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

Christiane Baumeister† Fiorella De Fiore‡ Jasmine Xiao§

March 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 New Keynesian model, we show that it endogenously generates inflation dynamics persistently above target and temporary output gap overshooting after a tail event, closely matching the data. Our results suggest that public misperceptions can potentially amplify volatility and undermine 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, Christian Matthes and Carolin Pflueger for very helpful comments and suggestions.

†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 interpretive 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, conventional models provide little guidance on the appropriate scale or timing of policy responses. Actions that would normally anchor expectations may instead be perceived 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 2008 financial crisis and 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 New Keynesian model with non-normally distributed 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 rare events.

When a large negative supply shock occurs, our model generates amplified and persistent dynamics that closely follows the data. Inflation and the policy rate rise sharply on impact, while the output gap contracts. Although the shock dissipates quickly, inflation, expected inflation, and the interest rate remain elevated for an extended period, and the output gap overshoots before slowly returning to steady state. 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 outside their 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 the expected inflation and output gap. The misperception of the policy rule amplifies the impact of the shock on inflation while simultaneously generating a temporary overshooting of the output gap.

Following a large negative supply shock, agents update their beliefs about the distribution of shocks, increasing the perceived likelihood of left-tail events. They learn much more 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.

We test the model implications empirically using U.S. macroeconomic data. 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 New Keynesian model. Furthermore, the nonlinear model 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 rises persistently, while the output gap overshoots following an initial contraction. At the same time,

perceived policy responsiveness declines, indicating that agents interpret part of the inflation surge as reflecting a weaker policy stance. 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.

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.[1]

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 nonlinear 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,

[1]Focusing on the actual rule rather than the perceived rule, Nakamura, Riblier, and Steinsson (2025) question whether the simple Taylor rule accurately describes actual monetary policy in the post-Covid era. We show that – even holding the true rule fixed – agents' misperceptions of that rule following large supply shocks can themselves generate substantial macroeconomic amplification.

the inference problem is nonlinear and non-Gaussian, and learning takes the form of particle filtering. 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 emphasize anchoring 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) 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 shocks in the economy, but estimate the distribution using real-time data. While their model requires several nonlinear features for learning to generate quantitatively meaningful effects, we show that simultaneous learning about the shock distribution and the policy rule can produce strong interaction effects.

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 (2023), which also emphasizes the role of oil shocks, though we differ significantly in model setup and focus.[2] 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 different information structures. Section 3 describes the algorithm and presents all model simulations. Section 4 provides empirical evidence for the U.S. Section 5 concludes.

[2]Baumeister, Huber, and Marcellino (2024) also emphasizes 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.

2 Model

We start with the three-equation New Keynesian model:

$$\pi_t = \delta \mathrm{E}_t \pi_{t+1} + \kappa x_t - v_t, \tag{1}$$

$$x_t = \mathrm{E}_t x_{t+1} - \sigma^{-1}(i_t - \mathrm{E}_t \pi_{t+1}), \tag{2}$$

$$i_t = \beta_t \pi_t + \varepsilon_t. \tag{3}$$

where (1) is the Phillips Curve, (2) is the Euler equation and (3) is the monetary policy rule. π_t denotes inflation in deviation from the inflation target, x_t the output gap, and i_t the interest rate. v_t is an independently and identically distributed (i.i.d.) supply shock. Let g denotes its density, which is potentially non-normal. The monetary policy shock ε_t is i.i.d. normal, $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$. $\delta \in (0,1)$ is the rate of time preference. σ denotes the private sector's risk aversion in the Euler equation. In the monetary policy rule (3), β_t is the time-varying response to inflation:[3]

$$\beta_t = (1-\rho)\overline{\beta} + \rho \beta_{t-1} + \xi_t, \tag{4}$$

where $\xi_t \sim iid\ N(0, \sigma_\xi^2)$ and $\rho \in (0,1)$. v_t, ε_t and ξ_t are mutually uncorrelated.

We consider three variants of the model where the differences lie in the agents' information set. Agents are fully rational in all setups.

1. Full information In the full information benchmark, all agents perfectly observe all shocks $(\xi_t, \varepsilon_t, v_t)$ and the time-varying policy rule β_t in each period. In addition, agents know the true density of the supply shock, g, and observing such a tail event does not change their beliefs of the density.

2. Learning about the policy rule (β_t) Private sector agents have a different information set from the central bank in each period: only the central bank knows the true Taylor rule and the realized policy shocks (ξ_t, ε_t). The private sector does not know the true β_t or policy shock ε_t.[4] Here we continue to assume that all agents know the true density of the supply shock, g. The realized macroeconomic outcomes (π_t, x_t, i_t) continue to be perfectly observed by all agents.

[3] In Section 3.7, we consider an alternative Taylor rule where the policy rate responds to output gap.

[4] For the purpose of dissecting the mechanism, we also consider an intermediate case: the symmetric but imperfect information model, where neither the central bank nor the private sector observes the realized shocks (ξ_t, ε_t) – see Figure A.9 in the appendix.

At the beginning of period t, the private sector agents are endowed with beliefs about the policy rule inherited from the previous period. Let $\widehat{\beta}_{t-1}$ denote the mean of this belief. They are also endowed with the (mean) belief for the policy shock, $\widehat{\varepsilon}_t$. Following much of the learning literature, we use the anticipated utility approach of Kreps (1998) and assume that the private sector makes economic plans at t as if they expect the policy parameters to stay fixed at $\widehat{\beta}_{t-1}$. The equilibrium given beliefs is given by:

$$\begin{pmatrix} \pi_t \\ x_t \end{pmatrix} = F(\widehat{\beta}_{t-1},\, v_t,\, \widehat{\varepsilon}_t; g) \tag{5}$$

$$\begin{aligned} i_t &= \widehat{\beta}_{t-1}\pi_t + \widehat{\varepsilon}_t \\ i_t &= \beta_t \pi_t + \varepsilon_t. \end{aligned} \tag{6}$$

The central bank sets the interest rate according to the true β_t and ε_t. The belief of private sector agents, $\widehat{\varepsilon}_t$, must reconcile the interest rate under the perceived rule (6) with the observed interest rate under the actual rule (3), given β_t, $\widehat{\beta}_{t-1}$ and ε_t. This implies that the following condition must hold:

$$\widehat{\varepsilon}_t = (\beta_t - \widehat{\beta}_{t-1})\pi_t + \varepsilon_t, \tag{7}$$

where π_t depends on $\widehat{\varepsilon}_t$ through F in (5).

After the equilibrium outcomes for (π_t, x_t, i_t) are realized, agents update these beliefs by solving the Bayesian filtering problem characterized by (1)-(4), and the result of this filtering problem is a posterior belief with mean $\widehat{\beta}_t$. Agents then use this to solve the equilibrium outcomes in the following period.[5]

3. Learning about the policy rule (β_t) and shock distribution (g) We continue to assume that the private sector has a different information set from the central bank, and that only the central bank knows the true policy rule β_t and observes the realized shocks (ζ_t, ε_t). The private sector continues to learn about β_t by observing the macroeconomic outcomes (π_t, x_t, i_t). In addition, the private sector does not know the true density of the supply shock, g, and must estimate the distribution using the history $\{v_j\}_{j=0}^{t}$. Let $\widehat{g}_t$ denote the estimated distribution in period t. The model is otherwise the same as the second setup.

[5]The assumption that beliefs about the rule parameters are updated after macroeconomic outcomes in the current period are realized is standard in the learning literature (see, for example, Cogley, Matthes, and Sbordone, 2015).

The timing of the problem is as follows. At the beginning of period t, agents are endowed with the beliefs $\{\widehat{\beta}_{t-1}, \widehat{g}_{t-1}\}$.[6] Agents use these to form their expectations, $\mathrm{E}_t[\pi_{t+1}]$ and $\mathrm{E}_t[x_{t+1}]$. Then all shocks $(v_t, \varepsilon_t, \xi_t)$ are realized. The equilibrium outcomes for (π_t, x_t, i_t) are characterized by:

$$\begin{aligned} \begin{pmatrix} \pi_t \\ x_t \end{pmatrix} &= F_t(\widehat{\beta}_{t-1},\ v_t,\ \widehat{\varepsilon}_t; \widehat{g}_{t-1}) \\ i_t &= \beta_t \pi_t + \varepsilon_t. \\ \widehat{\varepsilon}_t &= (\beta_t - \widehat{\beta}_{t-1})\pi_t + \varepsilon_t. \end{aligned} \tag{8}$$

At the end of period t, the private sector update their belief of the policy rule from $\widehat{\beta}_{t-1}$ to $\{\widehat{\beta}_t\}$ based on the macro outcomes (π_t, x_t, i_t), and their belief of the shock distribution from $\widehat{g}_{t-1}$ to $\widehat{g}_t$ given the current supply shock v_t. Notably, now the policy functions $F_t(\cdot)$ are time-varying since the estimated distribution $\widehat{g}_t$ is time-varying.

3 Simulations

3.1 Parameterization

We calibrate the model at the quarterly frequency. To facilitate comparison, we use the same parameter values across all three models.

Table 1: Calibration

Parameter	Description	Value
δ	rate of time preference	0.99
κ	inflation sensitivity to output gap	0.2
σ	risk aversion	20
α	rescaling supply shock series	0.15
$\overline{\beta}$	mean policy response to inflation	1.5
ρ	persistence of policy rule	0.8
σ_ξ	s.d. of policy shock	0.04
σ_ε	s.d. of policy rule shock	0.02

Parameters governing the rate of time preference δ, inflation sensitivity to output gap κ and the mean policy response to inflation $\overline{\beta}$ take common values from the liter-

[6] For robustness, we also consider an alternative assumption that the private sector makes decisions at t with beliefs $(1-\rho)\bar{\beta} + \rho\widehat{\beta}_{t-1}$, which is the rational forecast of β_t at $t-1$. The results are very similar to the baseline case (see Section 3.7).

ature. We have chosen a fairly high value for the risk aversion parameter σ due to the quasi-linear nature of the model. The quantitative impact on expectations is largely determined by the risk aversion parameter.[7] The rescaling parameter α for the supply shock is chosen to target a steady-state inflation rate of 2.5%. While much of the learning literature assumes a random walk process for the latent variable, we choose a fairly conservative value for the persistence of policy rule ρ, as we are interested in how much persistence our mechanism can generate endogenously. The standard deviations of the two policy shocks (σ_ξ and σ_ε) are jointly calibrated to target the volatilities of interest rate and inflation.

3.2 Algorithm

Here we provide an overview of the algorithm we use to solve the imperfect information model and relegate the detailed description to Appendix A. There are three distinctive features of our algorithm in this model setting.

First, the policy functions for inflation, output gap, and interest rate are time dependent. This is because agents' expectations, $\mathrm{E}_t\pi_{t+1}$ and $\mathrm{E}_t x_{t+1}$, depend on the estimated shock distribution $\widehat{g}_t$, which is time-varying. Therefore, before solving the model, we first estimate the supply shock density using a kernel estimator and an expanding window of the data, generating a sequence of estimates $\{\widehat{g}_t\}_{t=1}^{T}$, where T is the last date observed by the agent. Then, for each t, we solve for the equilibrium inflation, output gap, and interest rate, for different values of the state variables, v_t, ε_t and β_t. These are solutions in a quasi-full information model, where agents observe all state variables, and use the estimated shock density in each period to form rational expectations.

The second feature speaks to the asymmetric information assumption that agents do not know the monetary policy rule. As a result, now the macro outcome variables are functions of their beginning-of-period belief $\widehat{\beta}_{t-1}$ rather than the actual β_t. For each period t, we obtain values of π_t and x_t for different levels of $(\widehat{\beta}_{t-1}, v_t, \widehat{\varepsilon}_t)$. Then, using π_t, we find the equilibrium belief for the policy shock $\widehat{\varepsilon}_t^*$ that reconciles the interest rate under the actual rule (3) and the perceived rule (7). Thus, under asymmetric information, the equilibrium belief $\widehat{\varepsilon}_t^*$ is a function of β_t, $\widehat{\beta}_{t-1}$, v_t and ε_t. Finally, we evaluate π_t and x_t at $\widehat{\varepsilon}_t^*$, given $\widehat{\beta}_{t-1}$ and v_t.

Third, since the supply shock v_t has an arbitrary (potentially non-normal) distribution, we estimate agents' beliefs $\widehat{\beta}_t$ using a particle filter. This introduces a two-layer structure to the algorithm: an outer layer, where the particle filter updates beliefs,

[7] If β_t in the policy rule (3) were a constant, the model becomes the standard linearized New Keynesian model.

and an inner layer, where we solve for equilibrium outcomes. In each period t, the inner layer computes macroeconomic outcomes (π_t, x_t, i_t) and the equilibrium belief $\widehat{\varepsilon}_t^*$ given $(\beta_t, v_t, \varepsilon_t, \widehat{\beta}_{t-1})$. Then for each "particle" $\{\widehat{\beta}_t^j\}_{j=1}^N$, we compute the joint conditional density $f(\pi_t, x_t, i_t|\widehat{\beta}_t^j)$ for each $j = 1, 2, ..., N$, where N is the total number of particles. These conditional densities are subsequently used in the outer layer to update the mean belief $\widehat{\beta}_t$.

3.3 Beliefs about the density of supply shocks

We map the cost push shock v_t to an oil supply shock. The data on oil supply shocks is from Baumeister and Hamilton (2019) from 1975Q2 to 2024Q4 (Panel (a), Figure 1). In the model where agents do not know the distribution of v_t, they need to estimate it as new data arrive. When selecting an estimation procedure, we account for the presence of tail events, since assuming normally distributed shocks is inadequate for capturing changes in tail risk, as discussed in Kozlowski, Veldkamp, and Venkateswaran (2020). Thus, we take a nonparametric approach – using a kernel density estimation procedure – to let the data determine the shape of the distribution. Specifically, we employ the widely used normal kernel density estimator and the optimal bandwidth for normal densities (see Appendix A for details).

We use an expanding window of data to estimate the density of the supply shock, with the first and last estimated densities being $\widehat{g}_{1990Q1}$ and $\widehat{g}_{2024Q4}$, respectively. In Panel (b) of Figure 1, we plot the estimated density using data up to 2020Q1 ($\widehat{g}_{2020Q1}$), 2020Q2 ($\widehat{g}_{2020Q2}$), and 2024Q4 ($\widehat{g}_{2024Q4}$), respectively. Two observations are particularly noteworthy. First, due to the large negative supply shock in 2020Q2, the updated distribution $\widehat{g}_{2020Q2}$ has a fatter left tail, compared to $\widehat{g}_{2020Q1}$.[8] Second, although there were no more large shocks since 2020Q2, the one occurrence appears to have a permanent impact on the estimated distribution, as $\widehat{g}_{2024Q4}$ closely resembles $\widehat{g}_{2020Q2}$, especially in the tail. As discussed in Kozlowski et al. (2020), this stems from the martingale property of the estimated distribution.

