Evolving Reputation for Commitment: The Rise, Fall and Stabilization of US Inflation*

Robert G. King† Yang K. Lu‡

First draft: Feb 2019. This draft: November 2023

Abstract

We develop a computable recursive equilibrium for a dynamic game involving two types of purposeful policymakers – one with commitment capacity and the other without – and private agents who form expectations about future policies. Private agents are uncertain about policymaker type and their learning yields a time-varying reputation state. When applied to a New Keynesian setup with forward-looking inflation dynamics and a standard policy objective, our theory highlights the interplay between the reputation state and the differences in optimal policies of the two policymaker types. We provide a quantitative implementation of our theory via a nonlinear filter to show that active management of evolving reputation by committed policy is central to US inflation history.

Keywords: time inconsistency, reputation game, optimal monetary policy, forward-looking expectations

JEL classifications: E52, D82, D83.

*We would like to thank the audience at various conferences and seminars for helpful comments and discussions. All errors are ours. We acknowledge the financial support from the RGC of HKSAR (GRF HKUST-16504317).

†Corresponding author, Boston University and NBER. Email: ⟨rking@bu.edu⟩

‡Hong Kong University of Science and Technology. Email: ⟨yanglu@ust.hk⟩

1 Introduction

The inflation of the 1970s brought about a fundamental revolution in the theory of economic policy. Influential studies by Lucas and Sargent, as summarized in their manifesto (1997), showed that traditional econometric models were inappropriate for analysis of exogenous policy rules when rational expectations is coupled with forward-looking private sector behavior. Kydland and Prescott (1977) took the next step by incorporating purposeful policymakers into theoretical macroeconomic environments, formulated as dynamic games. They stressed the importance of policymaker commitment capacity,showing how its absence could radically change positive and normative outcomes.

In the extensive elaboration of these insights over the ensuing decades, there has been growing recognition that private agent learning is important and, indeed, that policymaker commitment capacity is inherently unobservable. A substantial body of literature now integrates private agent learning into the theory of economic policy.[1] Yet, an important gap remains as few models feature purposeful policymakers who actively seek to steer the learning of private agents.

This paper shows how to close this gap. We use the insights of modern contract theory (mechanism design) to develop a computable recursive equilibrium for a dynamic game with two types of purposeful policymakers, one which can commit and one which cannot, and private agents who learn policymaker type in a Bayesian manner. The forward-looking behavior of private agents, coupled with both types of policymakers being purposeful, necessitates our novel theoretical approach. In our recursive equilibrium, reputation – defined as private agents' likelihood that the policymaker can commit – emerges as a key endogenous state variable.

Our theoretical framework makes it possible to model rich strategic interactions between private agents and policymakers of differing commitment capacity that appear important in many contexts, including fiscal and monetary policy, sovereign borrowing and default, capital controls and exchange rate regimes, and regulation of banking and financial markets.

Our application Harking back to the subject that stimulated the revolution in the theory of economic policy, we show that our framework can be used to enhance understanding of the interplay of inflation expectations and inflation policy in the United States. To this end, we employ a variant of the textbook New Keynesian (NK) model with forward-looking inflation dynamics, purposeful policymakers with a dual mandate to stabilize inflation and output, and

[1]See for examples: Barro (1986), Backus and Driffill (1985), Phelan (2006), Dovis and Kirpalani (2022).

stochastic changes in regime.[2] A committed policymaker always follows an ex-ante optimal state-contingent plan for his intended inflation policy. An opportunistic policymaker chooses his intended inflation policy in a sequentially optimal way. Private agents do not observe policymaker type or intended inflation but only noisy inflation realizations, which they use to update their belief about policymaker type and to form expectations of future inflation. A central conceptual result is that high reputation narrows the equilibrium policy difference between the two policymaker types as the policymaker lacking commitment capacity is less tempted to deviate, whereas low reputation widens the equilibrium policy difference due to the incentive of the committed policymaker to prompt faster private agents learning. A central quantitative result is that evolving reputation is crucial for matching key aspects of U.S. inflation history.

Why new theory is necessary and the dynamic system it delivers Forward-looking NK inflation dynamics have largely replaced the inflation specifications employed by Lucas, Sargent, Kydland and Prescott in which private agents expectations are intra-temporal, i.e., the expected policy is chosen in the same period of expectation formation. In response to supply shocks, forward-looking inflation dynamics heighten the difference between optimal inflation policy with commitment and without.[3] Some prior literature has examined the interplay of optimal inflation policy and reputation with intra-temporal expectations, Cukierman and Liviatan (1991), King et al. (2008), Lu (2013), Dovis and Kirpalani (2021)). These studies exploit the fact that intra-temporal expectations make it possible to solve dynamic games using backward induction.

When expectations are forward-looking, strategic interactions become intertemporal and the earlier techniques no longer apply. To see why, consider the choice of period-t committed policy: the period-t payoff depends on private agent expectations, which are affected by future committed policy, future opportunistic policy, and reputation – how likely each policy will take place as perceived by private agents. But the future opportunistic policy cannot be taken as given because it optimally responds to future private agents expectations that change with how period-t committed and opportunistic policies affect the evolution of reputation.[4]

[2] A regime is time interval during which outcomes can be understood as choices of a single policymaker.

[3] See, for example, Clarida et al. (1999)

[4] One common way to avoid these strategic interactions is to assume that one type of policymaker being an automaton (Lu et al. (2016), Amador and Phelan (2021), Morelli and Moretti (2023)), or to assume that the committed policymaker ignores the effect of his policy on private sector learning (Clayton et al. (2022)). However, our analysis below indicates that these assumptions have considerable effects on outcomes, which

Our new mechanism design approach directly tackles these complications. To begin, we recast the equilibrium of the dynamic game as the solution to a dynamic principal-agent problem. The committed policymaker acts as principal to choose state contingent plans for his own policies, the policies of the opportunistic type subject to incentive compatibility constraints, and private agents expectations subject to rational expectation constraints. We then use the techniques of dynamic contract theory to formulate the principal-agent problem as a recursive optimization with only three state variables including a highly persistent reputation state,[5] a more temporary cost-push shock,[6] and a predetermined pseudo state.[7]

Our dynamic theory makes quantificative history feasible Based on the solution to the recursive optimization, we construct a calibrated quantitative theoretical model that maps structural shocks and latent states to observable macro data. Specifically, we require that private agent inflation expectations in the model match time series from the Survey of Professional Forecasters (SPF) starting in late 1968. Intuitively, the identification assumption is that short-term SPF forecasts should be more sensitive to temporary factors like cost-push shocks and longer-term forecasts should better capture persistent factors like reputation. Formally, we exploit the fact that our theoretical model's dynamic system suggests a nonlinear filter with hidden Markov-switching to jointly identify three structural shocks, the three state variables, and regime change events.[8]

We find novel empirical results that are exciting: estimated reputation emerges as powerful dynamic factor. It exhibits a big swing, declining throughout 1970s to near zero by the end of 1980 and gradually climbing back afterwards. Estimated probability of regime change spikes around 1981-2, with a committed regime unlikely before 1981 and increasingly likely afterward. Our nonlinear filter considers inflation as a latent state variable, resulting in estimated inflation values. Remarkably, these estimates align closely with the observed U.S. inflation, despite the fact that the observed inflation data is not used by the nonlinear

are – to our minds – undesirable.

[5]Reputation is a capital good for the committed policymaker but evolves as a martingale in the eyes of private agents.

[6]We use the common terminology for this shock, which shifts the output-inflation trade-off for the policymaker.

[7]As in other studies of optimal inflation policy, this variable is required to place the committed policy in recursive form, as discussed further below.

[8]We cannot use the standard Kalman filter since our model is not linear. As detailed below, we adopt a particular "sigma point" approximation method – the unscented Kalman filter – that has been shown to work well in nonlinear regime-switching models. Särkkä and Svensson (2023) describes the general Gaussian filtering. Recent macroeconomic applications of the unscented Kalman filter are Binning and Maih (2015), Benigno et al. (2020), and Foerster and Matthes (2022).

filter. Using our estimated shocks and states, we further compute the model-implied optimal inflation policies for both policymaker types and find that the U.S. inflation is tracked by the opportunistic policy before 1981 and by the committed policy after 1981.

To assess the importance of having optimal committed policy purposefully influence private agents learning, we also conduct a counterfactual exercise in which a naive committed policymaker optimizes but ignores the effect of his policy on reputation evolution. Such policymaker naivete results in a narrower policy difference between the committed and opportunistic policymakers, especially when the reputation is low. Using the history of estimated cost-push shocks and probabilities of regime changes from our benchmark quantitative model, we compute counterfactual time series of optimal committed and opportunistic policies by the naive policymakers, and reputation governed by the naive policymakers' past responses to shocks. The results show that a naive committed policymaker takes much longer to disinflate the economy than what is observed in post-1981 U.S. inflation history.

Links to the broader literature Our reputational equilibrium analysis adopts one of the two approaches in modern game theory, originated from Milgrom and Roberts (1982) and Kreps and Wilson (1982).[9] Based on Bayesian learning in a noisy environment, our reputational state variable is the likelihood that the current policymaker has commitment capability. Another familiar reputational approach, introduced by Barro and Gordon (1983) to macroeconomics, demonstrates that reputational forces may substitute for commitment capability, leading a "discretionary" policymaker to behave like a committed one as in the important modern literature on sustainable plans (Chari and Kehoe (1990)).[10] However, policymaker reputation does not vary over time in the sustainable plan literature: it is either excellent or nonexistent. Our learning-based framework permits *reputation building* by a policymaker that can commit and *reputation dissipation* by one that can't.

Our paper is related to a large literature studying the rise, fall and stabilization of US inflation, but our approach is quite different. Sargent (1999) stimulated a literature on the

[9]For a general discussion and specific examples see Mailath and Samuelson (2006). These leading theorists advocate for studying reputation as we do, writing "The idea that a player has an incentive to build, maintain, or milk his reputation is captured by the incentive that player has to manipulate the beliefs of other players about his type. The updating of these beliefs establishes links between past behavior and expectations of future behavior. We say 'reputations effects' arise if these links give rise to restrictions on equilibrium payoffs or behavior that do not arise in the underlying game of complete information."

[10]Within the NK framework, optimal policy under commitment involves time-varying inflation when there are Phillips curve shocks: Kurozumi (2008) and Loisel (2008) have shown that a policymaker without commitment capability can be led to follow such a policy so long as he is sufficiently patient and the shocks are not too large.

role of a purposeful policymaker's beliefs that does not require exogenous regime changes,[11] with Primiceri (2006) extending this approach and quantifying shifts in estimates of the Phillips curve slope and intercept. Bianchi (2013) and Debortoli and Lakdawala (2016) develop and estimate models in which private agents anticipate a possible exogenous policy regime change but do not face a learning problem. Our quantitative theory emphasizes the evolution of *private sector beliefs* and we use the SPF to extract the evolution of such beliefs. In seeking to recover the evolution of private sector beliefs about the commitment capacity of the Fed, our work is related to Matthes (2015), but policymakers in his study don't purposefully manage private sector learning.[12] Our model features interaction of private sector learning and optimal policies with and without commitment, which we see as essential to matching the pattern of actual inflation and its comovement with the SPF. Carvalho et al. (2022) and Hazell et al. (2022) attribute the Volcker disinflation and the inflation stabilization afterwards to a decline of long-term inflation expectations, highlighting that such expectations are anchored in the 1990s. Our theory rationalizes such long-term inflation expectations behavior as an equilibrium outcome.

Use of the SPF also links our research to the large and growing literature on survey measures of inflation (Coibion et al. (2018)). The SPF forecasts systematically underestimated inflation during its rise in the 1970s and then systematically overestimated it during its decline. Our explanation of persistent forecasting errors is consistent with the view that these SPF anomalies arise from agents not knowing the policy regime (Evans and Wachtel (1993), Coibion et al. (2018)) or the model generating the data (Farmer et al. (2021)). Our work differs from the existing literature by having unknown policy optimally evolving over time, rather than being generated by a random process or by exogenous policy rules.

Organization The balance of the paper is as follows. In section 2, we describe the economy. In section 3, we cast the macroeconomic equilibrium in game theoretic terms, defining a Bayesian perfect equilibrium. In section 4, we develop a recursive equilibrium and describe how to solve it. In section 5, we elaborate our new method of latent state extraction from the SPF and use it to construct quantitative measures of policies. Section 6 provides our model-based interpretation of U.S. inflation history and undertakes various exercises to shed light on our model's internal mechanisms. Section 7 concludes.

[11] See the Riksbank review article by Sargent and Soderstrom (2000) for an introduction.

[12] Other papers that investigate U.S. inflation history with private agent learning include Ball (1995), Erceg and Levin (2003), Orphanides and Williams (2005), Goodfriend and King (2005), Cogley et al. (2015), and Melosi (2016).

2 The Economy

A policymaker designs and announces a plan for current and future inflation. A private sector composed of atomistic forward-looking agents is uncertain whether the policymaker can commit or not. Their forward-looking decisions reflect the possibility that an announced policy plan may not be executed.

2.1 Private sector

Private agents' behavior is captured by a standard NK Phillips curve

$$\pi_t = \underbrace{\beta E_t \pi_{t+1}}_{e_t} + \kappa x_t + \varsigma_t, \tag{1}$$

where π_t is inflation, x_t is the output gap, and ς_t is a cost-push shock governed by an exogenous Markov chain with the transition probabilities $\varphi(\varsigma_{t+1}; \varsigma_t)$. Private agents' discount factor is β and $E_t\pi_{t+1}$ is their expectation about the next-period inflation, with e_t shorthand for discounted expected inflation.

2.2 Policymaker

The policymaker is responsible for the inflation rate, π, but cannot control it exactly.[13] There are two types of policymaker. A *committed* type ($\tau = 1$) chooses and announces an optimal state-contingent plan for intended inflation at all dates when he first takes office and executes it in all subsequent periods until replaced.[14] The committed inflation plan therefore shapes private sector's expected inflation. An *opportunistic* type ($\tau = 0$) makes the same announcements,[15] but chooses intended inflation on a period-by-period basis.

[13] We use "policymaker" rather than "central banker" to recognize that inflation policy may be the result of various actors. For example, DeLong (1996), Levin and Taylor (2013), and Meltzer (2014) stress various political influences on monetary policy outcomes, while other economists see direct connections of fiscal policy to inflation.

[14] We specify intended inflation rather than intended output for analytical convenience. If policy instead controlled intended real aggregate demand $\underline{x}_{\tau t}$ and $x_{\tau t} = \underline{x}_{\tau t} + \sigma_{x\tau}\varepsilon_t$, the Phillips curve $\pi_t = \kappa x_t + e_t + \varsigma_t$ implies that a choice of $\underline{x}_{\tau t} = \frac{1}{\kappa}[a_t - e_t - \varsigma_t]$ leads to identical intended inflation, although certain text expressions – particularly those for inflation expectations – are more cumbersome. We also abstract from policy instruments as in some other related studies (see, e.g., Faust and Svensson (2001) and Sargent (1999)).

[15] The opportunistic type makes the same announcements as the committed type to avoid revealing his type. This is consistent with a key conclusion made by Lu (2013) in a related fiscal model: the unique signalling equilibrium involves the truth-telling committed type announcing a policy that solves his optimal policy problem and the opportunistic type sending the same message. We therefore abstract from the analysis

At the start of each period, the policymaker may be replaced through a publicly observed event, occurring with probability q and denoted by $(\theta_t = 1)$. If no replacement occurs $(\theta_t = 0)$, the policymaker type remains unchanged. When a replacement does occur $(\theta_t = 1)$, the incoming policymaker inherits the same type from his predecessor $(\phi_t = 1)$ with probability δ_ρ; otherwise, he draws a new type and becomes a committed type with a probability $v_{\rho,t}$.

The private sector does not observe the policymaker's type (τ_t) or his intended inflation, denoted by a_t for the committed type and α_t for the opportunistic type. Yet, it observes an inflation rate π_t that deviates from the policymaker's intention with a random i.i.d. implementation error $v_{\pi,t} \sim N(0, \sigma^2_{v,\pi})$:[16]

$$\pi_t = \tau_t a_t + (1 - \tau_t)\alpha_t + v_{\pi,t}. \tag{2}$$

The policymaker's momentary objective depends on inflation π and output gap x.

$$u(\pi, x) = -\frac{1}{2}[(\pi - \pi^*)^2 + \vartheta_x (x - x^*)^2] \tag{3}$$

There is a long-run inflation target π^* and a strictly positive output target x^*.[17]

The committed type discount factor is β_a; the opportunistic type is myopic.[18]

2.3 Timing of events

Private agents start period t with a probability that the incumbent policymaker is the committed type, which we denote by ρ_t and call *reputation*. The within-period timing is shown in Figure 1. First, the public event of policymaker replacement may or may not occur. If it occurs $(\theta_t = 1)$, the regime clock t is set to zero and the new policymaker's initial

of signalling equilibria.

[16] We interpret random implementation error as a reduced-form representation for all unforeseeable factors that affect the inflation rate beyond the monetary policy, following Cukierman and Meltzer (1986), Faust and Svensson (2001), Atkeson and Kehoe (2006), etc. There is also ample evidence that realized inflation rates miss the intended inflation target, with examples including Roger and Stone (2005) and Mishkin and Schmidt-Hebbel (2007).

[17] The non-zero inflation target is common in central bank objectives. The output component in the objective can be written as $-\frac{\vartheta_x}{2}[x^2 + (x^*)^2] + (\vartheta_x x^*)x$ highlighting that there is a benefit to an additional unit of output. It is this composite coefficient $(\vartheta_x x^*)$ rather than its components that are important below. Our approach can easily handle publicly observable shocks to the targets π^* and x^*. But since these are not essential to our analysis and have been extensively explored elsewhere, we opt for simplicity in specification.

[18] A myopic opportunistic type is the most parsimonious modeling of an optimizing non-committed policymaker. Our framework and recursive method can be extended to a long-lived opportunistic type, but we leave that extension for future research.

reputation ρ_0 is a random draw from the distribution $\Xi(\rho_0|\rho_t)$ with support [0,1].[19] Second, the exogenous cost-push shock ς_t is realized. Third, there is a policy announcement. If there is a new policymaker, he announces a new inflation plan. Otherwise, either type of continuing policymaker simply reiterates that current economic conditions call for an intended inflation a_t. Fourth, private agents form their expectations about the next-period inflation, e_t. Fifth, the policymaker implements intended inflation, a_t or α_t, depending on his type. Sixth, this action leads to a random inflation rate π_t with a density $g(\pi_t|a_t)$ or $g(\pi_t|\alpha_t)$, and an output gap x_t determined by the Phillips curve.[20] New information leads private agents to update their beliefs about policymaker type.

[Figure 1 about here.]

3 Macro Equilibrium in a Dynamic Game

Our economy consists of a private sector and a policymaker that can be one of the two types, but whose actions do not directly reveal his type: a dynamic game with incomplete information. We now describe equilibrium in this game.

3.1 Public Equilibria

Define the public history of the current regime $h_t = \{h_{t-1}, \pi_{t-1}, \varsigma_t\}$ as the collection of all past realizations of inflation rates and exogenous states, with $h_0 = \{\rho_0, \varsigma_0\}$ being the public history of a new regime. We restrict our attention to equilibria in which all strategies depend only on the public history, i.e., "public strategies."[21] We denote the committed and opportunistic policymaker's equilibrium strategies as $\{a(h_t)\}_{t=0}^{\infty}$ and $\{\alpha(h_t)\}_{t=0}^{\infty}$, respectively. Comparably, we can write inflation expectations as $\{e(h_t)\}_{t=0}^{\infty}$.

[19] More specifically, $\rho_0 = \phi_t\rho_t + (1-\phi_t)v_{\rho,t}$, where $\phi_t \sim$ Bernoulli(δ_ρ) indicates whether the new policymaker inherits his predecessor's type (or equivalently, reputation), and $v_{\rho,t} \sim$ Beta($\overline{\rho}, \sigma_\rho$) is a random draw from a Beta distribution with mean $\overline{\rho}$ and standard deviation σ_ρ.

[20] We earlier specified that these densities are normal, but we use this notation to indicate the broader applicability of our analysis.

[21] This restriction is innocuous in our equilibrium analysis because: (1) the private sector's strategy is public since its information set is h_t; (2) the committed type's policy is public since it follows the announced policy plan, which needs to be verifiable by the private sector; and (3) given all the other player's strategies are public, it is also optimal for the opportunistic type to choose public strategies (Mailath and Samuelson (2006)).

3.2 Perfect Bayesian Equilibria

We further require the equilibrium of this incomplete information game to be perfect Bayesian. That is, the beliefs of the private sector are consistent and the strategies of the two types of policymakers satisfy sequential rationality.

3.2.1 Consistent beliefs: reputation

Consistency of beliefs requires the private sector's assessment of policymaker type is updated according to Bayes' rule (4) which depends on policymakers' equilibrium strategies and observed inflation π_t. Within a regime, the private sector's belief ρ is updated recursively,

$$\rho\left(h_{t+1}\right) = \rho\left(h_t, \pi_t\right) \equiv \frac{\rho\left(h_t\right) g(\pi_t | a\left(h_t\right))}{\rho\left(h_t\right) g(\pi_t | a\left(h_t\right)) + (1-\rho\left(h_t\right)) g(\pi_t | \alpha\left(h_t\right))} \tag{4}$$

With policymaker replacement, the regime clock t is reset to zero and reputation is $\rho_0 \sim \Xi(\rho_0 | \rho(h_t))$, given the inheritance mechanism for reputation discussed above.

3.2.2 Consistent beliefs: inflation expectations

Inflation expectations must be consistent with private sector beliefs about policymaker type and equilibrium strategies. With replacement, the consistent nowcast of inflation is:

$$z(h_t) = \int \left[\rho_0 a(\rho_0, \varsigma_t) + (1-\rho_0)\alpha(\rho_0, \varsigma_t)\right] d\Xi(\rho_0 | \rho(h_t)). \tag{5}$$

Within a regime, expectations of future inflation also reflect unknown policymaker type:

$$e(h_t) = \beta E(\pi_{t+1} | h_t) = \beta \rho(h_t) E(\pi_{t+1} | h_t, \tau_t = 1) + \beta(1-\rho(h_t)) E(\pi_{t+1} | h_t, \tau_t = 0) \tag{6}$$

with

$$E(\pi_{t+1} | h_t, \tau_t = 1) = \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) \left[(1-q)\, a(h_{t+1}) + qz\left(h_{t+1}\right)\right] g(\pi_t | a(h_t)) d\pi_t$$

$$E(\pi_{t+1} | h_t, \tau_t = 0) = \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) \left[(1-q)\, \alpha(h_{t+1}) + qz\left(h_{t+1}\right)\right] g(\pi_t | \alpha(h_t)) d\pi_t$$

Specifically, when private agents form date t inflation expectations, they know that (i) there is a committed type with $\rho_t = \rho(h_t)$,[22] and (ii) the committed type's intentions lead to stochastic inflation, with density $g(\pi_t|a(h_t))$, contributing to history $h_{t+1} = \{h_t, \pi_t, \varsigma_{t+1}\}$. Hence, if the regime continues next period, the committed type's intended inflation will be $a(h_{t+1})$. In the event of a regime change next period, the consistent belief is the history-dependent future nowcast $z(h_{t+1})$. Similarly, with probability $1-\rho_t$, the current policymaker is opportunistic and will generate stochastic inflation π_t with density $g(\pi_t|\alpha(h_t))$ and will implement $\alpha(h_{t+1})$ next period if the regime continues. In the event of a regime change next period, the expected inflation is $z(h_{t+1})$.

