Rationalizing Fed Interest Rate Decisions in the 2020's*

Thomas J. Sargent
New York University

Noah Williams†
Miami Herbert Business School, University of Miami

October 6, 2025

Abstract

After the COVID-19 pandemic, inflation surged in the United States. To help us understand the Fed's slow response, we create a model of what the Fed thought about when it set its interest policy rate. We assume that the Fed recurrently updated estimates of a drifting-coefficients model that it used to pose a sequence of Phelps (1967) control problems for setting the policy rate under a Kreps (1998) anticipated utility assumption. Our model tells how the Fed's decisions were shaped by (1) a decline in inflation persistence, (2) a flattening of the slope of a Phillips curve, and (3) mismeasurements of real-time output or unemployment gaps. The first two told the Fed that the inflation surge during COVID was transitory and that raising the policy rate aggressively would be costly, while the third suggested that aggregate output was below potential output. The data eventually pushed parameter estimates in directions that told the Fed that it could disinflate with smaller costs in terms of output and unemployment. We study how real-time Fed forecasts and policy statements align with prescriptions from our Phelps control problems.

Keywords: Model predictive control, drifting coefficients, inflation, Fed interest rate instrument.

*We thank David Andolfatto, Emi Nakamura, Giorgio Primiceri, and especially John Cochrane for helpful comments on and criticisms of an earlier draft.

†Corresponding author: noahwilliams@miami.edu

1 Introduction

The Fed responded belatedly to the COVID-19 surge in inflation. Despite widespread fears that doing so would lead to a recession, in 2022 the Fed raised its policy rate aggressively. To rationalize how the Fed set its target interest rate in the 2020s, we construct a model that tells us how the Fed's interest rate decisions were shaped by what it thought about the macroeconomy. In the spirit of Bernheim (1984) and Pearce (1984), we endow our "artificial Fed" with a flow of data, a statistical model, and a sequence of optimal control problems that determine its decisions. We pretend that the Fed uses "model predictive control" to set its policy rate. Our artificial Fed updates parameters of its statistical model each period, and then makes today's decisions by solving a new dynamic programming problem.[1] We use this structure to represent how the Fed expressed and managed its concerns that tightening monetary policy enough to arrest current or prospective inflation might provoke a recession. Our model rationalizes how the Fed struggled to balance its dual goals to keep inflation low and output growth high during the 2020s.[2]

The "model predictive control" behavior that we attribute to the Fed shares features with Phelps (1967), Sargent (1999), Cogley and Sargent (2005b), Primiceri (2006), and Sargent, Williams, and Zha (2006). Each period the Fed sets its interest rate instrument by solving a linear-quadratic version of a Phelps (1967) optimal control problem in light of the latest estimates of a drifting coefficients model that links outcomes to its policy instrument.[3] In formulating its Phelps problem, the Fed adopts an anticipated utility assumption presented by Kreps (1998) and Cogley and Sargent (2008): each period the Fed solves a new optimal control problem with its most recent estimate of a fixed-coefficient transition law as key components. As in Sargent (1999, ch. 5), the Fed's most recent estimate of the persistence of inflation plays a key role in shaping the advice that the Phelps problem provides. Our drifting coefficients statistical model describes how the Fed's beliefs about key determinants of its decision rule for the policy rate changed during the period of low and stable inflation that followed the 2008 financial crisis. Estimated persistence of inflation declined during this period, leading the Fed to believe that the post-COVID surge in inflation was transitory and would fade away without a more restrictive monetary policy action. The statistical model

[1] For descriptions and applications of model predictive control, see García et al. (1989) Mayne et al. (2000), Rawlings et al. (2017).

[2] See Summers (2024) and Staff (2024) for descriptions of diverse opinions about the likelihood of a "soft-landing" after the surge in US inflation after the COVID-19 pandemic.

[3] We recognize that, as Sargent (1999, p. 58) said, "... Phelps's control problem carries a tattered past. It fortified those in the 1970s who advocated learning to live with high inflation because of the unacceptably high costs of unemployment from disinflating." Cogley and Sargent (2005b) turn that defect to their advantage when they use a Bayesian race among three Phelps problems to rationalize why the Fed chose to preside over high inflation in the 1970s even when more natural-rate friendly models fit better.

that we impute to the Fed explains how, throughout the post-COVID rise in inflation, the Fed (and outside commentators from "team transitory") predicted a quick decline of inflation back to the 2% target.[4] The Fed's opinions about the slope of the Phillips curve also mattered. Post-2008 data led policymakers to believe that the Phillips curve had flattened, implying a weaker link between inflation and real economic activity.[5] Balancing consequences of declining persistence and declining slope, our artificial Fed believed that the post-COVID surge in inflation would decline quickly enough and that aggressive actions to accelerate its decline would provoke a serious recession.

We acknowledge that the drifting coefficients model that we impute to the Fed is what Koopmans (1947) called a purely descriptive "Kepler stage" model. It is not a structural model with parameters that a 1980's Minnesota monetary economist or a 2020's fiscal theorist of the price level could hope is invariant to a pertinent class of hypothetical policy interventions.[6] If nature were actually to generate outcomes according to a New Keynesian structural model that we describe in section 3, then the Fed's perceptions of low persistence of inflation and a flat Phillips curve were actually a *consequence* of past Fed actions that had responded aggressively to adverse shocks to inflation. That made the Fed think that an inflation surge would be easier to manage than it had been in Volcker's day.[7]

Section 2 describes three considerations that preoccupy our artificial Fed, namely, the slope of the Phillips curve, the persistence of inflation, and the accuracy of measures of an output gap. Section 3 takes a detour by describing a New Keynesian model in which the slope of the Phillips curve and the persistence of inflation that would be estimated by our artificial Fed are both functions of the "aggressiveness" parameter in the monetary authority's Taylor rule. We briefly describe a version of this model in which that parameter switches between a large hawkish value and a low dovish value according to a Markov chain. Simulations of the Markov switching version of our New Keynesian model yield time series that tell a drifting coefficients statistical model that cannot condition on the Markov state governing the aggressiveness parameter that the slope of the Phillips curve and the persistence of inflation both move over time. This sets the stage for Section 4 that presents the quantitative

[4] See Blinder (2021) and Krugman (2022) for some "team transitory" views.

[5] In the Lucas-style model in Sargent (1999, ch. 7), a flatter Phillips curve would mean that surprise inflations wouldn't affect real activity as much. That would motivate the Fed to want lower inflation. Our statistical model gives advice that would seem more appropriate in the context of a New Keynesian model like one in section 3. In that model, a flatter Phillips curve means that larger negative real outcomes are necessary to bring down inflation, leading the Fed to anticipate that a recession would be a consequence of aggressive application of its interest rate instrument.

[6] Thus, in our model, the Fed ignores the Lucas Critique in the same way that it does in the "vindication of econometric policy evaluation" story told in Sargent (1999, ch. 2)

[7] This process has features like those described by Minsky (1986), in which a sustained episode of stability set the stage for subsequent destabilization.

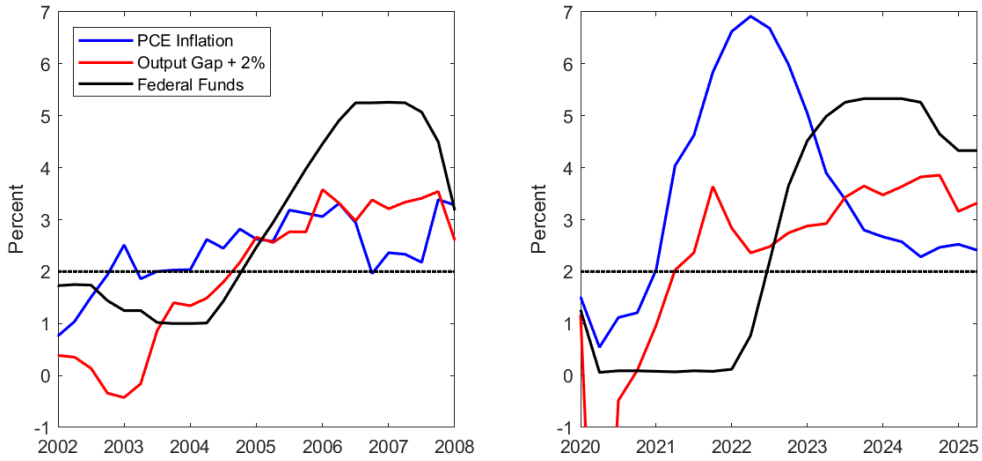


Figure 1: Inflation (PCE, year-over-year), the output gap (CBO), and the federal funds rate during two tightening episodes.

model that we use to rationalize the Fed's interest rate decisions during the 2020's. Section 5 compares actual projections by Fed decision makers contained in the Federal Reserve's *Summary of Economic Projections* (SEP) reports to the section 4 forecasts that we use to rationalize Fed. Section 6 offers concluding remarks. Appendix A provides proofs for some section 3 assertions about properties of the New Keynesian model, while section B provides statements about properties of the section 4 Phelps problems.

2 Elements of Our Story

2.1 Motivation and Overview

Figure 1 sets the stage by showing key economic indicators for monetary policy over two policy tightening episodes. It shows inflation (PCE inflation year-over-year), the output gap (percentage deviation of GDP from the CBO measure of potential GDP), and the federal funds rate on a quarterly frequency from 2002:Q1-2008:Q1 (left panel) and 2020:Q1-2025:Q2 (right panel). For ease of comparison, we add two percentage points to the output gap, so that 2% in the figure (shown with a dotted line) is where both inflation and the output gap are at their target levels.

The left panel of Figure 1 shows the type of policy response that has typified US monetary policy since the 1980s. After being flat right at the 2% target for a couple of quarters, inflation rose above target in the second quarter of 2004. The output gap was negative, but it had been moving up, and the Fed responded within one quarter by increasing the federal funds

rate. As both inflation and the output gap continued to increase, the Fed kept tightening until 2006, when the output gap stabilized and inflation fell back to target. While the Fed responded quickly to inflation rising above target, some commentators including John Taylor argued that even here the Fed was too late.[8]

The right panel of Figure 1 shows that the Fed's response to inflation was different after the COVID-19 pandemic. Economic activity collapsed with onset of the pandemic and shutdowns in early 2020, but it recovered quite rapidly.[9] The recovery was fueled by large transfers and other fiscal expenditures. Those and supply restrictions sparked inflation. By 2021:Q1 inflation had surpassed the Fed's 2% target, and the output gap turned positive the following quarter. But for more than a year, the Fed did not respond to the increases in inflation or output. Only in 2022:Q2 did the Fed start to tighten, at which point inflation was already at its peak of 6.9%.[10] We want to understand the Fed's delayed policy response. Our approach will be to construct a drifting-coefficients statistical model and an associated sequence of Phelps problems that rationalize the observed Fed choices of its interest rate target.

Our story emphasizes three key elements which together contributed to the slow response of policy to the inflation surge. The first two characterize changes in inflation dynamics that were well-documented and discussed starting in the mid-2000s: (i) inflation became much less persistent, and (ii) inflation had a weaker relation with real activity, or in other words, the Phillips curve became substantially flatter. The final component is that (iii) real-time uncertainty about economic conditions is especially acute in periods of rapid change. Each of these components is well documented by Fed staff reports and statements of Federal Reserve officials; each is also reflected in the belief estimates and real-time data we report below.

2.2 Element 1: Declining Inflation Persistence

From the 1970s into 1980s inflation became quite persistent, with shocks leading to long-lasting increases in inflation. A large body of empirical evidence, starting in the mid-2000s, documented a marked decline in persistence after the mid-1980s.

Cogley and Sargent (2002, 2005a) documented the decline in persistence of inflation during Volcker and Greenspan's conquest of U.S. inflation. Stock and Watson (2007) found that the magnitude of permanent shocks to inflation shrank in the 1980s–90s, so that "inflation

[8] Taylor (2007, 2009a,b) argued that policy was overly loose in the 2003-2004 period, as both inflation and the output gap were on an increasing path that was expected to continue.

[9] We crop the figure at -1% to make the key data more visible, thus eliding the -9.1% output gap in 2020:Q2, which would be plotted here at -7.1%.

[10] This timing refers to our quarterly data. The initial Fed tightening of a quarter point came on March 16, 2022, and so was very late in the first quarter of 2022.

became much less persistent on average." Williams (2006) showed that starting in the early 1990s, Phillips curve estimates yielded a sum of lag coefficients well below one, meaning that inflation reverted to mean much faster than in earlier decades.[11] In addition to these early papers, later contributions included Benati and Surico (2008), Carlstrom, Fuerst, and Paustian (2009), Cogley, Primiceri, and Sargent (2010), and Davig and Doh (2014). In a comparative analysis, Kurozumi and Zandweghe (2018) reported that during 1970–84, U.S. PCE inflation averaged 6.5% with an autocorrelation of 0.82, whereas in 1985–2008 inflation averaged 2.5% and the autocorrelation fell to 0.33.

While many of the papers related the decline in inflation persistence to changes in monetary policy, others including Pivetta and Reis (2007), cautioned that some of the decline might reflect smaller shocks rather than a true structural break. But there seems to be a broad consensus that inflation was far less persistent than in the 1970s. Summarizing this evidence, as well as providing an introduction to our next element, former Federal Reserve Chair Janet Yellen (2019) said:

> "The comparison I just described concerning the behavior of inflation as unemployment declined during the current expansion and that of the 1960s illustrates two robust empirical findings. First, the slope of the Phillips curve – a measure of the responsiveness of inflation to a decline in labor market slack —- has diminished very significantly since the 1960s. In other words, the Phillips curve appears to have become quite flat. And second, inflation has become much less persistent, because the impact of lagged inflation on current inflation has declined considerably."

2.3 Element 2: A Flatter Phillips Curve

The Phillips curve describes the sensitivity of inflation to economic slack, most commonly either the deviation of output from potential or the deviation of the unemployment rate from the natural rate. Since the 1990s, estimates of its slope in the U.S. have trended toward zero. Yellen (2019) noted that in the 1960s a fall in unemployment from 6.7% to 3.6% raised core inflation by 4 percentage points, while in the 2010s a comparable drop barely moved core inflation at all.

Kuttner and Robinson (2010) and Ball and Mazumder (2011) provided evidence that the slope of the New Keynesian Phillips curve was much lower in recent decades. This decline in the connection between inflation and real activity became evident in the "missing

[11] That sum of coefficients plays a key role in the policy prescriptions from a Phelps problem. See Sargent (1999, ch. 5).

disinflation" following the Great Recession, as surveyed by Simon, Matheson, and Sandri (2013). Davig (2016), McLeay and Tenreyro (2019), Jørgensen and Lansing (2019), and Occhino (2019) discuss how monetary policy and anchored inflation expectations may lead to the weakening of the observed (reduced form) Phillips curve relationship between inflation and real activity. Furlanetto and Lepetit (2024) survey the evidence on the slope of the Phillips curve, and discuss alternative estimation methods that attempt to deal with policy endogeneity. While we focus on the United States, many of these trends were international. In a summary of the evidence across 29 countries, Hobijn, Miles, Royal, and Zhang (2023) note that, "Before the pandemic, the estimated slopes of these countries' Phillips curves were small in magnitude and centered at approximately zero. In other words, most countries had a very flat Phillips curve."

