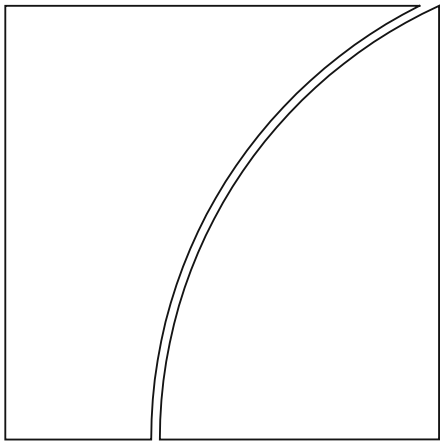


BIS Working Papers
No 1398

When the unthinkable happens: perceptions of monetary policy in tail events

by Christiane Baumeister, Fiorella De Fiore and Jasmine Xiao

Monetary and Economic Department

November 2026

JEL classification: D84, E31, E52, E58

Keywords: monetary policy rule, learning, tail risk, belief-driven business cycles

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in this publication are those of the authors and do not necessarily reflect the views of the BIS or its member central banks. The use of any map does not constitute, and should not be construed as constituting, an expression of a position by the BIS regarding the legal status or sovereignty of any territory or its authorities, the delimitation of international frontiers and boundaries, and/or the name and designation of any territory, city or area. References to individual firms and projects are for illustrative purposes and should not be construed as any statement on the soundness or the legal or regulatory status of any firm or project, or any form of endorsement by the BIS and/or the authors of this paper.

This publication is available on the BIS website (www.bis.org).

© *Bank for International Settlements 2026. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

When the Unthinkable Happens: Perceptions of Monetary Policy in Tail Events*

Christiane Baumeister† Fiorella De Fiore‡ Jasmine Xiao§

November 5, 2026

Abstract

We study how public perceptions of monetary policy evolve when the economy experiences rare and severe supply shocks. When agents learn in real time about the distribution of the shocks and the monetary policy reaction function, they tend to underestimate the strength of the response to inflation following large but infrequent negative supply shocks. Embedding this mechanism in a general equilibrium model with learning, we show that it endogenously generates inflation dynamics persistently above target and temporary output overshooting after a tail event. Using U.S. data, we find empirical evidence supporting our model's predictions about the distinctive role of tail events in shaping policy perceptions and macroeconomic dynamics. Our findings indicate that public misperceptions can amplify and prolong the aggregate effects of tail events, potentially undermining central banks' stabilization efforts.

JEL codes: D84, E31, E52, E58

Keywords: monetary policy rule, learning, tail risk, belief-driven business cycles

*The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Bank for International Settlements. We thank Laura Gáti, Brigitte Hochmuth, Bartosz Maćkowiak, Christian Matthes, Krisztina Molnar, Carolin Pflueger, Morten Ravn, Victor Rios Rull, Harald Uhlig, Venky Venkateswaran and Mirko Wiederholt for helpful comments and suggestions, and seminar and conference participants at 2026 Barcelona Summer Forum (Monetary Policy), Conference on New Challenges in Monetary Economics (Mannheim), 2026 International Research Forum on Monetary Policy (ECB), 2026 CEBRA annual meetings, Female Voices in Applied Economics Conference (NYUAD), 2026 Hydra Workshop on Dynamic Macroeconomics, NTU Behavioral Macro Workshop (Singapore), Paris-Oxford Conference on Macroeconomics of Expectations, SNB-BIS Research Workshop, So.Fi.E conference on DSGE Modelling (Brussels), St Andrews Workshop in Macroeconomics, 6th WE_ARE_IN Macroeconomics and Finance Conference (Rome), LMU Munich, and University of Amsterdam for useful feedback.

†University of Notre Dame, NBER and CEPR. Email: Christiane.Baumeister.2@nd.edu

‡Bank for International Settlements and CEPR. Email: Fiorella.DeFiore@bis.org

§University of Amsterdam. Email: y.xiao2@uva.nl

1 Introduction

Market perceptions of monetary policy play a central role in determining the effectiveness of policy actions. Central banks have made sustained efforts to improve communications but their reaction functions remain complex and are obscured by ever-changing economic conditions. As a result, the public inevitably holds an incomplete understanding of how policymakers respond to shocks and must infer the policy stance from observed decisions and macroeconomic outcomes.

These interpretative challenges become especially acute during "tail events" – when large and rare shocks materialize – which often lie outside prior experiences of both the public and policymakers. In such circumstances, incomplete understanding of the central bank's policy can heighten uncertainty surrounding the shock. Actions that would normally anchor expectations may instead be seen as inconsistent with the perceived policy reaction or insufficient to achieve the mandated price stability goal. Misperceptions about the central bank's reaction function can potentially amplify volatility and undermine stabilization efforts. Episodes such as the 2020 pandemic-induced supply disruptions demonstrate that when the unthinkable happens, the impact on inflation and output can often persist long after the initial shock has subsided, despite prompt and forceful policy interventions.

These observations motivate the central questions of this paper: How do public perceptions of monetary policy evolve after tail events, and can these perceptions, in turn, amplify and propagate shocks through belief-driven dynamics? We investigate these questions through the lens of a general equilibrium model with non-normally distributed supply shocks. Agents are fully rational but face informational frictions. The central bank follows a Taylor rule with a time-varying response to inflation that is unobserved by the public, who must infer it from realized macroeconomic outcomes. At the same time, they do not know the true distribution of supply shocks, which may feature fat tails, and estimate the distribution non-parametrically using real-time data. This dual learning problem – about the policy rule and the shock process – creates an endogenous expectation channel that amplifies and prolongs the effects of large but temporary shocks.

When a severe negative supply shock occurs, our model generates amplified and persistent dynamics. Inflation and the policy rate rise sharply on impact, while output contracts. Although the shock dissipates quickly, inflation, expected inflation, and the interest rate remain elevated for an extended period. Output overshoots before slowly returning to steady state, while remaining persistently above its flexible price

equilibrium level. These patterns closely resemble the behavior of U.S. macroeconomic variables that we document in the aftermath of a tail event — such as the major oil supply disruption in 2020Q2 — but differ from the observed dynamics following "small" supply shocks, which typically drive the business cycle.

We show that the macroeconomic dynamics observed after a large negative supply shock arise from a novel belief-driven channel. The public observes a sharp increase in inflation but cannot immediately discern whether it reflects an unusually severe shock or a weaker policy stance. Because tail events are, by definition, rare, the shock lies far in the tail of agents' perceived distribution, leading them to rationally misinterpret part of the inflation surge as reflecting a central bank more tolerant to inflation than previously perceived. This inference lowers the perceived inflation response coefficient in the policy reaction, raising expected inflation and output. The misperception of the policy rule thus amplifies the impact of the shock on inflation, while simultaneously contributing to output overshooting.

The reaction of labor also plays a role in this output overshooting. As agents observe real time data, they update their beliefs about the distribution of shocks, increasing the perceived likelihood of left-tail events. Expectations of persistently unfavorable future conditions induce consumption to fall below steady state and workers to supply more labor.

Private agents learn slowly about the policy rule, which can only be inferred indirectly from the central bank's actions and the resulting macroeconomic outcomes. As a result, expectations of inflation and output adjust slowly, causing both variables to remain away from their steady states for extended periods, long after the initial shock has dissipated. Crucially, agents' misperception of the policy rule is self-reinforcing: the ensuing inflation persistence and output overshoot confirm agents' mistaken view of a more lenient policy stance, further delaying the convergence of their beliefs to the true policy rule. When either the policy rule is fully observed or the shock distribution remains unchanged, this feedback loop disappears. Volatility falls and the macroeconomic variables quickly revert back to their steady states.

The amplification mechanism driven by learning about the shock distribution and policy rule is highly nonlinear. Tail events trigger significantly larger belief revisions about monetary policy and shocks, leading to unique dynamics such as output overshooting and more persistent inflation expectations — patterns that are absent in reaction to smaller business cycle shocks.

We empirically test the main implications of the model using U.S. macroeconomic data from 1985 to 2019. We estimate the dynamic responses of inflation, output, and

interest rates to oil supply shocks from Baumeister and Hamilton (2019) in a nonlinear VARX model. Specifically, the empirical specification does not impose assumptions on the probability distribution of shocks and includes measures of the evolution of the supply shock distribution over time, mirroring two key features of the model. Furthermore, the nonlinear specification allows the responses of inflation, output, and interest rates to depend on the magnitude of the shock, thereby distinguishing between small and large supply disruptions.

The estimated impulse responses echo the predictions of the model. After a temporary and large negative supply shock, inflation and the output gap rise persistently. At the same time, perceived policy responsiveness declines, indicating that agents interpret part of the inflation surge as reflecting a weaker policy reaction. Consistent with the model, these fluctuations in policy perceptions are associated with persistently elevated inflation expectations, which in turn amplify and prolong the inflationary response.

Related Literature Our paper is related to the literature that studies how the public learns and forms perceptions about the central bank's policy rule. The idea that the general public has an incomplete understanding of the Fed's policy rule and even experts – such as the professional forecasters – continuously update their perceptions about the rule from policy actions is consistent with the empirical findings by Bauer and Swanson (2023a,b) and Bauer, Pflueger, and Sunderam (2024a,b). Among others, Cieslak (2018), Schmeling, Schrimpf, and Steffensen (2022), Bianchi, Ludvigson, and Ma (2022), and Bianchi, Lettau, and Ludvigson (2022) have shown that (mis)perceptions about the policy stance matter for term premia and asset prices.

We contribute to this literature by providing an explanation of why public perceptions are especially important for the effectiveness of monetary policy during and after tail events. Our model explains why it is particularly difficult for the public to infer the central bank's response to large but infrequent negative shocks in real time, how such misperceptions can arise rationally, and how the resulting evolution of beliefs generates a novel propagation channel that sustains persistent inflation dynamics even in response to i.i.d. shocks.

Our paper also relates to the literature on the design of monetary policy rules when agents are learning (see, for example, Bullard and Mitra, 2002; Evans and Honkapohja, 2003; Schorfheide, 2005; Eusepi and Preston, 2010; Bodenstein, Hebden, and Winkler, 2022). Recently, Bocola, Dovis, Jørgensen, and Kirpalani (2025) develop a model in which agents are imperfectly informed about the central bank's objectives and learn

from its policy actions, so monetary policy has a reputation-building role. Together, this literature underscores the importance of policy perceptions for the macroeconomy and the conduct of monetary policy.

Like in this literature, we also assume that the public has a different information set from the central bank and solves a Bayesian filtering problem to learn about the policy rule. We contribute to this literature by developing a model in which agents learn about the policy rule while simultaneously learning about the distribution of supply shocks, which is unknown and potentially heavy-tailed. As a result, the inference problem is nonlinear and non-Gaussian, and learning is implemented using a nonlinear Bayesian filter. We show that in this dual-learning setting, misperceptions about the policy rule can arise rationally following extreme events. Therefore, we complement contemporaneous work (such as Bocola, Dovis, Jørgensen, and Kirpalani, 2025) that emphasizes the anchoring of expectations by highlighting the role of tail events, where de-anchoring can arise from imperfect inference.

Recent studies have examined mechanisms through which Bayesian learning can be slow, thereby generating endogenous persistence in macroeconomic dynamics (e.g. Kozlowski, Veldkamp, and Venkateswaran, 2020; Leduc, Moran, and Vigfusson, 2023; Gáti, 2026) or asset prices and forecasting anomalies (e.g. Johannes, Lochstoer, and Mou, 2016; Farmer, Nakamura, and Steinsson, 2024). Our model builds on these insights by combining learning about both the policy rule and the shock distribution, which interact to amplify the effects of rare events. We share the idea in Kozlowski, Veldkamp, and Venkateswaran (2020) that agents do not know the true distribution of supply shocks in the economy and update their beliefs about the distribution using incoming data. However, in their framework agents directly observe the realized shocks and estimate the distribution from those observations. By contrast, in our model agents never observe the supply shocks themselves. Instead, they must infer them jointly with the monetary policy rule from endogenous macroeconomic outcomes. As a result, beliefs about the shock distribution and beliefs about policy interact endogenously, generating a novel amplification mechanism through expectation formation and policy-rule misperception.

Finally, since we use the major oil supply disruption in 2020 as an example of a tail event, our results speak to the literature on the inflation surge after the pandemic, such as Guerrieri, Lorenzoni, Straub, and Werning (2022), di Giovanni, Şebnem Kalemli-Özcan, Álvaro Silva, and Yildirim (2022), Ball, Leigh, and Mishra (2022), Comin, Johnson, and Jones (2023), Benigno and Eggertsson (2023), Lorenzoni and Werning (2023), and Pflueger (2025). Closer to ours is Gagliardone and Gertler (2026), which also em-

phasizes the role of oil shocks, though we differ significantly in model setup and focus.[1] Furthermore, the key idea in our paper – that when the public must learn in real time about the central bank's response to large but infrequent shocks, they tend to underestimate the strength of the policy reaction to inflation – is not specific to pandemic-related shocks.

The remainder of the paper is organized as follows. Section 2 presents the model under private agents' dual learning process - about both the supply shock distribution and the monetary policy rule. Section 3 outlines the parameterization and the learning algorithm employed to solve for the equilibrium. Section 4 presents the quantitative analysis, including a comparison of the impulse responses to a tail supply shock under different informational assumptions, as well as an explanation of the main mechanism driving our results. Section 5 provides empirical evidence for the U.S. Section 6 concludes.

2 Model

This section presents a general equilibrium model where oil enters both household consumption and firms' production (Gagliardone and Gertler, 2026). Oil supply is stochastic and private agents neither observe the realization nor the distribution of oil supply shocks.

The economy is inhabited by a representative household comprising a family and a continuum of workers, as well as different types of firms. A continuum of competitive input producers combines labor and oil into differentiated production inputs. An aggregator sector transforms the differentiated inputs into a homogeneous composite input that is used by monopolistically competitive firms to produce differentiated production goods. These firms' price setting faces quadratic adjustment costs (Rotemberg, 1982). A continuum of competitive firms combine differentiated production goods into an aggregate composite consumption good.

The central bank sets the nominal policy rate according to a Taylor rule with a time-varying reaction coefficient to inflation. Private agents observe neither the central bank's reaction to inflation nor the shock to the monetary policy rule.

Private agents learn about the aggregate latent state – both the realization and distribution of the oil supply shocks as well as the strength of the central bank's reaction

[1] Baumeister, Huber, and Marcellino (2024) also emphasize the importance of tail events in the global market for crude oil. While their work focuses on forecasting tail risks, we examine how such risks influence the public's perceptions of monetary policy and their broader macroeconomic implications.

to inflation – from observing each period's macroeconomic outcome.

2.1 The environment

Oil is used in production and directly consumed by households. The supply of oil is exogenous and stochastic,

$$S_t = \overline{S}\exp(v_t), \tag{1}$$

where v_t is iid conditional on its true distribution g, which needs not be Gaussian.

Private agents do not know this distribution and instead use the recursively estimated density $\widehat{g}_t$.

2.1.1 Household

The representative household is composed of a family and a continuum of workers.

The family owns firms, pools income, trades the nominal bond, and makes consumption and saving decisions using common beliefs about aggregate conditions.

The household has preferences

$$\widehat{\mathbb{E}}_{0^-}\sum_{t=0}^{\infty}\beta^t\left[\frac{C_t^{1-\sigma}-1}{1-\sigma}-\int_0^1\frac{N_{h,t}^{1+\varphi}}{1+\varphi}dh\right], \tag{2}$$

where $\widehat{\mathbb{E}}_{t^-}[X] \equiv \mathbb{E}[X \mid \mathcal{I}_{t^-}]$ denotes the common private sector belief based on the information set I_{t^-} inherited at the beginning of period t. N_t denotes labor. Family consumption, C_t, is a CES basket of a composite consumption good $C_{q,t}$ and oil consumption $C_{o,t}$:

$$C_t = \left[\chi^{1/\psi}C_{o,t}^{1-1/\psi} + (1-\chi)^{1/\psi}C_{q,t}^{1-1/\psi}\right]^{\frac{1}{1-1/\psi}}. \tag{3}$$

The family budget constraint can be written as

$$P_tC_t + B_t = \int_0^1 W_{h,t}N_{h,t}\,dh + R_{t-1}B_{t-1} + \Xi_t, \tag{4}$$

where B_t denotes the risk free nominal bond which yields a return of R_t in period $t+1$. $W_{h,t}$ and $N_{h,t}$ represent, respectively, the wage and labor supply of each individual worker h within the family's continuum of workers. Ξ_t denotes profits distributed from the monopolistically competitive firms.

Let $P_{q,t}$ denote the price of the composite consumption good and $P_{o,t}$ the nominal

price of oil. The associated consumption-basket price index is

$$P_t = \left[\chi P_{o,t}^{1-\psi} + (1-\chi) P_{q,t}^{1-\psi}\right]^{\frac{1}{1-\psi}}. \tag{5}$$

Conditional demands are

$$C_{q,t} = (1-\chi)\left(\frac{P_{q,t}}{P_t}\right)^{-\psi} C_t, \tag{6}$$

$$C_{o,t} = \chi\left(\frac{P_{o,t}}{P_t}\right)^{-\psi} C_t. \tag{7}$$

Define the prices of oil and the consumption basket, expressed in terms of the composite consumption good, as

$$s_t \equiv \frac{P_{o,t}}{P_{q,t}} \quad \text{and} \quad p_t \equiv \frac{P_t}{P_{q,t}}. \tag{8}$$

Equation (5) can be rewritten as

$$p_t = \left[\chi s_t^{1-\psi} + (1-\chi)\right]^{\frac{1}{1-\psi}}. \tag{9}$$

The nominal-bond Euler equation is

$$1 = \beta R_t p_t C_t^{\sigma} \widehat{\mathbb{E}}_{t^-}\left[\frac{C_{t+1}^{-\sigma}}{p_{t+1}\Pi_{q,t+1}}\right]. \tag{10}$$

where $\Pi_{q,t} = \frac{P_{q,t}}{P_{q,t-1}}$ denotes inflation in the composite consumption good.