3.2.3 Sequential rationality of the committed type

The committed policymaker selects and announces a state-contingent plan for current and future intended inflation $\{a_t\}_{t=0}^{\infty}$ at the start of his term and then subsequently executes it.

The strategy of the committed type is *sequentially rational* if it maximizes his expected present discounted payoff at the beginning of his term,[23]

$$U_0 = \sum_{t=0}^{\infty} (\beta_a(1-q))^t \sum_{h_t} p(h_t) \underline{u}(a_t, e(h_t), \varsigma_t), \tag{7}$$

where $\underline{u}(a, e, \varsigma) \equiv \int u(\pi, x(\pi, e, \varsigma))\, g(\pi|a)\, d\pi$ is the expected momentary objective when the NK Phillips curve (1) is used to replace x with $x(\pi, e, \varsigma) = (\pi - e - \varsigma)/\kappa$. Note that (7) employs the probability of a specific history $h_t = [\varsigma_t, \pi_{t-1}, h_{t-1}]$ when inflation is generated by the committed type, i.e.,[24]

$$p(h_t) = \varphi(\varsigma_t; \varsigma_{t-1}) g(\pi_{t-1}|a(h_{t-1})) p(h_{t-1}) \tag{8}$$

combining the likelihood of the shock ς, the likelihood of inflation π given the committed type's decision, and the probability of the previous history.

In selecting the state-contingent plan at $t = 0$, the committed type takes into account the strategic power of his plan in shaping private sector inflation expectations. We consider

[22] With a slight abuse of notation, in the start of a new regime, $\rho(h_0) = \rho_0$.

[23] We assume the committed policymaker maximizes payoffs within his own term, so his discounting includes both the time discount factor β_a and the replacement probability q.

[24] There is a slight abuse of notation here by using summation Σ over history to capture the joint effects of continuous distribution of π and discrete Markov chain distribution of ς.

this crucial element further below.

3.2.4 Sequential rationality of the opportunistic type

An opportunistic policymaker chooses intended inflation α each period to maximize the expected objective, taking the response of expected inflation to history $\{e(h_t)\}_{t=0}^{\infty}$ as given:

$$\alpha(h_t) = \underset{\alpha}{\operatorname{argmax}}\, \underline{u}(\alpha, e(h_t), \varsigma_t) \tag{9}$$

where $\underline{u}(\alpha, e, \varsigma) \equiv \int u(\pi, x(\pi, e, \varsigma))\, g(\pi|\alpha)\, d\pi$ with $x(\pi, e, \varsigma) = (\pi - e - \varsigma)/\kappa$. The quadratic objective implies a linear best response of α to e and ς.

$$\alpha_t = Ae_t + B(\varsigma_t) \tag{10}$$

with $A = \vartheta_x/(\vartheta_x + \kappa^2)$, and $B(\varsigma_t) = (1 - A)\pi^* + A\kappa x^* + A\varsigma_t$.

Since Kydland and Prescott (1977), it has been understood that there is inflation bias when the central bank cannot commit. In our setup, the best response function (10) implies that the extent of inflation bias $\alpha - \pi^*$ varies with private sector's expected inflation $e_t = \beta(E_t\pi_{t+1})$. To highlight the sources of inflation bias, we have found it helpful to rewrite the best response function (10) as $\alpha_t - \pi^* = \iota + A\beta(E_t\pi_{t+1} - \pi^*)$ by denoting $\iota \equiv A(\kappa x^* - (1-\beta)\pi^*)$ and setting $\varsigma = 0$.[25] If private sector expects the inflation to be at target, i.e., $E_t\pi_{t+1} = \pi^*$, the optimal inflation bias is ι; we define this as *intrinsic inflation bias*. Figure 2 plots the best response function with the 45 degree line. The intersection of the two lines (the square marker) is the well-known *Nash equilibrium (NE) inflation bias* in which policy without commitment is fully expected (i.e., when $e = \beta\alpha$, $\alpha(e) - \pi^* = \iota/(1 - A\beta)$). The Figure highlights that Nash inflation bias can be much larger than intrinsic inflation bias (marked with a diamond) when $A\beta$ is close to one, as it will be in our quantitative model.

Our imperfect information framework will capture the dynamics of inflation and expectations in the 1970s as the outcome of expectations gradually increasing as private agents learn that there is an opportunistic policymaker behaving according to (10). Foreshadowing this finding, the Figure also includes two points that correspond to the one-quarter-head inflation forecasts by the professional forecasters (SPF) at two dates to illustrate the influence of rising expectations on opportunistic policy.

[Figure 2 about here.]

[25] As is conventional, these inflation bias measures are derived without any shock ς.

3.3 Public Perfect Bayesian Equilibrium

We now define our dynamic game's Public Perfect Bayesian Equilibrium (PBE).

Definition 1. A Public Perfect Bayesian Equilibrium is a set of functions in each history $\{z(h_t), e(h_t), \rho(h_t), \alpha(h_t), a(h_t)\}_{t=0}^{\infty}$ such that:

(i) given $\alpha(h_t)$, $a(h_t)$, and $\rho(h_t)$, the private sector's nowcast of inflation $z(h_t)$ conditional on a replacement satisfies (5);

(ii) given $\alpha(h_t)$, $a(h_t)$, and $z(h_t)$, the private sector's belief of policymaker type $\rho(h_{t+1})$ is updated according to (4); and its expected inflation function $e(h_t)$ satisfies (6);

(iii) given the expected inflation function, $e(h_t)$, the action of the opportunistic type policymaker $\alpha(h_t)$ maximizes his expected payoff (9);

and, at the start of a regime (t=0),

(iv) the strategy for the committed type policymaker $\{a(h_t)\}_{t=0}^{\infty}$ maximizes his expected payoff (7), taking into account the strategic power of $\{a(h_t)\}_{t=0}^{\infty}$ on $\{e(h_t)\}_{t=0}^{\infty}$ and $\{\alpha(h_t)\}_{t=0}^{\infty}$.

By "strategic power" of $\{a(h_t)\}_{t=0}^{\infty}$ on $\{e(h_t)\}_{t=0}^{\infty}$, we mean the influence that the committed policymaker's state-contingent plan – his *strategy* – has on the response of e_t to history h_t.

Given consistent private sector inflation expectations (6), there are three channels of influence. First, $e(h_t)$ is partially anchored by future committed policy $a(h_{t+1})$. Second, the extent of this anchoring depends on $\rho(h_t)$ which itself is affected by past committed policy $a(h_{t-1})$. Third, both $e(h_t)$ and $\rho(h_t)$ depend on intended inflation of a possible opportunistic policymaker $\alpha(h_{t+1})$ and $\alpha(h_{t-1})$. Sequential rationality of the opportunistic policymaker makes $\{\alpha(h_t)\}_{t=0}^{\infty}$ a best response to $\{e(h_t)\}_{t=0}^{\infty}$. Therefore, via shaping $\{e(h_t)\}_{t=0}^{\infty}$, the committed state-contingent plan also indirectly determines $\{\alpha(h_t)\}_{t=0}^{\infty}$.

4 Constructing the Equilibrium

Construction of the Public PBE is usefully viewed as inner and outer loops of a program. The inner loop builds a within-regime equilibrium $\{e(h_t), \rho(h_t), \alpha(h_t), a(h_t)\}$ taking as given beliefs $z(h_t)$ about the consequences of a regime change. The outer loop adjusts the beliefs $z(h_t)$ to be consistent with future regime outcomes, i.e., to attain a fixed point between $z(h_t)$ and $\{a(h_t), \alpha(h_t), \rho(h_t)\}$.

4.1 Our novel principal-agent approach

Solving the within-regime equilibrium may appear to be a formidable task, due to the strategic power of the committed policy plan $\{a(h_t)\}_{t=0}^{\infty}$ over private sector expectations and opportunistic policies. On one hand, the optimal choice for a committed policymaker depends on what the opportunistic type would do in the same history since private sector inflation expectations average across both types' future policy choices. On the other hand, the committed type's optimization cannot take future opportunistic policy as given since the opportunistic type responds to inflation expectations and in turn the committed policy plan.

To tackle these complications, we recast the within-regime equilibrium as the solution to a principal-agent problem. As principal, the committed policymaker maximizes (7) by choosing state contingent plans for his actions and those of two agents, the private sector and the opportunistic policymaker. Incentive compatibility (IC) constraints of two forms are relevant: (i) private sector consistent beliefs (4) and rational expectations (6); and (ii) opportunistic type optimal response to expected inflation (10).

4.2 Recursive formulation

Our framework is unusual because private agents disagree with the principal – the committed policymaker – in beliefs about the probability of a specific history. The private sector *thinks* that current inflation could be generated by the opportunistic policymaker, as captured in the third line of the expression for expected inflation (6) above. By contrast, the committed policymaker *knows* that current inflation is generated by his policy choices, as reflected in $p(h_t)$ in the intertemporal objective (7). A key necessary step in recursive formulation is to cast the Lagrangian component associated with the rational expectation constraint (6) into recursive form.[26] Disagreement in probability beliefs between principal and agent poses a challenge in this regard. We overcome it by a "change of measure". Attaching a multiplier $\gamma(h_t)$ and the committed type's probability of history $p(h_t)$ as weights to the constraint (6), we form the Lagrangian component as:

$$\Psi_0 = \sum_{t=0}^{\infty} (\beta_a(1-q))^t \sum_{h_t} p(h_t)\gamma(h_t)[e_t - e(h_t)], \tag{11}$$

where $e(h_t)$ is given by (6). Then, in (6), we write $E(\pi_{t+1}|h_t, \tau_t = 0)$ in terms of the committed type's probabilities, replacing $g(\pi_t|\alpha(h_t))$ with $\lambda(\pi_t, a_t, \alpha_t)g(\pi_t|a(h_t))$ where $\lambda(\pi_t, a_t, \alpha_t) \equiv$

[26] Following Kydland and Prescott (1980), Chang (1998), Phelan and Stacchetti (2001) and Marcet and Marimon (2019).

$g(\pi_t|\alpha_t)/g(\pi_t|a_t)$ is the likelihood ratio. This permits us to express Ψ recursively, so that the dynamic Lagrangian $U_t + \Psi_t$ is also recursive. Defining W_t as the optimized dynamic Lagrangian, we then establish:[27]

Proposition 1. The within-regime equilibrium is the solution to a recursive optimization problem, given $z(\varsigma, \rho)$ and the IC constraint $\alpha = Ae + B(\varsigma)$

$$W(\varsigma, \rho, \mu) = \min_{\gamma} \max_{a,\alpha,e} \{\underline{u}(a, e, \varsigma) + (\gamma e - \mu\omega) + \beta_a (1-q) \int \sum_{\varsigma'} \varphi(\varsigma'; \varsigma) W(\varsigma', \rho', \mu') g(\pi|a) d\pi\}, \tag{12}$$

$$\text{with } \omega \equiv (1-q)a + qz(\varsigma, \rho) + \frac{1-\rho}{\rho}[(1-q)\alpha + qz(\varsigma, \rho)] \tag{13}$$

$$\mu' = \frac{\beta}{\beta_a (1-q)} \gamma\rho, \text{ given } \mu_0 = 0 \tag{14}$$

$$\rho' = \frac{\rho g(\pi|a)}{\rho g(\pi|a) + (1-\rho) g(\pi|\alpha)} \text{ with prob } g(\pi|a), \text{ given } \rho_0 \tag{15}$$

This program enables us to analyze optimal choices of the principal – the committed policymaker – who faces private sector skepticism about his type. The optimal decisions a, α, e, γ each depend on the state vector $s = [\varsigma, \rho, \mu]$. The component $(\gamma e - \mu\omega)$ arises from the Lagrangian component of the forward-looking rational expectations constraints (11) expressed in the recursive form.[28] The pseudo state variable μ records past promises (contained in ω) made by the committed type.[29]

With two possible policymaker types and stochastic replacement, the composite promise term ω defined in (13) contains more than the committed type's promised a, because the expected inflation by private agents also depends on their perceived inflation α intended by the opportunistic type and their nowcast of inflation z in a new regime.[30] The weights attached to a, α, and z reflect the exogenous replacement probability q, the endogenous rep-

[27] Appendix A provides a detailed derivation of the recursive program.

[28] Our rational expectations constraint (6) is equivalent to the Phillips curve. Viewing it as an inequality constraint, with $x_t \leq (\pi_t - \beta E_t \pi_{t+1} - \varsigma_t)/\kappa$, the Phillips curve defines a set of feasible output gaps and inflation rates. Thus, the associated multiplier γ is nonnegative.

[29] The pseudo state variable terminology originates with Kydland and Prescott (1980). A new policymaker isn't held accountable for predecessor promises, so μ is initially zero.

[30] Note $\omega = a$ when $q = 0$, $\beta_a = \beta$, and $\rho = 1$. This is a textbook NK policy problem in recursive form. Appendix A.10 provides a fuller discussion.

utation state ρ, and the divergent probability beliefs about inflation π held by the committed policymaker and private agents.[31]

The evolution for the pseudo state μ and the reputation state ρ identifies two channels through which the committed policymaker manages inflation expectations by private agents.[32]

Expectation anchoring: The next-period pseudo state μ' evolves according to (14): a higher γ increases μ', making it costlier for the committed type to raise a' in the subsequent period period. For convenience, we express this as $\mu' = m(s)$. In this context, the committed policymaker selects the shadow price γ of promising a' and μ' accounts for this promise. The impact of μ' on a' is rationally anticipated by private agents, allowing the choice of γ to anchor inflation expectations. The effectiveness of this anchoring is moderated by private sector skepticism, as the influence of γ on μ' depends on the reputation state ρ.

Reputation building: The next-period reputation state ρ' evolves according to Bayes' rule (15). Since both a and α are functions of the state vector, it is convenient to express the Bayes' rule as $b(s, \pi)$. The committed policymaker affects ρ' by choosing a difference between his intended inflation (a) and the intended inflation of an opportunistic type (α). A larger policy difference, denoted by $\delta = a - \alpha$, accelerates private sector learning about the current policymaker type.[33] A higher ρ' influences the intended inflation for both policymaker type a' and α' in the subsequent period and increases the weight given to the committed policymaker's intended inflation in the private agents' expected inflation.

4.3 The PBE fixed point requirement

In a PBE, the nowcast of inflation $z^*(\varsigma, \rho)$ in a new regime must satisfy

$$z^*(\varsigma, \rho) = \int [\rho_0 a^*(\varsigma, \rho_0, 0; z^*(\varsigma, \rho)) + (1 - \rho_0)\alpha^*(\varsigma, \rho_0, 0; z^*(\varsigma, \rho))] d\Xi(\rho_0|\rho) \tag{16}$$

with $a^*(.)$ and $\alpha^*(.)$ obtained from the recursive program (12) given $z^*(\varsigma, \rho)$, and $\mu_0 = 0$ as a new policymaker is not held accountable for prior commitments made by his predecessor.[34]

[31] This final feature leads to $(1 - \rho)/\rho$ in ω. Appendix A.9 explains how we eliminate the likelihood ratio λ using Bayes' rule.

[32] Lemma 2 in Appendix B.2 formalizes the two channels.

[33] This channel is formalized when we simplify (15) to $\rho' = \rho'(v_\pi, \delta, \rho)$ by replacing $g(\pi|a) = g(v_\pi)$ and $g(\pi|\alpha) = g(\pi - a + a - \alpha) = g(v_\pi + \delta)$, where $g(\cdot)$ is the density of v_π.

[34] Schaumburg and Tambalotti (2007) impose a similar fixed point requirement in constructing an equilibrium in which a committed policymaker is randomly replaced.

4.4 Time series implications of the Public PBE

As a transition to quantitative analysis, we now consider the time invariant dynamic system that is implied by the Public PBE.[35] According to Proposition 1, the state vector $s = [\varsigma, \rho, \mu]$ determines the intended inflation policies, $a(s)$ and $\alpha(s)$, and the private sector expected inflation $e(s)$. At the end of a time period, the random inflation π is realized: $\pi = \tau a(s) + (1 - \tau)\alpha(s) + v_\pi$, where $\tau = 1$ indicates a committed policy regime and $\tau = 0$ indicates an opportunistic policy regime.

At the start of the next period, a new cost-push shock ς' will be drawn according to $\varphi(\varsigma'; \varsigma)$. The evolution of the reputation state ρ' and the pseudo state μ' will depend on the realizations of two random events. If the event of policymaker replacement does not occur ($\theta' = 0$), the reputation state will be updated via the Bayes' rule as $\rho' = b(s, \pi)$; and the pseudo state will evolve via $\mu' = m(s)$. If the replacement event occurs ($\theta' = 1$), the pseudo state $\mu' = 0$ and the reputation state $\rho' = \phi' b(s, \pi) + (1 - \phi')v'_\rho$. That is, if the new policymaker inherits his predecessor's reputation then ($\phi' = 1$), $\rho' = b(s, \pi)$. If the inheritance does not occur ($\phi' = 0$), the new policymaker's reputation is a random draw v'_ρ.

Thus, there will be a recursive evolution of $S = [s, \pi]$ and the recursion is conditional on the realizations of θ, ϕ, and τ. Private agents know the outcomes of θ, ϕ, and v_ρ, but not τ.

5 Building the quantitative model

We build the quantitative model in two stages. First, calibrating model parameters, we can compute the optimal decision functions using the recipe in Proposition 1, yielding $a(s)$, $\alpha(s)$ and $e(s)$. The Markovian structure also allows us to compute functions for private sector inflation forecasts at other horizons, $f(s_t, j) = E(\pi_{t+j}|s_t)$. Our theory reveals that there are three state variables $s_t = [\varsigma_t, \rho_t, \mu_t]$ – including the highly persistent reputation state ρ, a more temporary cost-push shock ς, and a predetermined pseudo state μ – but we must supply empirical counterparts. Second, as just discussed, our model has three structural shocks, $v_t = (v_\varsigma, v_\rho, v_\pi)$, to the process of cost-push shock, reputation in the event of replacement, and inflation, respectively. It also features three binary states (θ, ϕ, τ), indicating the event of policymaker replacement, whether or not there is reputation inheritance, and the type of current policymaker. To generate time series implications, we must develop and employ model functions of the three state variables that map the structural shocks and the binary

[35] This juncture also marks a shift in how we will use t. To this point is has been a regime clock. Now, it becomes a calendar indicator in time series analysis.

states to macro variables that can be directly measured using data.

We use a novel empirical strategy to jointly identify the continuous states s, the shocks v and the binary states (θ, ϕ, τ) by requiring that our model's expectations match time series from the Survey of Professional Forecasters. We convert our model to a state-space representation where the three state variables enter as latent states, and the three binary states enter as the outcome of an unobserved discrete-state Markov process, following Hamilton (1989) and Kim (1994). The transition probability matrix of the Markov process is designed to capture the interdependence of the three binary states. We then employ an efficient unscented Kalman filter with hidden Markov-switching to obtain the filtered and smoothed estimates of the latent continuous states and the probabilities of the discrete states. To validate our model, we use the identified states to construct model-implied variables that are not targeted in the filtering exercise and compare them to the observed data time series.

5.1 Calibration

Table 1 reports the calibrated values of important model parameters. One period is a quarter. The long-run inflation target π^* is 1.5%, which lies in the 1 to 2 percent range frequently cited by central bankers advocating price stability.[36] The private sector and committed type share a conventional quarterly discount factor based on a 2% annual real rate.

The slope of the Phillips curve and the policymaker's concerns about real activity are central elements in any study of inflation policy. In our setup, the PC slope κ relates the output gap x to the quarterly inflation π, holding expected inflation fixed. $\kappa = .08$ implies that an output gap of 3% leads to annualized inflation of -1%, a value compatible with diverse empirical evidence.[37]

[Table 1 about here.]

Turning to the preference parameters, we set the weight on output ϑ_x to 0.1, which is in the middle of the range used by prominent Fed researchers.[38] Together with $\kappa = .08$, it

[36] This value matches the estimate of Shapiro and Wilson (2019) in a careful and informative study of FOMC transcripts.

[37] U.S. data from the 1950s and 1960s suggests that a 1% decrease in unemployment led to about 0.54% - 0.65% increase in inflation. An estimate for Okun's coefficient is about 1.67 using U.S. data prior to 2008, implying a 1% increase in unemployment led to a 1.67% decrease in output. In a structural NKPC, the parameter is also consistent with an adjustment hazard leading to four quarters of stickiness on average and an elasticity of marginal cost with respect to output of unity.

[38] Brayton et al. (2014) after translating time units and using Okun's law.

implies $A = .94$ according to $A = \vartheta_x/(\vartheta_x + \kappa^2)$. The target output gap x^* is chosen to yield a relatively small intrinsic inflation bias $\iota = .5\%$ while yielding a NE bias large enough to capture the magnitude of the Great Inflation: $\iota/(1 - A\beta)$ is around 8%. Recall that x^* is linked to ι via $\iota = A(\kappa x^* - (1 - \beta)\pi^*)$. Hence, the implied value for $x^* = 1.73\%$.[39]

The replacement probability of $q = .03$ implies an average regime duration of 8 years. We have less empirical guidance about the inheritance mechanism for reputation: $\rho_0 = \phi_t\rho_t + (1-\phi_t)v_{\rho,t}$ with $\phi_t \sim \text{Bernoulli}(\delta_\rho)$ and $v_{\rho,t} \sim \text{Beta}(\overline{\rho}, \sigma_\rho)$. But our equilibrium policy functions are not sensitive to these parameters due to the small replacement probability q. We set $\delta_\rho = 0.9$, $\overline{\rho} = 0.1$, and $\sigma_\rho = 0.05$ so that the new policymaker inherits his predecessor's reputation with probability .9. Otherwise, his initial reputation is random with mean .1 and standard deviation 0.05.

Beginning in the 1970s, many studies of inflation use an observable "Food and Energy price shock"(FE shock hereafter).[40] We use the FE shock's serial correlation and its standard deviation as the cost-push shock's persistence δ_ς and innovation volatility $\sigma_{v,\varsigma}$. The transition probability matrix $\varphi(\varsigma';\varsigma)$ is calibrated to approximate $\varsigma' = \delta_\varsigma\varsigma + v_\varsigma$ where $v_\varsigma \sim N(0, \sigma^2_{v,\varsigma})$.[41] To calibrate the standard deviation of implementation errors, we combined the FE shock and the SPF one-quarter-ahead inflation forecast in an initial approximation to opportunistic intended inflation α, estimating the standard deviation of $(\pi - \alpha)$ over 1964Q4-1979Q2.

