Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Optimal Monetary Policy with Uncertain Private Sector Foresight

Christopher Gust and David Lopez-Salido

2024-059

Please cite this paper as:
Gust, Christopher, and David Lopez-Salido (2024). "Optimal Monetary Policy with Uncertain Private Sector Foresight," Finance and Economics Discussion Series 2024-059. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2024.059.

Optimal Monetary Policy with Uncertain Private Sector Foresight

Christopher Gust David López-Salido*

July 2024

Abstract

Central banks operate in a world in which there is substantial uncertainty regarding the transmission of its actions to the economy because of uncertainty regarding the formation of private-sector expectations. We model private sector expectations using a finite horizon planning framework: Households and firms have limited foresight when deciding spending, saving, and pricing decisions. In this setting, contrary to standard New Keynesian (NK) models, we show that "an inflation scares problem" for the central bank can arise where agents' longer-run inflation expectations deviate persistently from a central bank's inflation target. We formally characterize optimal time-consistent monetary policy when there is uncertainty about the planning horizons of private sector agents and a risk of inflation scares. We show how risk management considerations modify the optimal leaning-against-the-wind principle in the NK literature with a novel, additional preemptive motive to avert inflation scares. We quantify the importance of such risk management considerations during the recent post-pandemic inflation surge.

JEL Classification: C11, E52, E70
Keywords: Finite horizon planning, optimal time-consistent policy under uncertainty, leaning against the wind, attenuation principle.

*Federal Reserve Board. Gust: christopher.gust@frb.gov. López-Salido: david.lopez-salido@frb.gov. This paper presents our views and should not be interpreted as reflecting those of the Board or the Federal Reserve System.

1 Introduction

It has long been recognized that uncertainty is a pervasive feature affecting the design and conduct of monetary policy. Substantial research has been devoted to study how different forms of model uncertainty affect our understanding of the principles underlying the design of optimal monetary policy.[1] In this paper, we contribute to this literature by studying optimal time-consistent policy when policymakers are uncertain about the nature of expectations formation. We do this in the context of a microfounded model in which the cognitive ability of economic agents to solve complex infinite-horizon planning problems is limited. In particular, we use the New Keynesian, finite horizon planning (NK-FHP) framework developed in Woodford (2018) in which households and firms are boundedly rational because they have limited foresight: they use structural relationships to evaluate the full set of state-contingent paths along which the economy might evolve only up to a finite horizon.

An appealing feature of the NK-FHP model that we demonstrate in this paper is that it provides microfoundations for the "inflation scares" discussed in Goodfriend (1993) in which longer-term inflation expectations of the private sector can move persistently away from a central bank's inflation target. This feature makes the NK-FHP model suited to study the design of optimal policy when there is a risk of an inflation scare. To do so, we extend the NK-FHP framework of Woodford (2018) to an environment in which certainty equivalence no longer applies, because there is a distribution of agents with different planning horizons that changes over time. These fluctuations give rise to uncertainty that is non-additive and implies that optimal monetary policy depends not only on the means of output and inflation, but on the distribution of the output gap and inflation as well as agents' beliefs about future inflation.

There are additional benefits to studying optimal policy under uncertainty in a model in which agents have finite horizon planning. On the theoretical side, varying the foresight of households and firms allows us to flexibly approximate starkly different ways that the private sector agents may form expectations that are relevant for monetary policy. When agents make plans over very long horizons, expectations formation is rational, longer-run inflation expectations are well anchored at levels consistent with a central bank's inflation objective, disinflations are relatively costless, and the transmission of monetary policy occurs relatively quickly. In contrast, when agents have short planning horizons, agents are not fully rational: while they are sophisticated in thinking about events inside their planning horizon, they are less sophisticated in forming beliefs about events outside their planning horizons. In particular, they learn and update their longer-run beliefs averaging over past data that they have observed. This behavior gives rise to movements in trend inflation and output that reflect the private sector's longer-run beliefs (i.e., those pertaining to developments outside their planning horizon) and that change in response to realized data. With households and firms updating their longer-run beliefs based on past data, longer-run inflation expectations can become unanchored, disinflations can be costly, and the transmission lags of

[1] See Barlevy (2011) for a discussion of the Bayesian and robust control approaches to modeling uncertainty as well as a survey of the literature.

monetary policy can be long.

On the empirical side, the NK-FHP framework has been shown to be a fruitful way to model business cycle fluctuations in output, inflation, and interest rates as well as survey evidence on predictability of forecast errors. In particular, Gust, Herbst, and López-Salido (2022) show that the model fits the macroeconomic time series substantially better than other behavioral models as well as the "hybrid" NK model that features rational expectations, habit persistence in consumption, and exogenous price indexation. The model is also capable of generating substantial inflation persistence and realistic costs to an anticipated disinflation announced by a central bank. Moreover, because the model has been estimated, we can use its parameter estimates to help quantify how important limited and uncertain foresight is for the conduct of monetary policy. In addition, Gust, Herbst, and López-Salido (2024) extend this work and show–analytically and empirically–that such a model accounts for an initial underreaction and subsequent overreaction of inflation forecasts.

Our analysis also builds on the prominent work of Clarida, Gali, and Gertler (1999), who study optimal policy under discretion for a central bank with a dual mandate objective when private sector agents have rational expectations. They show that policy should "lean against the wind" by contracting aggregate demand whenever inflation is above a central bank's objective. This "lean against the wind policy" implies a short run tradeoff between inflation and output variability in an environment where longer-run inflation expectations are well anchored (at a level consistent with a central bank's target). This anchoring of inflation expectations in part reflects the rationality of private sector agents who fully understand the implications of how optimal policy under discretion acts to ensure that inflation converges to a central bank's objective.

We show that when households and firms planning horizons are long enough, optimal policy in the NK-FHP model is equivalent to the "leaning against the wind" strategy discussed in Clarida, Gali, and Gertler (1999). But, when agents have short planning horizons, agents' beliefs about longer-run inflation can move persistently away from a central bank's inflation target, and a policymaker acting under discretion follows a modified "lean against the wind" strategy that involves a forward-looking, anticipatory response to inflation. This anticipatory response reflects that a central bank realizes that inflationary pressure will boost private sector's future beliefs about trend inflation, leading to long-lasting departures of inflation from a central bank's objective. As a result, a central bank has a strong desire to act aggressively and preemptively to keep inflation close to target, which in turn helps anchor private-sector beliefs about trend inflation at their objective.

When policymakers are uncertain about the share of agents with different horizons, they are also uncertain about how agents' beliefs regarding longer-run inflation might evolve. We find that optimal, time-consistent, policy under uncertainty is such that the central bank acts more aggressively than when policymaker has perfect certainty about the nature of private-sector expectations formation. This results stands in contrast to the classic result in Brainard (1967), who showed that uncertainty about the effect of policy on the economy implies that policy should attenuate its response to shocks relative to the certainty-equivalent case. This more aggressive response reflects that uncertainty about expectations formation increases the likelihood that households' and firms' beliefs about longer-run inflation may move away from a central bank's target, leading to persistent departures of inflation from its objective. To avoid such undesirable outcome the central bank sets a relatively more aggressive policy path than under certainty.

Despite this more aggressive response, we show that an environment where inflation has been running above a central bank's objective, uncertainty about private-sector's forecast leads to a much wider dispersion of outcomes for inflation and the output gap than when expectations formation is known with certainty. In particular, with uncertain expectations formation, the distribution for inflation displays a long tail of above-target inflation and the distribution for output displays a long tail of below-potential activity. These less favorable outcomes under uncertainty reflect that the policy tradeoff between inflation and output gap stabilization worsens notably relative to the case in which the central bank is certain about how expectations are formed. We illustrate this feature by showing that the inflation-output variance frontier (Taylor (1999)) shifts as a result of uncertainty regarding agents' planning horizons.

We use the empirical estimates of the NK-FHP model of Gust, Herbst, and López-Salido (2024) to quantify the gains to incorporating risk-management considerations into the conduct of monetary policy when policymakers are uncertain about the foresight of the private sector. To evaluate these gains in an environment in which there is a heightened risk that above-target inflation could lead to an unanchoring of private-sector's longer-run inflation expectations, we use information from the FOMC's Survey of Economic Projections (SEP) in March of 2023. We use this information to construct a baseline scenario to compare optimal policy under uncertainty with alternative strategies, including one that would be optimal under certainty equivalence. We find that optimal policy rate response involves a considerably more aggressive response to the above-target inflation projected in the March 2023 SEP than the optimal response under certainty equivalence. Moreover, the gains associated with taking a risk management approach can be sizable in an environment where inflation has been running above a central bank's target and influences the private sector's beliefs about trend inflation.

Literature review. Building on Clarida, Gali, and Gertler (1999), this paper bridges two strands of the literature on optimal monetary policy. The first analyzes optimal monetary policy when expectations formation is imperfect and includes Woodford and Xie (2022), who study the coordination of optimal monetary and fiscal policy under commitment at the zero lower bound when private sector agents have finite planning horizons, but their beliefs about events outside their planning horizon are fixed. Relative to their work, we study optimal monetary policy under discretion when agents have finite planning horizons and agents learn and update their beliefs about events outside their planning horizons. In addition, our focus is on optimal monetary policy when there is uncertainty regarding the formation of private-sector expectations. Our paper is also related to papers studying optimal policy when agents are learning including Gáti (2023) and Molnár and Santoro (2014).[2] Like our paper, these papers emphasize that private-sector inflation expectations

[2]For a survey of the literature on optimal monetary policy when the private sector agents have imperfect expectations, see Eusepi and Preston (2018).

are more important when agents are learning and that there is an increased role for stabilizing inflation.[3] Our approach is distinct from these papers since we emphasize the role of uncertainty and in particular study the design of optimal policy when the central bank faces uncertainty about agents' planning horizons. Moreover, unlike in these papers, private-sector agents still take into account structural relationships over their finite-planning horizons, and thus announcements about future monetary policy still affect economic outcomes.

Our paper is also related to the literature studying monetary policy under uncertainty and is most closely related to those papers emphasizing uncertainty about inflation inertia using a Bayesian approach. Söderström (2002), Kimura and Kurozumi (2007), and Svensson and Williams (2008) incorporate inertia into the inflation process and show that the optimal policy response is not attenuated as in Brainard (1967) but is more aggressive than is the case under certainty equivalence. Our model differs from these earlier papers, as we explicitly model uncertainty about expectations formation and the unanchoring of agents' longer-run inflation beliefs as well as investigate the mechanism quantitatively during the high inflation that occurred in the United States in the aftermath of the pandemic.[4]

The rest of this paper proceeds as follow. The next section discusses our NK-FHP model including how we model the uncertainty a central bank faces about agents' planning horizons. We then discuss optimal discretionary policy when private sector agents have finite planning horizons. The fourth section discusses the results and the final section offers some conclusions and directions for future work.

2 Optimal Policy in a Finite-Horizon Planning Model

We use the NK-FHP model of Woodford (2018) to study optimal monetary policy when the central bank is uncertain about expectations formation and more specifically the planning horizon of agents in the model. To model uncertainty, we follow Svensson and Williams (2005), who model uncertainty using different "modes" or regimes that follow a Markov process. For the different modes, we assume that there are two types of agents that populate the economy: those with long planning horizons whose expectations take into structural relationships far into the future, and those with short planning horizons whose expectations about events beyond their planning horizons depend on past data. The distribution of agents is time-varying and governed by a Markov process. While a central bank observes the current distribution of agents, it is uncertain about this distribution in

[3]The dynamic target criteria implied by optimal policy in the FHP model is distinct from those implied by the adaptive learning models of Gáti (2023) and Molnár and Santoro (2014). In those models, the central bank adjusts the static targeting rule by responding to terms that reflect the expected discounted value of the future path of output gaps. In our model, the central bank adjusts the static targeting rule by responding to a term involving the expected discounted value of the future path of inflation gaps, as deviations of inflation from a central bank's target can lead to an undesirable drift in agent's longer-run inflation beliefs.