The family includes a continuum of workers, indexed by $h \in (0,1)$. Workers do not observe the aggregate wage in the economy. Each worker h makes static labor-supply decisions after observing only the local real wage $w_{h,t}^c \equiv \frac{W_{h,t}}{P_t}$, where

$$\log w_{h,t}^c = \log w_t^c + \eta_{h,t}^h \quad \text{where} \quad \eta_{h,t}^h \overset{iid}{\sim} N(0, \sigma_h^2) \tag{11}$$

Information about local wages is not pooled across households and therefore does not enter the common information set and learning problem. When making labor decisions, workers take as given family consumption. The intratemporal condition is

$$w_{h,t}^c = N_{h,t}^{\varphi} C_t^{\sigma}. \tag{12}$$

Thus household oil consumption affects labor supply through the consumption real

wage $W_{h,t}/P_t$. Aggregate labor supply is

$$N_t = \int_0^1 N_{h,t}\, dh. \tag{13}$$

2.1.2 Differentiated input producers

A continuum of competitive input producers indexed by $k \in [0,1]$ combines labor $L_{k,t}$ and oil $O_{k,t}$ into differentiated production inputs $X_{k,t}$,

$$X_{k,t} = \left[\alpha^{1/\zeta} L_{k,t}^{1-1/\zeta} + (1-\alpha)^{1/\zeta} O_{k,t}^{1-1/\zeta}\right]^{\frac{1}{1-1/\zeta}}. \tag{14}$$

Define the real transaction prices of labor and oil, in units of the composite consumption good, as $w_{k,t} \equiv \frac{W_{k,t}}{P_{q,t}} = w$ and $s_{k,t} \equiv \frac{P_{o,k,t}}{P_{q,t}}$. Each producer k observes only the local real transaction prices relevant for its labor and oil purchases, where

$$\log w_{k,t} = \log w_t + \eta^w_{k,t} \quad \text{where} \quad \eta^w_{k,t} \overset{iid}{\sim} N(0, \sigma_w^2), \tag{15}$$

$$\log s_{k,t} = \log s_t + \eta^o_{k,t} \quad \text{where} \quad \eta^o_{k,t} \overset{iid}{\sim} N(0, \sigma_o^2), \tag{16}$$

and makes static input decisions based on those prices. Observations about input prices are neither pooled across input producers nor incorporated into the common information set and learning problem.

Unit cost of the differentiated production input k is

$$q_{k,t} = \left[\alpha\, {w^q_{k,t}}^{1-\zeta} + (1-\alpha) s_{k,t}^{1-\zeta}\right]^{\frac{1}{1-\zeta}}. \tag{17}$$

Factor demands are

$$L_{k,t} = \alpha \left(\frac{w_{k,t}}{q_{k,t}}\right)^{-\zeta} X_{k,t}, \tag{18}$$

$$O_{k,t} = (1-\alpha) \left(\frac{s_{k,t}}{q_{k,t}}\right)^{-\zeta} X_{k,t}, \tag{19}$$

which aggregate to

$$L_t = \int_0^1 L_{k,t}\, dk \tag{20}$$

$$O_t = \int_0^1 O_{k,t}\, dk. \tag{21}$$

2.1.3 Composite-input aggregator

A competitive input aggregator combines the differentiated inputs, $X_{k,t}$, into a homogeneous composite input, X_t, using the technology

$$X_t = \left[\int_0^1 X_{k,t}^{\frac{\kappa-1}{\kappa}} dk\right]^{\frac{\kappa}{\kappa-1}}. \tag{22}$$

where κ is the elasticity of substitution across differentiated production inputs.

Let Q_t be the unit price of the composite input. Cost minimization implies that

$$Q_t = \left[\int_0^1 q_{k,t}^{1-\kappa} dk\right]^{\frac{1}{1-\kappa}}. \tag{23}$$

As local-price disturbances in equation (15) and (16) are iid across producers, Q_t can be written as

$$Q_t = Q(w_t, s_t; \sigma_w, \sigma_o),$$

where the expression for $Q(w_t, s_t; \sigma_w, \sigma_o)$ is provided in Appendix A, equation (A.1).

The problem of the composite-input aggregator is static and does not require forming beliefs about the aggregate latent state.

2.1.4 Monopolistically competitive production good firms

A continuum of monopolistically competitive production good firms indexed by $j \in [0,1]$ transforms the homogeneous composite input good into differentiated production goods.

Specifically, each firm j purchases a quantity $X_t(j)$ of the homogeneous composite input good at price Q_t and produces a differentiated production good $Y_t(j)$ using a linear technology

$$Y_t(j) = X_t(j). \tag{24}$$

Their real marginal cost is common and given by

$$mc_t = Q_t. \tag{25}$$

2.1.5 Competitive consumption good producers

A continuum of competitive firms produces a composite consumption good using the continuum of differentiated production goods j as inputs,

$$Y_t = \left[\int_0^1 Y_t(j)^{\frac{\eta-1}{\eta}} dj\right]^{\frac{\eta}{\eta-1}}, \qquad \eta > 1, \tag{26}$$

where η is the elasticity of substitution across varieties. Demand for variety j is

$$Y_t(j) = \left(\frac{P_{q,t}(j)}{P_{q,t}}\right)^{-\eta} Y_t, \tag{27}$$

and the corresponding aggregate price index is

$$P_{q,t} = \left[\int_0^1 P_{q,t}(j)^{1-\eta} dj\right]^{\frac{1}{1-\eta}}. \tag{28}$$

Common marginal costs imply that in a symmetric equilibrium,

$$P_{q,t}(j) = P_{q,t}, \qquad Y_t(j) = Y_t, \qquad \Pi_{q,t}(j) = \Pi_{q,t} \qquad \forall j, \tag{29}$$

implying that there is no endogenous distribution of firm-specific output prices to track over time.

2.1.6 Price setting

The monopolistically competitive intermediate-good firms set prices subject to Rotemberg adjustment costs (Rotemberg, 1982). When changing its price, firm j pays a quadratic cost equal to

$$\frac{\phi_p}{2}\left(\frac{P_{q,t}(j)}{P_{q,t-1}(j)} - 1\right)^2 Y_t, \tag{30}$$

where $\phi_p \geq 0$ governs the degree of nominal rigidity.

Real profits, measured in units of the composite consumption good, are

$$d_t(j) = \frac{P_{q,t}(j)}{P_{q,t}} Y_t(j) - Q_t Y_t(j) - \frac{\phi_p}{2}\left(\frac{P_{q,t}(j)}{P_{q,t-1}(j)} - 1\right)^2 Y_t. \tag{31}$$

The problem of firm j is to maximize the expected discounted value of these profits:

$$\max_{\{P_{q,t+s}(j)\}_{s=0}^{\infty}} \widehat{\mathbb{E}}_{t^-} \sum_{s=0}^{\infty} \Lambda_{t,t+s}^q d_{t+s}(j) \tag{32}$$

subject to the demand curve (27), where the stochastic discount factor relevant for payoffs measured in units of the composite consumption good is

$$\Lambda_{t,t+1}^{q} = \beta \left(\frac{C_t}{C_{t+1}} \right)^{\sigma} \frac{p_t}{p_{t+1}}. \tag{33}$$

As all intermediate-good firms purchase the same composite input at the price Q_t, they face a common marginal cost, $mc_t = Q_t$. Thus dispersed local input-price information affects inflation through the equilibrium composite-input price Q_t and hence marginal cost, while the expectations entering the forward-looking pricing condition are common.

The forward-looking pricing decisions use common beliefs held at the beginning of time t to form expectations,

$$(1-\eta) + \eta Q_t - \phi_p \Pi_{q,t} \left(\Pi_{q,t} - 1 \right) + \phi_p \widehat{\mathbb{E}}_{t^-} \left[\Lambda_{t,t+1}^{q} \frac{Y_{t+1}}{Y_t} \Pi_{q,t+1} \left(\Pi_{q,t+1} - 1 \right) \right] = 0, \tag{34}$$

where $\Pi_{q,t} = \frac{P_{q,t}}{P_{q,t-1}}$.

2.1.7 Central bank

The central bank has full information for the policy decision and sets the nominal policy rate according to a Taylor rule. It observes aggregate output Y_t, inflation $\Pi_{q,t}$, the natural level of output Y_t^n, and therefore the output gap.

Let Y_t^n denote flexible-price output in the otherwise identical economy. Under flexible prices,

$$mc_t^n = Q_t^n = \frac{\eta - 1}{\eta}. \tag{35}$$

Natural output depends on the oil supply shock that the private sector does not observe. Therefore Y_t^n is also unobserved by private agents, while the central bank observes it when implementing monetary policy. Hence, the central bank observes the true output gap, $\log\left(\frac{Y_t}{Y_t^n}\right)$.

The central bank follows a Taylor rule:[2]

$$R_t = R^* \left(\frac{\Pi_{q,t}}{\Pi_q^*} \right)^{\phi_{\pi,t}} \left(\frac{Y_t}{Y_t^n} \right)^{\phi_y} \exp(\epsilon_t), \tag{36}$$

where the monetary policy shock is

$$\epsilon_t \sim N(0, \sigma_\epsilon^2) \tag{37}$$

and it is unobserved by the private sector.

The reaction coefficient to inflation $\phi_{\pi,t}$ is time-varying and also unobserved by private agents. It evolves according to

$$\phi_{\pi,t} = (1 - \rho_\phi)\overline{\phi}_\pi + \rho_\phi \phi_{\pi,t-1} + \xi_t, \qquad \xi_t \sim N(0, \sigma_\xi^2). \tag{38}$$

2.2 Market clearing

Markets clear.

Oil market:

$$O_t + C_{o,t} = S_t = \overline{S} \exp(v_t). \tag{39}$$

Labor market:

$$N_t = L_t. \tag{40}$$

Intermediate goods market:

$$X_t = \int_0^1 X_t(j)\, dj. \tag{41}$$

Production goods market:

$$Y_t = C_{q,t} + \frac{\phi_p}{2}(\Pi_{q,t} - 1)^2 Y_t. \tag{42}$$

Bond market:

$$B_t = 0. \tag{43}$$

[2]We assume that the central bank responds to core inflation $\Pi_{q,t}$ rather than headline inflation $\Pi_t = \frac{P_t}{P_{t-1}}$, as the latter introduces an additional state variable. Specifically, as $\Pi_t = \Pi_{q,t} \frac{p_t}{p_{t-1}}$, headline inflation also reflects the relative price $\frac{p_t}{p_{t-1}}$. Our assumption excludes a direct central bank response to oil price increases - mirroring central banking focus on measures of inflation that exclude the most volatile components - while capturing their impact on consumption good prices through higher production costs. Since oil constitutes only 2% of households' expenditures (Gagliardone and Gertler, 2026), this assumption likely has a minimal impact on our quantitative results.

Appendix A characterizes the equilibrium.

2.3 Learning: information structure and timing

Private agents do not observe the central bank's reaction to inflation $\phi_{\pi,t}$, nor the shock to the monetary policy rule, ϵ_t. In addition, they do not observe the oil supply shock, v_t, or the distribution of the supply shocks, g.

The aggregate latent state relevant for private sector learning is

$$\theta_t = (v_t, \phi_{\pi,t}, \epsilon_t)\,. \tag{44}$$

Private agents learn about the latent state from observing the realization of the macroeconomic variables,

$$Z_t = \begin{pmatrix} \log Y_t \\ \log \Pi_{q,t} \\ \log R_t \end{pmatrix}. \tag{45}$$

Because the private sector does not observe the supply shock, it cannot measure the natural output Y_t^n either, and therefore does not observe the output gap, $\log\left(\frac{Y_t}{Y_t^n}\right)$.

Let

$$\mathcal{I}_{t-1} = (Z_0, \ldots, Z_{t-1}) \tag{46}$$

denote common private sector information at the beginning of period t, so that

$$\widehat{\mathbb{E}}_{t^-}[X] \equiv \mathbb{E}[X \mid \mathcal{I}_{t-1}]. \tag{47}$$

The perceived oil shock density available at the beginning of t, when household and firms' decisions are made, is $\widehat{g}_{t-1}$. The inherited belief about the policy coefficient is summarized by the predictive marginal distribution $\mu^{\phi_\pi}_{t|t-1}$. Period-t private decisions are therefore made using beliefs inherited from period $t-1$.

After current Y_t, $\Pi_{q,t}$, and R_t are observed, the information available to private agents is updated to

$$\mathcal{I}_t = \{\mathcal{I}_{t-1}, Z_t\}\,, \tag{48}$$

and private agents form the posterior

$$\mu_{t|t}(\theta_t) = p(\theta_t \mid \mathcal{I}_t)\,. \tag{49}$$

The posterior means used to summarize perceptions are

$$\widehat{\phi}_{\pi,t} = E[\phi_{\pi,t} \mid \mathcal{I}_t], \tag{50}$$

$$\widehat{v}_t = E[v_t \mid \mathcal{I}_t]. \tag{51}$$

The full marginal posterior of $\phi_{\pi,t}$ is propagated through its AR(1) transition law to construct the predictive distribution used in period $t+1$. The posterior estimate $\widehat{v}_t$ is appended to the historical oil shock sample and used to update the nonparametric density $\widehat{g}_t$.

Timing. At the beginning of period t, private agents enter with $(\widehat{g}_{t-1}, \mu^{\phi_\pi}_{t|t-1})$. The shocks (v_t, ξ_t, ϵ_t) are realized. The central bank observes the contemporaneous aggregate allocation and natural output and hence knows the true output gap; it also knows the realizations of $\phi_{\pi,t}$ and ϵ_t. Private agents do not observe these latent objects. Family saving and intermediate-firm pricing decisions use the beliefs inherited from period $t-1$, and current prices and quantities are jointly determined in equilibrium.

At the end of the period, private agents observe $Z_t = (Y_t, \Pi_{q,t}, R_t)'$ and update their beliefs. The current posterior therefore affects decisions beginning in period $t+1$, not the already-determined period-t allocation.

Expectations and future learning. When forming one-period-ahead expectations, agents use the common belief state inherited at the beginning of the period as if it remained fixed over that forecast. They do not internalize future posterior updating inside current expectations. This is the anticipated-utility convention typically used in the recursive learning environment (Kreps, 1998).

3 Parameterization and solution algorithm

The model is calibrated at a quarterly frequency.

3.1 Parameterization

Table 1 summarizes the baseline parameter values. We discipline the calibration using a combination of standard values from the literature, empirical estimates of the monetary policy process, and moments directly related to the 2020 oil supply shock.

Table 1: Parameterization

Parameter	Value	Description
Preferences, production, and price setting		
β	0.987	Discount factor
σ	5.00	Inverse intertemporal elasticity of substitution
φ	1.00	Inverse Frisch elasticity of labor supply
η	6.00	Elasticity of substitution across production goods
κ	6.00	Elasticity of substitution across differentiated inputs
α	0.97	Labor weight in production
ζ	0.10	Elasticity of substitution between labor and oil
ϕ_p	350	Rotemberg price adjustment cost
σ_h	0.05	Std. dev. of worker-level wage dispersion
σ_w	0.05	Std. dev. of input-producer wage dispersion
σ_o	0.10	Std. dev. of input-producer oil-price dispersion
Monetary policy and beliefs		
Π_q^*	1.00	steady-state core inflation
σ_ϵ	0.003	Std. dev. of monetary policy shock
$\underline{\phi}_y$	0.50	Monetary policy response to the output gap
$\overline{\phi}_\pi$	1.50	Long-run inflation response coefficient
ρ_ϕ	0.89	Persistence of the inflation response coefficient
σ_ξ	0.16	Std. dev. of innovations to $\phi_{\pi,t}$
Household oil demand and oil supply		
χ	0.02	Weight of oil in household consumption
$\underline{\psi}$	0.02	Elasticity of substitution between oil and produced consumption
$\overline{S}$	0.049	Reference oil supply
λ	0.016	Scaling parameter for the oil supply series

Preferences, production, and price setting We set the quarterly discount factor to $\beta = 0.987$. The inverse Frisch elasticity of labor supply is set to $\varphi = 1$, a conventional value in quantitative macroeconomic models. We set the inverse intertemporal elasticity of substitution to $\sigma = 5$. While this is higher than the values commonly adopted in standard calibrations, the corresponding elasticity of 0.2 lies within the range of estimates found in the empirical consumption literature (e.g. Guvenen, 2006). The elasticity of substitution across differentiated goods is set to $\eta = 6$. Under CES aggregation, this implies a steady-state price markup of 20 percent, within the range commonly used in quantitative New Keynesian models. We set the elasticity of substitution across differentiated production inputs to the same value, $\kappa = 6$.

The production parameters governing the role of oil are disciplined by the evidence in Gagliardone and Gertler (2026). We set the labor weight to $\alpha = 0.97$, following their calibration. Furthermore, consistent with their evidence that strong complementarity between oil and other inputs is important to reproduce the response of real activity to oil shocks, we set the elasticity of substitution between labor and oil to a low value, $\zeta = 0.10$.

Price rigidity is governed by the Rotemberg adjustment-cost parameter, which we set at $\phi_p = 350$. Log-linearizing the pricing condition around the zero-inflation steady state (see Appendix B) facilitates the comparison with standard New Keynesian calibrations:

$$\widehat{\pi}_t = \beta \widehat{\mathbb{E}}_t \widehat{\pi}_{q,t+1} + \kappa_{mc} \widehat{mc}_t, \qquad \kappa_{mc} = \frac{\eta - 1}{\phi_p}.$$

Our calibration implies $\kappa_{mc} = \frac{5}{350} \simeq 0.014$, corresponding to a relatively flat short-run Phillips curve. This is consistent with estimates from DSGE models that imply a low sensitivity of inflation to marginal cost (e.g., Del Negro, Giannoni, and Schorfheide, 2015) and with recent evidence of a substantial flattening of the Phillips curve before the COVID-19 pandemic (Del Negro, Lenza, Primiceri, and Tambalotti, 2020).

The remaining dispersion parameters govern cross-sectional heterogeneity in local transaction prices. Empirical evidence suggests substantial price dispersion even for narrowly defined products and markets. For example, Kaplan and Menzio (2015) document a standard deviation of prices of approximately 19 percent for identical products within the same local market and quarter, while evidence using more narrowly defined transactions generally points to dispersion on the order of 5–10 percent even after controlling for product and market characteristics. We therefore choose relatively conservative values for these parameters, $\sigma_h = 0.05$, $\sigma_w = 0.05$, and $\sigma_o = 0.10$, generating moderate heterogeneity in transaction prices while limiting the quantitative role of this source of dispersion. We verify that our main results do not rely on this variation.

Monetary policy We normalize steady-state producer inflation to zero, $\Pi_q^* = 1$, implying a steady-state gross nominal interest rate $R^* = \beta^{-1}$. We set the standard deviation of the monetary policy shock to $\sigma_\epsilon = 0.003$, in line with estimates from quarterly DSGE models (e.g. Smets and Wouters, 2007). The long-run response of the policy rate to inflation is set to $\overline{\phi}_\pi = 1.5$, the conventional Taylor-rule value, and the coefficient on the output gap is set to $\phi_y = 0.5$. The latter is also the standard coefficient in the original Taylor-rule calibration.

We set the persistence and standard deviation of the innovation of the time-varying inflation coefficient to $\rho_\phi = 0.89$ and $\sigma_{\tilde{\xi}} = 0.16$, respectively. These values are based on estimates of public perceptions of monetary policy responsiveness estimated by Bauer, Pflueger, and Sunderam (2024b) for the US, using professional forecasts of interest rates and macroeconomic conditions. Although Bauer et al. (2024b) estimate the dynamics of perceived monetary policy responsiveness, we use their estimates to calibrate the AR(1) process for $\phi_{\pi,t}$. This is consistent with our assumption that agents perceive $\phi_{\pi,t}$ to evolve according to the same law of motion as the true policy coefficient.