5.2 State extraction strategy

Figure 3 plots three-quarter-ahead SPF forecast of inflation (SPF3Q) against its one-quarter-ahead counterpart (SPF1Q), highlighting the smoother nature of SPF3Q relative to SPF1Q.[42] Taking a cue from literature on the term structure of interest rates, we form an SPF spread, plotted as the black dashed line and defined as SPF1Q-SPF3Q. Notice that the SPF spread rises during the first (1974-75) and the second (1978-80) inflation surges. This is consistent with our theory's implication that longer-term forecasts (SPF3Q) depend more on the persistent reputation variable ρ_t, while shorter-term forecasts are more sensitive to transitory cost-push shocks ς_t. We exploit this feature of data in our state extraction strategy, which is

[39] $x^* = 1.73\%$ is equivalent to targeting unemployment about 1% below the natural rate, if we use an Okun's law coefficient of 1.67.

[40] See R.J. Gordon (2013) and Watson (2014). It is constructed as the difference between the growth rate of the overall personal consumption deflator and its counterpart excluding food and energy.

[41] We use the Rouwenhorst (1995) method.

[42] Elmar Mertens guided us to the SPF term structure via Mertens and Nason (2020). We do not use SPF4Q due to missing observations, particularly important in 1975.

required because, as outside observers (econometricians), we do not know the state variables.

Using calibrated parameters, our theory provides a function $f(s, j)$ for private agent expectations at each horizon j.[43] Denote the SPF at horizon j as $f_{t+j|t}$. If the dates of policymaker replacement are known and an exact match between model and data expectations is assumed, it is possible to simply "invert" the theoretical relationship at each date to find an estimate of ς_t and ρ_t, since μ_t evolves deterministically as a function of these other states.[44] However, we instead adopt the more standard approach of treating the dates of policymaker replacement as unknown to econometricians along with the state variables and assuming that the model-implied variable differ from the data by a Gaussian observation error, $f_{t+j|t} = f(s_t, j) + \varepsilon_{jt}$.

With unknown dates of policymaker replacement, we share the challenges of the literature on Markov switching models (Hamilton (1989)), because the state evolution of μ depends on the realization of a replacement event (θ). As discussed in Section 4.4, our model has two more binary states ϕ and τ that are unobservable to us as econometricians but enter the time series recursion of $S = [s, \pi]$. Following the literature, we model these three binary states (θ, ϕ, τ) as the outcome of an unobserved discrete-state Markov process Θ. We define 6 discrete states of Θ: states $\{1, 3, 5\}$ corresponding to a continuing committed policymaker ($\theta = 0$, $\tau = 1$), a new committed policymaker with full reputation inheritance ($\theta = 1$, $\phi = 1$, $\tau = 1$), and a new committed policymaker with random reputation v_ρ ($\theta = 1$, $\phi = 0$, $\tau = 1$), and states $\{2, 4, 6\}$ corresponding to an opportunistic policymaker in a similar manner. We require that the transition probability matrix for the 6 discrete states respects the interdependence of (θ, ϕ, τ) imposed by our model structure, e.g., policymaker type cannot change without a replacement event, etc. Appendix C.3 displays this matrix.

[43] Appendix C.2 provides details for how to derive this function.

[44] In an earlier version of this research, King and Lu (2022), we used this approach, which Kollmann (2017) calls an "inversion filter" (see also Drautzburg et al. (2022)). We assumed a single replacement date at the start of the Reagan administration in 1981Q1, based on the narrative history of Goodfriend and King (2005).

5.3 The state space model with Markov-switching

We now detail the dynamics of continuous state variables $S_t = [\varsigma_t, \rho_t, \mu_t, \pi_t]$, taking as given the discrete state $\Theta_t = (\theta_t, \phi_t, \tau_t)$.

$$S_t = \begin{bmatrix} \delta_\varsigma \varsigma_{t-1} + v_{\varsigma,t} \\ (1-\theta_t+\theta_t\phi_t)b(s_{t-1},\pi_{t-1}) + \theta_t(1-\phi_t)v_{\rho,t} \\ (1-\theta_t)m(s_{t-1}) \\ \tau_t a(s_t) + (1-\tau_t)\alpha(s_t) + v_{\pi,t} \end{bmatrix} = F(S_{t-1}, v_t|\Theta_t)$$

The first entry specifies the process for the cost push shock ς. The second entry specifies that ρ_t is determined by the Bayes' rule $b(s_{t-1}, \pi_{t-1})$, if there is no replacement ($\theta = 0$) or if there is reputation inheritance ($\theta = 1$ and ϕ=1), while otherwise ρ_t is a random shock $v_{\rho,t}$ with support $[0, 1]$. The third entry indicates that the pseudo state variable evolves according to $\mu_{t+1} = m(s_t)$, except if there is replacement ($\theta = 1$) in which case it is set to zero. The final entry captures that inflation π_t depends on the type of policymaker in place.[45]

The one-quarter-ahead and three-quarter-ahead SPF forecasts are taken to be the model inflation forecasts corrupted by Gaussian observation errors ε_1 and ε_3.[46] That is, our observation equations are:

$$Y_t = \begin{bmatrix} f_{t+1|t} \\ f_{t+3|t} \end{bmatrix} = \begin{bmatrix} f(\varsigma_t, \rho_t, \mu_t, 1) \\ f(\varsigma_t, \rho_t, \mu_t, 3) \end{bmatrix} + \begin{bmatrix} \varepsilon_{1t} \\ \varepsilon_{3t} \end{bmatrix} = H(S_t) + \varepsilon_t$$

As our model is not linear, we cannot use the standard Kalman filter. We adopt a particular "sigma point" approximation method – the unscented Kalman filter – that has been shown to work well in nonlinear regime-switching models. Appendix C.3 provides details on the algorithm, which also employs the collapsing approach of Kim (1994) and Kim and Nelson (2017). For each element of S_t and Θ_t, our approach produces filtered estimates (based on $Y^t = [Y_1, Y_2, ...Y_t]$) and smoothed estimates (based on Y^T).

[45] Since the first three lines determine s_t as a function of S_{t-1}, one may use $a(s_t)$ and $\alpha(s_t)$ in the last line.

[46] ε_1 and ε_3 are i.i.d. normal random variables with mean zero and standard deviation 0.5% (annualized).

5.4 Estimates of continuous and discrete states

Figure 3 also plots the smoothed estimates of the cost-push shock $\widehat{\varsigma}_t$ (red) and the reputation state $\widehat{\rho}_t$ (cyan and measured on the right hand axis).[47] The estimated cost-push shock $\widehat{\varsigma}_t$ covaries positively with the SPF spread (SPF1Q-SPF3Q),[48] consistent with our strategy of exploiting greater sensitivity of near-term forecasts to transitory shocks. The estimated reputation state $\widehat{\rho}_t$ exhibits a big swing, declining from 0.7 in 1969 to near zero by the end of 1980 and finally climbing back to above 0.8 after 2000.

As an example of estimated conditional probabilities, Figure 4 plots the smoothed probabilities of a committed policy regime (blue) and of a policymaker replacement (red) in each period. The probability of a committed policy regime echos the dynamics of the estimated reputation state $\widehat{\rho}_t$ in Figure 3.[49] The probability is close to zero after 1975 and sharply increases to close to one in 1981-1982, suggesting that the most likely discrete state consistent with the observed SPF data switches from $\tau = 0$ (an opportunistic policy regime) to $\tau = 1$ (a committed policy regime). According to the model, policymaker's type can only switch in the event of a policymaker replacement. Our estimated probability of a replacement event peaks in the first quarter of 1982.

[Figure 3 about here.]

[Figure 4 about here.]

5.5 Targeted and untargeted variables

We now report the performance of the model-based non-linear Kalman method, in terms of fitting targeted time series, SPF1Q and SPF3Q, and matching untargeted time series.

5.5.1 Inflation expectations

Our extraction method produces a nearly perfect match for SPF1Q and SPF3Q. Using the extracted states, we can also compute model-implied inflation forecasts at horizons 2 and 4.

[47] The reported value is the probability-weighted average of smoothed estimates of state variables conditional on being in a discrete state: $\widehat{x}_t = \sum_{i=1}^{6} E(x_t|\Theta_t = i, Y^T)Pr(\Theta_t = i|Y^T)$.

[48] Appendix C.5 Figure 12 reports how our estimated cost-push shock covaries with the FE shock.

[49] The smoothed estimate of ρ_t is different from the smoothed probability of a committed policy regime. Our filter calculates the optimal estimates of $s_t = (\varsigma_t, \rho_t, \mu_t)$ for fitting the observed SPF data, given the assumption of being in a specific policy regime. Subsequently, it applies these regime-specific estimates to obtain the probability of that particular policy regime, taking into account the structure of shocks. The smoothed estimate of ρ_t is derived as a probability-weighted average of these regime-specific estimates.

Appendix C.5 Figure 11 shows that these additional forecasts lie almost entirely on top of the untargeted SPF2Q and SPF4Q, providing support for our state extraction approach.

5.5.2 Observed and estimated inflation

Our state-space model treats inflation as a latent state variable and therefore produces filtered and smoothed estimates for π_t. Because our extraction method only uses SPF data to obtain states, comparing the smoothed estimates $\widehat{\pi}_t$ with the observed inflation data serves as a model validation. To assess how well our estimates correspond to observed inflation data, we use the SPF1Q as a benchmark.

Figure 5 plots the observed inflation data (blue) against two series: one is the one-quarter-ahead SPF inflation forecast from the prior quarter (black) and the other is our method's smoothed estimates of inflation (red). The difference between observed inflation and the two series are plotted in dashed lines with corresponding colors. Therefore, the black dash line is the SPF forecast errors that are known to be persistent, with a serial correlation equal to 0.63 over our sample period. By contrast, the serial correlation of the errors of our smoothed estimates of inflation is only 0.46. We also compute the mean-squared-error (MSE) as another measure of fit: the MSE of the SPF1Q is 1.67 and the MSE of the smoothed estimates of inflation is only 0.99.

We conclude that our model-implied inflation captures the behavior of observed U.S. inflation. However, a skeptical reader might have two concerns. First, our smoothed measure of model inflation is based on the full sample, while the SPF is prepared in real time. In Appendix C.5 Figure 13, we therefore provide a version of Figure 5 with one-sided (filtered) estimates, revealing that the close correspondence of observed and model inflation is present even when our extraction method has no information advantage. Second, since our extraction is based on SPF1Q and SPF3Q, one reaction is that our method must work well because the SPF also tracks observed inflation. However, the extraction method chooses state estimates to produce model expectations close to the SPF, but does not guarantee that model inflation – governed by $\tau a(s) + (1-\tau)\alpha(s) + v_\pi$ – is close to observed inflation.

[Figure 5 about here.]

6 Evolving reputation: history and prospect

We have seen that our quantitative model yields time series that align closely with US inflation history over 1968 to 2007, in terms of fitting targeted expected inflation and matching untargeted observed inflation. We now further examine US inflation through the lens of our model, with core ingredients being a regime shift from opportunistic to committed policy around 1981; opportunistic policy optimally responding to expected inflation; and committed policy optimally influencing private sector learning.

In this section, we also conduct a counterfactual exercise to highlight the importance of having optimal committed policy influence private sector learning, a key mechanism that differentiates our paper from the literature. Drawing on some lessons from history and elements of our theory, we wrap up by discussing the current macroeconomic situation and prospects for inflation going forward.

6.1 Interpreting US inflation history 1969-2007

The smoothed estimates of states imply model-based *intended inflation policies*: $\widehat{a}_t$ for the committed type and $\widehat{\alpha}_t$ for the opportunistic type.[50] In Figure 6, we plot these optimal intended inflation policies, $\widehat{a}_t$ in green and $\widehat{\alpha}_t$ in red, together with their difference $\widehat{a}_t - \widehat{\alpha}_t$ (blue with circle markers), observed inflation data (black), and smoothed estimates of cost push shock $\widehat{\varsigma}_t$ (blue) and reputation $\widehat{\rho}_t$ (cyan and measured on the right hand axis).

Consistent with estimated probabilities in Figure 4, Figure 6 shows that US inflation is tracked by opportunistic policy before 1981 and by committed policy after 1982, with a transition during 1981-1982. This finding guides our interpretation of US inflation history, as we assume optimal opportunistic policy for the Great Inflation, and optimal committed policy for the Volcker Disinflation and the subsequent stabilization of inflation.

The Great Inflation: Our model portrays the Great Inflation a joint product of cost-push shocks and declining reputation. Advocates of the supply shock theory of the Great Inflation, such as Blinder and Rudd (2008), highlight the 1973-1975 surge and decline in inflation. These analysts point out that supply shocks – based on changes in relative prices – necessarily lead to temporary changes in inflation, so that they are well equipped to capture such "hills" as they do in our model.[51] But they acknowledge that their approach cannot explain why inflation is several percent higher *after* 1973-1975 than in the early 1970s.

Our framework captures this higher "plateau" of inflation as an optimal response of opportunistic policy to a decline in reputation. In Figure 6, estimated reputation (cyan line)

[50] $\widehat{a}_t = \sum_{i=1}^{6} E(a(s_t)|\Theta_t = i, Y^T)Pr(\Theta_t = i|Y^T)$ and $\widehat{\alpha}_t = \sum_{i=1}^{6} E(\alpha(s_t)|\Theta_t = i, Y^T)Pr(\Theta_t = i|Y^T)$

[51] Specifically, there are two relevant cases: (i) a temporary increase in a key price such as energy leads a high inflation period to be followed by a lower inflation period; and (ii) permanent changes in relative prices have at most a temporary effect on inflation.

is about .65 in 1972 and falls to around .3 after 1975, because a larger difference in policy responses (blue circled line) to cost-push shocks during the 1973-1975 episode makes it more likely that the policymaker in place is opportunistic. In the data and our model (Figure 3), expected inflation (SPF1Q) rises from 3.8% in 1972Q4 to 5.9% in 1976Q3.

A key feature of our framework is that an opportunistic policymaker responds to higher expected inflation by choosing a higher intended inflation, even in absence of cost-push shocks. Recall that our inflation bias diagram, Figure 2, plots the two points corresponding to 1972Q4 and 1976Q3 expected inflation, respectively, and optimal inflation bias is 2% higher in 1976 than in 1972.

In the late 1970s, another round of cost-push shocks (blue line) leads to further deterioration of reputation (cyan line) toward a trough by the end of 1980. The combination of cost-push shocks and declining reputation spurs a rapid rise in optimal opportunistic policy, culminating the second peak of inflation in the Great Inflation.

[Figure 6 about here.]

Volcker Disinflation: Our quantitative model estimates that the "Volcker Disinflation" did not start until 1981 and was the onset of a new committed policy regime. In Figure 6, the large difference between committed and opportunistic policies (blue circled line) persists through the Volcker disinflation and results in a rapid gain in reputation (cyan line) from a trough level of .02 in early 1981 to about .4 by the end of the recession in late 1982. The process of gaining reputation is difficult because as the reputation improves, expected inflation declines and brings down optimal opportunistic policy as well. As the policy difference shrinks with higher reputation, it becomes more difficult for private agents to determine the type in place. Once again, our optimizing approach to "alternative policy" is important, just as it was in the Great Inflation.

The Stabilization of Inflation starts with the end of the major recession in November 1982. From this point on, inflation is well known to be fairly stable and relatively low, particularly during the Greenspan years (1988-2005). In Figure 6 and after 1982, observed inflation (black line) roughly tracks the committed policy $\widehat{a}$ (green line), but the opportunistic policy $\widehat{\alpha}$ (red line) is also relatively low and stable. The policy difference (blue circled line) stabilizes around 0.6% during the period, leading to a slow rise of reputation (cyan line) from around .4 in early 1983 to around .85 in the 2000s.[52]

[52] Estimated reputation deteriorated during 1986-1987 and again during 2005-2006, corresponding to the end of chairmanship of Volcker and Greenspan, respectively. This suggests that anticipating a regime change

6.2 Counterfactual with naive committed policy

Our paper is not the first to investigate the evolution of private sector's beliefs about policymaker.[53] The main difference is that previous work abstracts from considering how a policymaker may want to affect those beliefs, which is a defining feature for optimal committed policy in our framework. In this subsection, we show that this new channel enhances our understanding of the evolution of private agents' beliefs.

We compute optimal policy functions when the committed policymaker acts naively, i.e. does not try to influence reputation evolution but simply views it as an exogenous leverage of his policy a on inflation expectations.[54] With these policy functions, we construct time series of naive intended inflation policies using the estimated cost-push shock from the benchmark model, the endogenous time-varying reputation governed by the naive policymakers' past responses to shocks, and the model-consistent pseudo state for the naive committed policymaker.

Policy Functions Figure 7 displays how committed policy a^* (bottom row), opportunistic policy α^* (middle row), and their policy difference $\delta^* = a^* - \alpha^*$ (top row) vary with reputation ρ. Each panel compares the policy function of two models: one where a committed policymaker acts naively (in red), and our benchmark model in which the committed policymaker actively manages his reputation (in blue). The policy functions are conditional on the values of two other state: μ and ς.[55] The cost push shock is set to be zero for the left column and to be 1% for the right column.

The gap between naive and benchmark policy behavior is most evident in the the policy difference functions (top panels). When reputation gets lower, the policy difference shrinks under naive policy but widens with benchmark policy. As a result, the policy difference is smaller in the naive policy model than it is in the benchmark policy model, especially at low levels of reputation. This sharp contrast is quite intuitive because a given policy difference incurs a larger output cost for the committed policymaker when his reputation is poorer. If the committed policymaker treats reputation as exogenous, he would prefer a smaller policy

is a feature of SPF but our model abstracts from it.

[53] See footnote 12 for examples.

[54] Appendix D explains the details of the naive optimization problem, which is related to work by Cogley and Sargent (2008) that builds on ideas of Kreps (1998). We thank Davide Debortoli for recommending the investigation of naive policy.

[55] Relative to ρ or ς, μ is less important in determining either the level or the shape of the policy functions. In this figure, we set μ at its steady state level when $\rho = 1$.

difference at lower reputation. Only when the committed policymaker understands that a larger policy difference will imply higher reputation in the future does he want to endure the contemporaneous output cost for future reputation gain: in this sense, the policymaker is sophisticated rather than naive.

Naive and sophisticated policymaker also respond differently to cost-push shocks. Comparing the left column with the right column reveals that the policy difference with naive policy does not increase in response to a positive cost-push shock, whereas it becomes larger in the benchmark (sophisticated) committed policy. The key point is that a positive cost-push shock provides a relatively cheaper opportunity for the committed policymaker to induce faster private sector learning.[56] So, the sophisticated committed policymaker responds to a positive cost-push shock with a larger policy difference. But such an altered incentive is absent when the committed policymaker takes reputation as exogenously given.

[Figure 7 about here.]

[Figure 8 about here.]

Time Series Figure 8 displays reputation and intended inflation policies when the committed policymaker is naive; as a reference, we replot time series from the benchmark model.

The two sets of time series share the same cost-push shocks (Figure 3) and probabilities of states as estimated using our benchmark model (Figure 4). That is, the opportunistic policy regime is more likely between 1973 and 1981 and the committed policy regime is more likely after 1981. We set realized implementation errors to zero in computing reputation dynamics under naive policy so as to focus on the effects of naive policymaker's past policies on private agents' beliefs.

First, observe that when the committed policymaker acts naively, reputation remains low for an extended period after 1981 when it is more likely that a committed policymaker is in charge. This is intuitive because, just as Figure 7 shows, a naive committed policymaker, lacking incentives to build reputation, tends to choose policies more similar to those of an opportunistic policymaker, particularly when confronted with a poor reputation. Correspondingly, the naive committed policymaker takes much longer to disinflate the economy than what is observed in post-1981 U.S. inflation history.

[56] It is cheaper in the sense that it takes a smaller deviation from the inflation target for the committed policymaker to induce a marginally larger policy difference because the opportunistic response to a positive cost-push shock is more inflationary.

Perhaps more surprising, the dynamics of reputation are affected by a committed policymaker acting naively, even during periods like 1973-1981 when an opportunistic policymaker is more likely to be in place. In the benchmark model, the rate of reputation decline is sensitive to cost-push shocks. It remains mostly moderate, accelerating only during two episodes of surge in cost-push shocks: 1973-1975 and 1978-1980. Analysis of policy functions above shows that this is because the optimal policy difference widens in response to positive cost-push shocks, as a sophisticated committed policymaker aims to influence private sector learning. Even though the committed policymaker was not present, rationally-expected changes in his policy response to cost-push shocks affect private sector learning.

By contrast, matters are different when there could be a committed policymaker that acts naively and private agents build that behavior into their expectations. Our analysis of naive policy indicates that the policy difference $\delta^* = a^* - \alpha^*$ does not respond to cost-push shocks or to reputation. Consequently, reputation dynamics is insensitive to cost push shocks during 1973-1981.

6.3 Looking Forward

Our quantitative analysis has so far focused on 1968Q4 through 2007Q4. We made this choice for several reasons. First, the sample matches that of leading studies of U.S. inflation's rise, fall and stabilization that were mainly undertaken prior to the Global Financial Crisis.[57] Second, our analysis abstracted from monetary policy instruments and fiscal actions, even as these varied through US history, because we viewed these as subordinated to intended inflation. Yet, the now-standard theory of policy with short-term interest rates at zero requires an aggregate demand specification and imposes additional constraints on our policy problem, so we avoided these complications.[58]

[Figure 9 about here.]

However, our abstaining from interpreting longer-term U.S. inflation history to this point does not mean that our model performs poorly beyond 2007. In fact, Appendix E shows that our model performs well in matching the SPF term structure and observed inflation through 2023Q1.

Since 2020Q4, the U.S. inflation has moved from the 1.5 to 2 percent range, increasing to over 8 percent in 2022Q2. This recent development has prompted many to make comparisons

[57] Examples include Sargent (1999) and Primiceri (2006) but there are many others.

[58] See for example Eggertsson and Woodford (2003).

with the 1973-4 upswing in inflation.[59] We thus include the model-based interpretation of the longer US inflation history (Figure 9), with a focus on the most recent inflationary episode, to consider our model's relevance for today's policy and the years ahead.

Figure 9 provides our SPF-based smoothed estimates of reputation (cyan, measured on right hand axis), the cost-push shock (blue), the probability of a committed policy regime (magenta, measured on right hand axis), the model-implied intended inflation policies $\widehat{a}$ (green) and $\widehat{\alpha}$ (red), and the observed inflation data (black). Two brown boxes highlight two episodes: 1973-1975 and 2020-2022.