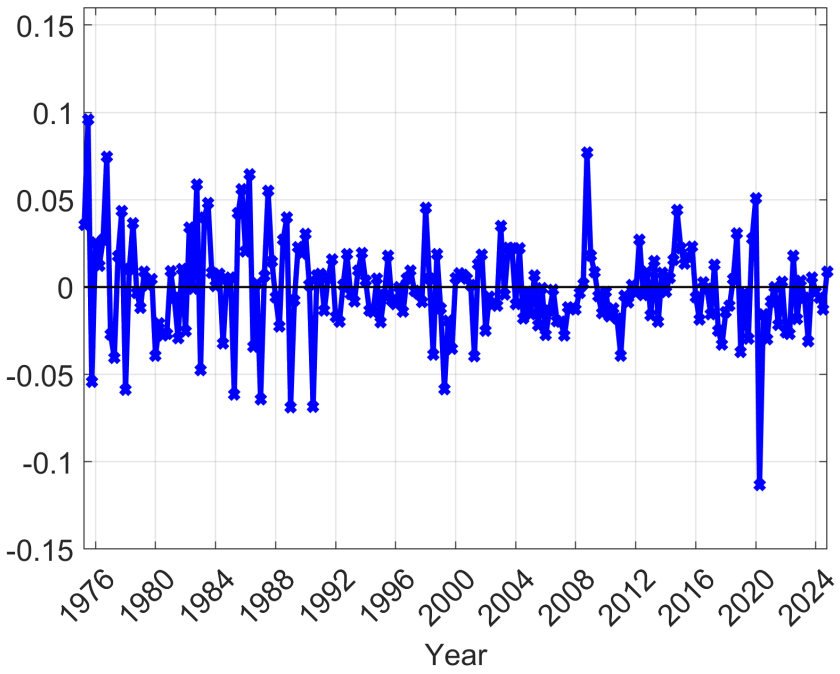
3.4 Impulse Response Functions

Figure 2 shows the impulse response functions following a one-time negative supply shock in period 7. The magnitude of the shock is about four standard deviations,

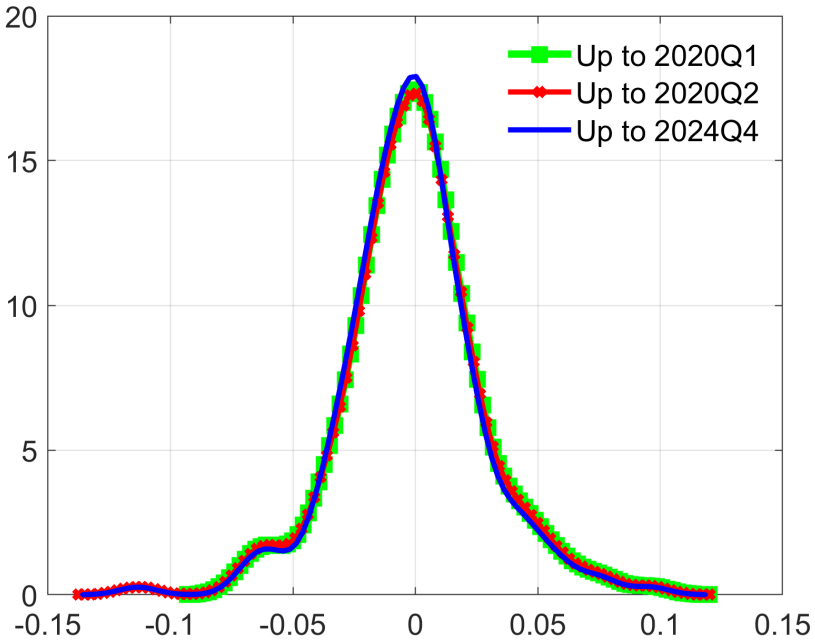
[8]The large negative supply shock in 2020Q2 reflects the joint agreement by OPEC+ at the end of April 2020 to jointl reduce oil output by around 10 million barrels per day, equivalent to around 10% of global fuel consumption, to lift oil prices off the floor (Baumeister, 2023).

Figure 1: Estimated density of oil supply shocks

(a) Oil supply shocks



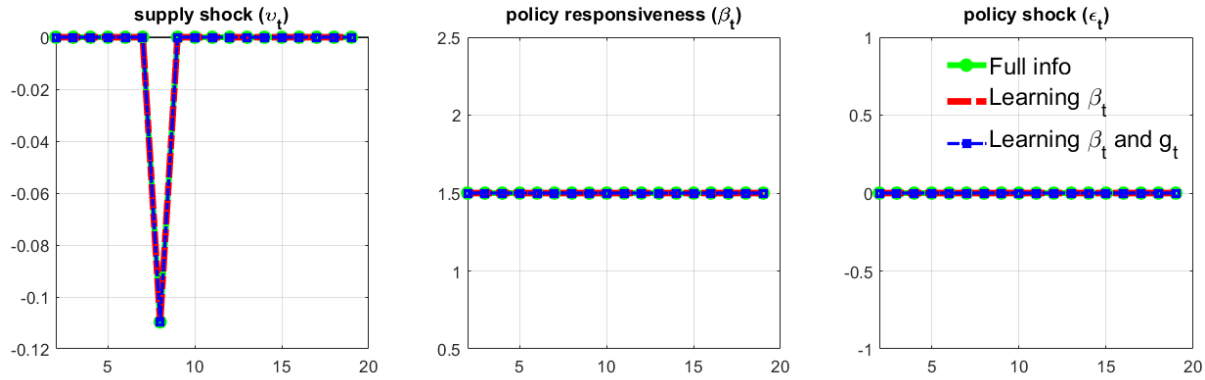
(b) Estimated beliefs $\widehat{g}_t(\cdot)$



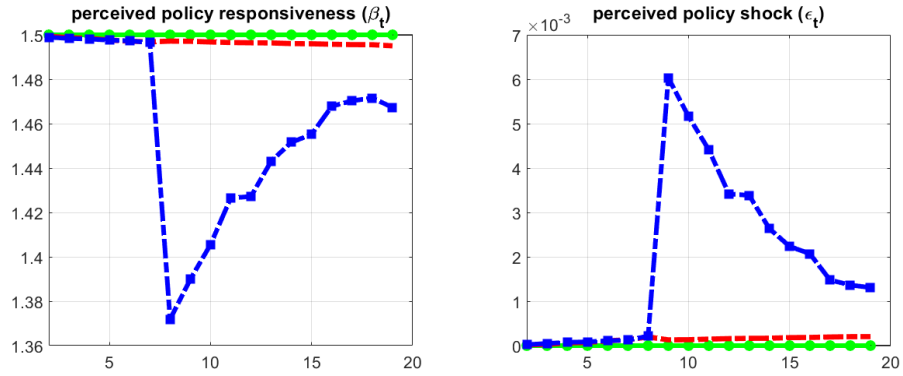
The top panel plots the quarterly oil supply shock series from Baumeister and Hamilton (2019). The bottom panel shows the kernel density, estimated from data available up to 2020Q1 (in green), 2020Q2 (in red) and 2024Q4 (in blue), respectively. The change in the left tail represents the effect of the large negative supply disruption in 2020Q2.

Figure 2: Impulse response functions following a large supply shock

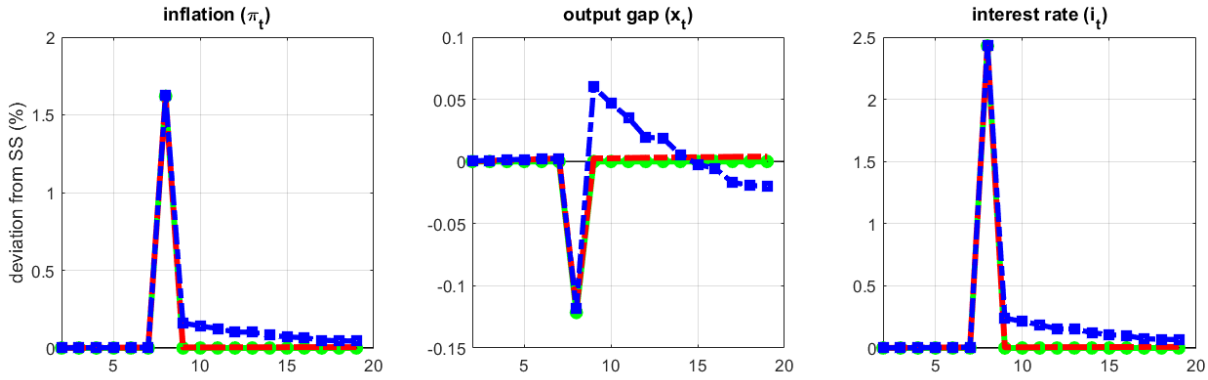
(a) Actual shocks



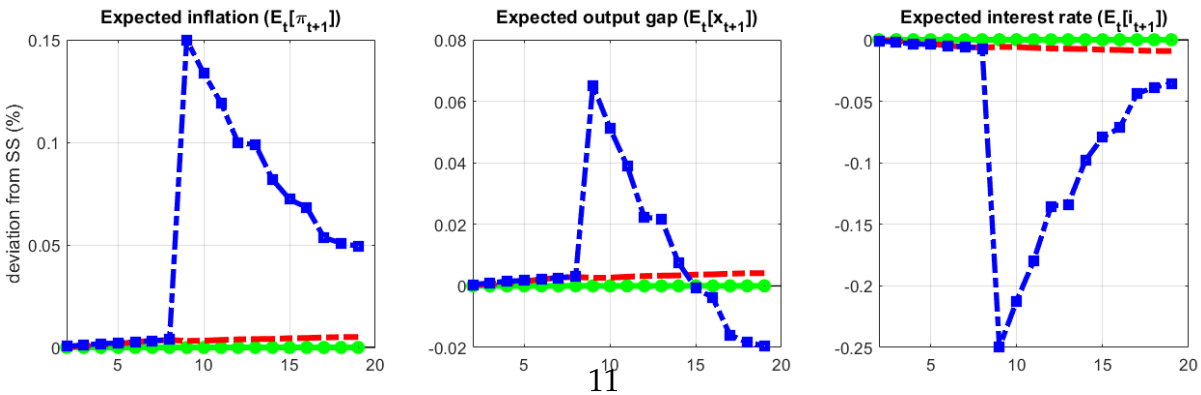
(b) Perceived shocks



(c) Macro variables



(d) Expectations



similar in size to the oil supply shock in 2020Q2 (Baumeister and Hamilton, 2019). Our baseline is the imperfect (asymmetric) information model where the private sector does not know the true β_t or the true shock distribution g. After the large shock, agents permanently revise their estimate of the distribution: the initial estimate corresponds to $\widehat{g}_{2020Q1}$, and the revised estimate is $\widehat{g}_{2020Q2}$ (see Figure 1). We also consider two counterfactuals. In the first counterfactual, we continue to assume that the private sector does not observe β_t, but they know the true shock distribution, corresponding to the full sample distribution, i.e. $g = \widehat{g}_{2024Q4}$. The second counterfactual is the full information model, where agents know both β_t and $g = \widehat{g}_{2024Q4}$.

In the model with imperfect (asymmetric) information about both β_t and g, agents persistently and significantly underestimate the central bank's responsiveness ($\widehat{\beta}_t < \beta_t$). A large negative shock shifts the distribution of v_t (see Figure 1), prompting agents to update their beliefs. However, since the private sector observes only aggregate macroeconomic outcomes, the source of high inflation during the shock period is ambiguous to them: it could reflect either elevated v_t or lower β_t, and agents assign some probabilities to both. When agents update their belief of the policy rule, they still have in mind that the supply shock distribution follows the last period's estimate $\widehat{g}_{t-1}$. Thus, their estimate of supply shock today v_t is higher than its true value and cannot fully account for the sharp rise in inflation. As a result, they mistakenly infer that β_t is lower than it actually is, i.e. $\widehat{\beta}_t < \beta_t$.

While $\widehat{\beta}_t$ falls below the true value (1.5), the perceived policy shock $\widehat{\varepsilon}_t$ rises above its true value (0). In equilibrium, $\widehat{\varepsilon}_t$ must reconcile the interest rate implied by the actual rule with that implied by the perceived rule; hence, its increase follows directly from equation 7. Intuitively, because the private sector observes only the interest rate and not the underlying policy shock, they mistakenly attribute part of the rate increase to a positive policy shock.

In terms of macroeconomic outcomes, inflation remains persistently elevated after the shock, even though the shock itself lasts only one period. This persistence arises from expectations. Since expected inflation is a decreasing function of $\widehat{\beta}_{t-1}$ (see Figure A.1), and $\widehat{\beta}_{t-1}$ adjusts only gradually, expectations decline slowly. As a result, inflation stays above its pre-shock level well after the shock subsides. The interest rate largely mirrors the dynamic of inflation, since the output gap is not in the interest rate rule.[9]

[9]Appendix Figure A.9, we solve the model assuming symmetrically imperfect information where neither the central bank nor the private sector observes β_t or ε_t in the Taylro rule (see A.3 of Appendix A for details). The IRFs are qualitatively and quantitatively similar for most variables, except for the interest rate. Since the perceived policy shock $\widehat{\varepsilon}_t$ is zero, the interest rate undershoots when the supply shock is over but $\widehat{\beta}_{t-1}$ has sharply fallen.

We relax this assumption in Section 3.7.

The output gap falls sharply during the shock period but rises above its steady-state level thereafter. This overshooting is driven by expectations: both expected inflation and the expected output gap increase sharply in the period following the shock, due to the lower perceived policy reaction. Since these expectations are decreasing functions of $\widehat{\beta}_{t-1}$, they adjust only with a lag to changes in perceived policy. The resulting output overshoot, in turn, reinforces the persistence of higher inflation after the shock. [10]

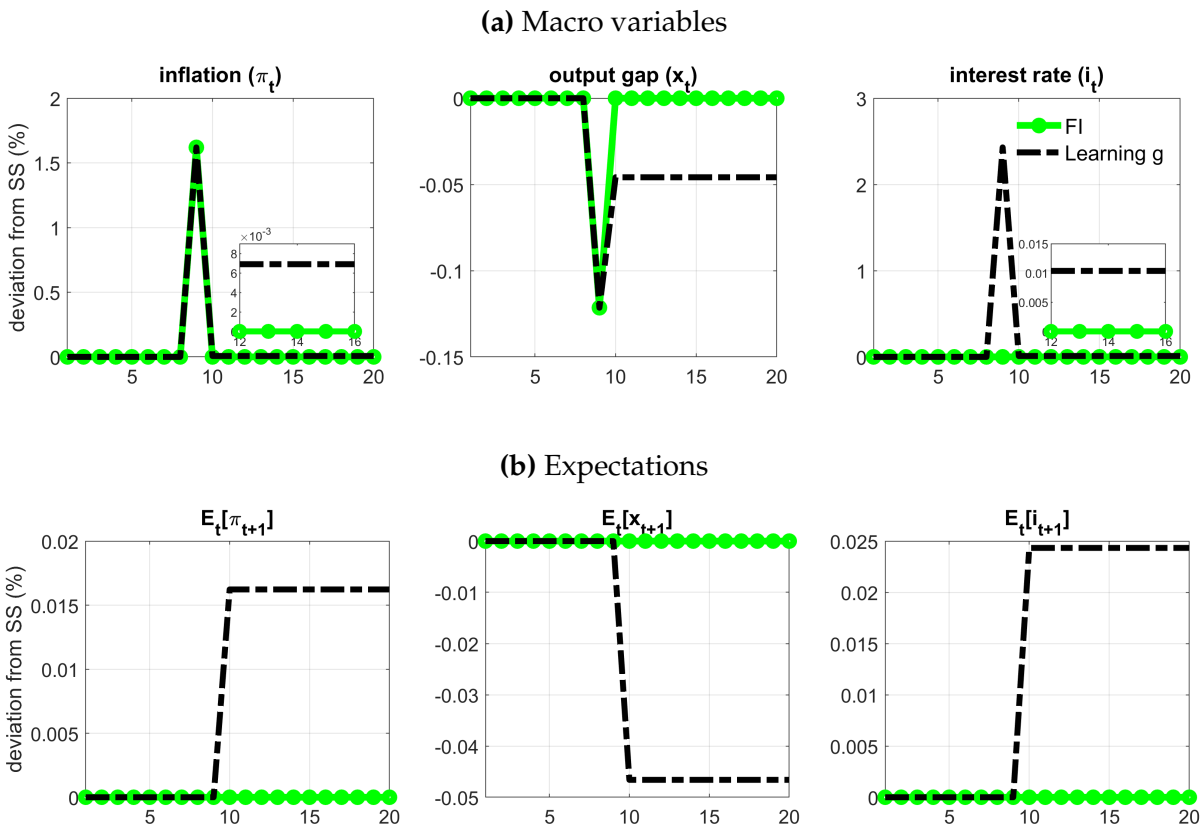
When comparing the two counterfactual models with the baseline, two observations stand out. First, even when the private sector knows the true shock distribution g, they still underestimate policy responsiveness ($\widehat{\beta}_t < \beta_t$), though their estimate is extremely close to the true value. Second, learning about $\widehat{\beta}_t$ alone generates little persistence in the macroeconomic variables. The impulse response functions for inflation, the output gap, and the interest rate when agents know g (and only learn about β_t) almost coincide with those under full information. Although this model displays a modest expectation effect once the shock subsides (see panel (d)), its quantitative impact on real variables is negligible.

Finally, to isolate the effect of learning about β_t from that of learning about g, Figure 3 presents impulse response functions to the same shock but agents update only their beliefs about the distribution g. As shown in the figure, both expected inflation and expected interest rate rise permanently after the shock, reflecting agents' revised estimates of the distribution. Nonetheless, since this is a quasi-linear model, these permanent shifts are quantitatively negligible. Increasing the risk aversion parameter value would somewhat amplify the expectation effect. In a highly nonlinear model – with nonlinearities from preferences and costly bankruptcy – Kozlowski et al. (2020) show that learning about the shock distribution generates quantitatively significant persistence after a temporary disaster shock. The same mechanism is present in our framework, but in a linearized model the magnitude of the effect remains very small.

