than the conventional 0.05 value, which the estimation rejects. Steady-state values for trend growth, the inflation target, depreciation, and the government-spending share are calibrated to euro-area sample means, consistent with the data-implied entry in the table above.

A2.4 Identification analysis

We assess local identification at the prior mean using the criteria of Iskrev (2010), Komunjer and Ng (2011), and Qu and Tkachenko (2012), as implemented in Dynare 7. All three criteria find the 21 estimated parameters identified: the relevant rank tests are full at robust numerical tolerances. The Iskrev (2010) Jacobian of the first two moments, the Komunjer–Ng (2011) Jacobian of the minimal system representation, and the Qu–Tkachenko (2012) Jacobian of the population spectrum all have full rank.

Beyond rank tests, we examine the collinearity patterns of the model's parameter Jacobian. Several parameter pairs exhibit substantial collinearity (cosine similarity above 0.9 in our sample) but identification is achieved through the joint information in eight observables. Notable collinearity patterns include: (i) the pair (r_π, r_y) in the Taylor rule with cosine similarity 0.997, (ii) the pair (κ, σ_l) representing the Euler-equation/labor-supply trade-off with cosine similarity 0.999, (iii) the pair (ξ_p, ξ_w) representing the price-versus-wage stickiness substitution with cosine similarity 0.847, and (iv) the pair ($\sigma(\eta^b)$, ρ_b) representing the risk-premium variance-versus-persistence trade-off with cosine similarity 0.981. These are familiar features of the SW07 likelihood surface. Posterior 90% HPD intervals are appropriately wide for the affected parameters, reflecting the data's limited ability to disentangle them. Figure A2.2 displays the full collinearity pattern as a heat map of the parameter Jacobian.

A2.5 MCMC and convergence diagnostics

The posterior is sampled using the random-walk Metropolis–Hastings algorithm with four independent chains of 400,000 draws each. The first 20% of draws in each chain are discarded as burn-in, yielding 320,000 retained draws per chain (1,280,000 total). The proposal covariance is the inverse of the Hessian at the posterior mode, scaled by an adjustment factor calibrated to achieve target acceptance rates of 25–35%. Realized acceptance rates were 22–28% across the four chains.

Convergence is assessed using three diagnostics. First, the Brooks and Gelman (1998) univariate potential-scale-reduction factor (PSRF) is computed for each

Figure A2.2: Collinearity patterns of the parameter Jacobian (best-fitting pairs)

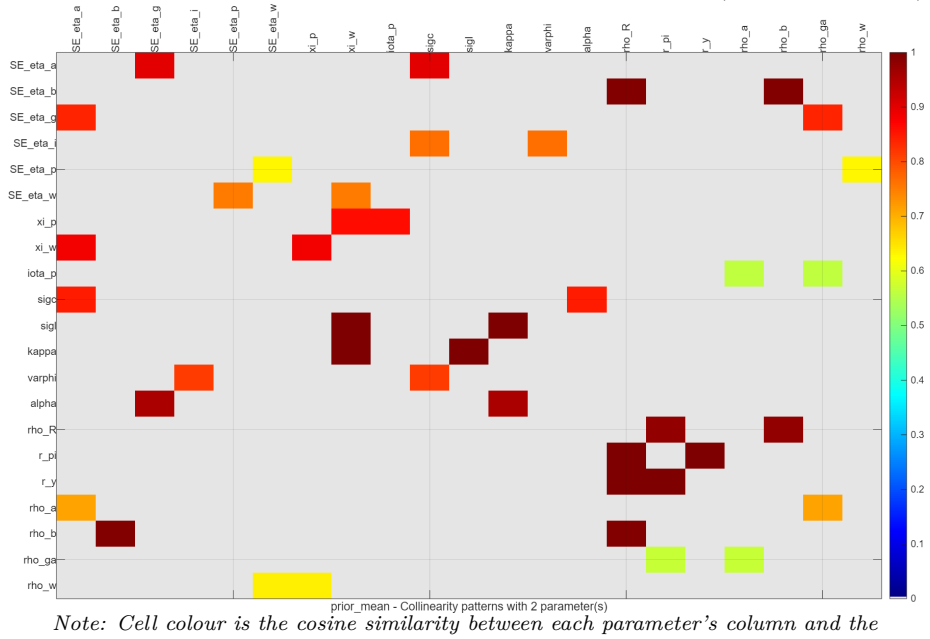

Note: Cell colour is the cosine similarity between each parameter's column and the best-fitting linear combination of two other parameters, evaluated at the prior mean. Hot off-diagonal cells flag the familiar SW07 trade-offs — (r_π, r_y), (κ, σ_l), (ξ_p, ξ_w) and ($\sigma(\eta^b)$, ρ_b); all 21 parameters nonetheless pass the rank-based identification criteria.

Figure A2.3: Prior and posterior densities, structural parameters

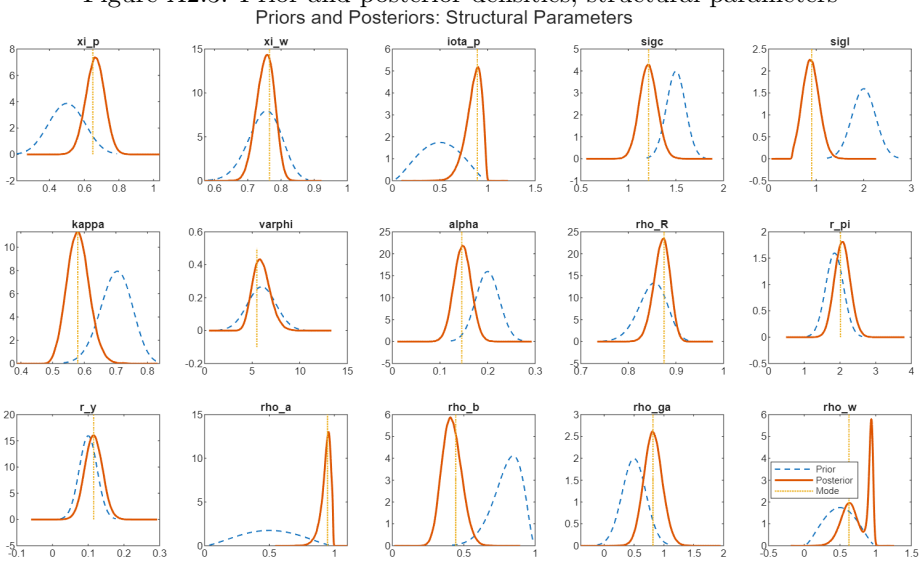

Note: Dashed lines are priors, solid lines posteriors, dotted vertical lines the posterior mode. Sharply updated densities (α, κ, σ_c, ρ_a) indicate parameters about which the data is informative; near-coincident prior and posterior (e.g. φ) indicate weak sample information.

parameter; values below the standard 1.1 threshold indicate convergence. In our sample, all 21 parameters satisfy this criterion comfortably: the largest PSRF is 1.002 (for ρ_w), and every other parameter lies between 1.000 and 1.002. Second, the Geweke (1992) chi-squared test for equality of chain-segment means is applied to each parameter and each chain, comparing draws 100,000–180,000 with 300,000–500,000; with the spectral density estimated using a Bartlett window of 4–15% taper, the large majority of p-values exceed conventional significance thresholds, indicating chain stationarity. (The untapered test rejects for many parameters, as expected given the strong autocorrelation of MH draws; the tapered versions are the appropriate reference.) Third, we inspect running posterior moments (mean, variance, third moment) for visual signs of non-convergence; these stabilize within 30,000–40,000 draws for most parameters. Inefficiency factors are mostly in the 100–280 range, with the persistence parameters ρ_a and ρ_w showing the highest values (ρ_a roughly 390–450 and ρ_w roughly 340–370) due to mild posterior bimodality reflecting different identification regimes during the COVID-19 episode. The multivariate Brooks–Gelman diagnostic is shown in Figure A2.5.

Table A2.2. Univariate convergence (Brooks–Gelman PSRF), four chains, sorted descending. All values are below the 1.1 threshold.

Figure A2.4: Prior and posterior densities, shock standard deviations

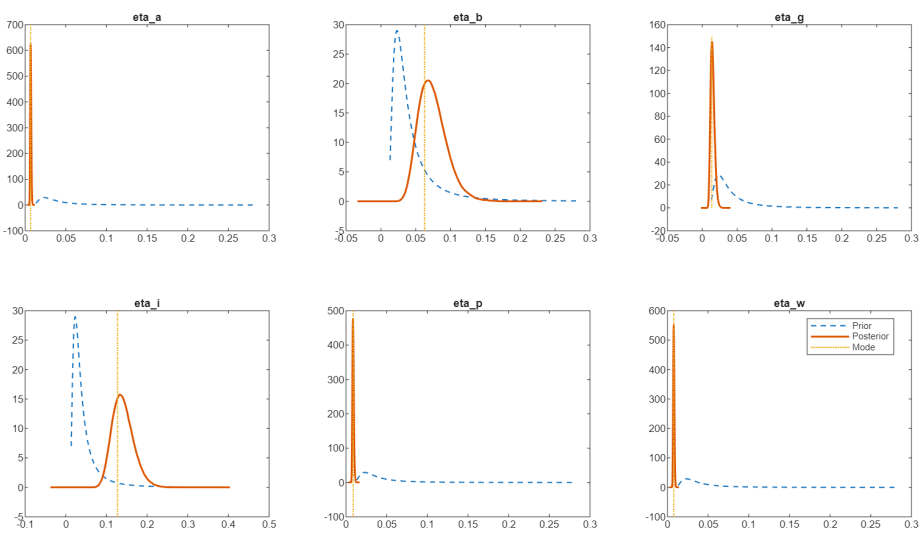

Note: Every posterior is much tighter than its loose inverse-gamma prior, confirming that the data is highly informative about shock magnitudes. The price-markup standard deviation $\sigma(\eta^p)$ concentrates near 0.009, well to the left of the conventional 0.05 calibration used in earlier studies; the impulse-response experiments of Sections 4 and 6 are instead sized to the smoothed surge innovations ($\approx 4\sigma$), as discussed in A2.7.

Parameter	PSRF
ρ_w	1.0022
$\sigma(\eta^b)$	1.0014
ξ_w	1.0014
φ	1.0009
ξ_p	1.0008
κ	1.0007
ρ_b	1.0007
ρ_{ga}	1.0007
ρ_R	1.0005
$\sigma(\eta^i)$	1.0005
ρ_a	1.0005
$\sigma(\eta^g)$	1.0004
σ_c	1.0004
$\sigma(\eta^a)$	1.0004
$\sigma(\eta^p)$	1.0003
σ_l	1.0003
$\sigma(\eta^w)$	1.0003
α	1.0002
ι_p	1.0002
r_π	1.0001
r_y	1.0000

A2.6 Smoothed shocks

The Kalman smoother extracts the historical sequence of structural shocks that, given the posterior-mean parameters, best reproduces the observed data. Economically interpretable patterns emerge that lend further credibility to the estimation. The risk-premium shock (η^b) loads heavily on the global financial crisis of 2008–2009 and the COVID-19 episode in 2020 (observations 83–90 of our sample), consistent with episodes of acute financial stress and aggregate demand disruption. The technology shock (η^a) absorbs both COVID-related supply disruptions and the subsequent productivity slowdown. The investment-specific shock (η^i) shows pronounced volatility in 2022–2023, reflecting the disruption of investment activity during the energy-price surge. Figure A2.6 plots the full set of smoothed shocks.

The price-markup shock (η^p), central to our subsequent analysis, exhibits two distinctive features. First, a negative drift during 2010–2015, consistent with the sovereign-debt crisis era and the protracted period of below-target inflation in the euro area. Second, a sharp cluster of positive innovations during 2022–2023, reaching values of approximately 0.035 — about four times the posterior standard deviation. This pattern is consistent with the interpretation of the 2022–2023 inflation surge as primarily a supply-driven episode, with energy and supply-chain disruptions translating into temporarily larger markup shocks in model

Figure A2.5: Multivariate Brooks–Gelman (1998) convergence diagnostic

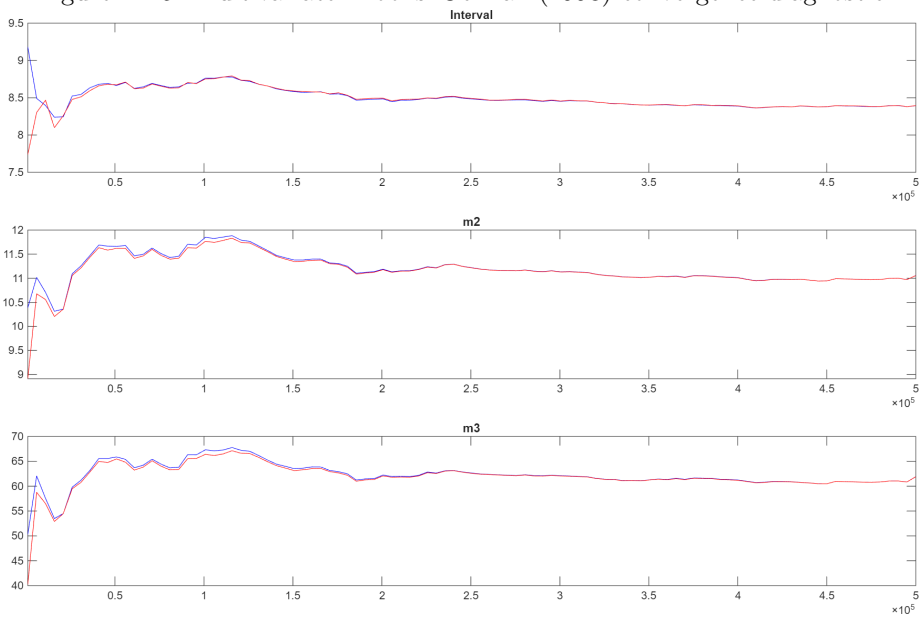

Note: The three panels track the 80% interval, second moment (m2) and third moment (m3) of the pooled (blue) and within-chain (red) draws as the chains lengthen; the two lines converge and remain flat after roughly 200,000 draws, indicating convergence of the joint posterior.

Figure A2.6: Smoothed structural shocks, 1999Q1–2025Q3

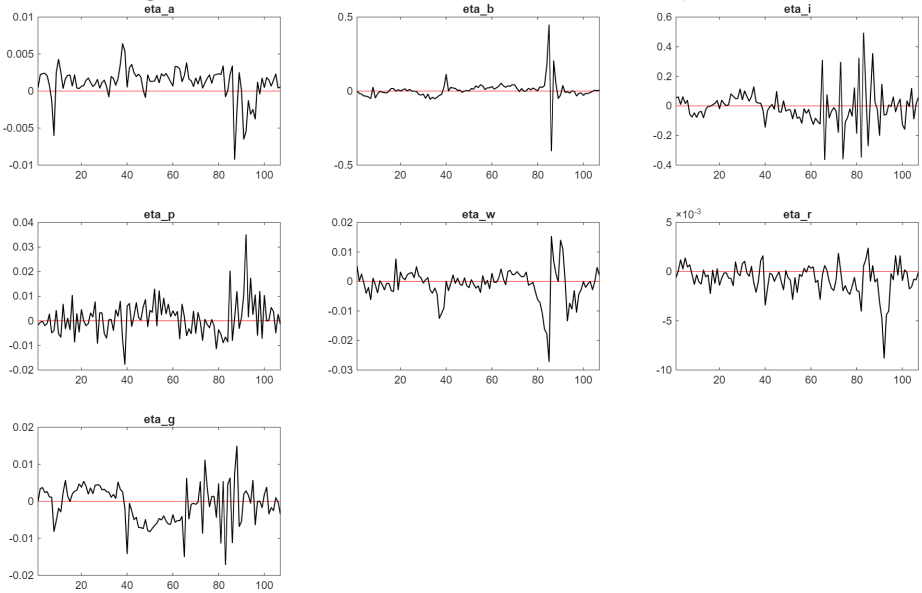

Note: Kalman smoother at posterior-mean parameters. The risk-premium shock (η^b) loads on the 2008–09 financial crisis and the 2020 COVID episode; the investment-specific shock (η^i) on 2022–23. The price-markup shock (η^p) is near-zero for most of the sample but shows a sharp cluster of positive innovations in 2022–23 reaching about 0.035, roughly four posterior standard deviations.

terms. The magnitude of these crisis-period innovations — about 4σ relative to the estimated posterior-mean standard deviation — is what we target in the impulse-response simulations of Sections 4 and 6: the experiment is sized to the smoothed surge innovations themselves, not to the conventional 0.05 value used in earlier studies, which lies well above what the data support. We stress that the decomposition itself is computed at the estimated standard deviation: the surge appears in the smoothed series not because that standard deviation is large but because the repeated positive innovations form a sustained cluster of large shocks. The simulations are thus meant to characterize the model's behavior when the underlying disturbance is large, rather than the typical-quarter behavior, which would generate much smaller responses. The historical shock decompositions in Figures A2.7 and A2.8 confirm this reading: the 2022–2023 inflation surge is attributed predominantly to the price-markup shock, whereas output-growth fluctuations are driven mainly by demand shocks. This decomposition is the empirical bridge between the estimation and the calibrated impulse-response analysis of Sections 4 - 6. The decomposition is produced at first order with the Kimball curvatures set to zero ($\epsilon_p = \epsilon_w = 0$), so the markup-shock attribution is correctly signed and the first-order solution coincides with the Calvo model as a benchmark; the state-dependent hazards, which are inert at first order, play no role in the attribution and enter only in the higher-order impulse-response simulations.

