Optimal Monetary Policy with Uncertain Private Sector Foresight*

Christopher Gust† Edward Herbst‡ David López-Salido§

July 2025

Abstract

We model private-sector expectations in a finite-horizon-planning framework: households and firms have limited foresight when making spending, saving, and pricing decisions. In this setting, contrary to standard New Keynesian (NK) models, we show that an "inflation scare" problem can arise in which agents' longer-run inflation expectations deviate persistently from a central bank's inflation target. We 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.

JEL Classification: C11, E52, E70
Keywords: Finite horizon planning, monetary policy under uncertainty.

*We thank the comments from an anonymous referee and Boragan Aruoba. Special thanks to Luminita Stevens for many constructive and detailed comments and suggestions. This paper presents our views and should not be interpreted as reflecting those of the Board or the Federal Reserve System.

†Federal Reserve Board. Email: christopher.gust@frb.gov
‡Federal Reserve Board. Email: edward.herbst@frb.gov
§Federal Reserve Board. Email: david.lopez-salido@frb.gov

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 studying 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, but 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 well 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. Fluctuations in the distribution of agents' planning horizons are a non-additive source of uncertainty that matters for optimal monetary policy and imply that optimal monetary policy depends on the distributions of economic variables such as inflation and the output gap and not just the means of these variables.

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 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, they 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 by 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 monetary policy can

[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.

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 et al. (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, Gust et al. (2022) show that the model's longer-run inflation expectations can deviate persistently away from the central bank's target, mimicking the empirical behavior of the longer-run inflation expectations seen in survey data in the 1970s and early 1980s. Gust et al. (2024) extend this work and show—analytically and empirically—that the NK-FHP model can account for the initial underreaction and subsequent overreaction of inflation forecasts emphasized in Angeletos et al. (2020).

Our theoretical analysis builds on the prominent work of Clarida et al. (1999) (hereafter CGG), 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 the 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 CGG. But, when agents have short planning horizons, their beliefs about longer-run inflation can move persistently away from the 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 the fact that the central bank realizes that inflationary pressure will boost the private sector's future beliefs about trend inflation, leading to long-lasting departures of inflation from the 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 longer-run expectations of inflation at a central bank's 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 result stands in contrast to the classic finding 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 an undesirable outcome, the central bank sets a relatively more aggressive policy path than under certainty.

This more aggressive policy response helps anchor private-sector expectations of longer-run inflation but results in a distribution of outcomes for output with a long tail of below-potential activity. These less favorable outcomes when planning horizons are uncertain reflect that the policy tradeoff between inflation and output gap stabilization worsens 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 estimate the NK-FHP model with uncertain planning horizons and use it to quantify the gains from guarding against inflation scares in March 2023 — a time when there was a heightened risk that above-target inflation could lead to an unanchoring of longer-run inflation expectations. We find that the response of the optimal policy rate is notably more aggressive than the response that would be optimal in the absence of inflation scares. Moreover, our analysis suggests that the gains associated with guarding against inflation scares are sizable when inflation is well above a central bank's target as it was in March 2023.

Literature review. Building on CGG, 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.[2] Relative to their work, we study optimal monetary policy under discretion when agents have finite planning horizons and 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).[3] Like our paper, these papers emphasize that private-sector inflation expectations are more important when agents are learning and that there is an increased role for stabilizing inflation.[4] Our approach is distinct from these papers since we emphasize the role of uncertainty and in particular study the design of optimal policy when

[2]See also Dupraz and Marx (2025) who study the properties of simple rules and characterize optimal policy in an NK-FHP model.

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

[4]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.

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 studies that emphasize uncertainty about inflation dynamics using a Bayesian approach. Söderström (2002), Kimura and Kurozumi (2007), and Svensson and Williams (2008) incorporate lagged dependence into the inflation process and show that the optimal policy response is not attenuated as in Brainard (1967) but is more aggressive than in the case under certainty equivalence. Our model differs from these earlier papers in that we explicitly model uncertainty about expectations formation and the unanchoring of agents' longer-run inflation beliefs, and we investigate the mechanism quantitatively during the high inflation that occurred in the United States in the aftermath of the pandemic.[5]

The rest of this paper proceeds as follows. The next section presents 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 presents the results and the final section offers conclusions and directions for future work.

2 A Finite-Horizon Model with Uncertain Planning

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. In this framework, households and firms make a complete set of state-contingent plans only up to the end of their planning horizon and use a value function based on past experience to assess the value of future events beyond that horizon. To model uncertainty, we follow Svensson and Williams (2005), who model uncertainty using different "modes" or regimes that follow a Markov process. We assume that there are two types of households and firms that differ only in the length of their planning horizons: those with short and long planning horizons. The distribution of agents across these types is time-varying and governed by a Markov process. While the central bank observes the current distribution of agents, it is uncertain about this distribution in 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 uses

[5] Similar to our paper, Kimura and Kurozumi (2007) also provide microfoundations for sluggish inflation dynamics. In their case, this sluggishness 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 learn about events outside of their planning horizons using past data.

its knowledge of the economy and sources of uncertainty to make interest-rate decisions over an infinite horizon.

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 that does not face this uncertainty. This allows us to isolate the effects of this particular source of uncertainty on economic outcomes.

2.1 Finite-Horizon Planning

In this subsection, we provide an abbreviated discussion of the NK-FHP model that we study, highlighting the source of model uncertainty faced by a central bank 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, $k \in \{k_0, k_1\}$. Although households and firms are infinitely lived, they do not formulate an infinite-horizon state-contingent plan. Instead, households and firms with planning horizon k make decisions at time t by 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 under the assumption that all other agents have the same k-period planning horizon as themselves. Accordingly, agents use a subjective expectations operator, $\mathbb{E}_t^k$, which reflects planning $\mathbf{k}$ periods ahead, taking into account the probabilistic evolution of events only through dates $t \leq \tau \leq t + \mathbf{k}$.

Because of their finite planning, an agent's plans at time $t+1$ will not generally coincide with the decisions it makes at time $t+1$, since this would imply that an agent making decisions at date t evaluate contingencies through period $t+\mathbf{k}+1$ rather than truncating its planning at period $t+\mathbf{k}$. Hence, an agent's expectations are not rational, which would require that an agent uses the model's structural relationships to evaluate contingencies in periods $\tau \geq t+\mathbf{k}+1$. Instead, an agent at date t makes plans in period $t+1$ by taking into account structural relationships and contingencies only $\mathbf{k}-1$ periods into the future. Similarly, in planning period $t+2$, an agent makes plans in period $t+2$ only taking into account structural relationships and contingencies through $\mathbf{k}-2$ periods into the future, and so on over the course of its planning horizon.

While households and firms use their knowledge of structural relationships and contingencies within their planning horizons, evaluating expectations of aggregate prices and quantities within those planning horizons requires forming beliefs about the planning horizons of other agents. In order for a household or firm to avoid considering what the model's structural relationships imply for states beyond its k-period planning horizon, it is assumed—as in Woodford (2018)—that an agent believes that other households, firms, and the central bank also have a k-period planning horizon. More generally, when making plans at date $t+k-j$, where j indexes the number of periods left in an agent's planning horizon, an agent assumes that the central bank and other households or firms will by making decisions as if they only had a j-period planning horizon. This assumption introduces a second way in which expectations in the model depart from rational expectations.

Because households and firms plan only over a k-period horizon, it is convenient to define π_t^k, y_t^k, and i_t^k as the model-consistent solutions for inflation, the output gap, and the policy rate (expressed in log-deviations from steady state) in a model where all decision-makers have planning horizons of length k. More generally, for an endogenous model variable Z_{t+k-j}, the following relationship holds:

$$\mathbb{E}_t^k Z_{t+k-j} = E_t Z_{t+k-j}^j. \tag{1}$$

Expression (1) provides a mapping between the subjective expectations operator of an agent with a k-period planning horizon and the model-consistent expectations operator. It reflects that agents formulate their plans in period $t+k-j$ using their full knowledge of the model's structural equations and contingencies under the (counterfactual) assumption that the central bank and other agents in the model have j periods left in their planning horizons–just like themselves. Thus, while E_t is the model-consistent expectations operator, the variable Z_{t+k-j}^j still reflects an agent's limited planning horizon and subjective beliefs regarding the planning horizons of other agents.

Woodford (2018) introduces finite-horizon planning into a standard NK model and shows that at date $\tau = t+k-j$ of their k-period planning horizon the beliefs of households and firms satisfy:

$$\pi_\tau^j = \beta E_\tau \pi_{\tau+1}^{j-1} + \kappa y_\tau^j + u_\tau \tag{2}$$

$$y_\tau^j = E_\tau y_{\tau+1}^{j-1} - \sigma\left(i_\tau^j - E_\tau \pi_{\tau+1}^{j-1} - r_\tau^e\right), \tag{3}$$

for $j = 1, 2, \ldots, k$ where j indexes the number of periods left in an agents' planning horizon for decisions made at date t. The variables π_τ^j, y_τ^j, and i_τ^j represent agents' beliefs for aggregate inflation, spending, and the policy rate at date τ of their planning horizon.

In the NK model, we have abstracted 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. In equation (2), κ determines the sensitivity of inflation to the output gap and depends on the Calvo parameter determining the frequency of price setting, while in equation (3), σ is the inverse of a household's relative risk aversion. The variables u_τ and r_τ^e represent exogenous aggregate shocks to firms' pricing and households' spending decisions, respectively.[6] These shocks are assumed to follow AR(1) processes:

$$u_t = \rho_u u_{t-1} + e_{ut}$$
$$r_t^e = \rho_r r_{t-1}^e + e_{rt},$$

with the parameters ρ_u and ρ_r, between zero and one, measuring the persistence of the shocks. The innovations to these shocks are assumed to be *i.i.d.* normal random variables with standard deviations, σ_u and σ_r, respectively. Agents are assumed to have a perfect understanding of the evolution of these shocks, and there is no need to index these variables by the length of time left

[6]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.

in an agent's planning horizon.