[4]Similar to our paper, Kimura and Kurozumi (2007) also provide microfoundations for inertial inflation dynamics. In their case, it arises because a subset of firms do not optimally choose their prices but set their prices based on lagged inflation. In contrast, the firms in our model choose their prices optimally but are boundedly rational because of their finite planning horizons and the learning they do about events outside of their planning horizons.

the future.

The central bank chooses the interest rate to minimize expected discounted losses consisting of squared deviations of inflation from the central bank's inflation objective and squared deviations of the output gap. The monetary authority acts in a time-consistent fashion taking as given the equilibrium conditions of the private sector. While private-sector agents have a finite horizon and thus are boundedly rational, the central bank is assumed to have rational expectations and thus uses its knowledge of the economy and sources of uncertainty to make interest-rate decisions. A key focus of our analysis is to compare the decisions of a central bank acting under uncertainty about expectations formation to a central bank who does not face this uncertainty. This allows us to isolate the effects of this particular source of uncertainty on economic outcomes.

2.1 Heterogeneous Planning

In this subsection we provide an abbreviated discussion of the NK-FHP model that we study, highlighting the source of model uncertainty that a central bank faces as well as the equilibrium beliefs of FHP households and firms. For a more detailed discussion of the model, see Woodford (2018).

The economy is populated by two groups of households and firms that differ in their planning horizons, $k \in \{k_0, k_1\}$. For the group of households and firms with planning horizon k, they make decisions at time t based on formulating state-contingent plans through period $t+k$. Within their planning horizon, they use the full knowledge of the model to formulate those plans except that they form beliefs about the aggregate variables assuming that all other agents have the same k-period planning horizon as themselves. Under these assumptions, Woodford (2018) shows that the time t beliefs of household and firm over their generic k period planning horizon satisfy:

$$\pi_\tau^j = \beta E_\tau \pi_{\tau+1}^{j-1} + \kappa y_\tau^j + u_\tau \tag{1}$$

$$y_\tau^j = E_\tau y_{\tau+1}^{j-1} - \sigma \left(i_\tau^j - E_\tau \pi_{\tau+1}^{j-1} - r_\tau^e \right), \tag{2}$$

for $\tau = t+k-j$ and $j > 0$ where j indexes the number of periods left in an agents' planning horizon and k denotes an agent's planning horizon. The variables π_τ^j and y_τ^j represent agents' beliefs for aggregate inflation and spending in period τ where both variables are expressed in log-deviation from steady state. We assume that the central bank's inflation target is fixed at the steady-state inflation rate, and we abstract from technology or other shocks that move the level of potential output so that it is constant and it corresponds to the steady state level of output around which the economy fluctuates. The variable, i_τ^j, (which is also in log-deviation from steady state) corresponds to agents' beliefs regarding the setting of the policy rate.

The expectations operator, E_t, in equations (1) and (2) denotes the model-consistent operator. However, the variables with the j superscripts reflect the subjective expectations of agents with finite planning horizons of length k. In particular, as shown in Woodford (2018), we can define the subjective expectations operator, $\mathbb{E}_t^k$ such that for variable Z_{t+k-j} with $k > j \geq 0$, the following relationship holds:

$$\mathbb{E}_t^k Z_{t+k-j} = E_t Z_{t+k-j}^j. \tag{3}$$

Expression (3) provides a mapping between the subjective expectations operator of an agent with a k-period planning horizon and the model-consistent expectations operator. The variables indexed by j reflect the subjective expectations of an agent with a k-period horizon, which depart from rational expectations in two ways. First, they only formulate fully-state contingent plans for k periods rather than their infinite lifetimes. Second, they view all other agents as having the same k-period horizon as themselves.

The variables u_τ and r_τ^e represent exogenous aggregate shocks to firm's pricing and household's spending decisions, respectively.[5] These shocks are assumed to follow AR(1) processes:

$$u_t = \rho_u u_{t-1} + e_{ut} \tag{4}$$

$$r_t^e = \rho_r r_{t-1}^e + e_{rt}, \tag{5}$$

with the parameters ρ_z and ρ_r, between zero and one, measuring the persistence of the shocks. The innovations to these shocks are assumed to be *iid* normal random variables with standard deviations, σ_u and σ_r, respectively.

Equation (1) reflects the beliefs of the firms who have the opportunity to change their prices at date t. Likewise, equation (2) represents household beliefs about their spending decisions. These equations can be iterated forward to show that the spending and pricing decisions at time t of households and firms with a k-period horizon satisfy:

$$\pi_t^k = E_t \sum_{i=0}^{k-1} \beta^i [\kappa y_{t+i}^{k-i} + u_{t+i}] + \beta^k E_t \pi_{t+k}^0 \tag{6}$$

$$y_t^k = -\sigma E_t \sum_{i=0}^{k-1} [i_{t+i}^{k-i} - \pi_{t+i+1}^{k-i} - r_{t+i}^e] + E_t y_{t+k}^0, \tag{7}$$

As $k \to \infty$, it can be shown that the 'canonical' NK model used to study optimal policy in Clarida, Gali, and Gertler (1999) is a special case of equations (6) and (7). In that case, households' and firms' planning horizon extends over their infinite lifetimes, and they use the model's structural relationships over their infinite lifetimes to make their spending and pricing decisions.

2.2 Longer-Run Learning

For households and firms with a finite horizon k, equations (6) and (7) indicate that their spending and pricing decision depend on their beliefs at the end of their planning horizons. As shown in

[5] The shock, r_t^e, affects a household's discount factor and is a departure from the preference shock used in Woodford (2018). The appendix in Gust et al. (2024) shows how to derive the log-linearized equilibrium conditions shown here from an FHP household's optimization problem in the presence of these shocks.

Woodford (2018), these beliefs satisfy:

$$\pi^0_{t+k} = \kappa y^0_{t+k} + \beta v_{pt} + u_{t+k} \tag{8}$$

$$y^0_{t+k} = -\sigma(i^0_{t+k} - r_{t+k}) + v_{ht}, \tag{9}$$

where v_{pt} and v_{ht} represents the continuation value functions of the economy's firms and households, respectively. These value functions reflect the beliefs of households and firms for events outside of their planning horizon, as they have infinite lifetimes and assign continuation values to events outside of their planning horizons. While they make their time t decisions taking v_{pt} and v_{ht} as fixed at date t, households and firms, as discussed in Woodford (2018) update these continuation values over time as part of their optimization problem.[6] Specifically, households and firms learn and update v_{pt} and v_{ht} according to:

$$v_{pt+1} = (1 - \gamma_p)v_{pt} + \gamma_p \pi_t \tag{10}$$

$$v_{ht+1} = (1 - \gamma_h)v_{ht} + \gamma_h(y_t + \sigma\pi_t), \tag{11}$$

where, π_t and y_t denote aggregate inflation and output (in log deviation from steady state), respectively; and γ_p and γ_h are parameters that determine how quickly firms and households update their beliefs in response to recent aggregate data.[7] In short, households and firms make sophisticated plans and forecasts within their planning horizons, using their full knowledge of the model's structural equations. However, for longer-run events (i.e., those outside of their planning horizons), households and firms are less sophisticated, updating their beliefs based on past economic outcomes.

While longer-run learning introduces two extra parameters, γ_p and γ_h, it has both theoretical and empirical benefits. On the theoretical side, it gives the FHP model attractive properties in response to long-lasting changes in policy or economic fundamentals. Without learning, v_{pt} and v_{ht} are fixed. Accordingly, these values do not change in response to long-lasting economic events and households and firms will continue to use outdated value functions if, for example, there is a permanent change in a central bank's inflation target. In contrast, with learning, the value functions of households and firms will change and eventually fully reflect the change in a central bank's inflation target. On the empirical side, Gust et al. (2022) show that the longer-run learning allows the FHP model to generate substantial aggregate persistence and fit macroeconomic data without resorting to additional features such as habit persistence or price contracts indexed to lagged inflation.

Inflation Scares. Firms' beliefs about events outside its planning horizon play a particularly important role in our analysis. In particular, v_{pt} can be interpreted as a firm's longer-run beliefs

[6] In making their time t decisions for prices and spending, households and firms ignore that their value functions change over time and thus the learning framework we adopt uses an 'anticipated utility' approach (e.g., Cogley and Sargent (2008)).

[7] For convenience, we have rescaled a firm's value function by the probability that the firm can re-optimize their price. Thus, relative to Woodford (2018), $v_{pt} = (1 - \theta_p)\widetilde{v}_t$, where $1 - \theta_p$ is the probability that a FHP firm has the opportunity to re-optimize its price and $\widetilde{v}_t$ is the continuation value at date t of such a firm.

about inflation, as a firm updates their value functions v_{pt} based on past inflation. Thus, agents' longer-run beliefs about inflation evolves slowly, as equation (10) implies:

$$v_{pt} = \gamma_p \sum_{i=0}^{t-1} (1-\gamma_p)^i \pi_{t-1-i}, \tag{12}$$

where the parameter γ_p determines the importance of recent lags of inflation relative to distant lags. Because v_{pt} depends on past deviations of inflation from the central bank's target, agent's longer-run inflation beliefs will evolve endogenously in response to the economy's (persistent) shocks and can potentially drift away from zero—the level consistent with their being anchored at a central bank's inflation target. Such adverse movements in longer-run inflation expectations are closely related to the "inflation scares" defined by Goodfriend (1993) and discussed in Orphanides and Williams (2005) and Orphanides and Williams (2022). Accordingly, the FHP model with longer-run learning can be viewed as providing microfoundations for the notion of inflation scares emphasized in the literature. Later, we discuss how optimal (time-consistent) policy is affected by the presence of inflation scares.

2.3 Model Uncertainty

We use the heterogeneity of households and firms as the key source of uncertainty faced by the central bank. Equations (6) and (7) describe the decisions for inflation and spending decision for two groups of households and firms, as $k \in \{k_0, k_1\}$. A central bank faces uncertainty about the fraction of agents with different planning horizons. Specifically, we denote $\omega_t \equiv \omega(m_t)$ as the share of households and firms with horizon k_0 at date t. The variable m_t is random with $m_t \in \{0, 1\}$. These values correspond to the different modes, and the modes follow a Markov process with constant transition probabilities:

$$P_{mn} = \Pr\{m_{t+1} = n | m_t = m\}, \qquad m, n = 0, 1 \tag{13}$$

with the matrix P denoting the 2×2 matrix $[P_{mn}]$. We assume there is a unique stationary distribution of nodes given by $\bar{p} = \bar{p}P$, where $\bar{p}$ is a row vector consisting of the ergodic probabilities.

Aggregate inflation and aggregate spending reflect the distribution of agents at time t:

$$\pi_t = \omega_t \pi_t^{k_0} + (1-\omega_t)\pi_t^{k_1} \tag{14}$$

$$y_t = \omega_t y_t^{k_0} + (1-\omega_t) y_t^{k_1}, \tag{15}$$

Accordingly, aggregate inflation and output reflect uncertainty arising from fluctuations in the length of agents' planning horizons. This form of uncertainty moves us beyond a linear framework with additive shocks so that certainty equivalence no longer holds. However, the analysis remains tractable since conditional on m_t, the model equations are linear.