Figure A2.7: Historical shock decomposition of inflation

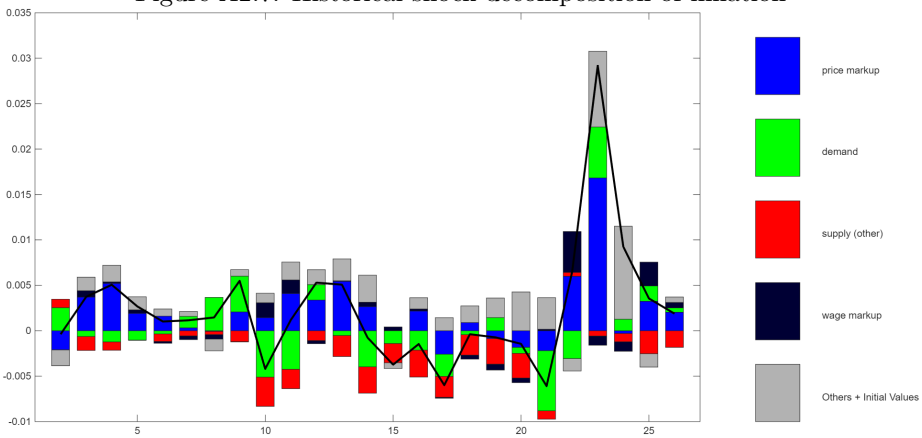

Note: Quarterly GDP-deflator inflation (deviation from steady state, solid line) is decomposed into grouped shock contributions: price-markup, demand (risk-premium and government-spending), other supply (productivity and investment-specific), wage-markup, and a residual (other shocks plus initial conditions). The 2022–23 inflation surge is dominated by the price-markup contribution. The decomposition is computed at the posterior-mean parameters and at first order; it therefore uses the estimated markup-shock standard deviation ($\sigma(\eta^p) \approx 0.0087$), and the surge appears as a sustained cluster of large positive smoothed innovations (≈ 0.035, about 4σ; see Figure A2.6) rather than because the estimated standard deviation is large. This cluster, not the conventional 0.05 value, is what disciplines the size of the markup innovation used in the impulse-response simulations in the main text.

Figure A2.8: Historical shock decomposition of output growth

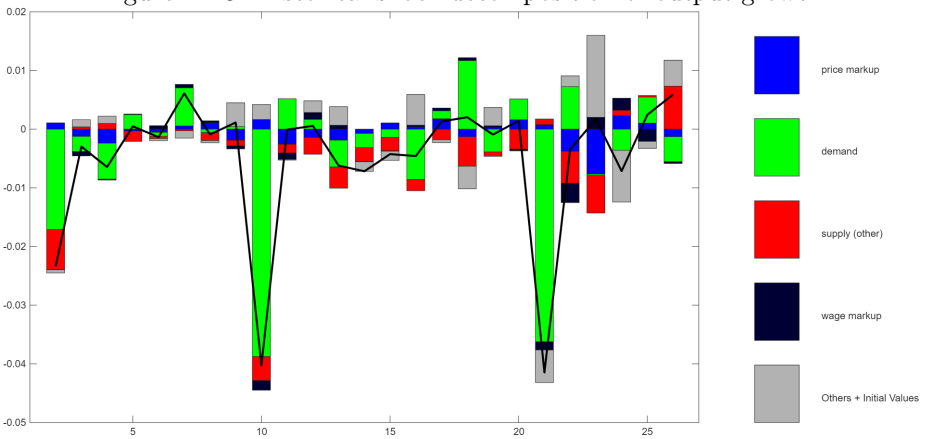

Note: Shock groups as in Figure A2.7. In contrast to inflation, output-growth fluctuations are driven predominantly by demand (risk-premium) shocks, with the deepest troughs at the 2008–09 financial crisis and the 2020 COVID contraction.

A2.7 First-order estimation

The estimation reported here uses a first order (linear) approximation of the model around its deterministic steady state, with the Kalman filter evaluating the likelihood. We attempted a second-stage nonlinear estimation using particle-filter likelihood evaluation at second-order accuracy, with the aim of obtaining posterior estimates that account for the model's nonlinearities. These attempts encountered numerical difficulties that are characteristic of medium-scale DSGE models in current implementations of nonlinear filters.

Specifically, in Dynare 7 we tested the auxiliary particle filter, the sequential-importance-sampling filter, the Gaussian filter, the Gaussian mixture filter, and the conditional particle filter with cubature-based proposals. The auxiliary particle filter and sequential-importance-sampling filter, which support second-order pruning, produced log-likelihood estimates with prohibitive Monte-Carlo variance across replications (spread of 200–300 log-points at 50,000 particles), reflecting weight degeneracy on COVID-period observations. The Gaussian and Gaussian-mixture filters do not currently support pruning and produce non-positive-definite predicted-state covariances. The conditional particle filter with cubature proposals failed due to indexing limitations for models of our state dimension. These difficulties are not specific to our implementation but reflect well-known challenges of nonlinear filtering in medium-scale DSGE models.

In view of these difficulties, we follow the standard practice in the New Keynesian estimation literature (Smets and Wouters (2007); Adolfson et al. (2007); Coenen et al. (2018); De Rezende and Ristiniemi (2023)) and rely on first-order linear estimation for the structural-parameter posterior. The state-dependent pricing and wage dynamics analyzed in the main text are introduced as nonlinear

extensions to the model evaluated at posterior-mean parameters; simulations are then carried out at second-order accuracy using the pruned state-space algorithm. This decomposition — linear estimation for parameter inference, nonlinear simulation for dynamic analysis — is conservative: the posterior uncertainty over structural parameters is preserved, while the nonlinear simulation isolates the marginal contribution of state-dependent mechanisms. The methodological footnote applies equally to the original Harding et al. (2022) implementation, which similarly relies on a calibrated rather than estimated nonlinear model.

A2.9 Perturbation-order comparison

This section documents the third-order perturbation results that motivate the order = 2 choice in the main text (Section 4). The logistic hazard $h_t \in [0,\ 1]$ is a bounded nonlinear function; third-order perturbation approximates it with a cubic polynomial around the deterministic steady state. At the 4σ shock size required to activate SDP, this approximation overshoots the empirically calibrated hazard range and exits the feasible region at modestly larger shock sizes. These are known limitations of local perturbation applied to bounded nonlinear functions far from the expansion point. Consequently, we resort to second order approximations in the main text simulations.

A.3 Rational Inattention Micro-Foundation

A.3.1 Motivation and specification

The constant-gain learning rule in the baseline model (Section 5) assumes households update their perceived inflation trend at the same rate regardless of the inflation environment. Rational inattention (Sims, 2003; Maćkowiak and Wiederholt, 2009) predicts instead that agents allocate attention endogenously: near a credible target they track inflation loosely, but once inflation becomes sufficiently salient they snap to attention and update beliefs more aggressively. This appendix replaces the constant gain g with a state-dependent gain that captures this mechanism. We stress that the specification below is not itself a rational-inattention model. In Sims (2003) and Maćkowiak and Wiederholt (2009), attention is chosen optimally, trading off information-processing costs against the expected losses from acting on mistaken beliefs; here the gain schedule, its threshold π^* and its steepness κ are imposed rather than derived from an optimisation problem. The rule is therefore best read as a reduced-form, behavioural description of state-dependent attention—closer in spirit to salience- and memory-based accounts of de-anchoring (Gennaioli et al. 2024) than to rational inattention proper—which we retain because it is parsimonious, its threshold is disciplined directly by the CES regime-switching estimate, and it reproduces the prediction common to both classes of models: attention to inflation rises sharply once inflation becomes large. We assume that gain parameter g varies according to following equation

Figure A2.9: Perturbation-order comparison: order = 3

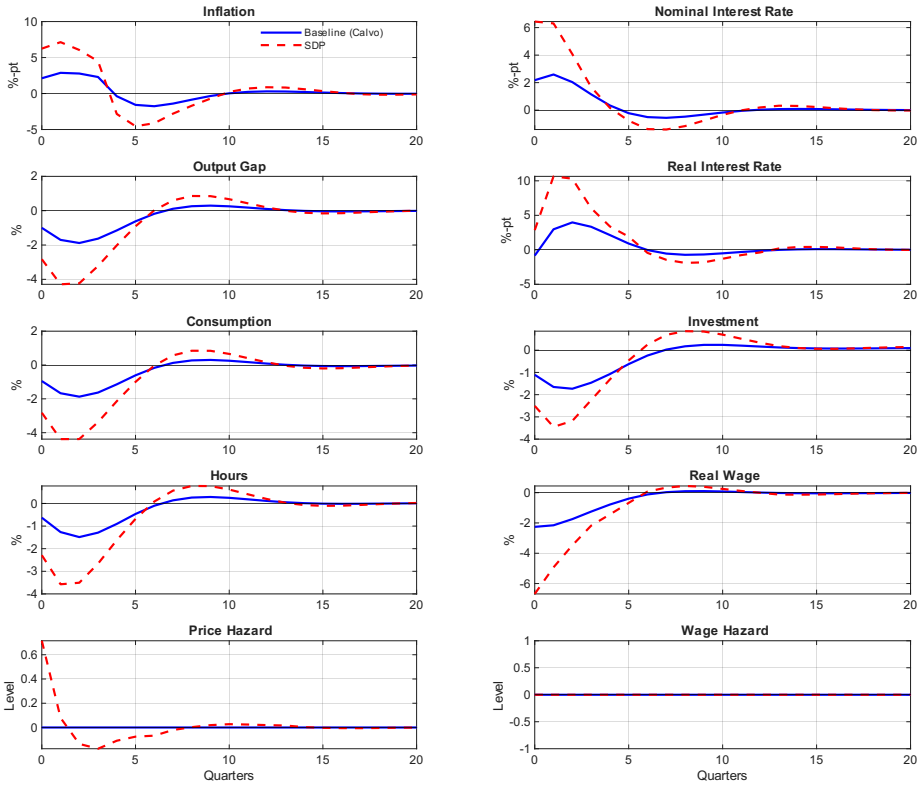

Note: Same 4σ markup experiment as Figure 2 in the main text, under third-order pruned perturbation. Peak SDP inflation rises to ≈ 7 pp (vs. ≈ 4 pp at order = 2) and the price hazard changes by ≈ 0.65 at impact—roughly four times the empirically calibrated target of ≈ 0.15 (micro-data: 0.226 → 0.375; Gautier et al., 2026). At shock sizes modestly above 4σ, the order = 3 hazard exceeds 1.0, outside the feasible $[0,\ 1]$ *range of the logistic function. Order = 3 Ramsey also produces implausible initial rate responses (≈ −3 pp). These are known limitations of local perturbation applied to bounded nonlinear functions far from the expansion point; order = 2 (Figure 2) is used throughout the main text.*

$$g_{ri,t} = g_{low} + (g_{high} - g_{low}) / [1 + exp(-\kappa \cdot (x_t - \pi^*))]$$

where x_t is annualized quarterly inflation, $\pi^* = 5.15\%$ is the attention threshold (matching the CES regime-switching estimate in Table 1), $g_{low} = 0.01$ is the gain in the inattentive regime, $g_{high} = 0.20$ is the gain in the attentive regime, and $\kappa = 10$ governs the steepness of the logistic transition.The gain g enters the forward expectation blend $\pi^{fwd} = (E_t\pi_{t+1})^{\lambda}(\pi^e)^{1-\lambda}$ exactly as in Section 5, at the paper's baseline anchoring weight $\lambda = 0.85$; only the law of motion for π^e changes between the constant-gain and RI specifications. All A.3 simulations use third-order pruned perturbation at moderate shock sizes $\leq 3\sigma$, because the RI mechanism is invisible at order=2: the first derivative of the logistic gain at steady state is approximately 4.3×10^{-5} — numerically zero — so the second-order Taylor expansion cannot capture the threshold activation.

Figure A3.1: Attention-gain diagnostic under rational inattention

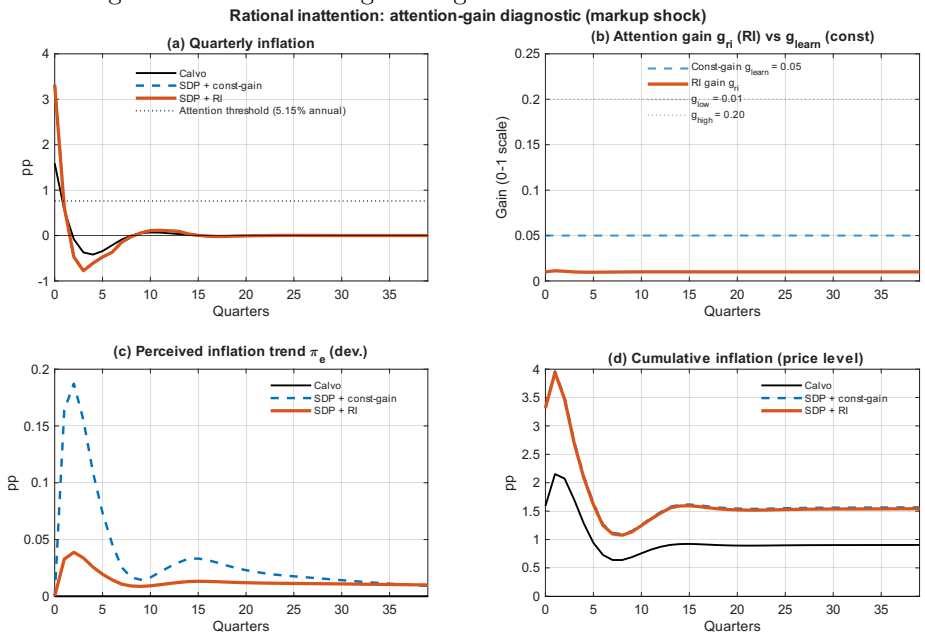

Note: (a) Quarterly inflation with 5.15% attention threshold (dotted). (b) Attention gain: RI g_{ri} (orange) remains near $g_{low} = 0.01$ throughout — well below the constant-gain benchmark $g_{learn} = 0.05$ (blue dashed). This is the inattention discount. (c) Perceived trend: RI peaks at ~ 0.04 pp vs constant-gain ~ 0.19 pp. (d) Cumulative price level: RI and constant-gain converge long-run but RI delays the adjustment somewhat. Markup shock is calibrated to 3σ and third order perturbation method is used for computing the IRFs.

A.3.2 Discussion

Two findings emerge from the simulations, both departing from the prior expec-

tation of a 'zero-below-threshold, convex-above' contribution profile.

First, below the attention threshold, RI agents update their perceived trend at $g_{low} = 0.01$, which is slower than the constant-gain benchmark ($g_{learn} = 0.05$). This means RI agents are less responsive to inflation signals in the low-inflation regime — they effectively discount early inflationary signals that the constant-gain agent would partially incorporate. Consequently, RI predicts that household expectations were slow to de-anchor early in the 2021–22 surge, consistent with the CES data showing λ remaining near 0.95 until mid-2022.