Equation (2) reflects the beliefs in planning period $t+k-j$ of a k-horizon firm who has the opportunity to change their prices at date t. Likewise, equation (3) represents household beliefs about their spending decisions. These equations can be iterated forward to determine an agent's beliefs about aggregate inflation and spending at time t:

$$\pi_t^k = E_t \sum_{i=0}^{k-1} \beta^i [\kappa y_{t+i}^{k-i} + u_{t+i}] + \beta^k E_t \pi_{t+k}^0 \tag{4}$$

$$y_t^k = -\sigma E_t \sum_{i=0}^{k-1} [i_{t+i}^{k-i} - \pi_{t+i+1}^{k-i} - r_{t+i}^e] + E_t y_{t+k}^0, \tag{5}$$

Equations (4) and (5) determine a k-horizon agent's beliefs about aggregate output and inflation at time t contingent on their expectations for future exogenous shocks and the future path of policy. In Woodford (2018), a k-period household or firm use a Taylor rule to form beliefs about the policy-rate path, which —along with equations (4) and (5)—determines the k-horizon agent's beliefs for π_t^k, y_t^k, and i_t^k.

We depart from this assumption, and instead follow CGG by specifying that the central bank choose the policy rate to minimize expected discounted losses involving squared deviations of inflation from a central bank's target and of output from potential. As discussed later, a k-horizon household or firm believes that the policy-rate path is set by a central bank that optimizes such a loss function over a k-period horizon. The central bank's k-horizon optimization problem along with equations (4) and (5) then determine the model-consistent beliefs of k-horizon households and firms for aggregate inflation, aggregate spending, and the policy rate in a model where the central bank and private sector agents all have planning horizons of length k.

2.2 Heterogeneous Planning and Model Uncertainty

Equations (4) and (5) represent a household or firm's beliefs regarding aggregate inflation and spending at time t. However, because of their limited planning horizons and beliefs about the planning horizons of other agents, households and firms make systematic errors in forecasting aggregate variables.[7] In particular, aggregate inflation and aggregate spending realized at time t satisfy:

$$\pi_t = \omega_t \pi_t^{k_0} + (1-\omega_t)\pi_t^{k_1} \tag{6}$$

$$y_t = \omega_t y_t^{k_0} + (1-\omega_t) y_t^{k_1}, \tag{7}$$

where ω_t is the share of households and firms with horizon k_0 at date t. Aggregate inflation and output reflect the presence of heterogeneous planning horizons, while each FHP household or firm

[7] As discussed in Coibion and Gorodnichenko (2015) and Angeletos et al. (2020) among others, there is considerable evidence from surveys that forecast errors of macroeconomic variables are systematic and predictable. Gust et al. (2024) study the predictability of inflation forecast errors in the NK-FHP model and show that the model can account for the initial underreaction and subsequent overreaction documented in Angeletos et al. (2020).

assumes that all other agents have the same planning horizon as themselves in order to determine their individual spending and pricing decisions.

To model uncertainty in planning horizons, we allow for exogenous fluctuations in ω_t, by specifying it evolve according to a two-state Markov process.[8] In particular, let $\omega_t \equiv \omega(m_t)$ where m_t is random variable with $m_t \in \{0, 1\}$. These values for m_t correspond to different regimes or modes for the distribution of planning horizons in the economy, and its transition probabilities satisfy:

$$P_{mn} = \Pr\{m_{t+1} = n | m_t = m\}, \qquad m, n \in \{0, 1\} \tag{8}$$

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

Because of movements in ω_t, aggregate inflation and output reflect uncertainty arising from fluctuations in the average planning horizon in the economy. 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 remain linear.

2.3 Longer-Run Learning

For households and firms with a finite horizon k, equations (4) and (5) indicate that beliefs about aggregate inflation and spending at time t depend on their expectation about inflation and spending at the end of their planning horizons. As shown in Woodford (2018), these beliefs satisfy:

$$\pi^0_{t+k} = \beta v_{pt} + \kappa y^0_{t+k} + u_{t+k} \tag{9}$$

$$y^0_{t+k} = v_{ht} - \sigma(i^0_{t+k} - r^e_{t+k}), \tag{10}$$

where v_{pt} and v_{ht} represent 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 households and firms 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.[9] Specifically, households and

[8]Shifts in the population-average planning horizon occur exogenously along the extensive margin, and in making their time t forward-looking plans an individual household or firm believes the length of their planning horizon is fixed. This assumption keeps the analysis of optimal policy under uncertain planning tractable. We leave to future work an extension of the analysis to an environment in which households and firms face cognitive costs to changing their planning horizons that would allow them to evolve endogenously in response to economic developments.

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

firms learn and update v_{pt} and v_{ht} according to:

$$v_{pt+1} = (1 - \gamma_p)v_{pt} + \gamma_p \pi_t \tag{11}$$

$$v_{ht+1} = (1 - \gamma_h)v_{ht} + \gamma_h(y_t + \sigma \pi_t), \tag{12}$$

where, π_t and y_t denote aggregate inflation and output (in log deviation from steady state), respectively.[10]

Equations (11) and (12) reflect that each firm uses a continuation value function that is a population average across firms (v_{pt}), while each household uses one that is a population average across households (v_{ht}). Accordingly, firms and households use past data on average, aggregate variables to update their value functions. The parameters γ_p and γ_h determine how quickly firms and households update their beliefs in response to incoming data. Overall, households and firms make relatively sophisticated plans and forecasts within their planning horizons. 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 additional parameters, γ_p and γ_h, it has both theoretical and empirical benefits. On the theoretical side, it gives the FHP model desirable properties in response to long-lasting changes in policy or economic fundamentals. Without learning, v_{pt} and v_{ht} would be fixed at their steady state values. Accordingly, in that case, they would not change in response to long-lasting economic events and households and firms would effectively continue to use outdated value functions if, for example, there were a permanent change in a central bank's inflation target. In contrast, with learning, the value functions of households and firms would evolve and eventually would fully reflect a change in a central bank's inflation target. On the empirical side, Gust et al. (2022) show that longer-run learning allows the FHP model to generate substantial aggregate persistence and fit macroeconomic data without relying on additional features such as habit persistence or price contracts indexed to lagged inflation.

Inflation Scares. Firms' beliefs about events outside their planning horizon play a particularly important role in our analysis. In particular, v_{pt} can be interpreted as a firm's longer-run beliefs about inflation, since firms update their value functions v_{pt} based on past inflation. Thus, agents' longer-run beliefs about inflation evolve slowly, as equation (11) implies:

$$v_{pt} = \gamma_p \sum_{i=0}^{t-1} (1 - \gamma_p)^i \pi_{t-1-i}, \tag{13}$$

where the parameter γ_p determines the weight placed on recent lags of inflation relative to distant lags. Because v_{pt} depends on past deviations of inflation from the central bank's target, agents' longer-run inflation beliefs will evolve endogenously in response to the economy's shocks and can potentially drift away from zero—the level consistent with their beliefs being anchored at the

[10]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.

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.

3 Optimal Monetary Policy

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

$$\mathcal{L}_t = \frac{1}{2} E_t \sum_{\tau=t}^{\infty} \beta^{\tau-t} \left[\pi_\tau^2 + \lambda y_\tau^2\right], \tag{14}$$

where the parameter λ is the central bank's relative weight on fluctuations in the output gap.[11] Unlike households and firms, the central bank has rational expectations and formulates policy over an infinite horizon using its full knowledge of private-sector behavior. The central bank is assumed to re-optimize each period in a time-consistent fashion.

3.1 Private Sector (Policy) Beliefs

Before describing the central bank's problem further, it is useful to discuss the beliefs of households and firms about future monetary policy. Because households and firms can not formulate plans over an infinite horizon, their perceptions about future monetary policy differ from the policy that the central bank implements. In particular, they perceive that the central bank, like themselves, has a k-period planning horizon. This belief shapes their expectations of future monetary policy.

For current policy, household and firm observe the policy decision at date t so $i_t^k = i_t$ for $k \in \{k_0, k_1\}$. For their forecasts of policy over the remainder of their k-period horizon, they perceive that the central bank chooses i_τ^j for $j = 1, 2, \ldots, k-1$ and $\tau = t + k - j$ to satisfy a sequence of minimization problems:

$$\mathcal{W}_\tau^j(s_\tau) = \min_{i_\tau^j} \frac{1}{2} \left[(\pi_\tau^j)^2 + \lambda(y_\tau^j)^2\right] + \beta E_\tau \left[\mathcal{W}_{\tau+1}^{j-1}(s_{\tau+1})\right] \tag{15}$$

subject to equations (2) and (3). In equation (15), $s_\tau = (u_\tau, r_\tau^e)$ denotes the vector of shocks. Private-sector agents perceive that 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.

[11] 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; thus, we use deviations of output from steady state and the output gap, interchangeably.

At the final period of their planning horizon ($j = 0$), households and firms perceive that the central bank solves the following problem:

$$\mathcal{W}^0_{t+k}(s_{t+k}) = \min_{i^0_{t+k}} \left[(\pi^0_{t+k})^2 + \lambda(y^0_{t+k})^2\right] \tag{16}$$

subject to equations (9) and (10), taking as given v_{pt} and v_{ht}. Households and firms believe that the central bank, like themselves, takes the continuation value functions as fixed at date t. However, when agents have a finite planning horizon, this perception is incorrect, as the central bank has an infinite horizon and takes into account the evolution of the value functions over time. In the special case where agents have an infinite planning horizon ($k_0, k_1 \to \infty$), their perceptions are correct, and there will be no difference between actual and perceived policy.

Iterating backward from the problem at the end of agents' planning horizons, the first order conditions from the sequence of problems imply that private-sector agents perceive the central bank to follow the targeting rule:

$$y^j_{t+k-j} = -\frac{\kappa}{\lambda}\pi^j_{t+k-j} \tag{17}$$

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

Agents' Perceived LAW. The agents' perceived targeting rule is similar to the targeting rule under optimal discretion in the canonical NK model discussed in CGG, 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 in future states *over their planning horizon*: whenever inflation is above target, they believe the central bank will contract demand below capacity (by raising the interest rate), and vice versa when it is below target. The central bank is less willing to combat high inflation by leaning against the wind and suffering large output losses the larger is their preference for output stabilization (i.e., λ is larger) and is more willing to combat high inflation the steeper is the slope of the Phillips curve (i.e., κ is larger).