Policymakers were well aware of the consequences of a flatter Phillips curve for monetary policy. Former Fed Chair Ben Bernanke (2022, p. 220) writes, "A flat Phillips curve means that inflation is a less reliable indicator of economic overheating. Should inflation get too high, the costs, in terms of unemployment, of bringing inflation back down to target could be higher than in the past." In addition, after reviewing the evidence on the Phillips curve, then Vice Chair of the Federal Reserve Richard Clarida (2019) discussed its implications:

> A flatter Phillips curve is, in a sense, a proverbial double-edged sword. It permits the Federal Reserve to support employment more aggressively during downturns—as was the case during and after the Great Recession—because a sustained inflation breakout is less likely when the Phillips curve is flatter. However, a flatter Phillips curve also increases the cost, in terms of economic output, of reversing unwelcome increases in longer-run inflation expectations.

2.4 Evidence about Changes in Inflation Dynamics

A key component of our section 4 model of Fed decision making is a drifting coefficients model of the Fed's perceived inflation dynamics. We construct the model by following Sargent (1999), Cho, Williams, and Sargent (2002), Sargent, Williams, and Zha (2006), Primiceri (2006), and many others by using a constant-gain recursive least squares procedure. In particular, we estimate the following Phillips curve relationship:

$$\pi_t = \alpha_{0,t} + \rho_t \, \pi_{t-1} + \kappa_t \, x_t + \varepsilon_t^{\pi}, \tag{1}$$

where π_t is inflation and x_t is the output gap. Here ρ_t is the perceived inflation persistence and κ_t is the perceived slope of the Phillips curve at date t. We group together the coefficients

into a vector θ_t, and the regressors (a constant, lagged inflation, and the output gap) into a vector X_t, and let R_t be an estimate of that vector's second moment matrix. Then updating follows recursive least squares with constant gain $\gamma > 0$:

$$\begin{aligned} \theta_{t+1} &= \theta_t + \gamma R_t^{-1} X_t \big(\pi_t - X_t' \theta_t\big) \\ R_{t+1} &= R_t + \gamma (X_t X_t' - R_t). \end{aligned}$$

The gain γ discounts past observations, allowing the algorithm to track drifting coefficients.

Backward-looking Phillips curves like equation (1) are widely used in the empirical literature referenced above, as well as in the Federal Reserve staff forecasts, as discussed by Peneva, Rudd, and Villar (2025). Figure 2 reports estimated coefficients. To link directly to the empirical results above, here we use current data (that is, ex-post revised data) on PCE inflation and the CBO measure of the output gap at a quarterly frequency using pre-pandemic data from 1960:Q1-2019:Q4.[12] Features we emphasize here are robust to adding more lags in the estimation and to using alternative measures of inflation or real activity. In the next section we discuss the implications of real-time data uncertainty.

Figure 2 shows that the evolution of the coefficients aligns with changes in inflation dynamics that we noted above. The middle panel shows that the persistence parameter ρ_t first increases substantially from the late 1960s through 1970s during the period of high and volatile inflation. Then persistence declines notably after the mid-1980s, with an especially precipitous fall following the 2008 recession. The bottom panel shows that the slope parameter κ_t follows a general downward tend from the early 1980s through 2019. This downward trend was interrupted by a short-lived upward spike in the slope of the Phillips curve in 2009 during the 2008-2009 recession, as the economy experienced a deflation during the worst periods of contraction.

In our section 4 model, lower perceived inflation persistence rationalizes moderation and delay in the Fed's response, since any increase in inflation can be expected to fade away quickly. At the same time, a flatter perceived Phillips curve implies that a larger sacrifice ratio would accompany disinflation, a consideration that discourages aggressive tightening.[13]

[12] The figure uses estimates from a pre-sample regression (1949-1959) as initial conditions for the beliefs, and sets the gain at $\gamma = 0.03$.

[13] In Appendix B we illustrate these comparative statics analytically in a simpler version of the model in section 4.

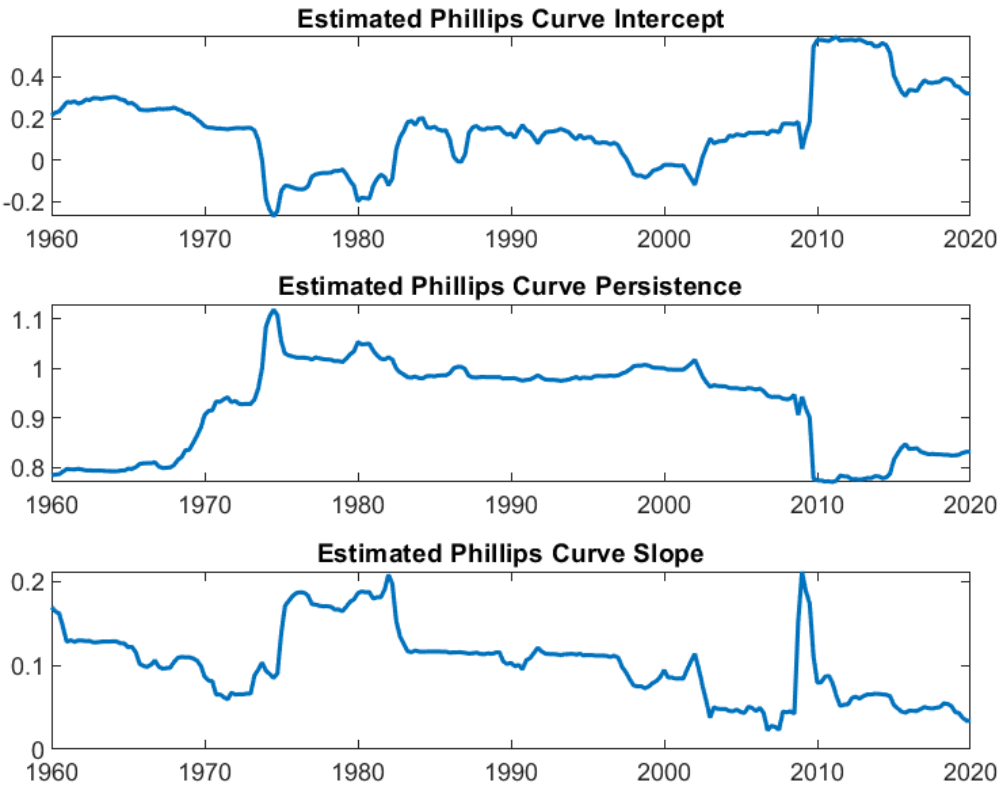


Figure 2: Constant-gain recursive estimates of the intercept, persistence ρ_t, and slope κ_t of the Phillips curve, 1960:Q1–2019:Q4.

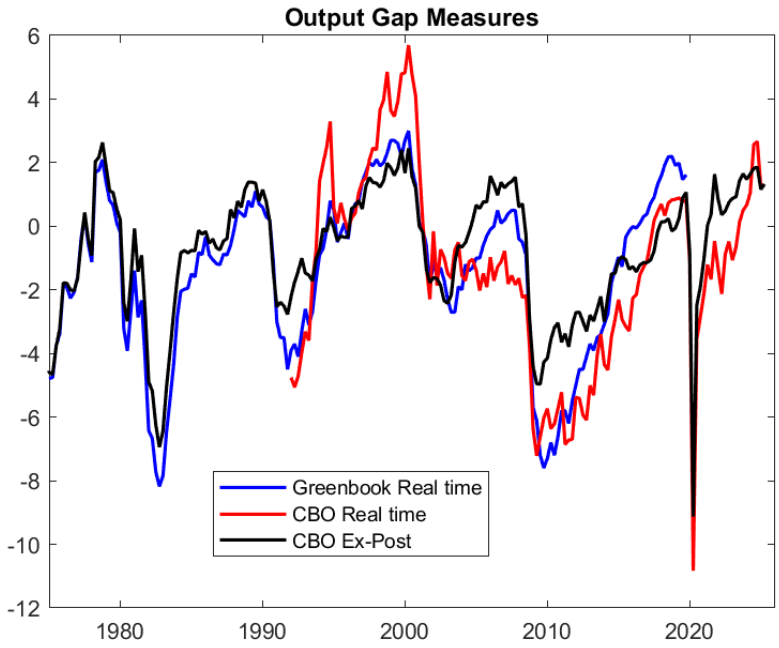


Figure 3: Output gap estimates, real-time vs. ex post. Real-time estimates come from the CBO contemporary vintages (1992-2025) data and Greenbook current-quarter assessments (1987-2019, with estimates for 1975-1986 provided in 1996).

2.5 Element 3: Real-Time Uncertainty about the Real Economy

A third part of our story is uncertainty about the state of the economy in real time. In particular, the Phillips curve relationships that underlie equation (1) use a measure of cyclical resource utilization, and hence rely on an estimate of the "gap" from the economy's potential. Those estimates of potential, whether potential output or the natural rate of unemployment, shift over time with structural changes in the economy and are inherently difficult to measure in real time.

Orphanides (2001) emphasized that policy based on real-time data can look very different from ex-post evaluations. He argued that mismeasured output gaps were central to the Fed's failures in the 1970s. Subsequent work, including Orphanides and van Norden (2002), Edge, Laubach, and Williams (2007), and Coibion and Gorodnichenko (2015), showed that real-time gaps are highly uncertain, subject to large revisions, and often misstate the degree of slack. Nakamura, Riblier, and Steinsson (2025) and Giannone and Primiceri (2025) provide recent analyses showing how real-time uncertainty affected policy decisions, including in the post-pandemic period. These findings underscore the importance of incorporating measurement error into models of policy and expectation formation.

While data collection and analysis have probably improved over time, real-time uncer-

tainty remains an important issue for policymaking. Figure 3 compares real-time estimates of the output gap with the most recent data provided by the CBO. In particular, the figure shows the Federal Reserve's Greenbook/Tealbook contemporary output gap estimates, which are available from the Philadelphia Fed from 1987-2019. That series is supplemented with the first available historical estimates from 1975-1986, which date from 1996. For comparison and also to cover the most recent years, we include the real-time estimates from the contemporary-vintage CBO estimates, which are available from 1992-2025.[14]

We see that both real-time output gap measures differed substantially from the current data estimates in ways that bear crucially on counter-cyclical policy. Before 2008, the Greenbook measures were usually more accurate, with the CBO real-time overstating the late-1990s boom and understating the mid-2000s expansion. From 2009-2019, both real-time measures provided a similar picture: overstating both the depth of the recession and strength of the recovery. For the pandemic and afterward, we only have the CBO measure, and from 2020 through at least 2023, the real-time gap was persistently below the ex-post revised measure, leading policymakers to perceive more slack than was ultimately present. This bias reinforced the belief that inflationary pressures would be temporary and that strong tightening was unnecessary.

While we mostly take the output gap as our cyclical measure of real economic activity, much of the literature also uses the unemployment gap, the gap between the current unemployment rate and the natural unemployment rate. It is widely recognized that real-time difficulties in estimating the natural rate of unemployment resemble those in estimating the output gap. Primiceri (2006) emphasized this in his analysis of the 1970s inflation. Figure 4 plots the measure of the real-time natural rate of unemployment that we use, along with the contemporary measure from the CBO (their estimate of the noncyclical rate of unemployment). Our real-time measure combines different estimates from the Federal Reserve. From 1989:Q1-2007:Q3 we take the contemporary Greenbook NAIRU estimates, provided by the Philadelphia Fed. Then starting in 2007:Q4 we use the median long-term unemployment rate provided in the Federal Reserve Survey of Economic Projections (SEP).

The graph shows notable deviations between the real-time and ex-post measures, particularly in the mid-1990s and the mid-2010s where the real-time measure overstated the natural unemployment rate, and hence understated the unemployment gap. Part of the difference may reflect differences in the natural-rate concepts, as the CBO also had earlier produced

[14] Edge and Rudd (2016) provide a longer series of real-time output gap measures from the Greenbook. But there seems to have been a conceptual change in the definition of the output gap. The reported real-time values were negative for every quarter from 1969:Q3–1988:Q1, inconsistent with the later cyclical measures which are centered at zero. Using this longer real-time series affected our results in the earlier periods, but not after 1995, which is our focus.

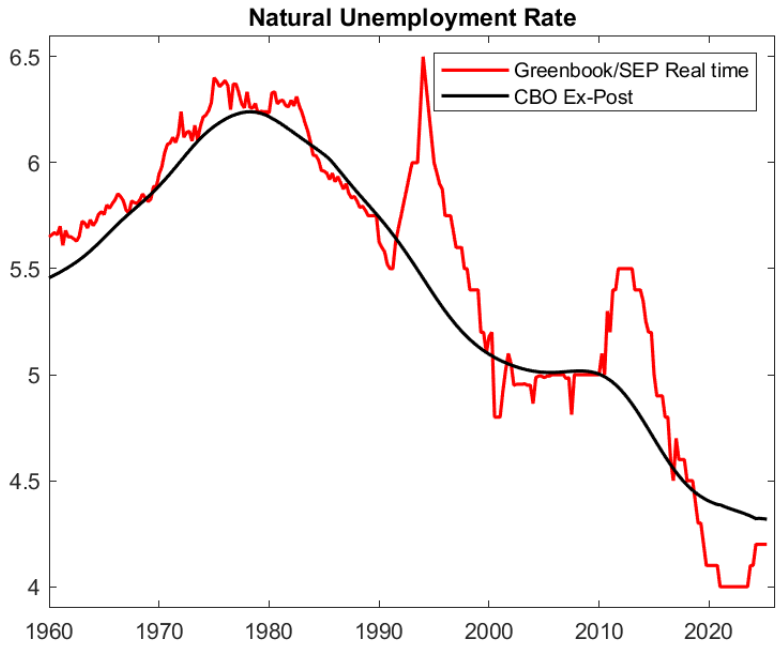


Figure 4: Natural unemployment rate: real-time vs. ex post. Real-time estimates come from the Greenbook current-quarter assessments (1989:Q1-2007:Q3, with 1960-1989 estimates from 1997) and the Federal Reserve SEP (2007Q4-2025). Ex-post is the current CBO noncyclical unemployment rate.

a short-term natural rate measure showing an increase similar to our real-time measure following the 2008 recession. Especially relevant for us, the real-time natural rate projection has been below the current CBO estimate since 2019, suggesting that policymakers have believed that they could continue to push for tighter labor market conditions. While we mostly focus on the output gap as our measure of real activity, below we use Okun's law to compare our model's output gap forecasts to the SEP unemployment projections.