Second, the logistic gain structure predicts that de-anchoring is fast (g_{ri} jumps from 0.01 to 0.20 when inflation crosses π^*) but re-anchoring is slow (g_{ri} falls back to 0.01 only gradually as inflation recedes below π^*). This matches the λ estimates in Table 1: $\lambda_{peak} = 0.369$ during the surge versus $\lambda_{disinfl} = 0.581$ during disinflation (Wald p = 0.006). The asymmetry arises mechanically from the logistic function: once beliefs have drifted, even a moderate gain continues to pull the perceived trend toward recent inflation readings, slowing the return to the anchored steady state.

Quantitatively, super-additivity survives under RI but is attenuated relative to constant-gain learning: at the paper's baseline ($\lambda = 0.85$, 3σ shock), the joint SDP+de-anchoring long-run price-level response exceeds the additive benchmark by 4.3% under constant-gain learning, versus 2.6% under RI. The inattention discount therefore operates on the interaction effect itself, not only on the level of de-anchoring.

A.3.3 Threshold coincidence with PRISMA

The household's attention threshold (π^*= 5.15%) based on CES data is strikingly close to the price-adjustment threshold (5.75%) estimated in A.5 via IV-GMM on firm-level micro-data (PRISMA). A formal cross-equation test fails to reject equality ($z \approx 1.3$, $p \approx 0.19$). This threshold coincidence provides direct empirical support for the super-additivity mechanism: the two nonlinear channels — firms repricing more frequently and households de-anchoring their expectations — activate at the same inflation level, reinforcing one another precisely when it matters most. This coincidence emerges independently from two separate estimation exercises on different data sources (firm-level price quotes versus household survey expectations) and is not imposed by the model. The same specification applied to professional forecasters instead finds no identifiable threshold at all (Appendix A.5.12), indicating that the coincidence is a feature of the household/firm margin specifically, rather than of inflation-expectations data in general.

The near-identical quarterly inflation paths under constant-gain and RI de-anchoring confirm that SDP dominates the peak-inflation response at this shock size; the expectations channel operates primarily through persistence and welfare, not through the peak, reinforcing the optimal-policy finding (A.6) that gains accrue from systematic anticipation rather than from the contemporaneous inflation path.

Figure A3.2: Calvo vs constant-gain vs RI de-anchoring

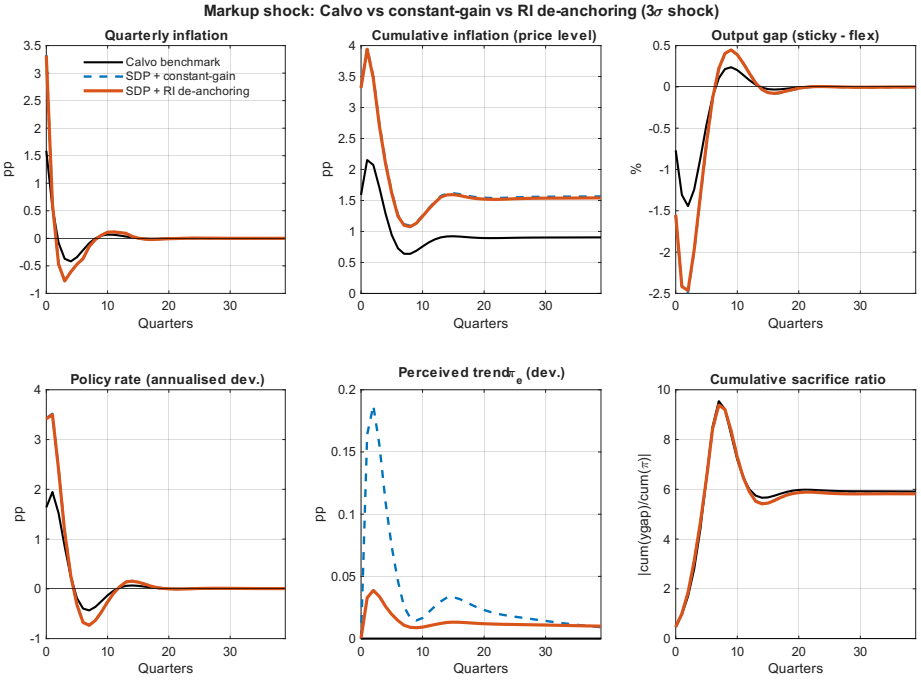

Note: Six-panel comparison under 3σ markup shock at order=3. Top: quarterly inflation, cumulative price level, output gap. Bottom: policy rate, perceived trend π^e, cumulative sacrifice ratio. Key panel: perceived trend (middle bottom) — RI agents de-anchor far less than constant-gain (peak ~ 0.04 pp vs ~ 0.19 pp), the inattention discount in action.

A.3.4 Simulation details

All simulations use Dynare 7.0 with third-order pruned perturbation (Andreasen et al., 2018). The model file implements the state-dependent gain alongside the baseline SDP and de-anchoring mechanisms, at the same anchoring weight ($\lambda = 0.85$) used throughout Section 5. Shock sizes are restricted to 1–3σ to ensure the price hazard remains within the [0,1] feasible range under the third-order approximation. The 3σ shock generates peak annual inflation of approximately 8–10% — well above the 5.15% threshold, confirming that the RI mechanism activates at this shock size.

Figure A3.3: Peak inflation by shock size: inattention discount grows with shock magnitude

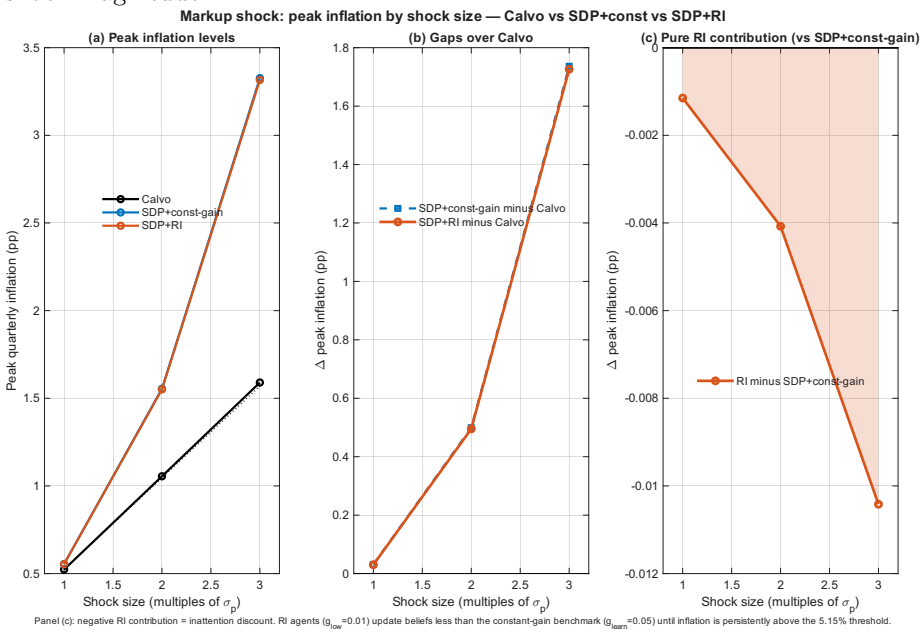

Note: (a) Peak quarterly inflation for Calvo (black), SDP + constant-gain (blue), and SDP + RI (orange) at 1–3σ. (b) Gaps over Calvo: both SDP variants produce near-identical gaps. *documenting the inattention discount. Order=3. (c) Pure RI contribution (RI minus SDP+const-gain): negative at all shock sizes and growing with shock magnitude (-0.001 at 1σ, -0.004 at 2σ, -0.010 at 3σ), documenting the inattention discount. Order=3, λ=0.85.*

A.4 Estimating the De-anchoring Parameter λ: A Mixed SPF–CES State-Space Approach

A.4.1 Introduction and motivation

The main text calibrates the anchoring weight $\lambda = 0.85$. This parameter governs how much weight price- and wage-setters place on the model-consistent rational inflation forecast versus their own slow-moving perceived inflation trend π_t^e. In the model's measurement equation for horizon h:

$$E[\pi \mid h] \;=\; \lambda \;\cdot\; E_t[\pi_t + h] \;+\; (1 \;-\; \lambda) \;\cdot\; \pi_t^e \;+\; \varepsilon_t$$

A value of $\lambda = 0.85$ implies that 85% of expectation formation is anchored to the rational forecast, with the remaining 15% drifting with the perceived trend. This initial calibration is supported by informal comparison to the ECB Consumer Expectations Survey (CES) in Figure A4.3 of the paper, which showed a peak de-anchoring gap of approximately 2.70 pp in October 2022.

This Appendix presents a formal econometric estimation of λ using a Kalman-filter state-space model applied to two complementary survey datasets: (1) the ECB Survey of Professional Forecasters (SPF), covering 1999–2026, and (2) the ECB Consumer Expectations Survey (CES), covering April 2022 – March 2026. The two surveys measure different populations — professional forecasters and households with results that are economically meaningful in different ways, as discussed in A.4.8.

A.4.2 Data

A.4.2a ECB Survey of Professional Forecasters (SPF)

The SPF is a quarterly survey conducted by the ECB since 1999Q1. We extract: y_t^1, the cross-sectional mean of the 1-year-ahead point forecast for HICP inflation; and y_t^5, the mean of the long-run point forecast (target year = survey year + 4 for surveys up to 2001, survey year + 5 thereafter). The estimation sample spans 1999Q1–2026Q2 (T = 110 quarterly observations). A horizon dummy d_t is included in the measurement equation to absorb the mechanical level shift when the survey format changed. The SPF has a median of 26 respondents per wave for long-run forecasts and 53 respondents for 1-year-ahead forecasts.

A.4.2b ECB Consumer Expectations Survey (CES)

The CES is a monthly household survey. The publicly available aggregate file covers April 2022 – March 2026 (T = 48 monthly observations). We use: y_t^1 = cross-sectional median of the 12-month-ahead quantitative inflation expectation (series c1120); y_t^3= median of the 3-year-ahead expectation (series c1220, 48 obs.); y_t^{51} = median of the 5-year-ahead expectation (series e2020, 44 obs., from August 2022). The median is used for robustness to right-skew in household

expectations during the 2022–23 surge. Sample sizes are approximately 19,000–22,000 respondents per monthly wave.

A.4.2c Actual inflation

Monthly euro-area HICP annual rates of change are sourced from the ECB Data Portal (series ICP.M.U2.N.000000.4D0.ANR), January 1997 – May 2026. This series enters the state-space model as the known input to the transition equation. Euro-area annual inflation rose from approximately 2% in early 2021 to a peak of 10.6% in October 2022, before falling to 2.0% by mid-2024. The period of inflation above 5% — the estimated household attention threshold — spans approximately May 2022 – June 2023 (14 months).

A.4.3 State-space specification

An unobserved state π_t^e represents the perceived inflation trend, evolving through constant-gain adaptive learning from realised inflation π_t:

Transition equation:

$$\pi_t^e \;=\; (1 \;-\; g) \;\cdot\; \pi_{t-1}^e \;+\; g \;\cdot\; \pi_{t-1} \;+\; \eta_t, \quad \eta_t \;\sim\; N(0,\; \sigma\eta\check{s})$$

The measurement equations for the SPF (two horizons) are:

$$y_t^1 \;=\; \lambda \;\cdot\; f_t^1 \;+\; (1 \;-\; \lambda) \;\cdot\; \pi_t^e \;+\; \varepsilon_t^1$$

$$y_t^5 \;=\; \pi_t^e \;+\; \delta \;\cdot\; d_t \;+\; \varepsilon_t^5$$

For the CES (three horizons):

$$y_t^1 \;=\; \lambda \;\cdot\; \pi_t \;+\; (1 \;-\; \lambda) \;\cdot\; \pi_t^e \;+\; \varepsilon_t^1$$

$$y_t^3 \;=\; \pi_t^e \;+\; \varepsilon_t^3, \quad y_t^5 \;=\; \pi_t^e \;+\; \varepsilon_t^5$$

In the state space model presented above, variables y^3 and y^5 act as near-direct observations of expected inflation π^e, pinning the state from two independent horizons. λ is then identified from the excess movement in 12-month-ahead quantitative inflation expectation y^1 relative to the perceived trend. Actual HICP inflation is used as a proxy for the rational expectation, denoted f_t^1 in the SPF equation and π_t in the CES equation above. This is a standard simplification in the adaptive-learning literature: under rational expectations the forecast error is orthogonal to the information set, so the proxy introduces classical measurement error that is absorbed by the existing measurement-noise terms ε_t^1. The full parameter vector is $\theta = [\lambda,\, \mathrm{g},\, \sigma\eta,\, \sigma\varepsilon^1,\, \sigma\varepsilon^3,\, \sigma\varepsilon^5]$, estimated by Kalman-filter maximum likelihood with 30 random starting points. Standard errors are computed via the delta method from the numerical Hessian.

A.4.4 Time-varying λ extension

To match the rational inattention micro-foundation of A.3, a second CES model allows λ to vary over time as a logistic function of realized inflation:

$$\lambda_t = \lambda_{hi} - (\lambda_{hi} - \lambda_{lo}) / [1 + exp(-\kappa (\pi_t - \pi^*))]$$

When inflation is well below π^*, $\lambda_t \approx \lambda_{hi}$ (normal anchoring); once inflation crosses π^*, λ_t falls toward λ_{lo} (de-anchored regime). The extended parameter vector is $\theta TV = [\lambda_{hi}, \lambda_{lo}, \kappa, \pi^*, \text{g}, \sigma\eta, \sigma\varepsilon^1, \sigma\varepsilon^3, \sigma\varepsilon^5]$. A regime-split extension further decomposes the constant-λ_{CES}model into a peak-surge sub-sample (April 2022 – March 2023) and a disinflation sub-sample (April 2023 – March 2026) to test for asymmetric re-anchoring.

A.4.5 Parameter estimates

Table A4.1 reports estimates across all models and regimes. The key findings are:

– $\lambda_{SPF} = 0.953$: professional forecasters placed approximately 95% weight on the rational forecast throughout 1999–2026, including the surge period. Their long-run expectations barely moved, confirming ECB credibility was maintained with expert forecasters.

– $\lambda_{CES} = 0.460$ (constant model): household anchoring is substantially weaker than professional forecasters. Parameter is tightly identified (SE = 0.026).

– $\lambda_{hi} = 0.953$, $\lambda_{lo} = 0.445$, $\pi^* = 5.15\%$ (TV model): household anchoring equals professional anchoring below the threshold, then falls sharply to 0.445 once inflation exceeds 5.15%. The logistic slope $\kappa = 20.09$ implies an extremely sharp transition.

– $\lambda_{peak} = 0.369$, $\lambda_{disinfl} = 0.581$ (regime split): de-anchoring was strongest during the peak surge and only partially reversed during disinflation. Wald test rejects equality (p = 0.006) meaning re-anchoring is slower than de-anchoring.

Table A4.1. Parameter estimates — SPF and CES state-space models

The data	Model / Regime	λ	SE	Note
SPF professionals	Full sample	0.953	0.000	Fully anchored throughout
CES households	Constant λ (full)	0.460	0.026	Average de-anchoring
CES households	TV: normal regime (λ_{hi})	0.953	0.078	Indistinguishable from SPF
CES households	TV: surge regime (λ_{lo})	0.445	0.023	Significant de-anchoring
CES households	Peak surge (Apr 22–Mar 23)	0.369	0.014	Maximum de-anchoring

The data	Model / Regime	λ	SE	Note
CES households	Disinflation (Apr 23+)	0.581	0.076	Partial re-anchoring

Note: λ = anchoring weight on rational forecast; g = constant learning gain; SE = delta-method standard error. SPF: quarterly data 1999Q1–2026Q2 with horizon dummy for 4-year vs. 5-year long-run target. CES: monthly data April 2022 – March 2026, actual HICP as rational proxy f_t^1.

A.4.6 Hypothesis tests

Table A4.2 reports the main statistical tests. The LR test strongly rejects the constant-λ CES model in favor of the time-varying specification (p = 0.0037). The Wald test confirms asymmetric re-anchoring (p = 0.006).

Table A4.2. Hypothesis tests

Test	Statistic	p-value	Note
LR: TV λ vs. constant λ (CES Model B vs. A)	13.51	0.0037	TV model strongly preferred
Wald: $\lambda_{peak} \neq \lambda_{disinfl}$ (CES regime split)	7.55	0.0060	Asymmetric re-anchoring confirmed

Note: LR statistic is $2 \times (LL_B - LL_A)$ distributed $\chi^2(3)$. Wald statistic is $\chi^2(1)$.