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 (2) and (9) implies that expected inflation by agents with horizon k is given by:

$$E_t\Pi^{k-1}_{t+1}(v_{pt}, u_t) \equiv \underbrace{\left[\frac{\lambda\rho_u}{\lambda+\kappa^2}\sum_{i=0}^{k-1}(\frac{\beta\lambda\rho_u}{\lambda+\kappa^2})^i\right]u_t}_{\text{Cyclical}} + \underbrace{a_p(k)v_{pt}}_{\text{Trend}} \tag{18}$$

where the use of $\Pi^{k-1}_{t+1}(v_{pt}, u_t)$ instead of π^{k-1}_{t+1} emphasizes that expected inflation is a function of the aggregate supply shock as well as the firm's continuation 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.[12] Intuitively, optimal policy is perceived as fully offsetting the effects of changes in aggregate demand on inflation that occur through movements in r_t^e and v_{ht}. Hence, expected inflation does not depend on r_t^e or v_{ht}, but does depend on the cost-push shock and price-setting firms' continuation value function, v_{pt}.

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 the 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 *i.i.d.* ($\rho_u = 0$) or if the 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 steeper 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 a firm's continuation value function—that is, its longer-run inflationary beliefs. As discussed earlier, these longer-run beliefs are backward-looking, which makes expected inflation sticky. In addition, if inflation is running above the central bank's target, it can lead to an inflation scare, as it results in a large value of v_{pt} that pushes trend inflation away from the central bank's target. The extent to which this occurs depends on $a_p(k)$, which governs the marginal effect of a 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. With reduced feedback, the likelihood of an inflation scare is also lower. A longer planning horizon also reduces this marginal effect: As $k_0, k_1 \to \infty$, households 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 no longer arise. Finally, note that, if the inflation-output trade off becomes 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 when choosing current policy. This rational behavior creates a key distinction between the optimal policy that is actually *implemented* and the 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 that is implemented includes a forward-looking component reflecting the

[12] In the appendix, we also characterize agents' beliefs about expected output and the policy rate. As shown there, agents with a k-period horizon believe that the future output gap is given by:

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

.

dependence of agents' longer-run inflation expectations on past inflation. In the NK-FHP model where households and firms learn about long-run inflation over time, optimal policy seeks to avoid the possibility of an inflation scare where longer-run inflation expectations drift away from the central bank's inflation target.

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

$$W(v_{pt}, s_t, m_t) = \min_{i_t} \frac{1}{2}\left[\Pi_t^2 + \lambda \mathbf{Y}_t^2\right] + \beta \sum_{n=0}^{1} \Pr(m_{t+1} = n | m_t) \int_{s_{t+1}} W(v_{pt+1}, s_{t+1}, n) f(s_{t+1} \mid s_t) ds_{t+1} \tag{19}$$

where the function $f(s_{t+1} \mid s_t)$ denotes the conditional density of the shocks to aggregate supply and demand. In addition, $\Pr(m_{t+1} = n \mid m_t)$ 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 respective targets. The central bank chooses the policy rate i_t, taking as given the private sector's equilibrium conditions and the functions determining agents' beliefs regarding future inflation and output. These equilibrium conditions can be written as:

$$\Pi_t = \beta E_t \Pi_{t+1} + \kappa \mathbf{Y}_t + u_t \tag{20}$$

$$\mathbf{Y}_t = E_t \mathbf{Y}_{t+1} - \sigma\left(i_t - E_t \Pi_{t+1} - r_t^e\right) \tag{21}$$

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

Expressions (20) and (21) resemble those that determine inflation and the output gap in the canonical NK model except that the expectations, $E_t \Pi_{t+1}$, 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' forecasts of future inflation and the output gap:

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

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

There are several noteworthy features about a central bank's optimization problem. First, in setting the optimal time-consistent policy, the central bank takes the expectational functions as given for $k \in \{k_0, k_1\}$. However, it recognizes that its interest-rate decision affects agents' beliefs indirectly since these expectational functions depend on v_{pt}, which in turn influences future outcomes and the central bank's future losses. Accordingly, when setting the current policy rate i_t, the central bank takes into account that v_{pt+1} depends on current inflation, which is itself inclucenced by today's policy decision. 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 future distribution of agents in the future and thus uses the probabilities across modes to weight future losses. Third, unlike the other macroeconomic shocks, fluctuations in ω_t are not additive. Instead, they interact multiplicatively with the model's endogenous variables. Accordingly, certainty equivalence does not hold. Nonetheless, the central bank's problem remains linear-quadratic *conditional on* ω_t, preserving tractability and allowing for a treatment of model uncertainty that follows the approach in Svensson and Williams (2005).

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

$$\pi_t + \gamma_p \beta E_t W_{pt+1} = -\frac{\lambda}{\kappa} y_t \tag{22}$$

$$W_{pt} = \beta a_{pt} \pi_t + \beta \left[1 - \gamma_p (1 - \beta a_{pt})\right] E_t W_{pt+1} \tag{23}$$

where aggregate inflation, π_t, and the output gap, y_t, are given by expressions (6) and (7), respectively. The term $a_{pt} = \omega(m_t) a_p(k_0) + (1 - \omega(m_t)) a_p(k_1)$ captures the marginal effect of a change in v_{pt} on a weighted average of agents' expectation of next period's inflation. The function, $W_{pt} \equiv W_p(v_{pt}, s_t, m_t)$ denotes the marginal effect of v_{pt} on the central bank's loss function. The expectation $E_t W_{pt+1}$ is given by:

$$E_t W_{pt+1} = \sum_{n=0}^{1} \Pr(m_{t+1} = n \mid m_t) \int_{s_{t+1}} W_p(v_{pt+1}, s_{t+1}, n) f(s_{t+1} \mid s_t) ds_{t+1}. \tag{24}$$

Expression (22) extends the LAW principle of CGG to an environment with finite-horizon planning and longer-run learning. The first component on the right-hand side of the targeting rule reflects the static LAW principle derived in CGG: if inflation is above the target ($\pi_t > 0$) due to the cost-push shock, optimal policy calls for a contractionary response that 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 implied by the CGG targeting rule, since it introduces an additional term, $E_t W_{pt+1}$. This term captures the marginal increase in the central bank's expected discounted losses arising from an increase in agents' longer-run inflation beliefs, v_{pt}. An increase in v_{pt} can be interpreted as an inflation scare in which both actual inflation and longer-run inflation expectations remain persistently above the central bank's target. The optimal response to such a scare is reflected in the condition $E_t W_{pt+1} > 0$. It implies that the central bank should act preemptively by pushing current output below potential ($y_t < 0$), thereby leaning more aggressively against the possibility of future inflation. This additional policy tightening puts further downward pressure on current inflation and helps anchor longer-run inflation expectations, mitigating the risk of an inflation scare.

A central bank's uncertainty about expectations formation also affects how they respond to inflation scares. According to equations (23) and (24), 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 (23) 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 when agents do not update their longer-run beliefs about inflation (i.e., $\gamma_p = 0$). In the first case, the model corresponds to the canonical NK model in which longer-run inflation beliefs (v_{pt}) become irrelevant and inflation scares are not possible.[13] As a result, the central bank has no incentive to act preemptively and optimal time-consistent policy satisfies the static LAW principle of CGG: $\pi_t = -\frac{\lambda}{\kappa} y_t$.

Similarly, when firms do not update their beliefs in response to past inflation (i.e, $\gamma_p = 0$), v_{pt} is constant and there is no variation in agents' longer-run inflationary beliefs. In this case, equation (22) reduces to the CGG rule.[14]

Fixed-Planning Economy. To isolate the role of uncertainty in expectation formation, we compare our baseline model to a version where there is no uncertainty about agents' planning horizons. Specifically, we assume $\omega_t = \bar{\omega}$ for all $t \geq 0$, implying a constant distribution of horizon lengths. We call this the fixed planning (FP) economy, and in this economy the population-average planning horizon is given by $k^{FP} = \bar{\omega} k_0 + (1-\bar{\omega}) k_1$. In the FP economy, certainty-equivalence holds, and we set k^{FP} to be the same as the unconditional population-average planning horizon in the baseline economy in which ω_t fluctuates. In the FP economy, a central bank's optimal targeting rule becomes:

$$\pi_t^{FP} + \gamma_p \beta E_t W_{pt+1}^{FP} = -\frac{\lambda}{\kappa} y_t^{FP} \tag{25}$$

$$W_{pt}^{FP} = \beta \bar{a}_p \pi_t^{FP} + \beta\left[1-\gamma_p\left(1+\beta\gamma_p\bar{a}_p\right)\right] E_t W_{pt+1}^{FP}, \tag{26}$$

where $E_t W_{pt+1}^{FP} = \int_{u_{t+1}} W_p^{FP}(v_{pt+1}, u_{t+1}) f(u_{t+1}|u_t) du_{t+1}$ and $\bar{a}_p = \omega a_p(k_0) + (1-\omega) a_p(k_1)$.

The role of uncertainty about agents' planning horizon can be seen by comparing expressions (22) and (23) with expressions (25) and (26). First, the marginal effect of v_{pt} on expected inflation next period, $\bar{a}_p$, is constant in the FP economy. Second, $E_t W_{pt+1}^{FP}$ depends only on the additive shock u_t, not on the multiplicative shock m_t. Finally, as $k_0, k_1 \to \infty$, inflation scares disappear and the optimal policy again satisfies the static LAW principle: $\pi_t^{FP} = -\frac{\lambda}{\kappa} y_t^{FP}$.

4 Results

In this section, we first present evidence of time variation in the length of the planning horizons of households and firms. Second, we use data from the FOMC's Survey of Economic Projections

[13] The appendix shows that the decision rule for optimal inflation converges to the rational expectations solution as $k_0, k_1 \to \infty$. As shown there, the decision rule for optimal inflation depends positively on both v_{pt} and u_t. As $k_0, k_1 \to \infty$, the response of optimal inflation to v_{pt} and u_t converge to their rational expectations' limits.

[14] Although our focus is on a rational central bank, the CGG targeting rule would also be optimal for a central bank with finite-planning horizon. Such a central bank, like households and firms, would treat v_{pt} and v_{ht} as fixed.

Figure 1: Estimated Fraction of Short-Horizon Planners (ω_t), 1966-2008

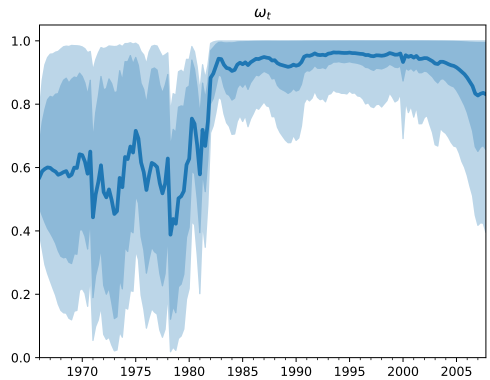

NOTE: Figure shows the posterior mean (solid line) and 68 percent (dark-shaded region) and 95 percent (light shaded region) uncertainty bands for ω_t from the estimated NK-FHP model with uncertain planning horizons.