Taken together with the results in Figure 2, we highlight a strong interaction effect from changing expectations: learning about either policy responsiveness or the shock distribution in isolation generates very little persistence, but learning about both simultaneously amplifies the effect substantially.

[10] The decline of the expected interest rate is the result of two forces (see equation 6): on the one hand, persistently lower $\widehat{\beta}_t$ leads to persistently lower expected interest rates (direct effect); on the other hand, higher expected inflation (through lower $\widehat{\beta}_t$) increases expected interest rate (indirect effect). Quantitatively, the direct effect dominates under the current calibration.

Figure 3: Impulse response functions when β_t is common information



3.5 Importance of a Large and Rare Shock

In this section, we emphasize that the presence of a large and rare shock – i.e., a tail event – is crucial for generating the mis-perception of the policy response and, consequently, for producing persistent macroeconomic outcomes. Intuitively, agents mistakenly attribute part of the rise in inflation to a lower β_t, since they do not anticipate the supply shock to be as severe as it actually is. Put differently, their perceived shock distribution excludes large, rare supply shocks, leading them to misinterpret the source of inflation.

To illustrate the importance of a tail event, we perform the following counterfactual analysis. First, we replace the large negative oil supply shock in 2020Q2 with a smaller value equal to one standard deviation of the sample observations (see Figure A.3), and re-estimate the kernel density of the oil supply shock series period-by-period. As shown in Panel (b) of Figure A.3, the much smaller shock in 2020Q2 has a negligible effect on the estimated distribution.

We solve the model again period-by-period, using the new shock distribution per period. Then, as before, we introduce a one-time negative supply shock in period 7, but now the magnitude of the shock is only one standard deviation. Figure 4 shows that versions of the model with and without learning are almost observationally equivalent. While agents still under-perceive the policy responsiveness ($\widehat{\beta}_t < \beta_t$), the degree of error is very small, and as a result, leads to almost no persistence in the macroeconomic variables.[11]

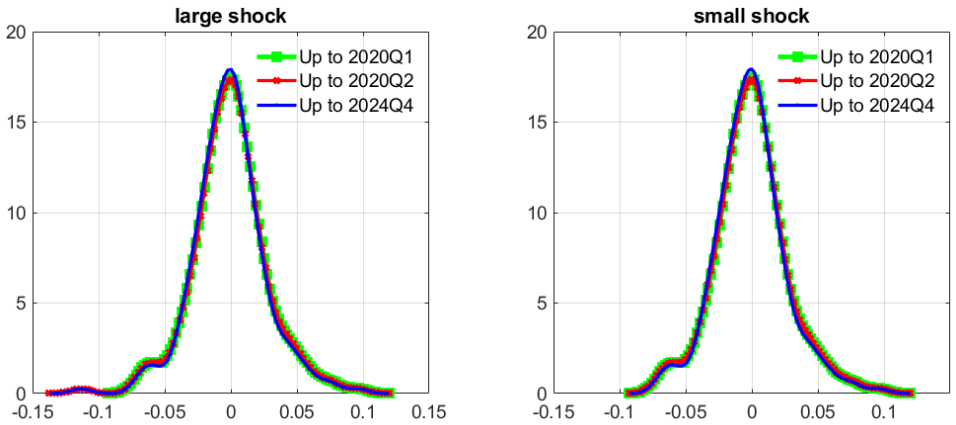
3.6 How Different is 2020 From the 1980s?

The high inflation episodes of the 1980s and 2020s were both driven by large negative supply shocks but differed in two key respects. First, inflation expectations were less well anchored in the 1980s than in the 2020s (see, for example, Blanchard and Gali, 2007; Blanchard and Riggi, 2013). Second, the negative supply shock experienced in

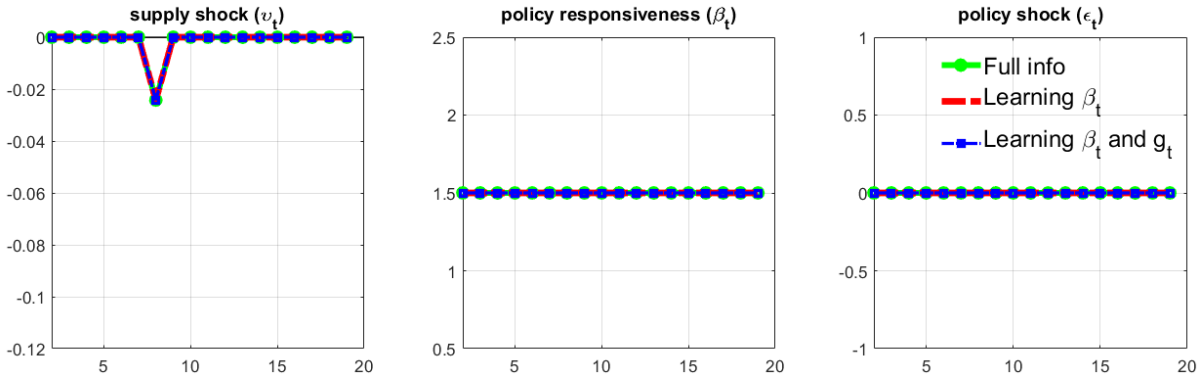
[11] A qualitative difference relative to Figure 2 emerges in the behavior of the expected output gap, $\mathrm{E}_t[x_{t+1}]$, when agents lack information about β_t or g. Following a large negative supply shock in period t, the expected output gap in $t+1$ increases (Panel (e), Figure 2), whereas after a small shock, $\mathrm{E}_{t+1}[x_{t+2}]$ declines (Panel (e), Figure 4). This contrast reflects two offsetting mechanisms. First, because the expected output gap decreases with $\widehat{\beta}_{t-1}$, a fall in $\widehat{\beta}_t$ – *ceteris paribus* – raises $\mathrm{E}_{t+1}[x_{t+2}]$. Second, a leftward shift in the estimated distribution (even slightly) reduces $\mathrm{E}_{t+1}[x_{t+2}]$ relative to $\mathrm{E}_t[x_{t+1}]$, for any values of $\widehat{\beta}_t$. Quantitatively, the first channel dominates when the supply shock is large, as $\widehat{\beta}_t$ declines sharply. We illustrate these two channels through policy functions in Figure A.1.

Figure 4: Impulse response functions following a small supply shock

(a) Estimated distribution $\widehat{g}_t(\cdot)$



(b) Actual shocks



(c) Perceived shocks

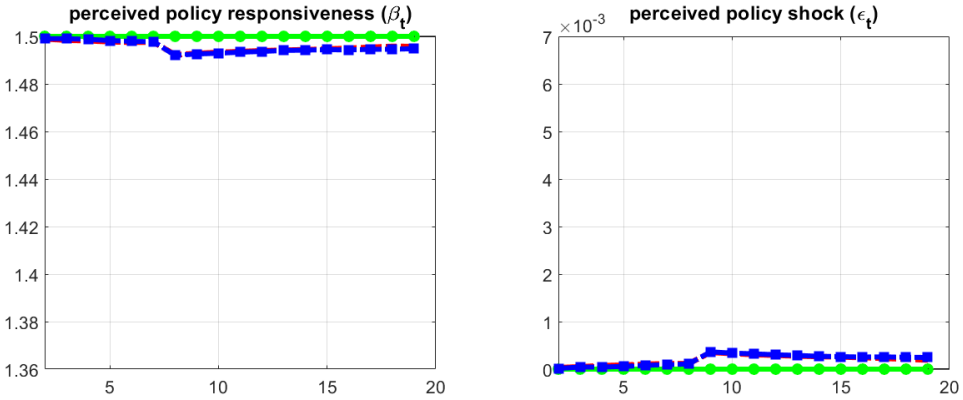
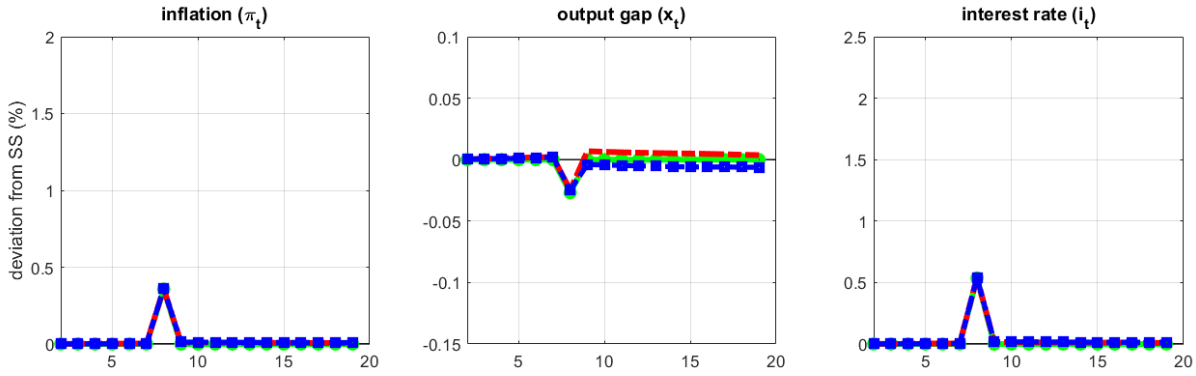


Figure 5: Impulse response functions following a small supply shock

(d) Macro variables



(e) Expectations

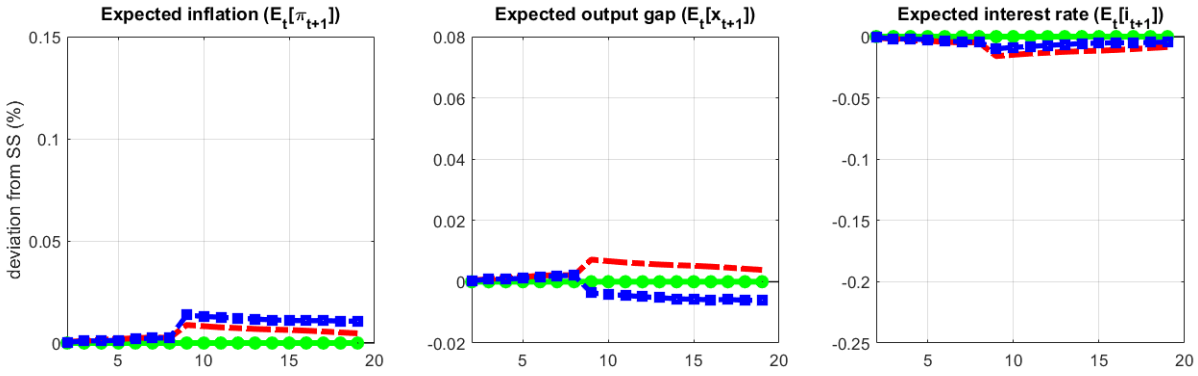
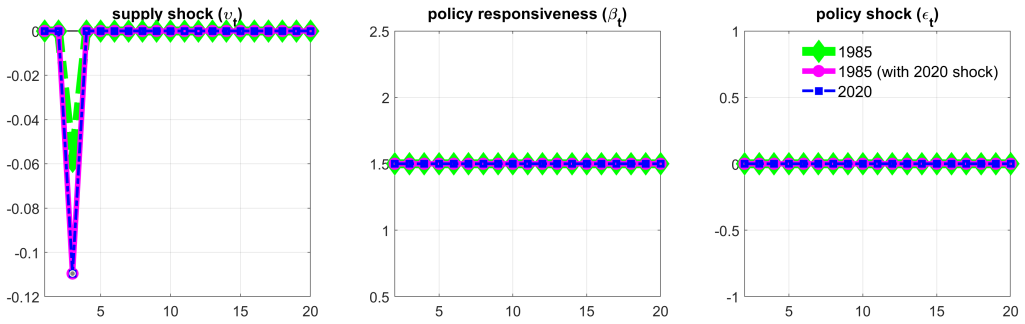
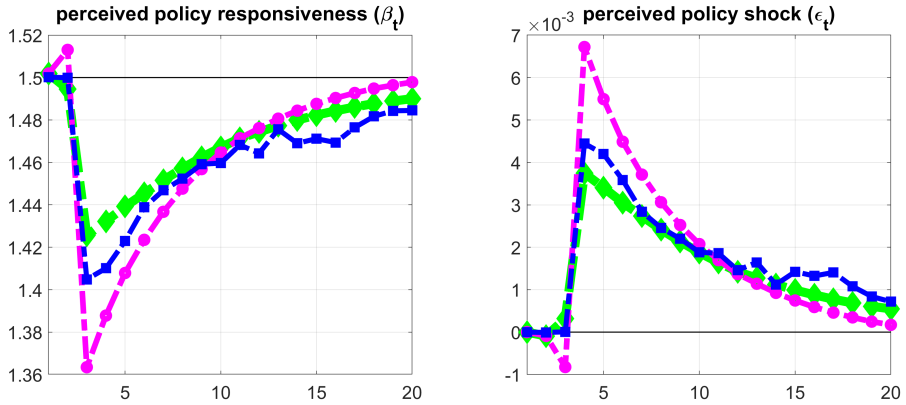


Figure 6: Impulse response functions comparing the 1980s and 2020

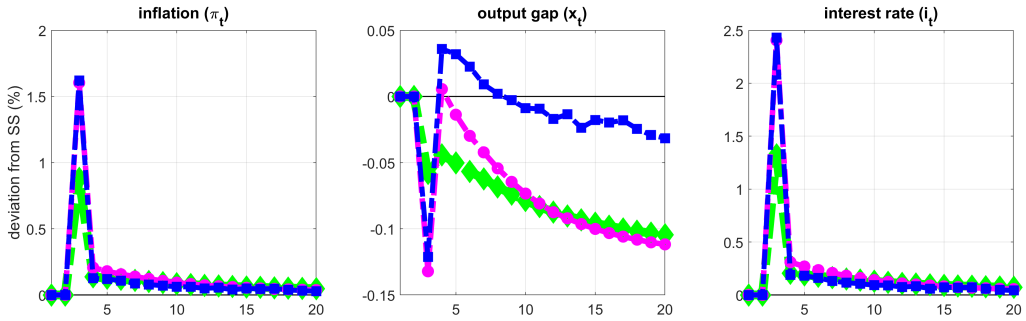
(a) Actual shocks



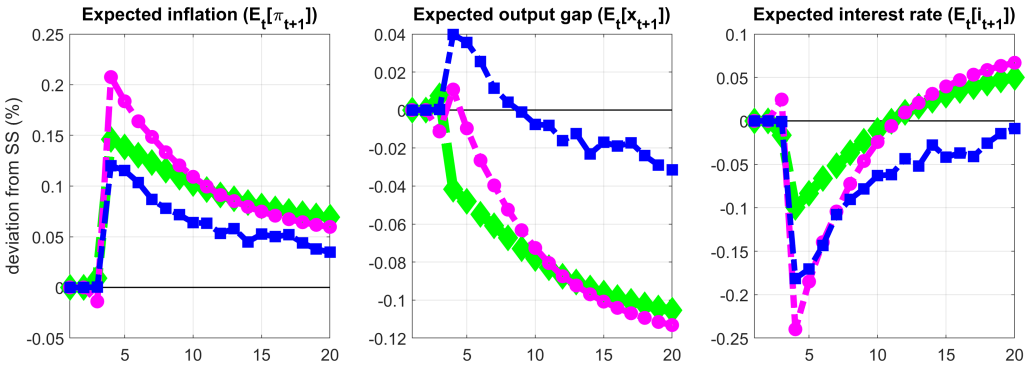
(b) Perceived shocks



(c) Macro variables



(d) Expectations



1985Q2 was smaller — approximately two standard deviations.[12]

In this section, we conduct an experiment designed to mimic the supply driven high-inflation environments of 1985Q2 and 2020Q2. For the 1980s episode, we consider two alternative specifications. In the first, we apply a supply shock of the same magnitude as in 2020 to facilitate a direct comparison of the underlying mechanisms; in the second, we impose a smaller shock — approximately two standard deviations — to better match the magnitude of the shock observed in 1985Q2. The results are presented in Figure 6. For the remainder of the section, we focus on the comparison of the baseline results from the 2020s episode with the first scenario for the 1980s, in which the shock size is held constant. This comparison allows us to highlight differences that arise from factors other than the magnitude of the shock itself.