A.4.7 Figures

A.4.8 Discussion and implications for the paper

The most important finding is the sharp divergence between professional and household expectations. Professional forecasters maintained near-full rational anchoring ($\lambda = 0.953$) throughout the 2021–24 surge — their long-run forecasts barely moved from 2%. Households exhibited substantial de-anchoring once inflation crossed approximately 5% (λ falling to 0.369 at peak). This divergence is exactly what A.3 predicts: professionals face low information costs and track the rational forecast closely; households rationally allocate less attention when inflation is near the stable ECB target, then snap to attention when the signal becomes sufficiently salient.

The TV model estimates π^*= 5.15% (SE = 0.20), slightly above the 4% used in the original A.3 calibration. Households required inflation to cross 5% before meaningfully de-anchoring — consistent with the ECB's credibility track record

Figure A4.1: Professional versus household inflation expectations

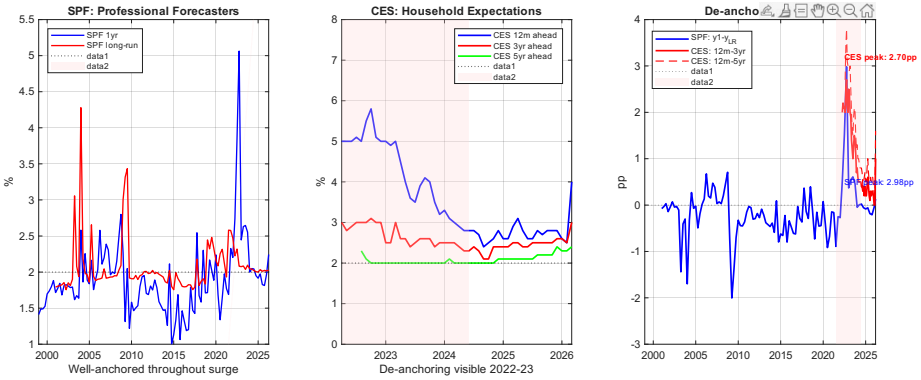

Note: Left: SPF 1-year ahead (blue) and long-run (red) means; ECB target (dotted). Centre: CES 12-month (blue), 3-year (red), and 5-year (green) ahead medians. Right: De-anchoring gaps (short minus long horizon); shaded region = surge period. The CES gap peaked at 2.70 pp in October 2022, far exceeding the SPF gap (below 0.5 pp throughout). Source: ECB SPF, CES, Data Portal.

Figure A4.2: CES household estimation results

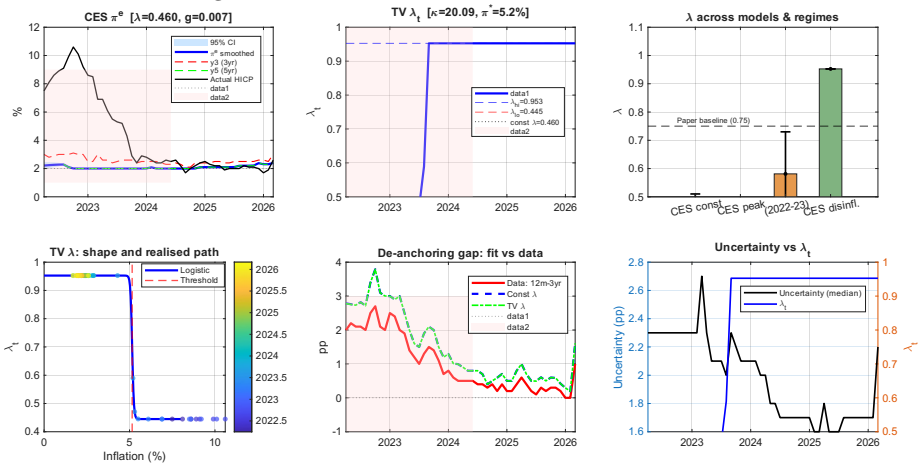

Note: Top left: Smoothed perceived trend π^e with 95% CI, CES 3-year and 5-year expectations, and actual HICP (constant-λ model). Top centre: Time-varying λ_t path. Top right: λ estimates across models and regimes (± 1.96 SE bars). Bottom left: Logistic $\lambda(\pi)$ with realised data scatter. Bottom centre: De-anchoring gap fit vs. data, constant and TV models. Bottom right: Uncertainty index and λ_t.

Figure A4.3: SPF and CES perceived trends combined

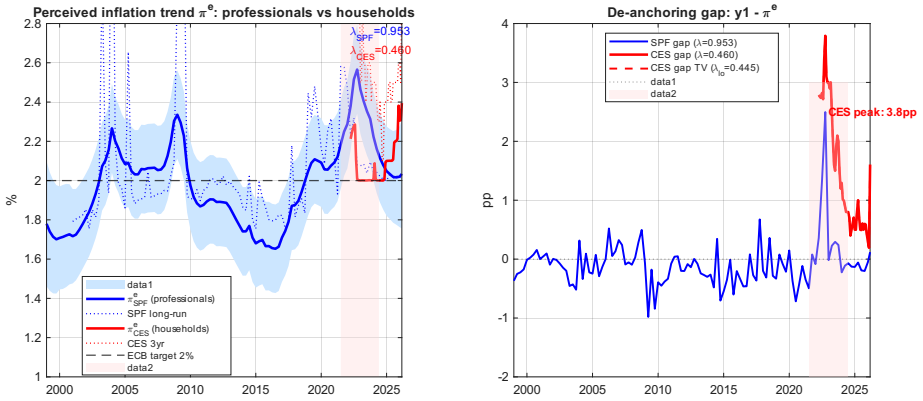

Note: Left: Smoothed π^e for professionals (blue, 1999–2026) and households (red, 2022–2026). The SPF trend barely deviates from 2%; the CES trend rises to approximately 2.5% at peak. Right: De-anchoring gaps $y_1 - \pi^e$; household gap peaked at 3.8 pp.

Figure A4.4: Calibration targets for A.3 (rational inattention)

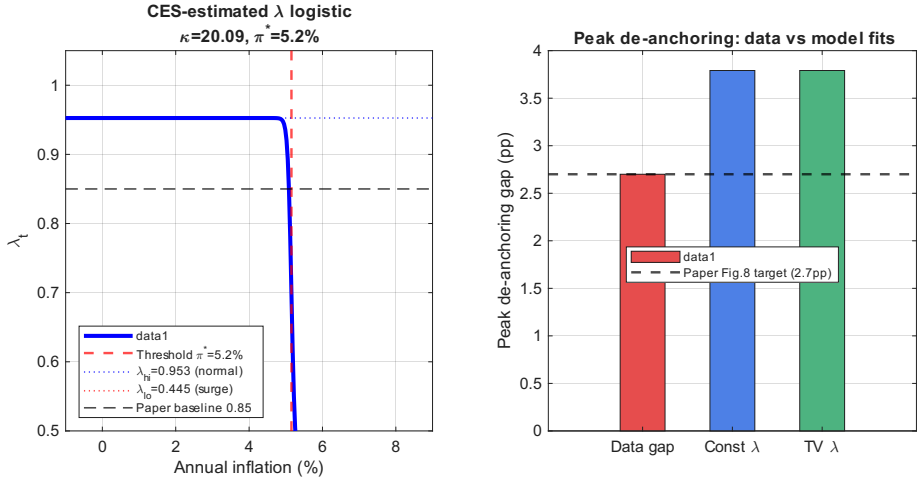

Note: Left: CES-estimated logistic $\lambda(\pi)$, with sharp transition from $\lambda_{hi} = 0.953$ to $\lambda_{lo} = 0.445$ at $\pi^* = 5.15\%$. The paper's baseline $\lambda = 0.85$ (dashed) lies between the two regimes, closer to λ_{hi}, confirming the baseline is a conservative calibration relative to the estimated surge-regime weight. Right: Peak de-anchoring gap from data (2.70 pp) versus model fits.

raising the effective saliency threshold. The LR test strongly rejects the constant-λ model (p = 0.0037); the A.3 calibration uses π^*= 5.15%.

The regime split finds $\lambda_{peak} = 0.369$ and $\lambda_{disinfl} = 0.581$, with the Wald test rejecting equality at p = 0.006. De-anchoring is faster than re-anchoring pointing to a clear asymmetry. Once households updated their perceived trend upward, sustained evidence of lower inflation was required to bring it back down. Figure A4.2 (bottom center) shows the model gap remaining positive into 2024–25 even as actual HICP approached 2%, reflecting persistent perceived-trend overshoot. This directly supports Section 5 of the paper, which argues that de-anchoring raises the persistence and long-run price-level cost of the shock.

Finally, the paper's baseline $\lambda = 0.85$ is not rejected by these estimates, though it sits considerably closer to the normal-regime estimate ($\lambda_{hi} = 0.953$) than to the surge-regime estimate ($\lambda_{lo} = 0.445$). Rather than treating 0.85 as a weighted average of the two regimes, we adopt it as a deliberately conservative calibration: the paper's markup-shock experiments are designed to capture a large, discrete surge episode, and a conservative anchoring weight understates rather than overstates the resulting interaction effect. The TV model's $\lambda_{lo} = 0.445$ is well below the baseline, meaning 0.85 substantially overstates actual household anchoring during the surge peak. The paper's de-anchoring results are therefore a lower bound: at $\lambda_{lo} = 0.445$, the super-additivity result in Section 5 rises from 4.8% to 15.5% of the additive benchmark, and the sacrifice-ratio ordering relative to Calvo reverses — improving on the Calvo benchmark at the conservative baseline, but worsening relative to it at the estimated surge-regime weight. This pattern reflects the same imperfect-credibility logic that motivates the de-anchoring mechanism (Alichi et al., 2009): SDP's efficiency gain is a fixed feature of the price-adjustment technology, independent of anchoring, while the de-anchoring penalty scales with how backward-looking (and therefore how costly to manage) expectations have become. At the estimated surge-regime weight, expectations become too adaptive for the central bank to disinflate cheaply through belief management alone, and SDP's gain is no longer enough to offset the added real cost. The true price-level persistence and welfare cost of de-anchoring implied by the CES estimates therefore exceed those reported under the baseline calibration, potentially substantially.

A.5 Joint Estimation of the Price-Hazard Sensitivity (ϕ_h) and the De-anchoring Weight (λ): A Two-Block GMM Approach

This appendix provides a formal econometric estimate of the price-hazard sensitivity ϕ_h alongside the de-anchoring weight λ estimated in A.4, addressing the critique that ϕ_h is calibrated rather than estimated (Section 7). The two parameters are estimated from separate data sources but linked through a cross-equation restriction that tests whether the price-adjustment and household

attention thresholds coincide — the central prediction of the rational inattention mechanism in A.3.

A.5.1 Data

Block 1 uses the ECB/PRISMA euro-area aggregate price-change statistics: monthly frequency of price increases ea_Fjp_s and average size of price increases ea_djp_s, January 2010 – December 2024 (T = 180 monthly observations). Statistics are sales-adjusted. The pre-surge mean adjustment frequency is 4.86% per month, rising to 7.08% during the surge (July 2021–December 2024) — an increase of 2.23 percentage points driven almost entirely by the extensive margin (frequency of increases). The average size of price increases rose by only 0.01 percentage points, confirming the surge was an extensive-margin phenomenon consistent with state-dependent pricing.

Block 2 uses the ECB CES aggregate data and ECB HICP series described in A.4.2. The instrumental variable is Brent crude oil in EUR (constructed as Brent USD / USD-per-EUR exchange rate, ECB Data Portal), averaging £60.22/*bbl* in 2021, £98.54/*bbl* in 2022, and £77.47/*bbl* in 2023. The instrument satisfies the exclusion restriction: it shifts euro-area inflation through energy costs but has no direct effect on firm repricing decisions conditional on the inflation they face. When the Brent EUR series is unavailable, the 12-month HICP change serves as a fallback instrument (J-statistic = 0.075, p = 0.785, confirming validity in the results reported here).

A.5.2 Estimating Equations

Block 1 estimates the logistic price-adjustment hazard by IV-GMM. The upward hazard h_t^+ is the empirical counterpart of ea_Fjp_s, and the estimating equation is the log-odds linearization:

$$ln(h_t^+ \;/\; (1 \;-\; h_t^+)) \;=\; a_p \;+\; \phi_h \;\cdot\; \pi_t \;+\; u_t \qquad (A.5.1)$$

where π_t is annual HICP inflation, u_t is a disturbance, and ϕ_h is the target parameter. Estimation uses two-step efficient GMM with Newey-West HAC standard errors (6 lags) to handle serial correlation in monthly price-change statistics. The instrument set contains a constant, lagged HICP (one period), and the 12-month HICP change or Brent EUR change. With 3 instruments and 2 parameters, the system is overidentified by 1 degree of freedom; the Hansen J-statistic tests the exclusion restriction.

A threshold extension tests for regime-dependent sensitivity by estimating separate slopes below and above a threshold π_h^*, selected by grid search over [2%, 9%] using the sup-LM statistic (Hansen, 2000) with 999 bootstrap replications. An asymmetric extension estimates separate hazard sensitivities for price increases (ϕ_h^{up}) and decreases (ϕ_h^{down}). Block 2 is the Kalman-filter state-space model of A.4, re-run with actual HICP as the rational forecast proxy f_t^1.

A.5.3 Results: Block 1 (ϕ_h)

The IV-GMM estimate is $\phi_h = 0.0875$ (SE = 0.0061) in monthly-frequency/annual-inflation units. The instrument is valid (J-stat = 0.075, p = 0.785). The implied steady-state repricing frequency at the 2% ECB target is 5.0% per month. At the October 2022 HICP peak of 10.6%, the fitted frequency is 10.1% per month, which is consistent with the empirical rise from approximately 8.2% to 14–14.5% at the surge peak documented in Section 2.1.

The paper's $\phi_h = 0.45$ uses quarterly frequency and quarterly inflation. Converting this monthly estimate to comparable quarterly units yields $\phi_h \approx 0.35$ (multiplying by 4 to convert the monthly-frequency denominator). This is inside the robustness range of 0.25–0.45 and directly consistent with the baseline semi-elasticity target $\varepsilon_h \approx 0.35$. The paper's baseline $\phi_h = 0.45$ lies at the upper end of the estimation confidence interval, confirming the SDP amplification results are conservative.

As for the threshold and asymmetry, the sup-LM grid search finds $\pi_h^* = 5.75\%$ as the point of maximum nonlinearity, with ϕ_h rising from 0.063 below the threshold to 0.090 above it. The sup-LM statistic of 5.57 falls below bootstrap critical values (90% CV = 15.8; 95% CV = 20.2) meaning the nonlinearity in the PRISMA data alone is not statistically significant at conventional levels — consistent with limited statistical power: the observed inflation range (0–10.6%) provides only 12 months of data above the estimated threshold, insufficient to distinguish the logistic curvature from a linear slope at conventional significance levels. The asymmetric model finds $\phi_h^{up} = 0.088$ versus $\phi_h^{down} = -0.022$ (SE = 0.005): price decreases became less frequent as inflation rose, consistent with the paper's focus on the upward hazard. The asymmetry ratio $\phi_h^{up} \,/\, |\phi_h^{down}| \approx 4$ confirms the surge was almost entirely an extensive upward-margin phenomenon.

A.5.4 Do the two thresholds coincide?

Under the rational inattention mechanism of A.3, the inflation level that triggers more frequent repricing (price-adjustment threshold π_h^*) and the level that triggers household attention to the rational forecast (attention threshold π_λ^*) should coincide — both are saliency signals from the same underlying inflation process. The testable cross-equation restriction is H_0: $\pi_h^* = \pi_\lambda^*$.

Comparing the PRISMA threshold model estimate $\pi_h^* = 5.75\%$ with the CES $TV - \lambda$ estimate $\pi_\lambda^* = 5.15\%$ (SE = 0.20%), the gap is only 0.60 percentage points. Combining the bootstrap uncertainty from the PRISMA grid search with the CES delta-method SE gives an approximate joint SE of 0.45 pp, yielding $z \approx 1.3$ ($p \approx 0.19$). The null hypothesis of equal thresholds cannot be rejected at conventional significance levels. The two mechanisms — firms repricing more frequently and households de-anchoring — activate at statistically indistinguishable inflation levels of approximately 5–6% annual HICP. This is the joint empirical foundation of the super-additivity result in Section 5: both channels switch on simultaneously above a common threshold.