While the real-time output gap measures differed substantially across data sources and relative to the ex-post revised measures, the picture was quite different for inflation. Figure 5 plots the real-time quarterly estimates of PCE inflation along with the contemporary revised series. The real-time series from 1979-2019 come from the Greenbook current-quarter estimates. For 1979-1999, the Greenbook reported CPI inflation projections. To convert these to comparable PCE inflation, we used the ex-post realized ratio of PCE inflation to CPI inflation. For 2020-2025, we used the real-time vintage estimates from the BEA provided in ALFRED. The figure clearly shows that inflation is much easier to recognize and diagnose in real time, as only the deflationary trough in 2009 and the inflationary spike in 2022 have notable, if minor, differences between the real-time and revised data. In both cases the

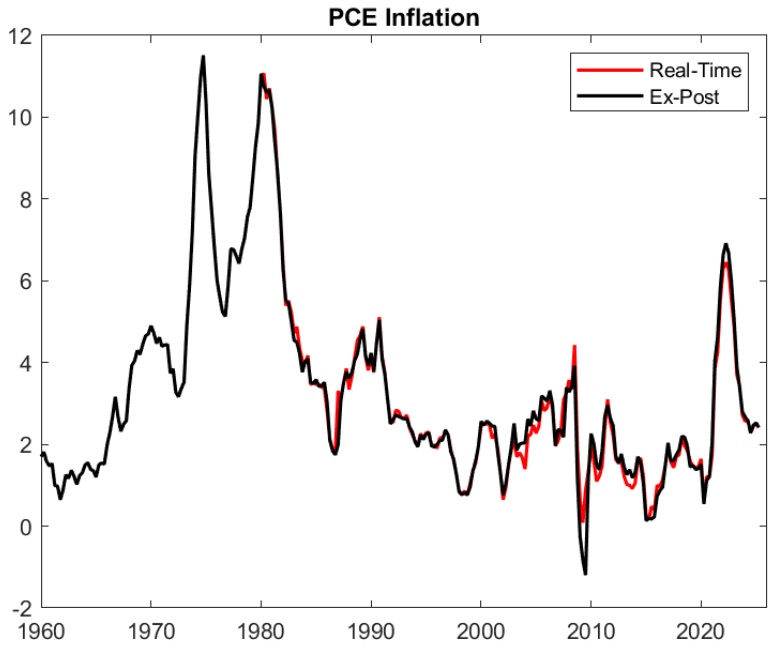


Figure 5: PCE Inflation: real-time vs. ex-post. Real-time estimates come from the Greenbook current-quarter assessments (1987-2019), and BEA contemporary vintages (2020-2025).

real-time data slightly understated the magnitude of the swings in inflation.

3 An Example of Nature's Model

We have assumed that because the Fed suspects that its Phillips curve drifts, it uses a constant-gain algorithm. This section provides an example of a structural model in which the Fed's actions shape inflation persistence and the slope of the Phillips curve, a feedback loop that our section 4 model ignores, but to which its drifting coefficients model would nevertheless respond.[15] Thus, this section puts a model on the table that differs from the section 4 model that we will use to rationalize the Fed's interest rate decisions in the 2020's. We acknowledge that putting multiple models on the table is alien to a rational expectations theorist, but it is an essential part of research about learning about a rational expectations equilibrium, self-confirming equilibria, and decision making in contexts with unrecognized model misspecifications.[16]

[15]Our model in this section is an example of what Koopmans (1947) called a "Newton stage" model with a subset of parameters that are posited to be invariant to alternative hypothetical settings of the parameters in the Fed's Taylor rule.

[16]For examples of such work, please see Cho and Kasa (2017), Cho and Kasa (2015), Evans, Honkapohja, Sargent, and Williams (2013), Cho, Williams, and Sargent (2002), Sargent, Williams, and Zha (2006), and

3.1 A New Keynesian Model

Consider a New Keynesian model in which inflation and the output gap satisfy

$$\pi_t = \beta\,\mathbb{E}_t\pi_{t+1} \;+\; \gamma\,\pi_{t-1} \;+\; \kappa\,x_t \;+\; u_t, \tag{2}$$

$$x_t = \mathbb{E}_t x_{t+1} \;-\; \sigma(i_t - \mathbb{E}_t\pi_{t+1} - r_t^n)\,, \tag{3}$$

where u_t is an i.i.d. cost-push disturbance and r_t^n is an i.i.d. the natural-rate disturbance. Assume that the monetary authority sets the interest rate according to

$$i_t \;=\; \phi_\pi\,\pi_t, \qquad \phi_\pi > 0, \tag{4}$$

so that a larger ϕ_π indicates a more aggressive response to inflation. The i.i.d. shocks u_t and r_t^n are mutually orthogonal, have means zero and variances σ_u^2 and σ_r^2. We focus on minimum-state-variable (MSV) rational-expectations equilibria.[17]

In the MSV equilibrium of (2)–(4), inflation expectations are proportional to current inflation,

$$\mathbb{E}_t\pi_{t+1} \;=\; \lambda\,\pi_t, \qquad |\lambda| < 1,$$

where λ is the unique stable root of a cubic equation $F(\lambda, \phi_\pi) = 0$ implied by (2)–(4). Imposing MSV restrictions in (2)–(3) and applying the method of undetermined coefficients yields rational expectations equilibrium laws of motion

$$\pi_t = \lambda\,\pi_{t-1} + \sigma_{\pi,u}u_t + \sigma_{\pi,r}r_t^n, \tag{5}$$

$$x_t = \psi\,\pi_{t-1} + \sigma_{x,u}u_t + \sigma_{x,r}r_t^n. \tag{6}$$

Here the collection $(\psi, \sigma_{\pi,u}, \sigma_{\pi,r}, \sigma_{x,u}, \sigma_{x,r})$ are constants that are functions of the stable root λ and the structural parameters, and their values are provided in appendix A.

Two propositions summarize how econometrically measured inflation persistence and Phillips curve slope depend on the Taylor rule coefficient ϕ_π in a stationary MSV rational expectations equilibrium of nature's model (2)-(3).

Proposition 1 (Aggressive policy lowers measured inflation persistence)**.** *In the determinacy region of the MSV equilibrium, the stable root $\lambda(\phi_\pi)$ is strictly decreasing in the policy aggressiveness ϕ_π:*

$$\frac{d\lambda}{d\phi_\pi} \;<\; 0.$$

Sargent and Williams (2005).

[17] See McCallum (1983, 2004) for the discussion of an MSV representation.

Consequently, the measured AR(1) persistence of inflation in (5) is lower when policy is more aggressive.

For a proof see appendix A.

We now state a finding about the slope of the Phillips curve similar to one of Bullard (2018). Empirical analysts often estimate a (backward-looking) Phillips curve by regressing inflation on its lag and the contemporaneous gap. For example, consider a time-invariant version of (1) where π_t and x_t are mean zero, so there is no intercept:

$$\pi_t = \rho\pi_{t-1} + kx_t + \epsilon_t$$

Here we use k for the time-invariant slope coefficient to distinguish it from the structural parameter κ. Accordingly, define $\widehat{k}$ as the population OLS slope coefficient from regressing π_t on (π_{t-1}, x_t).

$$\widehat{k} = \frac{\mathrm{Cov}(\pi_t, x_t \mid \pi_{t-1})}{\mathrm{Var}(x_t \mid \pi_{t-1})} = \frac{\mathrm{Cov}(\pi_t, x_t) - \mathrm{Cov}(\pi_t, \pi_{t-1})\,\mathrm{Cov}(x_t, \pi_{t-1})/\mathrm{Var}(\pi_{t-1})}{\mathrm{Var}(x_t) - \mathrm{Cov}(x_t, \pi_{t-1})^2/\mathrm{Var}(\pi_{t-1})} \tag{7}$$

In Appendix A we develop three sets of sufficient conditions which allow us to sign the effect of policy aggressiveness on the measured slope of the Phillips curve. For example, we require the lagged term γ to be sufficiently small and the policy coefficient ϕ_π to be sufficiently large.

Proposition 2 (Aggressive policy flattens the measured Phillips-curve slope)**.** *Under conditions the sufficient conditions (S1), (S2), or (S3) in Appendix A, the OLS slope $\widehat{k}$ from π_t on (π_{t-1}, x_t) is strictly decreasing in policy aggressiveness. In particular, for sufficiently small γ and sufficiently large ϕ_π, more aggressive rules imply a flatter Phillips curve.*

For a proof see Appendix A.

Thus, an aggressive monetary policy makes a Phillips curve look flat and reduces the persistence of inflation.

3.2 A Markov-Switching Version

In the spirit of Svensson and Williams (2008, 2009) we extend our subsection 3.1 model by assuming that the Taylor rule coefficient ϕ_π depends on a Markov state variable $s_t \in \{1, 2\}$. Let $\phi_{\pi,1}$ and $\phi_{\pi,2}$ denote the policy response coefficients in Markov states 1 and 2, respectively. The state follows a first-order Markov chain with transition matrix:

$$P = \begin{bmatrix} P_{1,1} & P_{1,2} \\ P_{2,1} & P_{2,2} \end{bmatrix}, \tag{8}$$

where $P_{i,j} = \Pr(s_{t+1} = j \mid s_t = i)$ and each row sums to unity.

If we denote the vector of endogenous variables as $y_t = [\pi_t, x_t]'$ and the shock vector as $v_t = [u_t, \sigma r_t^n]'$, our model now becomes:

$$\Gamma_{-1}\mathbb{E}_t y_{t+1} + \Gamma_{0,s_t} y_t + \Gamma_1 y_{t-1} + v_t = 0, \tag{9}$$

where the 2×2 matrix Γ_{0,s_t} now depends on the current state through ϕ_{π,s_t}:

$$\Gamma_{0,s_t} = \begin{bmatrix} 1 & \sigma\phi_{\pi,s_t} \\ -\kappa & 1 \end{bmatrix}, \tag{10}$$

where mathematical expectation $\mathbb{E}_t(\cdot) \equiv \mathbb{E}(\cdot \mid s_t, y_t)$ is conditioned on both the current state s_t and the history of endogenous variables.

Outcomes are now described by the following two Markov-state-dependent equations:

$$y_t = G_{s_t} y_{t-1} + H_{s_t} v_t,$$

where (G_1, G_2) are 2×2 matrices governing the endogenous state dynamics in each regime, and (H_1, H_2) are 2×2 matrices describing contemporaneous impacts of shocks. Appendix A presents further analysis of the Markov-switching version of our New Keynesian model, as well as an algorithm for computing an equilibrium.

Simulations of our Markov-switching New Keynesian model generate time series for which parameters of the drifting coefficients model of section 4 indicate lower persistence and flatter Phillips curves during times of a more aggressive Taylor rule, and higher persistence and steeper Phillips curves during less aggressive periods. But the decision makers in the model that we use to rationalize Fed policy do not know that is why those coefficients are drifting. We turn to that model next.

4 Quantitative Application

We now put the ideas discussed above to work. This involves two steps: (i) taking real-time data and using them to estimate policymakers' beliefs about the Phillips curve with constant–gain recursive least squares (RLS), and then (ii) each period solving the linear-quadratic dynamic programming problem that the Fed uses to set that period's interest rate.

4.1 Belief Dynamics

The Fed's statistical model consists of two equations: a Phillips curve with drifting coefficients and an "IS curve" with constant coefficients:

$$\pi_t = \alpha_{0,t} + \rho_t \pi_{t-1} + \kappa_t x_t + \varepsilon_{\pi,t}, \tag{11}$$

$$x_t = b_0 + b_1\, x_{t-1} + g(i_{t-1} - \pi_{t-1}) + \varepsilon_{x,t}. \tag{12}$$

We estimate the model on U.S. quarterly data from 1960:Q1-2025:Q2 using PCE inflation π_t, the real-time output gap x_t, and the federal funds rate i_t. For the real-time output gap measure, we combine the Greenbook/Tealbook estimates from 1975-2019 with the CBO real-time data from 2020-2025.[18] As we do not have a consistent real-time measure for the earlier period, we use the ex-post realized data for 1960-1974. This matters little, given our focus on recent years. For the Phillips curve, we run constant-gain recursive least squares on (11), similar to what we did in section 2. For the IS curve, we estimate (b_0, b_1, g) once using the full sample. Recursively updating these parameters had minimal impact on the results.

Figure 2 in section 2 focused on the pre-pandemic evolution of beliefs, stopping the estimation at 2019:Q4. Now we have to take a stand on how to treat the pandemic observations and the post-pandemic period. Directly incorporating the data from the pandemic leads to large and implausible movements in beliefs, as the pandemic observations were extreme outliers unlike any seen in history. Our baseline assumption is that policymakers did not believe that the pandemic observations were informative about the baseline structural estimates. In practice, this means that we set the gain to zero ($\gamma = 0$) for observations in 2020 and 2021, so the beliefs over this period are fixed at pre-pandemic levels. Then in 2022, policymakers again return to updating beliefs. Our assumption does not mean that policymakers didn't respond to incoming data during 2020-2021, only that their beliefs about inflation dynamics and monetary transmission were locked in place.

Figure 6 reports the time series of belief estimates. For the most part they replicate the patterns emphasized earlier using ex-post revised data: a downward drift in persistence rho_t since the mid-1980s (with a post-2010 trough), and a flattening of the slope κ_t toward zero in the 2010s–2020s (with a temporary spike around 2009–10). The intercept absorbs level shifts not explained by slack or lags.[19]

Unlike our earlier analysis, we now include the pandemic and post-pandemic periods. With the resumption of belief updating in 2022, the inflation persistence parameter ρ_t jumps

[18] Nakamura, Riblier, and Steinsson (2025) constructed a similar real-time output gap measure.

[19] If the intercept term were restricted to $\alpha_{0,t} = (1 - \rho_t)\pi^*$ then (11) would enforce that π_t be centered at π^*. Imposing such a restriction had minimal impact on our results.

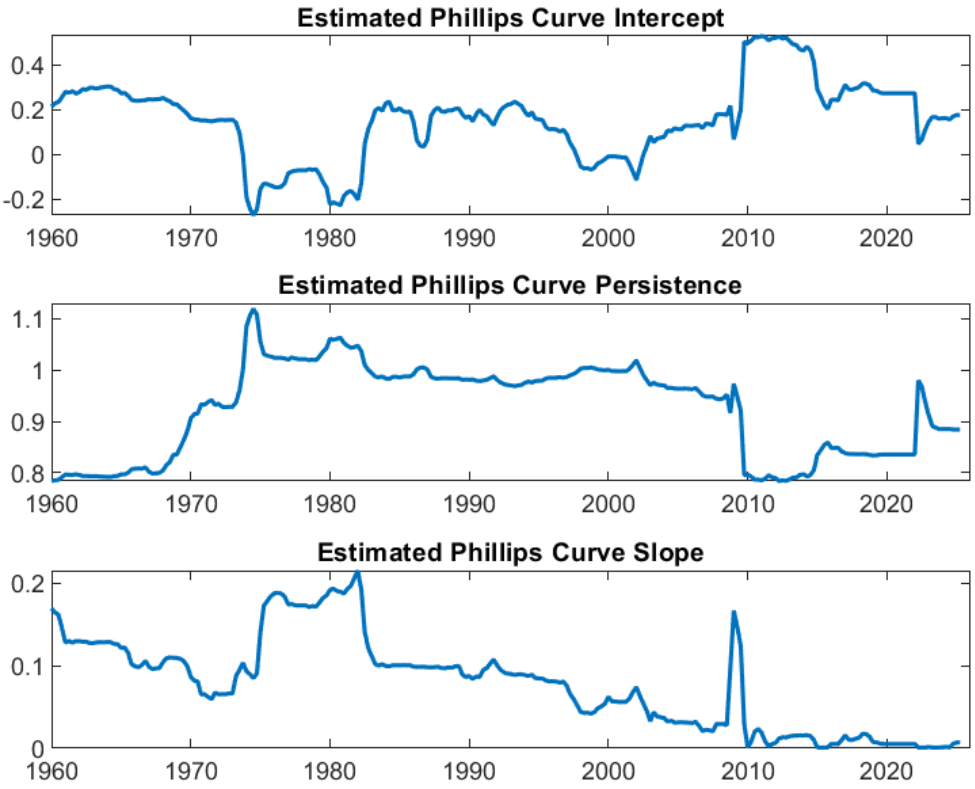


Figure 6: Beliefs estimated from U.S. data 1960:Q1–2025:Q2. Intercept (top), persistence ρ_t (middle), slope κ_t (bottom).

up to near one. As we see below, this was roughly in line with the Fed moving away from the characterization of the pandemic inflation as "transitory" and beginning to raise interest rates.