To capture the main features of the high-inflation episode of the 1980s, we set the priors of $\widehat{\beta}_t$ at $t = 1$ to be three times as dispersed as those in the baseline specification. In addition, we feed into the model the estimated density $\widehat{g}_{1985Q1}$ prior to the shock and $\widehat{g}_{1985Q2}$ thereafter.[13]

The main takeaway is shown in Figure 6. The model generates stagflation — persistently high inflation and low output gap — in the 1980s episode, but not in 2020. Inflation as well as expected inflation rise persistently in both experiments following the one-period shock, with greater persistence in the 1980s episode. On the contrary, the output gap shows a different behaviour. While it overshoots into the positive territory after the initial decline in the 2020 scenario, it remains persistently below its initial steady state in the 1980 scenario.

To understand why the model generates stagflation in the 1980 scenario, we disentangle the underlying mechanisms using two additional counterfactual experiments (Figure A.5). First, we consider an experiment in which the priors over $\widehat{\beta}_t$ are much more tightly anchored, so that the primary channel operates through agents' simultaneous learning about the shock distribution g and policy responsiveness β_t (in other words, the policy functions are time-varying). Second, we consider an experiment in which agents are assumed to know the shock distribution (i.e. with time-invariant policy functions), while $\widehat{\beta}_t$ remains unanchored; in this case, the dominant channel is learning about policy responsiveness β_t under dispersed priors.

The first observation is that more dispersed priors over β_t — corresponding to less

[12] We use the negative oil supply shock in 1985Q2 as an illustrative example, as it is among the largest shocks observed between 1975 and 1992. As shown in Figure 1.a, all supply shocks over this period lie within two standard deviations.

[13] See Figure A.4 for how the estimated shock distribution has changed from the 1980s through the early 1990s.

anchored inflation expectations — amplify the core mechanism of the model. Specifically, following a large supply shock, the perceived $\widehat{\beta}_t$ declines by more, *ceteris paribus*, leading to more persistent rise in inflation and greater output gap overshooting. This is because expected inflation and expected output gap react by more as the perceived $\widehat{\beta}_t$ declines by more.[14]

The second observation is that output gap does not overshoot in the counterfactual experiment in which the priors over $\widehat{\beta}_t$ are much more tightly anchored, but agents must learn about the shock distribution g and policy responsiveness β_t (Figure A.5). As expectations are much more tightly anchored, the perceived $\widehat{\beta}_t$ declines by much less and the expected output gap remains persistently below the steady state after the shock (Figure A.5.d, green line). Zooming into the policy functions in the 1980s scenario (Figure A.7), we see that the expected output gap $E_t[x_{t+1}]$ is lower in 1985Q2 than in 1985Q1, even if the perceived β_t returns to the mean (1.5). Overall, our 1980s simulations show stagflation: inflation rises more persistently in the former, and the output gap does not exhibit positive overshooting.

3.7 Robustness and Extensions

Adding output gap to the Taylor rule As a robustness check, we now allow the policy rate to respond directly to changes in the output gap, bringing the specification closer to a textbook Taylor rule (Galí, 2015):

$$i_t = \beta_t \pi_t + \theta x_t + \varepsilon_t. \tag{9}$$

We calibrate θ to be 0.5, and repeat the experiment using the same large negative supply shock as in Figure 2.[15] We report the results under the new Taylor rule in Figure A.8. The main difference from Figure 2 is quantitative: inflation, the output gap, and the interest rate all increase more sharply after the shock dissipates under the new Taylor rule. In other words, learning produces more persistent effects on macroeconomic variables compared to the baseline model. This additional persistence arises from stronger expectations effects (see panel (c)) under the new rule. Furthermore, it is worth noting the feedback from the macro variables to agents' beliefs about the policy responsiveness, $\widehat{\beta}_t$. As the macro variables exhibit more persistence, we also see

[14]This mechanism is illustrated in Figure A.6, which is a comparative static exercise of the baseline model (Figure 2) with more dispersed priors.

[15]To maintain tractability, we assume that that θ is constant here, and agents only learn about β_t. One natural (though computationally intensive) extension would be to allow θ to be time-varying and latent, and having agents to simultaneously learn about both θ_t and β_t.

slower changes in beliefs. In particular, the rate at which $\widehat{\beta}_t$ returns to the steady state is much slower here compared to the baseline model in Figure 2.

Anticipated utility approach In the baseline model, we assume that the private sector makes decisions in time t with the belief $\widehat{\beta}_{t-1}$, following the anticipated utility approach in the learning literature. Given the evolution of the latent variable in our setting (4), we now assume that the private sector makes decisions in time t with the belief $(1-\rho)\bar{\beta} + \rho\widehat{\beta}_{t-1}$, which is the rational forecast of β_t. Figure A.11 presents the impulse response functions to a large supply shocks. The results are both qualitatively and quantitatively similar to the baseline model.

Stronger inflation aversion We consider a case where the central bank realizes that a rare but large negative supply shock could feed into higher inflation expectations and hinder the achievement of the inflation target. For that reason, the central bank increases its reaction to inflation after observing the tail supply shock.

Figure A.10 illustrates the impulse response functions of the model variables following both a negative supply shock and, simultaneously, a positive shock to the central bank's reaction to inflation, i.e. the central bank responds more strongly to inflation when the large negative supply shock occurs. As before, we compare the impulse response functions in three different models, which differ in terms of the public's information set.

In the full information model, inflation rises upon impact of the negative supply shock, but falls below the steady state immediately afterwards due to the central bank's stronger response to inflation. Both shocks lower the output gap, which then converge back to the steady state without overshooting. The nominal interest rate rises upon impact and gradually goes back to the steady state.

As shown in panels (b) and (d), the two shocks have markedly different effects on the public's perceptions and expectations across models. Naturally, in the full-information model, the perceived policy responsiveness and policy shock coincide exactly with their actual counterparts. Expected inflation and output both decline in the full information model (since they are decreasing functions of β_t), while expected interest rate increases.

By contrast, when agents simultaneously learn about the policy responsiveness and the shock distribution, the policy shock appears much less effective, as the responses of all variables closely resemble those without a policy shock, particularly in the short

and medium run.[16] Despite the increase in actual β_t, agents still perceive it to have decreased when the large supply shock is realized. The dynamics of the perceived β_t continue to drive the behavior of expected inflation, the output gap, and the interest rate in the same way as before.

4 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 the public's 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.

4.1 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.[17] 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.[18] 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 monthly sample using cubic splines for potential real GDP and the Chow and Lin (1971) method for real GDP where 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, which is the Fed's preferred inflation gauge. 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

[16] See Figure A.12 for a clearer comparison.

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

[18] We obtain very similar results using alternative shadow rate series or the unadjusted fed funds rate (see Figure A.14 as an example).

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.[19] 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).[20] The real oil price is the West Texas Intermediate (WTI) spot price deflated by U.S. CPI which enters the model in growth rates.

4.2 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 Deboroli, 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 $u_t^{s,oil}$ 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 $\boldsymbol{\varepsilon}_t$ 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.[21] To differenti-

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

[20] 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.13 and A.14 in the appendix).

[21] 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.

ate between small and large shocks, we use the square of the shock as the nonlinear function: $f(u_t^{s,oil}) = (u_t^{s,oil})^2$.[22]

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.

4.3 Main Findings

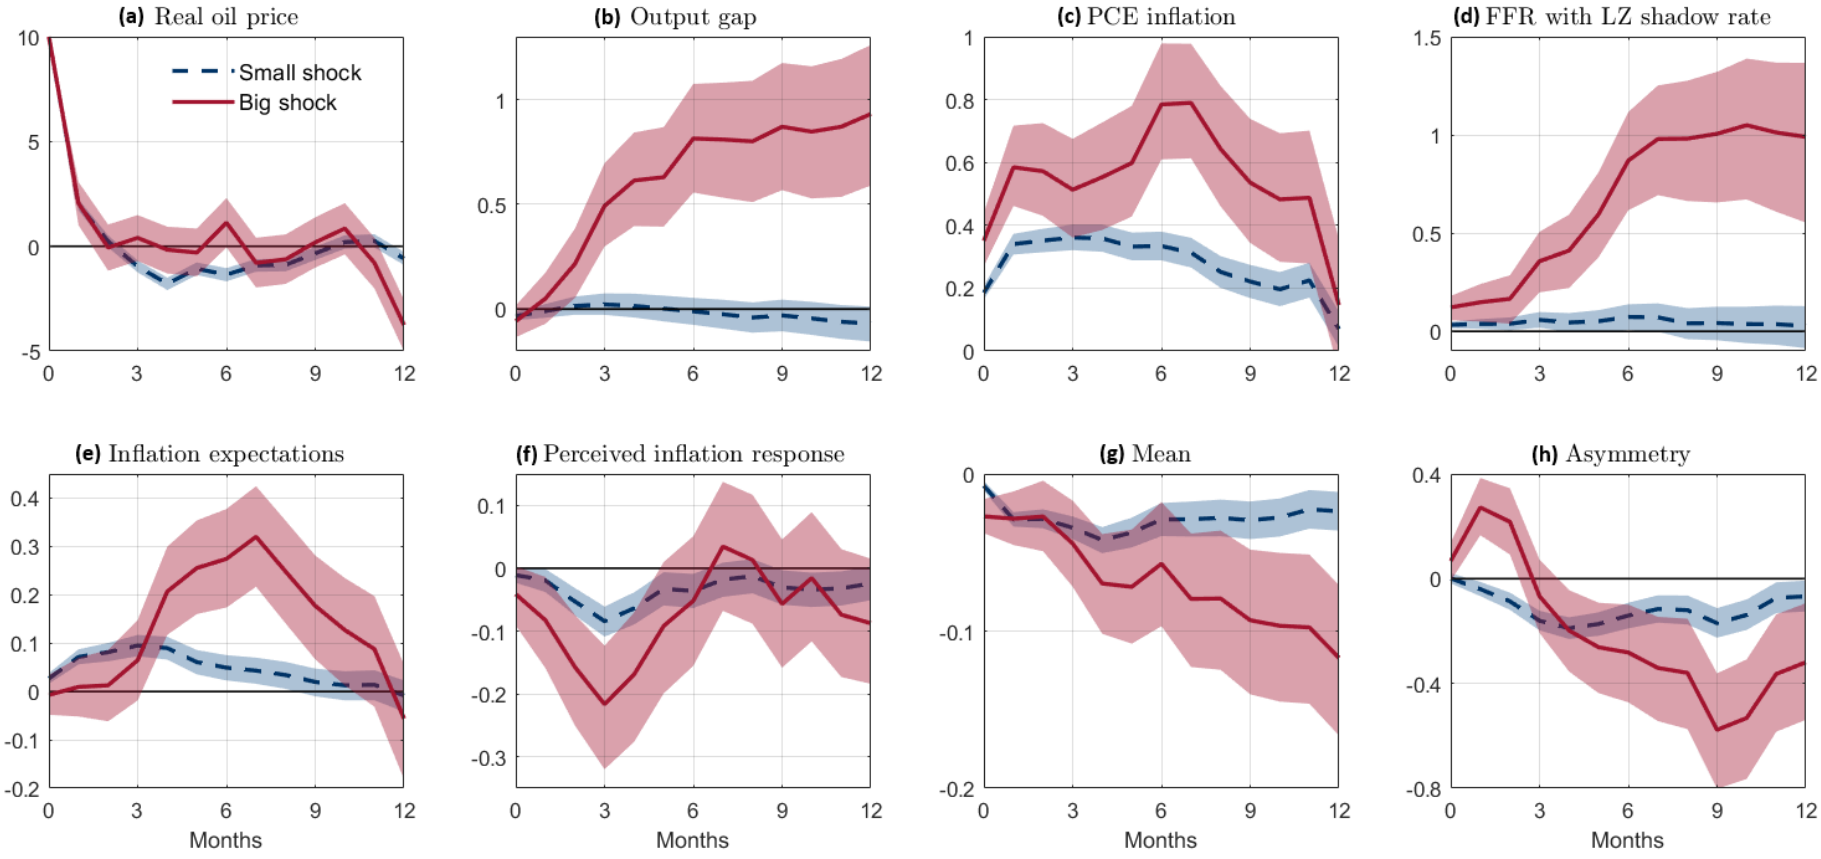
Figure 7 shows the impulse responses from the nonlinear VAR to oil supply shocks, with one standard deviation error bands.[23] We consider small and large shocks, which are one and six standard deviations, respectively. To make the impulse responses directly comparable, we normalize them to induce a 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 7, both the mean and the asymmetry measure in-

consumer price index. Their interest centers on the nonlinear transmission of positive versus negative oil supply shocks to the macroeconomy.

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

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

Figure 7: Impulse response functions from nonlinear VAR



dicate 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 2, this manifests as a decline in perceived policy responsiveness, even though the true policy rule remains unchanged. The VAR evidence closely mirrors this pattern. Following a large supply shock, the impulse response functions from the nonlinear VAR also show a decline in perceived inflation responsiveness (Figure 7, Panel (f)), indicating that agents interpret 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 model's prediction that a perceived weakening of policy feeds into 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 exhibits persistent overshooting after a temporary dip upon impact of the negative supply shock (Panel (b)).

Figure 7 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, and the nominal interest rate are substantially stronger and more persistent than those to small shocks; output gap overshooting only happens in the case of large shocks. These asymmetric patterns are consistent with the results in the model (Figure 4), 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 inflation and output gap overshooting.

5 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-gap overshooting after tail events. U.S. data from 1976–2024 supports this mechanism.

Our findings have important implications for monetary policy design and communication. When the unthinkable happens, private agents' beliefs about the policy rule is 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. *The 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 (2023). Pandemic, War, Inflation: Oil Markets At a Crossroads? NBER Working Paper No. 31496.

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.

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.

Cogley, Timothy, Christian Matthes, and Argia M Sbordone (2015). Optimized Taylor rules for disinflation when agents are learning. *Journal of Monetary Economics 72*, 131–147.

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.

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 (2023). Oil Prices, Monetary Policy and Inflation Surges. NBER Working Paper No. 31263.

Galí, Jordi (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton University Press.

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.

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

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.

Nakamura, Emi, Venance Riblier, and Jón Steinsson (2025). Beyond the Taylor Rule. NBER Working Paper No. 34200.

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.

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.

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

Online Appendix

Christiane Baumeister Fiorella De Fiore Jasmine Xiao

March 2026

A Algorithm

A.1 Estimating the supply shock density

First, we estimate the density of the supply shock $\widehat{g}_t$ for each t using the following normal kernel density estimator:

$$\widehat{g}_t = \frac{1}{n_t} \sum_{s_0}^{n_t-1} \Omega\big(\mathbf{x} - \mathbf{x}_{t-s}; \Theta_t\big),$$

where n_t is the number of available observations at t and $\Omega(\cdot)$ is the multivariate normal density function with covariance Θ_t, also known as the bandwidth matrix. We use the reference rule for the optimal bandwidth that minimizes the expected squared error.

We use an expanding window of the data to update their estimates, generating a sequence of $\{\widehat{g}_t\}$. The data is on oil supply shocks from Baumeister and Hamilton (2019) from 1975Q2, with the first and last estimated densities being $\widehat{g}_{1990Q1}$ and $\widehat{g}_{2024Q4}$, respectively. The shock series and estimated beliefs (at 2020Q1, 2020Q2 and 2024Q4) are shown in Figure 1.

As discussed in Kozlowski, Veldkamp, and Venkateswaran (2020), the kernel density estimator with the optimal bandwidth is approximately a martingale:

$$\mathrm{E}_t[\widehat{g}_{t+1}|\mathcal{I}_t] \approx \widehat{g}_t, \tag{A.1}$$

where $\mathcal{I}_t$ captures the information at time t. This means that agents expect her future belief about the distribution of v_t to be the same as her current beliefs, and as a result, any changes in beliefs (due to new information) are permanent in expectation.