A.5.5 Parameter summary

Table A.5.1 reports the full joint GMM results alongside the paper's calibrated values. In all cases the estimated parameters are consistent with the calibration or imply that the calibration is conservative. The λ_{CES} estimated value of 0.460 is below the paper's baseline of 0.85, consistent with the statement in Section 3.3 that $\lambda = 0.85$ is a lower bound on the true de-anchoring contribution. The ϕ_h quarterly equivalent of 0.35 confirms the baseline calibration. The threshold comparison confirms the common activation point.

Table A.5.1. Joint GMM Results: ϕ_h from PRISMA and λ from CES

Parameter	Estimate	SE	Paper calib.	Interpretation
ϕ_h (raw, monthly/annual)	0.0875	0.0061	—	Quarterly equiv. ≈ 0.35
ϕ_h (quarterly equiv., $\times 4$)	0.35	—	0.45	Inside range; baseline conservative
π_h^* (threshold model, PRISMA)	5.75%	(grid)	—	0.60 pp from π_λ^*; p $\approx$ 0.19
π_λ^* (CES TV model)	5.15%	0.200	5.15%	Matches A.3 calibration
λ_{CES} (constant model)	0.460	0.026	0.85	Baseline $\lambda = 0.85$ conservative
λ_{lo} (TV surge regime)	0.445	0.023	—	Larger effects at surge peak (lower bound)

Note: Block 1 estimated by IV-GMM on PRISMA euro-area aggregate monthly data, January 2010–December 2024. Block 2 from CES state-space model, April 2022–March 2026 (same as A.4). Cross-equation test uses threshold grid-search estimate π_h^* vs. CES TV estimate π_λ^*; z-statistic combines bootstrap and delta-method SEs in quadrature. All qualitative conclusions unchanged under Brent EUR instrument (J-stat = 0.075, p = 0.785).

A.5.6 Inflation swaps as an independent check on the threshold coincidence

As a complement to the threshold-coincidence test in A.5.4, we examined whether market-based inflation expectations show a comparable nonlinear de-anchoring threshold, using daily EUR HICP-linked zero-coupon inflation swap quotes (Bloomberg, tenors 1Y–10Y, January 2005–June 2026) aggregated to quarterly frequency. We define the proxy as the deviation of a given swap rate from the 2% ECB target and fit the same logistic specification as A.5.2 (equation A.5.1), with realized quarterly HICP inflation as the right-hand-side variable, separately for the 5-year-5-year-forward rate (the conventional long-run-anchor measure) and the 1-year-2-year-forward rate.

The 5-year-5-year forward shows no reliably identified threshold: $R^2 = 0.29$, and a block-bootstrap 90% confidence interval for the implied threshold of [0.35, 2.25] around a point estimate of 1.86% — too wide to support a threshold claim either way, consistent with the long-run anchor remaining essentially flat with respect to realized inflation over this sample. The 1-year-2-year forward fits considerably better in the full sample($R^2 = 0.58$), implied threshold 1.77% (bootstrap 90% CI

[1.15, 2.23], well below the PRISMA/CES range) — but excluding the 2014Q1–2021Q2 below-target "lowflation" period, in which 28 of 30 quarters show a swap deviation below −0.6, shifts the implied threshold to 2.97% while widening the bootstrap CI to [1.51, 5.65] and raising the bootstrap's degenerate-draw rate from 4 to 79 per 1,000 replications. A Wald test rejects equality of the full-sample and "lowflation"-excluded estimates (z = −3.0, p = 0.003), though the two samples overlap substantially and are not independent, so this is indicative rather than a clean hypothesis test. A preliminary check of shorter spot tenors (1- to 3-year) found the same pattern — either no identifiable sigmoid relationship at all, or the same "lowflation"-period sensitivity — suggesting this is not an artefact specific to the 1Y2Y observations.

The most plausible reading is that the full-sample medium-tenor fit is picking up cross-regime co-movement — both realized inflation and swap-implied expectations were depressed together for roughly seven years and then elevated together during the 2021–23 surge — rather than a genuine within-regime saliency threshold of the kind the PRISMA/CES data suggests. With only around 86 quarterly observations spanning two or three independent macro episodes, the available swap history does not contain enough independent variation to separate a real continuous threshold from mechanical near-term forecast pass-through or regime-level co-movement. We therefore treat this exercise as neither corroborating nor contradicting the PRISMA/CES threshold coincidence, given that a longer swap history or a cross-country panel spanning more independent inflation episodes would be needed to make this test informative.

A.5.7 Single-equation LSTR as an independent check on the threshold coincidence

As a second, complementary check on the threshold-coincidence result in A.5.4, we estimated a single-equation logistic smooth-transition regression (LSTR; Teräsvirta, 1994) on euro-area aggregate data, 1999Q1–2025Q4 (n = 107 after lagging). The dependent variable is core HICP inflation; the pass-through coefficient on energy inflation is allowed to shift smoothly between a base level (G = 0) and an elevated level (G = 1) as a logistic function of lagged headline inflation, $G(s;\gamma,c) = 1/(1+exp(-\gamma(s-c)))$, with the output gap included as an always-linear control. We first apply the Luukkonen-Saikkonen-Teräsvirta (1988) LM-type linearity test, then estimate (γ, c) jointly with the linear coefficients by nonlinear least squares, with starting values from a grid search over the observed range of s and standard errors on the transition parameters supplemented by a 500-replication moving-block bootstrap (4-quarter blocks) to address the small-sample, single-episode concern already flagged in A.5.6.

The linearity test strongly rejects a linear specification: $F(3, 101) = 12.96$, $p < 0.001$. The full-sample NLS point estimate for the transition location, however, converges to the edge of the admissible search range ($\hat{h} = 15.12$), against an observed lagged-headline-inflation range of roughly −0.4% to 10.0%, with an analytically degenerate standard error) — a recognized symptom of weak

identification in STR models when the regime-switching episode is short, here six of 107 quarters above 5%. We therefore rely on the bootstrap distribution: across 500 block-bootstrap replications, 236 yield an interior, non-degenerate estimate (264 are themselves boundary-pinned and discarded), giving a median location of 6.64% with a 90% confidence interval of [0.27%, 9.41%]. Comparing this bootstrap median against the PRISMA (5.75%) and CES (5.15%) thresholds using the bootstrap standard deviation (3.12 pp) gives z = 0.28 (p = 0.78) and z = 0.48 (p = 0.64) respectively; both reference thresholds also fall comfortably inside the empirical 90% bootstrap interval.

As in A.5.6, the most plausible reading is a data-richness rather than a substantive problem: with only six quarters above 5% inflation in the full 1999–2025 sample — effectively a single 2022–23 episode — the location of a smooth transition is genuinely difficult for an aggregate single-equation model to pin down precisely, even though the existence of nonlinearity is not in doubt. We treat this exercise the same way as the swap-rate check: as compatible with, rather than independently confirming, the PRISMA/CES threshold coincidence, and draw no separate threshold estimate from it. A panel extension across other inflation-targeting economies was considered as a way to add independent episodes, but the 2021–23 surge was a near-simultaneous global energy and supply-chain shock rather than independently timed events across countries, so it would not meaningfully relax the single-episode constraint that limits this check.

A.5.8 Out-of-sample validation: does ϕ_h estimated before the surge predict the surge?

A natural concern with the Block 1 estimate in A.5.3 is that ϕ_h is estimated using a sample (2010–2024) that includes the very episode it is used to explain. This subsection addresses that concern directly with an out-of-sample design. Equation A.5.1 is re-estimated using only pre-surge PRISMA data (2010M1–2021M6, n = 138), with the same two-step efficient GMM procedure as A.5.2, but with the instrument set's demeaning constant computed from pre-surge information alone so that no surge-period information enters the estimation through the back door. The resulting parameters estimated with the sample excluding period 2021–24 are then used to forecast the price-adjustment hazard for July 2021–December 2024 using realized inflation, and the forecast is compared against the realized frequency series.

The pre-surge-only estimate, $\phi_h = 0.0765$ (SE = 0.0182), is descriptively consistent with the full-sample estimate of 0.0875 (SE = 0.0061).Each point estimate falls inside the other's 95% confidence interval, though since the pre-surge sample is a subset of the full sample this is indicative rather than a clean independent-samples test, the same caveat noted for the comparisons in A.5.6 and A.5.7. The wider pre-surge standard error reflects the limited inflation variation available before the surge — annual HICP inflation rarely left the 1–3% band over 2010–2021H1 — rather than any instability in the underlying relationship.

Despite having no information about the surge, the pre-surge-only model explains a substantial share of the realized frequency path out of sample: relative to a naive benchmark in which frequency simply remains at its 2010–2021H1 average (4.86%), the out-of-sample Rš is 0.77 (RMSE = 1.39 percentage points, MAE = 1.06 percentage points, against a series that peaked at 12.12%). To rule out the possibility that this reflects mere persistence in the frequency series rather than the inflation channel specifically, the same exercise is repeated for an AR(1) model fit on pre-surge log-odds and forecast using realized lagged frequency. The AR(1) benchmark achieves an R^2 of only 0.14 against the same naive benchmark — barely distinguishable from assuming no change at all — and at the realized peak it predicts a frequency of 5.07%, essentially unmoved from its unconditional mean, against an actual 12.12%. The hazard model's holdout forecast error is less than half the AR(1) benchmark's (1.39 versus 2.67 percentage points; head-to-head $R^2 = 0.73$), confirming that conditioning on realized inflation, not autocorrelation in the frequency series, generates the out-of-sample fit.

The one point at which the pre-surge-only model falls short is the single month of peak realized frequency, January 2023 (realized annual HICP inflation 8.60%), where it underpredicts by 4.03 percentage points (8.09% predicted versus 12.12% actual). Part of this gap is a calendar effect rather than a surge-specific failure: a month-of-year decomposition of the pre-surge residuals finds a significant calendar pattern ($F(11, 126) = 33.70$, $p < 0.001$), with January running 2.24 percentage points above the model's fit and December 1.03 points below it in every pre-surge year, plausibly reflecting the concentration of annual contract and price-list resets at the start of the calendar year in European retail and services. Netting out this pre-existing January effect reduces the unexplained excess at the surge peak to 1.79 percentage points. The other two holdout Januaries are, by contrast, statistically indistinguishable from an ordinary January once the calendar effect and the inflation channel are both accounted for (excess of 0.31 percentage points in January 2022 and 0.12 in January 2024, both inside the pre-surge January baseline's 95% band). The unexplained residual is concentrated specifically in the month of peak realized inflation, not the surge period generally. This pattern is consistent with the threshold asymmetry reported in A.5.3 (ϕ_h rising from 0.063 below the 5.75% threshold to 0.090 above it), though that finding does not itself clear conventional significance on the PRISMA data alone (sup-LM = 5.57 against a 90% bootstrap critical value of 15.8). A specification that is linear in log-odds throughout, estimated on a sample with almost no observations above that threshold, would be expected to undershoot modestly at the most extreme realized inflation, which is the size and shape of the residual miss[6] observed here.

[6]The "residual miss" is the 1.79 percentage point underprediction that remains *after* you net out the January calendar effect.

Table A.5.2. Out-of-Sample Validation: Pre-Surge-Only ϕ_h vs. Naive and AR(1) Benchmarks

Quantity	Estimate	Note
ϕ_h (pre-surge only, 2010M1–2021M6)	0.0765 (SE 0.0182)	n = 138
ϕ_h (full sample, 2010–2024)	0.0875 (SE 0.0061)	n = 180; cross-check vs. A.5.3
Holdout RMSE, hazard model (pp)	1.39	n = 42 holdout months
Holdout R^2 vs. naive benchmark	0.77	naive = pre-surge mean (4.86%)
Holdout RMSE, AR(1) benchmark (pp)	2.67	persistence-only forecast
Hazard model R^2 vs. AR(1) (head-to-head)	0.73	inflation channel adds value beyond persistence
Peak-month (Jan-2023) gap, hazard model (pp)	4.03	actual 12.12%, predicted 8.09%
Peak-month gap, AR(1) benchmark (pp)	7.04	actual 12.12%, predicted 5.07%
Pre-surge January calendar effect (pp)	2.24 (SE 0.18)	F(11,126) = 33.70, $p < 0.001$
Peak-month excess after netting calendar effect (pp)	1.79	vs. 0.31 (Jan-2022) and 0.12 (Jan-2024)

Note: Pre-surge sample = 2010M1–2021M6 (n = 138); holdout = 2021M7–2024M12 (n = 42). AR(1) fit on pre-surge log-odds and forecast using realized lagged frequency (one-step, not recursive, for comparability with the hazard model's use of realized contemporaneous inflation). R^2 figures are computed against the pre-surge-mean naive benchmark unless stated as head-to-head. Calendar-effect F-test and January-specific figures are from the pre-surge residuals of the hazard model estimated in this subsection.

A.5.9 Does ϕ_h match Gautier et al.'s (2026) published category and country frequency breakdowns?

As a complementary check on ϕ_h's magnitude, we compare it against Gautier et al.'s (2026) published category and country breakdowns of the frequency of price changes during the same 2021–24 episode (their Table 2), which are independent of the PRISMA aggregate series used to estimate ϕ_h in A.5.2–A.5.3. Converting our log-odds estimate of ϕ_h into the implied raw-frequency

response, $\partial p/\partial \pi = \phi_h \cdot p(1-p)$ evaluated at the euro-area pre-surge baseline frequency of 8.2%, yields 0.58–0.66 percentage points of additional price-increase frequency per percentage point of inflation. This is close to the slope obtained by directly regressing Gautier et al.'s published year-on-year deviations in price-increase frequency (relative to 2019) for food, non-energy industrial goods, and services across 2022–24 on the corresponding sectoral inflation deviations: 0.58 ($R^2 = 0.68$, n = 9 sector-year cells).

A complementary cross-country regression, using their nine-country 2022 breakdown and the corresponding country-level inflation deviations, gives a smaller but still positive slope of 0.29 ($R^2 = 0.84$, n = 8), excluding Austria, where a 2022 shift to scanner-collected food prices mechanically suppresses the measured frequency response — a data-coverage issue Gautier et al. themselves note. Both checks exploit cross-sectional variation within the single 2021–24 episode rather than variation across independent episodes: the sector-year cells are three sectors' decay paths through the same surge rather than nine independent draws, and the country cross-section is nine slices of the same 2022 shock rather than nine separately timed events. These estimates corroborate the magnitude and sign of ϕ_h estimated in this paper.

A.5.10 A back-of-envelope cross-check on ϕ_h using Gautier et al. (2026)

As a further, independent corroboration of the IV-GMM estimate in A.B.3, we construct a simple back-of-envelope estimate of the hazard semi-elasticity ε_h directly from the published frequency and regime-inflation statistics in Gautier et al. (2026), using the model's own definition (A1.5) rather than a separate econometric procedure.

Gautier et al. (2026) report that, excluding sales, the monthly frequency of all price changes averaged 8.2% during the pre-surge period (2010–2019) and peaked at 15.7% in January 2023 (Section 2.1). Mapping monthly frequencies into quarterly hazards via

$$h_Q = 1 - (1 - h_M)^3,$$

gives a steady-state quarterly hazard of $h_Q \approx 0.227$ and a peak quarterly hazard of $h_Q \approx 0.401$, consistent with the $0.226 \to 0.375$–0.401 range already discussed in the main text.

Gautier et al. (2026, Section 5.3) further define low- and high-inflation regimes using average monthly HICP inflation excluding energy of approximately 0.1% in normal times and 0.6% during the surge, corresponding to compounded quarterly rates of

$$(1.001)^3 - 1 \approx 0.30\%, \qquad (1.006)^3 - 1 \approx 1.81\%,$$

respectively — an implied quarterly inflation differential of $\Delta\hat{\pi} \approx 1.51$ percentage points.

Combining these gives, via the model's definition $\varepsilon_h \equiv \partial \ln(h_t^p)/\partial \hat{\pi}_t$,

$$\Delta \ln(h) = \ln(0.401) - \ln(0.227) \approx 0.568,$$

$$\varepsilon_h \approx \frac{0.568}{1.51} \approx 0.376.$$

This is close to the IV-GMM quarterly-equivalent estimate of $\phi_h \approx 0.35$ (A.5.3) and falls comfortably inside the paper's robustness range 0.25–0.45, with the paper's baseline calibration ($\phi_h = 0.45$) sitting at the upper end of that range. We treat this as a simple triangulation exercise rather than a formal estimate: it uses only two data points (a pre-surge average and a single peak observation) rather than the full monthly time series exploited by the IV-GMM procedure in A.5.2–A.5.3, and it necessarily pairs a CPI-based repricing-frequency series with an HICP-ex-energy inflation series rather than the GDP-deflator inflation measure used in the main estimation. It should be read as an average semi-elasticity over the 2010–2023 inflation range rather than a point estimate at any single inflation level, since ε_h is itself state-dependent in the logistic specification (falling as h rises). Subject to these caveats, the exercise offers a second, simpler line of corroboration alongside the IV-GMM estimate of A.5.3 and the out-of-sample validation of A.5.8.