Household oil demand and oil supply Following the quantitative evidence in Gagliardone and Gertler (2026), we set the household oil expenditure weight to $\chi = 0.02$. We also allow oil to enter both household consumption and production with low elasticities of substitution, $\psi = 0.02$, consistent with their finding of very strong complementarity between oil and other household consumption.

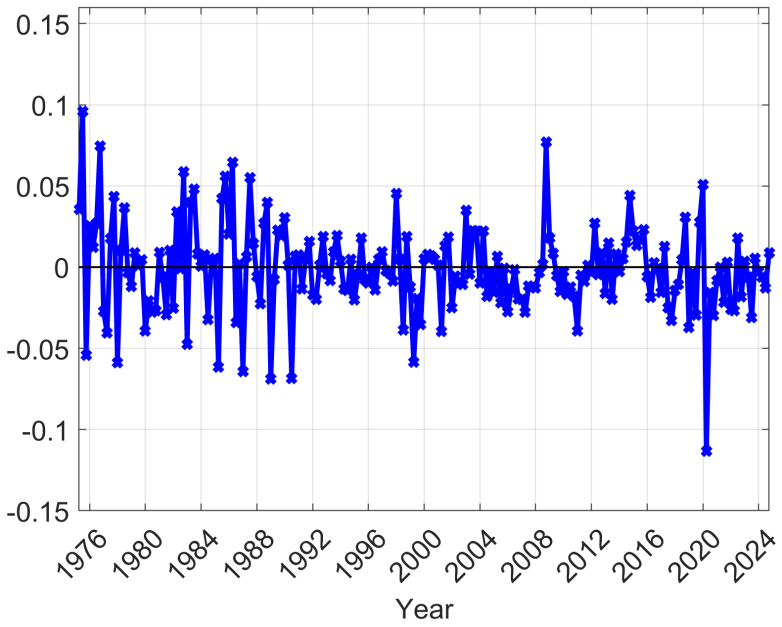
We set the reference level of oil supply to $\overline{S} = 0.049$, so that the steady-state aggregate oil expenditure relative to output, $\frac{s\overline{S}}{Y}$, is approximately 3 percent, consistent with the findings in Baumeister and Kilian (2016).

The historical oil supply shock series is based on the quarterly average of the identified oil supply shock series from Baumeister and Hamilton (2019) (see Figure 1). The parameter λ, which does not enter the equilibrium conditions directly, scales the empirical oil supply shock series into units of the model disturbance, v_t. We choose $\lambda = 0.016$ so that, conditional on the observed 2020Q2 oil supply shock, the model reproduces the magnitude of the contemporaneous increase in inflation attributed to the oil supply shock in the empirical analysis of section 5.

3.2 Solution and learning algorithm

At the beginning of period t, private agents inherit a belief state consisting of the perceived distribution $\widehat{g}_{t-1}$ of the oil supply shock and the predictive distribution $\mu_{t|t-1}^{\phi_\pi}$

Figure 1: Time series of oil supply shocks



Notes: The figure shows the historical oil supply shock series based on the identified oil supply shock from Baumeister and Hamilton (2019).

of the monetary policy coefficient. When making period-t decisions, agents take these beliefs as given when forming expectations of future outcomes. In particular, they do not internalize how future observations will affect their subsequent beliefs. After the period-t equilibrium is realized, agents update their beliefs using the newly observed data. The equilibrium is then re-solved under the updated belief state before period-$t+1$ decisions are made. This anticipated-utility treatment (Kreps, 1998) separates the nonlinear equilibrium problem conditional on a given belief state from the statistical learning problem and keeps the recursive learning problem computationally tractable.

Conditional on the inherited belief state, we solve the nonlinear equilibrium on grids over the latent oil shock, v_t, the monetary policy coefficient, $\phi_{\pi,t}$, and the monetary policy shock, ϵ_t. Expectations entering the household Euler equation and the forward-looking pricing condition are computed by integrating future equilibrium outcomes over the agents' current perceived distributions of the latent states. Because current decisions themselves depend on these expected future outcomes, the policy functions are obtained using an expectations fixed-point iteration. For the realized latent states, the resulting policy functions generate the actual period-t equilibrium.

Private agents subsequently observe aggregate output, inflation, and the nominal interest rate, $(Y_t, \Pi_{q,t}, R_t)$, but do not observe the oil supply shock, the monetary policy coefficient, the monetary policy shock, or the associated natural output. They therefore

use the observed aggregate outcomes to jointly infer the latent oil shock and monetary policy coefficient. For each candidate oil shock v, agents compute the corresponding natural level of output and hence a candidate-specific output gap. For each candidate pair (ϕ_π, v), the observed interest rate and the perceived Taylor rule then imply a monetary policy shock $\widehat{\epsilon}_t(\phi_\pi, v)$ that exactly rationalizes the observed policy rate:

$$\widehat{\epsilon}_t(\phi_\pi, v) = \log\left(\frac{R_t}{R^*}\right) - \phi_\pi \log\left(\frac{\Pi_{q,t}}{\Pi_q^*}\right) - \phi_y \log\left(\frac{Y_t}{Y_t^n(v)}\right). \tag{52}$$

Because the nonlinear mapping from the latent states to the aggregate observables need not be one-to-one, we do not invert this mapping exactly.[3] Instead, we implement inference using a nonlinear Bayesian filter. The posterior weights satisfy

$$\mu_{t|t}(\phi_\pi, v) \propto \mu_{t|t-1}^{\phi_\pi}(\phi_\pi)\, \widehat{g}_{t-1}(v)\, f_\epsilon(\widehat{\epsilon}_t(\phi_\pi, v))\, \mathcal{K}_t(\phi_\pi, v),$$

where $\mathcal{K}_t(\phi_\pi, v)$ is a Gaussian observation likelihood (see equation A.2 in Appendix) that assigns greater weight to candidate states whose model-implied inflation and output are closer to their observed values.[4] Thus, inference combines the inherited beliefs about the policy coefficient and the oil shock, the likelihood of the policy shock required to rationalize the observed interest rate, and the ability of each candidate state to reproduce the observed inflation and output outcomes.

Learning about the oil shock distribution is nonparametric (Kozlowski et al., 2020). The perceived density $\widehat{g}_{t-1}$ is constructed using an expanding window kernel density estimator (KDE) applied to the history of previously inferred oil shocks.[5] After filtering, the posterior mean $\widehat{v}_t$ is appended to the history of inferred shocks and the KDE

[3] In other words, different combinations of latent states can generate similar realizations of $(Y_t, \Pi_{q,t}, R_t)$. In particular, a more adverse supply disturbance and a weaker monetary policy response can both contribute to higher inflation, while different monetary policy innovations can rationalize the observed policy rate conditional on these candidate states. Agents therefore face a signal-extraction problem rather than a system of three equations that can be uniquely inverted. Nonetheless, this non-uniqueness in inference does not imply multiple equilibria: for a given latent state and inherited belief state, the equilibrium is uniquely determined.

[4] Since the observed policy rate pins down the policy innovation conditional on each candidate pair $(\phi_{\pi,t}, v_t)$, we compute the implied policy innovation for each pair and represent the posterior numerically over $(\phi_{\pi,t}, v_t)$.

[5] We follow Kozlowski et al. (2020) in estimating the perceived shock distribution nonparametrically using a kernel density estimator with an optimal bandwidth. A key difference is informational: in their setting, agents observe the realizations of the underlying shock and update the density using these observations. In our model, the corresponding oil supply shock v_t is latent. Agents therefore update the perceived distribution using the perceived shock $\widehat{v}_t$ inferred from the filtering step, based on the observables $(Y_t, \Pi_{q,t}, R_t)$, rather than using the true realization v_t.

is re-estimated, updating $\widehat{g}_{t-1}$ to $\widehat{g}_t$.[6] This feature is important following an unprecedented tail realization. When the true oil shock lies far in the tail of the perceived distribution – potentially beyond the range of shocks previously experienced – agents initially assign little probability to such a realization. As a result, their posterior estimate tends to understate the magnitude of the shock. Once the inferred shock is incorporated into the KDE, the perceived distribution adjusts, placing greater probability on adverse tail realizations.

Finally, the marginal posterior of the monetary policy coefficient is propagated through its perceived transition law to obtain $\mu^{\phi_\pi}_{t+1|t}(\phi_\pi)$. The updated oil shock distribution $\widehat{g}_t$ and the predictive distribution $\mu^{\phi_\pi}_{t+1|t}(\phi_\pi)$ constitute the belief state with which agents enter period $t+1$. The nonlinear equilibrium is then re-solved under these beliefs, and the procedure is repeated.

Appendix C provides further details on the numerical implementation.

4 Quantitative Analysis

We present a quantitative analysis to illustrate the model dynamics in response to a large negative supply shock under alternative informational assumptions.

We first explore the mechanism that generates private agents' misperceptions regarding both the severity of the supply shock and the monetary policy response.

We then compare the impulse responses to a tail disturbance under dual learning to those observed when private agents learn only about the shock distribution, only about the monetary policy rule, or when they operate under full information.

4.1 When the Unthinkable Happens

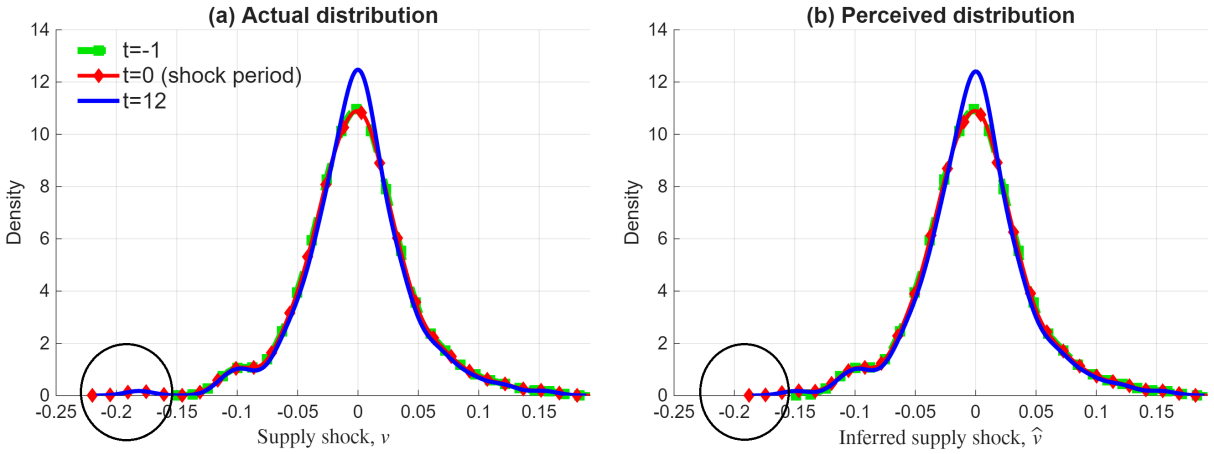
The model results depend crucially on the dual learning problem. We provide here an intuition of why this is the case.

In the model, agents estimate $\widehat{g}_t$ nonparametrically after observing $(Y_t, \Pi_{q,t}, R_t)$, using a kernel density estimator. Figure 2, panel (a), shows the actual supply shock distribution, namely the distribution that an econometrician would estimate after observing the history of realizations of the supply shock up to period t. Panel (b) shows instead the perceived distribution estimated by private agents in our model, based on

[6] Our baseline uses the posterior mean of the inferred shock to update the perceived shock distribution. Appendix C.6 considers an alternative specification that retains the full posterior uncertainty about v_t when updating the shock distribution. The resulting dynamics are shown in Figure A.3, and are very similar to the baseline results.

the macroeconomic variables observed at the end of each period t. The distributions are estimated at three different points in time: the period before the tail event, $t = -1$, the period in which the tail event occurs, $t = 0$, and 12 periods after the shock, $t = 12$, which coincides with the end of the sample covered in Figure 1.

Figure 2: Supply shock distributions



Notes: The left panel reports kernel density estimates of the actual supply shock, using an expanding window of realized shocks. The green line uses the pre-shock sample ($t = -1$); the red line adds the large supply shock at $t = 0$; and the blue line additionally includes the zero realizations from $t = 1$ through $t = 12$. The right panel is constructed analogously using agents' inferred shocks, $\widehat{v}_t$, from the joint-learning model.

Three observations stand out. First, when an exceptionally large negative supply shock occurs ($t = 0$), the left tail of both the actual and perceived distributions (red lines) increases relative to the pre-shock distributions (green lines). Even when agents do not observe v_t, they infer that a large negative supply shock has occurred, although they underestimate its magnitude. Second, the shift in both the actual and perceived distributions is persistent: at the end of the sample ($t = 12$), the left tail remains long (blue lines).[7] Third, under the baseline model where agents must jointly infer both the policy rule and the shock distribution, the perceived left tail of the supply-shock distribution remains smaller than the actual left tail.

The reason for this last observation is that agents do not observe the central bank's reaction to inflation either and slowly learn about it from observed outcomes. When they experience the inflation surge, they attribute it partly to a tail supply shock but partly also to a weaker monetary policy response. Consequently, the inferred supply

[7]Here we use an expanding window of observations to estimate the distribution. For robustness, we consider two alternative ways to estimate the shock distribution, and show that the results are very similar (see Figure A.1 in appendix). First, we use a rolling window of 25 years' observation. Second, we introduce a form of gradual forgetting into the learning process by assigning greater importance to more recent observations through exponential discounting.

shock $\widehat{v}_t$ is less negative than the true shock v_t. Because the perceived shock entering the kernel density estimator is attenuated relative to the true shock, the estimated distribution places less probability mass in the far left tail than the actual distribution. In other words, uncertainty about the policy rule causes agents to understate the severity of the realized supply disturbance when revising their beliefs.

A shock that was considered 'unthinkable' before the tail event — because it lay outside the perceived supply shock distribution — remains 'unthinkable' even after the event occurs, as the perceived distribution fails to align with the actual one.

4.2 Tail supply shocks and macroeconomic response

We analyse the response of the economy to a one-time negative supply shock of four standard deviations, similar in size to the oil supply shock observed in 2020Q2 (Figure 3). The shock hits the economy at time 0 and is immediately reabsorbed. To isolate the effect of the supply shock, we set the shock to the policy reaction coefficient, ξ_t, and the monetary policy shock, ε_t, to zero for the entire horizon of the analysis.

Figure 3: Actual shocks in impulse response functions



Our baseline economy is one under imperfect and asymmetric information, as described in Section 2. The private sector does not observe any shock and learns about the true $\phi_{\pi,t}$ and the shock distribution g by observing $(Y_t, \Pi_{q,t}, R_t)$.

To help dissect the mechanism, we also consider three counterfactuals. The first one is the full information model, where the agents observe all shocks as well as the shock distribution g. In the second counterfactual, the private sector does not observe $\phi_{\pi,t}$ or ε_t, but they do observe v_t and know the true shock distribution g. In the third counterfactual, the private sector observes $\phi_{\pi,t}$ and ε_t, but not v_t or g, and they continuously update their beliefs about g.

4.2.1 Baseline with dual learning

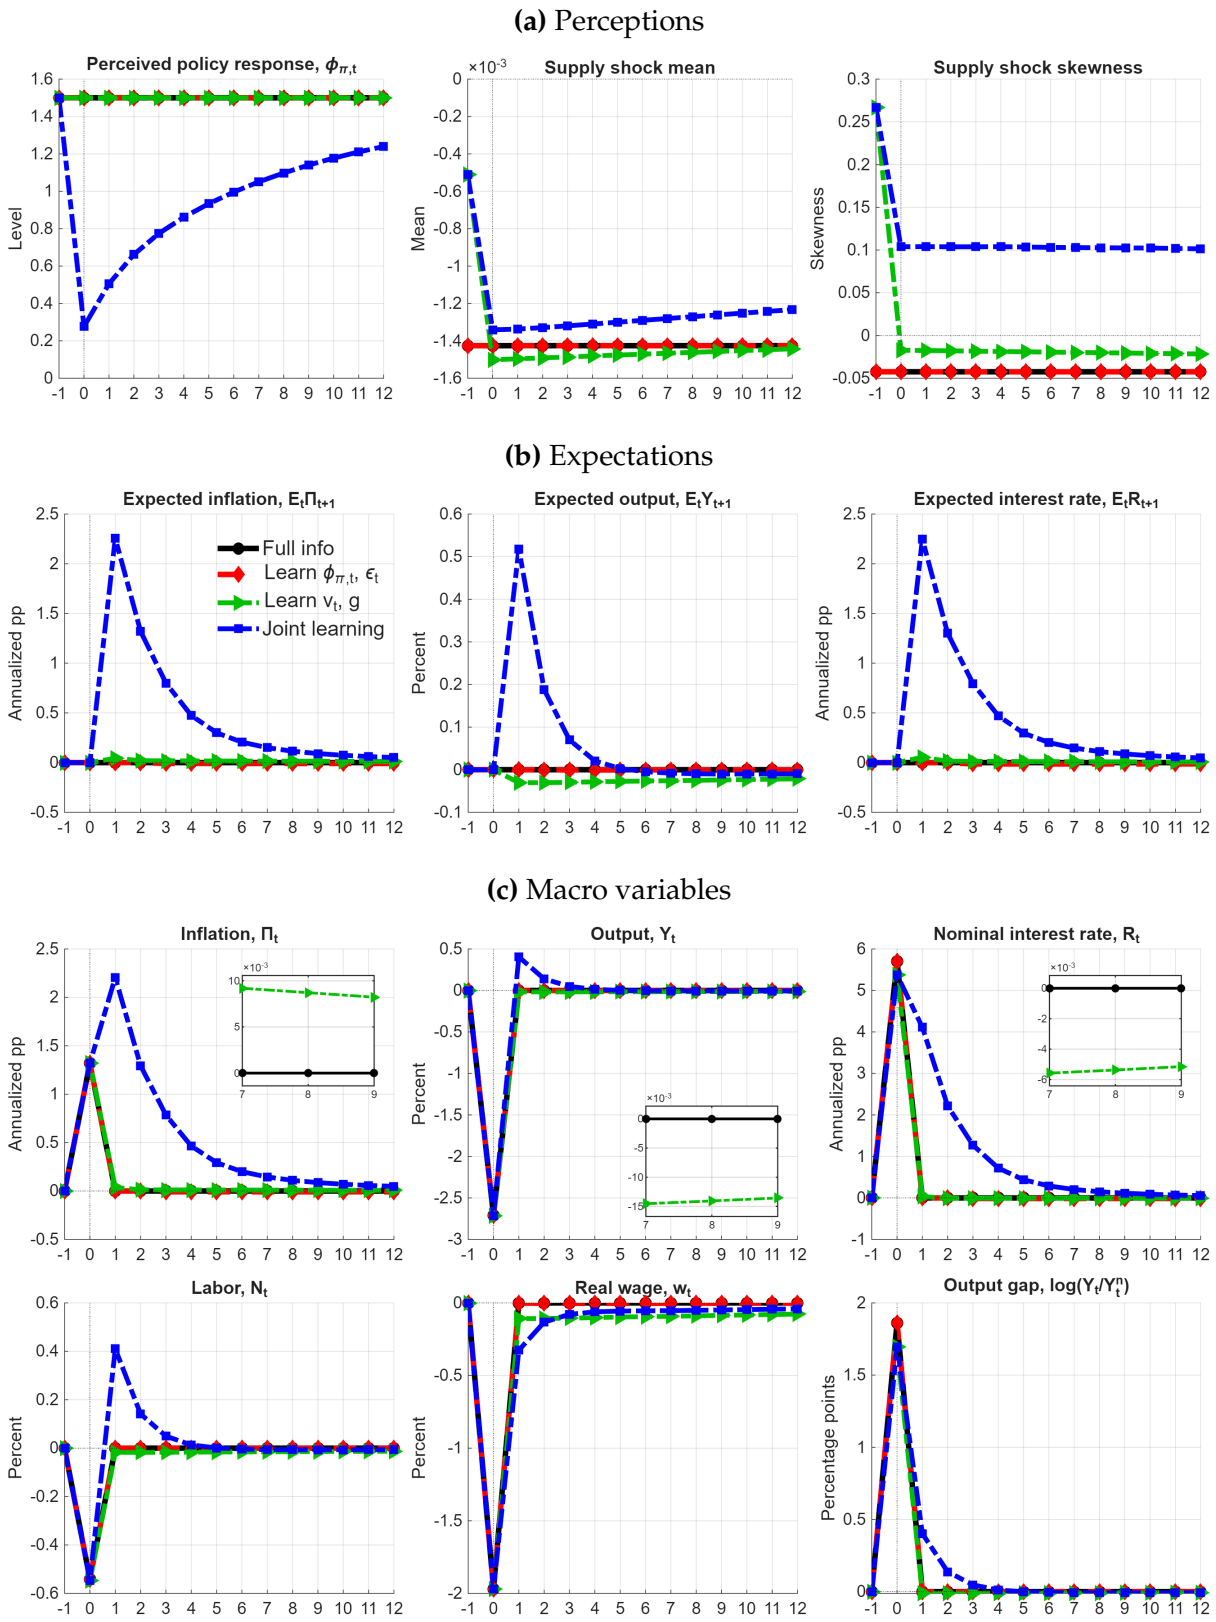
Figure 4 shows the impulse response functions following the four standard deviations negative supply shock. In the baseline case with imperfect information about both $\phi_{\pi,t}$ and g (blue dashed-dotted lines), a large negative shock outside the perceived distribution $\widehat{g}_{t-1}$ prompts agents to update their beliefs. However, since the private sector observes only aggregate macroeconomic outcomes, the source of high inflation after the shock is ambiguous to them. It could reflect either a more negative v_t or a lower $\phi_{\pi,t}$, and agents assign some probabilities to both possibilities. Moreover, when agents update their belief about the policy reaction to inflation, they still have in mind that the supply shock distribution follows the last period's estimate $\widehat{g}_{t-1}$. Thus, their estimate of today's supply shock $\widehat{v}_t$ is less negative than its true value and cannot fully account for the sharp rise in inflation. As a result, they mistakenly infer that $\phi_{\pi,t}$ is lower than it actually is, i.e. $\widehat{\phi}_{\pi,t} < \phi_{\pi,t}$.