3 Optimal Monetary Policy

A central bank minimizes an intertemporal loss function involving the squared deviations of inflation and the output from their respective targets:

$$\mathcal{L}_t = \frac{1}{2} E_t \sum_{\tau=t}^{\infty} \beta^{\tau-t} \left[\pi_\tau^2 + \lambda y_\tau^2 \right], \tag{16}$$

The parameter λ is the central bank's relative weight on fluctuations in the output gap.[8] Unlike households and firms, a central bank has rational expectations and formulates policy over an infinite horizon using its full knowledge of private-sector behavior. We also assume that policymakers can not commit and reoptimizes each period in a time-consistent fashion.

3.1 Private Sector (Policy) Beliefs

Before discussing the central bank's problem, it is useful to discuss the beliefs of households and firms about future monetary policy. Because households and firms can not formulate plans through an infinite horizon, their perceptions about future monetary policy differ from the policy that the central bank can implement. In particular, they perceive that in formulating policy the central bank has a k-period planning horizon like themselves. This impacts their future beliefs about monetary policy. For current policy, household and firm simply observe the policy decision at date t so $i_t^k = i_t$ for $k = k_0, k_1$. For their date t forecasts of policy over the remainder of their k-period horizon, they perceive that the central bank chooses i_τ^j for $j = 1, 2, ..., k-1$ and $\tau = t + k - j$ to satisfy a sequence of minimization problems:

$$\mathcal{W}_\tau^j(s_\tau) = \min_{i_\tau^j} \left[(\pi_\tau^j)^2 + \lambda (y_\tau^j)^2 \right] + \beta E_\tau \left[\mathcal{W}_{\tau+1}^{j-1}(s_{\tau+1}) \right] \tag{17}$$

subject to equations (1) and (2). In equation (17), $s_t = (u_t, r_t^e)$ denotes the vector of shocks. The private-sector agents' perceptions of the central bank's problem reflect that they believe the central banks has the same subjective expectations as themselves and work through future state-contingencies in a model-consistent fashion only through a finite k-period horizon.

At the last period of their planning horizon ($j = 0$), they perceive that the central bank minimizes:

$$\mathcal{W}_{t+k}^0(s_{t+k}) = \min_{i_{t+k}^0} \left[(\pi_{t+k}^0)^2 + \lambda (y_{t+k}^0)^2 \right] \tag{18}$$

subject to equations (1) and (2), taking as given v_{pt} and v_{ht}. Households and firms perceive that the central bank, like they do, takes the continuation value functions as fixed at date t. When agents have a finite planning horizon, this perception as we discuss below is incorrect, because the central

[8] We assume that inflation in the model's non-stochastic steady state is equal to a central bank's inflation target; hence, we use deviation from target and deviation from steady state, interchangeably. Similarly, the aggregate supply shock, u_t is assumed not to affect the level of potential output; hence, we use deviations of output from steady state and the output gap, interchangeeably.

bank has an infinite horizon and takes into account the evolution of the value functions. In the special case in which agents have an infinite planning horizon ($k \to \infty$), agents policy perceptions will be correct, and there will be no difference between actual and perceived policy.

Iterating backwards from the problem at the end of agents' planning horizon, the first order conditions from the sequence of problems imply that private-sector agents perceive that the central bank follows the targeting rule:

$$y^j_{t+k-j} = -\frac{\kappa}{\lambda}\pi^j_{t+k-j}. \tag{19}$$

for $j = 0, 1, ...k-1$.

Agents' Perceived LAW. The agents' perceived targeting rule is similar to the targeting rule under optimal discretion in the canonical NK model (e.g., Clarida, Gali, and Gertler (1999)), which features full stabilization of demand shocks and partial accommodation of supply shocks. In particular, this condition implies that agents believe that the central bank would pursue a *lean against the wind* (LAW) policy at each future date *of their planning horizon*: Whenever inflation is above target, they expect the central bank to contract demand below capacity (by raising the interest rate); and vice-versa when it is below target. The central bank's perceived aggressiveness in reducing y^j_τ to combat high inflation depends positively on the gain in reduced inflation per unit of output loss, κ, and inversely on the relative weight placed on output losses, λ.

Sticky Expectations of Inflation. To describe the actual policy that a central bank implements, it is useful first to characterize the beliefs of agents with horizon k about inflation next period. As shown in the appendix, combining agents' perceived targeting rule with equations (1) and (8) implies that their expected inflation is given by:

$$E_t\Pi^{k-1}_{t+1}(v_{pt}, u_t) \equiv \underbrace{\left[\frac{\lambda\rho_u}{\lambda+\kappa^2}\sum_{i=0}^{k-1}\left(\frac{\beta\lambda\rho_u}{\lambda+\kappa^2}\right)^i\right] u_t}_{\text{Cyclical}} + \underbrace{a_p(k)v_{pt}}_{\text{Trend}} \tag{20}$$

where the use of $\Pi^{k-1}_{t+1}(v_{pt}, u_t)$ instead of π^{k-1}_{t+1} is done to emphasize that this expected inflation is function of aggregate shocks as well as a firm's value function. The coefficient $a_p(k) = \left(\frac{\beta\lambda}{\lambda+\kappa^2}\right)^k$ determines the sensitivity of expected inflation next period to changes in agents' longer-run beliefs about inflation.[9] Intuitively, optimal policy is perceived as fully offsetting the effects on inflation of changes in aggregate demand that occur through movements in r^e_t and v_{ht}. Hence, expected inflation does not depend on r^e_t or v_{ht} but does depend on the cost-push shock and price-setting firms' continuation value function, v_{pt}.

[9] We can characterize agents' beliefs regarding expected output next period. As shown in the appendix, agents with a k-period horizon believe that the future output gap is given by:

$$E_t\mathbf{Y}^{k-1}_{t+1}(v_{pt}, u_t) \equiv -\frac{\kappa}{\lambda}\left\{a_p(k)v_{pt} + \left[\frac{\lambda\rho_u}{\lambda+\kappa^2}\sum_{i=0}^{k-1}\left(\frac{\beta\lambda\rho_u}{\lambda+\kappa^2}\right)^i\right]u_t\right\}$$

.

Following Woodford (2018), we decompose variables into a cyclical component that reflects movements in shocks and a trend component that reflects movements in agents' longer-run beliefs. The cyclical response of expected inflation to the cost-push shock depends on a central bank preference for stabilization of the output gap (λ), the slope of the Phillips curve (κ), the persistence of supply shocks (ρ_u), and the length of agents' planning horizons, k. If supply shocks are *iid* ($\rho_u = 0$) or a central bank is perceived as stabilizing only deviations of inflation from target ($\lambda = 0$), then expected inflation is unaffected by the cost-push shock, u_t. A more steeply sloped Phillips curve or a shorter planning horizon reduces the impact of cost-push shocks on expected inflation.

The trend component of expected inflation reflects movements in firm's continuation value function or their longer-run inflationary beliefs. As discussed earlier, these longer-run beliefs are backward-looking, and they make expected inflation sticky. In addition, inflation persistently above a central bank's target can lead to an inflation scare, since it can result in a large value of v_{pt} that pushes expected inflation away from a central bank's target. The extent to which this occurs depends on $a_p(k)$, which determines the marginal effect of change in v_{pt} on expected inflation. If a policymaker is perceived to be more aggressive towards inflation (i.e., a lower value of λ), this marginal effect is smaller, implying less feedback from agents' longer-run beliefs about inflation into expected inflation next period. With reduced feedback, the possibility of an inflation scare is also lower. A longer planning horizon also reduces this marginal effect: As $k \to \infty$, household and firms take into account the effects of the supply shock over their infinite lifetimes, and their longer-run beliefs (v_{pt}) become irrelevant. In that case, inflation scares are no longer possible. Finally, note that, if the inflation-output trade off gets increasingly small (i.e., κ decreases toward 0), $a_p(k)$ approaches β^k and becomes independent of the central bank's preference parameter (λ).

3.2 Optimal Targeting Rule

As discussed above, we assume that the central bank has rational expectations and optimizes policy under discretion. While the central bank does not commit to future actions, it understands how agents' longer-run beliefs about inflation, v_{pt}, depend on past inflation and takes that into account in choosing current policy. This rationality of the central bank creates a key difference between the optimal policy that is actually *implemented* and private-sector's beliefs regarding such policy (i.e., agents' perceptions of optimal time-consistent policy). In particular, as we show, the optimal time-consistent policy actually implemented has a forward-looking component that reflects that agents' expectations of longer-run inflation depend on past inflation. This forward-looking component reflects that optimal policy seeks to avoid the possibility of an inflation scare where longer-run inflation expectations drift away from a central bank's inflation target.

We now proceed to formalize these ideas. The central bank's problem at date t can be written as:

$$W(v_{pt}, s_t, m_t) = \min_{i_t} \frac{1}{2}\left[\Pi_t^2 + \lambda \mathbf{Y}_t^2\right] + \beta \sum_{n=0}^{1} \Pr(m_{t+1} = n|m_t) \int_{s_{t+1}} W(v_{pt+1}, s_{t+1}, n) f(s_{t+1}|s_t) ds_{t+1} \tag{21}$$

where the function $f(s_{t+1}|s_t)$ denotes the conditional density for the shocks to aggregate supply and demand. In addition, $\Pr(m_{t+1} = n|m_t)$ denotes the conditional probabilities in the transition matrix, P and $\Pi_t \equiv \pi(v_{pt}, s_t, m_t; i_t)$ and $\mathbf{Y}_t \equiv y(v_{pt}, s_t, m_t; i_t)$ denote functions that determine the deviations of aggregate inflation and output from their targets, respectively. The central bank chooses the policy rate, i_t, taking as given the private sector's equilibrium conditions and the functions determining agent's beliefs regarding future inflation and output. These equilibrium conditions can be written as:

$$\Pi_t = \beta E_t \Pi_{t+1} + \kappa \mathbf{Y}_t + u_t \tag{22}$$

$$\mathbf{Y}_t = E_t \mathbf{Y}_{t+1} - \sigma\left(i_t - E_t \Pi_{t+1} - r_t^e\right) \tag{23}$$

$$v_{pt+1} = (1-\gamma_p) v_{pt} + \gamma_p \Pi_t$$

where $\omega_t = \omega(m_t)$. Expressions (22) and (23) resemble the expressions that determine inflation and the output gap in the canonical NK model except that the functions, $E_t\Pi_{t+1}$, and $E_t\mathbf{Y}_{t+1}$ reflect the finite planning horizons of agents in the model. These expectations are population-weighted averages of agents' expectations for inflation and the output gap, respectively:

$$E_t\Pi_{t+1} = \omega_t E_t \Pi_{t+1}^{k_0}(v_{pt}, u_t) + (1-\omega_t) E_t \Pi_{t+1}^{k_1}(v_{pt}, u_t)$$
$$E_t\mathbf{Y}_{t+1} = \omega_t E_t \mathbf{Y}_{t+1}^{k_0}(v_{pt}, u_t) + (1-\omega_t) E_t \mathbf{Y}_{t+1}^{k_1}(v_{pt}, u_t)$$

To set the optimal time-consistent policy, the central bank takes the functions, $E_t\Pi_{t+1}^{k-1}$ and $E_t\mathbf{Y}_{t+1}^{k-1}$ as given for $k \in \{k_0, k_1\}$. However, a few additional comments are in order. First, the central bank takes into account that its interest-rate decision affects agents' beliefs indirectly since these expectational functions depend on v_{pt}, which in turn depends on past (aggregate) inflation. Accordingly, in setting the current policy rate, i_t, the central banks takes into account that v_{pt+1} depends on current inflation. Second, in choosing i_t, the central bank knows the current values of the aggregate demand and supply shocks as well as ω_t, the current distribution of households and firms. However, the central bank's problem is dynamic and the central bank does not know the distribution of agents in the future and uses the probabilities of the modes to weigh future losses. Third, unlike the other macroeconomic shocks, fluctuations in ω_t are not additive: fluctuations in ω_t interact multiplicatively with the model's endogenous variables. Accordingly, certainty equivalence does not hold. Still, the central bank's problem remains linear-quadratic *conditional on* ω_t, keeping the treatment of model uncertainty relatively tractable following the approach of Svensson and Williams (2005).