A.2 State-space representation

Substitute the Taylor rule (3) in the dynamic IS equation (2), and rearrange. Together with (1) we obtain a two-equation system in inflation and output gap:

$$y_t = \mathbf{A_t}^{-1}\mathbf{B}\mathrm{E}_t y_{t+1} + \mathbf{A_t}^{-1}\mathbf{C}\eta_t \tag{A.2}$$

where $y_t = (\pi_t,\ x_t)'$, $\eta_t = (v_t,\ \varepsilon_t)'$,

$$\mathbf{A_t} = \begin{pmatrix} 1 & -\kappa \\ \sigma^{-1}\beta_t & 1 \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} \delta & 0 \\ \sigma^{-1} & 1 \end{pmatrix}, \quad \mathbf{C} = \begin{pmatrix} -1 & 0 \\ 0 & -\sigma^{-1} \end{pmatrix}$$

(A.2) is the measurement equation, and the state equation is:

$$\widetilde{\beta}_t = \rho\widetilde{\beta}_{t-1} + \xi_t. \tag{A.3}$$

where $\widetilde{\beta}_t \equiv \beta_t - \overline{\beta}$. Recall that $\xi_t \sim iid\ N(0,\sigma_\xi^2)$ and $\varepsilon_t \sim iid\ N(0,\sigma_\varepsilon^2)$, but v_t is potentially non-normal with distribution g. v_t, ε_t and ξ_t are mutually uncorrelated. κ, δ, ρ, σ, $\overline{\beta}$, σ_ξ and σ_ε are parameters whose values are known.

Learning takes the form of particle filtering for two reasons. First, since β_t is inside a matrix inverse $\mathbf{A_t}^{-1}$, the mapping from the state β_t to the observable y_t is generally nonlinear. Second, as the cost-push shock that loads into the measurement equation v_t is potentially non-normal, the measurement noise is non-Gaussian.

A.3 Algorithm under asymmetric information

This section describes the numerical algorithm used to compute the learning equilibrium in the asymmetric information model using particle filtering:

1. Inherit $\{\widehat{\beta}^i_{t-1}, \omega^i_{t-1}\}_{i=1}^N$ from $t-1$, where $\{\omega^i_{t-1}\}_{i=1}^N$ is the sequence of weights for the particles $\{\widehat{\beta}^i_{t-1}\}_{i=1}^N$. Compute the weighted average:

$$\widehat{\beta}_{t-1} = \sum_{i=1}^{N} \omega^i_{t-1}\widehat{\beta}^i_{t-1}$$

2. For each $i = 1,2,...,N$, given the mean belief $\widehat{\beta}_{t-1}$, realized v_t and ε_t, and estimated $\widehat{g}_{t-1}$, find the equilibrium macro outcomes (π_t, x_t, i_t):

- Find (π_t, x_t) that solves:

$$\pi_t = \delta \mathrm{E}_t \pi_{t+1} + \kappa x_t - v_t,$$
$$x_t = \mathrm{E}_t x_{t+1} - \sigma^{-1}(\underbrace{\widehat{\beta}_{t-1}\pi_t + \widehat{\varepsilon}_t}_{\text{perceived rule}} - \mathrm{E}_t \pi_{t+1}).$$

  on a grid for $\widehat{\varepsilon}_t$. Denote the policy functions as $h_\pi(\widehat{\varepsilon}_t)$ and $h_x(\widehat{\varepsilon}_t)$. In computing $\mathrm{E}_t[\pi_{t+1}]$ and $\mathrm{E}_t[x_{t+1}]$ in (1) and (2), we use the property that the martingale property of the kernel density estimator (A.1).

- Since the central bank knows the true β_t, the observed interest rate follows the actual Taylor rule:

$$i_t = \beta_t \pi_t + \varepsilon_t.$$

  In equilibrium, agents' beliefs about the policy shock $\widehat{\varepsilon}_t$ must reconcile the interest rate under the perceived rule with the observed interest rate under the actual rule:

$$\widehat{\varepsilon}_t^* = (\beta_t - \widehat{\beta}_{t-1})\pi_t + \varepsilon_t$$

  Use bisection method to find the equilibrium belief, denoted by $\widehat{\varepsilon}_t^*$, from the $\widehat{\varepsilon}_t$-grid, given $(\beta_t, \widehat{\beta}_{t-1}, \varepsilon_t, v_t; \widehat{g}_{t-1})$.

- Evaluate $h_\pi(\widehat{\varepsilon}_t)$ and $h_x(\widehat{\varepsilon}_t)$ at $\widehat{\varepsilon}_t^*$ to obtain the equilibrium outcomes (π_t^*, x_t^*), given $(\beta_t, \widehat{\beta}_{t-1}, \varepsilon_t, v_t; \widehat{g}_{t-1})$. The equilibrium interest rate simultaneously satisfies:

$$i_t^* = \beta_t \pi_t^* + \varepsilon_t$$
$$i_t^* = \widehat{\beta}_{t-1}\pi_t^* + \widehat{\varepsilon}_t^*.$$

3. For each particle $\widehat{\beta}_{t-1}^i$, draw a particle $\widehat{\beta}_t^i$ according to:

$$\widehat{\beta}_t^i = (1-\rho)\overline{\beta} + \rho\widehat{\beta}_{t-1}^i + \xi_t^i, \quad \text{where} \quad \xi_t^i \sim iid\ N(0, \sigma_\xi^2).$$

4. Form the particle likelihood given the realized data (π_t^*, x_t^*, i_t^*):

   - Given the observed i_t^* and π_t^*, the monetary-policy shock implied by particle i is

$$\varepsilon_t^i = i_t^* - \widehat{\beta}_t^i \pi_t^*,$$

     and the policy shock perceived by agents under that same candidate particle

is

$$\widehat{\varepsilon}_t^i = (\widehat{\beta}_t^i - \widehat{\beta}_{t-1})\pi_t^* + \varepsilon_t^i = i_t^* - \widehat{\beta}_{t-1}\pi_t^*,$$

i.e. the perceived shock is actually pinned down by the data and does not vary across particles.

- Pick a grid $\{v^k\}_{k=1}^K$. Solve the model at each node k (for each particle i):

$$\begin{pmatrix} \pi_t^{i,k} \\ x_t^{i,k} \end{pmatrix} = F_t(\widehat{\beta}_{t-1},\, v^k,\, \widehat{\varepsilon}_t^i; \widehat{g}_{t-1})$$

- Compute the "distance" from the observed macro variables at each point:

$$d^2(v_t^k, \beta_t^i) = \frac{(\pi_t^* - \pi_t^{i,k})^2}{s_\pi^2} + \frac{(x_t^* - x_t^{i,k})^2}{s_x^2}$$

where s_π^2 and s_x^2 reflect the sample variance of inflation and output gap, respectively (otherwise the variable with larger scale dominates the likelihood).

- The particle likelihood is given by:

$$L_t^i = \phi(\varepsilon_t^i; 0, \sigma_\varepsilon^2) \sum_{k=1}^K w^k \exp\Big(-\frac{1}{2} d^2(v^k, \beta_t^i) \Big)$$

where w_k reflects agents' belief $\widehat{g}_{t-1}$ over v.

5. Update weights:

$$\widetilde{w}_t^i = w_{t-1}^i L_t^i, \qquad w_t^i = \frac{\widetilde{w}_t^i}{\sum_{j=1}^N \widetilde{w}_t^j}.$$

6. Resample if the effective sample size

$$ESS = \frac{1}{\sum_{i=1}^N (w_t^i)^2}$$

is too small. Resample particles according to w_t^i, and reset weights to $\frac{1}{N}$.

A.4 Symmetrically imperfect information model

In Figure A.9, we consider the symmetrically imperfect information case where neither the central bank nor the private sector observes β_t or ε_t in the Taylor rule (3).

At the beginning of period t, both the private sector and the central bank are endowed with the (mean) belief $\widehat{\beta}_{t-1}$ and $\widehat{\varepsilon}_t$. We continue to assume that beliefs about the rule pa-

rameters are updated via Baysian filtering after macro outcomes in the current period are realized. Then agents use the posterior belief with mean $\widehat{\beta}_t$ to determine the equilibrium outcomes in the following period.

If agents know the true distribution of v_t (g), the policy functions are given by the same relation as in the full information case, but using the subjective beliefs, $\widehat{\beta}_{t-1}$ and $\widehat{\varepsilon}_t$. The equilibrium given beliefs is characterized by:

$$\begin{pmatrix} \pi_t \\ x_t \\ i_t \end{pmatrix} = F(\widehat{\beta}_{t-1},\, v_t,\, \widehat{\varepsilon}_t). \tag{A.4}$$

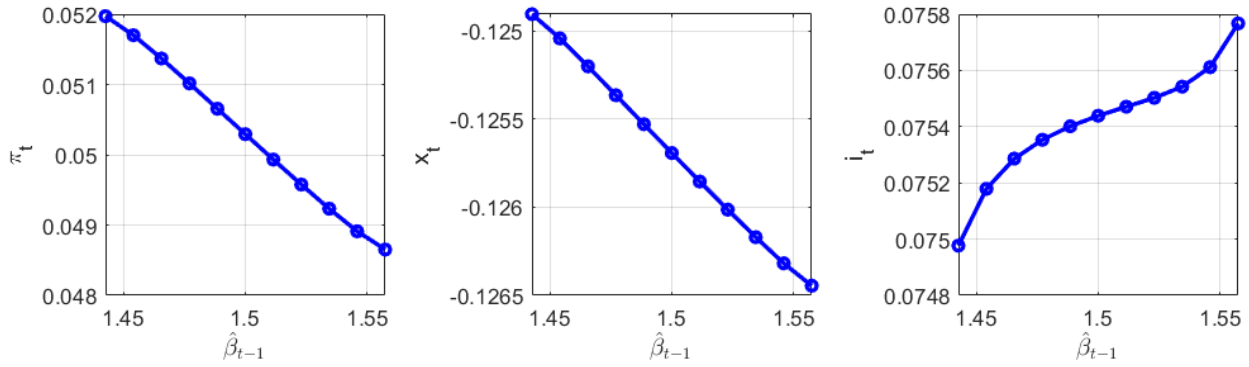
If agents do not know the true distribution g of v_t, and update their belief of the distribution in each period upon observing the realized v_t, then the equilibrium is characterized by:

$$\begin{pmatrix} \pi_t \\ x_t \\ i_t \end{pmatrix} = F_t(\widehat{\beta}_{t-1},\, v_t,\, \widehat{\varepsilon}_t; \widehat{g}_{t-1}). \tag{A.5}$$

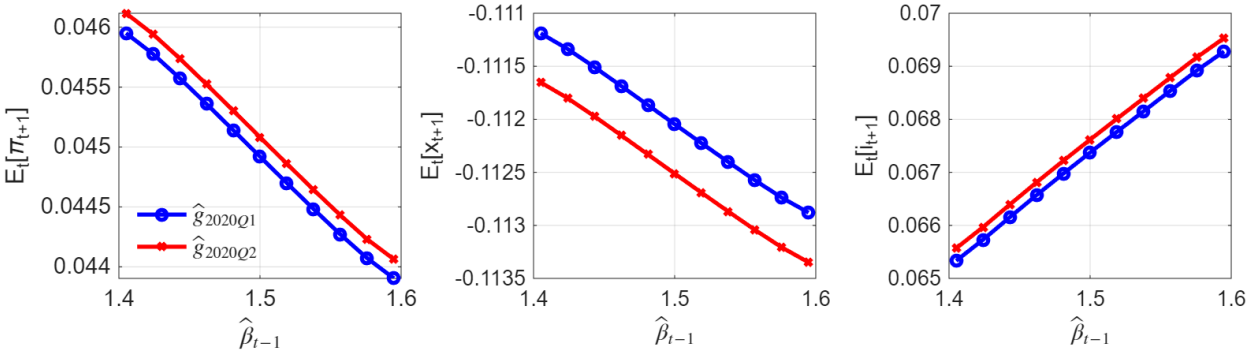
The main difference from the equilibrium under asymmetric information – either (5) or (8) – is that the macro variables are functions of $\widehat{\varepsilon}_t$ here. Unlike in the asymmetric information case, the belief $\widehat{\varepsilon}_t$ can be anything – it does not need to satisfy the equilibrium condition (7).

Figure A.1: Policy functions

(a) Macro variables



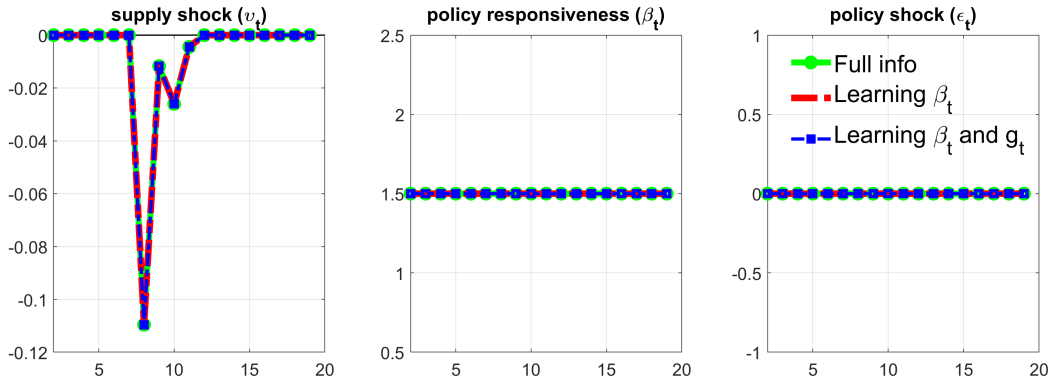
(b) Expectations



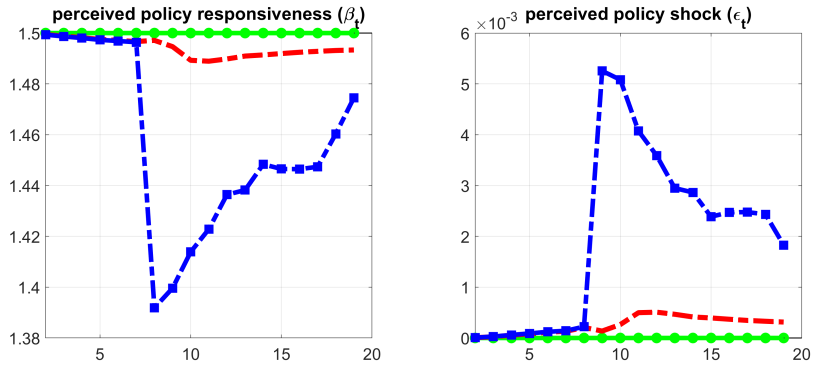
The macro variables (π_t, x_t, i_t) are time-varying functions of $(v_t, \widehat{\varepsilon}_t, \widehat{\beta}_{t-1})$. Panel (a) plots the policy functions at 2020Q2 for different values of $\widehat{\beta}_t$, fixing $\widehat{\varepsilon}_t$ at the mean and v_t at its value in 2020Q2. The expectation variables $(E_t[\pi_{t+1}], E_t[x_{t+1}], E_t[i_{t+1}])$ are time-varying functions of $\widehat{\beta}_{t-1}$. Panel (b) plots these for 2020Q1 (blue) and 2020Q2 (red), respectively. .