In terms of macroeconomic outcomes, inflation remains persistently elevated after the shock, even though the shock itself lasts only one period. This persistence reflects the formation of beliefs. Agents form beliefs $\widehat{\phi}_{\pi,t}$ and update $\widehat{g}_{t-1}$ to $\widehat{g}_t$ only after observing the macroeconomic outcomes in period t. Consequently, the perception of weaker policy reaction affects expectations beginning in the following period. The resulting increase in expected inflation and demand raises current marginal costs and, through the forward-looking Rotemberg pricing equation, produces the hump-shaped inflation response. Since the perceived policy reaction coefficient approaches its true value only gradually, expected inflation also adjusts slowly. As shown in Figure A.4, expected inflation is lower conditional on higher values of the latent policy responsiveness coefficient $\phi_{\pi,t}$. Hence, shifting the inherited belief distribution toward lower values of $\phi_{\pi,t}$, as occurs after the tail event, raises expected inflation.[8]

Output falls sharply during the shock period but rises above its steady-state level thereafter. This is driven by higher output expectations, which, like expected inflation, are higher conditional on lower values of the policy responsiveness coefficient, and by the response of labor. After the oil supply shock disappears, the oil input returns close to its normal level, while labor rises above steady state. The labor overshoot reflects a negative wealth effect. Joint learning leads households to expect persistently unfa-

[8]Figure A.4 plots equilibrium outcomes conditional on alternative values of the latent policy responsiveness coefficient $\phi_{\pi,t}$, holding fixed the inherited belief state at each date. Since $\phi_{\pi,t}$ is unobserved, these curves should not be interpreted as policy functions with respect to agents' posterior estimate $\widehat{\phi}_{\pi,t}$. Rather, private sector expectations integrate these conditional outcomes over the perceived distribution of $\phi_{\pi,t}$. The policy functions themselves also shift over time as the inherited belief state, including the perceived supply-shock distribution $\widehat{g}_t$, evolves.

Figure 4: Model-based impulse responses following a large supply shock



Notes: All four experiments use the same realized shock paths: a large negative supply shock at $t = 0$ followed by $v_t = 0$, with $\phi_{\pi,t} = 1.5$ and $\epsilon_t = 0$ throughout. Under "full information" (black), agents know g and perfectly observe $(\phi_{\pi,t}, v_t, \epsilon_t)$. Under "learn $\phi_{\pi,t}$" (red), they do not observe $\phi_{\pi,t}$ or ε_t, but they do observe v_t and know the true distribution g. Under "learn g" (green), they observe $\phi_{\pi,t}$ and ϵ_t, but not v_t or g. Under "joint learning" (blue, baseline), agents do not observe any shock and learn about $\phi_{\pi,t}$ and g by observing $(Y_t, \Pi_{q,t}, R_t)$.

vorable future conditions, and consumption remains below steady state. With CRRA utility and separable labor, lower consumption raises the marginal utility of consumption and induces households to supply more labor. Rotemberg adjustment costs contribute to the divergence between produced output and household consumption by absorbing part of output as a resource cost when inflation is elevated.

The output overshooting is amplified by agents' expectations: agents expect a less hawkish future monetary policy and therefore anticipate stronger future economic activity. These higher expectations feed directly into the forward-looking pricing and Euler equations, sustaining elevated inflation and generating output overshooting even after the original supply shock has dissipated.

In response to the tail event, the real wage falls on impact and remains persistently below steady state. The movement is explained by the change in the household marginal rate of substitution between consumption and labor. Although labor overshoots after the initial decline, the decline in consumption more than offsets the increase in labor supply, keeping the real wage persistently below its pre-shock level.

While output falls on impact, the output gap rises and remains persistently positive. The oil disruption sharply lowers the flexible-price level of natural output. Under Rotemberg price adjustment costs, firms do not immediately implement the price changes required to reproduce the flexible-price allocation, so actual output contracts by less than natural output. Thereafter, the persistent underestimation of policy responsiveness sustains expected inflation and aggregate demand. Actual output — and therefore the output gap — consequently remains elevated even after the oil shock has disappeared.

Through the Taylor rule, the persistent rise in inflation and the output gap also keeps the nominal interest rate persistently above steady state after the initial shock.

Overall, the baseline model delivers two central results. First, inflation remains persistently above target long after the initial supply disturbance has disappeared. Second, output overshoots, while the output gap remains persistently positive. These macroeconomic responses emerge endogenously from the interaction between learning about the policy rule and learning about the shock distribution. They constitute the key predictions that we later bring to the data.

The macroeconomic response under dual learning is qualitatively different from the sum of the responses under learning about the shock distribution and the policy reaction. To illustrate this, it is instructive to shut down one or both sources of imperfect information, which we do next.

4.2.2 The full-information economy

In the full-information counterfactual (indicated with black lines in Figure 4), agents observe all shocks as well as the true distribution g and the policy reaction to inflation $\phi_{\pi,t}$, for all t.

Since agents correctly notice that the rise in inflation is entirely driven by a temporary supply disturbance, perceived policy responsiveness remains anchored at the true value, $\phi_{\pi,t} = 1.5$. As a result, inflation expectations and output expectations remain stable, and all macroeconomic variables quickly return to steady state once the shock dissipates.

4.2.3 Learning about the policy rule only

In the economy under imperfect information about the policy rule but full information about the supply shock distribution (indicated with red lines in Figure 4), agents do not observe $\phi_{\pi,t}$ or ε_t, but they observe the supply shock v_t and its true distribution g.[9]

In response to the large negative supply shock, private agents correctly attribute the unusually severe inflation surge to the supply disturbance and no longer misinterpret the inflation surge as evidence of a weaker policy response. Consequently, perceived policy responsiveness remains close to the true value, expected inflation stays anchored, and the persistence in inflation and the output gap largely disappears.

4.2.4 Learning about the supply shock distribution only

The counterfactual of full information about the policy rule and the policy shock but imperfect information about the supply shocks v_t and their distribution g is indicated with green dashed lines in Figure 4.

In this scenario, private agents correctly understand that the central bank responds aggressively to inflation, so there is no decline in perceived policy responsiveness. Moreover, once agents observe both $\phi_{\pi,t}$ and ε_t, the supply shock can be inferred accurately from the equilibrium conditions. Since the policy component of the interest-rate movement is perfectly observed, any remaining movement in inflation and output must be attributed to the supply disturbance itself. As a result, the perceived supply shock $\widehat{v}_t$ lies very closely to the true shock v_t, and the perceived shock distribution $\widehat{g}_t$ evolves exactly like the actual distribution g_t.

[9] We consider an alternative counterfactual in Figure A.2, where agents do not observe $\phi_{\pi,t}$, ε_t or v_t, but they know the true distribution g. The results are very similar to this counterfactual here. Thus, it is imperfect knowledge of g, rather than of the current shock v_t, that is responsible for the policy-rule misperception.

The persistently higher inflation and lower output gap arise from the persistence of the left tail in the distribution of supply shocks.[10] Thus, learning about the supply shocks and their distribution does not generate the output overshooting observed in the baseline model.[11]

4.3 Importance of a Tail Event

Figure 5 compares the impulse response functions following a small and a large negative supply shock in the baseline model. The small shock corresponds to a one-standard-deviation supply disturbance, while the large shock is the four-standard-deviation event analyzed in the previous section. To facilitate comparison across the two episodes, we normalize all impulse responses such that the impact increase in the real oil price is identical across the two experiments. This normalization allows us to isolate the role of nonlinear belief dynamics and expectation formation, rather than simple differences in shock size.

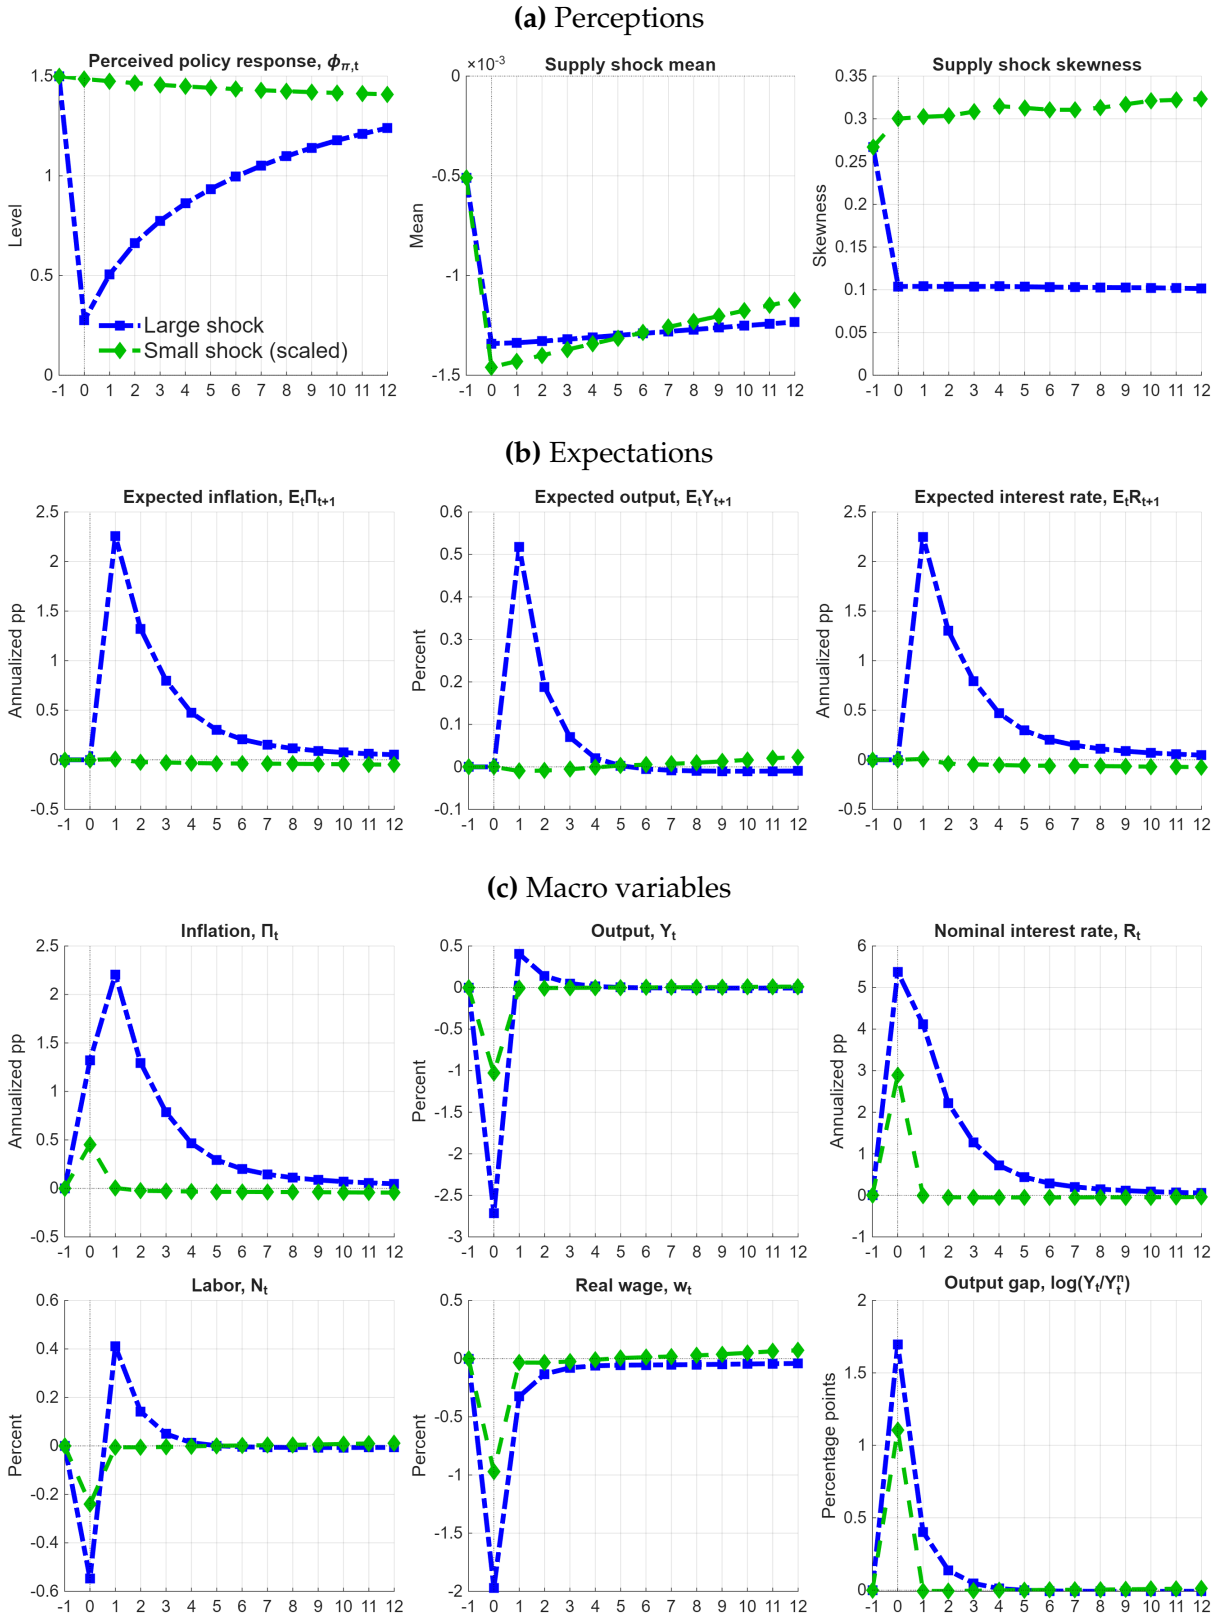
The figure shows that a large supply shock generates disproportionately larger and more persistent macroeconomic responses than a small shock. Following the small shock, agents' perceptions of the policy rule remain largely unchanged. The perceived inflation-response coefficient $\widehat{\phi}_{\pi,t}$ declines only modestly. As a result, expected inflation and expected output remain relatively stable, and inflation, output, and interest rates stay close to the steady state – though small deviations persist long after the shock has subsided. Notably, the output overshoot disappears.

By contrast, after the large supply shock, agents substantially revise downward their perceptions of the central bank's responsiveness to inflation. At the same time, they update their beliefs about the distribution of supply shocks, assigning greater probability to large adverse supply disturbances. These changes in beliefs lead to a persistent increase in inflation expectations and expected output. Since expectations

[10] As in Kozlowski et al. (2020), the tail event permanently shifts the estimated shock distribution, causing output to remain persistently below its full-information path. Quantitatively, however, the effect is small in our model. Kozlowski et al. (2020) use a production economy where firm decisions are highly sensitive to tail risk and where Epstein–Zin preferences and costly bankruptcy act as important amplifiers. These additional ingredients amplify the magnitude of the response, whereas its persistence comes from the lasting belief revision itself - consistent with our results. Notably, without uncertainty about the policy rule, the feedback loop between inflation outcomes and perceived monetary policy accommodation disappears.

[11] In the learning-g case, agents do not directly observe the supply shock v_t and update the perceived distribution using the inferred shock $\widehat{v}_t$. The results are very similar when agents observe v_t directly. Given known $\phi_{\pi,t}$ and ϵ_t, the joint observation of $(Y_t, \Pi_{q,t}, R_t)$ makes the remaining latent supply shock essentially a one-dimensional inversion problem, so agents can recover v_t quite accurately. Hence, replacing v_t with $\widehat{v}_t$ has little quantitative effect in this scenario.