As shown in the appendix, the optimal targeting rule is given by:

$$\pi_t + \gamma_p \beta E_t W_{pt+1} = -\frac{\lambda}{\kappa} y_t \tag{24}$$

$$W_{pt} = \beta a_{pt} \pi_t + \beta \left[1 - \gamma_p (1 - \beta a_{pt})\right] E_t W_{pt+1} \tag{25}$$

where aggregate inflation (π_t) and the output gap (y_t) are given by expressions (14) and (15), respectively. The term $a_{pt} = \omega_t a_p(k_0) + (1-\omega_t) a_p(k_1)$ reflects the marginal effect of a change in v_{pt} on a weighted average of agents' expectation of inflation next period. The function, $W_{pt} \equiv W_p(v_{pt}, s_t, m_t)$ is marginal effect of v_{pt} on the central bank's loss function and $E_t W_{pt+1}$ satisfies:

$$E_t W_{pt+1} = \sum_{n=0}^{1} \Pr(m_{t+1} = n | m_t) \int_{s_{t+1}} W_p(v_{pt+1}, s_{t+1}, n) f(s_{t+1}|s_t) ds_{t+1} \tag{26}$$

Expression (24) extends the the celebrated LAW principle of Clarida et al. (1999) to an environment with finite horizon planning and longer-run learning. The first component of the right-hand side of the targeting rule reflects the static LAW principle derived by Clarida et al. (1999): If inflation is above the target $\pi_t > 0$, as a result of the cost-push shock, then the optimal policy pushes the output gap into negative territory ($y_t < 0$). The new targeting criterion, however, is not static. It differs from the period-by-period tight connection between inflation and output as now incorporates an additional term, $E_t W_{pt+1}$. To understand this term, note that the variable W_{pt} represents the marginal increase in central bank's expected discounted losses coming from an increase in agents' longer-run inflation beliefs, v_{pt}. Underlying an increase in v_{pt} resides an inflation scare in which private-sector's longer-run inflation expectations and thus actual inflation can remain persistently above a central bank's target. The optimal response to such an inflation scare is reflected in the term $E_t W_{pt+1} > 0$. That term implies that it is optimal for a central bank to act preemptively by pushing current output below potential ($y_t < 0$) and leaning against future inflation in order to put additional downward pressure on current inflation and keep private-sector's expectations of longer-run inflation in check.

A central bank's uncertainty about expectations formation also affects how they respond to inflation scares. According to equations (25) and (26), a central bank's motive to preempt inflation scares is stronger the more likely is the regime in which households and firms have short planning horizons. In particular, if a large fraction of agents have short planning horizons, the coefficient a_{pt} is relatively high, inflation scares are more likely, and a central bank has a strong incentive to lean against future inflation to prevent such a scare. The presence of uncertainty also makes the discount factor in equation (25) stochastic, as the central bank needs to account for the time-variation in the share of agents with different planning horizons. If the share of agents with short planning horizons is greater, this discount factor, $(\beta \left[1 - \gamma_p(1 - \beta a_{pt})\right])$, is higher, intensifying a central bank's incentive to lean against the risk of an inflation scare.

Special cases. Two special cases of the model occur when all agents have infinite planning horizons ($k_0, k_1 \to \infty$) or if agents do not update their longer-run beliefs about inflation (i.e.,

$\gamma_p = 0$). If all of the private-sector agents had infinite planning horizons, v_{pt} becomes irrelevant and inflation scares are not possible. This model corresponds to the baseline NK model. As a result, the central bank does not need to act preemptively and optimal time-consistent policy satisfies the static LAW principle of Clarida, Gali, and Gertler (1999): $\pi_t = -\frac{\lambda}{\kappa} y_t$. Similarly, if firms do not update their beliefs in response to past inflation (i.e, $\gamma_p = 0$), there is no variation in agents' longer-run inflationary beliefs. As a result, equation (24) also simplifies and satisfies the LAW principle of Clarida, Gali, and Gertler (1999).

Certainty-Equivalent Economy. To assess the role of uncertain expectations formation, we compare the dynamics of the model to a version where there is no uncertainty about agents' planning horizons. To do that we consider an economy in which all agents have the same planning horizon: a no uncertainty benchmark. Because certainty-equivalence (CE) is satisfied in this version of the model, we call this the certainty-equivalent economy and denote the planning horizon of the CE economy as k^{CE}. We choose k^{CE} to be the same as the average planning horizon in the economy with uncertain expectations formation. In the CE economy, a central bank's optimal targeting rule satisfies:

$$\pi_t^{CE} = -\frac{\lambda}{\kappa} y_t^{CE} - \gamma_p \beta E_t W_{pt+1}^{CE} \tag{27}$$

$$W_{pt}^{CE} = \beta a_p(k^{CE}) \pi_t^{CE} + \beta \left[(1-\gamma_p) + \beta \gamma_p a_p(k^{CE}) \right] E_t W_{pt+1}^{CE} \tag{28}$$

$$E_t W_{pt+1}^{CE} = \int_{u_{t+1}} W_p^{CE}(v_{pt+1}, u_{t+1}) f(u_{t+1}|u_t) du_{t+1}$$

The role of uncertainty about agents' planning horizon can be seen by comparing expressions (24) and (25) with expressions (27) and (28). The marginal effect of v_{pt} on expected inflation next period, a_{pt}, does not vary in the certainty-equivalent economy. Instead, it is fixed at, $a_p(k^{CE})$, the average value of the planning horizons of the two types of agents. Second, in the certainty-equivalent economy, $E_t W_{pt+1}^{CE}$ only depends on the additive shock, u_t; but, it does not depend on the multiplicative shock, m_t. As discussed above, this multiplicative shock implies that certainty-equivalence does not hold in the two-agent economy. Finally, we note that in the CE economy, as $k^{CE} \to \infty$, inflation scares do not occur and the optimal targeting rule satisfies the static LAW principle: $\pi_t^{CE} = -\frac{\lambda}{\kappa} y_t^{CE}$.

4 Results

In the previous section, we characterized optimal policy when there are inflation scares and private sector foresight is uncertain. In this section we quantify the importance of inflation scares and the role of uncertain private-sector foresight for optimal policy. To do so, we use the estimated parameters of the NK-FHP model from Gust, Herbst, and López-Salido (2022) and Gust, Herbst, and López-Salido (2024). Those papers show that the NK-FHP model performs well in explaining some stylized facts about the predictability of consensus inflation forecasts as well as aggregate data on inflation, output, and interest rates. In our application, we find it useful to examine the role of uncertainty about agents' foresight in an environment in which inflation has been running above a policymaker's objective. To do so, we use information from the FOMC's Survey of Economic Projections (SEP) in March of 2023 to construct a baseline path of aggregate demand and supply shocks from 2023Q1 to 2027Q4. We conduct stochastic simulations around this baseline path to quantify the gains to optimal policy when there is a risk of an inflation scare because of uncertainty about expectations formation.

4.1 Parameter Values

For most of the parameters, we use the mean estimates reported in Gust, Herbst, and López-Salido (2022) and Gust, Herbst, and López-Salido (2024): We estimate $\kappa = 0.03$, $\sigma = 2.71$, $\gamma_p = 0.16$, $\gamma_h = 0.5$. For the persistence of the shocks, we set $\rho_r = 0.95$ and $\rho_y = 0.6$. For the volatility of the innovations, we use $\sigma_r = 0.03$ and $\sigma_u = 0.02$. Overall, these values are consistent with 95% credible sets presented in Gust, Herbst, and López-Salido (2022) and our results are robust to small variations in these parameters.

For the uncertainty regarding planning horizons, we allow the economy to fluctuate between an economy in which all agents have planning horizons corresponding to a year ($k_0 = 4$) and another in which all agents have rational expectations. Accordingly, ω_t satisfies: $\omega_t = 1$ if $m_t = 0$ and $\omega_t = 0$ if $m_t = 1$. The value of $k_0 = 4$ is higher than the mean estimate reported in Gust, Herbst, and López-Salido (2022) but corresponds to the value they find best matches fluctuations in inflation expectations from the survey of professional forecasters as recently discussed in Gust, Herbst, and López-Salido (2024). The mode with $m_t = 1$ corresponds to the canonical NK model under rational expectation. When planning horizons are uncertain, certainty equivalance no longer applies and the appendix describes how we solve the model in that case.

We use the NK model with RE as one of our modes, since it has been widely used to study monetary policy under discretion and allows us to focus on the risk of an inflation scare. In particular, if $m_t = 1$ for all t, there would be no risk of an inflation scare, as longer-term inflation expectations would remain well anchored at a central bank's target under optimal policy in the canonical NK model. However, because m_t is time-varying and can potentially switch to the mode in which private-sector agents have short planning horizons ($m_t = 0$), a central bank faces the risk of an inflation scare in which trend inflation can move persistently away from a central bank's inflation target. Such a risk was a prominent concern in March of 2023, as several participants at the FOMC meeting at the time "noted the importance of longer-term inflation expectations remaining anchored and remarked that the longer inflation remained elevated, the greater the risk of inflation expectations becoming unanchored" (FOMC minutes, March 21-22, 2023.)

We set the Markov transition probabilities of the two modes so that they are persistent: $P_{00} = 0.9$ and $P_{11} = 0.9675$. These probabilities imply that the likelihood of staying in the the regime of the rational expectations, canonical NK model is higher than the likelihood of staying in the regime of the NK-FHP model. They also imply that the unconditional or ergodic probabilities of the modes with rational expectations and FHP expectations are 75% and 25%, respectively.

For the remaining parameters, we set $\beta = 0.99875$, which is consistent with the a steady state (annualized) real interest rate of 0.5%. This value is consistent with the median longer-run estimate of the federal funds rate reported in the March 2023 SEP. We also set $\pi^{\star}$, the central bank's inflation target to 2% on annualized basis and report annualized values of inflation and trend inflation that use this value both as the central bank's inflation target and the model's steady state inflation rate. For a central bank's preference parameter in the loss function, we set $\lambda = \frac{1}{16}$. This value implies that the central bank equally weighs deviations of annualized inflation from target and deviations of output from potential in its loss function (Debortoli et al. (2019)).

4.2 SEP-Consistent Baseline

As noted above, we construct a path of aggregate demand and supply shocks to be consistent with the median projection of inflation, the unemployment rate, and the federal funds rate from the March 2023 Summary of Economic Projections (SEP). To construct these shocks, we assume that the federal funds rate follows a Taylor rule:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) \left(\phi_\pi \pi_t + \phi_y y_t \right) + e_{mt} \tag{29}$$

where $\rho_i = 0.85$, $\phi_\pi = 1.5$, and $\phi_y = 0.25$ and e_{mt} is an *iid* innovation.[10] We use this rule along with the canonical NK model to infer shocks to aggregate demand, aggregate supply, and the policy rule that are consistent with the paths of core PCE inflation, the federal funds rate, and the output gap implied by the SEP. To construct the implied path for these variables from the SEP, we linearly interpolate the annual median projections shown in the SEP and construct the output gap assuming an Okun's parameter equals to 2. This value relates the output gap and deviations of the unemployment rate from the median SEP participant's estimate of its longer-run value.