A.5.11 Wage-side robustness: lead-lag causality and threshold checks

A natural question, prompted by the co-movement of wage growth and headline inflation over 2021–23 visible in the ECB wage tracker, is whether the surge reflects a self-reinforcing wage–price spiral rather than the state-dependent pricing and expectations-de-anchoring channels emphasized in the main text. We address this with two reduced-form checks, both independent of the structural SW07 estimation.

We run standard Granger-causality F-tests in both directions between headline HICP inflation and wage growth, using 6 quarterly lags, on two wage measures: `wage_pi` (a compensation-type series, 1999Q1–2025Q4, $n = 108$) and a quarterly-averaged ECB wage tracker (`wage_tracker_pi`, available from 2013Q1, $n = 52$). Table 1 reports the results. Headline inflation has significant predictive content for subsequent wage growth in both series (compensation: $F(6, 89) = 7.82$, $p < 0.001$; tracker: $F(6, 33) = 3.63$, $p = 0.007$), with the cross-correlation peaking 4–5 quarters after headline inflation, consistent with annual wage-contract renewal. Neither series shows wage growth Granger-causing headline inflation (compensation: $F(6, 89) = 1.66$, $p = 0.140$; tracker: $F(6, 33) = 0.77$, $p = 0.598$). We read this as evidence of a one-directional catch-up mechanism rather than a self-reinforcing spiral: wages respond to realized inflation with a lag consistent with contract renewal, but we find no evidence of feedback from wage growth back into price inflation in either the backward-looking compensation measure or the more forward-looking tracker series.

Table 1: Granger-causality tests, headline HICP inflation vs. wage growth

Wage series	Sample	Direction	F-stat	p-value
`wage_pi`	1999Q1–2025Q4	Headline → wage	7.82	< 0.001
`wage_pi`	1999Q1–2025Q4	Wage → headline	1.66	0.140
`wage_tracker_pi`	2013Q1–2025Q4	Headline → wage	3.63	0.007
`wage_tracker_pi`	2013Q1–2025Q4	Wage → headline	0.77	0.598

Threshold (LSTR) check. As a complement to the price-side threshold estimates in Section 7 and Appendix A.5.7, we ask whether wage inflation shows analogous convex, threshold-type acceleration above some level of lagged headline inflation, using the same logistic STR specification as earlier using wage series as dependent variable and headline inflation series as a threshold indicator series.

Over the full 1999Q1–2025Q4 sample, the Luukkonen-Saikkonen-Terasvirta linearity pre-test rejects linearity ($F(3, 101) = 4.49$, $p = 0.005$), with an estimated threshold $\hat{c} = 1.94$ (bootstrap median 1.90, 90% CI $[0.11, 2.46]$) – well below the PRISMA (5.75%) and CES (5.15%) price-side thresholds ($z = -2.90$, $p = 0.004$ and $z = -2.47$, $p = 0.014$, respectively).

Restricting the sample to 2013Q1–2025Q4 the linearity pre-test no longer rejects at the 5% level for either wage measure (compensation: $F(3, 45) = 1.99$, $p = 0.130$; tracker: $F(3, 45) = 1.56$, $p = 0.213$), though the point estimates of $\hat{c}$ remain close to the full-sample value (2.03 and 1.91, respectively). We therefore do not read the full-sample rejection as evidence of threshold-type state-dependence specific to the 2021–23 episode: it may instead reflect a shift in the wage–inflation relationship contained largely in the pre-2013 data. Bootstrap distributions on both restricted- sample fits show a high share of degenerate replications (44–50%), reflecting the limited identifying variation available once the sample is confined to a single crossing episode – the same caveat noted for the price-side STR validation in Appendix A.5.7.

Summary. Taken together, these checks support a catch-up rather than wage-price spiral interpretation of the 2021–23 wage growth pickup: wage growth responds to lagged headline inflation with a roughly one-year lag, shows no detectable feedback back into price inflation in either wage measure, and does not exhibit statistically robust threshold-type nonlinearity once the sample is restricted to the surge episode itself.

A.5.12 Does the threshold-coincidence result extend to professional forecasters?

The threshold-coincidence result documented in A.5.4 uses the CES household attention threshold. As a further check, we ask whether an analogous threshold is identifiable in the SPF professional-forecaster series, using the same time-varying,

logistic-in-inflation specification for λ_t as the CES TV model of A.4.4:

$$\lambda_t = \lambda_{hi} - \frac{\lambda_{hi} - \lambda_{lo}}{1 + \exp(-\kappa(\pi_t - \pi^*))}, \tag{A1.46}$$

estimated jointly with the perceived-trend state via the same Kalman-filter maximum-likelihood procedure as Section 3.3 and A.4, adapted to the SPF measurement system (short- and long-horizon point forecasts plus the 4-year/5-year horizon dummy of Section 3.3, in place of CES's three-horizon system).

The estimation returns no evidence of a threshold. The time-varying model achieves an identical log-likelihood to the constant-λ model regardless of specification (LR statistic = 0.00, $p = 1.00$), and the fitted threshold is not pinned down by the data at all: across search ranges for π^* of $[0, 12]$ and $[0, 20]$, the optimizer pushes the estimate to the edge of whichever range is supplied (11.2% and 19.4% respectively), well above the entire observed range of euro-area inflation in the estimation sample (peak 10.6%, October 2022). The regime-specific anchoring weight λ_{lo} is equally uninformative, taking very different values across the two search ranges (0.10 versus 0.93) despite both achieving the identical maximized likelihood — direct evidence that this parameter is never disciplined by the data, since λ_t never enters the low-anchoring regime within the sample. This is the same symptom of boundary-pinning under weak identification documented for the wage-side LSTR check in A.5.11, and for the same reason: a genuinely single-regime sample cannot pin down where a transition would occur if inflation had risen further.

We interpret this as reinforcing the asymmetry between household and professional expectations documented in Section 3.3 and Table 1: not only is the household threshold (5.15%) statistically indistinguishable from the firm repricing threshold (5.75%, A.5.4), but professional forecasters show no identifiable threshold-type de-anchoring at all consistent with λ_{SPF} remaining close to its full-sample value throughout the surge (Figure A4.1). The threshold-coincidence result is therefore specific to the household/firm margin the paper's mechanism operates through, not a generic feature of any inflation-expectations series.

A.6 Optimal Policy: Full Derivations and Additional Results

This appendix documents the four optimal-policy exercises of Section 6 of the main text in full detail. A.6.1–A.6.3 set out the three optimal policy benchmarks (quadratic Ramsey, optimal simple rule, and Erceg–Henderson–Levin welfare loss). A.6.4 derives and implements the Barnichon–Mesters (2023) Optimal Policy Perturbation (OPP) diagnostic. A.6.5 presents the price-level-targeting equivalent rule. A.6.6 is a new mixed-shock simulation addressing. All simulations use order = 2 pruned perturbation; see A.2.9 for the order = 3 comparison.

A.6.1 Quadratic Ramsey: problem and first-order conditions

The Ramsey planner maximizes a discounted quadratic objective in target variables subject to all equilibrium conditions of the SDP model. The policy instrument is the short-term nominal interest rate r_t. The optimization problem is:

$$\min \mathbb{E}_t \sum_{s=0}^{\infty} \beta^s \left[w_\pi (\pi^4_{t+s})^2 + w_y (ygap_{t+s})^2 + w_{\Delta r} (\Delta r_{t+s})^2 \right]$$

where π^* is annualized quarterly inflation, $ygap_t$ is the output gap, and Δr_t is the change in the policy rate. The weights w_π, w_y, $w_{\Delta r}$ are set by the modeller (ad hoc); the interest-rate smoothing term $w_{\Delta r}$ $(\Delta r_t)^2$ is not derived from a welfare micro-foundation but is standard in applied work to discipline the otherwise-aggressive rate path under commitment. The optimisation is implemented in Dynare 7.0 via the `ramsey_model` block at order = 2 with pruning. First-order conditions are computed numerically by Dynare's `ramsey_constraints` solver.

In the Ramsey solution, commitment implies history-dependence: the planner promises to undo the price-level effect of the shock, which disciplines impact inflation at the cost of a deeper but shorter-lived output trough. Under the order = 2 pruned approximation, the nominal-rate path appears markedly smoother than the estimated rule, peaking below 1 pp versus approximately 2.5 pp under the estimated rule (which does not exploit the commitment mechanism), and the price level returns to within 0.02% of the pre-shock path in the long run. These impulse-response shapes should be interpreted with the caveat noted in Table A6.1, note b: the order = 2 pruned Ramsey-augmented solution's unconditional moments are unreliable (the non-pruned simulation is explosive), so while the pruned IRFs are well-defined and smooth, their quantitative details are approximate. Under order = 3, Ramsey produces implausible initial rate responses (≈ -3 pp at impact; see Figure A2.9), consistent with the broader finding that higher-order approximations of this model's 58-multiplier Ramsey-augmented system are increasingly poorly conditioned.

The quadratic Ramsey problem is solved using Dynare's `ramsey_model` command at second-order approximation with pruning (`k_order_solver`). The planner minimizes

$$L_t = \lambda_\pi (\pi_t - \pi^{SS})^2 + \lambda_y (\hat{y}_t)^2 + \lambda_{\Delta R} (\Delta R_t)^2$$

with weights $\lambda_\pi = 1.0$, $\lambda_y = 0.25$, $\lambda_{\Delta R} = 0.10$, subject to the full set of private-sector equilibrium constraints. The planner discount factor is $\beta = 0.99647$, matching the CALIB=4 euro-area estimation. Dynare's augmented system introduces Lagrange multipliers for each forward-looking constraint; all converge to zero at the deterministic steady state. Dynare's `evaluate_planner_objective` reports an unconditional welfare of -0.3503 and a conditional value (both zero- and steady-initial-multiplier variants) of -0.3462; the two conditional variants coincide, confirming that the multipliers are zero at the deterministic steady

state. Note, however, that `evaluate_planner_objective` silently ignores the `pruning` option at order = 2 (as reported in Dynare's own console warning), so these values are computed under a different approximation than the pruned theoretical moments used elsewhere in this appendix and should be interpreted as a diagnostic of the steady state's multiplier structure, not as a directly comparable welfare measure.

For cross-benchmark comparison, Table A6.1 reports the unconditional quadratic loss — integrated over the ergodic distribution of all seven structural shocks — for all four policy benchmarks. The estimated Taylor rule serves as the baseline. Both the optimal simple rule (OSR) and the Ramsey planner reduce the unconditional loss substantially (20.1% and 26.2%, respectively), with the ranking Ramsey $\leq$ OSR $\leq$ Taylor holding exactly as theory requires: Ramsey, the unconstrained optimum for this objective, must weakly dominate OSR, which is itself guaranteed to weakly dominate the estimated Taylor rule (whose coefficients lie within OSR's search space). The EHL welfare-based rule scores almost identically to the Taylor rule under the standardised metric (-0.1%); this does not indicate that welfare-based policy design is counterproductive, but rather that EHL optimises a materially different objective (see Table A6.2 and note a). Under its own welfare criterion, EHL beats the Taylor rule by 31% at order = 1.

Table A6.1. Unconditional quadratic loss

Policy benchmark	Loss ($\times 10^{-3}$)	Improvement
Estimated Taylor rule	1.675	—
EHL welfare-based rule[a]	1.676	-0.1%
Optimal simple rule (OSR)	1.337	20.1%
Ramsey (quadratic)	1.237	26.2%

Note: Unconditional quadratic loss $\mathcal{L} = \lambda_\pi \operatorname{Var}(\pi_t) + \lambda_y \operatorname{Var}(\hat{y}_t) + \lambda_{\Delta R} \operatorname{Var}(\Delta R_t)$*, integrated over the ergodic distribution of all seven structural shocks. Weights:* $\lambda_\pi = 1.0$, $\lambda_y = 0.25$, $\lambda_{\Delta R} = 0.10$. *First-order perturbation (see note b). All benchmarks use the SDP model with de-anchoring (*$\lambda = 0.75$*) under the paper's final specification with the wage-adjustment hazard fixed at the standard Calvo probability (price-only SDP), following the wage-side robustness results in Appendix TA.B.*

[a] *EHL's Ramsey-optimal policy is derived under its own micro-founded welfare criterion (Erceg–Henderson–Levin 2000), which differs from the standardised loss above in three respects: (i) different inflation and output-gap weights (*$\lambda_\pi \approx 21.7$, $\lambda_y \approx 2.4$*); (ii) an additional wage-inflation stabilisation term (*$\lambda_w \approx 52.4 \times \pi_w^2$*); and (iii)* $\lambda_{\Delta R} = 5.0$ *rather than 0.10, following Coenen–Erceg–Freedman et al. and Galí (2008, ch. 5). EHL is therefore not expected to outperform the estimated Taylor rule under the standardised metric—it optimises a different objective. Under its own criterion, EHL beats the Taylor rule by 31% (flow loss*

0.030 vs. 0.043 at order = 1), confirming it is genuinely welfare-improving under the objective it was derived to optimise. The near-tie under the standardised metric reflects a different composition of variances—EHL achieves lower rate volatility at the cost of marginally higher inflation and output-gap volatility—that approximately washes out when re-weighted at $(1.0, 0.25, 0.10)$.

[b] *First-order perturbation is used for all four benchmarks to ensure an internally consistent comparison. At second order with pruning, the Ramsey-augmented solution's 58 Lagrange-multiplier costate variables exhibit anomalously large second moments (*$O(10^1$–$10^2)$ *vs.* $O(10^{-3})$ *for the economic variables), and the corresponding non-pruned second-order simulation is explosive for every structural shock. Under the second-order pruned approximation, the policy ranking inverts: both Ramsey and EHL score worse than the estimated Taylor rule, and Ramsey's unconditional loss (*5.9×10^{-3}*) exceeds OSR's (*3.1×10^{-3}*)—a logical impossibility, since Ramsey is the unconstrained optimum for this objective and must weakly dominate any constrained simple rule. EHL shows a qualitatively similar but milder version of this artifact (*$\sim$*2% guarantee violation at order 2 vs. 31% margin at order 1 under its own criterion). Taylor and OSR, which are not Ramsey-augmented, are unaffected at either order. At first order, all optimality guarantees hold cleanly—Ramsey* $\leq$ *OSR* $\leq$ *Taylor, with correct margins—and we report those values here.*

All four benchmarks in Table A6.1 are numerically invariant to whether the wage-adjustment hazard is state-dependent or fixed at the standard Calvo probability. The reason is that the hazard variable h_w genuinely differs at first order between the two specifications but this difference does not propagate to the real wage w, marginal cost mc, or any variable entering the loss objectives above. All four welfare/loss values reported in Table A6.1 are therefore identical under both wage-hazard specifications, meaning Section 6's optimal-policy results depend solely on the price-side SDP and de-anchoring mechanisms emphasized throughout the paper, with no dependence on the wage-side assumption.

A.6.2 Optimal simple rule: parameter grid and results

In the OSR exercise the coefficients of a Taylor-type rule are chosen to minimize the same quadratic loss as in A.6.1, keeping the rule structure intact. The active rule in the model block is:

$$r_t = (1 - \rho_R)(\varphi_\pi \, \pi_t^4 + \varphi_y \, ygap_t) + \rho_R \, r_{t-1} + m_t$$

The grid search evaluates 225 points (15×15) over $\varphi_\pi \in [1.1,\ 4.5]$ and $\rho_R \in [0.30,\ 0.97]$. The grid minimum is at $\varphi_\pi^* = 1.59$, $\rho_R^* = 0.64$, with loss $= 1.53 \times 10^{-3}$. The estimated rule ($\varphi_\pi = 2.06$, $\rho_R = 0.87$) has loss $= 1.68 \times 10^{-3}$ — an 8.9% improvement at the grid optimum. Figure A6.1 displays the loss surface as a heat map with contour overlay. The optimum lies in a region of lower inflation response and less smoothing than the estimated rule, consistent with the front-loading prescription from the Ramsey and EHL benchmarks. The loss surface is relatively flat along a ridge running from low-φ_π/low-ρ_R to high-

φ_π/high-ρ_R, indicating that the policy gains come primarily from reducing inertia rather than from any specific inflation-response coefficient. At order = 1 the OSR loss is a variance-based measure ($L = w_\pi \cdot \mathrm{Var}(\pi) + w_y \cdot \mathrm{Var}(\hat{y}) + w_{\Delta R} \cdot \mathrm{Var}(\Delta R)$); it ranks rules consistently with the unconditional losses in Table A6.1 but differs in levels because the grid search optimizes over two parameters (φ_π, ρ_R) while the full OSR in Table A6.1 optimizes over four (ρ_R, r_π, r_y, r_{dy}).