Figure 5: Model-based impulse responses following a small vs. large supply shock



Notes: The figure compares the joint-learning responses to the large 2020Q2 supply shock with those to a one-standard-deviation negative supply shock. The small-shock responses are rescaled so that the impact increase in the real oil price is identical across the two experiments. Specifically, all small-shock responses are multiplied by the ratio of the large-shock to the small-shock impact response of the real oil price. For variables with nonzero pre-shock values, this factor is applied to their change from the pre-shock value.

enter directly into the Phillips curve and Euler equation, the economy exhibits persistent inflation dynamics and output overshooting even after the shock itself has dissipated.

Importantly, the amplification mechanism is highly nonlinear. The responses to the large shock are not simply scaled-up versions of the responses to the small shock. Instead, tail events generate qualitatively different dynamics because they induce much larger revisions in beliefs about both monetary policy and the shock distribution. In particular, output overshooting only emerges following the large shock. Likewise, inflation expectations become substantially more persistent only in the tail-event scenario. The non-linear responses arise because large shocks lie far into the perceived tail of the shock distribution, making it difficult for agents to disentangle whether elevated inflation reflects an unusually severe supply disturbance or a weaker monetary policy response.

5 Empirical Evidence

In this section, we assess whether the mechanisms highlighted in the model are borne out in the data. To this end, we estimate the nonlinear responses of U.S. macroeconomic variables as well as public perceptions about monetary policy and the shock distribution to an exogenous oil supply disruption. In particular, we contrast the dynamic effects of small versus large oil supply shocks.

5.1 Data Description

Our baseline empirical specification includes monthly data on the nominal interest rate, the output gap, inflation, and inflation expectations to characterize the U.S. economy as well as the coefficient on inflation in the monetary policy rule as perceived by the public, empirical measures of the perceived distribution of oil supply shocks, and the real price of crude oil.[12] For the nominal interest rate, we use the effective federal funds rate which we replace with the shadow rate proposed by Lombardi and Zhu (2018) to approximate the stance of monetary policy during the zero lower bound period.[13] The monthly output gap is defined as the (log) difference between real GDP and potential real GDP as estimated by the Congressional Budget Office. Since both series are only available at the quarterly frequency, they are interpolated to match our

[12]The data sources for the variables are provided in the appendix.

[13]We obtain very similar results using alternative shadow rate series or the unadjusted fed funds rate (see Figure A.6 as an example).

monthly sample. Cubic splines are used for potential real GDP, while the Chow and Lin (1971) method is applied for real GDP. For the latter, monthly data for industrial production, continued unemployment claims, a confidence indicator for manufacturing, and total nonfarm employment – all sourced from the FRED database – serve as higher-frequency indicators. Inflation is computed as the year-over-year change in the personal consumption expenditures price index, the Fed's preferred measure. Professional forecasts for headline consumer price (CPI) inflation from the Blue Chip survey form the basis for expected inflation. This choice is motivated by the fact that the same data source is used by Bauer, Pflueger, and Sunderam (2024b) to derive their time-varying estimate for the public's perceived short-run response of the Fed to inflation.[14] To summarize what people learn over time about the distribution of past oil supply shocks, we construct two measures based on the estimated time series of historical oil supply shocks of Baumeister and Hamilton (2019). Specifically, we use a sequence of oil supply shocks observed over a three-year moving window to compute the mean and an indicator of asymmetry that measures the distance of the 5% quantile and the 95% quantile from the median (see Baumeister, Huber, and Marcellino, 2024).[15] The real oil price is the West Texas Intermediate (WTI) spot price deflated by U.S. CPI which enters the model in growth rates.

5.2 Empirical Specification

To explore empirically whether the dynamic responses differ across small and large oil supply shocks, we rely on the nonlinear proxy-SVAR model proposed by Debortoli, Forni, Gambetti, and Sala (2023). This model is derived from a structural vector moving average representation of the economy which features both linear and nonlinear terms of the shock of interest. Debortoli et al. (2023) demonstrate that, under suitable conditions, a structural nonlinear model of this generic form can be written as a finite-order VARX model:

$$\mathbf{y}_t = \mathbf{c} + \boldsymbol{B}(\boldsymbol{L})\mathbf{y}_{t-1} + \alpha(L)u_t^{s,oil} + \beta(L)f(u_t^{s,oil}) + \boldsymbol{\varepsilon}_t$$

where $\mathbf{y}_t$ is an $(n \times 1)$ vector of endogenous variables, $\mathbf{c}$ is an $(n \times 1)$ vector of constants, and $\mathbf{y}_{t-1}$ contains p lags of the endogenous variables. The structural shock $u_t^{s,oil}$

[14] The data were retrieved from Carolin Pflueger's website.

[15] The perceived asymmetry measure is defined as $asy_t^{\tau} = (Q_t(1-\tau) - Q_t(0.5)) - (Q_t(0.5) - Q_t(\tau))$ where τ is set to 0.05 and Q_t refers to the quantile of the implied distribution of oil supply shocks. We also consider skewness as an alternative measure of the shock distribution, which yields qualitatively similar results (see Figures A.5 and A.6 in the appendix).

and its nonlinear transformation $f(u_t^{s,oil})$ enter as exogenous variables, $\alpha(L)$ and $\beta(L)$ are lag polynomials of order $q \leq p$, and ε_t is an $(n \times 1)$ vector of residuals. To estimate the VARX, a consistent estimate of the exogenous shock is needed. We follow the approach of Forni, Franconi, Gambetti, and Sala (2025) who treat the Baumeister and Hamilton (2019) oil supply shock as a valid instrument in a standard linear VAR model to obtain a noise-free shock estimate to be used in the VARX.[16] To differentiate between small and large shocks, we use the square of the shock as the nonlinear function: $f(u_t^{s,oil}) = (u_t^{s,oil})^2$.[17]

Impulse responses in this framework are obtained as the sum of the linear and nonlinear components. For a shock of magnitude $\overline{u}_M$, the generalized impulse response is thus given by

$$GIRF(h, u_t^{s,oil} = \overline{u}_M) = \alpha(L)\overline{u}_M + \beta(L)\overline{u}_M^2$$

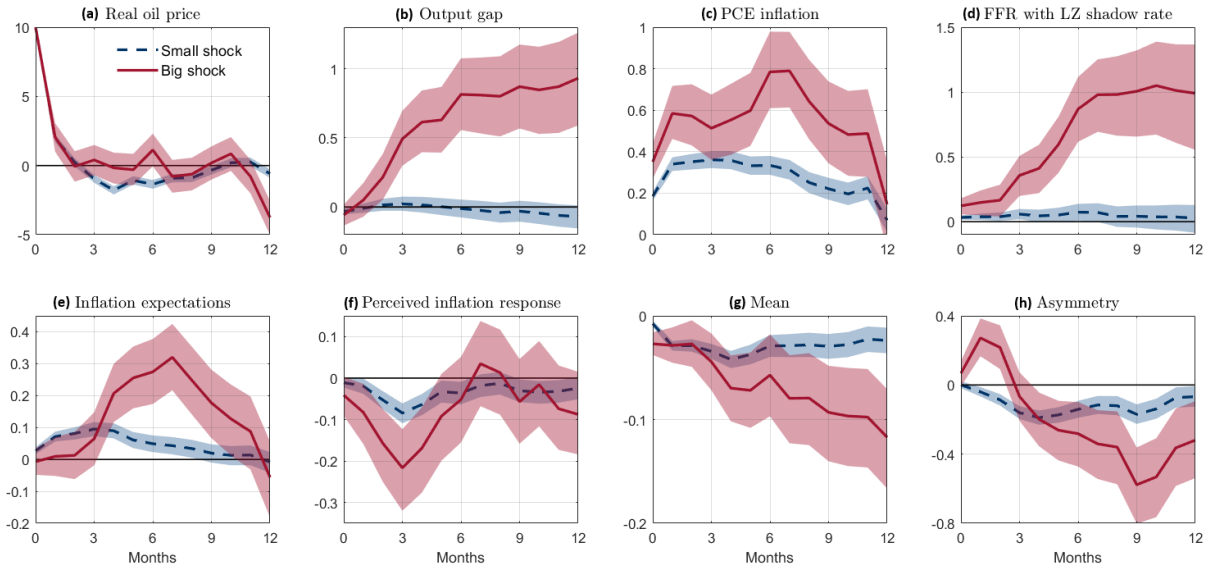
where $h = 0, 1, ..., h_{\max}$ is the response horizon with $h_{\max} \leq p$ in line with Baumeister (2025) and De Graeve and Westermark (2025) and $M = \{S, L\}$ with S referring to the small shock and L to the large shock. This formulation implies that responses scale nonlinearly with the size of the shock. In particular, while a small shock of one standard deviation yields a combined effect of $\alpha(L) + \beta(L)$, a larger shock amplifies both components, with the nonlinear term increasing quadratically. As a result, the model is able to capture the disproportionate impact of tail events.

The model is estimated over the period 1985M1 to 2019M12 where the start date is determined by the availability of the public's perceived inflation response. We end the sample before the onset of the COVID-19 pandemic to examine whether we can replicate the dynamics implied by the theoretical model without relying on the 2020 tail event in estimation. The sample period still contains several major oil supply disruptions such as the August 1990 invasion of Kuwait which corresponds to an eight-standard-deviation event, and the OPEC production-ceiling accord in early 1989, the joint OPEC and non-OPEC production cut in mid-1999, the Venezuelan general strike in late 2002, and the OPEC+ supply-cutting pact in early 2019, which ranged between three and five standard deviations.

[16] Specifically, they use the fitted value of a linear projection of the instrument on the residuals of a VAR(12) that contains the following six variables in logs: global oil production, the real price of oil, world industrial production, a measure of global oil inventories, U.S. industrial production and the U.S. consumer price index. Their interest centers on the nonlinear transmission of positive versus negative oil supply shocks to the macroeconomy.

[17] See Forni et al. (2024) for an application to financial shocks.

Figure 6: Empirical impulse responses to a small vs large oil supply shock



Notes: Impulse responses from the nonlinear VARX to small and large oil supply shocks, with bootstrapped one standard deviation error bands. The VARX includes 12 lags for both endogenous and exogenous variables. 'Small' and 'large' shocks are one and six standard deviations, respectively. The impulse responses are normalized to induce a 10% increase in real oil prices upon impact of the shock.

5.3 Main Findings

Figure 6 shows the impulse responses from the nonlinear VARX to oil supply shocks, with one standard deviation error bands.[18] We consider small and large shocks, which are one and six standard deviations, respectively.[19] To make the impulse responses directly comparable, we normalize them to induce a 10% increase in real oil prices upon impact of the shock (Panel (a)) so that we can focus on the role of nonlinearities. As shown in Panels (g) and (h) of Figure 6, both the mean and the asymmetry measure indicate a leftward shift of the perceived supply shock distribution, especially following a large shock.

A central prediction of the model is that large negative supply shocks generate systematic misperceptions of monetary policy, whereby the public underestimates the central bank's responsiveness to inflation. As shown in Figure 5, this manifests as a decline in perceived policy responsiveness, even though the true policy rule remains unchanged. The VARX evidence closely mirrors this pattern. Following a large supply shock, the impulse response functions from the nonlinear VARX also show a decline in perceived inflation responsiveness (Figure 6, Panel (f)), indicating that agents in-

[18]We include 12 lags for both endogenous and exogenous variables. The confidence bands are bootstrapped as described in Debortoli et al. (2023)

[19]The size of the 2020Q2 shock amounts to four standard deviations over the full sample considered in the model and six standard deviations over the estimation sample.

terpret the inflation increase as reflecting a weaker monetary policy stance; in other words, they perceive the tightening as insufficient. As a result, expected inflation rises persistently, as shown in Panel (e): Blue Chip inflation expectations rise significantly and remain elevated for several months after the shock. This aligns with the prediction of the model that a perceived weakening of policy leads to higher expected inflation, amplifying the initial inflationary impact of the supply disturbance.

The empirical evidence also supports our model predictions for the macro variables. Inflation increases significantly and remains elevated (Panel (c)), while the nominal interest rate rises gradually and persistently (Panel (d)). The sustained elevation of inflation is consistent with the persistent increase in expected inflation, which continues to amplify inflation dynamics even well after the shock has dissipated. Furthermore, echoing the model predictions, the output gap increases persistently after a temporary dip upon impact of the negative supply shock (Panel (b)).

Figure 6 also highlights the different impact of small versus large shocks, with the latter triggering larger changes in agents' beliefs about the policy rule (Panel (f)) and shock distribution (Panels (g) and (h)). Moreover, the responses of inflation, expected inflation, the output gap and the nominal interest rate are substantially stronger and more persistent than those to small shocks. These asymmetric patterns are consistent with the results in the model (Figure 5), where tail events disproportionately affect both agents' beliefs (about the shock distribution and policy behavior) and aggregate outcomes.

Taken together, the empirical IRFs validate the key elements of the model mechanism: (i) large supply shocks lead to a systematic underestimation of policy responsiveness, (ii) this misperception raises inflation expectations, and (iii) the resulting expectation dynamics generate persistent increases in inflation and the output gap.

6 Conclusion

This paper examines how rare and severe shocks reshape public perceptions of monetary policy and, in turn, affect macroeconomic dynamics. When agents must infer the central bank's response to large but infrequent shocks, they tend to underestimate the strength of policy reactions to inflation. Embedding this mechanism in a New Keynesian model, we show that belief updating creates an endogenous expectation channel that generates persistent inflation and temporary output overshooting after tail events. Evidence based on U.S. data supports this mechanism.

Our findings have important implications for monetary policy design and commu-

nication. When the unthinkable happens, private agents' beliefs about the policy rule are more likely to diverge from reality, generating persistent fluctuations that policy actions alone cannot immediately offset. Central banks operating in such environments may need to address not only current economic conditions but also the public's evolving perceptions of policy strength and credibility. Managing these perceptions – through clear communication, consistent responses, and credible frameworks – becomes a central component of stabilization policy when the economy faces the unthinkable.

References

Ball, Laurence, Daniel Leigh, and Prachi Mishra (2022). Understanding US inflation during the COVID-19 era. *Brookings Papers on Economic Activity 2022*(2), 1–80.

Bauer, Michael D, Carolin E Pflueger, and Adi Sunderam (2024a). Changing Perceptions and Post-Pandemic Monetary Policy. In *Jackson Hole Economic Policy Symposium Proceedings*, pp. 226–285.

Bauer, Michael D, Carolin E Pflueger, and Adi Sunderam (2024b). Perceptions about Monetary Policy. *Quarterly Journal of Economics 139*(4), 2227–2278.

Bauer, Michael D and Eric T Swanson (2023a). A Reassessment of Monetary Policy Surprises and High-Frequency Identification. *NBER Macroeconomics Annual 37*(1), 87–155.

Bauer, Michael D and Eric T Swanson (2023b). An Alternative Explanation for the "Fed Information Effect". *American Economic Review 113*(3), 664–700.

Baumeister, Christiane (2025). Local Projections or VARs? A Primer for Macroeconomists: Comment. In *NBER Macroeconomics Annual, Volume 40*.

Baumeister, Christiane and James D Hamilton (2019). Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks. *American Economic Review 109*(5), 1873–1910.

Baumeister, Christiane, Florian Huber, and Massimiliano Marcellino (2024). Risky Oil: It's All in the Tails. NBER Working Paper No. 32524.

Baumeister, Christiane and Lutz Kilian (2016). Lower Oil Prices and the US Economy: Is This Time Different? *Brookings Papers on Economic Activity 2016*(2), 287–357.

Benigno, Pierpaolo and Gauti B Eggertsson (2023). It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve. NBER Working Paper No. 31197.

Bianchi, Francesco, Martin Lettau, and Sydney C Ludvigson (2022). Monetary Policy and Asset Valuation. *Journal of Finance 77*(2), 967–1017.

Bianchi, Francesco, Sydney C Ludvigson, and Sai Ma (2022). Belief Distortions and Macroeconomic Fluctuations. *American Economic Review 112*(7), 2269–2315.

Bocola, Luigi, Alessandro Dovis, Kasper Jørgensen, and Rishabh Kirpalani (2025). Monetary Policy without an Anchor. NBER Working Paper No. 34436.

Bodenstein, Martin, James Hebden, and Fabian Winkler (2022). Learning and Misperception of Makeup Strategies. *Journal of Economic Dynamics and Control 139*, 104417.

Bullard, James and Kaushik Mitra (2002). Learning about Monetary Policy Rules. *Journal of Monetary Economics 49*(6), 1105–1129.

Chow, Gregory and An-loh Lin (1971). Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series. *The Review of Economics and Statistics*, 372–375.

Cieslak, Anna (2018). Short-Rate Expectations and Unexpected Returns in Treasury Bonds. *Review of Financial Studies 31*(9), 3265–3306.

Comin, Diego, Robert C. Johnson, and Callum Jones (2023). Supply Chain Constraints and Inflation. NBER Working Paper No. 31179.

De Graeve, Ferre and Andreas Westermark (2025). Long-Lag VARS. *Journal of Monetary Economics*, 103831.

Debortoli, Davide, Mario Forni, Luca Gambetti, and Luca Sala (2023). Asymmetric Monetary Policy Tradeoffs. CEPR Discussion Paper 18438.

Del Negro, Marco, Marc P. Giannoni, and Frank Schorfheide (2015). Inflation in the Great Recession and New Keynesian Models. *American Economic Journal: Macroeconomics 7*(1), 168–196.

Del Negro, Marco, Michele Lenza, Giorgio E Primiceri, and Andrea Tambalotti (2020). What's Up with the Phillips Curve? *Brookings Papers on Economic Activity*.

di Giovanni, Julian, Şebnem Kalemli-Özcan, Álvaro Silva, and Muhammed A. Yildirim (2022). Global Supply Chain Pressures, International Trade, and Inflation. NBER Working Paper No. 30240.

Eusepi, Stefano and Bruce Preston (2010). Central Bank Communication and Expectations Stabilization. *American Economic Journal: Macroeconomics 2*(3), 235–271.

Evans, George W and Seppo Honkapohja (2003). Adaptive Learning and Monetary Policy Design. *Journal of Money, Credit and Banking*, 1045–1072.

Farmer, Leland E, Emi Nakamura, and Jón Steinsson (2024). Learning about the Long Run. *Journal of Political Economy 132*(10), 3334–3377.

Forni, Mario, Alessandro Franconi, Luca Gambetti, and Luca Sala (2025). Asymmetric Transmission of Oil Supply News. *Quantitative Economics 16*(3), 947–979.

Forni, Mario, Luca Gambetti, Nicolò Maffei-Faccioli, and Luca Sala (2024). Nonlinear Transmission of Financial Shocks: Some New Evidence. *Journal of Money, Credit and Banking 56*(1), 5–33.