4.2 The Fed's Phelps Problem

First we define the augmented state vector

$$X_t \equiv [1,\ \pi_t,\ x_t,\ i_{t-1}]',$$

where the constant is necessary because the belief dynamics have nonzero intercepts. Then given beliefs $(\alpha_{0,t}, \rho_t, \kappa_t)$ and the constant IS curve parameters (b_0, b_1, g), we can combine (11)-(12) to write the state dynamics at time t as

$$X_{s+1} = A_t X_s + B_t i_s + C\varepsilon_{s+1} \tag{13}$$

for $s \geq t$. Here ε_t is the combined shock vector and the matrices A_t and B_t depend on the time-t belief estimates.

As in Kreps (1998), Cogley and Sargent (2008), and many other papers that bring drifting coefficients models into macroeconomics, we study an anticipated utility problem where the beliefs are anticipated to remain constant looking forward. Each period, policymakers update their beliefs based on observed data, then use these fixed beliefs to forecast future outcomes. The policymaker chooses the policy rate i_t to solve:

$$\min_{\{i_s\}_{s\geq t}} \ \mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} \Big[\ (\pi_s - \pi^*)^2 + \ \lambda_x x_s^2 + \eta \, (i_s - i_{s-1})^2 \ \Big], \tag{14}$$

subject to the perceived linear state dynamics (13). Here $\beta \in (0,1)$ is a discount factor, π^* is the inflation target, $\lambda_x > 0$ is the loss function weight on output gap, $\eta \geq 0$ is the weight on interest rate changes. We include this smoothing term in the objective function to match the tendency for incremental and smooth policy adjustments in practice.

If we denote the period loss function ℓ_t, we can write:

$$\ell_t = (\pi_t - \pi^\star)^2 + \lambda_x x_t^2 + \eta \, (i_t - i_{t-1})^2 = X_t' Q X_t \ + \ 2X_t' N i_t + i_t' R i_t + L_0,$$

where the constant term L_0 is irrelevant for policy. This problem is a standard linear–quadratic regulator with a cross-term in the state and control. The optimal value function takes the

form

$$V_t(X) = X'P_tX + V_{0t}$$

where the matrix P_t solves a discounted algebraic Riccati equation (with the t indexing the state matrices). The constant V_{0t} depends on the perceived volatility of the shock vector, but is irrelevant for the policy choice. The optimal policy rule is

$$\begin{aligned} i_t &= -F_tX_t \\ &= \phi_{0,t} + \phi_{i,t}i_{t-1} + \phi_{\pi,t}\pi_t + \phi_{x,t}x_t. \end{aligned}$$

where

$$F_t = (R + \beta B_t'P_tB_t)^{-1}\big(B_t'P_tA_t + N'\big),$$

and the second line uses the definition of the state vector to write the solution as an interest rate rule. Thus at each date t, given the current estimates, the optimal interest rate policy is a smoothed Taylor rule. As the Fed's beliefs change, the coefficients of its policy rule change.

4.3 Phelps Problem Details

We assume that each period policymakers compute the optimal policy rule, given their beliefs and assuming commitment. Then the current policy recommendation is taken from the current period's optimal decision rule. Specifically, we assume that at each date t policymakers take the following steps:

1. Update the Phillips curve coefficients $(\alpha_{0,t}, \rho_t, \kappa_t)$ via recursive least squares.
2. Form the state matrices (A_t, B_t) using the current estimates.
3. Solve the dynamic programming problem to obtain F_t and the policy recommendation:

$$i_t^{\text{opt}} = \max\{-F_tX_t, 0\}.$$

Note that we impose a hard zero-lower bound ex-post, meaning that, looking forward policymakers do not anticipate hitting up against the zero lower bound. In the next sections, we compare the subjectively optimal policy i_t^{opt} to the actual federal funds rate.

4.4 Results

The following figures use the parameters: the gain $\gamma = 0.03$, discount factor $\beta = 0.95$, inflation target $\pi^* = 2\%$, weight on the output gap in the loss function $\lambda_x = 0.2$, interest

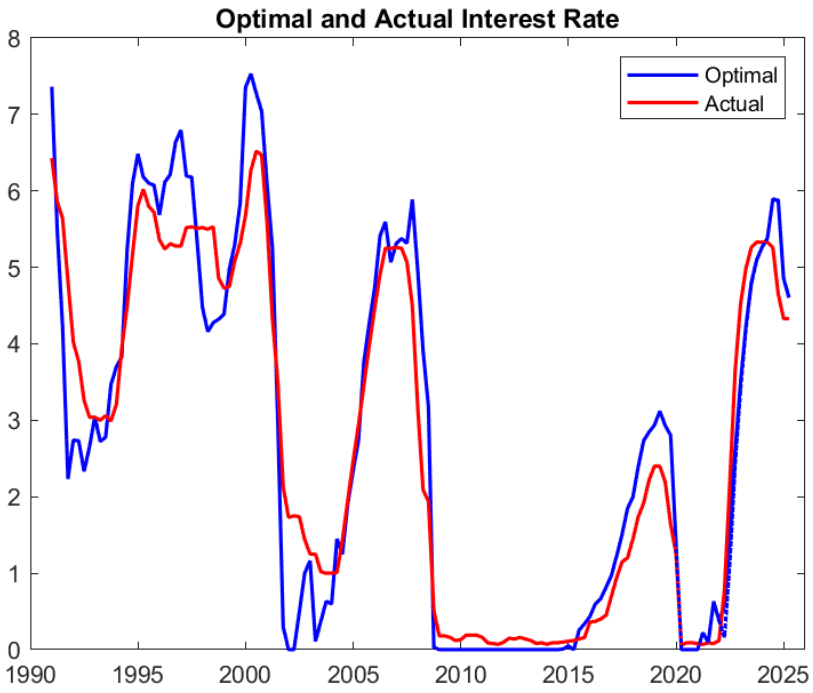


Figure 7: Optimal vs. actual policy. Interest rate recommendations from the Phelps problem (blue) vs. the actual federal funds rate (red), 1991:Q1–2025:Q2.

rate smoothing penalty in the loss function $\eta = 0.5$. The loss function parameters are loosely tuned to provide a reasonable match to the observed policy outcomes. The interest smoothing parameter η seems to be one of the most important parameters. A lack of smoothing gives an implausibly volatile interest rate series that responds very abruptly and aggressively to movements in inflation or the output gap. On the other hand, a larger smoothing penalty gives a better fit to the observed fed funds data. But this fit comes at the cost of interpretability, as the recommended policy comes to just approximate the lagged interest rate. The chosen value strikes a balance between fit and interpretability.

Figure 7 plots the recommended rate i_t^{opt} and the actual funds rate from 1991-2025. The subjectively optimal policy tracks the level and turning points of policy closely over the last three decades. The figure shows that the policy recommendations capture the mid-1990s tightening cycle, the easing following the 2001 recession and subsequent tightening, the zero bound era following the 2008 recession, and the normalization of rates beginning in 2015. Notably, the subjectively optimal policy shows minimal change in the recommended policy rate during the period of increasing inflation in 2021, slightly lagging the actual policy during the tightening which began in earnest in 2022. The overshoot in the recommended policy rate at the end of the sample came from the rapid recovery in the real-time output gap

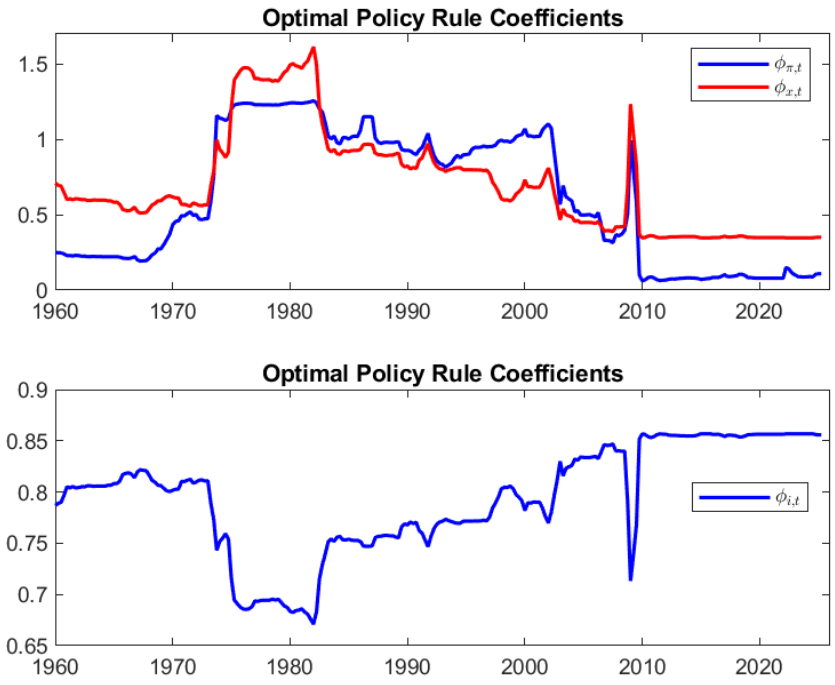


Figure 8: Optimal policy rule coefficients from the Phelps problem. Top panel: response to inflation ($\phi_{\pi,t}$) and the output gap ($\phi_{x,t}$). Bottom panel: smoothing coefficient on the lagged interest rate ($\phi_{i,t}$).

measure, which overshot the ex-post output gap in late 2024. While differences are visible, the belief-driven rule provides a remarkably good summary of actual policy.

Figure 8 shows the time series of the optimal policy rule coefficients from the Phelps problem. The top panel shows the response to inflation ($\phi_{\pi,t}$) and the output gap ($\phi_{x,t}$), while the bottom panel shows the smoothing coefficient on the lagged interest rate ($\phi_{i,t}$). In general, they follow the trends which we expected given the patterns of the belief coefficients. Broadly speaking, there is a downward trend after 1980 in the response of policy to both the output gap and inflation, and an increase in the interest rate smoothing. After 2010 recession there has been relatively little change in these policy response coefficients. Over this time span, policy has been characterized by strong interest rate smoothing, a mild response to output gap fluctuations, an especially weak response to inflation. There was a mild uptick in the response to inflation $\phi_{\pi,t}$ in 2022 during the tightening episode. But a notable part of the interest rate changes ($i_t = -F_t X_t$) came from the observed changes in the inflation and output gap state variables (X_t), not the policy reaction function (F_t).

To quantify the role of changing beliefs about inflation dynamics, we recompute policy recommendations over the same time period, but now keep beliefs fixed at their January

2000 values throughout. From Figure 6, we see that in early 2000 the inflation persistence parameter ρ_t was still quite high, having declined somewhat since the 1980s but not yet fallen dramatically as it would thereafter. Similarly, the slope of the Phillips curve κ_t was about halfway through its long-term decline, well below values from the 1980s but well above those in the 2010s. From our discussion and illustration in Section 3 above, we thus expect that the (now fixed) optimal policy rule from January 2000 should exhibit a stronger response to inflation than the policy with beliefs from 2021.

Figure 9 shows the time series of this counterfactual policy rule with fixed beliefs, along with the subjectively optimal rule with updated beliefs that we had seen in Figure 7. For January 2000, the two time series are equal by construction, and in fact for much of the sample the counterfactual and baseline rules are similar. Early in the sample, the counterfactual policy was looser than the baseline, due largely to the lower trend inflation estimate (from the lower intercept $\alpha_{0,t}$) from the fixed beliefs. But then from 1999-2019 the two time series closely match each other.

However, the two policy rules differ dramatically in their response to the post-pandemic inflation. The counterfactual policy rule moves much more quickly and sharply to tighten policy in response to the rise in inflation. Even with real-time uncertainty about the state of the economy, there is no delay in the policy response to the inflation surge. Instead, the recommended policy jumps rapidly to a federal funds rate of 4.5% by the end of 2021, increasing to over 6% in 2022. This illustrates that perceptions of a less persistent inflation and a flatter Phillips curve led to a much weaker and delayed policy response.

4.5 Using Ex-Post Revised Data

To illustrate the role of real-time data uncertainty, we re-run the entire exercise replacing the real-time output gap series with latest-vintage (ex-post revised) data. The Phillips-curve beliefs $\{\alpha_{0,t}, \rho_t, \kappa_t\}$ are re-estimated each quarter with the same constant-gain recursive least squares procedure. The IS relation (b_0, b_1, g) is re-estimated with the ex-post output gap series, but kept constant over time as before. The belief estimates through 2019 are those that we reported in Figure 2 above. The rest of the parameters in the loss function and belief estimation are left unchanged. We then compute Phelps recommendations as above, where now the state X_t includes the ex-post output gap measure.

With revised data, the estimated belief paths display the broad trends that we have already outlined above: a decline in the persistence of inflation and the slope of the Phillips curve. Relative to what we saw discussing Figure 3 in Section 2, the key difference is that with the real-time data the pandemic recession appeared longer-lasting than it appeared

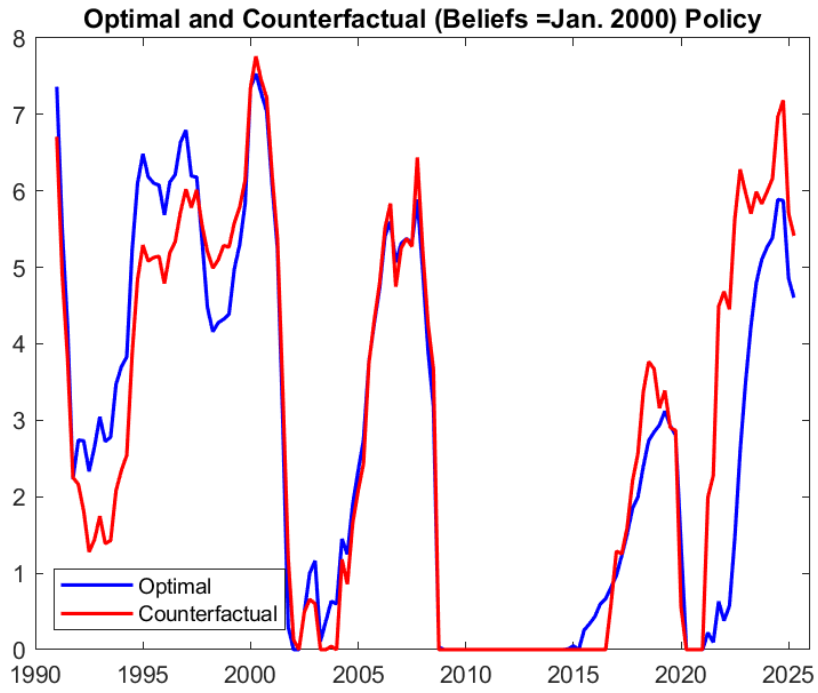


Figure 9: Counterfactual with beliefs fixed at their values from January 2000. The baseline Phelps recommendations (blue) vs. counterfactual with fixed beliefs (red).

to have been in the ex-post data. The ex-post output gap is consistently higher than its real-time counterpart in the early 2020s. Feeding this higher gap into the system raises the model's recommended policy somewhat earlier in 2021.