The start of the recent inflationary episode is marked by the Spring 2020 onset of the COVID-19 pandemic and the Summer 2020 Fed announcement of a shift to a flexible average inflation targeting (FAIT) with a short-run inflation target above 2 percent. Prior to 2020Q3, estimated reputation $\widehat{\rho}_t$ fluctuated above around 0.9. In the subsequent 8 quarters, $\widehat{\rho}_t$ declined toward 0.8. At the same time, the estimated cost push shock $\widehat{\varsigma}_t$ increased from around 0 to 1.2%. According to our model, the combination of deteriorating reputation and positive cost-push shocks will lead to rising intended inflation policy for both types of policymaker: the committed policy $\widehat{a}_t$ reaches 4.6% and the opportunistic policy $\widehat{\alpha}_t$ reaches 5.1% by 2022Q2. These qualitative features of the 2020-2022 episode resemble those of the 1973-1975 one.

However, there is a notable quantitative difference between the two episodes. While observed inflation surged sharply during 2020-2022, the decline in estimated reputation (from roughly 0.9 to 0.8) was considerably smaller than during 1973-1975 (from 0.65 to 0.3). This indicates that professional forecasters are relatively optimistic, attributing a significant portion of the inflation increase to positive but transitory implementation errors. Our model suggests that this quantitative difference is due to the starting level of reputation in each episode. As shown in Figure 7, a better reputation narrows the optimal policy difference in the benchmark model, implying slower private sector learning when a policymaker is more reputable. The model also offers a cautionary projection: if inflation remains high, reputation will further deteriorate, potentially creating a prolonged period of elevated inflation and inflation expectations, akin to the aftermath of 1973-1975. At that juncture, the cost of disinflation could resemble that of the 1980s.

[59] A Fed staffer from that period recalls Arthur Burn's fixation on special factors in inflation that were assumed to be transitory. He observes that "The US Federal Reserve is insisting that recent increases in (prices of specific goods) reflect transitory factors that will quickly fade with post-pandemic normalization. But what if they are a harbinger, not a "noisy" deviation?" Roach 2021.

7 Summary, Conclusions and Final Remarks

We present a novel theoretical approach to study the equilibrium of a dynamic policy game that features two types of purposeful policymakers, a committed type which can commit and an opportunistic type which cannot, and private agents who are Bayesian learners about policymaker type and form forward-looking expectations of future policies.

The committed policymaker strategically uses its policy plan to influence private agents learning and inflation expectations, understanding that (i) private agents inflation expectations include future policy of an opportunistic type; and (ii) an opportunistic type's optimal policy depends on private agents inflation expectations. We use the insights of modern contract theory to develop a computable recursive equilibrium for the dynamic game. This permits calculation of equilibrium policies of both policymaker types and the rational expectations of private agents as functions of only three state variables, including an important reputation state that captures the evolution of private agents belief about the commitment capacity of current policymaker.

Putting our theory to work, we show how our parsimonious model can simultaneously capture the expected and actual inflation in the U.S. We use our theoretical model's dynamic system to build a nonlinear filter with hidden Markov-switching and extract latent states of the model from just the SPF inflation data. The estimates from the nonlinear filter suggest a regime switch from an opportunistic regime to a committed regime around 1981. The model-implied inflation also tracks US inflation's rise, fall, and stabilization between 1970 and 2005 to a surprising high degree, even though the observed inflation is not used by the nonlinear filter.

Our quantitative exercise reveals that evolving reputation is very important in accounting for actual inflation. In particular, endogenous policy differences help explain why private sector learning is slow in early 1970s; why cost-push shocks in the mid 1970s sped up learning, intensifying and prolonging the Great Inflation; and why the Volcker disinflation may be understood as a committed policymaker rebuilding reputation lost during the Great Inflation.

These lessons from the 1970s and 1980s appear particularly relevant for the ongoing fight against inflation in the U.S. Small but persistent deviations of inflation from targets can eventually lead to run-away inflation expectations, even though such expectations may appear very sticky early on. Explicitly committing to inflation targets – including flexible inflation targeting – helps the central bank to acquire and maintain credibility for attaining its monetary policy objectives. Our theory highlights that reputation for commitment, a measure of long-term credibility, can be gained or lost.

Our model is deliberately stark. But it yields results that have surprised us and others. We believe its success in matching U.S. time series indicates great promise to further research on models that feature agents learning about the commitment capacity of purposeful policymakers.

References

Amador, Manuel and Christopher Phelan, "Reputation and sovereign default," *Econometrica*, 2021, *89* (4), 1979–2010. Publisher: Wiley Online Library.

Atkeson, Andrew and Patrick J. Kehoe, "The advantage of transparency in monetary policy instruments," *Federal Reserve Bank of Minneapolis Staff Report 297*, 2006.

Backus, David A. and John Driffill, "Inflation and Reputation.," *American Economic Review*, 1985, *75(3)*, 530–538.

Ball, Laurence, "Disinflation with Imperfect Credibility," *Journal of Monetary Economics*, 1995, *35* (1), 5–23.

Barro, Robert J., "Reputation in a model of monetary policy with incomplete information," *Journal of Monetary Economics*, January 1986, *17* (1), 3–20.

_ **and David B. Gordon**, "Rules, discretion and reputation in a model of monetary policy," *Journal of Monetary Economics*, 1983, *12* (1), 101–121.

Benigno, Gianluca, Andrew Foerster, Christopher Otrok, and Alessandro Rebucci, "Estimating Macroeconomic Models of Financial Crises: An Endogenous Regime-Switching Approach," April 2020.

Bianchi, Francesco, "Regime Switches, Agents' Beliefs, and Post-World War II U.S. Macroeconomic Dynamics," *The Review of Economic Studies*, oct 2013, *80* (2), 463–490.

Binning, Andrew and Junior Maih, "Sigma Point Filters for Dynamic Nonlinear Regime Switching Models," *SSRN Electronic Journal*, 2015.

Blinder, Alan and Jeremy Rudd, "The Supply-Shock Explanation of the Great Stagflation Revisited," Technical Report w14563, National Bureau of Economic Research, Cambridge, MA December 2008.

Brayton, Flint, Thomas Laubach, and David Reifschneider, "Optimal-Control Monetary Policy in the FRB/US Model," *FEDS Notes*, November 2014, *2014* (0035).

Carvalho, Carlos, Stefano Eusepi, Emanuel Moench, and Bruce Preston, "Anchored Inflation Expectations," *American Economic Journal: Macroeconomics*, 2022.

Chang, Roberto, "Credible Monetary Policy in an Infinite Horizon Model: Recursive Approaches," *Journal of Economic Theory*, 1998, *81*, 431 – 461.

Chari, V. V. and Patrick J. Kehoe, "Sustainable Plans," *Journal of Political Economy*, 1990, *98* (4), 783–802.

Clarida, Richard, Jordi Gali, and Mark Gertler, "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 1999, *37* (4), 1661–1707.

Clayton, Christopher, Amanda Dos Santos, Matteo Maggiori, and Jesse Schreger, "Internationalizing Like China," August 2022.

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

_ , Christian Matthes, and Argia M. Sbordone, "Optimized Taylor rules for disinflation when agents are learning," *Journal of Monetary Economics*, may 2015, *72*, 131–147.

Coibion, Olivier, Yuriy Gorodnichenko, and Rupal Kamdar, "The Formation of Expectations, Inflation, and the Phillips Curve," *Journal of Economic Literature*, December 2018, *56* (4), 1447–1491.

Cukierman, Alex and Allan H. Meltzer, "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information," *Econometrica*, September 1986, *54* (5), 1099.

_ and Nissan Liviatan, "Optimal accommodation by strong policymakers under incomplete information," *Journal of Monetary Economics*, February 1991, *27* (1), 99–127.

Debortoli, Davide and Aeimit Lakdawala, "How Credible Is the Federal Reserve? A Structural Estimation of Policy Re-Optimizations," *American Economic Journal: Macroeconomics*, jul 2016, *8* (3), 42–76.

DeLong, J. Bradford, "America's Only Peacetime Inflation: The 1970s," Technical Report h0084, National Bureau of Economic Research, Cambridge, MA May 1996.

Dovis, Alessandro and Rishabh Kirpalani, "Rules without Commitment: Reputation and Incentives," *The Review of Economic Studies*, feb 2021.

_ and _ , "Reputation, Bailouts, and Interest Rate Spread Dynamics," *American Economic Journal: Macroeconomics*, July 2022, *14* (3), 411–449.

Drautzburg, Thorsten, Jesus Fernandez-Villaverde, Pablo Guerron-Quintana, and Dick Oosthuizen, "Filtering with limited information," 2022.

Eggertsson, Gauti B. and Michael Woodford, "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 2003, *2003* (1), 139–211.

Erceg, Christopher J. and Andrew T. Levin, "Imperfect credibility and inflation persistence," *Journal of Monetary Economics*, May 2003, *50* (4), 915–944.

Evans, Martin and Paul Wachtel, "Inflation Regimes and the Sources of Inflation Uncertainty," *Journal of Money, Credit and Banking*, 1993, *25* (3), 475–511.

Farmer, Leland, Emi Nakamura, and Jón Steinsson, "Learning About the Long Run," Working Paper 29495, National Bureau of Economic Research November 2021.

Faust, Jon and Lars E. O. Svensson, "Transparency and Credibility: Monetary Policy with Unobservable Goals," *International Economic Review*, 2001, *42* (2), 369–397.

Foerster, Andrew and Christian Matthes, "Learning About Regime Change," *International Economic Review*, 2022, *63* (4), 1829–1859.

Goodfriend, Marvin and Robert G. King, "The incredible Volcker disinflation," *Journal of Monetary Economics*, July 2005, *52* (5), 981–1015.

Gordon, Robert, "The Phillips Curve is Alive and Well: Inflation and the NAIRU During the Slow Recovery," Technical Report aug 2013.

Hamilton, James D., "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 1989, *57* (2), 357–384.

Hazell, Jonathon, Juan Herreño, Emi Nakamura, and Jón Steinsson, "The Slope of the Phillips Curve: Evidence from U.S. States*," *The Quarterly Journal of Economics*, 02 2022, *137* (3), 1299–1344.

Kim, Chang-Jin, "Dynamic linear models with Markov-switching," *Journal of Econometrics*, January 1994, *60* (1-2), 1–22.

— and Charles R. Nelson, *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, The MIT Press, November 2017.

King, Robert G. and Yang K. Lu, "Evolving Reputation for Commitment: The Rise, Fall and Stabilization of US Inflation," December 2022.

—, —, and Ernesto S. Pastén, "Managing Expectations," *Journal of Money, Credit and Banking*, December 2008, *40* (8), 1625–1666.

Kollmann, Robert, "Tractable likelihood-based estimation of non-linear DSGE models," *Economics Letters*, December 2017, *161*, 90–92.

Kreps, David M., "Anticipated Utility and Dynamic Choice (1997)," in Donald P. Jacobs, Ehud Kalai, Morton I. Kamien, and Nancy L. Schwartz, eds., *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983–1997*, Econometric Society Monographs, Cambridge: Cambridge University Press, 1998, pp. 242–274.

Kreps, David M and Robert Wilson, "Reputation and imperfect information," *Journal of Economic Theory*, August 1982, *27* (2), 253–279.

Kurozumi, Takushi, "Optimal sustainable monetary policy," *Journal of Monetary Economics*, 2008, *55* (7), 1277–1289.

Kydland, Finn E. and Edward C. Prescott, "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 1977, *85* (3), 473–491.

— **and** — , "Dynamic optimal taxation, rational expectations and optimal control," *Journal of Economic Dynamics and Control*, January 1980, *2*, 79–91.

Levin, Andrew and John B. Taylor, "Falling Behind the Curve: A Positive Analysis of Stop-Start Monetary Policies and the Great Inflation," in "The Great Inflation: The Rebirth of Modern Central Banking," University of Chicago Press, June 2013, pp. 217–244.

Loisel, Olivier, "Central bank reputation in a forward-looking model," *Journal of Economic Dynamics and Control*, 2008, *32* (11), 3718–3742.

Lu, Yang K., "Optimal policy with credibility concerns," *Journal of Economic Theory*, September 2013, *148* (5), 2007–2032.

— **, Robert G. King, and Ernesto Pasten**, "Optimal reputation building in the New Keynesian model," *Journal of Monetary Economics*, December 2016, *84*, 233–249.

Lucas, Robert E. and Thomas J. Sargent, "After Keynesian macroeconomics," *Quarterly Review*, 1979, *3* (Spr). Publisher: Federal Reserve Bank of Minneapolis.

Mailath, George J. and Larry Samuelson, *Repeated Games and Reputations: Long-run Relationships*, Oxford University Press, July 2006.

Marcet, Albert and Ramon Marimon, "Recursive Contracts," *Econometrica*, 2019, *87* (5), 1589–1631.

Matthes, Christian, "Figuring Out the Fed-Beliefs about Policymakers and Gains from Transparency," *Journal of Money, Credit and Banking*, January 2015, *47* (1), 1–29.

Melosi, Leonardo, "Signalling Effects of Monetary Policy," *The Review of Economic Studies*, sep 2016, p. rdw050.

Meltzer, Allan H., *A history of the Federal Reserve. vol. 2, book 2: 1970 - 1986*, paperback ed ed., Chicago: Univ. of Chicago Press, 2014.

Mertens, Elmar and James M. Nason, "Inflation and professional forecast dynamics: An evaluation of stickiness, persistence, and volatility," *Quantitative Economics*, November 2020, *11* (4), 1485–1520.

Milgrom, Paul and John Roberts, "Limit pricing and entry under incomplete information: An equilibrium analysis," *Econometrica*, 1982, *50* (2), 443–460.

Mishkin, Frederic S. and Klaus Schmidt-Hebbel, "Does Inflation Targeting Make a Difference?," NBER Working Papers 12876 January 2007.

Morelli, Juan M. and Matias Moretti, "Information Frictions, Reputation, and Sovereign Spreads," *Journal of Political Economy*, March 2023, pp. 000–000. Publisher: The University of Chicago Press.

Orphanides, Athanasios and John C. Williams, "The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations," *Journal of Economic Dynamics and Control*, 2005, *29* (11), 1927–1950.

Phelan, Christopher, "Public trust and government betrayal," *Journal of Economic Theory*, September 2006, *130* (1), 27–43.

_ **and Ennio Stacchetti**, "Sequential Equilibria in a Ramsey Tax Model," *Econometrica*, 2001, *69* (6), 1491–1518.

Primiceri, Giorgio E, "Why Inflation Rose and Fell: Policy-Makers's Beliefs and U. S. Postwar Stabilization Policy," *Quarterly Journal of Economics*, aug 2006, *121* (3), 867–901.

Roger, Scott and Mark R. Stone, "On Target? the International Experience with Achieving Inflation Targets," *IMF Working Paper No. 05/163*, 2005.

Rouwenhorst, K. Geert, "Asset Pricing Implications of Equilibrium Business Cycle Models," in Thomas F. Cooley, ed., *Frontiers of Business Cycle Research*, Princeton University Press, 1995, pp. 294–330.

Sargent, Thomas J., *The conquest of American inflation*, 1st print ed., Princeton, NJ: Princeton University Press, 1999.

_ **and Ulf Soderstrom**, "The conquest of American inflation: A summary," *Sveriges Riksbank Economic Review*, 2000, *3*, 12–45.

Schaumburg, Ernst and Andrea Tambalotti, "An investigation of the gains from commitment in monetary policy," *Journal of Monetary Economics*, 2007, *54* (2), 302–324.

Shapiro, Adam and Daniel J. Wilson, "The Evolution of the FOMC's Explicit Inflation Target," *FRBSF Economic Letter*, April 2019.

Särkkä, Simo and Lennart Svensson, *Bayesian Filtering and Smoothing*, 2023 ed., Cambridge University Press, 2023.

Watson, Mark W., "Inflation Persistence, the NAIRU, and the Great Recession," *American Economic Review*, May 2014, *104* (5), 31–36.

Appendices

A Recursive optimal policy design

The optimal policy problem for the committed type at the start of its tenure involves forward-looking constraints, which must be transformed to yield a recursive specification. Conceptually, this involves casting Lagrangian components in recursive form, relying on (i) application of the law of iterated expectation and (ii) appropriate rearrangement of expected discounted sums. In the current model, the transformation to recursive form must also take into account that the committed policymaker and the private sector have different discount factors and probability beliefs, so that the law of iterated expectation must be applied carefully.

This appendix's derivation of the recursive program in Proposition 1 incorporates three structural features described in section 2 of the text: (1) informational subperiods; (2) different information sets for the committed policymaker and the private sector; and (3) private sector learning. It also generalizes the section 2 framework so that (a) it can be used with constant reputation or a mechanical alternative type; (b) it can be used when the opportunistic type of policymaker is forward-looking with time discount factor β_α. Various elements from the main text are repeated, so that the appendix may be read separately.

The detailed derivation of the recursive form is a slow-moving proof, designed for readers with various degrees of prior exposure to recursive optimal policy design. A key new feature relative to other macro applications is a "change of measure" in the expectations constraint on the committed policymaker, which arises because private agents understand that inflation may come from the decisions of an optimizing alternative type.[1]

As we develop the optimal policy for the committed type, we assume that the committed type takes as given a function governing private agents' expected inflation in the event of its replacement, which may depend on events during its tenure and, in particular, on its terminal reputation. But in the background, there is an equilibrium requirement that private agents form rational beliefs about inflation in the event of a replacement next period. We discuss imposing this requirement at the end of this appendix.

A.1 Intended and actual inflation

At each date, the policymaker chooses intended inflation, denoted as a for the committed type (τ_a) and α for the alternative type (τ_α). Intended inflation is not observed by the

[1] This feature will play an even more important role in future research that makes the alternative type care more about the future than in the current case of a myopic alternative.

private sector. Actual inflation is randomly distributed around this intention, with density $g(\pi|a)$ if there is a committed type and $g(\pi|\alpha)$ if there is an alternative type. We assume

$$
\begin{aligned}
a &= \int \pi g(\pi|a)d\pi \\
\alpha &= \int \pi g(\pi|\alpha)d\pi
\end{aligned}
$$

Implementation errors are $\varepsilon_a = \pi - a$ and $\varepsilon_\alpha = \pi - \alpha$ for the two types. While we allow for different continuous distributions on the same range of inflation outcomes, we do not separately include type τ as an argument to avoid notation clutter in the balance of this appendix (i.e., we write $g(\pi|a)$ and $g(\pi|\alpha)$).

A.2 Measures of history

We use period t as the time index within a regime, so period 0 is the date of last regime change. The committed type begins with a reputation, ρ_0, known to private agents.

Private agents at the end of period t know the entire history of inflation (π), output (x), and inflation shocks (ς) since period 0 (the last regime change date). After the next period starts, the ς shock is realized. The policymaker's intended inflation (a or α) is conditioned on this information, as is the expectations shifter in the output-inflation trade-off, e. We write the information history as

$$h_t = [\varsigma_t, \{\varsigma_{t-s}\}_{s=1}^{t}, \{\pi_{t-s}\}_{s=1}^{t}]$$

After the policymaker chooses his intended inflation, actual inflation and output are realized. Other variables, notably private agents' updated belief about policymaker type, are conditioned on this extended information,

$$h_t^+ = [\pi_t, h_t].$$

Note that

$$h_{t+1} = [\varsigma_{t+1}, h_t^+] = [\varsigma_{t+1}, \pi_t, h_t]$$

A word on notation: In the Public Perfect Bayesian Equilibrium of our dynamic game, variables depend just on the relevant history (e.g., $a(h_t)$) and not separately on the date (e.g., $a_t(h_t)$). To further streamline some formulas, we will sometimes condense variables even further, writing $a(h_t)$ as a_t.

A.3 Beliefs about current inflation

Although private agents do not know the type of policymaker that is in place, at the start of period t, they have a prior belief ρ_t that there is a committed type which will choose a_t and a complementary prior belief $1-\rho_t$ that there is an alternative type which will choose α_t. Accordingly, their rational likelihood of the outcome π_t is

$$g(\pi_t|a_t)\rho_t + g(\pi_t|\alpha_t)(1-\rho_t) \tag{A17}$$

A.4 Beliefs about policymaker type

On observing inflation within a regime, private agents use Bayes' law to update their conditional probability that the current policymaker is the committed type

$$\begin{aligned} \rho(h_t^+) &= \frac{g(\pi_t|a(h_t))\rho(h_t)}{g(\pi_t|a(h_t))\rho(h_t) + g(\pi_t|\alpha(h_t))(1-\rho(h_t))} \\ &\equiv b(\pi_t, a(h_t), \alpha(h_t), \rho(h_t)) \end{aligned} \tag{A18}$$

where the b function is a convenient short-hand and $h_t^+ = [\pi_t, h_t]$. As there is no information about type revealed by ς_{t+1}, $\rho(h_{t+1}) = \rho(h_t^+)$. This updating may be written

$$\rho(h_t^+) = \frac{\rho(h_t)}{\rho(h_t) + \lambda(\pi_t, h_t)(1-\rho(h_t))} \tag{A19}$$

using the likelihood ratio $\lambda(\pi_t, h_t) \equiv \frac{g(\pi_t|\alpha(h_t))}{g(\pi_t|a(h_t))}$.

A.5 Constructing expected inflation

We now construct the private sector's expected inflation, $E\pi_{t+1}$, working backwards from the start of next period to the start of this period. We take into account that there will be a regime change ($\theta_{t+1} = 1$) with probability q and won't ($\theta_{t+1} = 0$) with probability $1-q$.

If the committed type is known to be in place, with decision rule $a([\varsigma_{t+1}, h_t^+])$, then

$$E(\pi_{t+1}|h_{t+1}, \tau_{t+1} = 1) = a([\varsigma_{t+1}, h_t^+])$$

since intended inflation is the mean of realized inflation. Similarly,

$$E(\pi_{t+1}|h_{t+1}, \tau_{t+1} = 0) = \alpha([\varsigma_{t+1}, h_t^+])$$

Since the private sector will not know the type of policymaker in place at the start of next period, expected inflation will be

$$E(\pi_{t+1}|h_{t+1},\theta_{t+1}=0)=\rho(h_{t+1})a(h_{t+1})+(1-\rho(h_{t+1}))\alpha(h_{t+1}) \tag{A20}$$

if there isn't a regime change. Without taking a stand on the details of reputation inheritance, we simply define

$$E(\pi_{t+1}|h_{t+1},\theta_{t+1}=1)=z(h_{t+1}) \tag{A21}$$

as the private sector's expectation of inflation conditional on a replacement.