Gagliardone, Luca and Mark Gertler (2026). Oil Prices, Monetary Policy and Inflation Surges. *American Economic Journal: Macroeconomics (Forthcoming)*.

Gáti, Laura (2026). Away from FIRE. Working Paper.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning (2022). Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages? *American Economic Review 112*(5), 1437–1474.

Guvenen, Fatih (2006). Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective. *Journal of Monetary Economics 53*(7), 1451–1472.

Johannes, Michael, Lars A Lochstoer, and Yiqun Mou (2016). Learning about Consumption Dynamics. *Journal of Finance 71*(2), 551–600.

Kaplan, Greg and Guido Menzio (2015). The Morphology of Price Dispersion. *International Economic Review 56*(4), 1165–1206.

Kozlowski, Julian, Laura Veldkamp, and Venky Venkateswaran (2020). The Tail That Wags the Economy: Beliefs and Persistent Stagnation. *Journal of Political Economy 128*(8), 2839–2879.

Kreps, David M (1998). Anticipated utility and dynamic choice. *Econometric Society Monographs 29*, 242–274.

Leduc, Sylvain, Kevin Moran, and Robert J Vigfusson (2023). Learning in the Oil Futures Markets: Evidence and Macroeconomic Implications. *Review of Economics and Statistics 105*(2), 392–407.

Lombardi, Marco J and Feng Zhu (2018). A Shadow Policy Rate to Calibrate US Monetary Policy at the Zero Lower Bound. *International Journal of Central Banking*.

Lorenzoni, Guido and Iván Werning (2023). Wage-Price Spirals. *Brookings Papers on Economic Activity 2023*(2), 317–393.

Pflueger, Carolin (2025). Back to the 1980s or Not? The Drivers of Inflation and Real Risks in Treasury Bonds. *Journal of Financial Economics 167*, 104027.

Rotemberg, Julio J (1982). Sticky Prices in the United States. *Journal of Political Economy 90*(6), 1187–1211.

Schmeling, Maik, Andreas Schrimpf, and Sigurd AM Steffensen (2022). Monetary Policy Expectation Errors. *Journal of Financial Economics 146*(3), 841–858.

Schorfheide, Frank (2005). Learning and Monetary Policy Shifts. *Review of Economic Dynamics 8*(2), 392–419.

Smets, Frank and Rafael Wouters (2007). Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach. *American Economic Review 97*(3), 586–606.

Appendix

A Equilibrium

We now define the recursive equilibrium conditional on private sector beliefs. Let

$$\mathcal{B}_t = \left(\widehat{g}_{t-1}, \mu_{t|t-1}^{\phi_\pi}\right)$$

denote the belief state inherited at the beginning of period t, where $\widehat{g}_{t-1}$ is the perceived distribution of the oil supply shock and $\mu_{t|t-1}^{\phi_\pi}$ is the predictive distribution of the monetary policy coefficient. Let the realized latent state be

$$\theta_t = (v_t, \phi_{\pi,t}, \epsilon_t).$$

Definition. Given the inherited belief state $\mathcal{B}_t$, the perceived distribution of shocks, and the perceived laws of motion for the exogenous states, a recursive equilibrium consists of policy functions for quantities and prices,

$$\left\{C, C_q, C_o, N, L, O, Y, X, p, s, w, Q, \Pi_q, R, Y^n\right\},$$

together with a private sector belief-updating rule, such that, for every admissible $(\mathcal{B}_t, \theta_t)$, the following conditions hold.

1. *Household optimization.* Given prices and the inherited belief state, household choices satisfy the conditional demand equations

$$C_{q,t} = (1-\chi)\left(\frac{1}{p_t}\right)^{-\psi} C_t, \qquad C_{o,t} = \chi\left(\frac{s_t}{p_t}\right)^{-\psi} C_t,$$

where

$$p_t = \left[\chi s_t^{1-\psi} + (1-\chi)\right]^{\frac{1}{1-\psi}},$$

and the intertemporal optimality condition is

$$1 = R_t \beta C_t^{\sigma} p_t \widehat{\mathbb{E}}_{t^-}\left[\frac{C_{t+1}^{-\sigma}}{p_{t+1}\Pi_{q,t+1}}\right].$$

Aggregate labor supply satisfies

$$w_t = p_t C_t^{\sigma} N_t^{\varphi}.$$

2. *Optimal input demand and aggregation.* For each input producer k, $(L_{k,t}, O_{k,t})$ minimizes the cost of producing $X_{k,t}$ subject to

$$X_{k,t} = \left[\alpha^{1/\zeta} L_{k,t}^{1-1/\zeta} + (1-\alpha)^{1/\zeta} O_{k,t}^{1-1/\zeta}\right]^{\frac{1}{1-1/\zeta}}.$$

The competitive input aggregator minimizes the cost of producing X_t using the differentiated inputs $\{X_{k,t}\}_{k\in[0,1]}$. The resulting composite-input price Q_t equals the real marginal cost faced by intermediate-good firms,

$$mc_t = Q_t,$$

where

$$Q_t = \left[\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\left\{\alpha\left(w_t e^{\eta_{k,t}^w}\right)^{1-\zeta} + (1-\alpha)\left(s_t e^{\eta_{k,t}^o}\right)^{1-\zeta}\right\}^{\frac{1-\kappa}{1-\zeta}} f_w(\eta_{k,t}^w) f_o(\eta_{k,t}^o)\, d\eta_{k,t}^w\, d\eta_{k,t}^o\right]^{\frac{1}{1-\kappa}} \tag{A.1}$$

where f_w and f_o denote the Gaussian densities of $\eta_{k,t}^w$ and $\eta_{k,t}^o$, respectively.

3. *Optimal price setting.* Given marginal cost and the inherited belief state, production goods firms optimally choose prices. In a symmetric equilibrium, their first-order condition is

$$(1-\eta) + \eta Q_t - \phi_p \Pi_{q,t}(\Pi_{q,t} - 1) + \phi_p \widehat{\mathbb{E}}_{t^-}\left[\Lambda_{t,t+1}^q \frac{Y_{t+1}}{Y_t}\Pi_{q,t+1}(\Pi_{q,t+1} - 1)\right] = 0,$$

where

$$\Lambda_{t,t+1}^q = \beta\left(\frac{C_t}{C_{t+1}}\right)^{\sigma}\frac{p_t}{p_{t+1}}.$$

4. *Market clearing.* Oil supply satisfies

$$S_t = \overline{S}\exp(v_t),$$

and the oil market clears:

$$O_t + C_{o,t} = S_t.$$

The labor market clears,

$$N_t = L_t,$$

and the produced-good market clears:

$$Y_t = C_{q,t} + \frac{\phi_p}{2}(\Pi_{q,t} - 1)^2 Y_t.$$

5. *Natural allocation.* For every realization of the oil supply shock v_t, natural output Y_t^n is the output associated with the equilibrium of the otherwise identical flexible-price economy. The flexible price equilibrium satisfies the condition for marginal cost

$$Q_t^n = \frac{\eta - 1}{\eta},$$

together with household optimality conditions, optimal input demand, and oil, labor, and goods-market clearing.

6. *Monetary policy.* Given the true natural allocation, the central bank sets the nominal interest rate according to

$$R_t = R^* \left(\frac{\Pi_{q,t}}{\Pi_q^*} \right)^{\phi_{\pi,t}} \left(\frac{Y_t}{Y_t^n} \right)^{\phi_y} \exp(\epsilon_t),$$

where

$$\phi_{\pi,t} = (1 - \rho_\phi)\overline{\phi}_\pi + \rho_\phi \phi_{\pi,t-1} + \xi_t.$$

The central bank observes the contemporaneous latent state and Y_t^n when implementing this rule.

7. *Expectations are consistent with private sector beliefs.* All expectations entering household saving and firms' price-setting conditions are computed using the policy functions of the equilibrium and the inherited belief state $\mathcal{B}_t$. Under anticipated utility, agents hold this belief state fixed when evaluating one-period-ahead expectations.

8. *Beliefs are updated consistently with equilibrium outcomes.* After the period-t equilibrium has generated $(Y_t, \Pi_{q,t}, R_t)$, private agents update their beliefs about $(v_t, \phi_{\pi,t}, \epsilon_t)$ using the filtering rule described in Appendix C. The resulting posterior determines $\widehat{v}_t$ and the marginal posterior of $\phi_{\pi,t}$. The former is incorporated into the updated oil-shock density $\widehat{g}_t$, while the latter is propagated through the AR(1) transition law. The resulting belief state

$$\mathcal{B}_{t+1} = \left(\widehat{g}_t, \mu_{t+1|t}^{\phi_\pi} \right)$$

is the belief state used in period $t+1$.

Thus, conditional on the inherited belief state, quantities and prices satisfy private optimality, firm optimality, monetary policy, and all market-clearing conditions. The equilibrium outcomes in turn generate the public information used to update beliefs, and the resulting belief state determines expectations and equilibrium decisions in the following period.

B Mapping Rotemberg Cost into the Phillips-Curve Slope

Starting from the first-order condition for price setting, evaluated in a symmetric equilibrium in which $P_{q,t}(j) = P_{q,t}$ for all j:

$$1 - \eta + \eta mc_t - \phi_p(\Pi_{q,t} - 1)\Pi_{q,t} + \widehat{\mathbb{E}}_{t^-}\left[\Lambda^q_{t,t+1}\frac{Y_{t+1}}{Y_t}\phi_p(\Pi_{q,t+1} - 1)\Pi_{q,t+1}\right] = 0.$$

Log-linearize this around the zero-inflation steady state, $\overline{\Pi}_q = 1$. In steady state, the price-adjustment terms vanish, so the pricing condition implies

$$1 - \eta + \eta\overline{mc} = 0,$$

and hence

$$\overline{mc} = \frac{\eta - 1}{\eta}.$$

Let

$$\widehat{mc}_t \equiv \log\left(\frac{mc_t}{\overline{mc}}\right), \qquad \widehat{\pi}_{q,t} \equiv \log\Pi_{q,t}.$$

Since

$$mc_t \simeq \overline{mc}(1 + \widehat{mc}_t),$$

the marginal-cost component of the pricing condition satisfies

$$\begin{aligned} 1 - \eta + \eta mc_t &\simeq 1 - \eta + \eta\overline{mc} + \eta\overline{mc}\widehat{mc}_t \\ &= (\eta - 1)\widehat{mc}_t. \end{aligned}$$

For the current inflation term, log linearization around $\overline{\Pi}_q = 1$ gives

$$(\Pi_{q,t} - 1)\Pi_{q,t} \simeq \widehat{\pi}_{q,t}.$$

Similarly,

$$(\Pi_{q,t+1} - 1)\Pi_{q,t+1} \simeq \widehat{\pi}_{q,t+1}.$$

Because the latter expression is zero in steady state, the other terms multiplying it can be evaluated at their steady-state values. Therefore,

$$\widehat{\mathbb{E}}_{t^-}\left[\Lambda^q_{t,t+1}\frac{Y_{t+1}}{Y_t}\phi_p(\Pi_{q,t+1}-1)\Pi_{q,t+1}\right] \simeq \phi_p\widehat{\mathbb{E}}_{t^-}\left[\widehat{\pi}_{q,t+1}\right].$$

The log-linearized pricing condition is consequently

$$(\eta-1)\widehat{mc}_t - \phi_p\widehat{\pi}_{q,t} + \beta\phi_p\widehat{\mathbb{E}}_{t^-}\left[\widehat{\pi}_{q,t+1}\right] = 0.$$

Rearranging gives

$$\widehat{\pi}_{q,t} = \beta\widehat{\mathbb{E}}_{t^-}\left[\widehat{\pi}_{q,t+1}\right] + \kappa_{mc}\widehat{mc}_t,$$

where

$$\kappa_{mc} = \frac{\eta-1}{\phi_p}.$$

With $\eta = 6$ and $\phi_p = 350$, this gives

$$\kappa_{mc} = \frac{5}{350} \simeq 0.0143.$$

C Algorithm

This appendix describes the numerical solution and learning algorithm used to compute the model.

C.1 Overview

The computation has two nested components. First, for a given private sector belief state, the nonlinear equilibrium is solved on grids for the latent oil shock, the monetary policy coefficient, and the policy shock. Second, after the current aggregate outcomes $(Y_t, \Pi_{q,t}, R_t)$ are observed, private agents use these policy functions to update their beliefs about the current oil shock and the policy coefficient. The updated beliefs are then carried into the next period.

A key feature of the implementation is that the perceived oil-shock distribution is estimated nonparametrically. Its numerical support is therefore allowed to change over time. In particular, private agents evaluate the current oil shock only over the support implied by the KDE estimated from information available before the current observation.

C.2 Initialization

Denote the initial belief state as:

$$\mathcal{B}_t = \left(\widehat{g}_{t-1}, \mu^{\phi_\pi}_{t|t-1}\right),$$

where $\widehat{g}_{t-1}$ is the perceived oil-shock density inherited from the previous period and $\mu^{\phi_\pi}_{t|t-1}$ is the predictive distribution of the monetary policy coefficient.

Before the simulation begins, the historical oil-shock observations are used to construct the initial perceived density $\widehat{g}_{-1}$ using a kernel density estimator. This density provides the initial private sector beliefs about the distribution of oil shocks.

The initial distribution of the monetary policy coefficient is discretized on a finite grid. The prior is centered around the long-run policy coefficient, while its dispersion is specified separately from the transition innovation. The transition matrix for the policy coefficient is constructed from its AR(1) process.

The remaining exogenous disturbances are represented by finite numerical grids and quadrature weights. The deterministic steady state is used to initialize the nonlinear policy functions.

C.3 Solving the equilibrium for a given belief state

At the beginning of each period, private agents have an inherited belief state consisting of the perceived oil-shock density and the predictive distribution of the policy coefficient. These beliefs determine the expectations entering the household Euler equation and the forward-looking pricing condition. More generally, for any future equilibrium object X_{t+1}, expectations take the form

$$\widehat{\mathbb{E}}_{t^-}[X_{t+1}] = \int X_{t+1}(\phi_{\pi,t+1}, v_{t+1}, \epsilon_{t+1}; \mathcal{B}_t)\, dF_t(\phi_{\pi,t+1}, v_{t+1}, \epsilon_{t+1}),$$

where F_t denotes the joint perceived distribution over next-period latent states implied by the inherited belief state and the perceived laws of motion for the latent variables.

For each point on the state grids, the code solves the current nonlinear equilibrium conditions for consumption, labor, firm oil use, household oil consumption, the oil price, inflation, output, and the other endogenous variables. Household oil consumption and firm oil use jointly satisfy the oil-market clearing condition, while the produced-good resource constraint accounts for household consumption of the produced good and Rotemberg price-adjustment costs.

Because current decisions depend on expected future consumption, the consumption

price index and inflation, the solution is obtained by an expectations fixed-point iteration. Starting from the previous policy functions, the algorithm alternates between:

1. evaluating expected future objects under the current belief state;
2. solving the nonlinear current-period equilibrium at each grid point; and
3. updating the policy functions.

The iteration continues until the Euler and pricing objects converge or the maximum number of iterations is reached. The resulting policy functions map candidate latent states into equilibrium outcomes.

Natural output is computed from the corresponding flexible-price economy. In the actual data-generating equilibrium, the central bank observes true natural output and therefore uses the true contemporaneous output gap when setting the policy rate.

C.4 Generating the current observation

For the realized oil shock, policy coefficient, and policy shock, the solved policy functions generate the actual period-t equilibrium. Private agents subsequently observe

$$(Y_t, \Pi_{q,t}, R_t),$$

while they do not observe natural output and the output gap.

The actual equilibrium is generated at the true realization of the oil shock, whereas private sector inference is based on the perceived oil-shock distribution inherited from the previous period.

C.5 Private sector filtering

After observing $(Y_t, \Pi_{q,t}, R_t)$, private agents evaluate possible explanations of the observation. Candidate oil shocks are restricted to the numerical support of the previous-period density $\widehat{g}_{t-1}$.

For each candidate oil shock v, the flexible-price model provides a candidate level of natural output. Combining this with observed output gives the private sector's candidate-dependent inferred output gap,

$$x_t^{\inf}(v) = \log Y_t - \log Y_t^n(v).$$

Thus, the output gap used in inference is not observed directly and varies across candidate oil shocks.

For every candidate pair (ϕ_π, v), the observed policy rate is used to calculate the implied perceived monetary policy shock:

$$\widehat{\epsilon}_t(\phi_\pi, v) = \log\left(\frac{R_t}{R^*}\right) - \phi_\pi \log\left(\frac{\Pi_{q,t}}{\Pi_q^*}\right) - \phi_y x_t^{\inf}(v);$$

in other words, the implied policy shock is chosen so that the perceived Taylor rule reproduces the policy rate set by the central bank exactly. Conditional on each candidate pair (ϕ_π, v), the observed policy rate therefore pins down the implied policy innovation. We consequently concentrate out ϵ_t and represent the posterior numerically over (ϕ_π, v).

The candidate state and its implied policy shock are then passed through the nonlinear policy functions. This produces candidate inflation and output, denoted by $\Pi_t^m(\phi_\pi, v)$ and $Y_t^m(\phi_\pi, v)$.

The nonlinear mapping from the latent state $(v_t, \epsilon_t, \phi_{\pi,t})$ to the observed outcomes $(Y_t, \Pi_{q,t}, R_t)$ does not always admit a unique and well-conditioned exact inversion. Different combinations of latent states can generate very similar observable outcomes. For example, a more adverse supply disturbance and weaker monetary policy responsiveness can both contribute to higher inflation, while the policy innovation provides an additional source of variation in the observed interest rate. We therefore treat inference as a signal-extraction problem rather than attempting to recover the latent state through exact inversion.

Instead, we implement the inference problem using a grid-based nonlinear Bayesian filter. For each candidate (ϕ_π, v), the filter assigns greater weight to states whose model-implied inflation and output are closer to the observed outcomes. Specifically, the matching kernel is

$$\mathcal{K}_t(\phi_\pi, v) = \exp\left\{-\frac{1}{2}\left[\frac{\log \Pi_{q,t} - \log \Pi_{q,t}^m(\phi_\pi, v)}{\sigma_\pi^K}\right]^2 - \frac{1}{2}\left[\frac{\log Y_t - \log Y_t^m(\phi_\pi, v)}{\sigma_Y^K}\right]^2\right\}. \tag{A.2}$$

The bandwidths σ_π^K and σ_Y^K determine how sensitive the matching weight is to differences between the model-implied and observed outcomes. Candidate latent states whose implied inflation and output are close to the observed values receive greater weight, with the weight declining as the discrepancy increases.