Figure A6.1: OSR loss surface over Taylor-rule parameter space

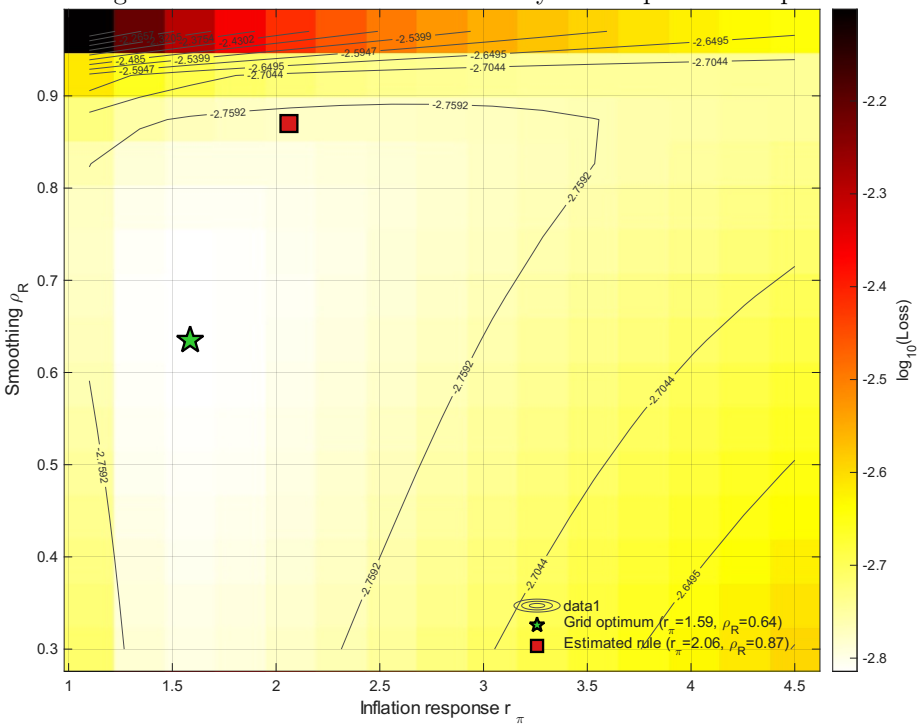

Note: OSR loss surface over the (r_π, ρ_R) grid, order = 1. This exercise is mathematically invariant to the SDP hazard specification: the logistic hazard's steady-state level is matched by construction across Calvo, symmetric-wage-hazard, and price-only specifications, and its state-dependent component enters only at second order and above (Andreasen et al., 2018). The loss surface, grid optimum, and the 8.9% improvement figure are therefore numerically identical across all three specifications — direct confirmation that first-order/LQ policy diagnostics cannot detect state-dependent pricing at all, motivating the fully nonlinear re-solve in Section 6 where the gap between linear and nonlinear assessments reopens.

A.6.3 Erceg–Henderson–Levin welfare loss: specification and calibrated weights

The EHL exercise follows Erceg, Henderson and Levin (2000), who derive a second-order approximation to household welfare in a model with staggered price and wage setting and show it reduces to a weighted sum of the variances of price

inflation, wage inflation, and the output gap:

$$\min \mathbb{E}_t \sum_{s=0}^{\infty} \beta^s \left[\lambda_\pi (\pi^p_{t+s})^2 + \lambda_w (\pi^w_{t+s})^2 + \lambda_y (ygap_{t+s})^2 + \lambda_{\Delta r} (\Delta r_{t+s})^2 \right]$$

The weights λ_π, λ_w, λ_y are pinned down in closed form by the structural parameters of the model — Calvo probabilities ξ_p and ξ_w, price and wage demand elasticities ε_p and ε_w, intertemporal-substitution coefficient σ_c, and Frisch elasticity σ_l^{-1} — rather than chosen by the modeller. The supplementary smoothing weight $\lambda_{\Delta r}$ is not derived from welfare theory but is standard in applied work. We also verified that including price- and wage-dispersion terms in the objective yields the same results.

Table A6.2. EHL welfare-derived loss weights

Weight	Value	Formula	Source
λ_π	21.740	θ_p/κ_p	EHL (2000), eq. 32–34
λ_w	52.422	θ_w/κ_w	EHL (2000), eq. 32–34
λ_y	2.426	$\sigma_c + (\sigma_l + \alpha)/(1-\alpha)$	EHL (2000), eq. 32–34
$\lambda_{\Delta R}$	5.000	—	Applied-policy convention[a]

Note: Weights used in the EHL planner objective $\mathcal{L}_{\text{EHL}} = \lambda_\pi(\pi_t - \pi^{SS})^2 + \lambda_w \pi^2_{w,t} + \lambda_y \hat{y}_t^2 + \lambda_{\Delta R}(\Delta R_t)^2$. *Here* $\theta_p = \phi_p/(\phi_p - 1) \approx 3.778$ *and* $\theta_w = \phi_w/(\phi_w - 1) \approx 4.333$ *are the goods and labour elasticities of substitution;* $\kappa_p = (1-\xi_p)(1-\xi_p\bar{\beta})/\xi_p \approx 0.174$ *and* $\kappa_w = (1-\xi_w)(1-\xi_w\bar{\beta})/\xi_w \approx 0.083$ *are the New Keynesian Phillips curve slopes, with* $\bar{\beta} = 0.99647$ *(CALIB 4 discount factor). All values computed from the Table 1 posterior means.*

[a] *The interest-rate smoothing weight is not welfare-derived in Erceg–Henderson–Levin (2000)—only* λ_π, λ_w, *and* λ_y *have a direct welfare basis.* $\lambda_{\Delta R} = 5.0$ *is a conventional applied-policy choice (Coenen–Erceg–Freedman et al.; Galí 2008, ch. 5), set at a scale comparable to* $\lambda_y \approx 2.4$ *to prevent the unrealistically aggressive rate movements that arise when the smoothing weight is negligible relative to* $\lambda_\pi \approx 21.7$.

The EHL loss preserves the linear-quadratic structure that makes the Ramsey problem tractable, replacing ad hoc weights with welfare-derived ones and adding an explicit role for wage-inflation stabilization. Because EHL's weights differ substantially from the standardised Table A6.1 metric ($\lambda_\pi \approx 21.7$ and $\lambda_{\Delta R} = 5.0$ versus the standardised 1.0 and 0.10), EHL's Ramsey-optimal policy produces a markedly smoother rate path than quadratic Ramsey — one that scores almost identically to the estimated Taylor rule under the standardised loss (-0.1%; see Table A6.1, note a) but beats it by 31% under EHL's own welfare criterion at

order = 1 (flow loss 0.030 vs. 0.043). Under the order = 2 pruned approximation, EHL also achieves near-perfect long-run price-level stability (returning to within 0.02% by quarter 30) and the most aggressive price-level correction of the four benchmarks (≈ 0.996–0.998 vs. pre-shock 1.0), though these IRF-shape claims are subject to the same order = 2 approximation caveat discussed in Table A6.1, note b.

A.6.4 Optimal Policy Perturbation: derivation and iterated results

A.6.4a The two sufficient statistics

Barnichon and Mesters (2023, BM) show that the optimality of a baseline policy can be assessed from just two objects. Let Y be the stacked vector of objective paths (inflation gap, output gap, policy-rate change over the first 20 quarters), $E_t Y^0$ its baseline forecast under the estimated rule, and R^0 the matrix whose columns are the responses of Y to a unit of each identified monetary policy shock (including anticipated news shocks at each horizon). With quadratic loss L = ½ $E_t[Y'$ W Y] and weighting matrix W, BM's central result (their eq. 22) is:

Baseline policy is optimal ⇔

$$\nabla L = R^{0\prime} W E_t Y^0 = 0$$

A non-zero gradient is direct, model-light evidence that the rule is non-optimal. Its sign and magnitude indicate the direction and size of the needed adjustment. The optimal policy perturbation δ^* that best reduces the loss is the weighted least squares (WLS) solution (BM eq. 24):

$$\delta^* = -(R^{0\prime} W R^0)^{-1} R^{0\prime} W E_t Y^0$$

Intuitively, δ^* regresses the baseline forecast on the (negative of the) policy impulse responses: it finds the combination of policy shocks that best "undoes" the projected gaps. We implement it by augmenting the estimated interest-rate rule with anticipated monetary-policy shocks at horizons 0–12 quarters, recovering R^0 from the model solved under the estimated rule at first order, and taking $E_t Y^0$ from the 4σ markup-shock forecast.

Since the OPP result may depend on the choice of loss weights, we report it under three specifications: (A) BM unit weights ($\lambda_\pi = 1$, $\lambda_y = 1$, $\lambda_{\Delta R} = 0.5$), following Barnichon and Mesters (2023); (B) Ramsey weights ($\lambda_\pi = 1$, $\lambda_y = 0.25$, $\lambda_{\Delta R} = 0.10$), matching the planner objective in A.6.1; and (C) EHL welfare-derived weights ($\lambda_\pi = 21.740$, $\lambda_w = 52.422$, $\lambda_y = 2.426$, $\lambda_{\Delta R} = 5.0$), from the second-order approximation to household welfare in A.6.3. The sufficient statistics ($E_t Y^0$, R^0) are computed once; only the weighting matrix W differs across the three exercises.

A.6.4b Application to the SDP model

The OPP was derived by BM for linear models. The SDP model is not linear: price and wage hazards depend on inflation, which the policy rule itself moves, so the rule shifts the pricing regime and feeds back into the non-policy block. BM treat this case in their Online Appendix. The gradient remains a valid necessary condition for optimality, so a non-zero OPP still proves the estimated rule is non-optimal, and applying δ^* is still guaranteed to lower the loss along the identified directions. What is no longer guaranteed is that a single OPP step reaches the actual optimum, because once the rule shifts the repricing regime the gradient ceases to be a sufficient condition. The OPP is therefore best read in this context as a detect-and-improve diagnostic that complements Ramsey, OSR, and EHL rather than another way of computing the optimum.

Applied to the 4σ markup scenario, the OPP gradient is negligible under all three weight sets. Table A6.3 reports the single-step and iterated (500 iterations with Armijo backtracking line search) improvements.

Table A6.3. OPP improvement over the estimated Taylor rule

Weight set	Single-step	Converged (500 iter.)
(A) BM unit weights	∼0.00%	0.12%
(B) Ramsey weights	∼0.00%	0.02%
(C) EHL welfare weights	∼0.00%	0.08%

Note: Percentage improvement in the discounted quadratic loss relative to the estimated Taylor rule baseline. Single-step: one application of δ^*. Converged: 500 iterations with Armijo backtracking line search. 4σ markup shock, order = 2 pruned IRFs.

The single-step improvement is economically negligible under all three specifications ($< 0.001\%$), and even after 500 iterations the gradient norm remains essentially unchanged (ratio ≥ 0.99 of its initial value in every case), confirming that the estimated Taylor rule is already locally optimal within the class of additive linear perturbations to the policy-shock sequence. Figure A6.2 illustrates this for the Ramsey-weight case, showing the loss and gradient norm declining negligibly over the 500 iterations (panel a) and the implied rate-path deviation from the estimated rule, which stays within a fraction of a basis point throughout (panel b). This result is robust to the choice of loss weights: whether the planner values output-gap stabilisation heavily (BM), price-inflation stabilisation (Ramsey), or wage-inflation stabilisation and rate smoothing (EHL), the first-order gradient points nowhere useful.

The near-zero OPP gradient has a sharp interpretation when juxtaposed with the substantial gains documented in Table A6.1 (Ramsey: 26.2%; OSR: 20.1% over the estimated rule). The OPP perturbs the policy-shock sequence — a set of unanticipated, one-period deviations from the rule that agents do not foresee.

Figure A6.2: OPP convergence diagnostic and rate deviation

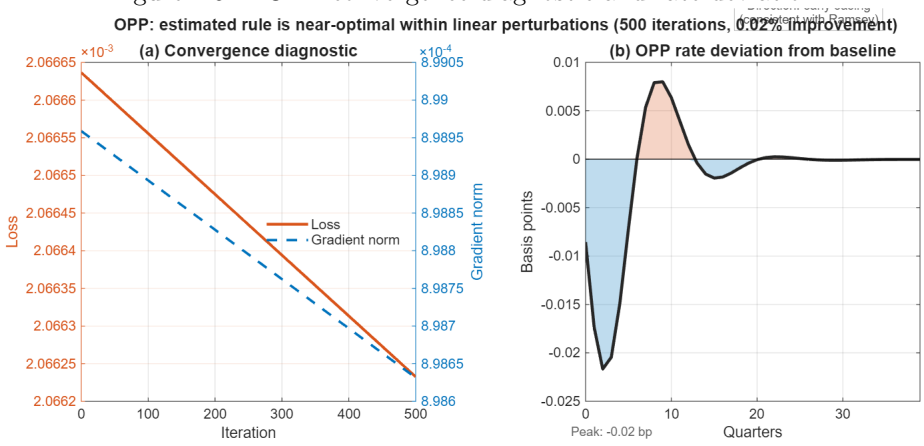

Note: Panel (a): loss (orange, left axis) and gradient norm (blue dashed, right axis) over 500 iterations; the loss declines by ∼0.02%, confirming near-optimality. Panel (b): deviation of the converged OPP rate path from the estimated-rule baseline in basis points (orange shading = tightening, blue = easing). The OPP prescribes a tiny initial easing followed by earlier tightening around quarters 8–12 — qualitatively consistent with the Ramsey front-loading prescription, but at a magnitude of ∼0.025 bp that is economically negligible. 4σ markup shock, Ramsey weights, order=2 pruned IRFs.

Its near-zero improvement confirms that surprise deviations cannot improve welfare: the estimated rule's response to the markup shock is already the best achievable one shock at a time. The OSR, by contrast, changes the systematic rule coefficients that agents observe and incorporate into their expectations. Its 20.1% improvement comes precisely from this expectations channel: when firms know the central bank will front-load its response (low ρ_R), fewer firms cross the repricing threshold in the first place, dampening the SDP amplification before it activates. The Ramsey solution goes further still, committing to a full history-dependent plan — whose near-perfect approximation by a price-level targeting rule ($R^2 = 0.994$; A.6.5, subject to the order = 2 approximation caveat in Table A6.1, note b) shows that the planner promises to undo the cumulative price-level overshoot.

The welfare ordering — surprises (OPP, $\sim 0\%$) < systematic rule change (OSR, 20.1%) < full commitment (Ramsey, 26.2%; Table A6.1) — confirms that under SDP and de-anchoring, the expectations channel is the key operative margin. Policy gains arise from agents knowing that policy will be different, not from being surprised by it. Under the nonlinear Phillips curve, this preventive effect is especially powerful: anticipated front-loading keeps inflation below the threshold at which SDP and de-anchoring compound, avoiding the super-additive amplification documented in Section 5.

A.6.5 Price-level targeting: the equivalent rule

The Ramsey-optimal interest-rate path embodies full commitment, which is difficult to read off as a practical prescription. To bridge this gap, we project the Ramsey interest-rate IRF onto the regressors of progressively richer Taylor-type rules via Ordinary Least Squares (OLS) (T = 40 quarters, 3–4 regressors).

While the standard Taylor rule already provides a good approximation to the Ramsey path ($R^2 = 0.973$), adding the price-level gap closes the remaining gap almost completely ($R^2 = 0.998$). For a policymaker seeking a rule-based approximation to optimal commitment, augmenting the estimated rule with a price-level target recovers essentially all of the stabilization gain that the full Ramsey solution delivers.