Figure A.2: Impulse response functions following consecutive shocks

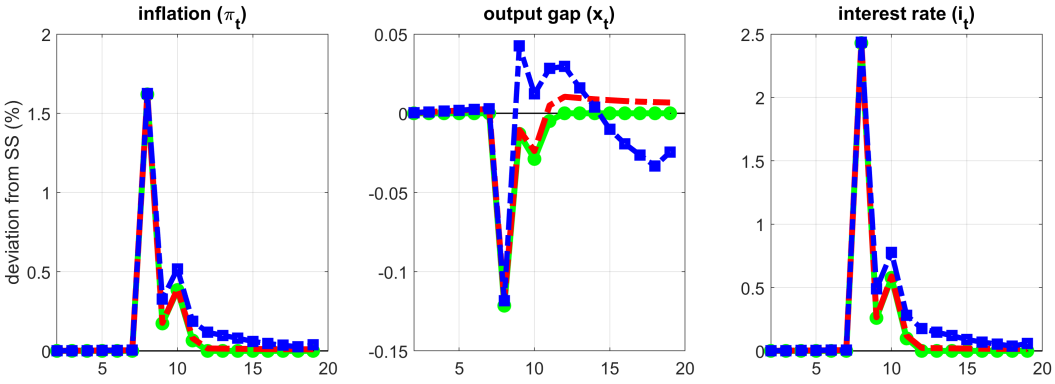
(a) Actual shocks



(b) Perceived shocks



(c) Macro variables



(d) Expectations

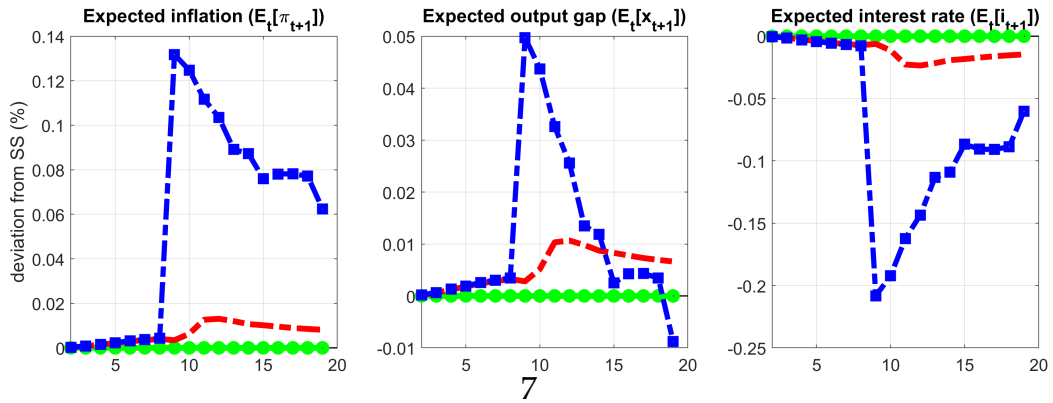
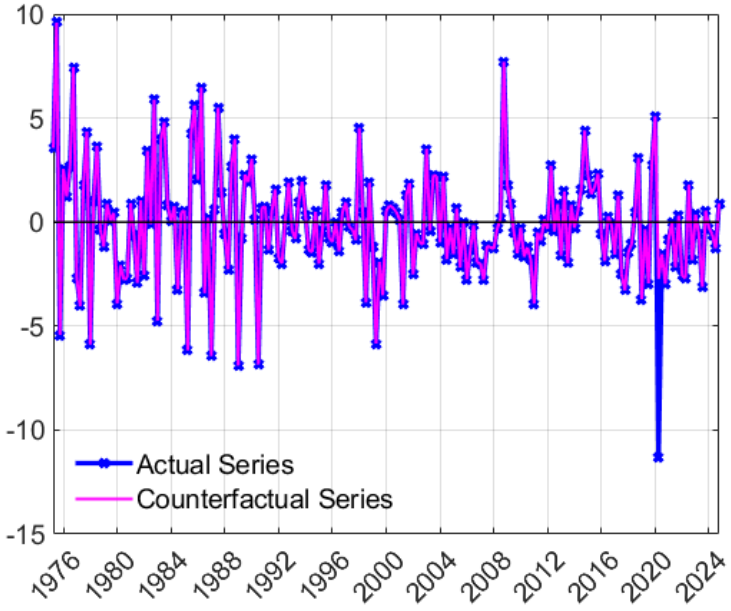
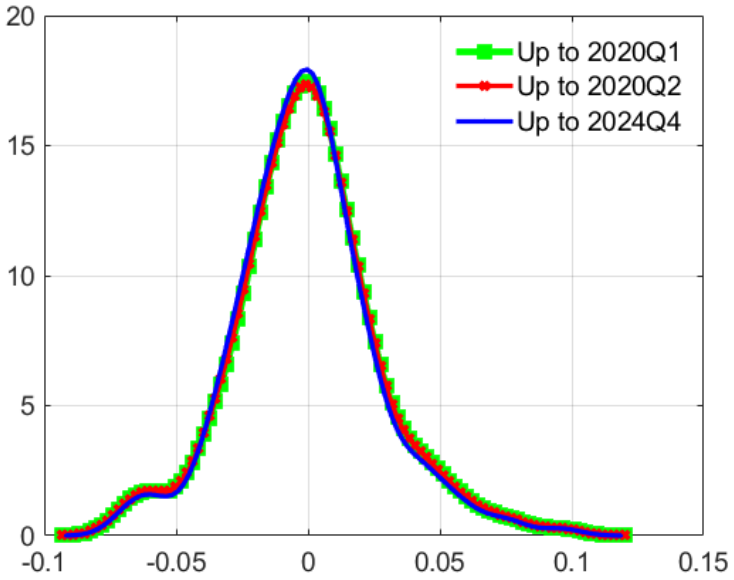


Figure A.3: Small shock counterfactual

(a) Shock series

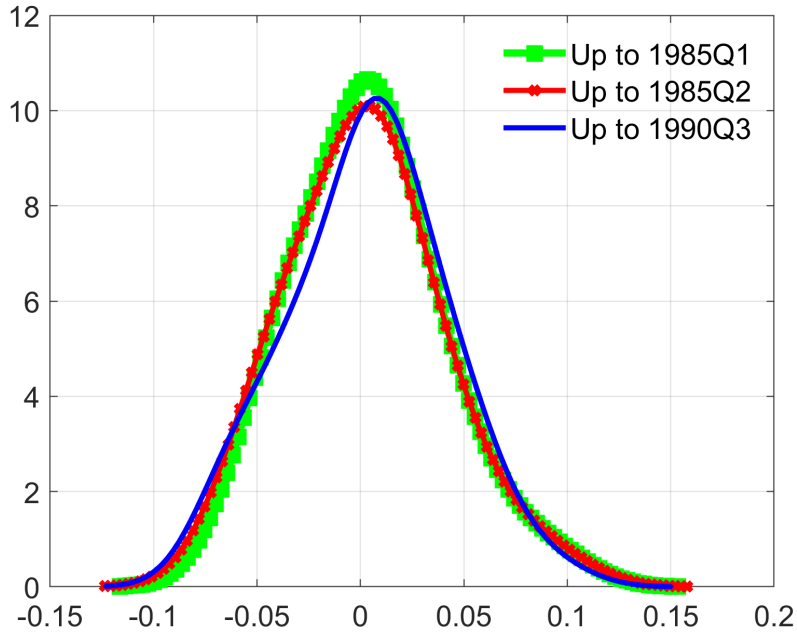


(b) Estimated density



The top panel plots the actual shock series from Baumeister and Hamilton (2019) (in blue) and the counterfactual series (in red) where we replace the large negative shock values in 2020Q2 with a smaller value that is equal to (negative) one standard deviation of the sample. The bottom panel shows the kernel density estimated from full sample for each series.

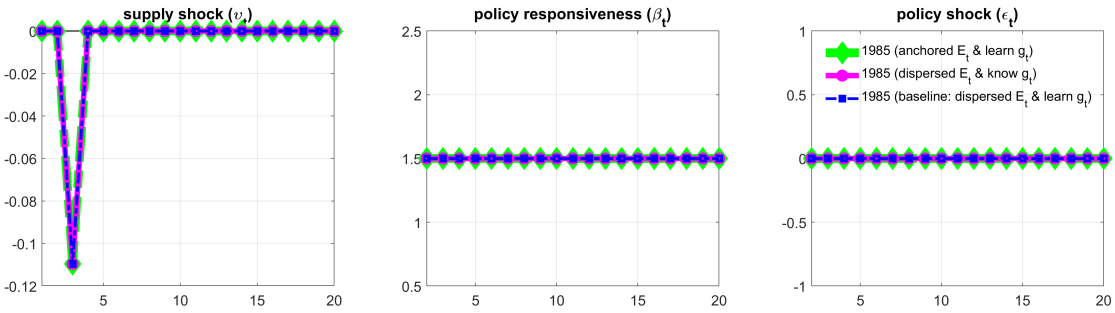
Figure A.4: Estimated density of oil supply shocks in the 1980s/90s



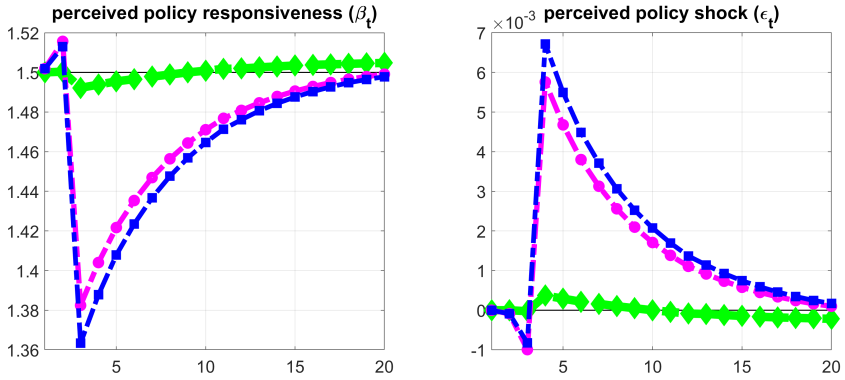
The figure shows the kernel density, estimated from the quarterly oil supply shock series from Baumeister and Hamilton (2019). We use an expanding window of data from 1975Q2 to 1985Q1 (in green), to 1985Q2 (in red), and to 1990Q3 (in blue), respectively.

Figure A.5: Decomposing 1980s' Impulse Response Functions

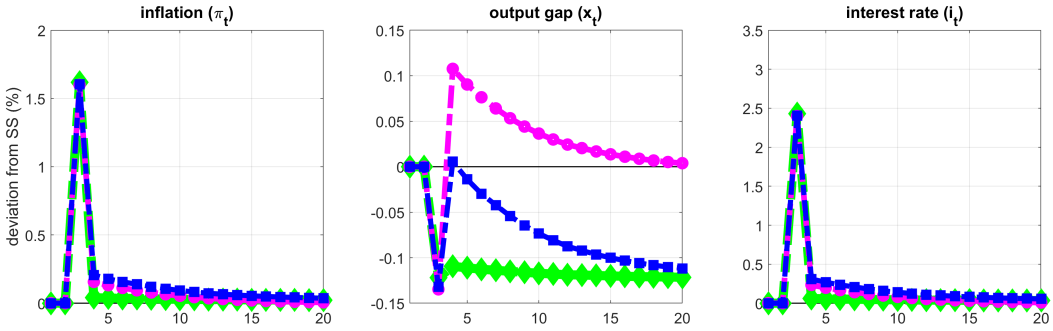
(a) Actual shocks



(b) Perceived shocks



(c) Macro variables



(d) Expectations

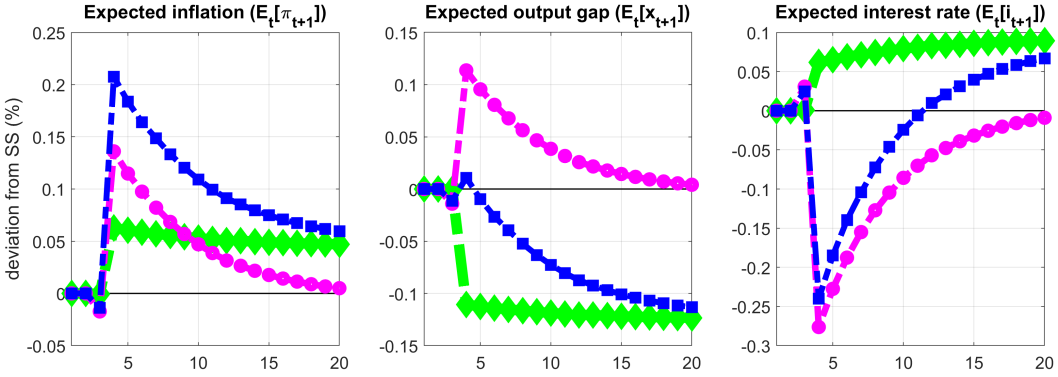
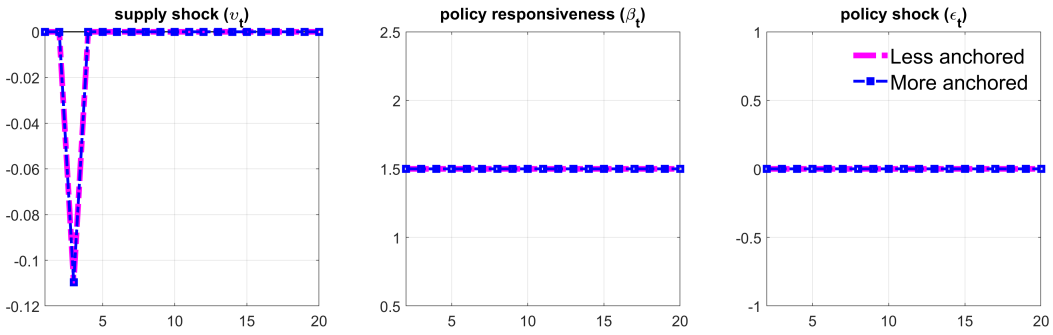
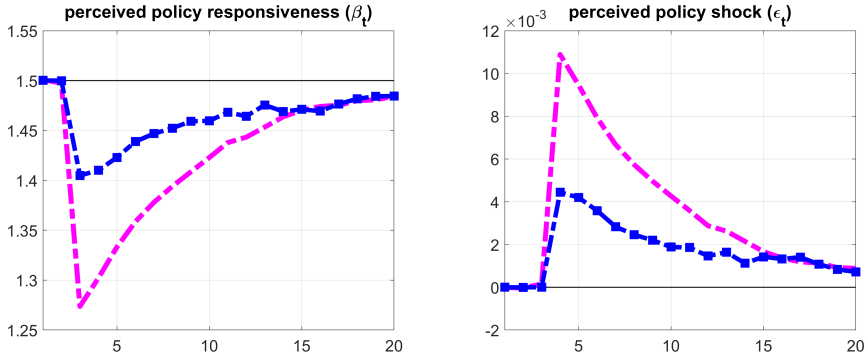


Figure A.6: Impulse Response Functions to a large supply shock with less anchored perceptions of β_t

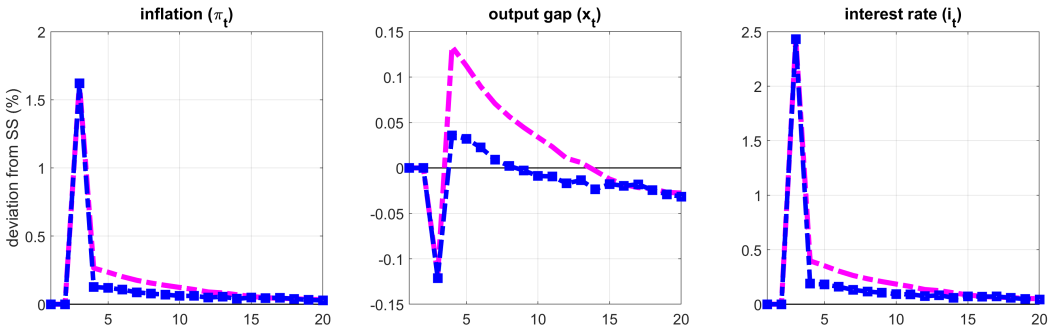
(a) Actual shocks



(b) Perceived shocks



(c) Macro variables



(d) Expectations

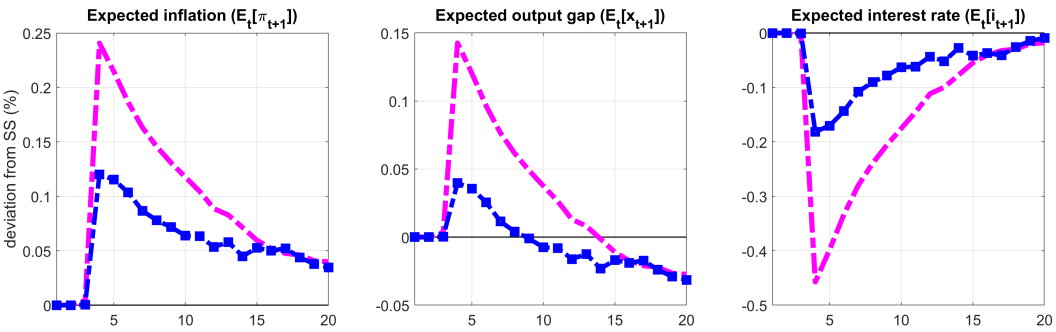
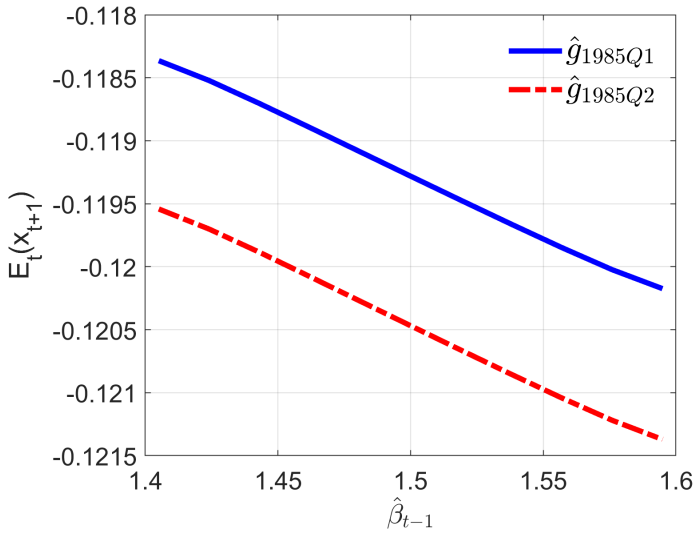