Stepping back now to period t, expected inflation conditional on h_t is

$$\begin{aligned} E(\pi_{t+1}|h_t)=\rho(h_t)&\int\sum_{\varsigma_{t+1}}\varphi(\varsigma_{t+1};\varsigma_t)\left[(1-q)\,a(h_{t+1})+qz\,(h_{t+1})\right]g(\pi_t|a(h_t))d\pi_t \\ &+(1-\rho\,(h_t))\int\sum_{\varsigma_{t+1}}\varphi(\varsigma_{t+1};\varsigma_t)\left[(1-q)\,\alpha(h_{t+1})+qz\,(h_{t+1})\right]g(\pi_t|\alpha(h_t))d\pi_t \end{aligned} \tag{A22}$$

There may appear to be a conflict between this expression and (A20) that contains reputation at t+1. But there is not. Weighting (A20) and (A21) by $(1-q)$ and q and then integrating over the private sector's belief about inflation (A17) leads directly to it. The simplicity arises because (A17) also occurs in the denominator of the Bayesian updating expression (A18).

A.6 Intertemporal objective

We assume that the policymaker's intertemporal objective involves discounting at $\beta_a(1-q)$, where β_a is its structural discount factor and $(1-q)$ reflects discounting due to replacement.

$$U_t=\underline{u}(a_t,e_t,\varsigma_t)+(\beta_a(1-q))E_t^cU_{t+1}$$

where $\underline{u}\,(a,e,\varsigma)\equiv\int u\,(\pi,x(\pi,e),\varsigma)\,g\,(\pi|a)\,d\pi$ is the expected momentary objective with x replaced by $x(\pi,e)=(\pi-e-\varsigma)\,/\kappa$, and the conditional expectation operator $E_t^c(\cdot)$ is using the committed type's probability $p(h_{t+j})$ of a specific history h_{t+j} when his actions generate inflation.

More specifically, at any date t given the history h_t, the intertemporal objective is

$$U_t=\sum_{j=0}^{\infty}(\beta_a(1-q))^j\sum_{h_{t+j}}\frac{p(h_{t+j})}{p(h_t)}\underline{u}(a(h_{t+j}),e(h_{t+j}),\varsigma(h_{t+j})) \tag{A23}$$

Given $h_{t+j} = [\varsigma_{t+j}, \pi_{t+j-1}, h_{t+j-1}]$, the committed type's probability of a specific history is:

$$p(h_{t+j}) = \varphi(\varsigma_{t+j}; \varsigma_{t+j-1}) \times g(\pi_{t+j-1}|a(h_{t+j-1})) \times p(h_{t+j-1}) \tag{A24}$$

That is, it combines the likelihood of inflation π given the committed type's decision, the likelihood of the shock ς and the probability of the previous history.[2]

A.7 Rational expectations constraint

To develop the desired recursive form, we construct the Lagrangian component using the committed type's probabilities as weights on the multipliers

$$\Psi_t = \sum_{j=0}^{\infty} (\beta_a(1-q))^j \sum_{h_{t+j}} \frac{p(h_{t+j})}{p(h_t)} \gamma(h_{t+j})[e(h_{t+j}) - \beta E(\pi_{t+j+1}|h_{t+j})] \tag{A25}$$

and then express it recursively. We detailed $E(\pi_{t+1}|h_t)$ in (A22), but the expression involved the probability of inflation under the alternative type. So, we undertake a "change of measure" and rewrite it as

$$\begin{aligned} &\rho(h_t) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t)[\beta(1-q)a(h_{t+1}) + \beta q z(h_{t+1})] g(\pi|a(h_t)) d\pi \\ &+(1-\rho(h_t)) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t)[\beta(1-q)\alpha(h_{t+1}) + \beta q z(h_{t+1})] \boldsymbol{\lambda}(\mathbf{h_{t+1}}) g(\pi|a(h_t)) d\pi \end{aligned} \tag{A26}$$

where $\lambda(h_{t+1})$ is the likelihood ratio discussed above in the context of Bayesian updating.

$$\frac{g(\pi_t|\alpha(h_t))}{g(\pi_t|a(h_t))} = \lambda(h_t^+) = \lambda(h_{t+1}) \tag{A27}$$

As the notations emphasize, this is a random variable from the standpoint of h_t but it is known as of $h_t^+ = [\pi_t, h_t]$ and $h_{t+1} = [\varsigma_{t+1}, h_t^+]$.

We now return to (A25) and replace $E(\pi_{t+1}|h_t)$ with the expression in (A26). Note that $a(h_{t+1}), \alpha(h_{t+1})\lambda(h_{t+1})$, and $z(h_{t+1})$ are multiplied by $\varphi(\varsigma_{t+1}; \varsigma_t) g(\pi|a(h_t)) p(h_t)$ and by $\gamma(h_t)$, which is $p(h_{t+1})\gamma(h_t)$. So, just as in simpler models, it is possible to eliminate expectations at future dates, essentially by applying the law of iterated expectation. Adjusting for different

[2] We ask for the reader's patience in using a sum over histories to capture the joint effects of the possibly continuous distribution of π and the discrete Markov chain distribution for ς.

discount factors, we can write (A25) as

$$\Psi_t = E_t^c[\sum_{j=0}^{\infty}(\beta_a(1-q))^j \psi_{t+j}] \tag{A28}$$

with

$$\psi_t = \gamma_t e_t - \frac{\beta}{\beta_a(1-q)}\gamma_{t-1}\{\rho_{t-1}[(1-q)a_t + qz_t] + (1-\rho_{t-1})\lambda_t[(1-q)\alpha_t + qz_t]\} \tag{A29}$$

This latter expression captures past commitments about current state-contingent decisions as these were relevant to past expectations of inflation.[3] Note that at the start of the regime, when $t = 0$, $\gamma_{t-1} = 0$ by assumption. The initial condition on reputation specifies ρ_0.

A.8 The basic recursive specification

The preceding derivations suggest a recursive version of $U_t + \Psi_t$ with states $(\varsigma_t, \gamma_{t-1}, \rho_{t-1}, \lambda_t)$. For algebraic convenience, we define $\eta_t = \frac{\beta}{\beta_a(1-q)}\gamma_{t-1}$. Then, the recursive form as in Marcet and Marimon (2019) is

$$\begin{aligned} W(\varsigma_t, \eta_t, \rho_{t-1}, \lambda_t) = \min_{\gamma}\max_{a,\alpha,e}\{&\underline{u}(a_t, e_t, \varsigma_t) + \gamma_t e_t \\ &- \eta_t[\rho_{t-1}((1-q)a_t + qz_t) + (1-\rho_{t-1})\lambda_t((1-q)\alpha_t + qz_t)] \\ &+ \beta_a(1-q)\int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t) W(\varsigma_{t+1}, \eta_{t+1}, \rho_t, \lambda_{t+1}) g(\pi_t|a_t) d\pi_t\} \end{aligned} \tag{A30}$$

subject to the IC constraint

$$\alpha_t = Ae_t + B(\varsigma_t)$$

with state dynamics (from the perspective of the committed type)

$$\begin{aligned} \eta_{t+1} &= \frac{\beta}{\beta_a(1-q)}\gamma_t \text{ with } \gamma_{-1} = 0 \\ \rho_t &= \frac{\rho_{t-1}}{\rho_{t-1} + (1-\rho_{t-1})\lambda_t} \text{ given } \rho_0 \\ \lambda_{t+1} &= \lambda(\pi_t, a_t, \alpha_t) \text{ with probability } g(\pi_t|a_t) \end{aligned}$$

[3] Our short hand notation replaces $\lambda(h_t)$ with λ_t. Given (A27), the likelihood ratio λ_t is predetermined in period t by actions and inflation outcome in period $t-1$.

A.9 State space reduction

For computational and analytical benefits, it is desirable to reduce the state space. We now show how to eliminate the likelihood ratio (λ) from the state vector so that we only need three state variables instead of four. Start by rewriting (A29) as

$$\text{(A31)} \qquad \psi_t = \gamma_t e_t - \frac{\beta}{\beta_a(1-q)}\gamma_{t-1}\rho_{t-1}\{[(1-q)a_t + qz_t] + \frac{(1-\rho_{t-1})\lambda_t}{\rho_{t-1}}[(1-q)\alpha_t + qz_t]\}$$

Then, note that $\rho_t = \frac{\rho_{t-1}}{\rho_{t-1}+(1-\rho_{t-1})\lambda_t}$ implies that $\frac{(1-\rho_{t-1})\lambda_t}{\rho_{t-1}} = \frac{1-\rho_t}{\rho_t}$ so that Bayes' rule can be used to eliminate λ_t. Substitution of this expression into that above yields

$$\text{(A32)} \qquad \psi_t = \gamma_t e_t - \frac{\beta}{\beta_a(1-q)}\gamma_{t-1}\rho_{t-1}\{[(1-q)a_t + qz_t] + \boxed{\frac{(1-\rho_t)}{\rho_t}}[(1-q)\alpha_t + qz_t]\}$$

which indicates that the states $(\varsigma_t, \eta_t, \rho_{t-1}, \lambda_t)$ can be reduced to ς_t, $\mu_t = \frac{\beta}{\beta_1(1-q)}\gamma_{t-1}\rho_{t-1}$ and ρ_t with the following transition rules for the endogenous states given ρ_0:

$$\text{(A33)} \qquad \mu_{t+1} = \frac{\beta}{\beta_a(1-q)}\gamma_t\rho_t \text{ with } \mu_0 = 0$$

$$\text{(A34)} \qquad \rho_{t+1} = b(\pi_t, a_t, \alpha_t, \rho_t) \text{ with probability } g(\pi_t|a_t)$$

The recursive optimization (A30) can now be written with only three state variables $(\varsigma_t, \rho_t, \mu_t)$ as stated in Proposition 1.

Proposition 1. The within-regime equilibrium is the solution to a recursive optimization problem, given $z(\varsigma, \rho)$ and the IC constraint $\alpha = Ae + B(\varsigma)$

$$\text{(A35)} \qquad W(\varsigma, \rho, \mu) = \min_{\gamma}\max_{a,\alpha,e}\{\underline{u}(a, e, \varsigma) + (\gamma e - \mu\omega) + \beta_a(1-q)\int\sum_{\varsigma'}\varphi(\varsigma';\varsigma)\, W(\varsigma', \rho', \mu')\, g(\pi|a)\, d\pi\},$$

$$\text{(A36)} \qquad \text{with } \omega \equiv (1-q)\,a + qz(\varsigma,\rho) + \frac{1-\rho}{\rho}\left[(1-q)\,\alpha + qz(\varsigma,\rho)\right]$$

$$\text{(A37)} \qquad \mu' = \frac{\beta}{\beta_a(1-q)}\gamma\rho, \text{ given } \mu_0 = 0$$

$$\text{(A38)} \qquad \rho' = \frac{\rho g(\pi|a)}{\rho g(\pi|a) + (1-\rho)\, g(\pi|\alpha)} \text{ with prob } g(\pi|a), \text{ given } \rho_0$$

A.10 A special case

If $q = 0$, $\beta_a = \beta$, and $\rho = 1$ always, our recursive program collapses to a textbook NK policy problem in recursive form. For example, in Clarida et al. (1999), the policymaker maximizes

$E_0 \sum_{t=0}^{\infty} \beta^t u(\pi_t, x_t)$ subject to $\pi_t = \kappa x_t + \beta E_t \pi_{t+1} + \varsigma_t$.

To create a dynamic Lagrangian one attaches $E_0 \sum_{t=0}^{\infty} \beta^t \gamma_t [\pi_t - \kappa x_t - \beta E_t \pi_{t+1} - \varsigma_t]$ to the objective. The law of iterated expectation and rearrangement of terms allow this expression to be written as $E_0 \sum_{t=0}^{\infty} \beta^t \{(\gamma_t - \gamma_{t-1})\pi_t - \gamma_t \kappa x_t - \gamma_t \varsigma_t\}$ with $\gamma_{-1} = 0$. Defining the pseudo state variable $\mu_t = \gamma_{t-1}$, the recursive optimization problem is

$$W(\varsigma_t, \mu_t) = \min_{\gamma_t} \max_{\pi_t, x_t} \{u(\pi_t, x_t) + \gamma_t(\pi_t - \kappa x_t - \varsigma_t) - \mu_t \pi_t + \beta E_t W(\varsigma_{t+1}, \mu_{t+1})\}$$

with $\mu_{t+1} = \gamma_t$ and $\mu_0 = 0$.

B Consolidation

The recursive program in Proposition 1 is valuable, as it sheds light on the relevant state variables. But it contains many choice variables, making it inefficient for computation. This appendix explains how we consolidate choice variables by exploring the implications of private sector's rational expectation constraint.

B.1 Relation between W and U

Taking the first order condition of the recursive optimization problem (A35) with respect to γ and using an envelope theorem result for W_μ, we recover the rational inflation expectations constraint (A22). That is, the optimization imposes the sequence of rational inflation expectations constraints, leading to the following lemma that relates the value function $W(s)$ to the committed policymaker's optimized intertemporal objective $U^*(s)$

Lemma 1. Let $U^*(s)$ and $\omega^*(s)$ be the intertemporal objective (A23) and the composite promise term in (A36) evaluated at optimal decision rules, then

$$W(\varsigma, \rho, \mu) = U^*(\varsigma, \rho, \mu) - \mu \omega^*(\varsigma, \rho, \mu) \tag{B1}$$

Proof. The envelope theorem implication for μ is

$$W_\mu(\varsigma_t, \rho_t, \mu_t) = -\{[(1-q)a_t + qz_t] + \boxed{\frac{(1-\rho_t)}{\rho_t}}[(1-q)\alpha_t + qz_t]\} = -\omega_t$$

The first order necessary condition for γ_t is

$$\begin{aligned} 0 &= e_t + \beta_a(1-q) \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) \int W_\mu(\varsigma_{t+1}, \rho_{t+1}, \mu_{t+1}) \frac{\partial \mu_{t+1}}{\partial \gamma_t} g(\pi_t | a_t) d\pi_t \\ &= e_t + \beta \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) \int W_\mu(\varsigma_{t+1}, \rho_{t+1}, \mu_{t+1}) \rho_t g(\pi_t | a_t) d\pi_t \end{aligned}$$

where the state evolution equation (A33) implies $\partial\mu_{t+1}/\partial\gamma_t = \rho_t\beta/(\beta_a(1-q))$.

When combined with an updated version of the envelope theorem implication, this FOC recovers the private sector's rational expectation constraint as in (A26):

$$e_t = \beta \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t) \left[[(1-q)a_{t+1} + qz_{t+1}] + \boxed{\frac{(1-\rho_{t+1})}{\rho_{t+1}}} [(1-q)\alpha_{t+1} + qz_{t+1}] \right] \rho_t g(\pi_t|a_t) d\pi_t$$

where

$$\frac{1-\rho_{t+1}}{\rho_{t+1}} = \frac{(1-\rho_t)\,\lambda_{t+1}}{\rho_t}.$$

Hence, in equilibrium where the rational expectation constraint must hold, we obtain

$$e_t^* = \beta \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t)\omega_{t+1}^*\rho_t g(\pi_t|a_t^*) d\pi_t$$

Utilizing this equilibrium condition, we now show by "guess and verify" that in equilibrium: $W(\varsigma_t,\rho_t,\mu_t) = U^*(\varsigma_t,\rho_t,\mu_t) - \mu_t\omega_t^*$. The following recursion must hold:

$$\text{(B2)} \qquad \begin{aligned} W(\varsigma_t,\rho_t,\mu_t) + \mu_t\omega_t^* = {} & \underline{u}(a_t^*, e_t^*, \varsigma_t) + \gamma_t e_t^* \\ & + \beta_a(1-q) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t) W(\varsigma_{t+1},\rho_{t+1},\mu_{t+1}) g(\pi_t|a_t^*) d\pi_t \end{aligned}$$

Suppose $W(\varsigma_{t+1},\rho_{t+1},\mu_{t+1}) = -\mu_{t+1}\omega_{t+1}^* + U^*(\varsigma_{t+1},\rho_{t+1},\mu_{t+1})$, the right hand side can be written as

$$\begin{aligned} & \underline{u}(a_t^*, e_t^*, \varsigma_t) + \gamma_t e_t^* - \beta_a(1-q) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t) [\frac{\beta}{\beta_a(1-q)}\gamma_t\rho_t\omega_{t+1}^*] g(\pi_t|a_t^*) d\pi_t \\ & + \beta_a(1-q) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t) U^*(\varsigma_{t+1},\rho_{t+1},\mu_{t+1}) g(\pi_t|a_t^*) d\pi_t \\ = & \underline{u}(a_t^*, e_t^*, \varsigma_t) + \gamma_t [e_t^* - \beta \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t)\omega_{t+1}^*\rho_t g(\pi_t|a_t^*) d\pi_t] \\ & + \beta_a(1-q) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t) U^*(\varsigma_{t+1},\rho_{t+1},\mu_{t+1}) g(\pi_t|a_t^*) d\pi_t \\ = & \underline{u}(a_t^*, e_t^*, \varsigma_t) + \beta_a(1-q) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t) U^*(\varsigma_{t+1},\rho_{t+1},\mu_{t+1}) g(\pi_t|a_t^*) d\pi_t \\ = & U^*(\varsigma_t,\rho_t,\mu_t) \end{aligned}$$

which implies $W(\varsigma_t,\rho_t,\mu_t) = U^*(\varsigma_t,\rho_t,\mu_t) - \mu_t\omega_t^*$. ∎

The value function of the committed policymaker is therefore his optimized intertemporal objective net the cost of delivering on his past promises, captured by the term $\mu\omega^*$.

B.2 The operational expectation function

We now show that imposing the rational expectation constraint (A26) on the choice of e_t implies an operational expectation function:

Lemma 2. Given (ς, ρ), and that future policymakers follow the equilibrium strategies $a^*(\varsigma', \rho', \mu')$, $\alpha^*(\varsigma', \rho', \mu')$ and $z^*(\varsigma', \rho')$, rationally expected inflation is uniquely determined by the contemporaneous policy difference $\delta = a - \alpha$, and the future pseudo-state variable μ'.

(B3)
$$\begin{aligned} e = e(\delta, \mu'; \varsigma, \rho) \;&=\; \beta\rho \int \widehat{M}_1(\varsigma, b(v_\pi, v_\pi + \delta, \rho), \mu') g(v_\pi) dv_\pi + \\ &\quad\; \beta(1-\rho) \int \widehat{M}_2(\varsigma, b(v_\pi - \delta, v_\pi, \rho), \mu') g(v_\pi) dv_\pi; \end{aligned}$$

where $g(v_\pi)$ denotes the density function of v_π;

$$\begin{aligned} \widehat{M}_1(\varsigma, \rho', \mu') \;&:=\; \sum_{\varsigma'} \varphi(\varsigma'; \varsigma) \left[(1-q)\, a^*(\varsigma', \rho', \mu') + q z^*(\varsigma', \rho')\right]; \\ \widehat{M}_2(\varsigma, \rho', \mu') \;&:=\; \sum_{\varsigma'} \varphi(\varsigma'; \varsigma) \left[(1-q)\, \alpha^*(\varsigma', \rho', \mu') + q z^*(\varsigma', \rho')\right]; \end{aligned}$$

Proof. Recall that (A26) comes from (A22) before undertaking a "change of measure". So the original form of the rational expectation constraint on e_t is:

(B4)
$$\begin{aligned} e_t = \beta\rho_t &\int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) \left[(1-q)\, a_{t+1} + q z_{t+1}\right] g(\pi_t | a_t) d\pi_t \\ &+ \beta(1-\rho_t) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) \left[(1-q)\, \alpha_{t+1} + q z_{t+1}\right] g(\pi_t | \alpha_t) d\pi_t \end{aligned}$$

with a_{t+1}, α_{t+1}, and z_{t+1} determined by the three states $(\varsigma_{t+1}, \rho_{t+1}, \mu_{t+1})$ through the equilibrium strategies: $a^*(\cdot)$, $\alpha^*(\cdot)$, and $z^*(\cdot)$.

Recall $\rho_{t+1} = b(\pi_t, a_t, \alpha_t, \rho_t)$ from (A34) and $b(\cdot)$ is the Bayes' learning rule specified in (A18). The inflation distribution is $\pi = a + v_\pi$ under the committed type and $\pi = \alpha + v_\pi$ under the opportunistic type, with v_π being zero mean random variables. We can therefore rewrite the Bayes' learning rule (A18) as

(B5)
$$\begin{aligned} \rho_{t+1} \;&=\; \frac{g(\pi_t - a_t)\rho_t}{g(\pi_t - a_t)\rho_t + g(\pi_t - \alpha_t)(1-\rho_t)} \\ &\equiv\; b(\pi_t - a_t, \pi_t - \alpha_t, \rho_t) \end{aligned}$$

where the b function is a version of our general convenient short-hand which is identified by its three argument nature.

Then, in terms of the policy difference $\delta = a - \alpha$, future reputation is

$$\rho' = b\left(v_\pi, v_\pi + \delta, \rho\right) \text{ conditional on } \tau = 1 \tag{B6}$$

$$\rho' = b\left(v_\pi - \delta, v_\pi, \rho\right) \text{ conditional on } \tau = 0 \tag{B7}$$

Replacing $g(\pi|a)$ and $g(\pi|\alpha)$ in (B4) with $g(v_\pi)$, ρ_{t+1} with (B6) and (B7), and realizing choosing γ_t is equivalent to choosing μ_{t+1} due to $\mu_{t+1} = \frac{\beta}{\beta_a(1-q)}\gamma_t\rho_t$, we obtain the operational expectation function in (B3). ∎

B.3 Simplified recursive program

Using Lemma 1 and 2, we simplify the recursive program (A35), moving from choosing (γ, a, α, e) to merely choosing (δ, μ'):

Proposition 2. Given $z^*(\varsigma, \rho)$ and $U^*\left(\varsigma, \rho, \mu\right)$, a simplified program is

$$W(\varsigma, \rho, \mu) = \max_{\delta, \mu'} \left[\underline{\underline{u}}(\delta, \mu'; \varsigma, \rho) - \mu\underline{\underline{\omega}}(\delta, \mu'; \varsigma, \rho) + \beta_a\left(1-q\right)\Omega(\delta, \mu'; \varsigma, \rho)\right] \tag{B8}$$

with $\Omega(\delta, \mu'; \varsigma, \rho) = \int \sum_{\varsigma'} \varphi(\varsigma'; \varsigma) U^*\left(\varsigma', b(v_\pi, v_\pi + \delta, \rho), \mu'\right) g(v_\pi) dv_\pi$.