Figure 10 plots the Phelps problem's recommended rate (blue) against the actual funds rate (red). The overall picture is much the same as in Figure 7: the Phelps problem's decision rule tracks the level and turning points of the actual policy closely. Relative to the real-time version, the model with ex-post data more closely matches the policy normalization from 2015-2019, suggesting that perhaps the Greenbook output gap measure did not fully capture policymaker assessments. But more relevant for our purposes, the Phelps problem's recommended policy now suggests an immediate response to the inflation surge that began in 2021. The initial response is mild, with interest rates hitting 1.5% by the end of 2021, which again reflects the subjective inflation dynamics that we have emphasized. Only later in 2022 does the Phelps problem recommend more substantial rate increases, with timing and magnitude roughly in line with observed Fed choices.

Doing the same counterfactual exercise as above, fixing beliefs at their January 2000 values, provides even stronger contrasts with both the baseline from the model and actual policy. Figure 11 shows the same comparison as in Figure 9 above, but now with the ex-post

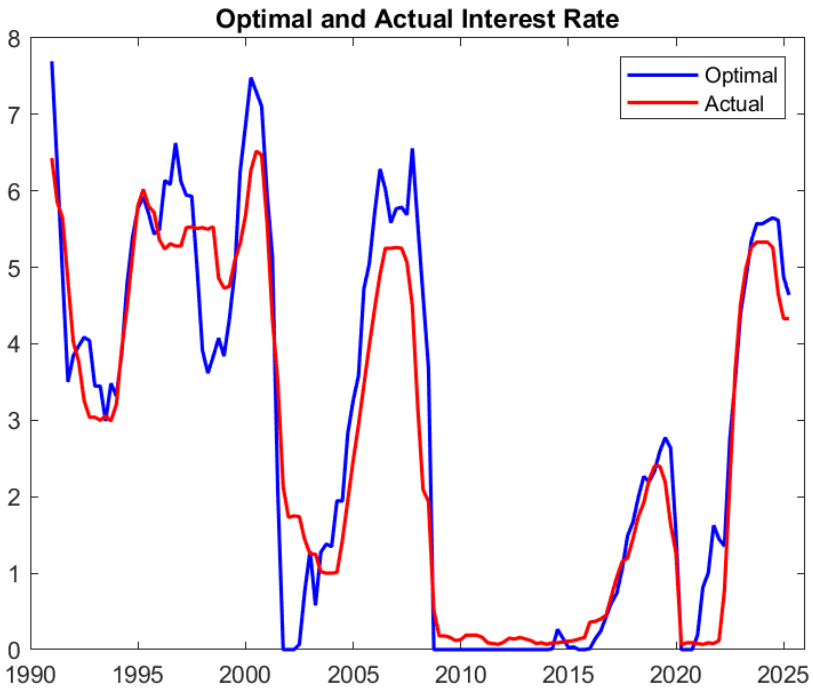


Figure 10: Optimal vs. actual policy using ex-post data. Phelps recommendations (blue) computed with revised data vs. actual federal funds rate (red), 1991:Q1–2025:Q2.

data. Again, from 1999-2019 the fixed-belief counterfactual matches the Phelps problem recommendations with updated beliefs. But now the response with fixed beliefs to the post-pandemic inflation is even larger and more dramatic, relative both to the baseline with real-time beliefs and to the ex-post baseline in Figure 10. The policy rate under January 2000 beliefs now hits 7% by the end of 2021 and eventually tops out at nearly 8%.

The first two elements of our story relate to changes in inflation dynamics, and these hold true whether we use real-time or ex-post data. Low inflation persistence and a flat slope of the Phillips rationalize a mild response to an increase in inflation. But even with changed subjective beliefs, an accurate picture of the economy would have recommended *some* policy response, however muted, to the inflation in 2021. Only by incorporating real-time uncertainty about the economy, which tends to be elevated during business cycle turning points, do we rationalize the observed delay in the policy response. Because the real-time assessments suggested that the economy was in worse shape than the ex-post data revealed, policymakers delayed tightening. Both real-time uncertainty and changing beliefs shaped the timing and aggressiveness of the Fed's policy responses.

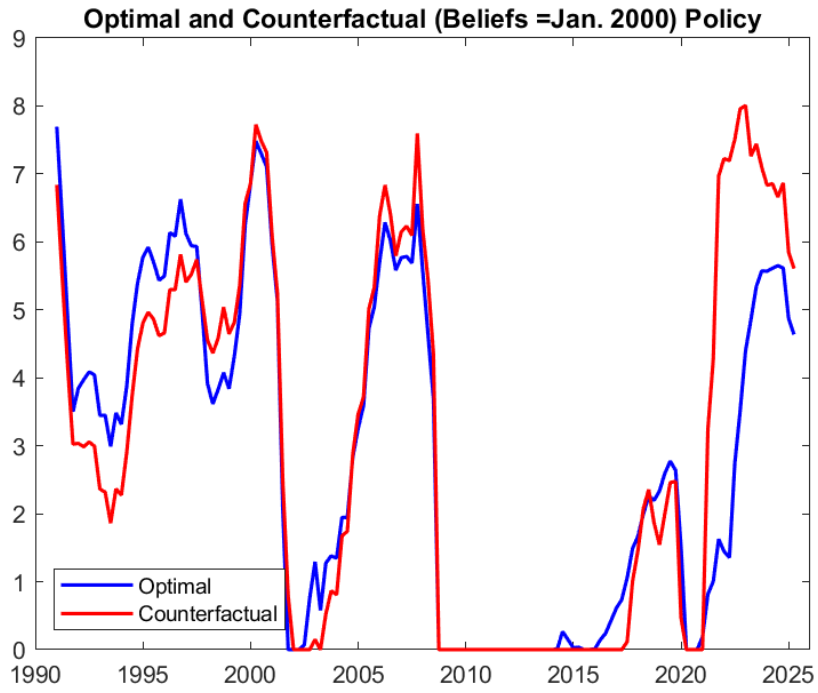


Figure 11: Counterfactual with fixed beliefs using ex-post data. Baseline Phelps recommendations (blue) vs. recommendations if beliefs are frozen at January 2000 values (red).

5 Policy Makers' Forecasts

The previous section showed that our recommended policy from the model was able to match observed policy relatively well. Here we show that belief dynamics from our model lead to time-varying forecasts that match the forecasts reported by policymakers.

5.1 The FOMC Summary of Economic Projections

The Federal Reserve's *Summary of Economic Projections* (SEP) reports, for each projection round, the distribution of FOMC participants' judgments for key indicators for monetary policy. In each release, the SEP reports the end-of-year annual growth rates (4th quarter over the preceding year's 4th quarter) of real GDP and the PCE price index, along with the end-of-year values of the unemployment rate and (since 2012) the appropriate federal funds rate. Each SEP report provides projections for three or four years, along with a value for the longer run.

The SEP provides a measure of forecasts that underlie the Federal Reserve policymakers' decisions. We compiled the median SEP projections from 2007 to the present, which are shown along with the actual data in Figures 12 and 13. In each panel of each figure, we plot

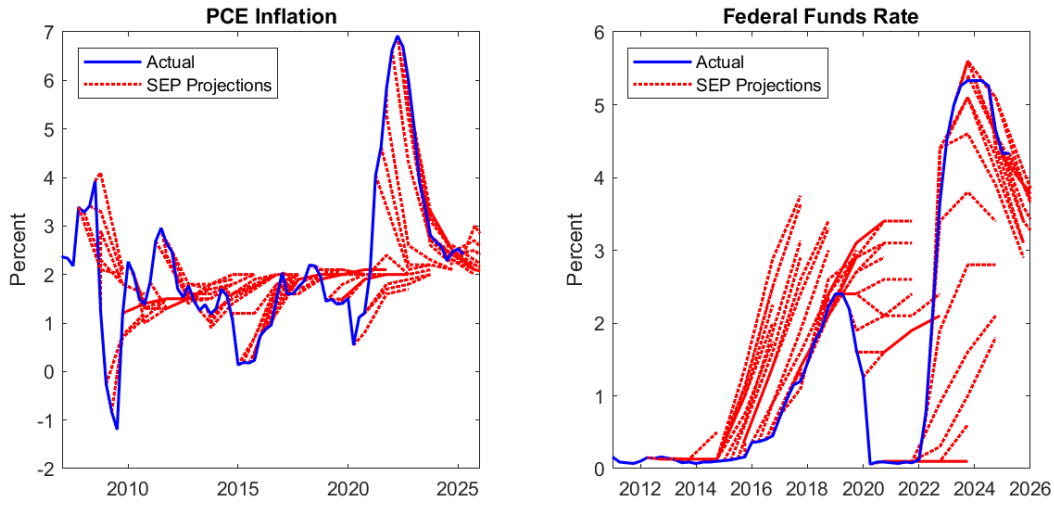


Figure 12: Median from FOMC Summary of Economic Projections (red dash) for each projection date, along with the actual data (blue). PCE inflation (left panel) and federal funds rate (right panel).

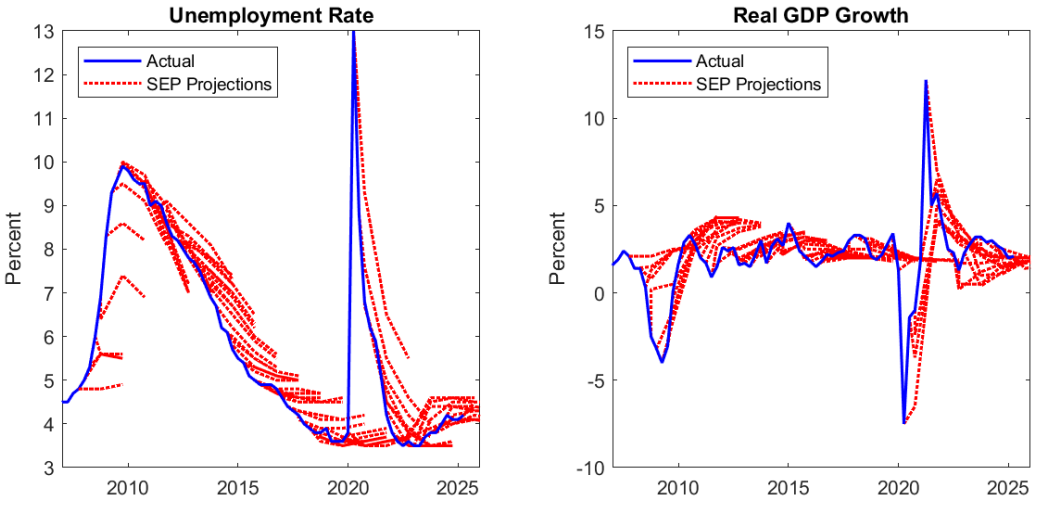


Figure 13: Median from FOMC Summary of Economic Projections (red dash) for each projection date, along with the actual data (blue). Unemployment rate (left panel) and real GDP growth rate (right panel).

the actual quarterly data as a solid blue line. The red dashed lines in the figures are the SEP medians, which branch off from the actual data at the date of the SEP release, with the projection values plotted at the end of the appropriate years.[20]

In Figure 12, we see that both the inflation and federal funds rate projections show strong reversion to long-run levels. For inflation, at almost every point in the sample the median projection showed reversion to near the 2% target level by the end of the projection period. In the recent inflation episode, even as the current inflation rate was racheting upward, the Fed was projecting it to fall going forward.

For the federal funds rate, the right panel shows that throughout 2014 and 2015 the FOMC participants expected to normalize policy and raise rates off the zero lower bound, but that normalization only began gradually in 2016. More recently, in March 2021, the Fed was projecting that the federal funds rate would remain at zero through the end of 2023. Later in 2021, the out-year projections increased slightly. Then throughout the tightening phase that began in 2022, the projections consistently underestimated the pace and degree of policy tightening that would take place.

While the projections of the monetary variables show strong reversion, the projections of the real variables in Figure 13 show much less reversion. For most of the sample, the unemployment rate projections in the left panel are closer to a random walk. During the labor market recoveries following the Great Recession and the COVID recession, the SEP consistently underestimated the magnitude and the rapidity of the declines in unemployment. In the right panel, we see that the SEP showed a stronger and more prolonged real GDP growth rate bounce-back following both of those recessions.

5.2 Comparisons at Key Dates

We now compare the subjective forecasts implied by our model with the projections from the SEP. To compare the SEP with our model, we freeze beliefs at the projection date and then forecast forward using the subjective state dynamics (13) under the optimal policy implied by those beliefs. This mirrors the information set of participants at each SEP round. Since PCE inflation and the federal funds rate are variables in our model, these provide the most direct comparisons.

Figure 14 plots the model's forecasts (solid lines) against SEP medians (dashed with markers) at six representative dates during the recent inflation episode. For each projection date, beliefs are frozen and the current optimal policy is applied to generate 5 years of model forecasts, shown as solid lines. The blue curves are inflation paths, the red curves are federal

[20] Farmer, Nakamura, and Steinsson (2024) show a similar plot for the SPF forecasts of the 3-month T-bill rate.

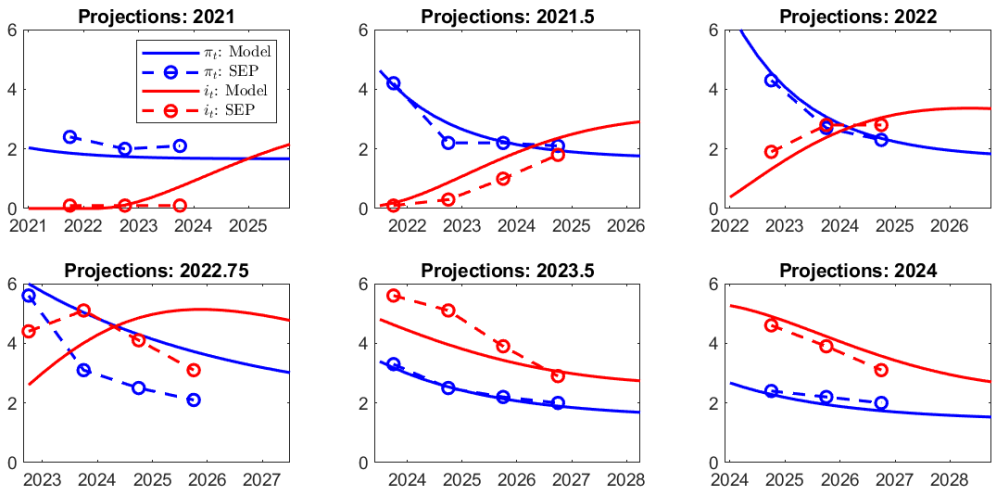


Figure 14: Model forecasts and SEP projections at selected dates. For each projection date, beliefs are frozen and the current optimal policy is applied to generate 5 years of model paths (solid). SEP medians (dashed circles) report Q4/Q4 PCE inflation (blue) and year-end funds rate (red) at the years shown.

funds paths. The SEP medians (dashed circles) report Q4/Q4 PCE inflation (blue) and the year-end funds rate (red) at the years shown.

The model's forecasts closely track the SEP medians throughout the episode, capturing both the level and shape of the projections. The first panel shows that in 2021:Q1 both the model and the SEP projected inflation to remain flat near the 2% target, and the policy rate to remain near zero for an extended period, with the model forecasting a policy liftoff in later years. The next two panels show that as inflation increased (the blue lines start progressively higher), it was projected to be transitory and return to target without substantial policy tightening. Once tightening began, the model's projections lagged a bit (as seen the bottom left panel), and both the model and SEP projected a smooth decline in inflation with a mild hump-shaped policy response. By 20203 and into 2024, the model caught back up to the SEP, with both projecting inflation on a glide path back to 2% with gradual policy easing.