The left-hand side of Figure 1 shows the SEP-consistent path of these variables. As shown there, SEP participants expected inflation to fall and eventually converge to 2 percent. The level of output implied by SEP participants projections of the unemployment-rate is above potential in early 2023 but is projected to decline in 2023 and remain below potential next year. The median SEP participant also projects that under appropriate monetary policy the federal funds rate will peak later this year and then fall in line with the projected decline in inflation.

We use the canonical NK model with RE to infer the path of shocks. That model, like the median path of the SEP, is consistent with a relatively costless return of inflation to target.[11] The right-hand side of Figure 1 shows the shocks over the 2023Q1-2027Q4 period that allow the model paths of output, inflation, and the policy rate to match the SEP-consistent baseline. The

[10]The view that Federal Reserve policy procedures have generally involved interest rate smoothing was introduced by Mankiw and Miron (1986). On this issue, see also the discussion in Goodfriend (1987) and the references therein. The rest of the parameters of this interest rate rule are from Taylor (1999).

[11]See Gust, Herbst, and López-Salido (2022) for a comparison of the cost of a pre-announced, permanent disinflation in the NK-FHP model to the canonical NK model. They show that canonical NK model implies a relatively costless disinflation and that the model with finite horizon planning can help account for historical estimates of the output cost of a disinflation.

Figure 1: The SEP-Consistent Baseline

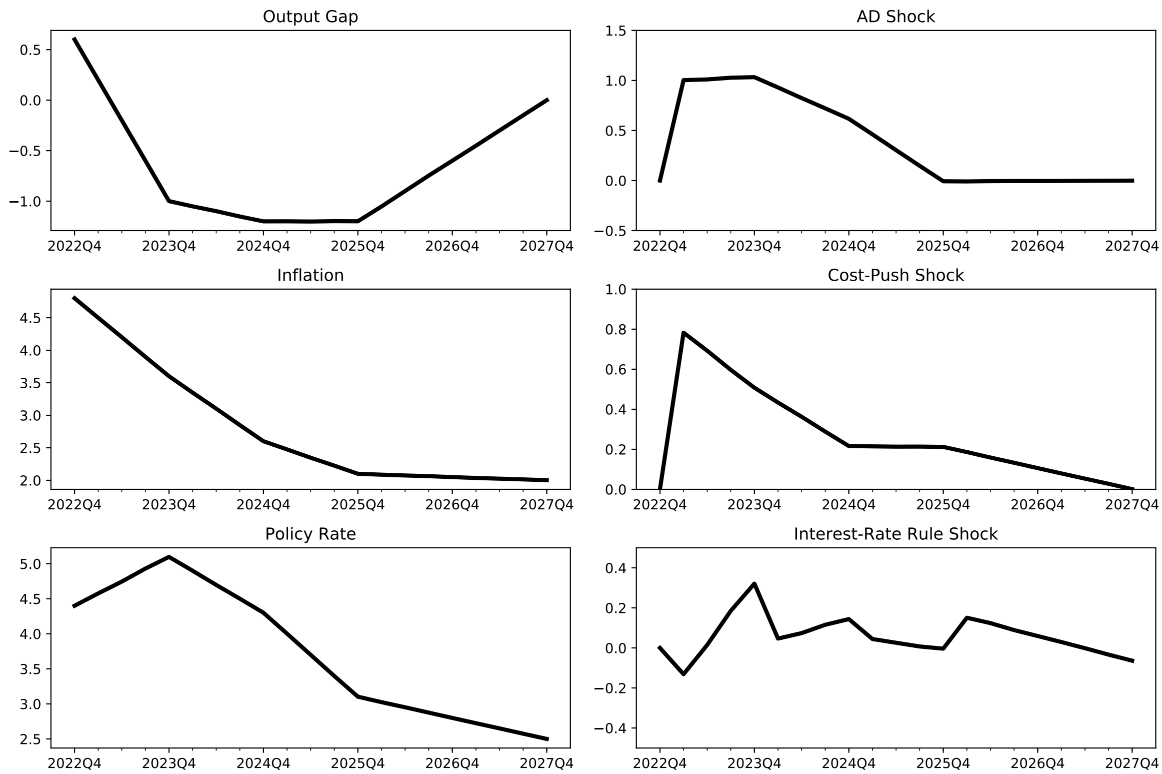

NOTE: The SEP-consistent baseline is constructed by interpolating quarterly data using the March 2023 SEP. The canonical NK model is used to infer the shocks to aggregate demand, aggregate supply (cost-push), and the interest-rate rule over the 2023Q1-2027Q4 period.

combination of aggregate demand and supply shocks shown there contributes to projections of high inflation along with output below potential in 2023. As these shocks fall back toward zero, inflation comes down and output converges toward potential. The lower right panel of Figure 1 shows that the monetary policy shocks are relatively small, less than 25 basis points over most of the projection period, suggesting that an inertial version of a Taylor rule fits the SEP-consistent baseline reasonably well.[12]

For simulations of the NK-FHP model, we also need initial conditions for household and firm's continuation value functions, v_{pt} and v_{ht}. Given the above parameter estimates for γ_p, we can determine v_{pt} in 2022Q4, the period before the start of the simulations, using equation (12) and past inflation data. Using this equation along with past data on core PCE inflation implies $v_{pt} = 0.41$ in 2022Q4. In the model with uncertainty about expectations formation, this value implies trend inflation around 2.4% in 2022Q4. We follow a similar procedure to initialize v_{ht}. In particular,

[12]Interestingly, the canonical NK model generates a combination of demand and supply shocks that resonates with the evidence presented in Blanchard and Bernanke (2023). These authors find an important role for pandemic-induced supply constraints as well as persistently higher aggregate demand as setting off the inflation in 2022.

equation (11) implies that v_{ht} depends on past values of the output gap and inflation. Using our parameter estimates along with the Congressional Budget Office's estimates of the output gap implies $v_{ht} = 3.1$ in 2022Q4.

4.3 The Transmission of Shocks in the NK-FHP Model

Figure 2: The Effect of an Aggregate-Supply Shock in the NK Model with Rational and FHP Expectations

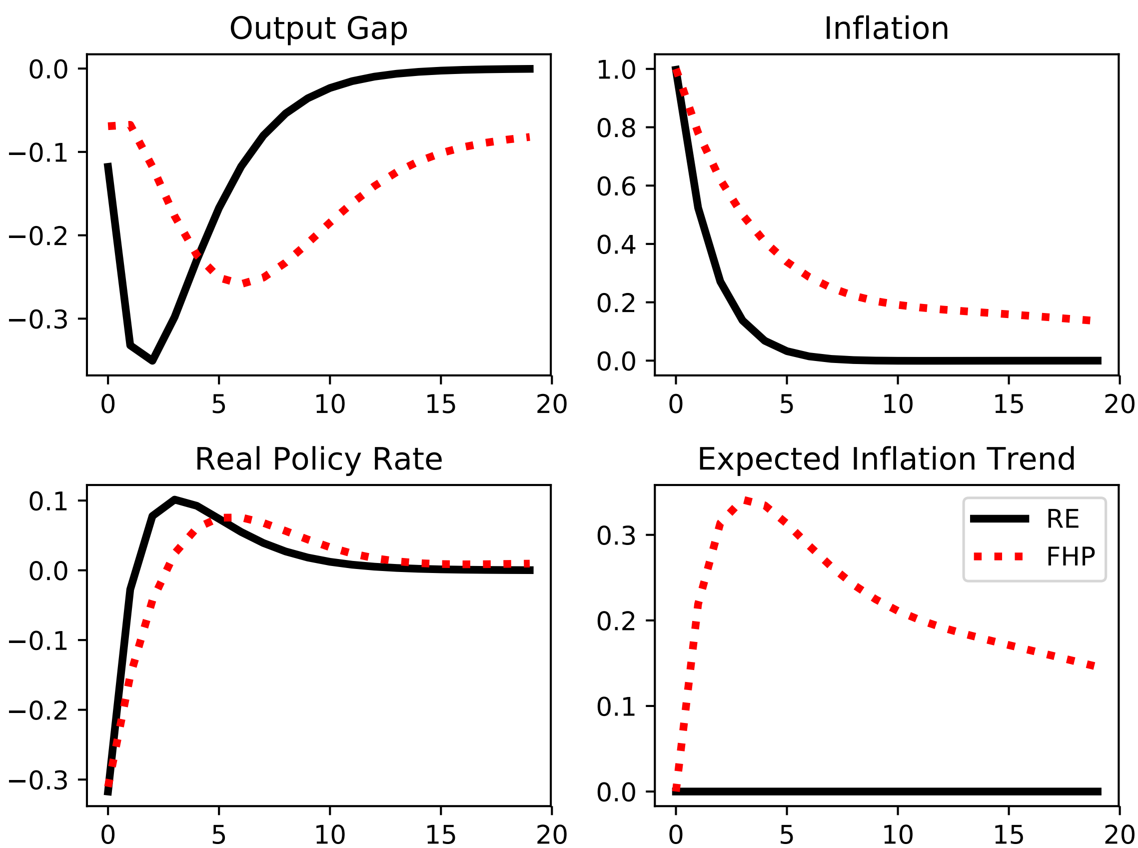

NOTE: The figure shows the effect of a shock to aggregate supply that increases annualized inflation by 1 percentage point in the canonical NK model with rational expectations (dark line) and the NK model with FHP expectations (red line). Monetary policy is assumed to follow a Taylor rule.

We use the aggregate demand and aggregate supply shocks associated with our SEP-consistent baseline to simulate optimal monetary policy over the 2023Q1-2027Q4. Before doing so, it is useful to examine how the transmission of shocks differs in the FHP model from the canonical NK. Because the aggregate supply shock plays an important role in our optimal policy simulations, Figure 2 shows the effects of an aggregate supply shock in both models. In the simulation, the policy rate is assumed to follow the inertial version of the Taylor (1993) rule, and the shock is constructed so that it increases inflation by 1 percentage point in both models. While the inertia in the policy rule implies an initial fall in the real policy rate in the canonical NK model, the real policy rises above steady state within the first year and remains slightly above steady state in the first two years after the shock. This tightening in policy along with the temporary nature of the shock implies that inflation quickly returns to steady state. Moreover, the trend component of expected inflation remains fixed in the canonical NK model.

In the NK-FHP model, the effects of the shocks on inflation, output, and the policy rate are noticeably more persistent than in the canonical NK model, reflecting that household and firm's beliefs about longer-run events depend on past economic outcomes. The dependence of longer-run beliefs on past economic outcomes implies that the trend component of expected inflation remains persistently above steady state in response to the adverse aggregate supply shock. The lower right panel of Figure 2 shows that the mean response of the trend component of expected inflation rises persistently in the FHP model.[13] This rise in the trend component leads to an inflation rate that remains above its steady state level of 2 percent for an extended period of time. With inflation rising persistently, the policy rate rises higher than its response in the canonical NK model and stays elevated well for longer as well. As a result, output falls and remains persistently below potential. In short, an unfavorable aggregate supply shock in the NK-FHP model results an increase in longer-run expectations of inflation that induces an unfavorable tradeoff between inflation and output that persists much longer than in the canonical NK model. The aggregate demand and monetary shock are also notably more persistent in the NK-FHP model than in the canonical NK model.[14]

4.4 Optimal Policy under Uncertainty

The black and red lines in Figure 3 display the expected paths of inflation, the output gap, the real interest rate, and trend inflation in the canonical NK model (labelled RE) and FHP model under optimal policy using the baseline aggregate demand and supply shocks shown in Figure 1. In both cases, there is no uncertainty about expectations formation so each model mode applies with perfect certainty. For the canonical NK model, the only difference in outcomes between this figure and the ones shown in Figure 1 reflects the difference in monetary policy. In this figure, the optimal targeting rule applies, which in the case of the canonical NK model satisfies the LAW relationship: $\kappa\pi_t = -\lambda y_t$. This targeting rule and the SEP-consistent policy result in broadly similar outcomes. Under the optimal targeting rule, the policy rate peaks at a higher level and declines faster in 2023 than under the SEP-consistent baseline. Output remains closer to potential under optimal policy, but the decline in inflation is more gradual relative to the SEP-consistent baseline.