Proof. Lemma 1 implies that the objective of the recursive optimization (A35) can be reduced to

$$\underline{u}(a, e, \varsigma) - \mu\omega(a, \alpha) + \beta_a(1-q) \int \sum_{\varsigma'} \varphi(\varsigma'; \varsigma) U^*(\varsigma', \rho', \mu') g(\pi|a) d\pi$$

Lemma 2 implies that (δ, μ') determines $e = e(\delta, \mu'; \varsigma, \rho)$, $\alpha = Ae + B(\varsigma)$, and $a = \alpha + \delta$. Because $\underline{u}(\cdot)$ and $\omega(\cdot)$ are functions of (e, α, a), they can be written as functions of (δ, μ'):

$$\underline{\underline{u}}(\delta, \mu') := \underline{u}(Ae + B(\varsigma), e, \varsigma) \tag{B9}$$

$$\underline{\underline{\omega}}(\delta, \mu') := \frac{1}{\rho}\left[(1-q)\left(Ae + B(\varsigma)\right) + qz^*\left(\varsigma, \rho\right)\right] + (1-q)\,\delta \tag{B10}$$

ρ' in $U^*(\cdot)$ is determined by (B6) and $g(\pi|a) = g(v_\pi)$. We obtain the simplified program. ∎

Lemma 2 and Proposition 2 facilitate our computation. With a guessed function $z\left(\varsigma, \rho\right)$ specified in the outer loop, we can (i) use $a(\varsigma, \rho, \mu), \alpha(\varsigma, \rho, \mu)$ and $U(\varsigma, \rho, \eta)$ functions to obtain $e(\delta, \mu'; \varsigma, \rho)$ and $\Omega(\delta, \mu'; \varsigma, \rho)$; (ii) optimize over (δ, μ'); (iii) construct new a and α functions from optimal e and δ; and (iv) construct a new U function. Within the inner loop, we iterate until the policy functions converge.[4] We then calculate a new $z(\varsigma, \rho)$ and repeat the process until the outer loop has reached a fixed point in z.

[4] Bayesian learning makes this not a linear-quadratic problem. In view of Proposition 2, we use direct maximization as part of a projection method to obtain a global solution. Overall, we employ a variant of the "dynamic programming with a rational expectations constraint" as sometimes advocated for calculating optimal policy under commitment.

C Forecasting Functions and Matching the SPF

C.1 SPF Data

We construct the SPF inflation data from "individual responses" file for the *level* of the GDP deflator available at https://www.philadelphiafed.org/surveys-and-data/pgdp. The sample starts from the fourth quarter of 1968.

In the middle of each quarter, each survey participant submits a forecast for the price level in that quarter and the next four. We first calculate inflation forecasts for each individual forecaster j, using the continuously compounded growth rate: $400 \times \ln(P^j_{t+k|t}/P^j_{t+k-1|t})$. We then take the median of these inflation forecasts.

Alternatively, one can use the summary data files constructed by the Federal Reserve Bank of Philadelphia, particularly the "annualized percent change of median responses" file from https://www.philadelphiafed.org/surveys-and-data/pgdp, as a measure for the SPF inflation data. This file includes an inflation "nowcast" and forecasts at the 1,2,3, and 4 quarter horizons. The nature of these inflation series is explained by Stark (2010). The FRBP first constructs a median price level for each horizon from "individual responses", say $P_{t+k|t}$ for k=0,1,...4. It then constructs an annualized percentage growth rate using the formula $100 \times ([P_{t+k|t}/P_{t+k-1|t}]^4 - 1)$.

Our procedure yields time series that are less prone to transitory outliers than the standard FRBP constructions. Each difference matters, i.e., (i) the median of the inflation rates is less prone than is the change in the median price level; and (ii) the continuously compounded inflation rate is less prone than is the FRBP inflation rate.

Figure 10 contrasts the two measures.

[Figure 10 about here.]

C.2 Recursive forecasting in our theory

This appendix describes the calculation of private agents expectations of inflation at each horizon j: $E(\pi_{t+j}|h_t)$.

The information set is assumed to be the start of period information of the private sector, $s_t = (\varsigma_t, \rho_t, \mu_t)$. We denote the forecast function using $f(s_t, j) = E(\pi_{t+j}|s_t)$.

Given s_t Private agents know the intended inflation policies of the committed and the opportunistic policymakers:

$$a(\varsigma_t, \rho_t, \mu_t)$$
$$\alpha(\varsigma_t, \rho_t, \mu_t)$$

Because implementation errors have mean zero, the private sector "nowcast" of inflation is

$$f(\varsigma_t, \rho_t, \mu_t, 0) = \rho_t a(\varsigma_t, \rho_t, \mu_t) + (1-\rho_t)\alpha(\varsigma_t, \rho_t, \mu_t)$$

Utilizing the law of iterated expectation, today's forecast of π_{t+j} is today's forecast of tomorrow's forecast of π_{t+j}. We can compute multistep forecasts of inflation recursively:

$$\text{(C1)} \qquad E(\pi_{t+j}|s_t) = f(\varsigma_t, \rho_t, \mu_t, j) = E[E(\pi_{t+j}|s_{t+1})|s_t] = E[f(\varsigma_{t+1}, \rho_{t+1}, \mu_{t+1}, j-1)|s_t]$$

The pseudo state variable μ_{t+1} evolves according to:

$$\mu_{t+1} = \begin{cases} \mu'^{*}(\varsigma_t, \rho_t, \mu_t) & \text{with prob } 1-q \\ 0 & \text{with prob } q \end{cases}$$

The reputation state variable ρ_{t+1} evolves according to:

$$\rho_{t+1} = \begin{cases} b(v_\pi, v_\pi + \delta, \rho_t) & \text{with prob } (1-q)\rho_t \\ b(v_\pi - \delta, v_\pi, \rho_t) & \text{with prob } (1-q)(1-\rho_t) \\ \phi_{t+1} b(v_\pi, v_\pi + \delta, \rho_t) + (1-\phi_{t+1})v_{\rho,t+1} & \text{with prob } q\rho_t \\ \phi_{t+1} b(v_\pi - \delta, v_\pi, \rho_t) + (1-\phi_{t+1})v_{\rho,t+1} & \text{with prob } q(1-\rho_t) \end{cases}$$

where $\phi_{t+1} \sim \text{Bernoulli}(\delta_\rho)$ and $v_{\rho,t+1} \sim \text{Beta}(\overline{\rho}, \sigma_\rho)$. Therefore:

$$f(\varsigma_t, \rho_t, \mu_t, j) = \sum \varphi(\varsigma_{t+1}; \varsigma_t)\Big\{ q(1-\delta_\rho) \int f(\varsigma_{t+1}, v_\rho, 0, j-1) d\text{Beta}(v_\rho|\overline{\rho}, \sigma_\rho)$$
$$(1-q)\rho_t \int f\big(\varsigma_{t+1}, b(v_\pi, v_\pi+\delta, \rho_t), \mu'^{*}(\varsigma_t, \rho_t, \mu_t), j-1\big) g(v_\pi) dv_\pi$$
$$+(1-q)(1-\rho_t) \int f\big(\varsigma_{t+1}, b(v_\pi-\delta, v_\pi, \rho_t), \mu'^{*}(\varsigma_t, \rho_t, \mu_t), j-1\big) g(v_\pi) dv_\pi$$
$$+q\rho_t\delta_\rho \int f\big(\varsigma_{t+1}, b(v_\pi, v_\pi+\delta, \rho_t), 0, j-1\big) g(v_\pi) dv_\pi$$
$$+q(1-\rho_t)\delta_\rho \int f\big(\varsigma_{t+1}, b(v_\pi-\delta, v_\pi, \rho_t), 0, j-1\big) g(v_\pi) dv_\pi \Big\}$$

C.3 Matching the SPF: motivation and mechanics

From the standpoint of modern econometrics, our theory is a very simple one that is easily rejected: conditional on the dates of policymaker replacement and the policymaker type within each regime: we have just three random inputs – cost-push shocks ς_t, implementation errors $v_{\pi,t}$, and reputation shocks $v_{\rho,t}$ – that drive many observable macro time series, including the policies a_t and α_t, inflation π_t, and, as we just discussed, expectations at various horizons $E_t(\pi_{t+j})$.

Our work in this paper is quantitative theory and, following early RBC analyses, we fix model parameters and use a transparent strategy for extracting the unobserved states. Then, with the states in hand, we calculate the historical behavior of observables.[5] But the literature has stressed that one of the difficulties with this RBC strategy is that the technology state is measured by the Solow residual, which is based on observable variables (output, capital, and labor) whose behavior is ultimately to be explored.

We therefore develop a strategy for extracting state information that does not use the behavior of the GDP deflator. It relies on the fact that our model provides a mapping between states and private sector inflation expectations at various horizons, the latter of which are measured by the SPF.

The state-space representation of our model can be written as follows

(C2)
$$S_t = [\varsigma_t, \rho_t, \mu_t, \pi_t]' = F(S_{t-1}, v_t | \theta_t, \phi_t, \tau_t)$$
$$= \begin{bmatrix} \delta_\varsigma \varsigma_{t-1} + v_{\varsigma,t} \\ (1-\theta_t+\theta_t\phi_t) b(\varsigma_{t-1}, \rho_{t-1}, \mu_{t-1}, \pi_{t-1}) + \theta_t(1-\phi_t) v_{\rho,t} \\ (1-\theta_t) m(\varsigma_{t-1}, \rho_{t-1}, \mu_{t-1}) \\ \tau_t a(\varsigma_t, \rho_t, \mu_t) + (1-\tau_t)\alpha(\varsigma_t, \rho_t, \mu_t) + v_{\pi,t} \end{bmatrix}$$

(C3)
$$Y_t = \begin{bmatrix} f_{t+1|t} \\ f_{t+3|t} \end{bmatrix} = \begin{bmatrix} f(\varsigma_t, \rho_t, \mu_t, 1) \\ f(\varsigma_t, \rho_t, \mu_t, 3) \end{bmatrix} + \begin{bmatrix} \varepsilon_{1t} \\ \varepsilon_{3t} \end{bmatrix} = H(S_t) + \varepsilon_t$$

The state vector collects the three state variables $(\varsigma_t, \rho_t, \mu_t)$ identified in Proposition 1 and inflation π_t. The state evolution equations are the stochastic processes of shocks and the

[5] Prescott (1986) constructs Solow residuals as productivity indicators and then calculates moment implications for many variables of a model with calibrated parameters. Our work is closer to Plosser (1989), who uses the Solow residual time series and a basic calibrated model to construct time series of many variables, including consumption, investment and so on.

equilibrium policy functions, conditional on $(\theta_t, \phi_t, \tau_t)$, representing the event of policymaker replacement $(\theta_t = 1)$, continuing type in a new regime $(\phi_t = 1)$, and committed type in place $(\tau_t = 1)$.

The observable vector consists of the SPF at one quarter and three quarter horizons ($f_{t+j|t}$, j=1,3). The measurement equations are model-implied one-period and three-period ahead inflation forecasts by private agents. $\varepsilon_{j,t}$ is the normal measurement error with mean zero and standard deviation 0.5% at annualized rate.

We model $(\theta_t, \phi_t, \tau_t)$ as the outcome of an unobserved discrete-state Markov process Θ_t, with six discrete states:[6] $\{(\theta_t = 0, \tau_t = 1), (\theta_t = 0, \tau_t = 0), (\theta_t = 1, \phi_t = 1, \tau_t = 1), (\theta_t = 1, \phi_t = 1, \tau_t = 0), (\theta_t = 1, \phi_t = 0, \tau_t = 1), (\theta_t = 1, \phi_t = 0, \tau_t = 0)\}$. The transitional probability matrix $T_{i,j} = Pr(\Theta_t = j|\Theta_{t-1} = i)$ is determined by the structure of our model: 1) when $\theta_t = 0$, i.e., no replacement of policymaker, the policymaker type remains the same in period $t-1$ and t; 2) when $\theta_t = 1$ and $\phi_t = 1$, i.e., there is a new policymaker whose type is the same as his predecessor, the probability that a committed type will be in place in period t is the private agents' posterior belief at the end of period $t-1$, $b^i_{t-1} := b(s_{t-1}, \pi_{t-1}|\Theta_{t-1} = i)$; 3) when $\theta_t = 1$ and $\phi_t = 0$, i.e., there is a new policymaker whose type is a random draw, the probability that a committed type will be in place in period t is the unconditional mean $\overline{\rho}$ of the reputation shock v_ρ.

$$\text{(C4)} \qquad T = \begin{bmatrix} 1-q & 0 & \delta_\rho b^{i=1}_{t-1} q & \delta_\rho(1-b^{i=1}_{t-1})q & (1-\delta_\rho)\overline{\rho}q & (1-\delta_\rho)(1-\overline{\rho})q \\ 0 & (1-q) & \delta_\rho b^{i=2}_{t-1} q & \delta_\rho(1-b^{i=2}_{t-1})q & (1-\delta_\rho)\overline{\rho}q & (1-\delta_\rho)(1-\overline{\rho})q \\ 1-q & 0 & \delta_\rho b^{i=3}_{t-1} q & \delta_\rho(1-b^{i=3}_{t-1})q & (1-\delta_\rho)\overline{\rho}q & (1-\delta_\rho)(1-\overline{\rho})q \\ 0 & (1-q) & \delta_\rho b^{i=4}_{t-1} q & \delta_\rho(1-b^{i=4}_{t-1})q & (1-\delta_\rho)\overline{\rho}q & (1-\delta_\rho)(1-\overline{\rho})q \\ 1-q & 0 & \delta_\rho b^{i=5}_{t-1} q & \delta_\rho(1-b^{i=5}_{t-1})q & (1-\delta_\rho)\overline{\rho}q & (1-\delta_\rho)(1-\overline{\rho})q \\ 0 & (1-q) & \delta_\rho b^{i=6}_{t-1} q & \delta_\rho(1-b^{i=6}_{t-1})q & (1-\delta_\rho)\overline{\rho}q & (1-\delta_\rho)(1-\overline{\rho})q \end{bmatrix}$$

[Figure 11 about here.]

C.4 Unscented Kalman filter with Markov-switching

This subsection describes the detailed algorithm we employ to obtain filtered and smoothed estimates of latent states in the state space model (C2) and (C3). Relative to a standard nonlinear system with additive Gaussian errors, our model has three complications.

[6] In general, there will be eight discrete states constructed from combinations of three binary variables. In this case, the state ϕ_t is only relevant in a new regime, i.e., $\theta_t = 1$.

First, the shocks v_ς and v_ρ enter the evolution equation of π nonlinearly because the policy function $a(\cdot)$ and $\alpha(\cdot)$ are nonlinear functions of ς and ρ. Moreover, the shock v_ρ follows a Beta distribution instead of a Gaussian one. Following Särkkä and Svensson (2023), this complication can be dealt with by: 1) approximating the Beta random variable v_ρ using a nonlinear transformation of a Gaussian random variable $\tilde{v}_\rho$:

$$v_\rho = R(\tilde{v}_\rho) = \frac{exp(\tilde{v}_\rho)}{1 + exp(\tilde{v}_\rho)}$$

2) forming sigma points for the state vector augmented by v_ς and $\tilde{v}_\rho$.

Second, the reputation state ρ is bounded between 0 and 1. To enforce the boundary condition, we use "constrained unscented Kalman filter" (Kandepu et al. (2008), Rouhani and Abur (2018)) that projects the sigma points outside the feasible region to the nearest points within the region.

Third, the state evolution equations depend on the outcome of an unobserved discrete-state Markov process Θ_t. We follow Kim (1994) and Kim and Nelson (2017) to obtain the conditional probability of Θ_t being in each discrete state and to collapse state estimate and covariance.

To ease the notation, we rewrite the state space model (C2) and (C3) as follows:

$$S_t = F_{\Theta_t}\big(S_{t-1}, [v_{\varsigma,t}, \tilde{v}_{\rho,t}]\big) + [0, 0, 0, v_{\pi,t}]'$$
$$Y_t = H(S_t) + \varepsilon_t$$

where $\Theta_t \in \{1, ..., 6\}$ corresponding to $\{(\theta_t = 0, \tau_t = 1), (\theta_t = 0, \tau_t = 0), (\theta_t = 1, \phi_t = 1, \tau_t = 1), (\theta_t = 1, \phi_t = 1, \tau_t = 0), (\theta_t = 1, \phi_t = 0, \tau_t = 1), (\theta_t = 1, \phi_t = 0, \tau_t = 0)\}$ with transitional probability matrix $T_{i,j} = Pr(\Theta_t = j|\Theta_{t-1} = i)$ defined in (C4).

Notation:

- Covariance of $[0, 0, 0, v_\pi]'$: Q
- Covariance of measurement noise ε: R
- Mean of the shock vector $[v_{\varsigma,t}, \tilde{v}_{\rho,t}]'$: $\hat{v} = [0, \tilde{\rho}]'$
- Covariance of the shock vector $[v_{\varsigma,t}, \tilde{v}_{\rho,t}]'$: $V = diag(\sigma_\varsigma^2, \tilde{\sigma}_\rho^2)$
- Initial state estimate: $\hat{s}_0^j$, $j = 1, 2, ..., 6$
- Initial state covariance: P_0^j, $j = 1, 2, ..., 6$

Parameters related to sigma points

- Number parameter: L
- Scaling parameters: α, β, κ
- Weight parameter: $\lambda = \alpha^2(L + \kappa) - L$
- $w_{m,0} = \frac{\lambda}{L+\lambda}$, $w_{m(n)} = \frac{1}{2(L+\lambda)}$, $n = 1, ..., 2L$
- $w_{c,0} = \frac{\lambda}{L+\lambda} + (1 - \alpha^2 + \beta)$, $w_{c(n)} = \frac{1}{2(L+\lambda)}$, $n = 1, ..., 2L$

Prediction Step: $L = 6$, conditional on $\Theta_{t-1} = i$, $\Theta_t = j$:

- Augment the state vector:

$$\hat{x}^i_{t-1} = \begin{bmatrix} \hat{s}^i_{t-1} \\ \hat{v} \end{bmatrix}; \quad \tilde{P}^i_{t-1} = \begin{bmatrix} P^i_{t-1} & 0 \\ 0 & V \end{bmatrix}$$

- Generate $2L + 1$ sigma points:
  - $X^i_{t-1,(0)} = \hat{x}^i_{t-1}$
  - $X^i_{t-1,(n)} = \hat{x}^i_{t-1} + \sqrt{(L + \lambda)}\left[\sqrt{\tilde{P}^i_{t-1}}\right]_n$
  - $X^i_{t-1,(n+L)} = \hat{x}^i_{t-1} - \sqrt{(L + \lambda)}\left[\sqrt{\tilde{P}^i_{t-1}}\right]_n, \quad n = 1, ..., L$
- Propagate sigma points through the state transition function:
  - $S'^{(i,j)}_{t(n)} = F_j(X^i_{t-1,(n)}), \quad n = 0, ..., 2L$
- Compute the predicted state estimate:
  - $\hat{s}^{-(i,j)}_t = \sum_{n=0}^{2L} w_{m(n)} S'^{(i,j)}_{t(n)}$
- Compute the predicted state covariance:
  - $P^{-(i,j)}_t = \sum_{n=0}^{2L} w_{c(n)}(S'^{(i,j)}_{t(n)} - \hat{s}^{-(i,j)}_t)(S'^{(i,j)}_{t(n)} - \hat{s}^{-(i,j)}_t)^\top + Q$

Update Step: $L = 4$, conditional on $\Theta_{t-1} = i$, $\Theta_t = j$:

- Generate sigma points:
  - $S_{t,(0)}^{-(i,j)} = \hat{s}_t^{-(i,j)}$
  - $S_{t,(n)}^{-(i,j)} = \hat{s}_t^{-(i,j)} + \sqrt{(L+\lambda)}\big[\sqrt{P_t^{-(i,j)}}\big]_n$
  - $S_{t,(n+L)}^{-(i,j)} = \hat{s}_t^{-(i,j)} - \sqrt{(L+\lambda)}\big[\sqrt{P_t^{-(i,j)}}\big]_n, \quad n = 1, ..., L$
- Propagate sigma points through the measurement function:
  - $Y_{t(n)}^{-(i,j)} = H(S_{t(n)}^{-(i,j)}), \quad n = 0, ..., 2L$
- Compute the predicted measurement mean and covariance:
  - $\hat{y}_t^{-(i,j)} = \sum_{n=0}^{2L} w_{m(n)} Y_{t(n)}^{-(i,j)}$
  - $P_{yy,t}^{-(i,j)} = \sum_{n=0}^{2L} w_{c(n)} (Y_{t(n)}^{-(i,j)} - \hat{y}_t^{-(i,j)})(Y_{t(n)}^{-(i,j)} - \hat{y}_t^{-(i,j)})^\top + R$
- Compute the cross-covariance between state and measurement:
  - $P_{sy,t}^{-(i,j)} = \sum_{n=0}^{2L} w_{c(n)} (S_{t(n)}^{-(i,j)} - \hat{s}_t^{-(i,j)})(Y_{t(n)}^{(i,j)} - \hat{y}_t^{-(i,j)})^\top$
- Compute the Kalman gain:
  - $K_t^{(i,j)} = P_{sy,t}^{-(i,j)}(P_{yy,t}^{-(i,j)})^{-1}$
- Update the state estimate:
  - $\hat{s}_t^{(i,j)} = \hat{s}_t^{-(i,j)} + K_t^{(i,j)}(Y_t - \hat{y}_t^{-(i,j)})$
- Update the state covariance:
  - $P_t^{(i,j)} = P_t^{-(i,j)} - K_t^{(i,j)} P_{yy,t}^{-(i,j)} (K_t^{(i,j)})^\top$

Conditional Probability Step:

- Start from $Pr(\Theta_{t-1} = i | Y^{t-1})$
  - $Pr(\Theta_{t-1} = i, \Theta_t = j | Y^{t-1}) = Pr(\Theta_t = j | \Theta_{t-1} = i) Pr(\Theta_{t-1} = i | Y^{t-1})$

- Update using Bayes' rule

$$Pr(\Theta_{t-1} = i, \Theta_t = j|Y^t) = \frac{f(Y_t|\Theta_{t-1} = i, \Theta_t = j, Y^{t-1})Pr(\Theta_{t-1} = i, \Theta_t = j|Y^{t-1})}{\sum_{j=1}^{6}\sum_{i=1}^{6} f(Y_t, \Theta_{t-1} = i, \Theta_t = j|Y^{t-1})}$$

where $f(Y_t|\Theta_{t-1} = i, \Theta_t = j, Y^{t-1}) \sim N(\hat{y}_t^{-(i,j)}, P_{yy,t}^{-(i,j)})$

- Collapse $Pr(\Theta_t = j|Y^t) = \sum_{i=1}^{6} Pr(\Theta_{t-1} = i, \Theta_t = j|Y^t)$

Collapse Step:

$$\hat{s}_t^j = \frac{\sum_{i=1}^{6} Pr(\Theta_{t-1} = i, \Theta_t = j|Y^t)\hat{s}_t^{(i,j)}}{Pr(\Theta_t = j|Y^t)}$$

$$P_t^j = \frac{\sum_{i=1}^{6} Pr(\Theta_{t-1} = i, \Theta_t = j|Y^t)\{P_t^{(i,j)} + (\hat{s}_t^j - \hat{s}_t^{(i,j)})(\hat{s}_t^j - \hat{s}_t^{(i,j)})^\top\}}{Pr(\Theta_t = j|Y^t)}$$