Across the projection dates shown in Figure 14, the model replicates the general shape and level of SEP projections for inflation and the federal funds rate. Overall, these projections illustrate what our mechanism implies: low perceived persistence and a flat Phillips curve slope initially rationalize patience. As realized inflation persists and beliefs update, the implied optimal response strengthens, helping to guide inflation back to target. An optimizing model with drifting-beliefs not only reasonably matches the realized policy path, but also provides an informative representation of stated policymaker beliefs.

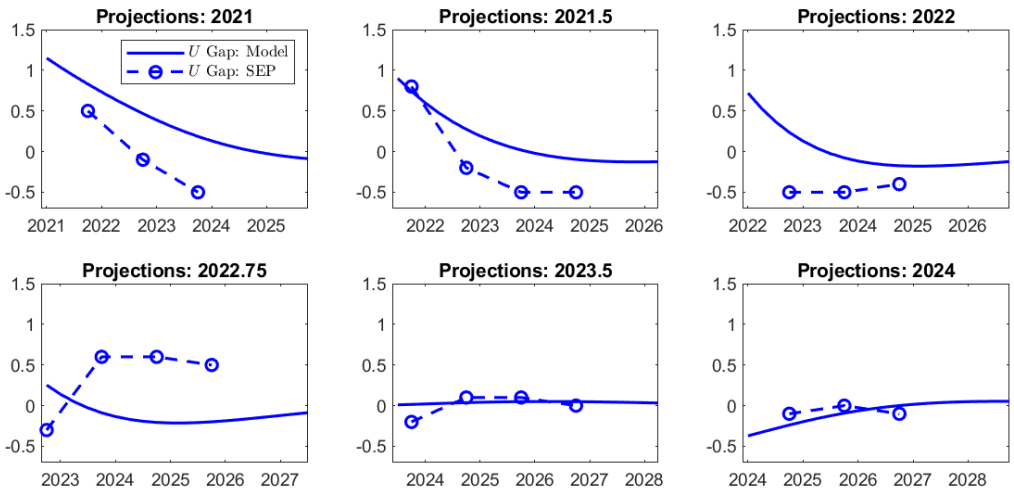


Figure 15: Model vs. SEP unemployment gap projections at selected dates. The model's output gap forecasts are converted to the unemployment gap using an estimated Okun's law relationship (solid lines). The SEP median unemployment projection is converted to an unemployment gap by subtracting the long-term projected unemployment rate (dashed circles).

5.3 Unemployment Projections

Our model produces a path for the output gap, measured as percent deviations of real GDP from potential output. To compare with the FOMC's projections for the unemployment rate, we map the output gap into the unemployment gap using an estimated Okun's law.

In particular, using our real-time output gap time series x_t, along with data on the unemployment rate U_t and the real-time measure of the natural rate of unemployment U_t^*, we estimate:

$$U_t - U_t^* = \theta_0 + \theta_1 x_t + \epsilon_t^U$$

For this exercise, we run the regression over the period 1990-2019, during which this Okun's law relationship was relatively stable. Different time windows had minimal impact on our results.

For the SEP, we take the reported median long-term unemployment rate as a measure of the natural rate and compute the implied unemployment gap projections. Then we use the estimated (θ_0, θ_1) coefficients to convert the model's output gap forecasts into the unemployment gap.

Figure 15 reports the results. In broad outline, the model's converted unemployment gap forecasts match the general trends of those from the SEP. But the fit is much less close than in the variables where we have direct comparisons. One discrepancy is that the model

predicts stronger return to a zero unemployment gap at all horizons, while for many of the projection dates the FOMC members projected unemployment to deviate from the natural rate for an extended period of time. For example, the shape of the paths from the model and the SEP in the top row, for 2021 to early 2022, look similar, but the levels differed. Both were generally forecasting a decline in unemployment, but the Fed was projecting unemployment to remain half of a percentage point below its long-term rate at the end of the projection period. Then in 2022:Q4 (the bottom left panel) the Fed was projecting an increase in unemployment in the later projection years which did not materialize. By late 2023 both the model and SEP were projecting little further change in the unemployment rate.

While our model does well in capturing the Fed's subjective forecasts of inflation and its policy rate, this simple Okun's law conversion misses some aspects of the Fed decision makers' beliefs about post-pandemic labor market prospects.

6 Concluding Remarks

We have constructed a quantitative model in which the Fed's evolving opinions about the parameters of the Phillips curve affect its choices about how to keep inflation low and aggregate output growth high. Our model resides within a 21st century "monetary economics with no money" tradition that Lagos (2025) challenges. Statements about the money supply, government deficits, and government debt are absent not only from our drifting coefficients, anticipated utility Phelps problem quantitative model of section 4 but also from the section 3 New Keynesian model that we use to indicate how, though unrecognized by the Phelps problems that determine the Fed's actions, those actions eventually influence the Phillips curve slope and persistence parameters that shape them. Our story aligns with Phillips-curve-centric accounts by Fed insiders, some of which we have quoted, as well as with prominent Fed officials' public dismissals concerns about constraints that government deficits impose on Fed decisions. Thus, after describing Alan Greenspan's forceful public expressions of concerns about Federal budget deficits, Ben Bernanke wrote "From today's perspective, Greenspan's involvement in fiscal matters looks not only like a political overreach but also like an analytical error, as recent experience, as well as academic research, suggests that, in an advanced economy like the United States, the economic risks of moderate government deficits are low." Bernanke (2022, p. 52). He went on to write: "With fiscal dominance, helicopter money is inflationary. However, we are far from such a situation in the United States today." Bernanke (2022, p. 363). For better or worse, along with Bernanke (2022), our rationalization of Fed interest rate setting in the 2020's has much company in sweeping

concerns about fiscal policy into the background.[21]

Other accounts of U.S. inflation in the 2020's put fiscal policy back on the table. Thus, Bianchi, Faccini, and Melosi (2023), Hall and Sargent (2022, 2025), Andolfatto and Martin (2025), and Cochrane (2025) emphasize fiscal consequences and purposes of 2020's monetary accommodation. Those papers raise the possibility that the Fed purposefully used both unanticipated and anticipated inflation to accommodate prospective fiscal deficits. That is the starting point for an alternative quantitative rationalization of the Fed's decisions during the 2020's.

Appendixes

A New Keynesian Model

This appendix presents further analysis of our section 3 New Keynesian model.

A.1 Minimum State Variable Solution

To find the minimum state variable (MSV) solution, we first zero out the shocks from the New Keynesian model in section 3, with $i_t = \phi_\pi \pi_t$:

$$\begin{aligned}\pi_t &= \beta\, \mathbb{E}_t \pi_{t+1} + \gamma\, \pi_{t-1} + \kappa x_t,\\ x_t &= \mathbb{E}_t x_{t+1} - \sigma(\phi_\pi \pi_t - \mathbb{E}_t \pi_{t+1})\,.\end{aligned}$$

From the New Keynesian Phillips curve:

$$x_t = \frac{\pi_t - \beta \mathbb{E}_t \pi_{t+1} - \gamma \pi_{t-1}}{\kappa} \quad \Rightarrow \mathbb{E}_t x_{t+1} = \frac{\mathbb{E}_t \pi_{t+1} - \beta \mathbb{E}_t \pi_{t+2} - \gamma \pi_t}{\kappa}.$$

We then substitute into the IS equation:

$$\pi_t - \beta \mathbb{E}_t \pi_{t+1} - \gamma \pi_{t-1} = \mathbb{E}_t \pi_{t+1} - \beta \mathbb{E}_t \pi_{t+2} - \gamma \pi_t - \kappa\sigma\phi_\pi \pi_t + \kappa\sigma \mathbb{E}_t \pi_{t+1}.$$

Collect terms to obtain

$$\beta\, \mathbb{E}_t \pi_{t+2} - (\beta + 1 + \kappa\sigma)\mathbb{E}_t \pi_{t+1} + (1 + \gamma + \kappa\sigma\phi_\pi)\pi_t - \gamma \pi_{t-1} = 0.$$

[21] This was not always so. Silber (2012, chs. 12-16) contains extended accounts of how Paul Volcker and Alan Greenspan cared and were about outspoken about constraints that fiscal policy put on the Fed. For good examples, see Silber (2012, pp. 207, 214, 215, 235, 250, 270). See Bassetto and Sargent (2020) for a presentation of the "old-time religion" that emphasizes connections between fiscal policy and money supplies.

Impose the MSV conjecture $\mathbb{E}_t\pi_{t+1} = \lambda\pi_t$, $\mathbb{E}_t\pi_{t+2} = \lambda^2\pi_t$, and $\pi_t = \lambda\pi_{t-1}$, divide by π_t, and multiply by λ:

$$F(\lambda, \phi_\pi) \equiv \beta\lambda^3 - (\beta + 1 + \kappa\sigma)\lambda^2 + (1 + \gamma + \kappa\sigma\phi_\pi)\lambda - \gamma = 0. \tag{15}$$

Let the roots of $F(\lambda, \phi_\pi) = 0$ be $\{\lambda_i\}_{i=1}^3$. In this model the only predetermined state is π_{t-1}, so the MSV criterion requires exactly one root inside the unit circle, $|\lambda_i| < 1$.

To derive the shock loadings, we now consider the stochastic model with $i_t = \phi_\pi\pi_t$:

$$\begin{aligned}\pi_t &= \beta\,\mathbb{E}_t\pi_{t+1} + \gamma\,\pi_{t-1} + \kappa\,x_t + u_t,\\ x_t &= \mathbb{E}_t x_{t+1} - \sigma(\phi_\pi\pi_t - \mathbb{E}_t\pi_{t+1} - r_t^n)\,.\end{aligned}$$

Combining equations as above gives:

$$\beta\,\mathbb{E}_t\pi_{t+2} - (\beta + 1 + \kappa\sigma)\,\mathbb{E}_t\pi_{t+1} + (1 + \gamma + \kappa\sigma\phi_\pi)\,\pi_t - \gamma\,\pi_{t-1} \;=\; u_t - \kappa\sigma\,r_t^n.$$

Then using the MSV form to eliminate expectations gives:

$$\Big[\beta\lambda^2 - (\beta + 1 + \kappa\sigma)\lambda + (1 + \gamma + \kappa\sigma\phi_\pi)\Big]\pi_t - \gamma\,\pi_{t-1} = u_t - \kappa\sigma\,r_t^n.$$

Using $F(\lambda, \phi_\pi) = 0$ and dividing by λ gives $\beta\lambda^2 - (\beta + 1 + \kappa\sigma)\lambda + (1 + \gamma + \kappa\sigma\phi_\pi) = \gamma/\lambda$, so the term in the bracket equals γ/λ. Hence the inflation law of motion (5) is:

$$\pi_t = \lambda\,\pi_{t-1} \;+\; \frac{\lambda}{\gamma}\,u_t \;-\; \frac{\kappa\sigma\,\lambda}{\gamma}\,r_t^n$$

Using the Phillips curve to solve for x_t then substitute for π_t yields (6):

$$x_t = \frac{(1 - \beta\lambda)\pi_t - \gamma\pi_{t-1} - u_t}{\kappa} = \frac{(1 - \beta\lambda)\lambda - \gamma}{\kappa}\,\pi_{t-1} + \frac{(1 - \beta\lambda)\lambda/\gamma - 1}{\kappa}\,u_t - \frac{(1 - \beta\lambda)\lambda}{\gamma}\,\sigma\,r_t^n.$$

A.2 Measured Inflation Persistence

Proposition 3 (Aggressive policy lowers measured inflation persistence)**.** *In the determinacy region of the MSV equilibrium, the stable root $\lambda(\phi_\pi)$ is strictly decreasing in the policy aggressiveness ϕ_π:*

$$\frac{d\lambda}{d\phi_\pi} \;<\; 0.$$

Consequently, the measured AR(1) persistence of inflation in (5) is lower when policy is more aggressive.

Sketch. We have defined $F(\lambda, \phi_\pi) = 0$ defined in (15). By the Implicit Function Theorem,

$$\frac{d\lambda}{d\phi_\pi} = -\frac{\partial F/\partial \phi_\pi}{\partial F/\partial \lambda} = -\frac{\kappa\sigma\,\lambda}{3\beta\lambda^2 - 2(\beta + 1 + \kappa\sigma)\lambda + (1 + \gamma + \kappa\sigma\phi_\pi)}.$$

In the standard determinacy region (unique stable root), the denominator at the stable root is positive, so

$$\frac{d\lambda}{d\phi_\pi} < 0$$

□

A larger ϕ_π tightens the feedback on expectations, shrinking the MSV root λ toward zero. Measured inflation falls accordingly.

A.3 Measured Phillips Curve Slope in Regressions

Empirically, as in the literature discussed above, analysts often estimate a (backward-looking) Phillips curve by regressing inflation on its lag and the contemporaneous gap.

Let $\widehat{k}$ be the population OLS coefficient on x_t in the regression of π_t on (π_{t-1}, x_t).

$$\pi_t = \rho\pi_{t-1} + kx_t + \varepsilon_t.$$

By the Frisch-Waugh-Lovell Theorem

$$\widehat{k} = \frac{\mathrm{Cov}(\pi_t, x_t) - \mathrm{Cov}(\pi_t, \pi_{t-1})\,\mathrm{Cov}(x_t, \pi_{t-1})/\mathrm{Var}(\pi_{t-1})}{\mathrm{Var}(x_t) - \mathrm{Cov}(x_t, \pi_{t-1})^2/\mathrm{Var}(\pi_{t-1})}.$$

Using the MSV laws of motion (5)-(6) and the orthogonality of shocks:

$$\mathrm{Cov}(\pi_t, \pi_{t-1}) = \lambda\,\mathrm{Var}(\pi_{t-1}), \qquad \mathrm{Cov}(x_t, \pi_{t-1}) = \psi\,\mathrm{Var}(\pi_{t-1}),$$

$$\mathrm{Cov}(\pi_t, x_t) = \lambda\psi\,\mathrm{Var}(\pi_{t-1}) + \sigma_{\pi,u}\sigma_{x,u}\,\sigma_u^2 + \sigma_{\pi,r}\sigma_{x,r}\,\sigma_r^2, \qquad \mathrm{Var}(x_t) = \psi^2\,\mathrm{Var}(\pi_{t-1}) + \sigma_{x,u}^2\,\sigma_u^2 + \sigma_{x,r}^2\,\sigma_r^2.$$

The $\mathrm{Var}(\pi_{t-1})$ terms cancel, giving the closed form

$$\widehat{k} = \frac{\sigma_{\pi,u}\sigma_{x,u}\,\sigma_u^2 + \sigma_{\pi,r}\sigma_{x,r}\,\sigma_r^2}{\sigma_{x,u}^2\,\sigma_u^2 + \sigma_{x,r}^2\,\sigma_r^2} = \frac{S_u\,w_u + S_r\,w_r}{w_u + w_r},$$

where

$$S_u = \frac{\sigma_{\pi,u}}{\sigma_{x,u}} = \frac{\lambda\kappa}{\lambda - \beta\lambda^2 - \gamma}, \qquad S_r = \frac{\sigma_{\pi,r}}{\sigma_{x,r}} = \frac{\kappa}{1 - \beta\lambda}, \qquad w_u = \sigma_{x,u}^2\sigma_u^2, \quad w_r = \sigma_{x,r}^2\sigma_r^2.$$

When there is only one source of shocks, the expression for $\widehat{k}$ simplifies to either S_u or S_r, but otherwise it is a variance–weighted average of the two shock-specific slopes.