The red line shows the results if economic outcomes were determined by the NK-FHP model with perfect certainty. Here, the economy starts with the trend component of expected inflation above 3% and policy needs to be substantially more aggressive to bring inflation down. With trend

[13]The trend component is defined by ignoring the effects of the shocks on $E_t\Pi_{t+1}$ so that the trend component only reflects the effects of agents' value functions on expected inflation.

[14]Gust, Herbst, and López-Salido (2022) show that the transmission of monetary shocks occurs more gradually in the NK-FHP model than in the canonical NK model with the peak effect on inflation occurring considerably later.

Figure 3: Optimal Policy under Different Expectational Assumptions

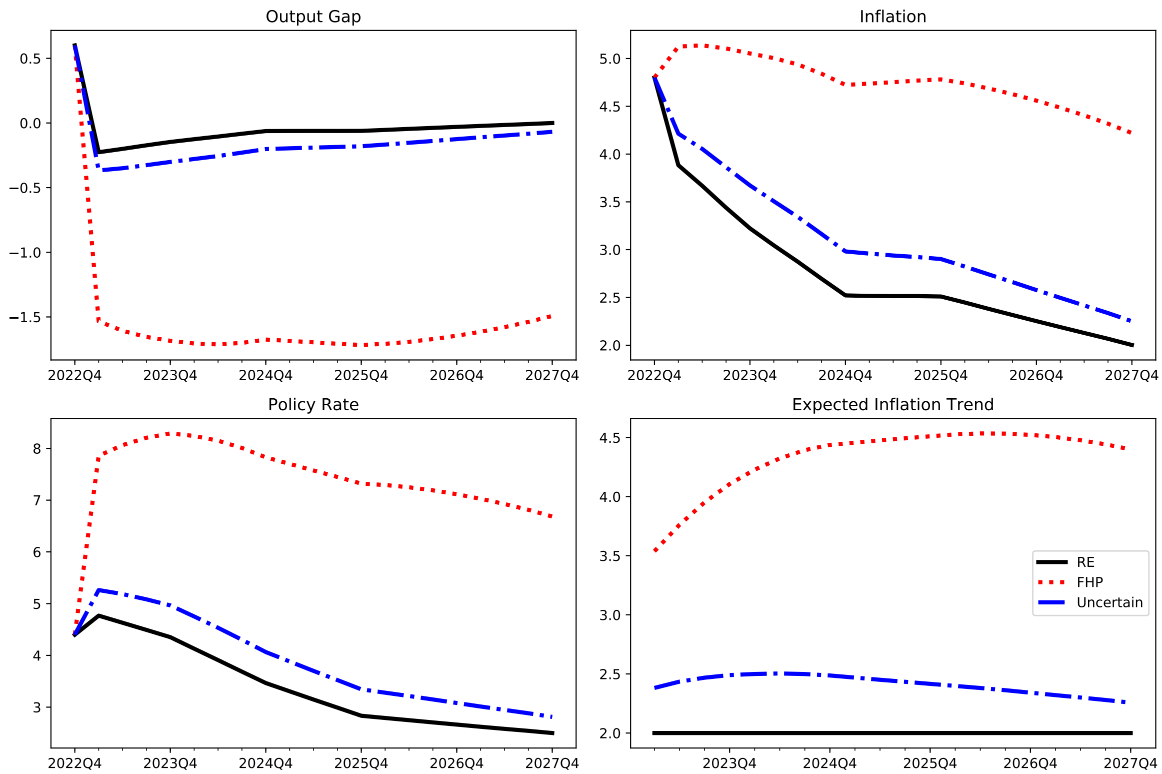

NOTE: The figure shows the mean responses under optimal discretionary policies using 50,000 draws of shocks centered around the aggregate demand and supply shocks associated with the SEP-consistent baseline. The black line shows the optimal policy paths under canonical NK model mode with rational expectations; the red dashed line shows the optimal policy paths using the NK-FHP model mode with inertial expectations; and, the dashed blue shows the optimal policy paths when there is uncertainty about expectations formation.

inflation inertial and costly to reduce, the optimal policy response involves raising the policy rate above 8%. This policy response results in a much deeper decline in the output gap than in the NK model with RE but is necessary to put inflation on a downward trajectory.

The dashed-dotted blue line shows the expected path of outcomes under optimal policy when the central bank is uncertain about expectations formation. We start the economy from its unconditional distribution so that at each date there is a 25% chance the economy is in the NK-FHP model mode where expectations formation is inertial. Accordingly, at each date, the expected path of outcomes lie in between the paths where the modes are known with certainty. When the central bank is uncertain about expectations formation, inflation comes down slower than it does under the canonical NK model and is still above target at the end of the 2027. This reflects the inertia in private sector's longer-run inflation beliefs, as the trend component of expected inflation remains above 2% at the end of 2027. In effect, progress in returning inflation back to the central bank's target is slow, because a policy that resulted in faster progress would result in an even greater decline in economic activity, which the central bank views as undesirable.

Figure 4: The Policy Tradeoff Frontier Under Optimal Discretion

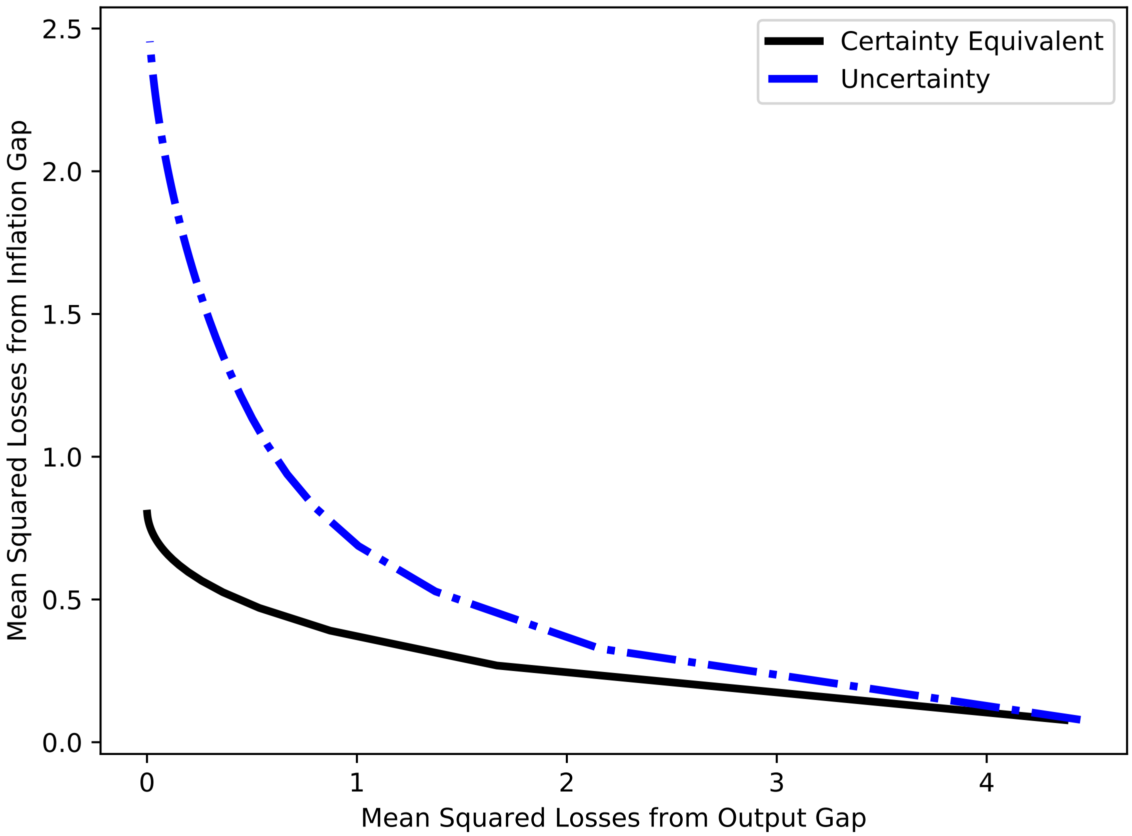

NOTE: For the two models, each point along the policy tradeoff frontier is constructed by computing the mean-squared deviations of the output and inflation gaps from stochastic simulations around the aggregate demand and supply shocks associated with the SEP-consistent baseline for a given value of λ. The black line shows the tradeoff in the certainty-equivalent economy, and the dashed-dotted blue line shows the tradeoff when there is uncertainty about expectations formation.

Figure 3 indicates that policy responds more aggressively under uncertainty than in the CE economy. Because the mode with rational expectations corresponds to $k_1 \to \infty$, the average planning horizon in the economy in which the planning horizon is uncertain is the same as for that mode.[15] Thus, the canonical NK model with rational expectations has the same average planning horizon as the economy with an uncertain planning horizon and we use it as our certainty-equivalent benchmark. Figure 3 shows that, relative to that CE benchmark economy, the central bank tightens policy notably more than under certainty equivalence because in the economy with

[15]The average planning horizon in the economy with uncertain expectations formation is given by $k_A = \bar{p}_0 k_0 + (1 - \bar{p}_0)k_1$, where $\bar{p}_0$ denotes the ergodic probability of the mode with short planning horizons. With $p_0 = 0.25$, $k_0 = 4$, and using a very large value (e.g., $k_1 = 10000$) to approximate RE for this mode, there is little difference between k_A and k_1.

uncertain expectations formation there is a significant risk of an inflation scare where private sector's beliefs about longer-run inflation can remain persistently above the central bank's inflation objective.

Policy Tradeoff Frontier. The more aggressive response of monetary policy under uncertainty relative to the response under CE reflects that uncertainty about expectations formation increases the likelihood of an inflation scare. Such a scare is particularly pernicious in terms of a central bank's losses because it can persistently worsen the tradeoff between the central bank's objectives. Accordingly, optimal policy under uncertainty responds preemptively and more aggressively towards inflation to reduce the risk that inflationary pressure increases private-sector beliefs about longer-run inflation.

Figure 4 highlights how uncertainty about the planning horizon and trend inflation affects the tradeoff between a central bank's objectives. It traces out a central bank's tradeoff between inflation and the output gap under uncertainty about expectations formation and under certainty equivalence. In the CE economy of the canonical NK model, if a central bank would like to reduce the losses associated with deviations of output from potential, the cost in terms of increased deviations of inflation from target is relatively low. In contrast, in the economy with uncertain expectations formation, reducing output fluctuations could set off an inflation scare, leading to substantial losses from high inflation. Thus, the tradeoff between central bank objectives is considerably steeper under uncertainty than in the certainty-equivalent economy.