Combining the matching kernel with the prior distributions over the monetary policy coefficient and the oil shock, and with the density of the implied policy shock, gives the unnormalized joint posterior

$$\mu_{t|t}(\phi_\pi, v) \propto \mu_{t|t-1}^{\phi_\pi}(\phi_\pi)\, \widehat{g}_{t-1}(v)\, f_\epsilon(\widehat{\epsilon}_t(\phi_\pi, v))\, K_t(\phi_\pi, v).$$

Here, $\mu^{\phi_\pi}_{t|t-1}(\phi_\pi)$ is the predictive distribution of the monetary policy coefficient, $\widehat{g}_{t-1}(v)$ is the perceived oil-shock density inherited from the previous period, and f_ϵ is the density of the monetary policy innovation. Thus, candidate states receive greater posterior weight when they are more likely under the inherited beliefs, require a more plausible monetary policy innovation to rationalize the observed policy rate, and generate inflation and output closer to their observed values.The normalized joint posterior is

$$\mu_{t|t}(\phi_\pi, v) = \frac{\widetilde{\mu}_{t|t}(\phi_\pi, v)}{\sum_{\phi_\pi}\sum_{v}\widetilde{\mu}_{t|t}(\phi_\pi, v)}.$$

The reported perceptions $\widehat{\phi}_{\pi,t}$ and $\widehat{v}_t$ are posterior means. The full marginal posterior of the policy coefficient is retained for the transition to the next period.

C.6 Updating beliefs

After filtering, the posterior distribution of the policy coefficient is propagated through its AR(1) transition law. In the discretized ϕ_π grid, the predictive distribution is

$$\mu^{\phi_\pi}_{t+1|t}(\phi^j_\pi) = \sum_i p(\phi^j_\pi \mid \phi^i_\pi)\, \mu^{\phi_\pi}_{t|t}(\phi^i_\pi),$$

where $p(\phi^j_\pi \mid \phi^i_\pi)$ denotes the transition probability from grid point ϕ^i_π to ϕ^j_π. This produces the predictive distribution of the policy coefficient used for period $t+1$ decisions.

Beliefs about the oil-shock distribution are updated nonparametrically. In the baseline specification, agents summarize their posterior about the current oil shock by its posterior mean,

$$\widehat{v}_t = \sum_j v^j\, \mu^v_{t|t}(v^j),$$

where $\mu^v_{t|t}(v^j)$ denotes the marginal posterior probability assigned to oil-shock grid point v^j. The posterior mean $\widehat{v}_t$ is appended to the history of inferred oil shocks, and the perceived shock distribution is re-estimated using a kernel density estimator (KDE). Let $\mathcal{H}_{-1}$ denote the pre-event history of oil-shock realizations, T_{pre} the number of observations in this history, and $K_h(\cdot)$ a kernel with bandwidth h. The baseline update can be written as

$$\widehat{g}_t(v) = \frac{1}{T_{\text{pre}} + t + 1}\left[\sum_{\tau\in\mathcal{H}_{-1}} K_h(v - v_\tau) + \sum_{\tau=0}^{t} K_h(v - \widehat{v}_\tau)\right].$$

We implement the KDE using a Gaussian kernel and the data-dependent normal-reference bandwidth, i.e. the bandwidth is re-estimated as the history of inferred shocks

expands, following Kozlowski et al. (2020). A key difference is that shocks are directly observed in their setting, whereas in our model the distribution is estimated from inferred shocks $\widehat{v}_t$.

This updating rule allows an unusually large inferred realization to alter the estimated tail of the perceived oil-shock distribution. Because the realization entering the KDE is itself inferred rather than observed, the extent of this revision depends on agents' joint inference about the oil shock and monetary policy responsiveness.

As a robustness check, we consider an alternative updating rule that retains the full posterior uncertainty about each inferred oil shock. Rather than assigning the period-τ observation entirely to its posterior mean $\widehat{v}_\tau$, its contribution to the KDE is averaged over the marginal posterior distribution of v_τ. On the discretized oil-shock grid, the perceived density is

$$\widehat{g}_t^{\text{robust}}(v) = \frac{1}{T_{\text{pre}} + t + 1} \left[\sum_{\tau \in \mathcal{H}_{-1}} K_h(v - v_\tau) + \sum_{\tau=0}^{t} \sum_j \mu^v_{\tau|\tau}(v^j) K_h(v - v^j) \right].$$

Thus, the baseline treats the posterior mean as the newly inferred shock realization, whereas the alternative allows the entire marginal posterior of the inferred shock to contribute to the estimation of $\widehat{g}_t$. All other aspects of the information structure and learning algorithm are unchanged. Figure A.3 compares the two specifications and shows that the resulting dynamics are very similar.

C.7 Transition to the next period

The updated KDE and the predictive distribution of the policy coefficient form the belief state for the next period. The nonlinear equilibrium is then re-solved under these new beliefs. Consequently, current inference affects future allocations through the expectations entering household saving and firms' price-setting decisions.

The algorithm can be summarized as follows:

1. Start period t with the inherited oil-shock KDE and predictive distribution of the policy coefficient.
2. Solve the nonlinear equilibrium under these beliefs.
3. Generate the actual equilibrium at the realized latent state; the fully informed central bank sets R_t using the true output gap.
4. Private agents observe $(Y_t, \Pi_{q,t}, R_t)$.

5. Evaluate candidate oil shocks under the inherited perceived oil-shock distribution. For each candidate oil shock, compute candidate natural output and the corresponding latent output gap.

6. For each (ϕ_π, v), use the observed policy rate to obtain the implied policy shock $\widehat{\epsilon}_t(\phi_\pi, v)$.

7. Evaluate the nonlinear policy functions at each candidate state and evaluate the Gaussian matching likelihood for candidate inflation and output to measure how closely candidate inflation and output reproduce the observations.

8. Combine the matching kernel with the prior weights and the density of the implied policy shock to obtain the joint posterior over $(\phi_{\pi,t}, v_t)$.

9. Propagate the policy-coefficient posterior and re-estimate the oil-shock KDE after appending $\widehat{v}_t$.

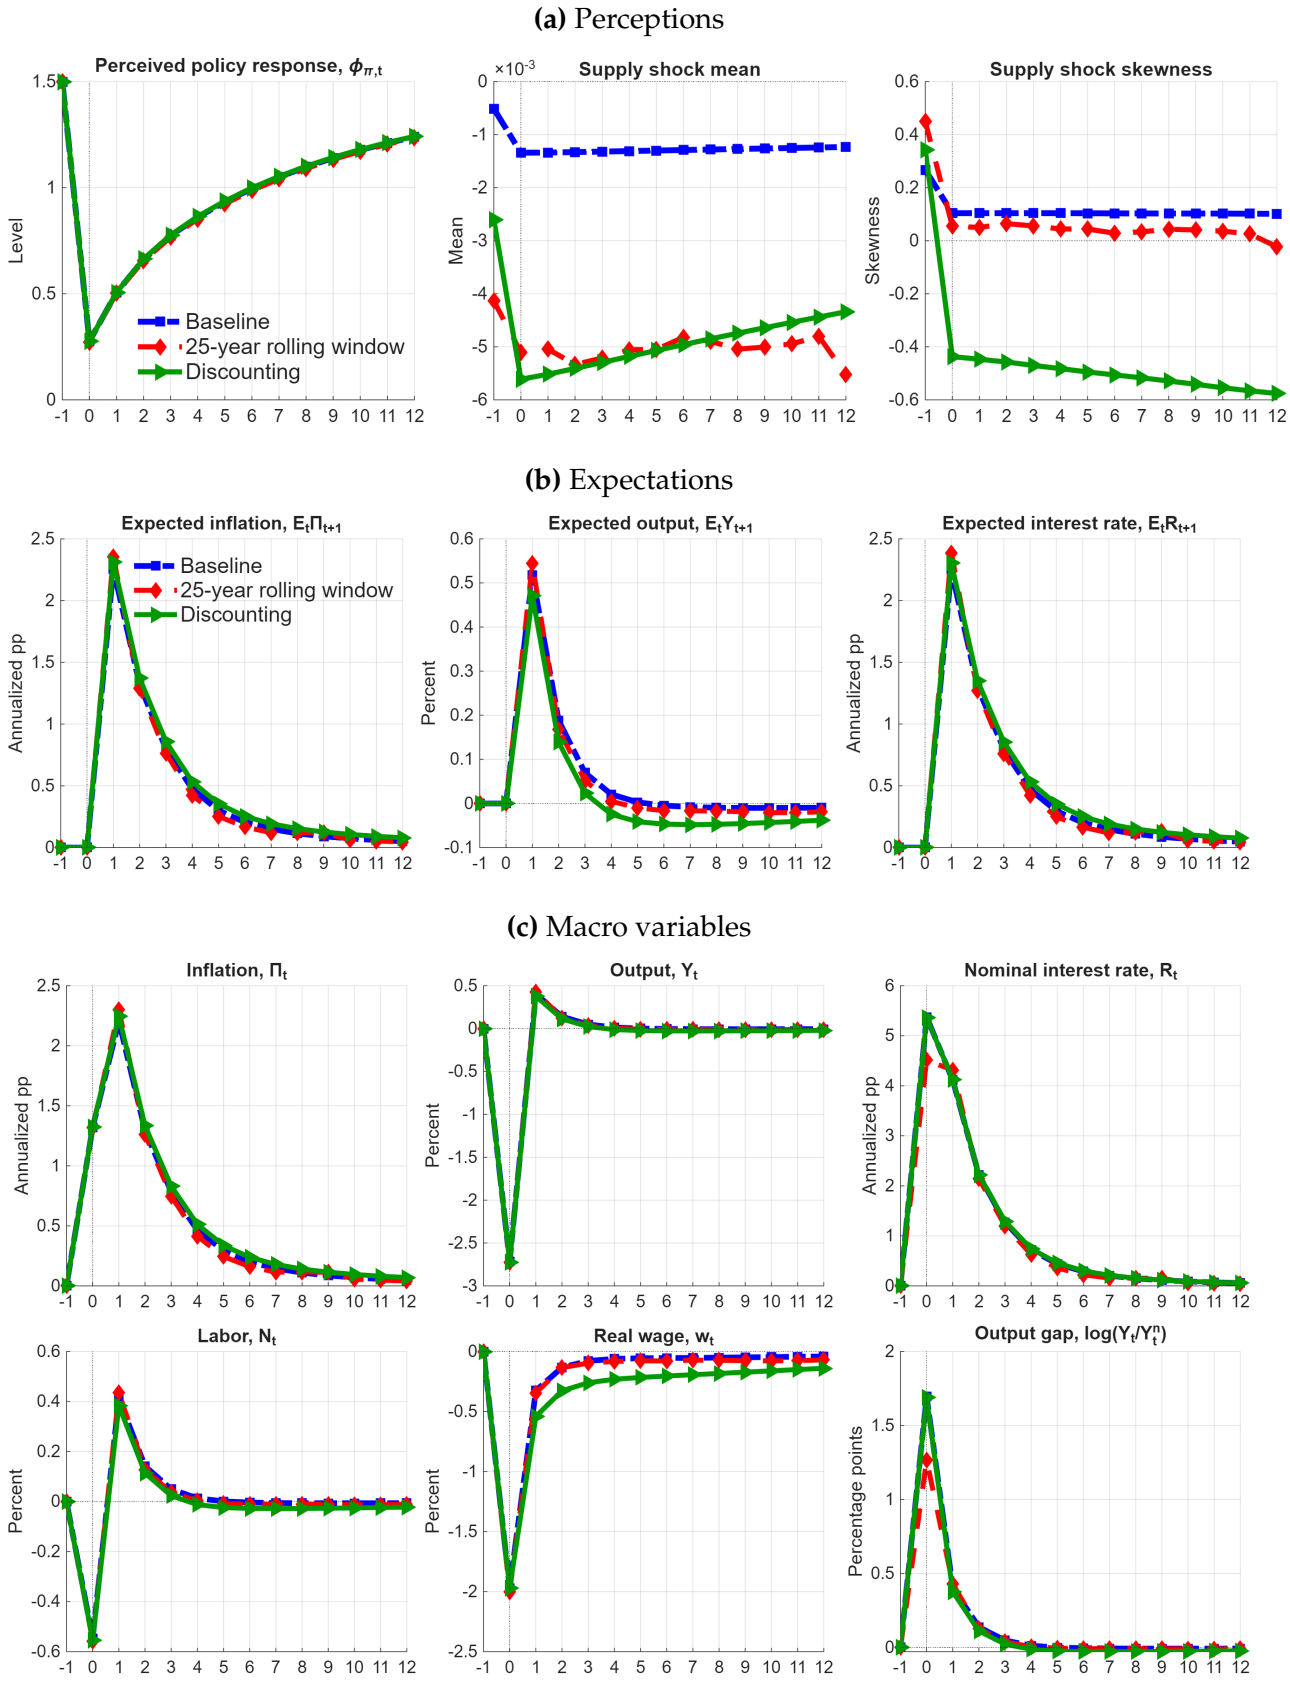
10. Carry the resulting belief state into period $t+1$ and repeat.

C.8 Anticipated-utility approach

When forming period-t expectations of future equilibrium outcomes, agents condition on the inherited belief state and do not integrate over the endogenous evolution of beliefs that would result from future observations. Future outcomes are therefore evaluated under the current perceived distributions. After the period-t realization is observed, beliefs are updated according to the filtering procedure described above, and the equilibrium is re-solved before period-$t+1$ decisions are made.

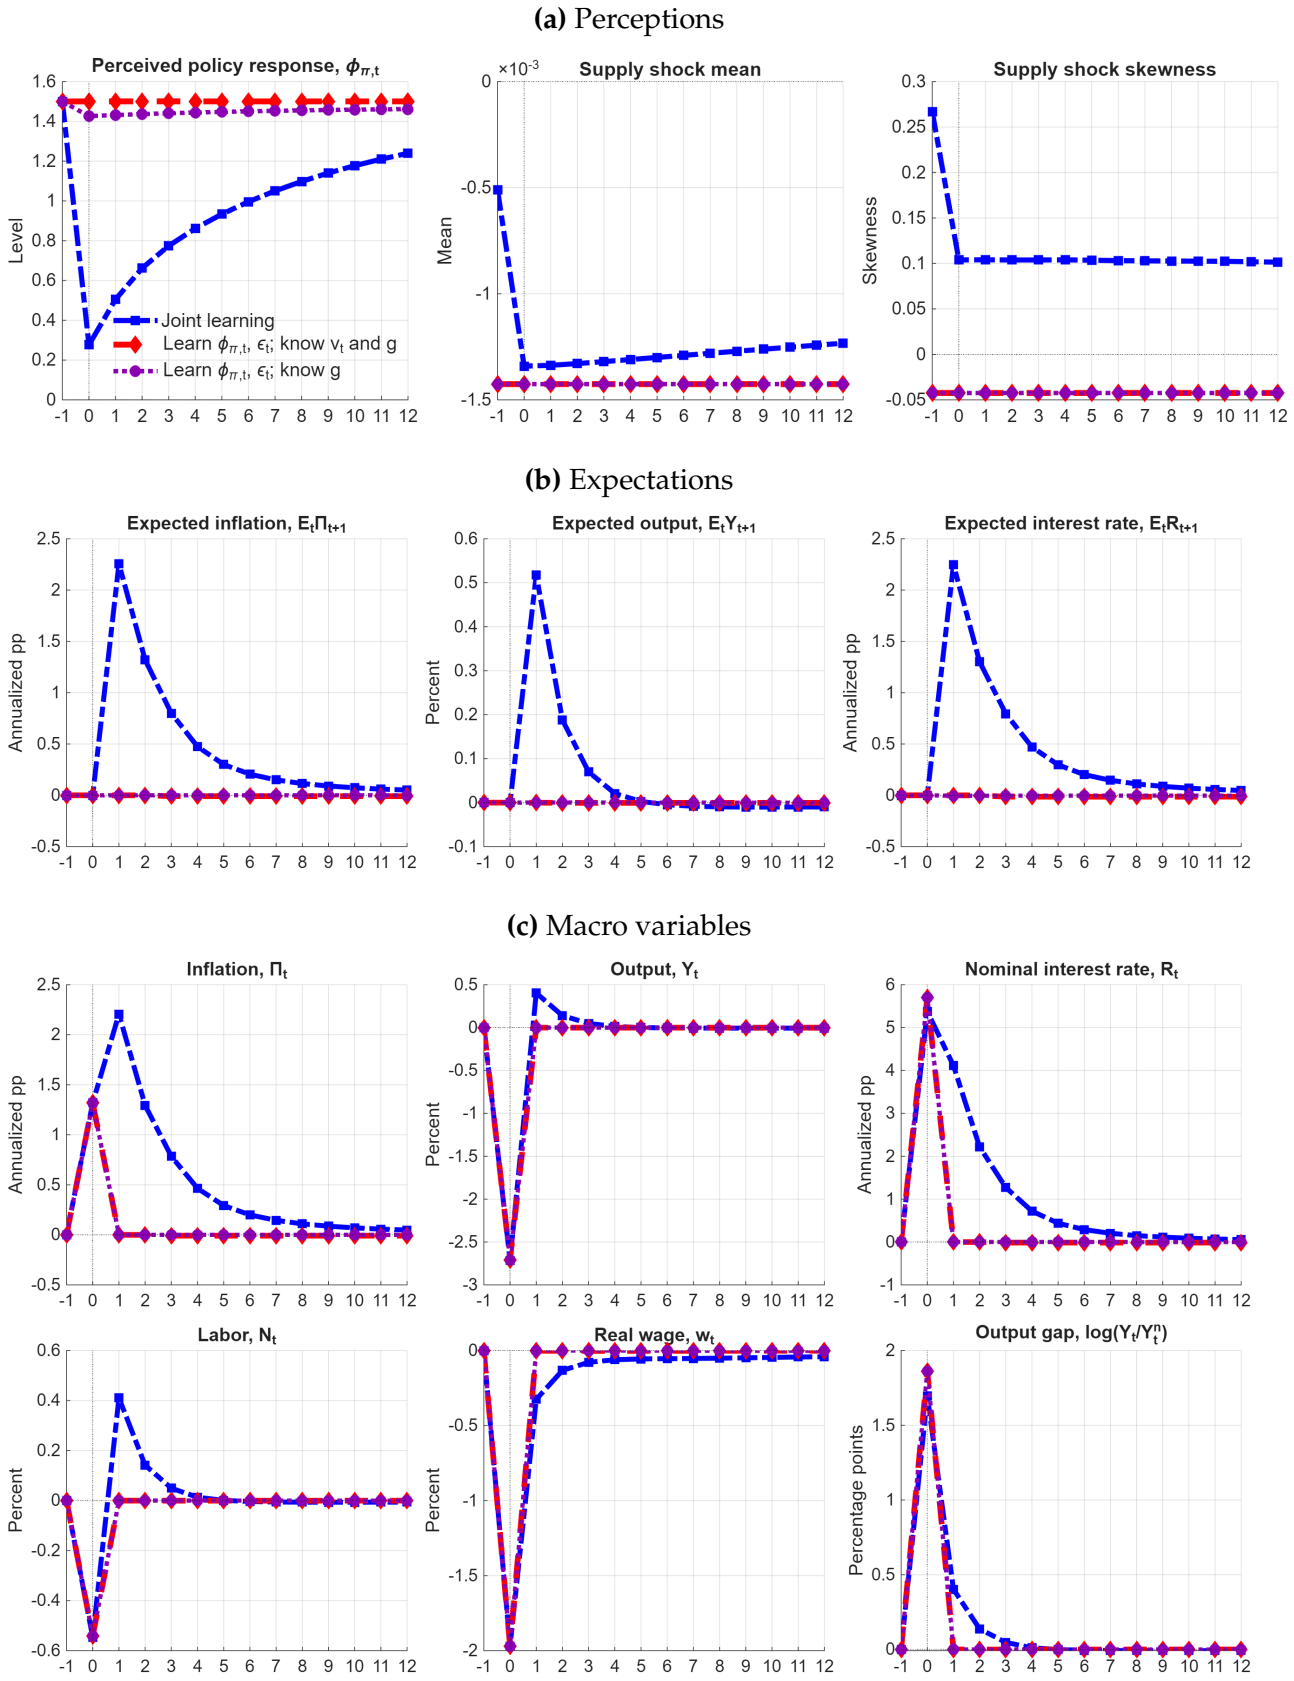
This anticipated-utility treatment separates the equilibrium solution conditional on a belief state from the statistical updating of that belief state and keeps the recursive learning problem computationally tractable.