Each of the following conditions is sufficient to guarantee

$$\frac{d\widehat{k}}{d\phi_\pi} \;<\; 0.$$

(S1) (Natural-rate-only shocks)

$\sigma_u^2 = 0$ (cost-push shock variance zero). Then

$$\widehat{k} = \frac{\kappa}{1-\beta\lambda(\phi_\pi)}, \qquad \frac{d\widehat{k}}{d\phi_\pi} = \frac{\beta\kappa}{(1-\beta\lambda)^2}\frac{d\lambda}{d\phi_\pi} \;<\; 0,$$

because, from above, in the determinacy region $\dfrac{d\lambda}{d\phi_\pi} = -\dfrac{\kappa\sigma\,\lambda}{\partial F/\partial\lambda} < 0.$

(S2) (Mild backwardness + sufficiently aggressive rule)

Assume

$$\gamma \leq \;\beta \;\text{ and }\; \phi_\pi \geq \bar{\phi}, \text{ where } F\Big(\tfrac{1}{2\beta}, \bar{\phi}\Big) = 0\ .$$

Equivalently, the stable root satisfies $\lambda(\phi_\pi) \leq \frac{1}{2\beta}$. Under $\gamma \leq \beta$, both shock-specific slopes $S_u = \frac{\lambda\kappa}{\lambda-\beta\lambda^2-\gamma}$ and $S_r = \frac{\kappa}{1-\beta\lambda}$ are increasing in $\lambda \in [0,1]$. Moreover, when $\lambda \leq \frac{1}{2\beta}$ we have $(1-2\beta\lambda)/\gamma \geq 0$, which implies the weight ratio $w_r/w_u = (\sigma_{x,r}^2\sigma_r^2)/(\sigma_{x,u}^2\sigma_u^2)$ is nondecreasing in λ. Therefore the mixture $\widehat{k} = (S_u w_u + S_r w_r)/(w_u + w_r)$ is increasing in λ, and since $d\lambda/d\phi_\pi < 0$ in the determinacy region,

$$\frac{d\widehat{k}}{d\phi_\pi} \;=\; \frac{d\widehat{k}}{d\lambda}\cdot\frac{d\lambda}{d\phi_\pi} \;<\; 0.$$

(S3) (Cost-push only + mild backwardness).

If $\sigma_r^2 = 0$ and $\gamma \leq \beta$, then $\widehat{k} = S_u$ and

$$\frac{\partial S_u}{\partial\lambda} = \frac{\kappa(\beta\lambda^2-\gamma)}{(\lambda-\beta\lambda^2-\gamma)^2} \;\geq\; 0 \quad (\text{since } \gamma \leq \beta \text{ and } \lambda^2 \leq 1).$$

So $\widehat{k}$ increases in λ and hence decreases in ϕ_π.

(S1) and (S3) are cases where there is only one source of shocks. (S2) covers the empirically relevant hybrid case with both shocks. It only requires a mild bound on the backward-looking weight ($\gamma \leq \beta$) and a lower bound on policy aggressiveness, namely $\phi_\pi \geq \bar{\phi}$ defined

by the cubic at $\lambda = 1/(2\beta)$. The policy threshold is computable from primitives by solving $F\left(\frac{1}{2\beta}, \bar{\phi}\right) = 0$. Thus for sufficiently small γ and sufficiently large ϕ_π, the result is guaranteed to hold.

Stronger feedback on inflation reduces the co-movement of π_t with x_t conditional on π_{t-1}. Intuitively, policy offsets gap-driven inflation pressure more promptly, so regression methods that treat x_t as an exogenous driver recover a smaller slope. This is the sense in which good policy can make the Phillips curve look flat.

A.4 Markov-Switching Version

We provide further analysis of the subsection 3.2 version of our New Keynesian model. To compute a rational expectations equilibrium, we guess the Markov state-dependent linear form:

$$y_t = G_{s_t} y_{t-1} + H_{s_t} v_t, \tag{16}$$

where (G_1, G_2) are 2×2 matrices governing the endogenous state dynamics in each regime, and (H_1, H_2) are 2×2 matrices describing contemporaneous impacts of shocks. Under guess (16), the conditional expectation of y_{t+1} is:

$$\mathbb{E}_t y_{t+1} = P_{s_t,1} G_1 y_t + P_{s_t,2} G_2 y_t = \left(P_{s_t,1} G_1 + P_{s_t,2} G_2\right) y_t.$$

Substitute equation into system (9) and use (16):

$$\Gamma_{-1} \left(P_{s_t,1} G_1 + P_{s_t,2} G_2\right) y_t + \Gamma_{0,s_t} y_t + \Gamma_1 y_{t-1} + v_t = 0.$$

Evidently, for state $s_t = 1$:

$$\left[\Gamma_{-1} \left(P_{1,1} G_1 + P_{1,2} G_2\right) + \Gamma_{0,1}\right] y_t + \Gamma_1 y_{t-1} + v_t = 0.$$

For state $s_t = 2$:

$$\left[\Gamma_{-1} \left(P_{2,1} G_1 + P_{2,2} G_2\right) + \Gamma_{0,2}\right] y_t + \Gamma_1 y_{t-1} + v_t = 0.$$

Now define state-dependent matrices:

$$\Theta_1 \equiv \Gamma_{-1} \left(P_{1,1} G_1 + P_{1,2} G_2\right) + \Gamma_{0,1}, \tag{17}$$

$$\Theta_2 \equiv \Gamma_{-1} \left(P_{2,1} G_1 + P_{2,2} G_2\right) + \Gamma_{0,2}. \tag{18}$$

Then our system becomes:

$$\Theta_1 y_t + \Gamma_1 y_{t-1} + v_t = 0,$$
$$\Theta_2 y_t + \Gamma_1 y_{t-1} + v_t = 0.$$

Now solve for y_t:

$$y_t = -\Theta_1^{-1}\Gamma_1 y_{t-1} - \Theta_1^{-1} v_t,$$
$$y_t = -\Theta_2^{-1}\Gamma_1 y_{t-1} - \Theta_2^{-1} v_t.$$

Comparing these equations with (16), we obtain the fixed point conditions:

$$G_1 = -\Theta_1^{-1}\Gamma_1, \tag{19}$$
$$G_2 = -\Theta_2^{-1}\Gamma_1, \tag{20}$$
$$H_1 = -\Theta_1^{-1}, \tag{21}$$
$$H_2 = -\Theta_2^{-1}. \tag{22}$$

Iterative Algorithm

The fixed-point equations (19)–(22) are implicit because the Θ matrices depend on the unknown G matrices, as shown in (17)-(18). We can solve them via the following iterative procedure:

Step 1: Initial Guess (Anticipated Utility Solution)

Compute "anticipated utility" solutions that ignore regime switching. For each state $i \in \{1, 2\}$, solve the single-regime model:

$$\Gamma_{-1}\mathbb{E}_t y_{t+1} + \Gamma_{0,i} y_t + \Gamma_1 y_{t-1} + v_t = 0,$$

assuming permanent residence in state i. The solution $y_t = G_i^{AU} y_{t-1} + H_i^{AU} v_t$ satisfies:

$$\left(\Gamma_{-1} G_i^{AU} + \Gamma_{0,i}\right) G_i^{AU} y_{t-1} + \Gamma_1 y_{t-1} = 0,$$

which can be solved iteratively for G_i^{AU} using:

$$G_i^{AU,(n+1)} = -\left(\Gamma_{-1} G_i^{AU,(n)} + \Gamma_{0,i}\right)^{-1} \Gamma_1.$$

These provide initial guesses: $G_1^{(0)} = G_1^{AU}$, $G_2^{(0)} = G_2^{AU}$, $H_1^{(0)} = -\left(\Gamma_{-1}G_1^{AU} + \Gamma_{0,1}\right)^{-1}$, and $H_2^{(0)} = -\left(\Gamma_{-1}G_2^{AU} + \Gamma_{0,2}\right)^{-1}$.

Step 2: Fixed Point Iteration with Relaxation

Given current iterates $G_1^{(n)}, G_2^{(n)}, H_1^{(n)}, H_2^{(n)}$ do the following:

1. Update the Θ matrices:

$$\begin{aligned}\Theta_1^{(n)} &= \Gamma_{-1}\left(P_{1,1}G_1^{(n)} + P_{1,2}G_2^{(n)}\right) + \Gamma_{0,1},\\ \Theta_2^{(n)} &= \Gamma_{-1}\left(P_{2,1}G_1^{(n)} + P_{2,2}G_2^{(n)}\right) + \Gamma_{0,2}.\end{aligned}$$

2. Compute new policy matrices:

$$\begin{aligned}\tilde{G}_1^{(n+1)} &= -\left(\Theta_1^{(n)}\right)^{-1}\Gamma_1,\\ \tilde{G}_2^{(n+1)} &= -\left(\Theta_2^{(n)}\right)^{-1}\Gamma_1,\\ \tilde{H}_1^{(n+1)} &= -\left(\Theta_1^{(n)}\right)^{-1},\\ \tilde{H}_2^{(n+1)} &= -\left(\Theta_2^{(n)}\right)^{-1}.\end{aligned}$$

3. Apply a relaxation parameter $\nu \in (0,1)$:

$$\begin{aligned}G_i^{(n+1)} &= (1-\nu)B_i^{(n)} + \nu\tilde{G}_i^{(n+1)},\\ H_i^{(n+1)} &= (1-\nu)C_i^{(n)} + \nu\tilde{H}_i^{(n+1)},\end{aligned}$$

   for $i \in \{1,2\}$.

4. Check convergence:

$$\max_{i,j,k}\left|G_{i,j,k}^{(n+1)} - G_{i,j,k}^{(n)}\right| < \epsilon,$$

   where ϵ is a small tolerance (e.g., 10^{-8}). If not converged, return to step 1.

The relaxation parameter $\nu \in (0,1)$ facilitates convergence. Mean square stability of the solution can be verified using the following criterion. Define the block matrix:

$$\mathcal{G} = \begin{bmatrix} P_{1,1}G_1 & P_{1,2}G_1 \\ P_{2,1}G_2 & P_{2,2}G_2 \end{bmatrix}.$$

Our solution is mean square stable if and only if all eigenvalues of $\mathcal{G}$ lie strictly inside the unit circle.

B The Phelps Problem

This appendix provides additional analysis a slightly simpler of our section 4 quantitative model. The simplifications permit analytic results, which align with the numerical results we saw above.

B.1 Setup

In particular, we suppose that policymakers believe that the economy satisfies:

$$\pi_{t+1} = a_1 \pi_t + k x_{t+1} + \varepsilon_{\pi,t+1}, \tag{23}$$

$$x_{t+1} = b_1 x_t + g(i_t - \pi_t) + \varepsilon_{x,t+1}, \tag{24}$$

where $|b_1| < 1$, $k > 0$, and $g < 0$, and the shocks $\varepsilon_{\pi,t}$ and $\varepsilon_{x,t}$ are i.i.d. Gaussian random variables. These are versions of (11)-(12) in our quantitative application. There we consider a slightly generalized model with more dynamics, nonzero target terms and intercepts, and an interest-rate smoothing penalty in the loss function. Here we let the control be the real interest rate

$$r_t \equiv i_t - \pi_t.$$

Substituting (24) into (23) gives the linear state-space form

$$y_{t+1} = A y_t + B r_t + w_{t+1}$$

where $y_t = [\pi_t, x_t]'$, and the state matrices and composite shock are

$$A = \begin{bmatrix} a_1 & k b_1 \\ 0 & b_1 \end{bmatrix}, \quad B = \begin{bmatrix} kg \\ g \end{bmatrix}, \quad w_{t+1} = \begin{bmatrix} \varepsilon_{\pi,t+1} + k\,\varepsilon_{x,t+1} \\ \varepsilon_{x,t+1} \end{bmatrix}.$$

We assume here that the policymaker minimizes the discounted quadratic loss:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \lambda_x\, x_t^2\right), \qquad 0 < \beta < 1, \quad \lambda_x > 0.$$

This is a standard discounted linear quadratic regulator with no control cost. Defining $Q = \mathrm{diag}(1, \lambda_x)$, the value function is $V(y) = y'Py + p_0$, where p_0 is a constant (dependent

on the shock variances) and $P \succeq 0$ solves the discounted algebraic Riccati equation

$$P \;=\; Q \;+\; \beta\Big(A'PA - A'PB\,(B'PB)^{-1}B'PA\Big).$$

The optimal real rate rule is $r_t = -Fy_t$ with

$$F \;=\; (B'PB)^{-1}B'PA \;=\; \begin{bmatrix} f_\pi & f_x \end{bmatrix}.$$

Since $i_t = r_t + \pi_t$, the implementable interest-rate rule is

$$i_t \;=\; \big(1 - f_\pi\big)\pi_t - f_x x_t \equiv \phi_\pi \pi_t + \phi_x x_t.$$

Solving the Riccati system yields two simple identities:

$$P_{12} = 0, \qquad P_{22} = \lambda_x,$$

and $P_{11} > 0$ is the stabilizing root of the quadratic

$$k^2 P_{11}^2 \;+\; \Big(\lambda_x\,(1 - \beta a_1^2) - k^2\Big)P_{11} \;-\; \lambda_x \;=\; 0,$$

namely

$$P_{11} \;=\; \frac{-\Big(\lambda_x(1-\beta a_1^2) - k^2\Big) + \sqrt{\Big(\lambda_x(1-\beta a_1^2) - k^2\Big)^2 + 4\lambda_x k^2}}{2k^2}. \tag{25}$$

Using P in $F = (B'PB)^{-1}B'PA$ and simplifying gives

$$f_\pi \;=\; \frac{k\,P_{11}\,a_1}{g\,(k^2 P_{11} + \lambda_x)}, \qquad f_x \;=\; \frac{b_1}{g},$$

so the optimal policy coefficients are

$$\phi_\pi(a_1, k) \;=\; 1 \;-\; \frac{a_1}{g\,k} \cdot \frac{Z}{Z + \lambda_x}, \qquad \phi_x \;=\; -\,\frac{b_1}{g}, \tag{26}$$

where we define $Z \equiv k^2 P_{11}$ and, from (25),

$$Z \;=\; \frac{-\,\Delta \;+\; \sqrt{\Delta^2 + 4\lambda_x k^2}}{2}, \qquad \Delta \;=\; \lambda_x(1 - \beta a_1^2) - k^2. \tag{27}$$

Note ϕ_x depends only on (b_1, g) and is independent of (a_1, k).

B.2 Comparative Statics

We now establish how ϕ_π moves with the perceived persistence a_1 and slope k.