Distribution of Macroeconomic Outcomes. Figure 3 compared the mean outcomes of optimal policy with uncertainty about expectations formation to the mean outcomes of its certainty-equivalent counterpart. Figure 5 compares the distribution of outcomes in 2024Q1 for the two models. In the CE economy of the canonical NK model, shown by the black lines, the distributions for the output gap, inflation, and the policy rate are normally distributed and centered around the mean outcomes for these variables in 2024Q1. There is no distribution for trend inflation since it is degenerate, as trend inflation remains anchored at 2% at all times in the canonical NK model.

When the central bank faces uncertainty about expectations formation, the distribution of outcomes is wider with the distributions displaying long tails. These long tails reflect the draws in which an inflation scare is realized and higher trend inflation emerges. This risk is shown in the lower right panel, which shows that outcomes with trend inflation well above the central bank's inflation target are possible. This higher upside risk to inflation heightens the likelihood of lower output, as the policy rate may need to go much higher to reduce trend inflation. Consistent with the upside risk to inflation, the distributions for both inflation and the policy rate display long right tails.

Gains from a Risk-Management Approach. In an environment where there is a heightened risk of an inflation scare, the gains from a risk-management approach can be considerable. To demonstrate this, Table 1 shows the additional losses from the central bank following the targeting rule that is optimal in the CE economy in an economy in which there is uncertainty about expectations formation. In particular, we compare the effects on a central bank's expected discounted

Figure 5: The Distribution of Outcomes under Optimal Discretion in 2024Q1

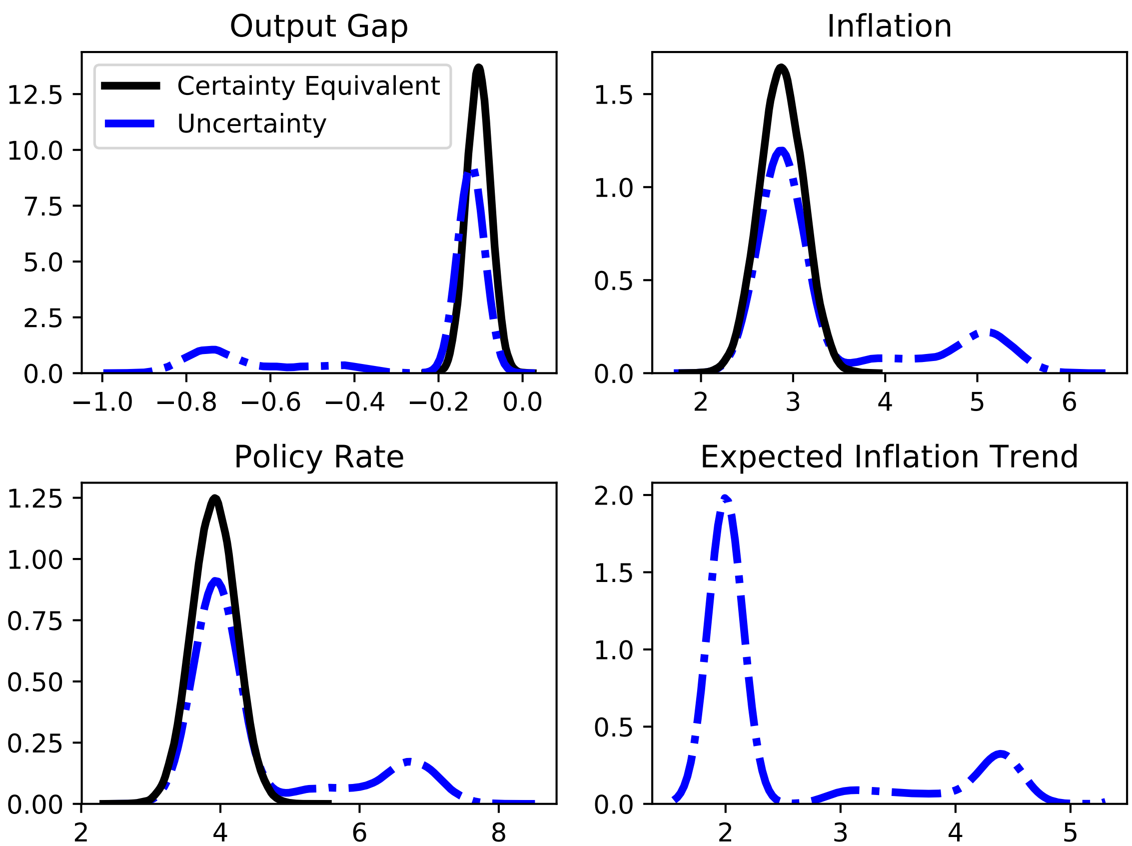

NOTE: The figure shows the distribution of responses under optimal discretionary policies in 2024Q1 using 50,000 draws of shocks centered around the aggregate demand and supply shocks associated with the SEP-consistent baseline.

losses from following the CE targeting rule of the canonical NK model in which $\kappa\pi_t = -\lambda y_t$ instead of the optimal targeting rule under uncertainty given by equation (24). Table 1 shows how much higher a central bank's expected discounted losses are in percentage terms from adopting the CE targeting rule for different degrees of uncertainty and for different central bank preferences regarding the tradeoff between inflation and economic activity. It also shows how much higher optimal policy under uncertainty increases the policy rate in 2023 than the policy rate increases under the CE targeting rule.

If the likelihood of the FHP model mode and thus inflation scares is low, then the losses associated with following a targeting rule that ignores risk-management considerations are small. However, these additional losses can grow quickly as that likelihood increases. For a loss function that equally weights inflation and output-gap deviations ($\lambda = \frac{1}{16}$), then the additional losses from ignoring risk-management considerations can be high. For example, with an equal-weighted loss

Table 1: Additional Losses from Ignoring Risk Management

Probability of FHP Model Mode	Additional Losses (Percent) $\lambda = 0.01$	Additional Losses (Percent) $\lambda = \frac{1}{16}$	Policy Rate Difference (basis points) $\lambda = 0.01$	Policy Rate Difference (basis points) $\lambda = \frac{1}{16}$
10%	0.1%	0.3%	12	7
15%	0.3%	0.7%	18	11
20%	0.4%	1.4%	24	15
25%	0.6%	2.8%	30	19
30%	1.1%	7.2%	36	23
33.3%	1.6%	20.8%	41	26

Note: The table shows the average percent increase in a central bank's losses over the 2023-2027 period from adopting the optimal, certainty-equivalent policy instead of the optimal policy under uncertainty. It also shows the average policy rate increase in 2023 under optimal policy relative to the certainty-equivalent policy response. These averages are computed from stochastic simulations using 50,000 draws of shocks centered around the aggregate demand and supply shocks associated with the SEP-baseline.

function and a 33% chance of being in a regime in which longer-run inflationary beliefs can deviate from a central bank's target, the losses from following the CE targeting rule are about 20 percent higher than following the optimal policy that takes into uncertainty about expectations formation.

A policymaker, who takes into account risk management considerations, is able to reduce their losses by raising the policy rate higher than a policymaker who ignores the risk of an inflation scare. If the policymaker judges there to be a 30% probability of being in the regime where inflation scares are possible, Table 1 shows that a policymaker with an equal-weighted loss function, and who takes into account uncertainty about expectations formation, would raise the policy rate about 25 basis points higher in 2023 than a policymaker who acts with certainty about expectations formation. If the policymaker had a stronger preference for inflation stabilization ($\lambda = 0.01$), then the policy rate would be even more responsive to risk-management considerations, as the differences between the policy rate under optimal policy and the certainty equivalent policy are even larger in that case.

5 Concluding Remarks

We analyzed optimal monetary policy when the central bank faces uncertainty about the foresight of private sector agents. Because agents learn adaptively about developments beyond their planning horizon, inflation scares in which longer-run inflation expectations of the private sector can move persistently away from a central bank's inflation objective are possible. A central bank with a dual-mandate loss function, who faces such a risk, reacts forcefully and preemptively to contain inflationary pressure, responding more aggressively than would be the case under certainty-equivalence. In situations where inflation has been running above target, we find the gains to taking a risk-management approach in response to such risks can be sizable.

We focused on the case in which a central bank faces future uncertainty to keep the analysis tractable. While more computationally demanding, a natural extension would be to analysis optimal policy when the central bank does not observe the current state of the economy and updates their beliefs according to Bayes' law (Svensson and Williams (2008)). It would also be interesting to explore whether the transition probabilities for the uncertain regimes should depend on the economic state (Davig and Leeper (2008)). Long spells of inflation deviating from the central bank's objective may increase the chance of an unfavorable shift in private sector's expectations, which would reinforce a central bank's desire to preemptively act against emerging inflationary pressures.

References

Barlevy, G. (2011): "Robustness and Macroeconomic Policy," *Annual Review of Economics*, 3, 1–24.

Blanchard, O. J. and B. S. Bernanke (2023): "What Caused the US Pandemic-Era Inflation?" Working Paper 31417, National Bureau of Economic Research.

Brainard, W. (1967): "Uncertainty and the Effectiveness of Policy," *American Economic Review*, 57, 411 – 425.

Clarida, R., J. Gali, and M. Gertler (1999): "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 37, 1661–1707.

Cogley, T. and T. J. Sargent (2008): "Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making," *International Economic Review*, 49, 185–221.

Davig, T. and E. M. Leeper (2008): "Endogenous Monetary Policy Regime Change," in *NBER International Seminar on Macroeconomics 2006*, National Bureau of Economic Research, Inc, NBER Chapters, 345–391.

Debortoli, D., J. Kim, J. Lindé, and R. Nunes (2019): "Designing a Simple Loss Function for Central Banks: Does a Dual Mandate Make Sense?" *The Economic Journal*, 129, 2010–2038.

Eusepi, S. and B. Preston (2018): "The Science of Monetary Policy: An Imperfect Knowledge Perspective," *Journal of Economic Literature*, 56, 3–59.

Foerster, A., J. F. Rubio-Ramírez, D. F. Waggoner, and T. Zha (2016): "Perturbation Methods for Markov-Switching Dynamic Stochastic General Equilibrium Models," *Quantitative Economics*, 7, 637–669.

Goodfriend, M. (1987): "Interest Rate Smoothing and Price Level Trend-Stationarity," *Journal of Monetary Economics*, 19, 335–348.

——— (1993): "Interest Rate Policy and the Inflation Scare Problem: 1979-1992," *Economic Quarterly*, 1–24.

Gust, C., E. Herbst, and D. López-Salido (2024): "Inflation Expectations with Finite Horizon Planning," Finance and Economics Discussion Series 3556, Federal Reserve Board of Governors.

Gust, C., E. Herbst, and D. López-Salido (2022): "Short-Term Planning, Monetary Policy, and Macroeconomic Persistence," *American Economic Journal: Macroeconomics*, 14, 174–209.

Gáti, L. (2023): "Monetary Policy and Anchored Expectations—An Endogenous Gain Learning Model," *Journal of Monetary Economics*, 140, S37–S47.

Kimura, T. and T. Kurozumi (2007): "Optimal Monetary Policy in a Micro-Founded Model with Parameter Uncertainty," *Journal of Economic Dynamics and Control*, 31, 399–431.

Mankiw, N. G. and J. A. Miron (1986): "The Changing Behavior of the Term Structure of Interest Rates," *The Quarterly Journal of Economics*, 101, 211–228.

Molnár, K. and S. Santoro (2014): "Optimal Monetary Policy When Agents are Learning," *European Economic Review*, 66, 39–62.

ORPHANIDES, A. AND J. WILLIAMS (2022): "Taming Inflation Scares," Prepared for a Federal Reserve Bank of Richmond project "Marvin Goodfriend: Economist and Central Banker".