(SEP) from March 2023 to construct a baseline path for aggregate demand and supply shocks from 2023Q1 to 2027Q4. We then conduct stochastic simulations around this baseline to illustrate the potential gains from optimal policy when there is a risk of an inflation scares due to elevated inflation.

4.1 Evidence on Fluctuations in Planning Horizons

To provide evidence on time-varying planning horizons, we estimate the NK-FHP model using a Bayesian, full-information likelihood-based approach with U.S. data on output growth, inflation, and nominal interest rates from 1966:Q1 through 2007:Q4, a time period marked by notable, low-frequency variation in both inflation and output growth. The details of our estimation strategy, including the estimated policy rule, follow Gust et al. (2024) and are described in the appendix.

A key difference from our earlier work is that we now allow for fluctuations in the fraction of short-horizon agents, ω_t. To do so, we set $k_0 = 4$ and $k_1 = 32$. These values are chosen to be consistent with the empirical evidence in Gust et al. (2024), which suggests that most agents have short planning horizons of a year or less, while still allowing for meaningful differences in planning and expectations formation across the two types of households and firms. This speification ensures that the model captures both the tendency for households and firms to evaluate contingencies and plan in a sophisticated manner only a few quarters ahead, as well as periods when they may do so over much longer horizons.

Figure 1 shows the time variation in the estimated fraction of short horizon agents, ω_t. The time series exhibits two distinct regimes. Prior to the early 1980s, ω_t hovered around 0.5, indicating that the weights assigned to $k_0 = 4$ and $k_1 = 32$ are approximately equal in the economy. During this period, planning horizons were relatively long, on average, at a time when inflation was high and volatile. In the early 1980s, as inflation fell, the estimated population-average planning horizon declines significantly, as implied by the rise in the fraction of short-horizon planners. In the post-1980s period, this fraction increases, and is more precisely estimated, and hovers just below 1. This suggests that economic dynamics during this period were dominated by agents with planning horizons of roughly a year—consistent with the estimates in Gust et al. (2024). Overall, this analysis supports the presence of time variation in agents' planning horizons, and suggests that longer planning horizons were more prevalent when inflation became high and volatile.

Additional evidence suggestive of fluctuations in planning horizons can be found in Coibion and Gorodnichenko (2015). They show that forecast errors of macroeconomic variables can be predicted by forecast revisions and that this relationship fluctuates over time. Gust et al. (2024) show that inflation forecasts systematically underreact to forecast revisions in the short run in the NK-FHP model, consistent with the predictability regressions of Coibion and Gorodnichenko (2015). Interestingly, the fluctuations in ω_t shown in Figure 1 would generate fluctuations in the relationship between forecast errors and forecast revisions for inflation and output, broadly in line with the results of Coibion and Gorodnichenko (2015).[15] In addition, Korenok et al. (2023) and Pfäuti (2025) find that the level of the public's attentiveness to inflation news is time varying and increases with the level of inflation. This evidence is also consistent with the fluctuations in ω_t shown in Figure 1, assuming households and firms that engage in more sophisticated, forward-looking planning are also more likely to be attentive to inflation news.

4.2 Parameter Values

To calibrate the discrete, two-state Markov process governing agents' uncertain planning horizon, we use the properties of the estimated time series behavior of ω_t shown in Figure 1. This series displays two distinct regimes, and in line with those regimes, we set $\omega(0) = 0.98$ and $\omega(1) = 0.45$ with $k_0 = 4$ and $k_1 = 32$. This calibration implies that the population-average planning horizon is just over four quarters in the first regime and about $4\frac{1}{2}$ years in the second regime.

For the transition probabilities of the Markov process, we use the estimated autocorrelation of ω_t in Figure 1. Specifically, our estimates imply an autocorrelation of 0.96 for ω_t, and as discussed in Kopecky and Suen (2010), setting the transition probabilities of the matrix P such that $P_{mm} = \frac{1+0.96}{2} = 0.98$ for $m = 0, 1$ effectively mimics the persistence of a highly autocorrelated continuous processes such as ω_t. With $P_{00} = P_{11}$, the ergodic (unconditional) probability of either regime is 50% ($\bar{p}_0 = \bar{p}_1 = 0.5$), and the unconditional population-average planning horizon is just under three years. To understand the role of uncertain planning, we compare the outcomes of the NK-FHP model with fluctuations in ω_t to those of the fixed-planning economy where $\bar{\omega} = 0.72$. This value implies that the unconditional population-average planning horizon in the fixed-planning economy is the same as in the economy with uncertain planning.

The remaining model parameters were chosen as follows. We set $\sigma = 2.5$, $\gamma_p = 0.12$, $\gamma_h = 0.56$, $\rho_r = 0.87$, $\rho_u = 0.74$, and $\sigma_r = 0.34$, all of which are near their estimated mean values shown in the appendix. We chose $\kappa = 0.03$, which is higher than its mean estimate, but we later also report results for $\kappa = 0.01$, its mean estimate. This somewhat higher value helps illustrate the role of

[15] As shown in the appendix, the NK-FHP model displays the same systematic, positive relationship between forecast errors and forecast revisions as documented in Coibion and Gorodnichenko (2015). As the length of planning horizons increases in the model, the relationship between these two variables weakens because forecast errors become less predictable.

uncertain planning in the simulations that we show. Similarly, we set $\sigma_u = 0.56$, which implies that the aggregate supply shock accounts for 50% of inflation volatility under optimal policy, with the remaining volatility explained by fluctuations in planning horizons. This shock size is higher than its estimated value but is chosen for illustrative purposes; we also report results for its mean estimate.

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

4.3 SEP-Consistent Baseline

We study inflation scares in the NK-FHP model using the March 2023 SEP. At the time, inflation was well above target, and the risk of an inflation scare was a prominent concern among policymakers. 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). To use the March 2023 SEP in conjunction with the NK-FHP model, we construct a path of aggregate demand and supply shocks that is consistent with the median projections of inflation, the unemployment rate, and the federal funds rate from the SEP. The appendix describes how we use the NK-FHP model to recover these shocks. We run stochastic simulations of optimal policy around the aggregate demand and supply shock implied by our SEP-consistent baseline.

4.4 The Transmission of Shocks in the NK-FHP Model

Before proceeding, it is useful to examine how the transmission of shocks differs in the NK-FHP model with fixed planning from the canonical NK model. Figure 2 shows the effects of the aggregate supply shocks in which all of agents have planning horizons of 4 quarters ($\bar{\omega} = 1$) and when all agents have planning horizons of 32 quarters ($\bar{\omega} = 0$). Compared to the canonical NK model, the effects of the shocks on inflation, output, and the policy rate are noticeably more persistent, particularly when all agents have short horizons.

This persistence reflects that thouseholds' and firms' beliefs about events outside their planning horizons depend on past economic outcomes, and this dependence intensifies as planning horizons shorten. It also implies that the trend component of expected inflation (lower right panel) remains persistently above steady state in response to the adverse aggregate supply shock, and that trend inflation rises noticeably more when agents have short planning horizons than when they are long.[16]

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

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

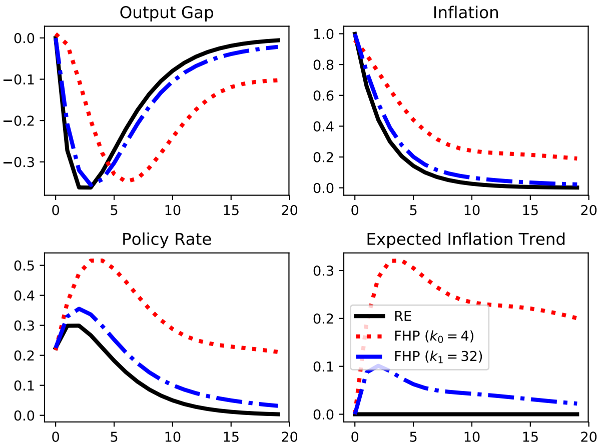

NOTE: The figure shows the effect of a shock to aggregate supply that increases annualized inflation by 1 percentage point in the NK-FHP model and the canonical NK model with rational expectations (dark line). For the NK-FHP model, the red dotted line shows the results for $\bar{\omega} = 1$ when all agents have planning horizons of one year ($k_0 = 4$) and the blue dashed-dotted line shows the results for $\bar{\omega} = 0$ when all agents have planning horizons of eight years ($k_1 = 32$). In the simulations, monetary policy is assumed to follow a Taylor rule with interest-rate smoothing. See the appendix for details.

This rise in the trend component causes the inflation rate to remain above its steady state level of 2 percent for an extended period of time in the NK-FHP model with four-quarter planning. With inflation rising persistently, the policy rate rises higher than its response in the canonical NK model and stays elevated for longer. As a result, output falls and remains persistently below potential.

4.5 Optimal Policy under Uncertain Planning

We now turn to illustrating the implications of uncertain planning for optimal policy. To do so, we compare stochastic simulations around the aggregate demand and supply shocks underlying the 2023 SEP baseline (shown in the appendix). The simulations of optimal policy under uncertain and fixed planning differ only because of the volatility of ω_t in the uncertain planning case, as the unconditional mean of the population-average planning horizon is the same in both cases.[17] Figure 3 compares the distribution of outcomes for the output gap, inflation, the policy rate, and the trend component of expected inflation in 2024Q1 under optimal policy with fixed and uncertain planning. For comparison, the figure also shows the distribution of outcomes in the canonical NK model in which the CGG targeting rule is optimal. In the canonical model and the NK-FHP model with fixed planning, the distributions of outcomes are symmetric, since these models are linear and the shocks to aggregate demand and supply are normally distributed. Figure 3 shows that the distribution of inflation outcomes in the NK-FHP model lies well above those of the canonical NK model. This difference reflects that, in the NK-FHP model, agents believe that trend inflation will remain persistently elevated. As a consequence, the distribution of outcomes for the policy rate is notably higher in 2024Q1 in the NK-FHP model than in the canonical model. This tighter policy

[17] We start the simulations of optimal policy under uncertain planning from the ergodic probabilities; accordingly, at each point in time, there is an equal chance of either mode occurring, and the mean of the population-average planning horizon under uncertain planning is the same as under fixed planning at each point in time.