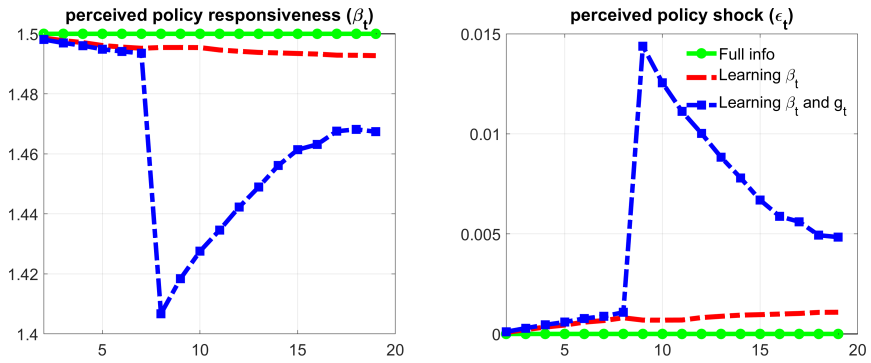
Figure A.7: Policy functions in the 1980s simulation



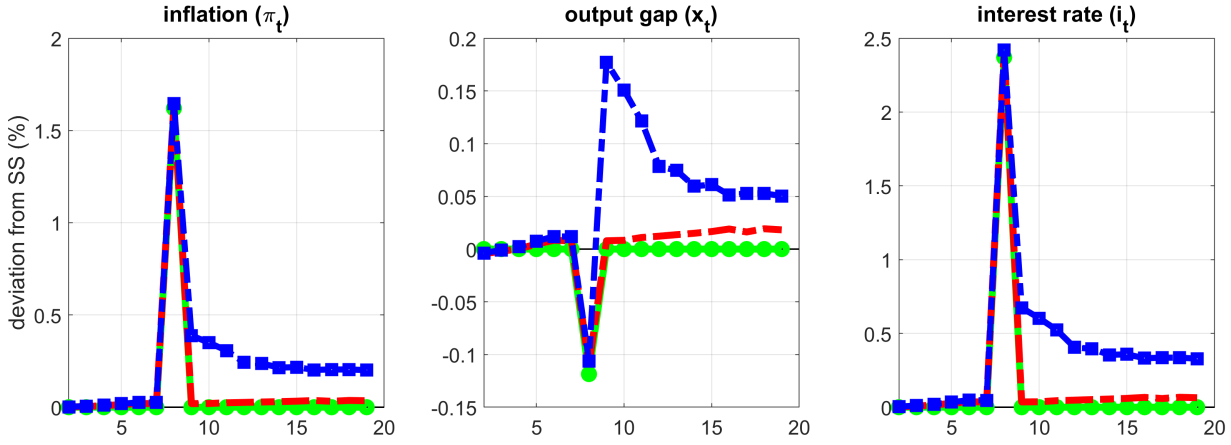
This figure plots the expectation variable $E_t[x_{t+1}]$ as a function of $\widehat{\beta}_{t-1}$ at 1985Q1 (blue) and 1985Q2 (red), respectively.

Figure A.8: Model with output gap-augmented Taylor rule

(a) Perceived shocks



(b) Macro variables



(c) Expectations

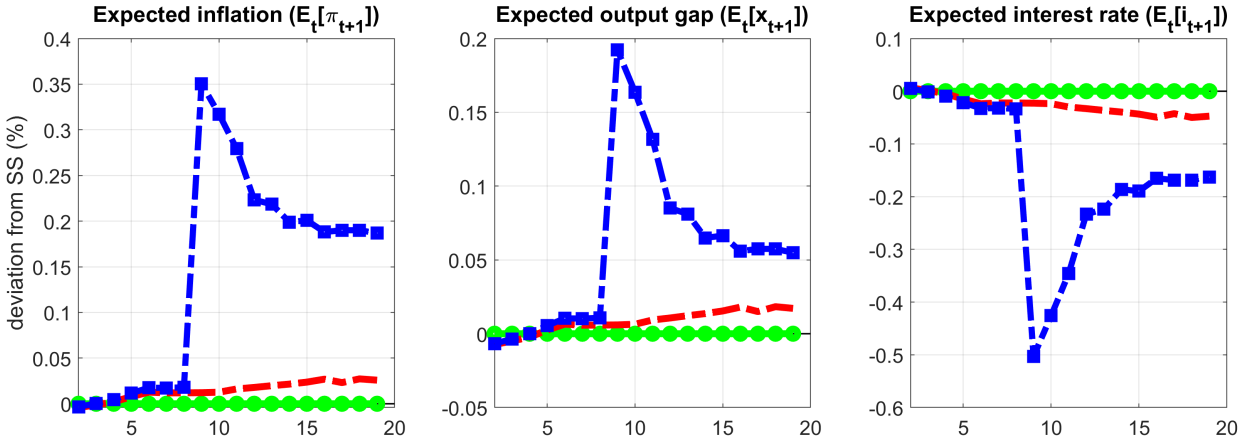
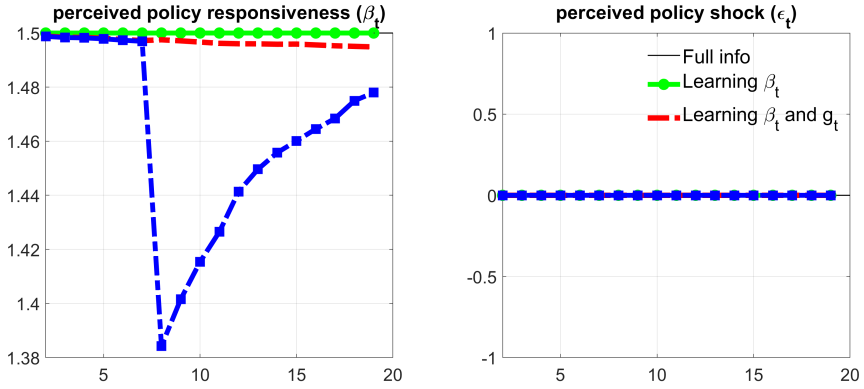
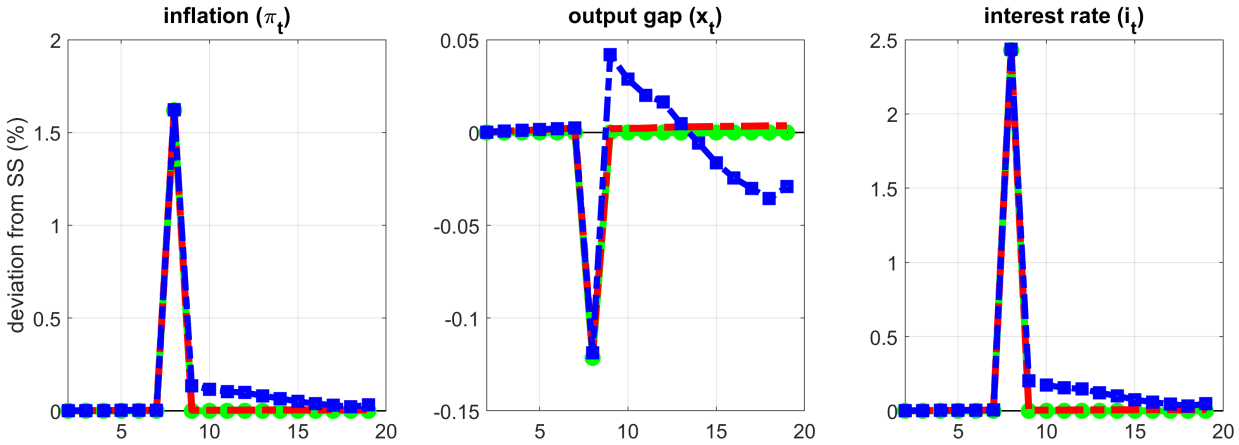


Figure A.9: Impulse response functions following a large supply shock when neither the central bank nor the private sector knows β_t

(a) Perceived shocks



(b) Macro variables



(c) Expectations

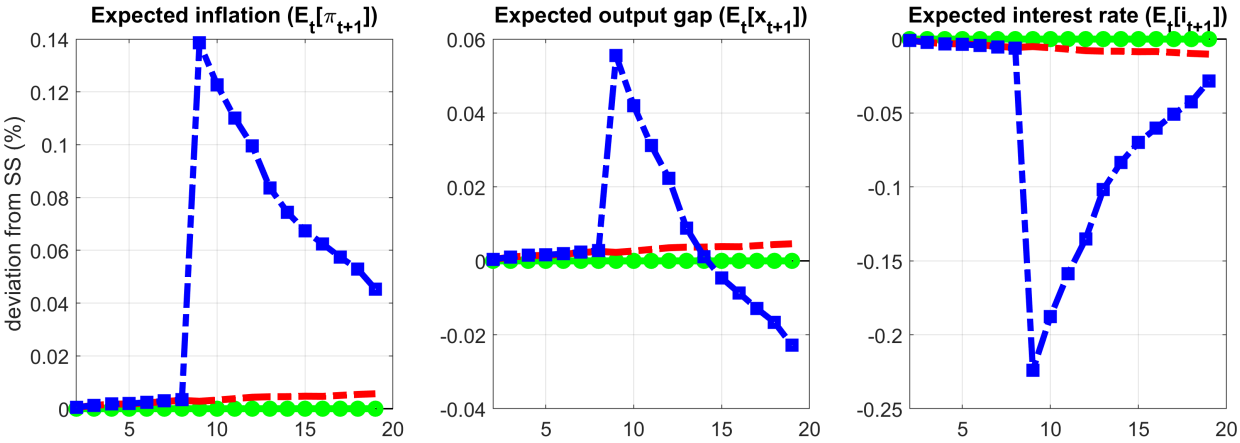


Figure A.10: Impulse response functions following supply and policy shocks

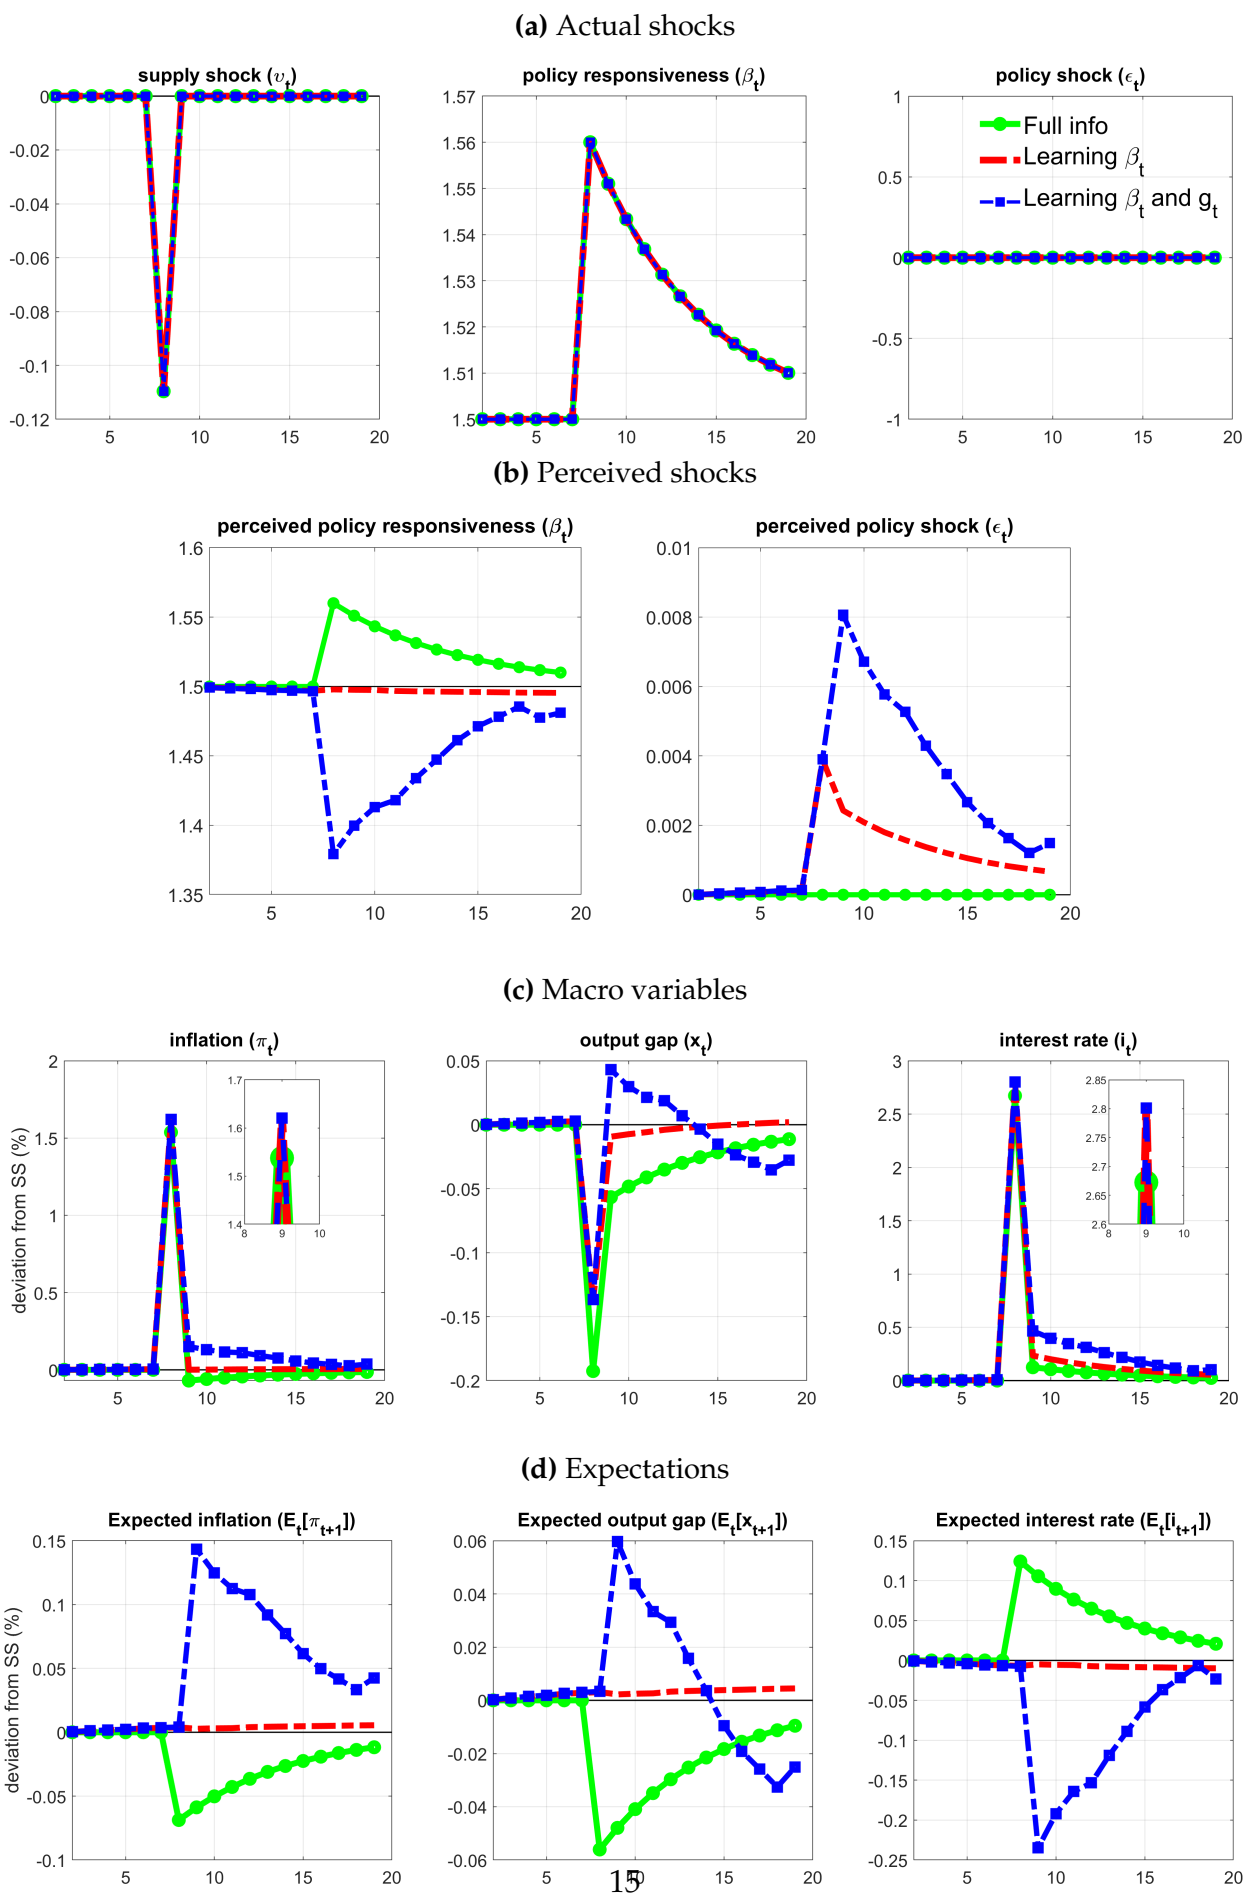
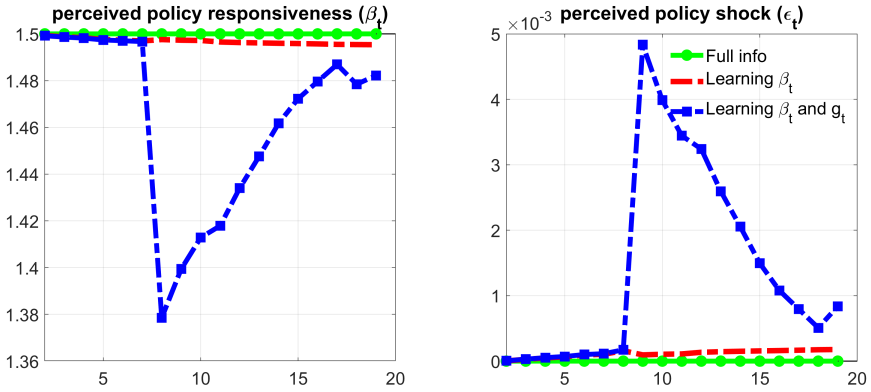
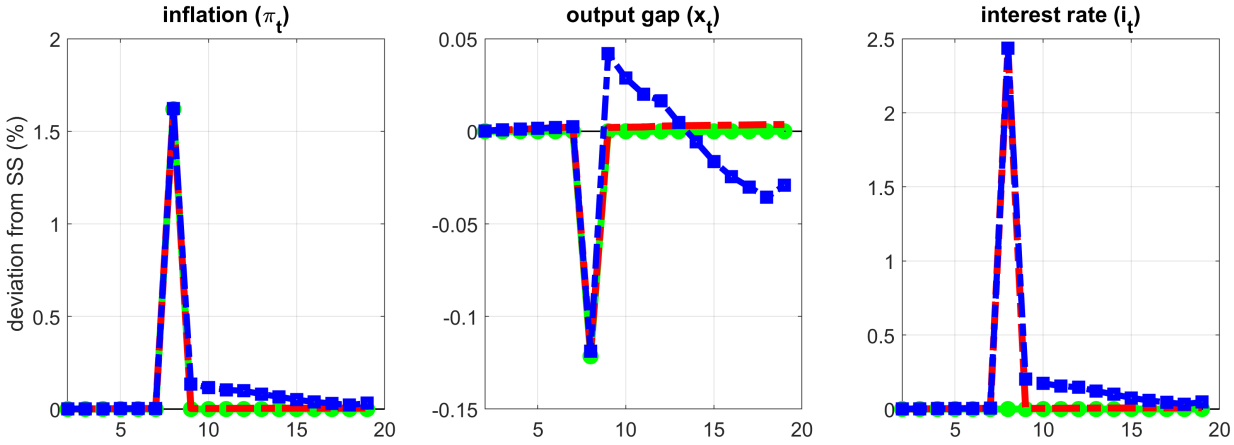