Figure A.1: Impulse response functions following a large supply shock: Updating the shock distribution with a rolling window or discounting



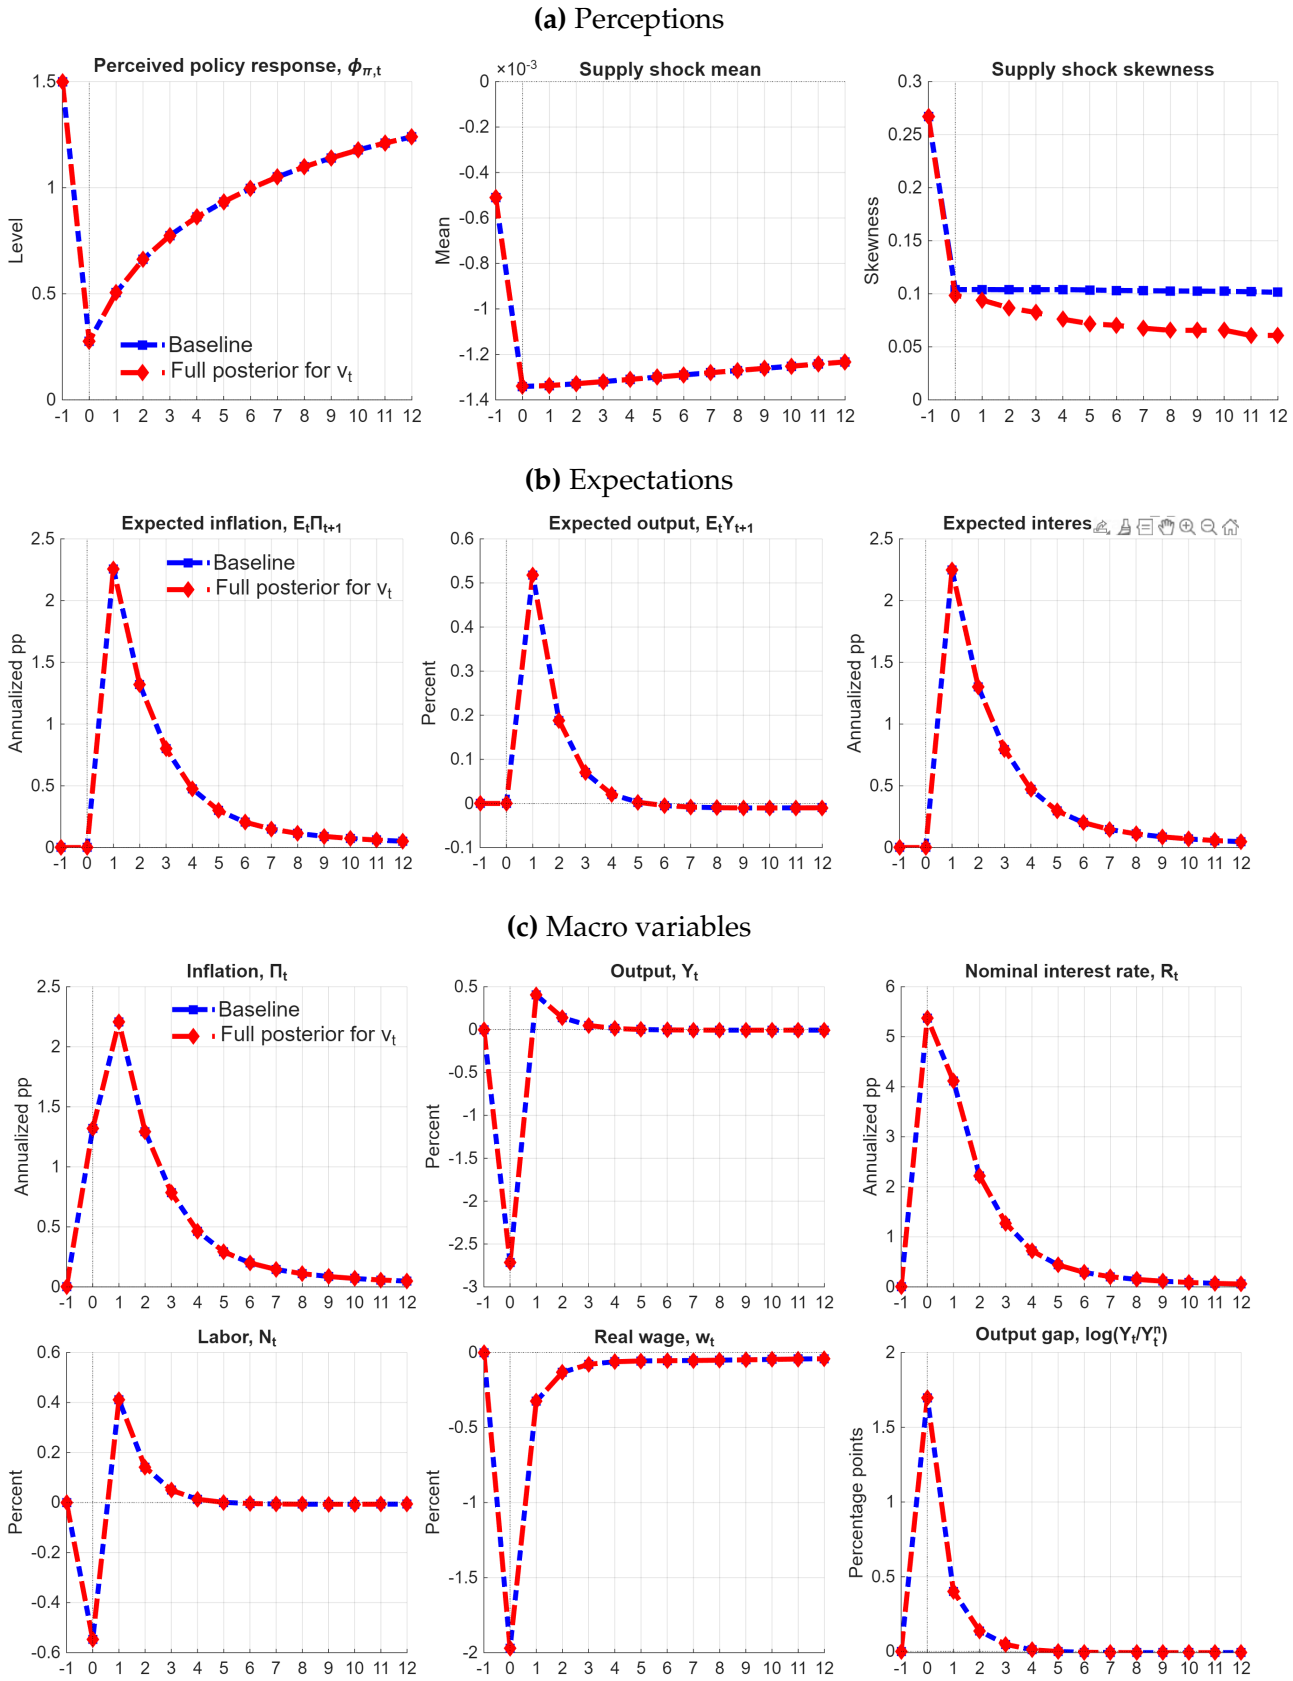
Notes: The blue line denotes impulse responses to a large negative supply shock under the baseline learning algorithm. The red line shows the case where learning occurs using a rolling window of 25 years' observations. The green line corresponds to gradual forgetting in the learning process, i.e. assigning greater importance to more recent observations through exponential discounting.

Figure A.2: Impulse response functions following a large supply shock: Alternative information structures



Notes: The figure compares the baseline joint-learning model (blue) with two counterfactual information structures. In the first counterfactual (red), agents know the supply-shock distribution g and observe the realization of v_t, but do not observe $\phi_{\pi,t}$ or ϵ_t (as in Section 4.2.3). In the second counterfactual (purple), agents know g but observe neither v_t, $\phi_{\pi,t}$, nor ϵ_t.

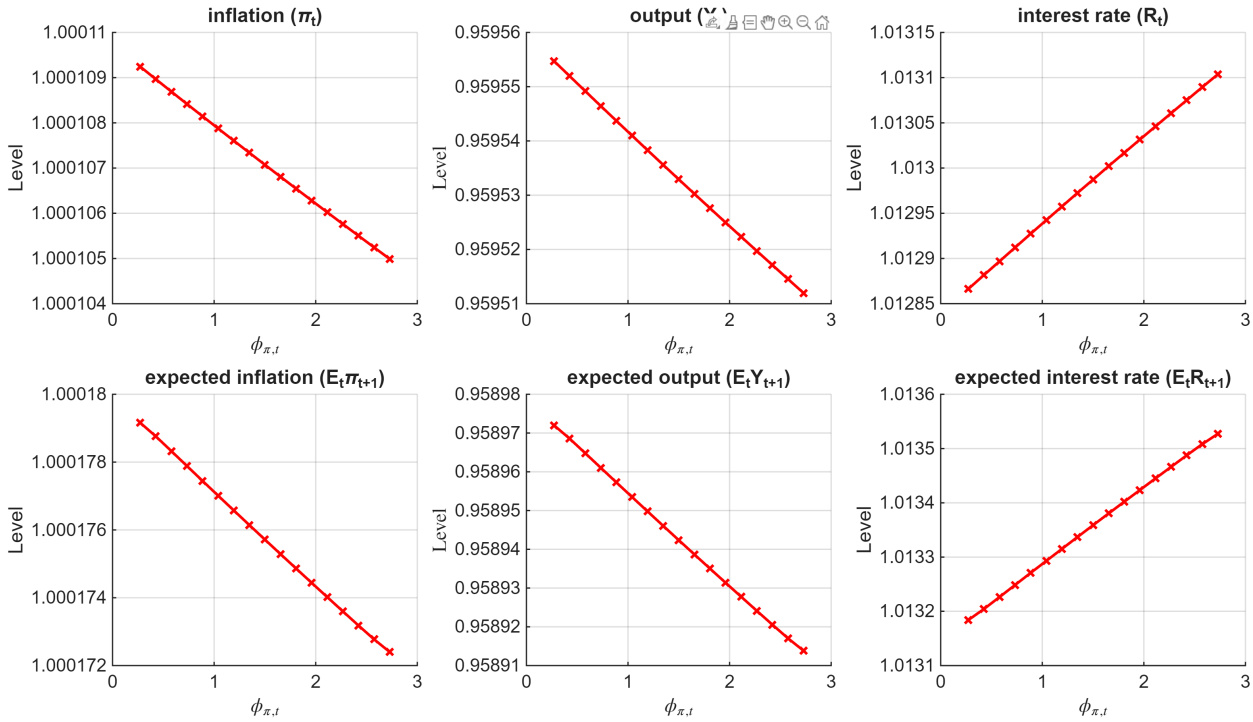
Figure A.3: Impulse response functions following a large supply shock: Alternative updating rule for the perceived supply shock distribution



Note: The blue line shows the baseline specification, in which the perceived supply-shock distribution is updated using the posterior mean $\widehat{v}_t$. The red line shows the alternative specification, in which it is updated using the full marginal posterior distribution of v_t (see Appendix C.6 for details). All other aspects of the model and learning algorithm are identical.

Figure A.4: Equilibrium outcomes conditional on alternative values of $\phi_{\pi,t}$

(a) Before the shock ($t = -1$)



(b) After the shock ($t = 1$)

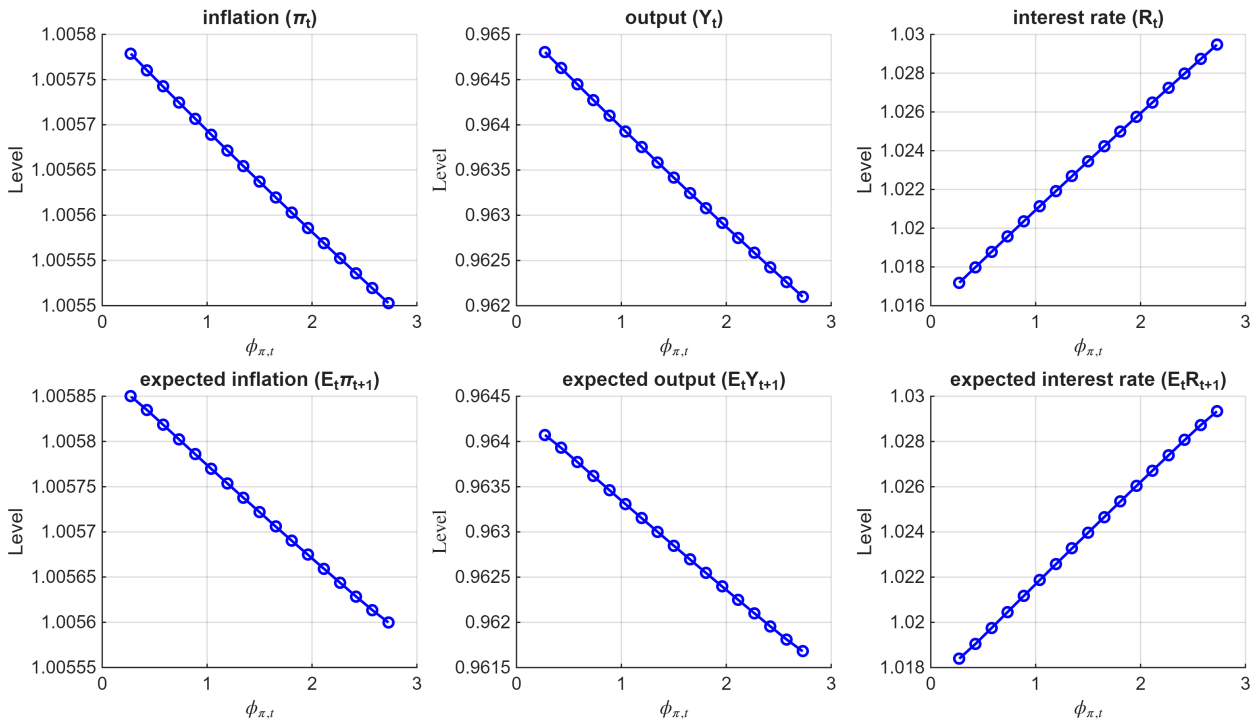


Figure A.5: Impulse response functions from nonlinear VAR: Alternative measure for inflation expectations

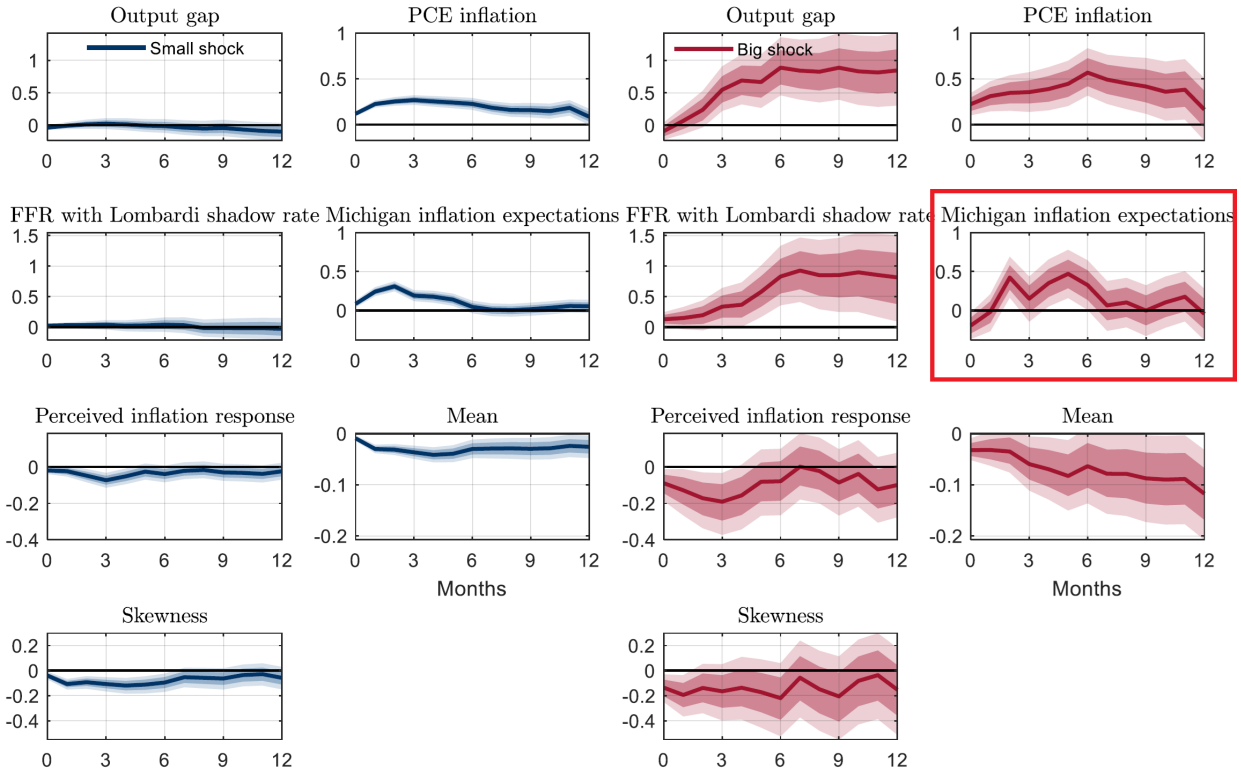


Figure A.6: Impulse response functions from nonlinear VAR: Alternative measure of interest rate