Figure 3: The Distribution of Outcomes under Optimal Discretion in 2024Q1

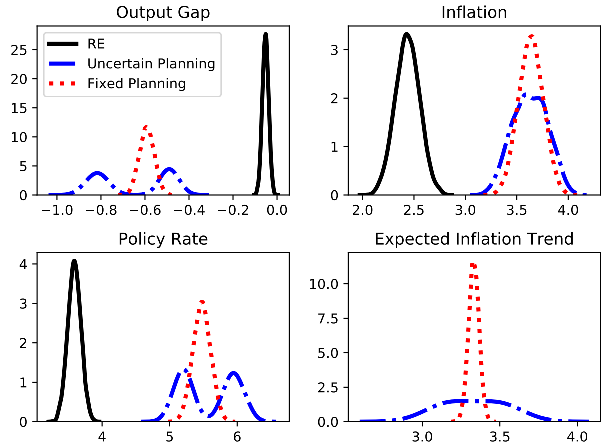

NOTE: The figure shows the distribution of responses under optimal discretionary policies in 2024Q1 using 50,000 draws of shocks centered around the aggregate demand and supply shocks associated with the SEP-consistent baseline.

puts trend inflation on a downward trajectory over time but comes at the cost of a persistently larger output gap than in the canonical NK model.[18]

Figure 3 also compares the distribution of outcomes in 2024Q1 under uncertain planning to fixed planning. With uncertain planning, the distribution of ω_t is symmetric around $\bar{\omega}$, but the distributions of outcomes for endogenous variables are asymmetric, reflecting that the ω_t enters the model multiplicatively. As shown in the figure, the policy rate is skewed to the upside, while the output gap is skewed to the downside. This skewness reflects that a central bank faces a more unfavorable tradeoff between stabilizing the output gap or inflation in states of the world where inflation scares can be severe (i.e., when inflation is above target and there is a high fraction of short-horizon agents.) In such situations, a more aggressive policy response is necessary to bring inflation down; however, it comes at the cost of larger deviations of output from potential.[19]

Policy Tradeoff Frontier. To illustrate the more unfavorable tradeoff for policymakers under uncertain planning, Figure 4 traces out a central bank's tradeoff between inflation and the output gap under both fixed and uncertain planning. It also does so for the canonical NK model with rational expectations. As shown there, the policy tradeoff frontier lies much lower in the canonical NK model than in the NK-FHP models with either fixed or uncertain planning, because inflation scares are precluded in the canonical NK model. The policy tradeoff frontier also shifts further to the right under uncertain planning than under fixed planning, implying a less favorable tradeoff between inflation and output stabilization. This shift reflects the increased risk of a more severe inflation scare in an uncertain environment.

Gains from Guarding Against Inflation Scares. In an environment in which inflation is above target and inflation scares are possible, the gains from guarding against them can be

[18] In the appendix, we show how the corresponding mean outcomes for the output gap, inflation, the policy rate, and trend inflation evolve over time.

[19] As shown in the appendix, the extent to which the distributions of output and inflation under optimal policy are skewed depends on the relative importance of the aggregate supply shock (u_t) to fluctuations in agents' planning horizons (ω_t). Less volatility in agents' planning horizons reduces the skewness of these distributions.

Figure 4: The Policy Tradeoff Frontier Under Optimal Discretion

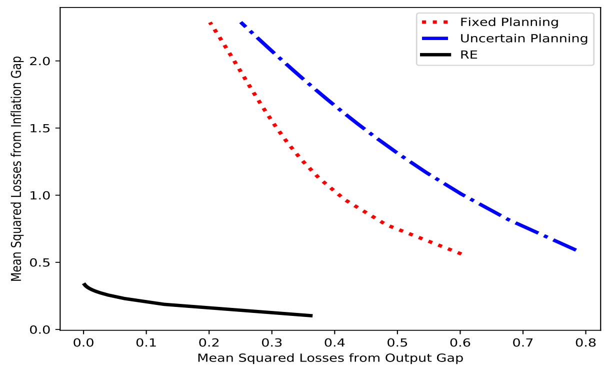

NOTE: For the two models, each point along the policy tradeoff frontier is constructed by computing the mean-squared deviations of the output and inflation gaps from stochastic simulations around the aggregate demand and supply shocks associated with the SEP-consistent baseline for a given value of λ. The black line shows the tradeoff in the fixed planning economy, and the dashed-dotted blue line shows the tradeoff in the uncertain planning economy. The solid line shows optimal policy under rational expectations in the canonical NK model.

Table 1: The Gains of Guarding Against Inflation Scares

	Contribution of ω_t to Inflation Volatility		Additional Losses (Percent)		Policy Rate Difference (basis points)	
σ_ω	$\kappa = 0.03$	$\kappa = 0.01$	$\kappa = 0.03$	$\kappa = 0.01$	$\kappa = 0.03$	$\kappa = 0.01$
0	0	0	22.4%	31.7%	66	102
0.08	9%	1%	23.2%	32.0%	67	103
0.13	21%	2.6%	24.6%	32.5%	69	103
0.23	44%	7.3%	30.1%	34.3%	74	105
0.26	50%	9.3%	32.9%	35.2%	76	106
Addendum (Mean Parameter Estimates):						
0.26		72%		35.5%		106

NOTE: The table shows the average percent increase in a central bank's losses over the 2023-2027 period from adopting the LAW policy of CGG, which is optimal in the canonical NK model. It also shows the average policy rate increase in 2023 under optimal policy relative to their targeting rule. 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. Contribution of ω_t to inflation volatility refers to the contribution in 2024Q1 of fluctuations in the average planning horizon to inflation volatility under the optimal targeting rule.

considerable. To demonstrate this, Table 1 shows the additional losses from the central bank following the targeting rule of CGG in which $\kappa\pi_t = -\lambda y_t$ instead of the optimal targeting rule under uncertainty given by equation (22). The CGG targeting rule is suboptimal in the NK-FHP model, because it does not take into account how agents' longer-run beliefs can evolve into inflation scares. Table 1 shows the cost of ignoring inflation scares using the shocks underlying the SEP baseline. Specifically, it simulates stochastic shocks around that path and computes the increase in the central bank's expected discounted losses from adopting the CGG targeting rule instead of the optimal targeting rule. It does so for the NK-FHP model for different degrees of uncertainty about agents' planning horizons, as well as for different values of κ, the slope of the Phillips curve.

As shown in the first row of the table, losses with $\kappa = 0.03$ are about 22% higher if a policy-maker follows the CGG targeting rule instead of the optimal targeting rule, ignoring the possibility of inflation scares in the NK-FHP model with a fixed planning horizon. The table also indicates that the optimal policy calls for about 65 basis points more tightening than the suboptimal targeting rule in 2023 when SEP participants projected that inflation would be well over 3%. This additional tightening results in better inflation outcomes and reduces central bank's losses.[20] Table 1 also shows that as uncertainty about agents' planning horizons increases, the losses from ignoring inflation scares grow larger. This result reflects that inflation scares are more likely to be severe when uncertainty is high, increasing the cost of disregarding them.

Table 1 also reports losses for $\kappa = 0.01$, our mean estimate for the slope of the Phillips curve. Lower sensitivity of inflation to the output gap makes it more costly for a central bank to combat high inflation, and the losses from following the CGG targeting rule rise. The last row of the table reduces the size of the markup shock to its mean estimate and shows the losses with all the estimated parameters set to their mean values. Using these estimates, optimal policy prescribes about 105 basis points more tightening than the CGG targeting rule, and the central bank's losses are about 35% higher if the possibility of inflation scares is ignored.

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 horizons, inflation scares in which longer-run inflation expectations can deviate persistently from a central bank's inflation objective can arise. A central bank with a dual-mandate loss function that faces such a risk reacts forcefully to contain inflationary pressure, responding more aggressively than if planning horizons were known with certainty. In situations where inflation has been running above target, we find that the gains from mitigating the risk of an inflation scare 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 analyze optimal

[20] As shown in the appendix, the mean outcomes for inflation are persistently closer to central bank's inflation target under optimal policy, as trend inflation comes down faster. However, because the CGG targeting rule does not tighten as much as the optimal targeting rule, it delivers outcomes for output that lie closer to potential.

policy when the central bank does not observe the current state of the economy and updates its beliefs according to Bayes' law. It would also be interesting to consider approaches in which the uncertainty regarding agents' planning horizons depend on the state of the economy.

References

ANGELETOS, G.-M., Z. HUO, AND K. A. SASTRY (2020): "Imperfect Macroeconomic Expectations: Evidence and Theory," Working Paper 27308, National Bureau of Economic Research.

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.

COIBION, O. AND Y. GORODNICHENKO (2015): "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 105, 2644–78.

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.

DUPRAZ, S. AND M. MARX (2025): "Keeping Control Over Boundedly Rational Expectations," Bank de France manuscript.

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," *Federal Reserve Bank of Richmond Economic Quarterly*, 79, 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.

HERBST, E. AND F. SCHORFHEIDE (2014): "Sequential Monte Carlo Sampling for DSGE Models," *Journal of Applied Econometrics*, 29, 1073–1098.

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.

Kopecky, K. A. and R. M. Suen (2010): "Finite state Markov-chain approximations to highly persistent processes," *Review of Economic Dynamics*, 13, 701–714.

Korenok, O., D. Munro, and J. Chen (2023): "Inflation and Attention Thresholds," *The Review of Economics and Statistics*, 1–28.

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.

Pfäuti, O. (2025): "The Inflation Attention Threshold and Inflation Surges," University of Texas at Austin manuscript.

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, Edward Herbst, and David López-Salido

This appendix contains additional results and derivations for the paper, "Optimal Monetary Policy with Uncertain Private Sector Foresight." It derives the expressions that determine private-sector expectations formation under optimal time-consistent policy and the targeting rule implemented under optimal policy. It also describes the solution algorithm, model estimation, and details regarding the use of the Summary of Economic Projections. Finally, it includes some additional simulation results of the New Keynesian, Finite Horizon Planning (NK-FHP) model.