Figure A.11: Impulse response functions following a large supply shock when the private sector makes decisions with beliefs $(1-\rho)\bar{\beta}+\rho\widehat{\beta}_{t-1}$

(a) Perceived shocks



(b) Macro variables



(c) Expectations

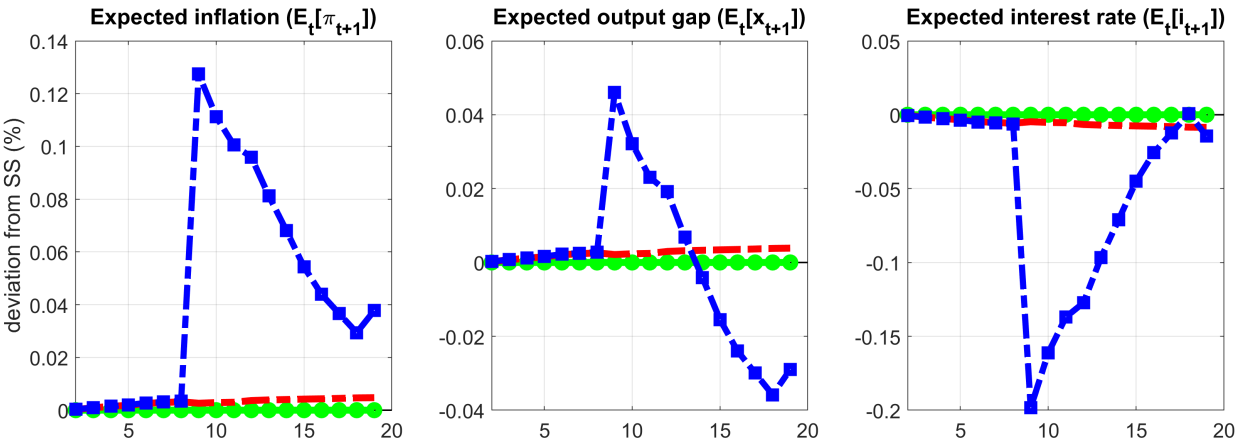


Figure A.12: Policy (in)effectiveness when agents do not observe β_t or g_t

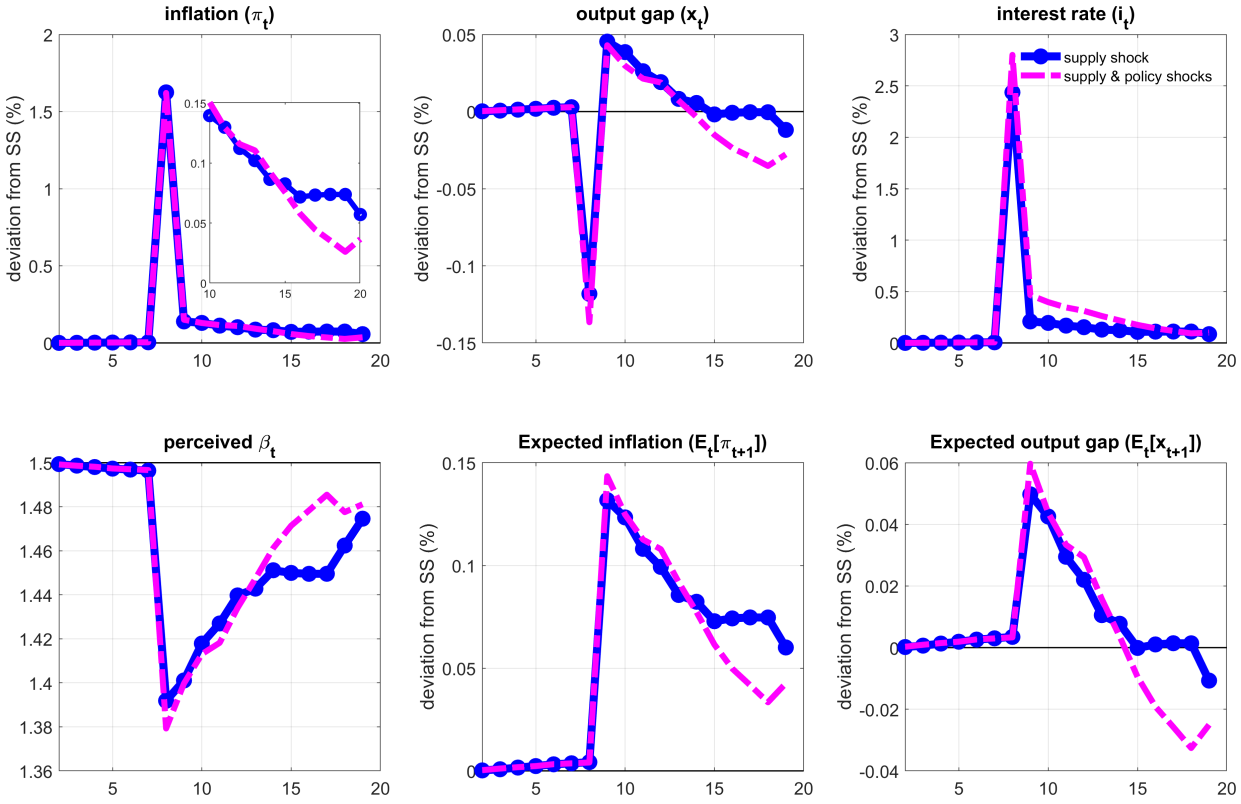


Figure A.13: Impulse response functions from nonlinear VAR (alternative specifications)

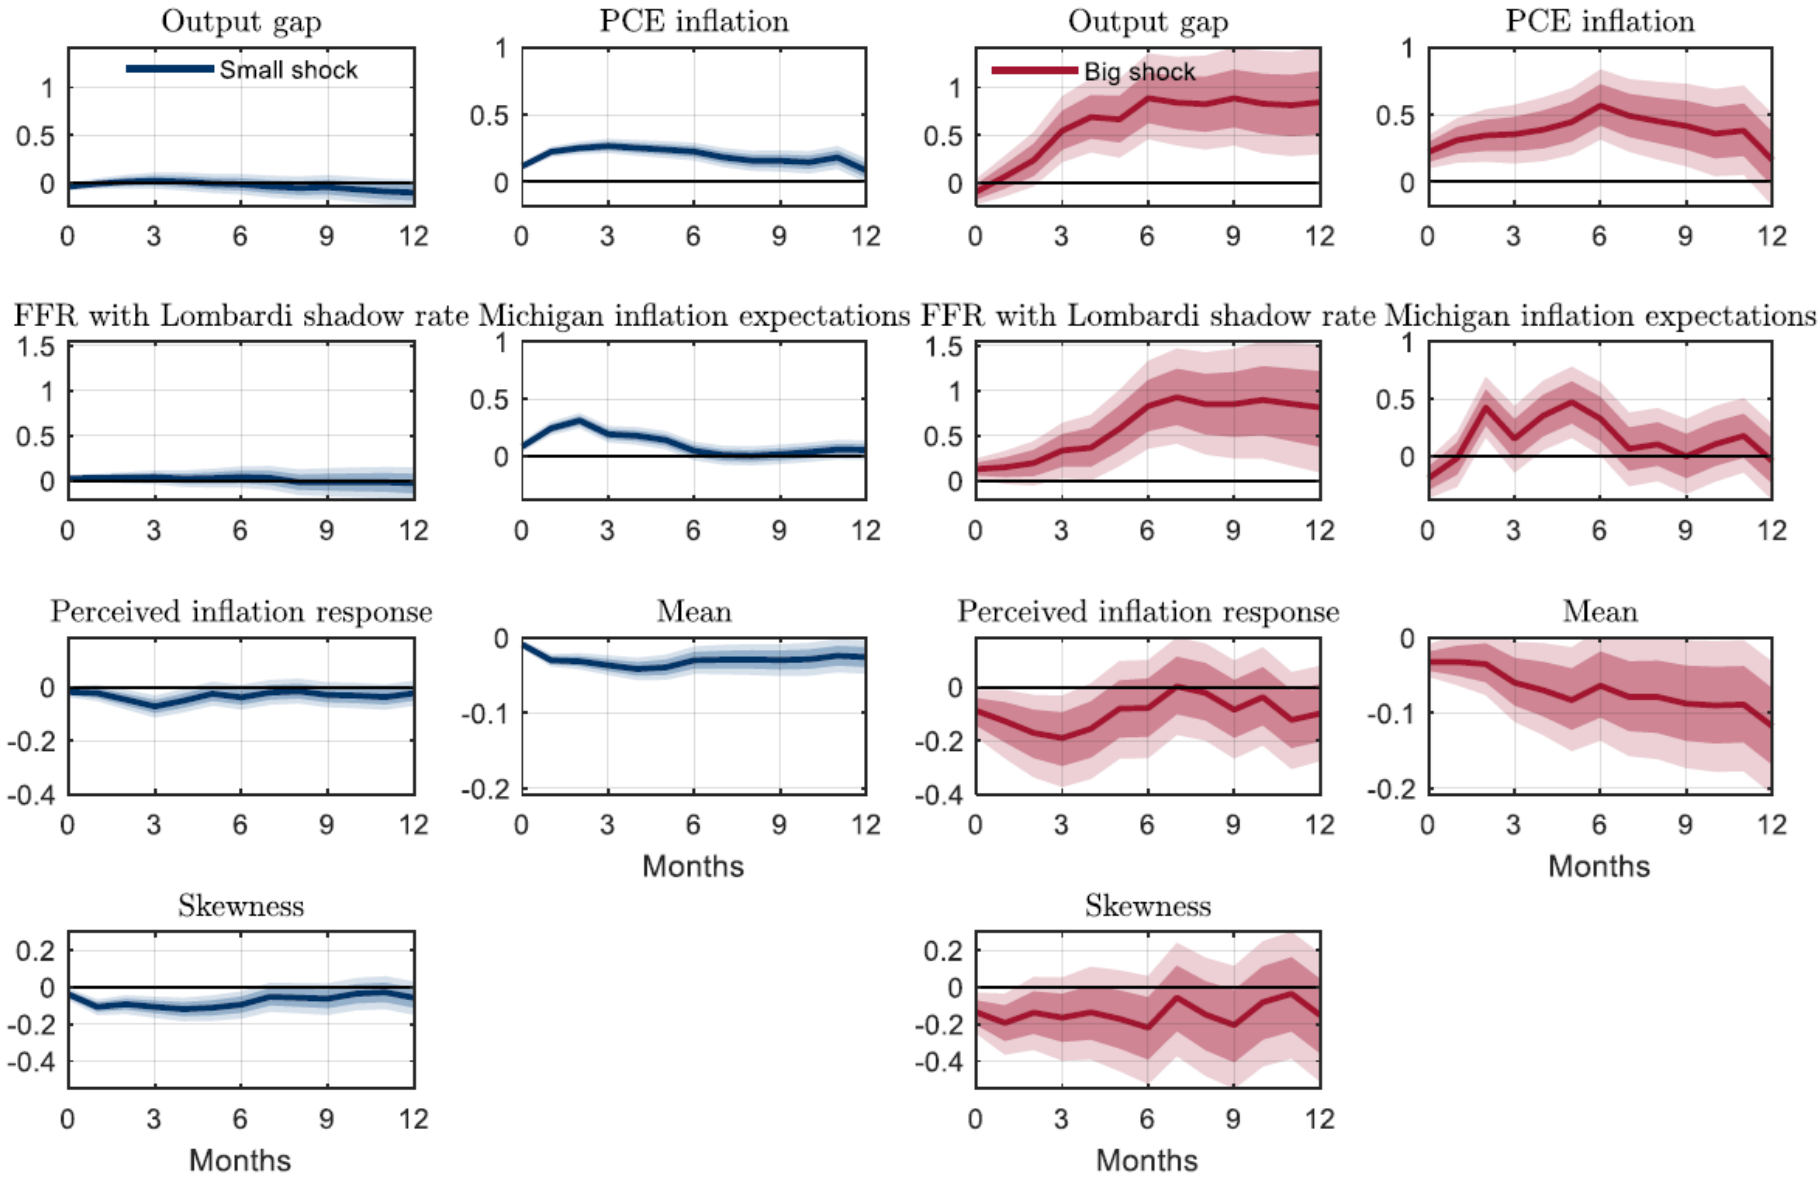


Figure A.14: Impulse response functions from nonlinear VAR (alternative specifications)