Smooth Step:

- Initialize the smoothed state estimate and covariance at the last time step:
  - $\hat{s}_T^{s,j} = \hat{s}_T^j$
  - $P_T^{s,j} = P_T^j$
  - $Pr(\Theta_T = j|Y^T)$
- Smooth probability for $\Theta_t = j$ and $\Theta_{t+1} = k$ from $t = T-1, ..., 1$:

$$\begin{aligned}
&Pr(\Theta_t = j, \Theta_{t+1} = k|Y^T) \\
&\quad = Pr(\Theta_{t+1} = k|Y^T)Pr(\Theta_t = j|\Theta_{t+1} = k, Y^T) \\
&\quad \approx Pr(\Theta_{t+1} = k|Y^T)Pr(\Theta_t = j|\Theta_{t+1} = k, Y^t) \\
&\quad = \frac{Pr(\Theta_{t+1} = k|Y^T)Pr(\Theta_t = j, \Theta_{t+1} = k|Y^t)}{Pr(\Theta_{t+1} = k|Y^t)} \\
&\quad = Pr(\Theta_{t+1} = k|Y^T)\frac{Pr(\Theta_t = j|Y^t)Pr(\Theta_{t+1} = k|\Theta_t = j)}{\sum_{j=1}^{6} Pr(\Theta_t = j|Y^t)Pr(\Theta_{t+1} = k|\Theta_t = j)}
\end{aligned}$$

- Smooth probability for $\Theta_t = j$ for $t = T-1, ..., 1$:

$$Pr(\Theta_t = j|Y^T) = \sum_{k=1}^{6} Pr(\Theta_t = j, \Theta_{t+1} = k|Y^T)$$

- Perform the smoothing recursion from $t = T-1, ..., 1$, conditional on $\Theta_t = j, \Theta_{t+1} = k$:

– Augment the state vector:

$$\hat{x}_t^j = \begin{bmatrix} \hat{s}_t^j \\ \hat{v} \end{bmatrix}; \quad \tilde{P}_t^j = \begin{bmatrix} P_t^j & 0 \\ 0 & V \end{bmatrix}$$

– Generate $2L+1$ sigma points given $L = 6$:

  * $X_{t,(0)}^j = \hat{x}_t^j$
  * $X_{t,(n)}^j = \hat{x}_t^j + \sqrt{(L+\lambda)}\big[\sqrt{\tilde{P}_t^j}\big]_n$
  * $X_{t,(n+L)}^j = \hat{x}_t^j - \sqrt{(L+\lambda)}\big[\sqrt{\tilde{P}_t^j}\big]_n, \quad n = 1, ..., L$

– Propagate sigma points through the state transition function:

  * $S_{t+1,(n)}^{\prime(j,k)} = F_k(X_{t,(n)}^j)$

– Compute the predicted state mean and covariance:

  * $\hat{s}_{t+1}^{-(j,k)} = \sum_{n=0}^{2L} w_{m(n)} S_{t+1,(n)}^{\prime(j,k)}$
  * $P_{t+1}^{-(j,k)} = \sum_{n=0}^{2L} w_{c(n)} (S_{t+1,(n)}^{\prime(j,k)} - \hat{s}_{t+1}^{-(j,k)})(S_{t+1,(n)}^{\prime(j,k)} - \hat{s}_{t+1}^{-(j,k)})^\top + Q$

– Compute the cross-covariance:

  * $D_{t+1}^{-(j,k)} = \sum_{n=0}^{2L} w_{c(n)} (X_{t,(n)}^{j,S} - \hat{s}_t^j)(S_{t+1,(n)}^{\prime(j,k)} - \hat{s}_{t+1}^{-(j,k)})^\top$
  * where $X_{t,(n)}^{j,S}$ denotes the part of sigma point n which corresponds to S_t

– Compute the smoothed state gain:

  * $K_t^{s,(j,k)} = D_{t+1}^{-(j,k)} (P_{t+1}^{-(j,k)})^{-1}$

– Compute the smoothed state estimate:

  * $\hat{s}_t^{s,(j,k)} = \hat{x}_t^j + K_t^{s,(j,k)} (\hat{s}_{t+1}^{s,k} - \hat{s}_{t+1}^{-(j,k)})$

– Compute the smoothed state covariance:

  * $P_t^{s,(j,k)} = P_t^j + K_t^{s,(j,k)} (P_{t+1}^{s,k} - P_{t+1}^{-(j,k)})(K_t^{s,(j,k)})^\top$

– Collapse the smoothed state estimate and covariance

$$\hat{s}_t^{s,j} = \frac{\sum_{k=1}^{6} Pr(\Theta_t = j, \Theta_{t+1} = k | Y^T) \hat{s}_t^{s,(j,k)}}{Pr(\Theta_t = j | Y^T)}$$

$$P_t^{s,j} = \frac{\sum_{k=1}^{6} Pr(\Theta_t = j, \Theta_{t+1} = k | Y^T) \{P_t^{s,(j,k)} + (\hat{s}_t^{s,j} - \hat{s}_t^{s,(j,k)})(\hat{s}_t^{s,j} - \hat{s}_t^{s,(j,k)})^\top\}}{Pr(\Theta_t = j | Y^T)}$$

Filtered and smoothed estimates of states and observables:

- Filtered estimates for states conditional on $\Theta_{t-1} = i, \Theta_t = j$

$$\hat{s}_t = \sum_{j=1}^{6} Pr(\Theta_t = j | Y^t) \hat{s}_t^j$$

$$P_t = \sum_{j=1}^{6} Pr(\Theta_t = j | Y^t) \{ P_t^j + (\hat{s}_t - \hat{s}_t^j)(\hat{s}_t - \hat{s}_t^j)^\top \}$$

- Filtered estimates for observables
  - Generate sigma points, $L = 4$:
    * $S_{t,(0)}^{(i,j)} = \hat{s}_t^{(i,j)}$
    * $S_{t,(n)}^{(i,j)} = \hat{s}_t^{(i,j)} + \sqrt{(L+\lambda)} \big[\sqrt{P_t^{(i,j)}} \big]_n$
    * $S_{t,(n+L)}^{(i,j)} = \hat{s}_t^{(i,j)} - \sqrt{(L+\lambda)} \big[\sqrt{P_t^{(i,j)}} \big]_n, \quad n = 1, ..., L$
  - Propagate sigma points through the measurement function:
    * $Y_{t(n)}^{(i,j)} = H(S_{t(n)}^{(i,j)}), \quad n = 0, ..., 2L$
  - Compute the predicted measurement mean and covariance:
    * $\hat{y}_t^{(i,j)} = \sum_{n=0}^{2L} w_{m(n)} Y_{t(n)}^{(i,j)}$
    * $P_{yy,t}^{(i,j)} = \sum_{n=0}^{2L} w_{c(n)} (Y_{t(n)}^{(i,j)} - \hat{y}_t^{(i,j)})(Y_{t(n)}^{(i,j)} - \hat{y}_t^{(i,j)})^\top + R$
  - Collapse

$$\hat{y}_t = \sum_{i=1}^{6} \sum_{j=1}^{6} Pr(\Theta_{t-1} = i, \Theta_t = j | Y^t) \hat{y}_t^{(i,j)}$$

$$P_{yy,t} = \sum_{i=1}^{6} \sum_{j=1}^{6} Pr(\Theta_{t-1} = i, \Theta_t = j | Y^t) \{ P_{yy,t}^{(i,j)} + (\hat{y}_t - \hat{y}_t^{(i,j)})(\hat{y}_t - \hat{y}_t^{(i,j)})^\top \}$$

- Smoothed estimates for states

$$\hat{s}_t^s = \sum_{j=1}^{6} Pr(\Theta_t = j | Y^T) \hat{s}_t^{s,j}$$

$$P_t^s = \sum_{j=1}^{6} Pr(\Theta_t = j | Y^T) \{ P_t^{s,j} + (\hat{s}_t^s - \hat{s}_t^{s,j})(\hat{s}_t^s - \hat{s}_t^{s,j})^\top \}$$

- Smoothed estimates for observables conditional on $\Theta_t = j, \Theta_{t+1} = k$
  - Generate sigma points, $L = 4$:
    - $S^{s,(j,k)}_{t,(0)} = \hat{s}^{s,(j,k)}_t$
    - $S^{s,(j,k)}_{t,(n)} = \hat{s}^{s,(j,k)}_t + \sqrt{(L+\lambda)}\big[\sqrt{P^{s,(j,k)}_t}\big]_n$
    - $S^{s,(j,k)}_{t,(n+L)} = \hat{s}^{s,(j,k)}_t - \sqrt{(L+\lambda)}\big[\sqrt{P^{s,(j,k)}_t}\big]_n, \quad n = 1, ..., L$
  - Propagate sigma points through the measurement function:
    - $Y^{(i,j)}_{t(n)} = H(S^{s,(j,k)}_{t(n)}), \quad n = 0, ..., 2L$
  - Compute the predicted measurement mean and covariance:
    - $\hat{y}^{s,(j,k)}_t = \sum_{n=0}^{2L} w_{m(n)} Y^{s,(j,k)}_{t(n)}$
    - $P^{s,(j,k)}_{yy,t} = \sum_{n=0}^{2L} w_{c(n)} (Y^{s,(j,k)}_{t(n)} - \hat{y}^{s,(j,k)}_t)(Y^{s,(j,k)}_{t(n)} - \hat{y}^{s,(j,k)}_t)^\top + R$
  - Collapse

$$\hat{y}^s_t = \sum_{j=1}^{6}\sum_{k=1}^{6} Pr(\Theta_t = j, \Theta_{t+1} = k|Y^T)\hat{y}^{s,(j,k)}_t$$

$$P^s_{yy,t} = \sum_{j=1}^{6}\sum_{k=1}^{6} Pr(\Theta_t = j, \Theta_{t+1} = k|Y^T)\{P^{s,(j,k)}_{yy,t} + (\hat{y}^s_t - \hat{y}^{s,(j,k)}_t)(\hat{y}^s_t - \hat{y}^{s,(j,k)}_t)^\top\}$$

C.5 Fitting performance

Inflation expectations: As discussed in Section 5.3 of the main text, we extract latent states by matching model-implied inflation forecasts at horizons 1 and 3 with SPF one-quarter-ahead and three-quarter-ahead forecasts. The left panels in Figure 11 shows our match is nearly perfect given that we choose two state variables ς and ρ each period to match two data points SPF1Q and SPF3Q. Using the extracted states, we can also compute model-implied inflation forecasts at horizons 2 and 4, and compare them with SPF two-quarter-ahead and four-quarter-ahead forecasts. The comparison is shown in the right panels of Figure 11. It is notable that our model-implied forecasts lie almost entirely on top of the SPF data for both forecasting horizons, which are not explicitly targeted. We view this figure as evidence in support of our state extraction approach.

Cost-push shocks: We have explored two indicators of the cost-push shock. One is a food and and energy shock constructed along the lines of Watson (2014). The other is the estimated shock series from matching the SPFs. Figure 12 displays these two series. Note first that both measures rise during the famous "oil price shock" of late 1973 and early 1974 and also during the late 1970s interval that preceded Volcker's appointment. Note next that the estimated shocks are more persistent. Contemporary sources, such as the January 1975 Economic Report of the President prepared by Alan Greenspan and his CEA colleagues, point to other price shocks in addition to oil during the preceding year. Econometric studies such as those of Gordon (2013) and Watson (2014) estimate price shocks, including those from price decontrols in the 1970s, of more lasting form. Our estimated shock series echo their findings.

[Figure 12 about here.]

Filtered inflation Section 5.5.2 demonstrates that the smoothed estimates of inflation by our state-space model fit the observed U.S. inflation well without explicitly targeting it. The benchmark we use to measure the fitting performance is to compare the smoothed estimates with the SPF one-quarter-ahead forecast, as shown in Figure 5. A skeptical reader may concern that our smoothed measure performs better simply because it is based on the full sample of SPF, while the SPF1Q is prepared with information up to the period t. We therefore provide a filtered version Figure 13, where no information after the period t is used to obtain the period-t filtered measure. Our filtered estimates for inflation continue to outperform SPF1Q in both measures of fit: lower persistence of fitting error and lower mean-squared error.

[Figure 13 about here.]

D Counterfactual with Naive Committed Policy

D.1 Optimization of a naive committed policymaker

The key departure from the benchmark model is that the committed type optimizes as if the reputation is a given parameter ρ. When the reputation is no longer a function of the inflation shock π (at least in the committed type's optimization), there is no channel for the current π_t to affect future state variables.[7]

This observation helps us to reduce the forwarding expectation constraint to:

$$e_t = \beta E_t \pi_{t+1} = \beta (1-q) \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) \left[\rho a(h_{t+1}) + (1-\rho)\alpha(h_{t+1})\right] + \beta q \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) z(h_{t+1})$$

[7] Recall that the lagrangian multiplier γ_t is chosen before the realization of π_t and it will determine the next-period pseudo state variable.

because a_{t+1}, α_{t+1}, and z_{t+1} are independent of π_t. As a result, we avoid carrying the likelihood ratio $\lambda\left(h_{t+1}\right) := \frac{g(\pi_t|\alpha_t)}{g(\pi_t|a_t)}$ as a state variable.

The recursive form of the naive optimization of the committed policymaker is

$$\begin{aligned} W\left(\varsigma_t, \eta_t; \rho\right) &= \min_{\gamma} \max_{a,e} \underline{u}(a_t, e_t, \varsigma_t) + \gamma_t e_t - (1-q)\,\eta_t\left[\rho a_t + (1-\rho)\,\alpha_t\right] - q\eta_t z\left(\varsigma_t, \rho\right) \\ &\quad + \beta_a (1-q) \sum_{\varsigma_{t+1}} \varphi\left(\varsigma_{t+1}; \varsigma_t\right) W\left(\varsigma_{t+1}, \eta_{t+1}; \rho\right) \end{aligned}$$

subject to

$$\alpha_t = Ae_t + B\left(\varsigma_t\right)$$

with

$$\eta_{t+1} = \frac{\beta}{\beta_a (1-q)} \gamma_t \text{ with } \gamma_{-1} = 0.$$

Given $z\left(\varsigma_t, \rho\right)$, the optimization yields the following policy rules: $a\left(\varsigma_t, \eta_t; \rho\right)$, $e\left(\varsigma_t, \eta_t; \rho\right)$, and $\gamma\left(\varsigma_t, \eta_t; \rho\right)$. The fixed point requires

$$z\left(\varsigma_t, \rho\right) = \rho a\left(\varsigma_t, 0; \rho\right) + (1-\rho)\left[Ae\left(\varsigma_t, 0; \rho\right) + B\left(\varsigma_t\right)\right]$$

The policy function under the setup of naive committed policymaker are denoted by

$$a^N(\varsigma, \rho, \mu)$$
$$\alpha^N(\varsigma, \rho, \mu)$$
$$\mu'^N(\varsigma, \rho, \mu)$$

D.2 Constructing counterfactual time series

Initialization step for $t = 1$: $\rho_1^{N,j} = \hat{\rho}_1^j$ and $\mu_1^{N,j} = \hat{\mu}_1^j$ for $\Theta_1 = j$. $\{\hat{\varsigma}_t^j\}_{t=1}^T$ and $\{Pr(\Theta_t = j|Y^T)\}_{t=1}^T$ are smoothed estimates of the cost-push shocks and smoothed probabilities of $\Theta_t = j$ obtained from the benchmark model.

Conditional on $\Theta_t = j$ and $\Theta_{t+1} = k$, we obtain

$$
\begin{aligned}
a_t^{N,j} &= a^N(\hat{\varsigma}_t^j, \rho_t^{N,j}, \mu_t^{N,j}) \\
\alpha_t^{N,j} &= \alpha^N(\hat{\varsigma}_t^j, \rho_t^{N,j}, \mu_t^{N,j})) \\
\rho_{t+1}^{N,(j,k)} &= \begin{cases} b(\pi_t^j; a_t^{N,j}, \alpha_t^{N,j}, \rho_t^{N,j}) & \text{if } k = 1,2,3,4 \\ \hat{\rho}_{t+1}^k & \text{if } k = 5,6 \end{cases} \\
\mu_{t+1}^{N,(j,k)} &= \begin{cases} \mu'^N(\hat{\varsigma}_t^j, \rho_t^{N,j}, \mu_t^{N,j}) & \text{if } k = 1,2 \\ 0 & \text{if } k = 3,4,5,6 \end{cases}
\end{aligned}
$$

where $\pi_t^j = a_t^{N,j}$ if $j = 1,3,5$ and $\pi_t^j = \alpha_t^{N,j}$ if $j = 2,4,6$. Notice that we shut down the implementation errors, $v_{\pi,t} = 0$, to focus on the effect of past policies on reputation evolution.

We then perform the collapsing step:

$$
\begin{aligned}
\rho_{t+1}^{N,k} &= \sum_{j=1}^{6} \rho_{t+1}^{N,(j,k)} Pr(\Theta_t = j|\Theta_{t+1} = k, Y^T) \\
\mu_{t+1}^{N,k} &= \sum_{j=1}^{6} \mu_{t+1}^{N,(j,k)} Pr(\Theta_t = j|\Theta_{t+1} = k, Y^T)
\end{aligned}
$$

where

$$
\begin{aligned}
Pr(\Theta_t = j|\Theta_{t+1} = k, Y^T) &= \frac{Pr(\Theta_t = j, \Theta_{t+1} = k|Y^T)}{\sum_{j=1}^{6} Pr(\Theta_t = j, \Theta_{t+1} = k|Y^T)} \\
&= \frac{Pr(\Theta_{t+1} = k|\Theta_t = j)Pr(\Theta_t = j|Y^T)}{\sum_{j=1}^{6} Pr(\Theta_{t+1} = k|\Theta_t = j)Pr(\Theta_t = j|Y^T)}
\end{aligned}
$$

The transitional probability $Pr(\Theta_{t+1} = k|\Theta_t = j)$ are the same as the one in the benchmark model (C4) except that b_{t-1}^i is replaced with the naive-policy version $b(\pi_t^j; a_t^{N,j}, \alpha_t^{N,j}, \rho_t^{N,j})$.

The reported counterfactual time series t=1,...T are constructed as follows:

$$
\begin{aligned}
\rho_t^N &= \sum_{j=1}^{6} \rho_t^{N,j} Pr(\Theta_t = j|Y^T) \\
a_t^N &= \sum_{j=1}^{6} a_t^{N,j} Pr(\Theta_t = j|Y^T) \\
\alpha_t^N &= \sum_{j=1}^{6} \alpha_t^{N,j} Pr(\Theta_t = j|Y^T)
\end{aligned}
$$

E Model performance with a longer sample

This appendix reports the performance of our quantitative model using the SPF sample from 1968Q4 to 2023Q1.

Figure 14 plots the smoothed estimates for inflation forecasts against the SPF data. Again, the filtering exercise only uses the information of SPF1Q and SPF3Q, but the model-implied inflation expectations also fit the untargeted SPF2Q and SPF4Q very well.

Figure 15 demonstrates that our quantitative model performs well in matching the U.S. inflation data through 2023Q1. Compared with the one-quarter-head inflation forecast (SPF1Q), our model estimates for inflation yield lower mean-squared-error and error-persistence. This result holds both for the smoothed estimates (using the full sample information) and for the filtered estimates (using the sample up to t).

[Figure 14 about here.]

[Figure 15 about here.]

Appendix References

Gordon, Robert, "The Phillips Curve is Alive and Well: Inflation and the NAIRU During the Slow Recovery," Technical Report aug 2013.

Kandepu, Rambabu, Lars Imsland, and Bjarne A. Foss, "Constrained state estimation using the Unscented Kalman Filter," in "2008 16th Mediterranean Conference on Control and Automation" June 2008, pp. 1453–1458.

Kim, Chang-Jin, "Dynamic linear models with Markov-switching," *Journal of Econometrics*, January 1994, *60* (1-2), 1–22.

_ **and Charles R. Nelson**, *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, The MIT Press, November 2017.

Marcet, Albert and Ramon Marimon, "Recursive Contracts," *Econometrica*, 2019, *87* (5), 1589–1631.

Rouhani, Alireza and Ali Abur, "Constrained Iterated Unscented Kalman Filter for Dynamic State and Parameter Estimation," *IEEE Transactions on Power Systems*, May 2018, *33* (3), 2404–2414.

Särkkä, Simo and Lennart Svensson, *Bayesian Filtering and Smoothing*, 2023 ed., Cambridge University Press, 2023.

Watson, Mark W., "Inflation Persistence, the NAIRU, and the Great Recession," *American Economic Review*, May 2014, *104* (5), 31–36.

Table 1: Parameters

π^*	Inflation target	1.5%
β, β_a	Discount factor (private, committed type)	0.995
κ	PC output slope	0.08
ϑ_x	Output weight	0.1
x^*	Output target	1.73%
q	Replacement probability	0.03
δ_ρ	prob of reputation inheritance	0.9
$\overline{\rho}$	mean of reputation draw	0.1
σ_ρ	std of reputation draw	0.05
δ_ς	Persistence of cost-push shock	0.7
$\sigma_{v,\varsigma}$	Std of cost-push innovation	0.7%
$\sigma_{v,\pi}$	Std of implementation error v_π	1.2%

One period is a quarter. Inflation target π^*, std of cost-push innovation $\sigma_{v,\varsigma}$, and std of implementation error $\sigma_{v,\pi}$ are all annualized rates.