Proposition 4 (Lower measured persistence $\Rightarrow$ less aggressive policy)**.** *For any $\beta \in (0,1)$, $\lambda_x > 0$, $k > 0$, $g < 0$, the optimal inflation coefficient $\phi_\pi(a_1, k)$ in (26) is strictly increasing in a_1. Equivalently, a reduction in the measured persistence a_1 implies a less aggressive policy (lower ϕ_π).*

Proof. From (27), $\phi_\pi = 1 - \frac{a_1}{gk}\frac{Z}{Z+\lambda_x}$. Since $\frac{\partial \Delta}{\partial a_1} = -2\lambda_x \beta a_1 < 0$ (for $a_1 \geq 0$) and $\partial Z/\partial \Delta \in (-1, 0)$, we have $\partial Z/\partial a_1 > 0$. Because $Z/(Z+\lambda_x)$ increases in Z and $gk < 0$, it follows that $\partial \phi_\pi/\partial a_1 > 0$. □

Proposition 5 (Flatter perceived slope $\Rightarrow$ less aggressive policy (sufficient condition))**.** *Let $\phi_\pi(a_1, k)$ be given by (26). Then $\partial \phi_\pi/\partial k > 0$ whenever*

$$k^2 \lambda \left(1 + \frac{-\Delta + 2\lambda_x}{\sqrt{\Delta^2 + 4\lambda_x k^2}}\right) > Z\,(Z+\lambda_x), \qquad \Delta = \lambda_x(1-\beta a_1^2) - k^2, \;\; Z = \frac{-\Delta + \sqrt{\Delta^2 + 4\lambda_x k^2}}{2}. \tag{28}$$

Sketch. Differentiate $\phi_\pi = 1 - (a_1/(gk))\, f(Z)$ with $f(Z) = Z/(Z+\lambda)$. Then

$$\frac{\partial \phi_\pi}{\partial k} = -\frac{a_1}{g}\Big(-\frac{f(Z)}{k^2} + \frac{f'(Z)}{k}\frac{dZ}{dk}\Big), \quad f'(Z) = \frac{\lambda_x}{(Z+\lambda)^2},$$

and $dZ/dk = \frac{k}{2}\Big(2 + \frac{-2\Delta + 4\lambda_x}{\sqrt{\Delta^2 + 4\lambda_x k^2}}\Big)$. Because $-a_1/g > 0$, the sign reduces to the inequality (28). □

The key inequality holds, in particular, for moderate k (relative to λ and a_1), and it holds globally once an arbitrarily small rate-smoothing term $\eta > 0$ is added to the loss.[22]

Equation (26) shows that $\phi_x = -b_1/g$ is pinned down by the persistence of the output gap and the interest sensitivity parameter g. The inflation coefficient ϕ_π depends on persistence (a_1) and the perceived slope (k) only through the weight $Z/(Z+\lambda_x)$, where Z in (28) aggregates both the gap channel (k) and the discounted persistence term (βa_1^2). Proposition 4 is global: when inflation is believed to be less persistent ($a_1 \downarrow$), the policymaker optimally moderates the response to π_t (lower ϕ_π). Proposition 5 formalizes the slope effect: a flat Phillips curve (small k) makes policy less aggressive. The result provides a sufficient condition, but with any small control penalty $\eta > 0$, this holds for all k.

[22] With $\eta > 0$ the problem has a loss function matrix $R = \eta I$, and the derivative $\partial \phi_\pi/\partial k$ is strictly positive for all parameter values (see above). As $\eta \downarrow 0$, the inequality (28) becomes the boundary condition separating the regions with $\partial \phi_\pi/\partial k > 0$, over the empirically relevant case of a small slope, and a possible weak reversal for very large k.

References

Andolfatto, D. and F. M. Martin (2025). Monetary Policy and the COVID-19 Price Level Shock. Working Paper 2025-004, Federal Reserve Bank of St. Louis.

Ball, L. and S. Mazumder (2011). Inflation Dynamics and the Great Recession. *Brookings Papers on Economic Activity Spring*, 337–405.

Bassetto, M. and T. J. Sargent (2020). Shotgun Wedding: Fiscal and Monetary policy. *Annual Review of Economics 12*(1), 659–690.

Benati, L. and P. Surico (2008). Evolving U.S. Monetary Policy and the Decline of Inflation Persistence. *Journal of the European Economic Association 6*(2–3), 634–646.

Bernanke, B. S. (2022). *21st Century Monetary Policy: The Federal Reserve from the Great Inflation to COVID-19*. New York: WW Norton & Company.

Bernheim, B. D. (1984). Rationalizable Strategic Behavior. *Econometrica 52*, 1007–1028.

Bianchi, F., R. Faccini, and L. Melosi (2023). A Fiscal Theory of Persistent Inflation. *Quarterly Journal of Economics 138*(4), 2127–2179.

Blinder, A. S. (2021, December). When It Comes to Inflation, I'm Still on Team Transitory. *Wall Street Journal*.

Bullard, J. B. (2018, June). The Case of the Disappearing Phillips Curve: A Presentation at the 2018 ECB Forum on Central Banking Macroeconomics of Price- and Wage-Setting, Sintra, Portugal. Speech 314.

Carlstrom, C. T., T. S. Fuerst, and M. Paustian (2009). Inflation Persistence, Monetary Policy, and the Great Moderation. *Journal of Money, Credit and Banking 41*(4), 767–786.

Cho, I.-K. and K. Kasa (2015). Learning and Model Validation. *The Review of Economic Studies 82*(1), 45–82.

Cho, I.-K. and K. Kasa (2017). Gresham's Law of Model Averaging. *American Economic Review 107*(11), 3589–3616.

Cho, I.-K., N. Williams, and T. J. Sargent (2002). Escaping Nash Inflation. *Review of Economic Studies 69*(1), 1–40.

Clarida, R. H. (2019, May). Sustaining Maximum Employment and Price Stability. Speech at the Economic Club of New York, New York, NY, May 30, 2019. Vice Chair, Board of Governors of the Federal Reserve System.

Cochrane, J. H. (2025). Monetary-Fiscal Interactions. Working paper, Hoover Institution. Prepared for the 50-Year Retrospective On The Shadow Open Market Committee And Its Role In Monetary Policy, Hoover Institution, October 13, 2024.

Cogley, T., G. E. Primiceri, and T. J. Sargent (2010, January). Inflation-Gap Persistence in the US. *American Economic Journal: Macroeconomics 2*(1), 43–69.

Cogley, T. and T. J. Sargent (2002). Evolving Post-World War II U.S. Inflation Dynamics. In B. S. Bernanke and K. Rogoff (Eds.), *NBER Macroeconomics Annual 2001, Volume 16*, pp. 331–373. Cambridge, MA: MIT Press.

Cogley, T. and T. J. Sargent (2005a). Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII US. *Review of Economic Dynamics 8*(2), 262–302.

Cogley, T. and T. J. Sargent (2005b). The Conquest of US Inflation: Learning and Robustness to Model Uncertainty. *Review of Economic Dynamics 8*(2), 528–563.

Cogley, T. and T. J. Sargent (2008). Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making. *International Economic Review 49*(1), 185–221.

Coibion, O. and Y. Gorodnichenko (2015). Is the Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation. *Brookings Papers on Economic Activity 2015*(2), 197–255.

Davig, T. (2016). Phillips Curve Instability and Optimal Monetary Policy. *Journal of Money, Credit and Banking 48*(1), 239–274. First published online: 19 January 2016.

Davig, T. and T. Doh (2014). Monetary Policy Regime Shifts and Inflation Persistence. *The Review of Economics and Statistics 96*(5), 862–875.

Edge, R. M., T. Laubach, and J. C. Williams (2007). Learning and Shifts in Long-Run Productivity Growth. *Journal of Monetary Economics 54*(8), 2421–2438.

Edge, R. M. and J. B. Rudd (2016). Real-Time Properties of the Federal Reserve's Output Gap. *Review of Economics and Statistics 98*, 785–791.

Evans, G. W., S. Honkapohja, T. J. Sargent, and N. Williams (2013). Bayesian Model Averaging, Learning and Model Selection. In T. Sargent and J. Vilmunen (Eds.), *Macroeconomics at the Service of Public Policy*, pp. 99–119. Oxford University Press Oxford.

Farmer, L. E., E. Nakamura, and J. Steinsson (2024). Learning About the Long Run. *Journal of Political Economy 132*(10), 3334–3377.

Furlanetto, F. and A. Lepetit (2024, June). The Slope of the Phillips Curve. FEDS Working Paper 2024-43, Board of Governors of the Federal Reserve System.

García, C. E., D. M. Prett, and M. Morari (1989). Model Predictive Control: Theory and Practice – A Survey. *Automatica 25*(3), 335–348.

Giannone, D. and G. Primiceri (2025). Demand Driven Inflation. Working paper.

Hall, G. J. and T. J. Sargent (2022). Three World Wars: Fiscal–Monetary Consequences. *Proceedings of the National Academy of Sciences 119*(18), e2200349119.

Hall, G. J. and T. J. Sargent (2025). Fiscal Consequences of the US War on COVID. *International Economic Review*. Forthcoming.

Hobijn, B., R. A. Miles, J. Royal, and J. Zhang (2023). The Recent Steepening of Phillips Curves. Technical Report 475, Federal Reserve Bank of Chicago.

Jørgensen, P. L. and K. J. Lansing (2019). Anchored Inflation Expectations and the Slope of the Phillips Curve. Technical Report 2019-27, Federal Reserve Bank of San Francisco.

Koopmans, T. C. (1947). Measurement Without Theory. *The Review of Economics and Statistics 29*(3), 161–172.

Kreps, D. M. (1998). Anticipated Utility and Dynamic Choice. In D. P. Jacobs, E. Kalai, and M. I. Kamien (Eds.), *Frontiers of Research in Economic Theory*, pp. 242–274. Cambridge: Cambridge University Press.

Krugman, P. (2022, July). I Was Wrong About Inflation. *The New York Times*.

Kurozumi, T. and W. V. Zandweghe (2018). Why Has Inflation Persistence Declined? Economic Bulletin.

Kuttner, K. N. and T. Robinson (2010). Understanding the Flattening of the Phillips Curve. *North American Journal of Economics and Finance 21*(2), 110–125.

Lagos, R. (2025, August). onetary Economics at 30: A Reexamination of the Relevance of Money in Cashless Limiting Monetary Economies. Working Paper 34155, National Bureau of Economic Research, Cambridge, MA.

Mayne, D. Q., J. B. Rawlings, C. V. Rao, and P. O. Scokaert (2000). Constrained Model Predictive Control: Stability and Optimality. *Automatica 36*(6), 789–814.

McCallum, B. T. (1983). On Non-Uniqueness in Rational Expectations Models: An Attempt at Perspective. *Journal of Monetary Economics 11*(2), 139–168.

McCallum, B. T. (2004). On the Relationship Between Determinate and MSV Solutions in Linear RE Models. *Economics Letters 84*(1), 55–60.

McLeay, M. and S. Tenreyro (2019). Optimal Inflation and the Identification of the Phillips Curve. In M. Eichenbaum, E. Hurst, and J. A. Parker (Eds.), *NBER Macroeconomics Annual 2018, Volume 33*, pp. 199–255. Chicago: University of Chicago Press.

Minsky, H. P. (1986). *Stabilizing an Unstable Economy*. New Haven: Yale University Press.

Nakamura, E., V. Riblier, and J. Steinsson (2025, August). Beyond the Taylor Rule. Working paper.

Occhino, F. (2019). The Flattening of the Phillips Curve: Policy Implications Depend on the Cause. Federal Reserve Bank of Cleveland Economic Commentary.

Orphanides, A. (2001). Monetary Policy Rules Based on Real-Time Data. *American Economic Review 91*(4), 964–985.

Orphanides, A. and S. van Norden (2002). The Unreliability of Output-Gap Estimates in Real Time. *Review of Economics and Statistics 84*(4), 569–583.

Pearce, D. G. (1984). Rationalizable Strategic Behavior and the Problem of Perfection. *Econometrica 52*, 1029–1050.

Peneva, E., J. Rudd, and D. Villar (2025). Retrospective on the Federal Reserve Board Staff's Inflation Forecast Errors since 2019. Staff Report 2025-069, Board of Governors of the Federal Reserve System.

Phelps, E. S. (1967). Phillips Curves, Expectations of Inflation and Optimal Unemployment over Time. *Economica 35*, 254–281.

Pivetta, F. and R. Reis (2007). The Persistence of Inflation in the United States. *Journal of Economic Dynamics and Control 31*(4), 1326–1358.

Primiceri, G. E. (2006). Why Inflation Rose and Fell: Policymakers' Beliefs and U.S. Postwar Stabilization Policy. *Quarterly Journal of Economics 121*(3), 867–901.

Rawlings, J. B., D. Q. Mayne, and M. Diehl (2017). *Model Predictive Control: Theory, Computation, and Design* (2nd ed.). Madison, WI: Nob Hill Publishing.

Sargent, T. J. (1999). *The Conquest of American Inflation*. Princeton, NJ: Princeton University Press.

Sargent, T. J. and N. Williams (2005). Impacts of Priors on Convergence and Escapes from Nash Inflation. *Review of Economic Dynamics 8*(2), 360–391.

Sargent, T. J., N. Williams, and T. Zha (2006). Shocks and Government Beliefs: The Rise and Fall of American Inflation. *American Economic Review 96*(4), 1193–1224.

Silber, W. L. (2012). *Volcker: The Triumph of Persistence*. Bloomsbury Publishing USA.

Simon, J., T. Matheson, and D. Sandri (2013). The Dog That Didn't Bark: Has Inflation Been Muzzled or Was It Just Sleeping? In *World Economic Outlook, April 2013: Hopes, Realities, Risks*. International Monetary Fund.

Staff, W. E. (2024, December). The Fed Official who Correctly Predicted Inflation Could Fall Without a Recession. *Washington Examiner*.

Stock, J. H. and M. W. Watson (2007). Why Has U.S. Inflation Become Harder to Forecast? *Journal of Money, Credit and Banking 39*(S1), 3–33.

Summers, L. H. (2024, January). Soft Landing Is 'Very Real Possibility' Now. Remarks at Economic Club of New York event, as reported by Bloomberg.

Svensson, L. E. and N. Williams (2008). Optimal Monetary Policy Under Uncertainty: A Markov Jump-Linear-Quadratic Approach. *Federal Reserve Bank of St. Louis Review 90*(4), 275–293.

Svensson, L. E. and N. Williams (2009, March). Optimal Monetary Policy under Uncertainty in DSGE Models: A Markov Jump-Linear-Quadratic Approach. In K. Schmidt-Hebbel and C. E. Walsh (Eds.), *Monetary Policy under Uncertainty and Learning*, Volume 13 of *Central Banking, Analysis, and Economic Policies Book Series*, Chapter 3, pp. 77–114. Central Bank of Chile.

Taylor, J. B. (2007). Housing and Monetary Policy. In *Housing, Housing Finance, and Monetary Policy*, pp. 463–476. Kansas City, MO: Federal Reserve Bank of Kansas City. Proceedings of the Economic Policy Symposium, Jackson Hole, Wyoming.

Taylor, J. B. (2009a). *Getting Off Track: How Government Actions and Interventions Caused, Prolonged, and Worsened the Financial Crisis*. Stanford, CA: Hoover Institution Press.

Taylor, J. B. (2009b, January). The Financial Crisis and the Policy Responses: An Empirical Analysis of What Went Wrong. Working Paper 14631, National Bureau of Economic Research, Cambridge, MA.

Williams, J. C. (2006, October). Inflation Persistence in an Era of Well-Anchored Inflation Expectations. Technical Report 2006-27, Federal Reserve Bank of San Francisco.

Yellen, J. L. (2019). Former Fed Chair Janet Yellen on why the answer to the inflation puzzle matters. hhttps://www.brookings.edu/articles/former-fed-chair-janet-yellen-on-why-the-answer-to-the-inflation-puzzle-matters. Speech at the Brookings Instiution, October 3, 2019.