A Private Sector Expectations

We begin by deriving expression (12) in the main text, which determines a k-horizon agent's expectation of inflation. This expression can be derived from substituting equation (17) at $j = 0$ into equation (9) to write an agent's beliefs about 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 (17) at $j > 0$ into equation (2):

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

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 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 (12) in the main text with

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

We can substitute expression (A-4) into (17) 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 (22), we differentiate the Bellman equation (19) 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 (24). This first-order condition is the first part of the optimal targeting rule shown in the main text. The second part comes from the envelope condition associated with the Bellman equation (19). This envelope condition satisfies:

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

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 becomes:

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

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

C Solution Algorithm

To solve for the outcomes implied by 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:

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

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

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

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

where private sector expectations for inflation next period in equation (A-10) are given by:

$$g_\pi(v_{pt}, u_t, m_t) = a_{pt} v_{pt} + b_{pt} u_t \tag{A-14}$$

where $a_{pt} = \omega(m_t) a_p(k_0) + (1 - \omega(m_t)) a_p(k_1)$ with $a_p(k) = \left(\frac{\beta\lambda}{\lambda+\kappa^2}\right)^k$. Also, $b_{pt} = \omega(m_t) b_p(k_0) + (1 - \omega(m_t)) b_p(k_1)$ where $b_p(k)$ is given by:

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

With expressions (A-10)-(A-13) determining the optimal outcomes for inflation, the output gap, W_{pt}, and v_{pt+1}, we then determine the optimal policy rate using equation (21).

Equations (A-10)-(A-13) are linear conditional on a value for m_t and we can use them to 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, though with a very important distinction: the central bank knows the planning horizons of the two types of agents in the economy and knows ω_t as well as the process governing its future evolution. (In contrast, a finite planning household or firm, even if it had a very long horizon forms beliefs under the assumption that all other agents have the same planning horizon as themselves.) 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 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 effect 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 $T^i_{j_\tau}$ and $R^i_{j_\tau}$ 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. For an alternative approach to solving dynamic stochastic general equilibrium models with Markov-switching processes, see Foerster et al. (2016).

Properties of Optimal Inflation. When agents have the same planning horizon ($k = k_0 = k_1$), the decision rule under optimal policy simplifies to:

$$X_t = T v_{pt} + R u_t.$$

To show how the optimal response of inflation depends on agents' planning horizons, Figure A-1 plots the first elements of these vectors, T_1 and R_1, which determine the optimal response of inflation, for different values of k. As shown there, when agents have a finite planning horizon, inflation depends on agents' beliefs about longer-run inflation, v_{pt}. Because v_{pt} evolves endogenously in response to past inflation, agents' longer-run beliefs can drift away from the inflation target, giving rise to inflation scares. As k increases, inflation becomes less responsive to v_{pt}, and agents' longer-run beliefs and thus inflation scares become a less important factor in determining aggregate inflation. As $k \to \infty$, the model solution converges to the optimal policy solution of the canonical NK model with rational expectations (e.g., the targeting rule of Clarida et al. (1999)).

D Estimation Procedure

We modify the estimation procedure discussed in Gust, Herbst, and López-Salido (2022) to allow for time-variation in the population-average planning horizon by allowing ω_t to follow a flexible autoregressive process. Specifically, we model an underlying latent process λ_t as an AR(1) process:

$$\lambda_t = (1 - \rho_\lambda)\mu_\lambda + \rho_\lambda \lambda_{t-1} + \sqrt{1 - \rho_\lambda^2}\epsilon_{\lambda t}, \quad \epsilon_{\lambda t} \overset{\text{i.i.d.}}{\sim} N\left(0, \sigma_\lambda^2\right). \tag{A-21}$$

The weight ω_t is obtained using:

$$\omega_t = \Phi(\lambda_t),$$

where $\Phi(\cdot)$ denotes the standard normal CDF. This transformation ensures that ω_t remains within the unit interval $[0, 1]$. While the discrete Markov process is analytically convenient for characterizing the solution to the model under optimal policy, for estimation purposes equation (A-21) allows for a more flexible evolution for ω_t, while remaining computationally simple and analytically tractable when the model is solved assuming that monetary policy follows a policy rule.

In estimating the model, we use the same policy rule discussed in Gust, Herbst, and López-Salido (2022):

$$i_t = \bar{i}_t + \phi_\pi \pi_t + \phi_y y_t + i_t^*,$$

where the intercept is given by $\bar{i}_t = \overline{\phi}_\pi \overline{\pi}_t + \overline{\phi}_y \overline{y}_t$ and i_t^* is an exogenous shock to the nominal interest rate which follows an AR(1) process with innovation ϵ_{it}. These trend variables are constructed using the agents' longer run beliefs, as described in Gust, Herbst, and López-Salido (2022). Given a planning horizon k, we write the decision rule as

$$x_t = T_k x_{t-1} + R_k e_t,$$

where x_t is a vector collecting the relevant variables of the model including the endogenous variables, estimates of the value functions, and exogenous shocks. The vector e_t collects the innovations to

Figure A-1: Properties of the Optimal Decision Rule for Inflation

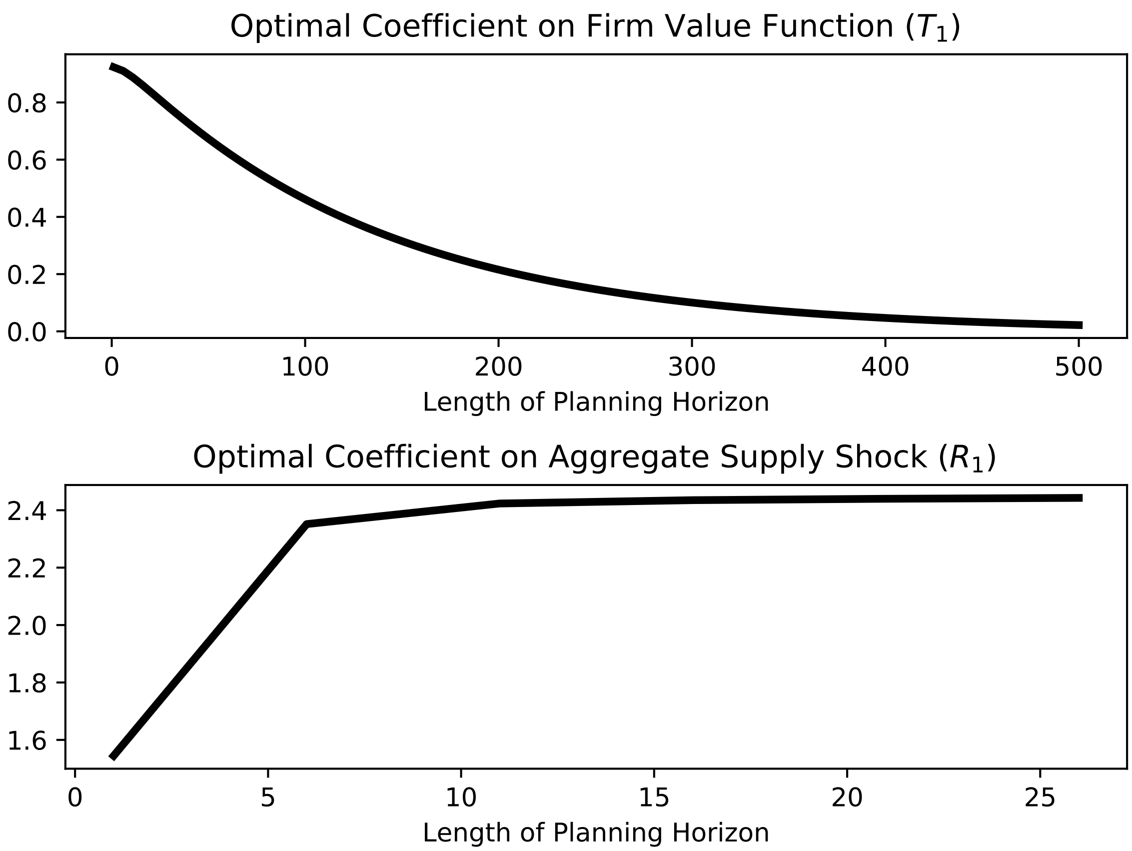

NOTE: The figure shows how the coefficients in the optimal decision rule for inflation depend on agents' planning horizons ($k = k_0 = k_1$) in the NK-FHP model.

the supply, demand, and monetary policy shocks. To aggregate the decision rules, we use the time-varying weights ω_t to construct a weighted combination of the two sets of decision rules. Specifically, we construct time-varying matrices for the state transition equations as:

$$T_t = \omega_t T_{k_0} + (1 - \omega_t) T_{k_1} \text{ and } R_t = \omega_t R_{k_0} + (1 - \omega_t) R_{k_1},$$

for $k_0 = 4$, $k_1 = 32$, and $t = 1, \ldots, T$. (In a slight abuse of notation, T stands for the autoregressive component of the decision rule and the sample size.) After combining the model solution with a set of observables, we have a time-varying linear Gaussian state space system, whose likelihood function can be evaluated using the Kalman filter. The model is estimated as a Bayesian model, incorporating the same observables and prior distributions as in Gust, Herbst, and López-Salido (2024)[21]. For the parameters μ_λ, ρ_λ, and σ_λ, we adopt Normal(0, 1), Beta(8.1, 0.9), and Inverse-Gamma(1, 4) prior distributions, respectively. These priors are chosen to provide moderate informativeness for the latent AR(1) process λ_t, while ensuring that the implied marginal distribution of the weight

[21]The estimated value for the variance of the aggregate supply shock reported in A-1 differs from Gust et al. (2024) because in that paper the reported variance is scaled by κ.

Table A-1: Posterior Distribution of Mixture Model

Parameter	Mean	[5, 95]	Parameter	Mean	[5,95]
γ_h	0.56	[0.36, 0.79]	γ_p	0.12	[0.08, 0.17]
κ	0.01	[0.01, 0.02]	σ	2.46	[1.69, 3.37]
ρ_r	0.87	[0.79, 0.94]	σ_r	0.34	[0.29, 0.39]
ρ_u	0.74	[0.39, 0.97]	σ_u	0.08	[0.03, 0.16]
μ_λ	1.11	[0.07, 2.17]			
ρ_λ	0.95	[0.89, 0.98]	σ_λ	1.21	[0.65, 2.25]

$\omega_t = \Phi(\lambda_t)$ is approximately uniform over the unit interval. Rather than filtering ω_t, we simply add the series $\{\omega_t\}_{t=1}^T$ to the parameter vector, using the implied prior from the process for λ_t.

We elicit draws from the posterior using a Sequential Monte Carlo algorithm as described in Herbst and Schorfheide (2014). We use $45,000$ particles and 500 tempering stages. For the mutation step, we use a block random walk Metropolis-Hastings algorithm with 14 blocks. Table A-1 displays the posterior distribution for selected parameters. The estimates are very similar to those for the model with $k = 4$ for the entire sample in Gust et al. (2024), though the persistence of the supply shock is slightly higher here, reflecting the influence of the $k = 32$ agents.