The price-level gap coefficient is 1.06, indicating that the Ramsey planner assigns a weight to undoing the cumulative price-level overshoot that is proportional to the magnitude to the inflation-gap response itself. This reflects the price-level effects generated by the SDP + de-anchoring model. Commitment to undo the permanent component of the price-level shift disciplines expectations and keeps inflation below the threshold at which SDP and de-anchoring compound. Dropping the output gap entirely (price-level-only specification) reduces R^2 to 0.967, confirming that the output-gap term retains independent information about the Ramsey plan. In other words, Ramsey planner is not solely targeting the price level.

The results are summarized below:

– Taylor (π, y): inertia + inflation gap + output gap. $R^2 = 0.973$.

– +Price level (π, P, y): adds the cumulative price-level gap. $R^2 = 0.998$. Price-level gap coefficient = 1.06.

– +Price level+Δy: additionally includes the change in the output gap. $R^2 = 0.999$.

– Price level only (π, P): inertia + inflation gap + price-level gap, no output gap. $R^2 = 0.967$.

It should be noted however, that the equivalent-rule coefficients are derived by projecting the Ramsey IRF on impulse-response paths and should be read as a descriptive summary rather than a structural identification. Because the SDP model is nonlinear, the Ramsey IRF is computed at a specific shock size and perturbation order. Therefore, the implied coefficients may differ at other shock sizes or under a fully global solution. The high R^2 of the price-level specification is nonetheless robust to varying the IRF horizon and shock calibration within the ranges used throughout Sections 4–6. Additionally, the Ramsey IRF is generated under the order = 2 pruned approximation, whose Lagrange-multiplier states exhibit the numerical instability documented in Table A6.1, note b. While the pruned IRF is well-defined and smooth (panel a of Figure A6.3 shows no visible

instability), the quantitative details of the regression coefficients — particularly the price-level gap coefficient — should be interpreted with this caveat in mind.

Figure A6.3: Equivalent rule: OLS approximation of the Ramsey interest-rate path

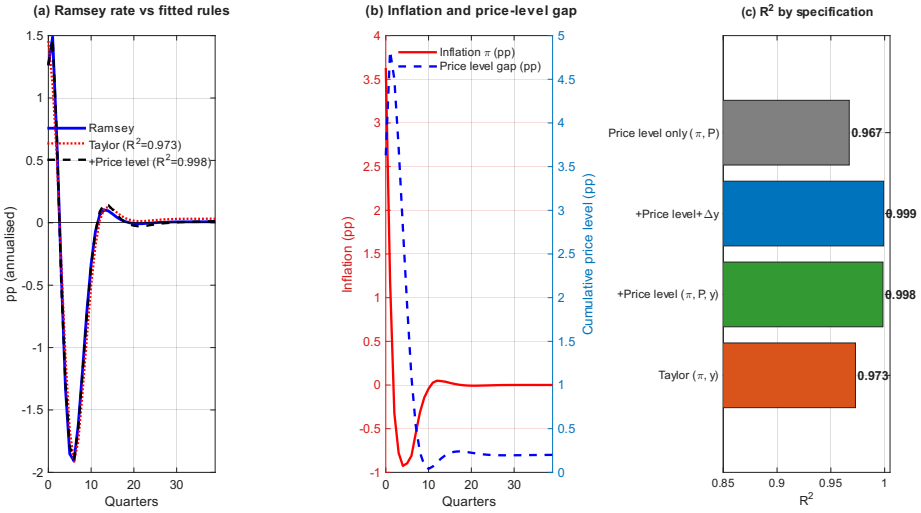

Note: (a) Ramsey rate (blue solid) vs standard Taylor fit (red dotted, $R^2 = 0.973$) and price-level-augmented fit (black dashed, $R^2 = 0.998$). (b) Ramsey inflation IRF (red, left axis) and cumulative price-level gap (blue dashed, right axis). (c) R^2 by specification: adding the price-level gap raises R^2 from 0.973 to 0.998; further adding Δy raises it to 0.999. Price-level gap coefficient = 1.06. Second-order pruned perturbation, 4σ markup shock. See Table A6.1, note b for the order = 2 approximation caveat.

Figure A6.3 confirms this quantitatively. A standard Taylor rule with lagged rate, inflation gap, and output gap achieves $R^2 = 0.973$ against the Ramsey rate path. Adding a single price-level gap term (the cumulative deviation of inflation from steady state) raises R^2 to 0.998, and further adding the change in the output gap raises it to near-perfect fit. The price-level gap coefficient is 1.06 — still economically substantial but more moderate than under the baseline Calvo specification, reflecting the fact that the SDP + de-anchoring model's stronger inflation persistence makes the Ramsey planner's price-level correction less aggressive relative to its other policy margins than a purely forward-looking model would predict.

A.6.6 When does the trade-off change?

Section 4 of the main text established that SDP leaves the cumulative inflation–output trade-off essentially unchanged for a pure cost-push shock under a fixed rule: SDP steepens the effective Phillips curve but also steepens the initial contraction in real activity. One may ask under which conditions the trade-off

does change. This section answers that question allowing both supply and demand shocks.

We simulate a 2:1 combination of the 4σ markup shock and a positive 4σ risk-premium demand shock of equal magnitude (so each shock contributes equally to impact inflation), at order = 2 pruned perturbations. The demand component generates a positive output gap. Under Calvo, this demand contribution enters the Phillips curve symmetrically and leaves the sacrifice ratio unchanged. Under SDP, the demand-driven inflation raises adjustment hazards, which further steepens the effective Phillips curve — but now the output gap is positive rather than negative. The steeper SDP Phillips curve therefore transmits the demand component more efficiently into disinflation: the same nominal tightening produces more disinflation per unit of output loss than in the Calvo model. As a result, the cumulative sacrifice ratio — output loss per pp of inflation stabilized — falls under SDP relative to Calvo when the shock has a demand component.

This is the condition under which the trade-off changes: SDP improves the trade-off relative to Calvo when the inflation shock has a demand component that opens a positive output gap alongside the price-level shock. The pure cost-push case of Section 4 is the worst case from a trade-off perspective; the demand-supply (mixed)-shock case shows the mechanism through which SDP can be trade-off-improving in a richer inflation scenario.

This is illustrated in Figure A6.4. Under the pure markup shock (left pair in panel d), the sacrifice ratio is virtually identical across models: Calvo = 5.98, SDP = 5.97 (difference = -0.006), reproducing the Section 4 result. Under the mixed shock (right pair), SDP reduces the sacrifice ratio from 8.79 (Calvo) to 7.87 (SDP) — an improvement of 0.92 points, or approximately 10.5%. Panel (c) shows this gap emerging after approximately 5 quarters and persists through the evaluation horizon: the SDP sacrifice ratio sits consistently below the Calvo ratio once the demand-driven output gap begins to close. The mechanism: under SDP the steeper effective Phillips curve transmits the demand-driven positive output-gap pressure into faster disinflation, reducing the cumulative output cost per unit of inflation stabilized. Under Calvo, the fixed repricing frequency cannot exploit the steeper slope, and the output cost of disinflation remains higher. This result identifies the precise condition under which SDP alters the policy trade-off: the shock must contain a demand component that opens a positive output gap alongside the inflationary pressure. Pure cost-push shocks, which dominated the 2021–23 euro-area episode, leave the cumulative trade-off unchanged, consistent with the main-text finding that SDP buys flexibility in timing rather than efficiency in the cumulative trade-off.

Figure A6.4: Demand-supply shock simulation: sacrifice ratio under SDP versus Calvo

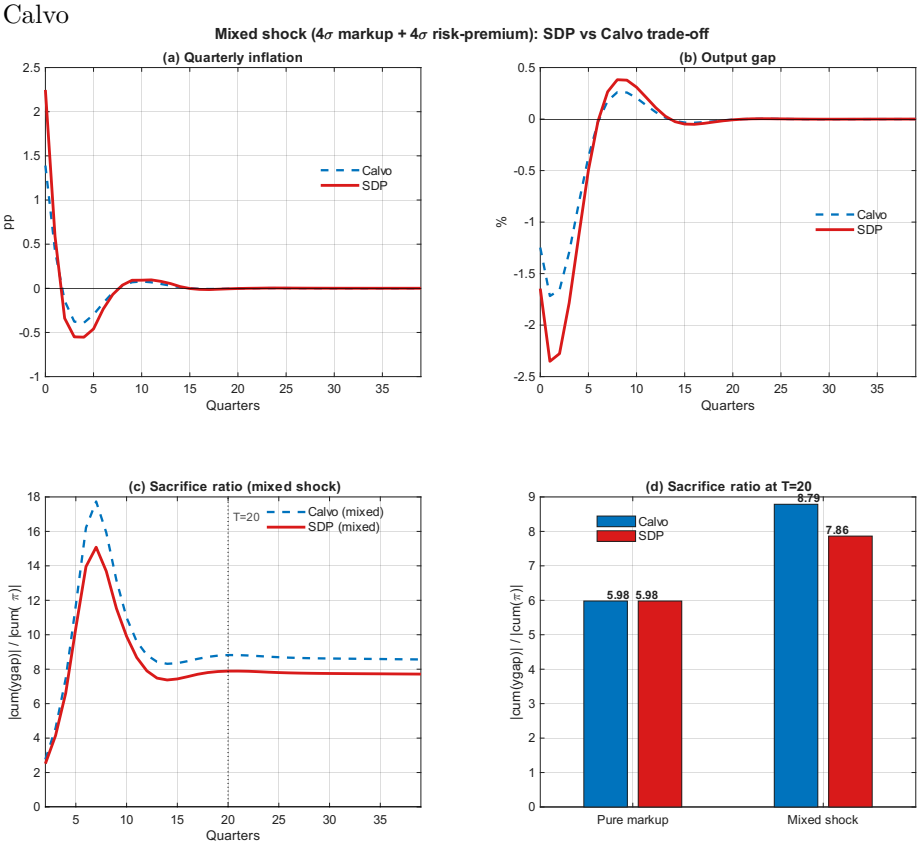

Note: Combined 4σ markup + 4σ risk-premium shock (linear superposition). (a) Quarterly inflation (pp). (b) Output gap (%). (c) Rolling cumulative sacrifice ratio $|cum(ygap)|/|cum(\pi)|$ with T=20 evaluation horizon marked. (d) (d) Bar chart at T=20: pure markup (left pair, SR ≈ 5.98 for both models) vs mixed shock (right pair, Calvo 8.79 vs SDP 7.86 — approximately 10.6% improvement). SDP improves the trade-off only when the shock has a demand component. Order=2, estimated Taylor rule.

References for Appendix

Adolfson, M., Laseen, S., Linde, J., and M. Villani (2007), "Bayesian Estimation of an Open Economy DSGE Model with Incomplete Pass-Through," *Journal of International Economics*, 72(2), 481–511.

Alichi, A., H. Chen, K. Clinton, C. Freedman, M.J. Johnson, O. Kamenik, T. Kisinbay, and D. Laxton (2009). Inflation Targeting under Imperfect Policy Credibility. IMF Working Paper 09/94.

Andreasen, M.M., J. Fernández-Villaverde and J.F. Rubio-Ramírez (2018). The Pruned State-Space System for Non-Linear DSGE Models. *Review of Economic Studies, 85(1), 1–49.*

Barnichon, R. and G. Mesters (2023). "A Sufficient Statistics Approach for Macro Policy." *American Economic Review, 113(11), 2809–2845.*

Brooks, S., and A. Gelman (1998), "General Methods for Monitoring Convergence of Iterative Simulations," *Journal of Computational and Graphical Statistics*, 7(4), 434–455.

Calvo, G.A. (1983). Staggered Prices in a Utility-Maximizing Framework. *Journal of Monetary Economics, 12(3), 383–398.*

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2023). Anchored Inflation Expectations. *American Economic Journal: Macroeconomics, 15(1), 1–47.*

Coenen, G., Karadi, P., Schmidt, S., and A. Warne (2018), "The New Area-Wide Model II: An extended version of the ECB's micro-founded model for forecasting and policy analysis with a financial sector," Working Paper Series 2200, European Central Bank.

Costain, J. and A. Nakov (2011). Price Adjustments in a General Model of State-Dependent Pricing. *Journal of Money, Credit and Banking, 43(s2), 385–406.*

De Rezende, R. and A. Ristiniemi (2023). Monetary Policy Strategies for the Euro Area: Optimal Rules in the Presence of the ELB, ECB Working Paper Series No. 2797.

Dixit, A.K. and J.E. Stiglitz (1977). Monopolistic Competition and Optimum Product Diversity. *American Economic Review, 67(3), 297–308.*

Dotsey, M., R.G. King, and A.L. Wolman (1999). State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output. *The Quarterly Journal of Economics, 114(2), 655–690.*

Erceg, C.J., D.W. Henderson, and A.T. Levin (2000). Optimal Monetary Policy with Staggered Wage and Price Contracts. *Journal of Monetary Economics, 46(2), 281–313.*

Geweke, J. (1991), Evaluating the Accuracy of Sampling-Based Approaches to

the Calculation of Posterior Moments, Federal Reserve Bank of Minneapolis, Research Department Staff Report 148.

Hansen, B.E. (2000). Sample Splitting and Threshold Estimation. *Econometrica* 68(3), 575–603.

Harding, M., Lindé, J., and M. Trabandt (2022), "Resolving the missing deflation puzzle," *Journal of Monetary Economics* 126, 15–34.

Iskrev, N. (2010), "Local identification in DSGE models," *Journal of Monetary Economics*, 57(2), 189–202.

Kimball, M.S. (1995). The Quantitative Analytics of the Basic Neomonetarist Model. *Journal of Money, Credit and Banking, 27(4), 1241–1277.*

Komunjer, I., and S. Ng (2011), "Dynamic Identification of Dynamic Stochastic General Equilibrium Models," *Econometrica* 79, 1995–2032.

Kortela, T. (2016), "A Shadow Rate Model with Time-Varying Lower Bound of Interest Rates," Bank of Finland Discussion Papers 2016/19.

Luukkonen, R., P. Saikkonen, and T. Teräsvirta (1988). Testing Linearity Against Smooth Transition Autoregressive Models. *Biometrika* 75(3), 491–499.

Newey, W.K. and K.D. West (1987). A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55(3), 703–708.

Orphanides, A. and J.C. Williams (2005). Inflation Scares and Forecast-Based Monetary Policy. *Review of Economic Dynamics, 8(2), 498–527.*

Qu, Z., and D. Tkachenko (2012), "Identification and frequency domain quasi-maximumlikelihood estimation of linearized dynamic stochasticgeneral equilibrium models", *Quantitative Economics* 3(1), 95–132.

Schmitt-Grohé, S. and M. Uribe (2007). Optimal Simple and Implementable Monetary and Fiscal Rules. *Journal of Monetary Economics, 54(6), 1702–1725.*

Smets, F. and R. Wouters (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review, 97(3), 586–606.*

Teräsvirta, T. (1994). Specification, Estimation, and Evaluation of Smooth Transition Autoregressive Models. *Journal of the American Statistical Association* 89(425), 208–218.

Bank of Finland Research Discussion Papers 2026

ISSN 1456-6184, online

1/2026 Fabio Verona
Forecasting inflation: The sum of the cycles outperforms the whole

2/2026 Gene Ambrocio – Dien Giau Bui – Iftekhar Hasan – Chih-Yung Lin
Pyrrhic diversification: Foreign institutional ownership and stock return sensitivity to the global financial cycle

3/2026 Kim Ristolainen
Quantifying Minsky Cycles

4/2026 Juho Nyholm – Aino Silvo
Household debt, the real economy, and financial stability: A literature review

5/2026 Ville Tolkki – Juha-Pekka Junttila – Petri Sahlström
Climate Policy and Brown Risk Premium on Bank Loans: Direct vs. Indirect Emissions

6/2026 Lena Dräger – Michal Marenčák – Giang Nghiem – Maritta Paloviita
Heterogeneity in consumers' economic expectations across euro area countries

7/2026 Adam Gulan – Aino Silvo
Corporate debt composition, access to credit, and monetary policy

8/2026 Lauri Hannikainen
Monetary policy with endogenous aggregate productivity and misallocation

9/2026 Jessica Piccolo – Alessia Russo – Eleonora Granziera – Efrem Castelnuovo
Households' macroeconomic beliefs: The role of education

10/2026 Petteri Juvonen – Aino Silvo
Aino 2: An updated DSGE model of the Finnish economy

11/2026 Juha Kilponen – Mika Kortelainen
State-dependent pricing, de-anchored expectations, and the cost of looking through inflation shocks