ORPHANIDES, A. AND J. C. WILLIAMS (2005): "Inflation Scares and Forecast-Based Monetary Policy," *Review of Economic Dynamics*, 8, 498–527.

SVENSSON, L. E. AND N. WILLIAMS (2005): "Monetary Policy With Model Uncertainty: Distribution Forecast Targeting," Working Paper 11733, National Bureau of Economic Research.

——— (2008): "Optimal Monetary Policy under Uncertainty in DSGE Models: A Markov Jump-Linear-Quadratic Approach," Working Paper 13892, National Bureau of Economic Research.

SÖDERSTRÖM, U. (2002): "Monetary Policy with Uncertain Parameters," *The Scandinavian Journal of Economics*, 104, 125–145.

TAYLOR, J. B. (1999): "A Historical Analysis of Monetary Policy Rules," in *Monetary Policy Rules*, National Bureau of Economic Research, Inc, NBER Chapters, 319–348.

WOODFORD, M. (2018): "Monetary Policy Analysis When Planning Horizons are Finite," Working Paper 24692, National Bureau of Economic Research.

WOODFORD, M. AND Y. XIE (2022): "Fiscal and monetary stabilization policy at the zero lower bound: Consequences of limited foresight," *Journal of Monetary Economics*, 125, 18–35.

Appendix for
Optimal Monetary Policy with Uncertain Private Sector Foresight

Christopher Gust and David López-Salido

In the appendix, we derive the expressions that determine private-sector expectations formation under optimal time-consistent policy. We also derive the targeting rule implemented under optimal policy and describe the solution algorithm we use to determine the associated equilibrium outcomes.

A Private Sector Expectations

We begin by deriving expression (9) in the main text, which determines expected inflation of a k-horizon agent. This expression can be derived from substituting equation (19) at $j = 0$ into equation (8) to write an agent's beliefs for inflation at the end of their planning horizon as:

$$\pi^0_{t+k} = \frac{\lambda}{\lambda + \kappa^2} \left[\beta v_{pt} + u_{t+k}\right]. \tag{A-1}$$

For $\tau = t + k - j$ with $j > 0$, we can derive a similar condition by substituting equation (19) at $j > 0$ into equation (1):

$$\pi^j_\tau = \frac{\lambda}{\lambda + \kappa^2} \left[\beta \pi^{j-1}_{\tau+1} + u_\tau\right]. \tag{A-2}$$

Substituting equation (A-1) into (A-2) at $j = 1$ yields:

$$\pi^1_{t+k-1} = \frac{\beta\lambda}{\lambda + \kappa^2} v_{pt} + \frac{\lambda}{\lambda + \kappa^2} \left[\sum_{i=0}^{1} (\beta\rho_u)^i\right] u_{t+k-1} \tag{A-3}$$

Continuing with these substitutions back to $j = k - 1$ yields:

$$E_t \pi^{k-1}_{t+1} = a_p(k) v_{pt} + \frac{\lambda\rho_u}{\lambda + \kappa^2} \left[\sum_{i=0}^{k-1} \left(\frac{\beta\lambda\rho_u}{\lambda + \kappa^2}\right)^i\right] u_t \tag{A-4}$$

which is expression (9) in the main text with

$$a_p(k) = \left(\frac{\beta\lambda}{\lambda + \kappa^2}\right)^k$$

We can substitute expression (A-4) into (19) at $j = k - 1$ to determine the expected output gap for an agent that looks k-periods ahead:

$$E_t y^{k-1}_{t+1} = -\frac{\kappa}{\lambda} \left\{ a_p(k) v_{pt} + \frac{\lambda\rho_u}{\lambda + \kappa^2} \left[\sum_{i=0}^{k-1} \left(\frac{\beta\lambda\rho_u}{\lambda + \kappa^2}\right)^i\right] u_t \right\} \tag{A-5}$$

B Derivation of Optimal Targeting Rule

To derive the optimal targeting rule, equation (24), we need to differentiate the Bellman equation (21) with respect to i_t. Doing so, taking into account the dependence of $\Pi_t \equiv \pi(v_{pt}, s_t, m_t; i_t)$ and $\mathbf{Y}_t \equiv y(v_{pt}, s_t, m_t; i_t)$ on i_t yields the first order condition:

$$\pi_t + \gamma_p \beta E_t W_{pt+1} = -\frac{\lambda}{\kappa} y_t \tag{A-6}$$

$$\tag{A-7}$$

where $E_t W_{pt+1}$ satisfies equation (26). This first order condition is the first part of the optimal targeting rule shown in the main text. The other part comes from the envelope condition associated with the the Bellman equation (21). This envelope condition satisfies:

$$W_{pt} = (\beta+\kappa\sigma)(a_{pt}+\kappa a_{yt})\pi_t + \lambda y_t(a_{yt}+\sigma a_{pt}) + \beta(1-\gamma_p)E_t W_{pt+1} + \beta\gamma(\beta+\kappa\sigma)(a_{pt}+\kappa a_{yt})E_t W_{pt+1} \tag{A-8}$$

where $a_{pt} = \omega_t a_p(k_0) + (1-\omega_t)a_p(k_1)$ and $a_{yt} = -\frac{\kappa}{\lambda}a_{pt}$. We can simplify expression (A-8) by combining it with the first order condition for i_t and writing it as:

$$W_{pt} = \beta a_{pt}\pi_t + \beta(1-\gamma_p)E_t W_{pt+1} + \beta^2\gamma_p a_{pt}E_t W_{pt+1} \tag{A-9}$$

Collecting terms associated with $E_t W_{pt+1}$, this expression can be rewritten as:

$$W_{pt} = \beta a_{pt}\pi_t + \beta\left[1-\gamma_p(1-\beta a_{pt})\right]E_t W_{pt+1}$$

which is the expression for W_{pt} shown in the text.

C Solution Algorithm

To solve for the outcomes associated with the optimal targeting rule, note that the following system of equations can be used to determine the outcomes for inflation, the output gap, and W_{pt}, and v_{pt+1} as a function of v_{pt} and u_t:

$$\pi_t = \beta g_\pi(v_{pt}, u_t, m_t) + \kappa y_t + u_t \tag{A-10}$$

$$\pi_t + \gamma_p\beta E_t W_{pt+1} = -\frac{\lambda}{\kappa}y_t \tag{A-11}$$

$$W_{pt} = \beta a_{pt}\pi_t + \beta\left[1-\gamma_p(1-\beta a_{pt})\right]E_t W_{pt+1} \tag{A-12}$$

$$v_{pt+1} = (1-\gamma_p)v_{pt} + \gamma_p\pi_t \tag{A-13}$$

With these outcomes in hand, the optimal policy rate can then be determined using equation (23).

Our calibration implies that all the agents have short-planning horizons when $m_t = 0$ so that $\omega_t = 1$ and $a_{pt} = a_p(k_0)$ in that case. When $m_t = 1$, all agents have long-planning horizons so that $\omega_t = 0$ and $a_{pt} = 0$. Using this calibration, the function determining private-sector expectations of inflation can be written as:

$$g_\pi(v_{pt}, u_t, m_t) = \begin{cases} a_p(k_0)v_{pt} + b_p(k_0)u_t \text{ for } m_t = 0 \\ \frac{1}{1-a_p(1)\rho_u}u_t \text{ for } m_t = 1. \end{cases} \tag{A-14}$$

where $a_p(k) = \left(\frac{\beta\lambda}{\lambda+\kappa^2}\right)^k$. Expression (A-14) is consistent with expression (A-4), as the coefficient $b_p(k)$ satisfies:

$$b_p(k) = \frac{\lambda\rho_u}{\lambda+\kappa^2}\sum_{i=0}^{k-1}\left(\frac{\beta\lambda\rho_u}{\lambda+\kappa^2}\right)^i = \frac{\lambda\rho_u}{\lambda+\kappa^2}\left[\frac{1-a_p(k)\rho_u^k}{1-a_p(1)\rho_u}\right].$$

Equations (A-10)-(A-13) are linear conditional on a value for m_t and we solve for a solution of the form:

$$X_t(m) = T_m v_{pt} + R_m u_t, \tag{A-15}$$

with $m \in \{0, 1\}$ and $X_t(m) = (\pi_t(m), y_t(m), W_{pt}(m))'$. Accordingly, the solution is conditionally linear in m, with the solution matrices, T_m and R_m, varying depending on whether $m = 0$ or $m = 1$ at time t.

To determine T_m and R_m, we write the system of equations, (A-10)-(A-13), over a long horizon, truncating the horizon at $t + K_{CB}$. For the periods before this truncation point, the equilibrium conditions in matrix form can be written as:

$$C_m X_\tau(m) = F_m E_\tau\left[X_{\tau+1}(m')|m\right] + B_m X_{\tau-1} + D_m u_\tau, \tag{A-16}$$

for $\tau = t + K_{CB} - i$ with $i = 1, 2..., K_{CB}$ and $m \in \{0, 1\}$. For the terminal period, we impose

$$C_m X_{t+K_{CB}}(m) = B_m X_{t+K_{CB}-1}(m) + D_m u_{t+K_{CB}}. \tag{A-17}$$

Relative to the infinite-horizon problem, equation (A-17) is truncated at date $t + K_{CB}$ as it omits the expected future endogenous variables, effectively treating the central bank as if it had a finite planning horizon. As $K_{CB} \to \infty$, the equilibrium conditions for this problem converge to those given by equations (A-10)-(A-13). For the simulations in the paper, we checked that as $K_{CB} \to \infty$, the solution matrices converged and set $K_{CB} = 10,000$, which we found to be sufficiently large to ensure convergence.

The matrices $\{C_m, F_m, B_m\}$ capture the contemporaneous, forward-looking, and backward-looking relationships of the model equations, respectively, while D_m captures the effects of the aggregate supply shock, u_t. The conditional expectations operator in equation (A-16) reflects the transition probabilities for the modes and satisfies:

$$E_\tau\left[X_{\tau+1}(m')|m\right] = \left[P_{m,0}T_0 + P_{m,1}T_1\right] X_\tau(0) + \left[P_{m,0}R_0 + P_{m,1}R_1\right] \rho_u u_t \tag{A-18}$$

To iterate backwards on this system, we first solve equation (A-17) for each mode, $m \in \{0, 1\}$, which provides an initial guess for the solution matrices:

$$T_m^0 = (C_m)^{-1} B_m,$$
$$R_m^0 = (C_m)^{-1} D_m,$$

With the initial guess and the conditional expectations operator defined in (A-18), we can iterate backwards on equation (A-16) to solve for $T_{j_\tau}^i$ and $R_{j_\tau}^i$ at each date τ for $i = 1, \ldots, K_{CB}$:

$$T_m^i = \left\{C_m - F_m\left[\sum_{n=0}^{1} P_{m,n} T_n^{i-1}\right]\right\}^{-1} B_m \tag{A-19}$$

$$R_m^i = \left\{C_m - F_m\left[\sum_{n=0}^{1} P_{m,n} T_n^{i-1}\right]\right\}^{-1} \left\{D_m + F_m\left[\sum_{n=0}^{1} P_{m,n} R_n^{i-1}\right]\rho_u\right\}. \tag{A-20}$$

We solve for $T_m^{K_{CB}}$ and $R_m^{K_{CB}}$ and check that as $K_{CB} \to \infty$, $T_m^{K_{CB}} \to T_m$ and $R_m^{K_{CB}} \to R_m$. Because the solution algorithm truncates the central bank's planning horizon, it yields a unique solution that corresponds to one that is consistent with the central bank having a finite but very long planning horizon. For a general approach to solving dynamic stochastic general equilibrium models with Markov-switching processes, see Foerster et al. (2016).