Figure A-2: Inflation Forecast Predictability and Planning Horizons in the NK-FHP Model

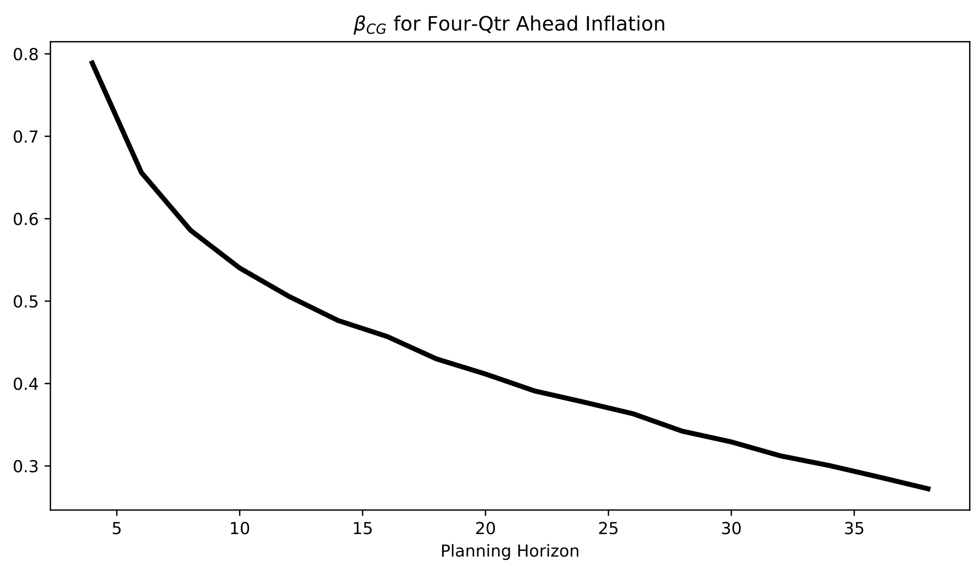

Note: The figure shows the univariate regression statistic from a regression of four-quarter ahead inflation forecast errors on forecast revisions using data generated from simulating the NK-FHP model in which all households and firms have the same planning horizon. Monetary policy follows the Taylor rule with interest rate smoothing (equation (A-22)).

E Forecast Error Predictability in the NK-FHP Model

Coibion and Gorodnichenko (2015) document the predictability of forecast errors from regressions using survey data on expectations. They regress the median forecast error of inflation on the median forecast revision and show that there is a positive correlation between the forecast error and forecast revision, implying that forecast errors can be predicted by forecast revisions and that forecasts tend to underreact to new information. The regression statistic that they use can be written as:

$$\beta_{CG} = \frac{\text{cov}(\mathbb{R}_t,\ \mathbb{F}_{t+4})}{\text{var}(\mathbb{R}_t)},$$

where $\mathbb{F}_{t+4}$ denotes the four-quarter ahead forecast error defined as the realized variable minus its forecast and $\mathbb{R}_t$ denotes the revision to the four-quarter ahead forecast. Coibion and Gorodnichenko (2015) show that β_{CG} varies over time with the high volatility period of the 1970s and early 1980s corresponding to a decline in β_{CG}, which remained low before rebounding in the 1990s. As discussed in Gust, Herbst, and López-Salido (2024), the NK-FHP model generates values of β_{CG} for inflation forecasts in line with the evidence in Coibion and Gorodnichenko (2015). Moreover, it can generate fluctuations in β_{CG} through variations in agents' planning horizons broadly in line with their results. To demonstrate this, Figure A-2 shows β_{CG} computed from inflation forecasts simulated from the NK-FHP model for different planning horizons. β_{CG} has a positive sign in the NK-FHP model, and as agents' planning horizons lengthen, β_{CG} declines. This decline reflects that in the NK-FHP model a longer planning horizon implies that agents' forecasts become more rational and thus forecast errors become less predictable.

F Details Regarding Using the SEP

To use the March 2023 SEP in conjunction with the NK-FHP model, we construct a path of aggregate demand and supply shocks that is consistent with the median projections for inflation, the unemployment rate, and the federal funds rate from the SEP. We assume that the median path of the federal funds rate follows a Taylor rule:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i)(\phi_\pi \pi_t + \phi_y y_t) + e_{mt} \quad \text{(A-22)}$$

where $\rho_i = 0.85$, $\phi_\pi = 1.5$, and $\phi_y = 0.25$ and e_{mt} is an *iid* innovation.[22] We use this rule along with the NK-FHP model in which ω_t is fixed to infer the 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. We also use this rule to simulate the effects of the aggregate supply shock shown in Figure 2 of the main text.

To construct the implied path of core PCE inflation, the federal funds rate, and the output gap from the SEP, the annual median projections are linearly interpolated. The output gap is constructed assuming an Okun's law relationship between the output gap and unemployment rate gap. A value of 2 is used to translate the median's participant's projection for the unemployment rate minus the median estimate of its longer-run value into an output-gap series.

To simulate the NK-FHP model using the March 2023 SEP, we also need to set initial values in 2022Q4 for the household and firm's continuation value functions, v_{pt} and v_{ht}. For v_{pt}, we use equation (13) along with past data on core PCE inflation to determine its initial value. Setting

[22]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).

$\gamma_p = 0.12$, its estimated value, implies $v_{pt} = 0.31$ in 2022Q4. We follow a similar procedure to initialize v_{ht}. In particular, equation (12) implies that v_{ht} depends on past values of the output gap and inflation. Using the estimated value of γ_h from Gust et al. (2022) along with the Congressional Budget Office's estimates of the output gap implies $v_{ht} = 3.2$ in 2022Q4.

Figure A-3: The SEP-Consistent Baseline

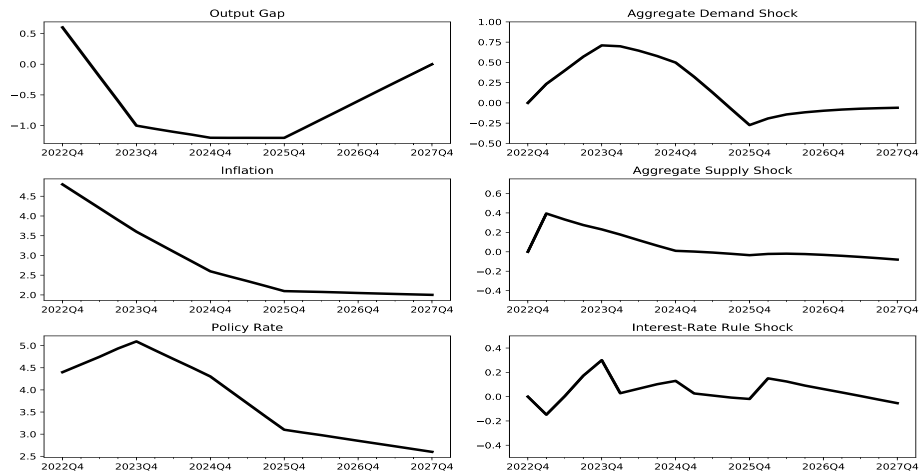

NOTE: The SEP-consistent baseline is constructed by interpolating quarterly data using the March 2023 SEP. The NK-FHP model with a fixed planning horizon is used to infer the shocks to aggregate demand, aggregate supply (cost-push), and the interest-rate rule over the 2023Q1-2027Q4 period.

The left-hand side of Figure A-3 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 was above potential in early 2023 but expected to decline in 2023 and remain below potential in 2024. The median SEP participant also projected that under appropriate monetary policy the federal funds rate would peak later in 2023 and then fall in 2024 in line with a projected decline in inflation.

The right-hand side of Figure A-3 shows the shocks over the 2023Q1-2027Q4 period that allow the model-implied paths of output, inflation, and the policy rate to match the SEP-consistent baseline. The combination of aggregate demand and supply shocks shown there underlie the high inflation and restrictive monetary policy that contributes to a level of output below potential in 2023.[23] As the aggregate demand and supply shocks fall back toward zero, inflation comes down, policy becomes less restrictive, and output converges toward potential. The lower right panel of Figure A-3 shows that the monetary policy shocks are relatively small, less than 25 basis points over most of the projection period, suggesting that a Taylor rule with interest-rate smoothing fits the SEP-consistent baseline reasonably well.

G Additional Policy Simulations

Figure A-4 shows the corresponding path of mean outcomes for the simulations that underlie the distributions of outcomes shown in Figure 3. Specifically, the red dotted lines display the expected

[23] Interestingly, the NK-FHP model generates a combination of aggregate 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.

paths of inflation, the output gap, the real interest rate, and trend inflation in the NK-FHP model under optimal policy with fixed planning using the baseline aggregate demand and supply shocks shown in Figure A-3. As shown there, the mean outcomes for inflation and the policy rate in the NK-FHP model are notably higher than in the canonical NK model, while the path of the output gap is considerably lower as a result of the higher path of the policy rate. The figure also shows that the policy rate under uncertain planning needs to be slightly higher than under fixed planning in order to achieve similar inflation outcomes. This slightly tighter policy, on average, results in mean outcomes for the output gap that lie below those under fixed planning.

Figure A-4: Optimal Policy Under Fixed and Uncertain Planning

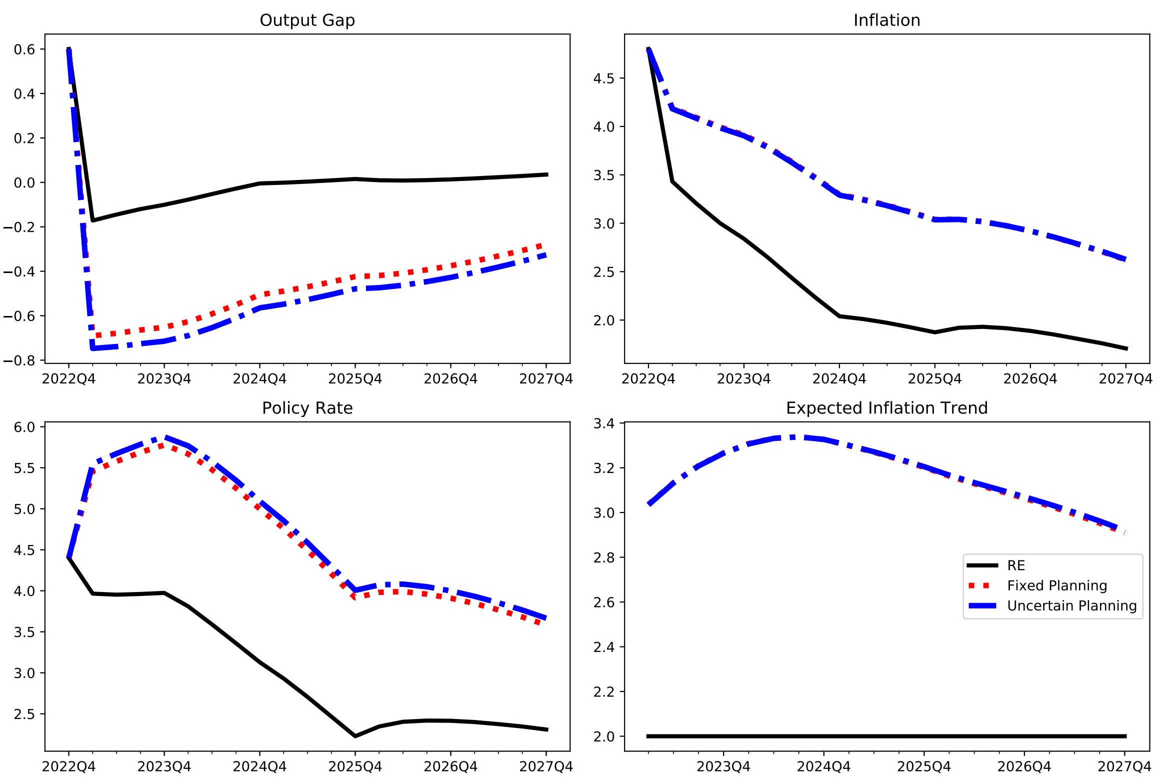

Note: The figure shows the mean responses under optimal discretionary policies using 50,000 draws of shocks centered around the aggregate demand and supply shocks associated with the SEP-consistent baseline. The red dashed line shows the optimal policy paths using the NK-FHP model with a fixed planning horizon; and, the dashed blue line shows the optimal policy paths when the policymaker is uncertain about agents' planning horizons.

In Figure 3, the outcome distributions reflected that the volatility in planning horizons (ω_t) accounted for about half of the volatility in inflation in 2024Q1. In Figure A-5, we reduce the volatility of ω_t so that it only accounts for about 20 percent of inflation volatility with the markup shock accounting for the rest. In this case, the distributions of the policy rate and the output gap are less skewed; however, the distribution for the policy rate is still skewed to the upside while the distribution for the output gap remains skewed to the downside. The distribution of outcomes for inflation is very similar to that under fixed planning when the volatility of ω_t is relatively low.

Comparison of Alternative Policies. Figures A-6 shows simulations around the SEP base-

Figure A-5: Distribution of Outcomes With Smaller Fluctuations in Planning Horizons

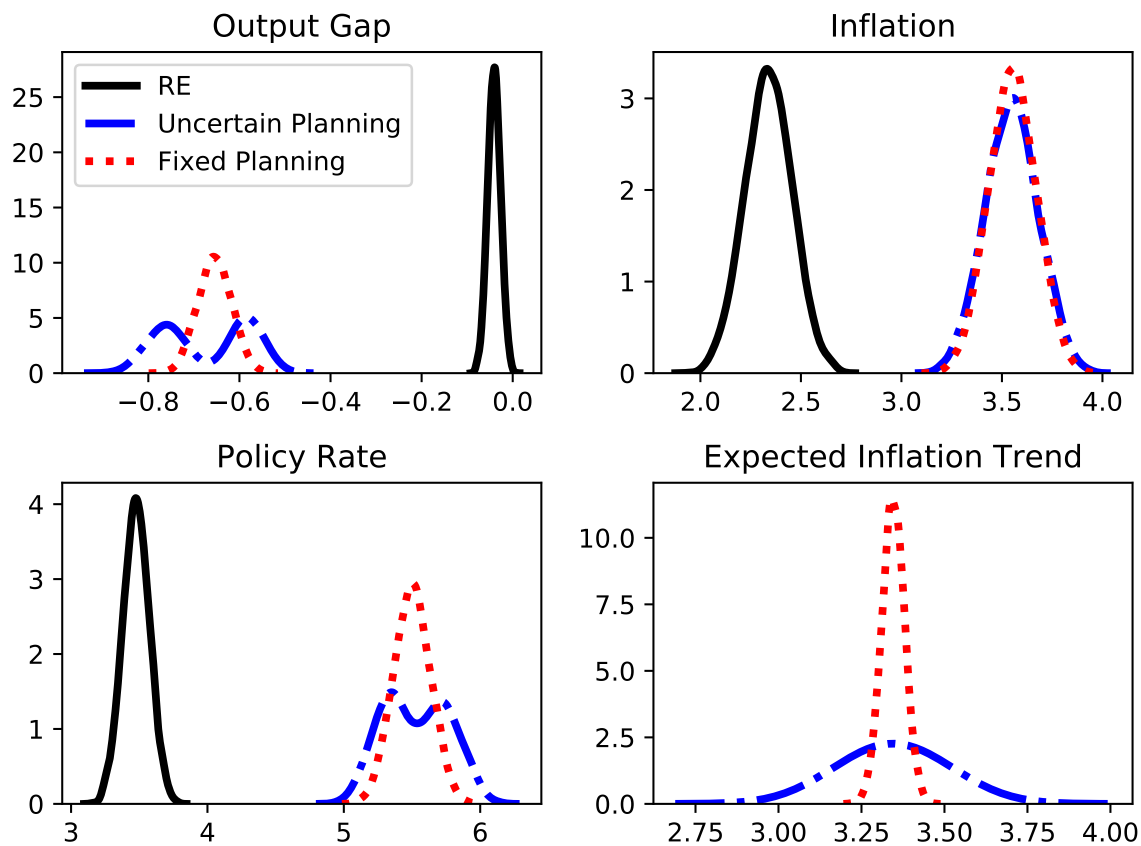

NOTE: The figure shows the distribution of responses under optimal discretionary policies in 2024Q1 using 50,000 draws of shocks centered around the aggregate demand and supply shocks associated with the SEP-consistent baseline.

line comparing the mean outcomes under CGG targeting rule to those under optimal policy in an economy with uncertain planning horizons. These simulations correspond to the last row of Table 1 in which uncertain planning contributes to 50% to the volatility of inflation in 2024Q1 under optimal policy, the value used in Figures A-4 and 3. Accordingly, the dashed blue lines in this figure are the same as in Figure A-4. As discussed in the paper, the CGG targeting rule is suboptimal in the NK-FHP model with uncertain planning because it does not explicitly take into account the possibility of inflation scares. As a result, policy is less restrictive under the CGG targeting rule than under optimal policy despite the high level of inflation projected by SEP participants, resulting in higher trend inflation. This difference in trend inflation between the two policies increases through 2027 and persists for many years after the period shown in the figure. Still, the CGG targeting rule is restrictive enough to put inflation on a downward trajectory toward the central bank's inflation target of 2%. Although the CGG targeting rule results in larger deviations of inflation from target than the optimal policy, it compensates for this by generating a mean path for output that lies closer to potential.

Figure A-6: A Comparison of Alternative Policies in an Economy with Uncertain Planning

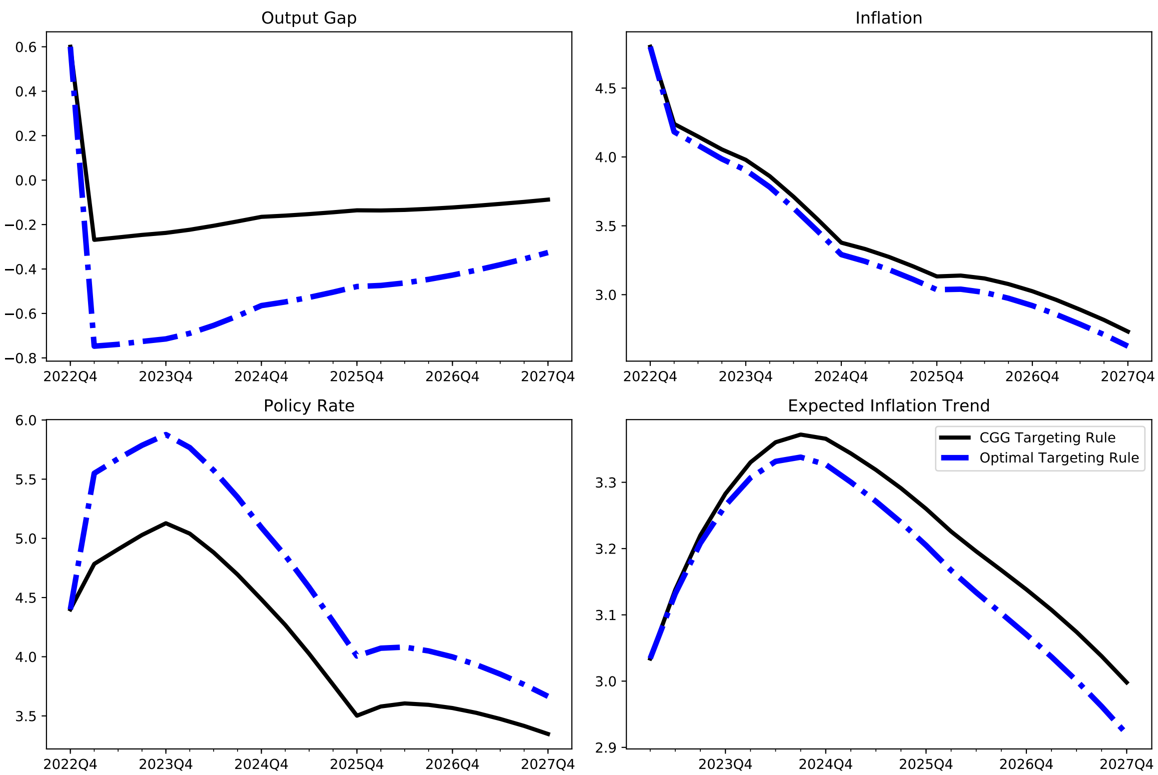

NOTE: The figure compares the mean responses under optimal policy to the CGG targeting rule in an economy with uncertain planning horizons using 50,000 draws of shocks centered around the aggregate demand and supply shocks associated with the SEP-consistent baseline.