Figure 1: Timing of events within a period

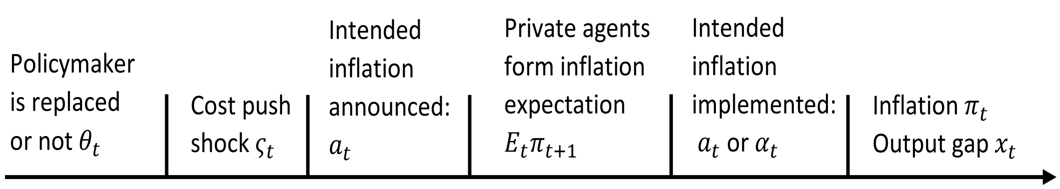

Figure 2: Optimal Response of Opportunistic Policy to Inflation Expectations

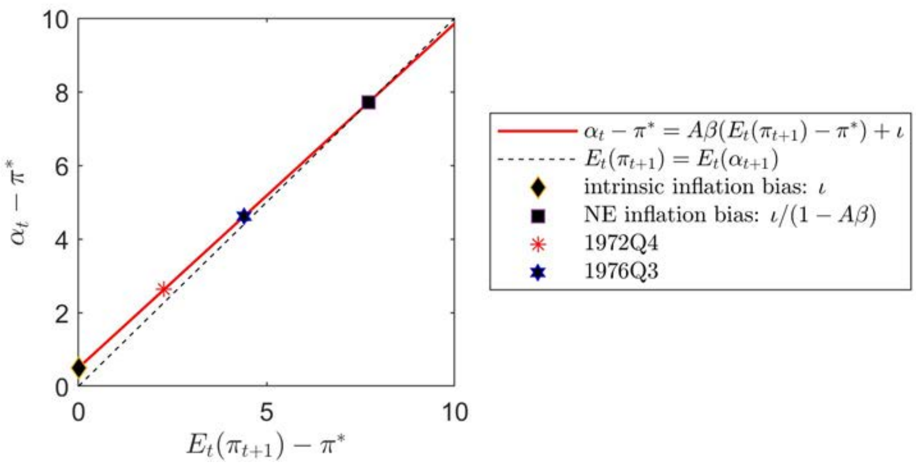

Figure 3: Targeted SPFs, SPF spread, and Estimated state variables

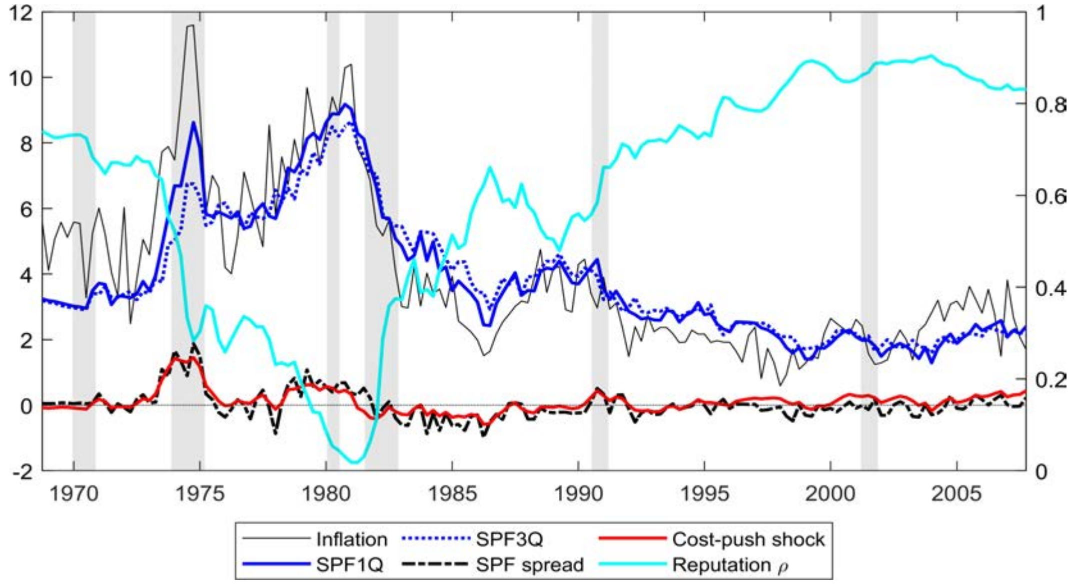

The SPF spread is the difference between the one and three quarter forecasts. All variables are continuously compounded annualized rates of change. Appendix C provides details on our SPF constructions.

Figure 4: Smoothed probabilities

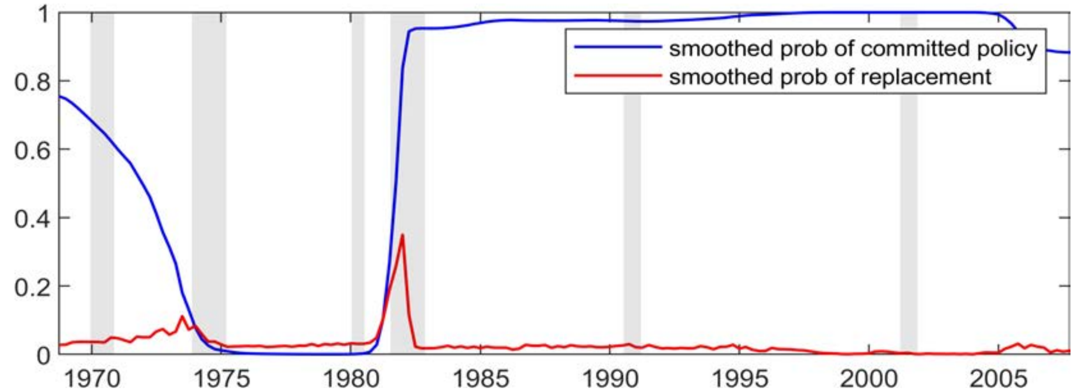

The smoothed probability of committed policy $Pr(\tau_t = 1|Y^T)$ is the sum of three conditional probabilities $Pr(\theta_t = 0, \tau_t = 1|Y^T)$ and $Pr(\theta_t = 1, \phi_t = 0, 1, \tau_t = 1|Y^T)$. The smoothed probability of replacement $Pr(\theta_t = 1|Y^T)$ is the sum of four conditional probabilities $Pr(\theta_t = 1, \phi_t = 0, 1, \tau_t = 0, 1|Y^T)$.

Figure 5: Inflation history and model-implied inflation

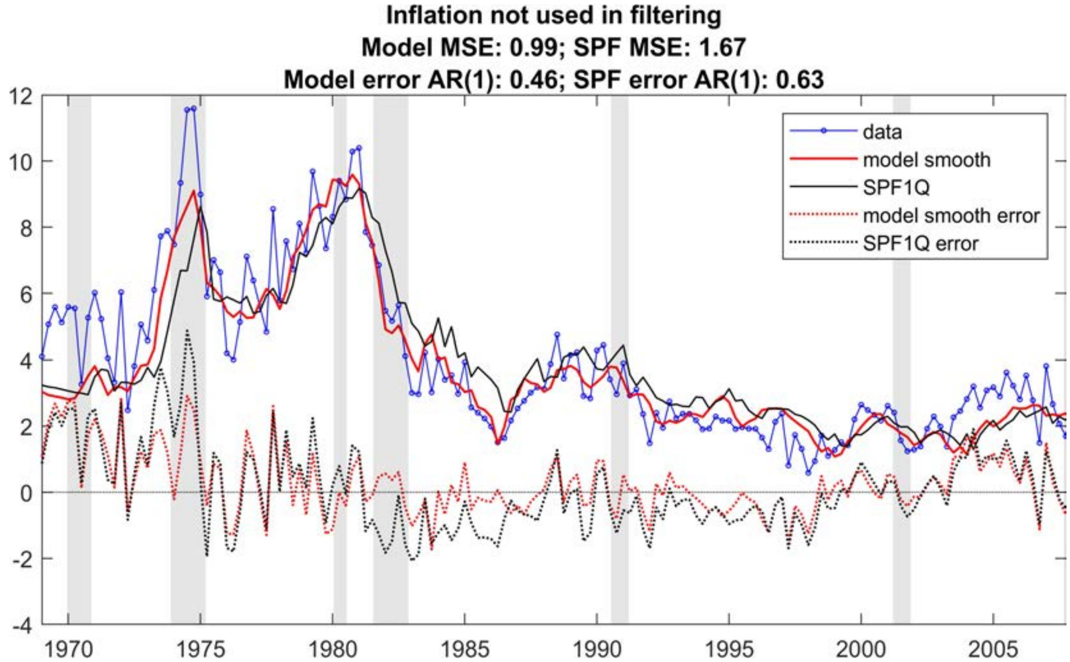

Figure 6: Model-based interpretation of US inflation

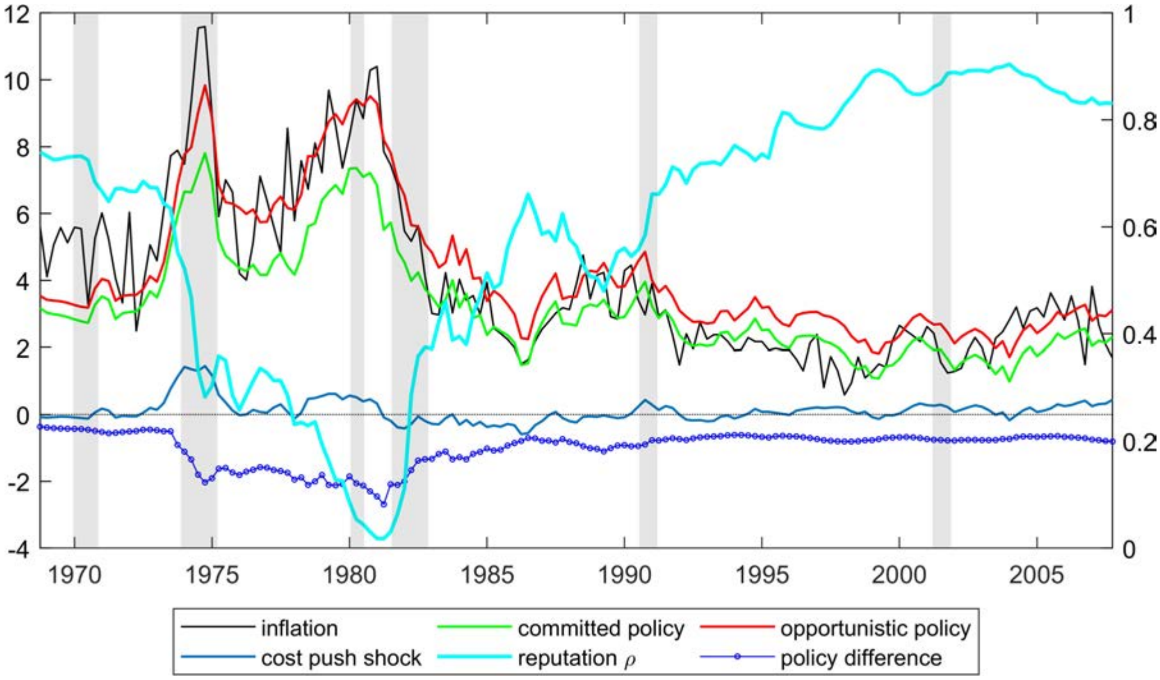

Figure 7: Optimal policy functions

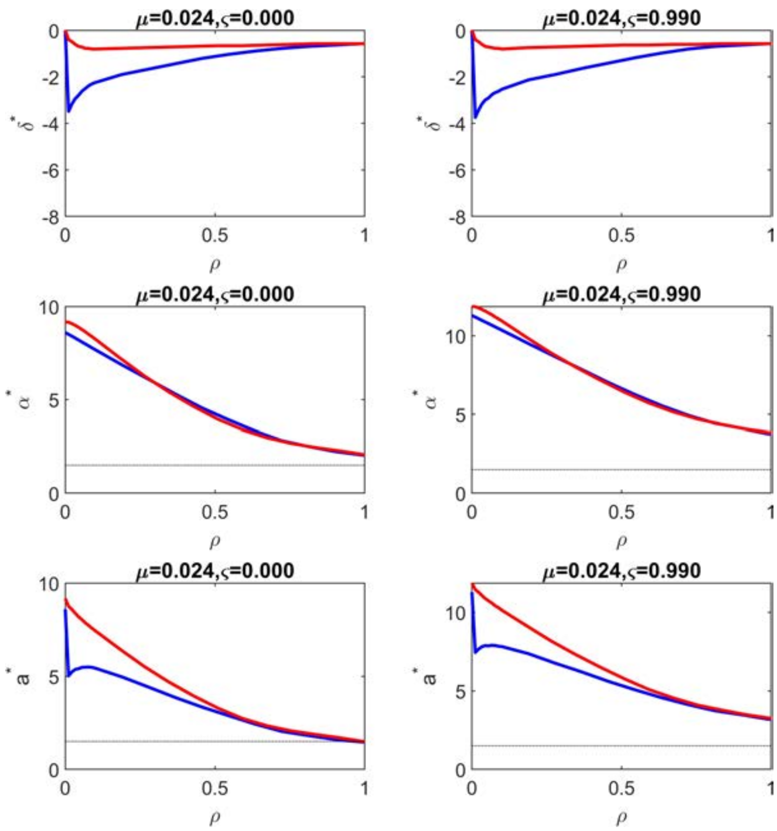

The blue lines are policy functions of the benchmark model where the optimal committed policy takes into account its influence on private sector's learning. The red lines are policy functions of a model where a naive committed policymaker treats reputation an exogenous process.

Figure 8: Counterfactual

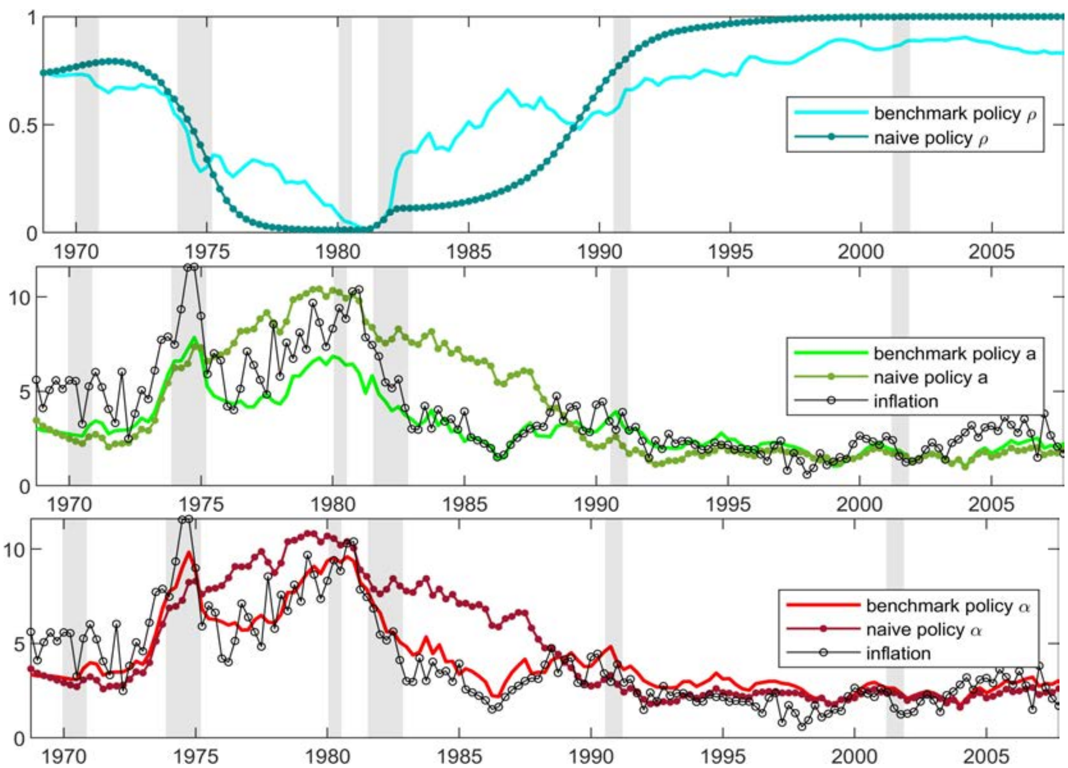

Figure 9: Model-based interpretation of recent US inflation

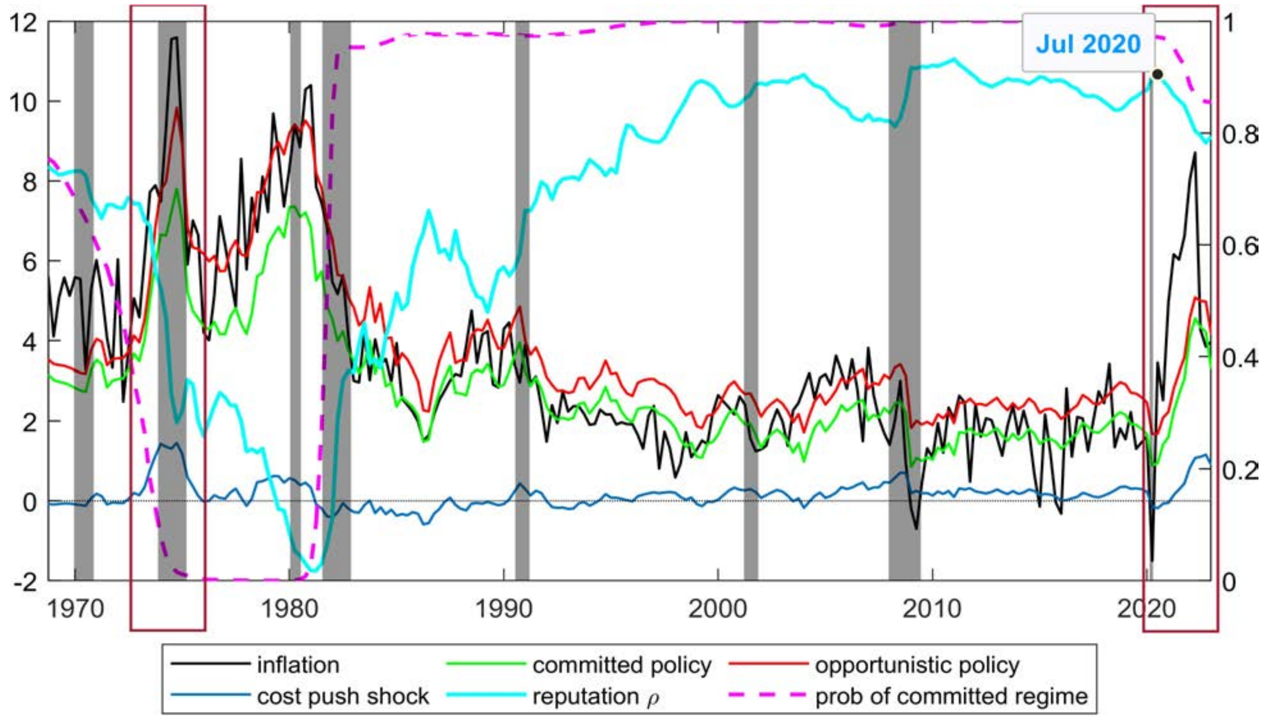

Figure 10: Contrasting median inflation and change in median price

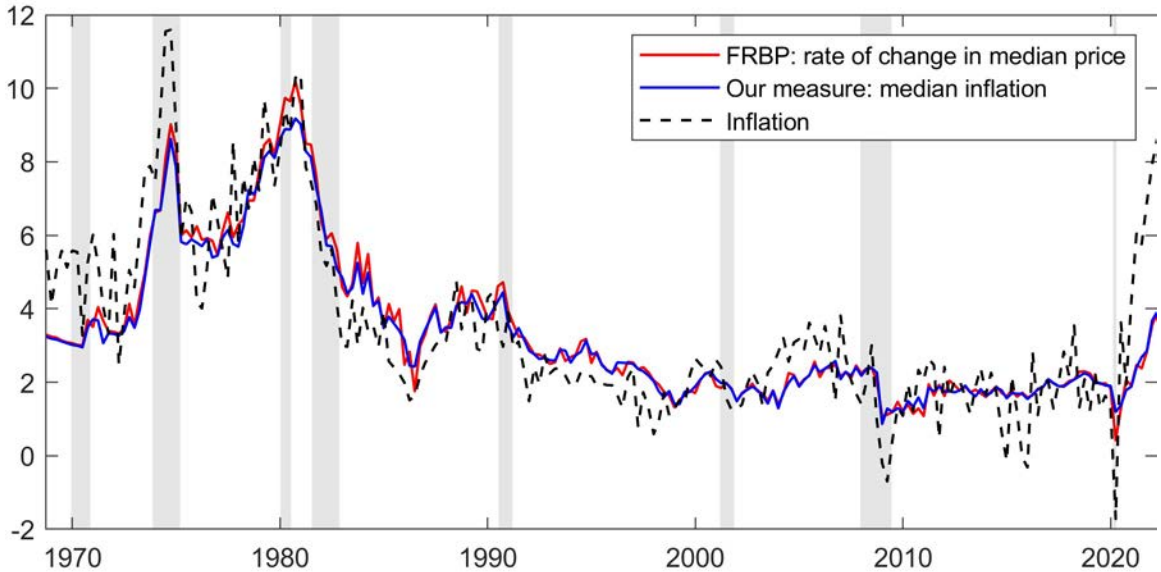

Figure 11: Model-implied and SPF forecasts of inflation

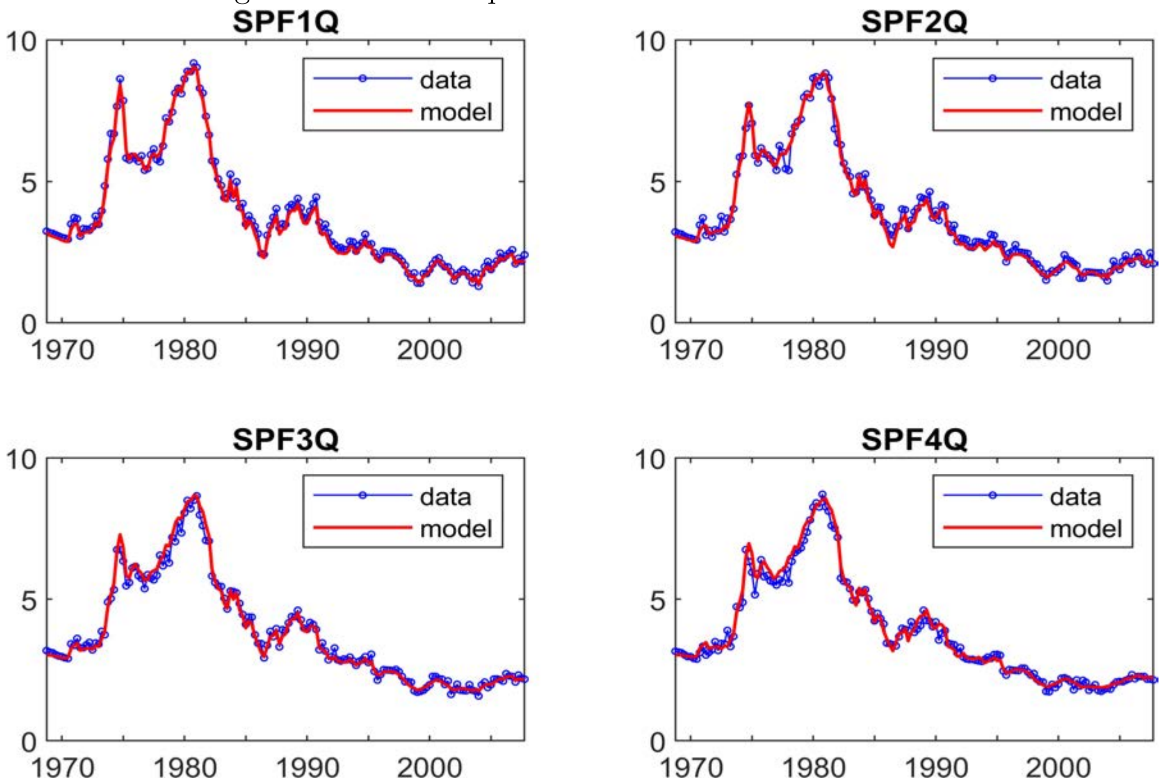

Figure 12: Untargeted cost-push shocks

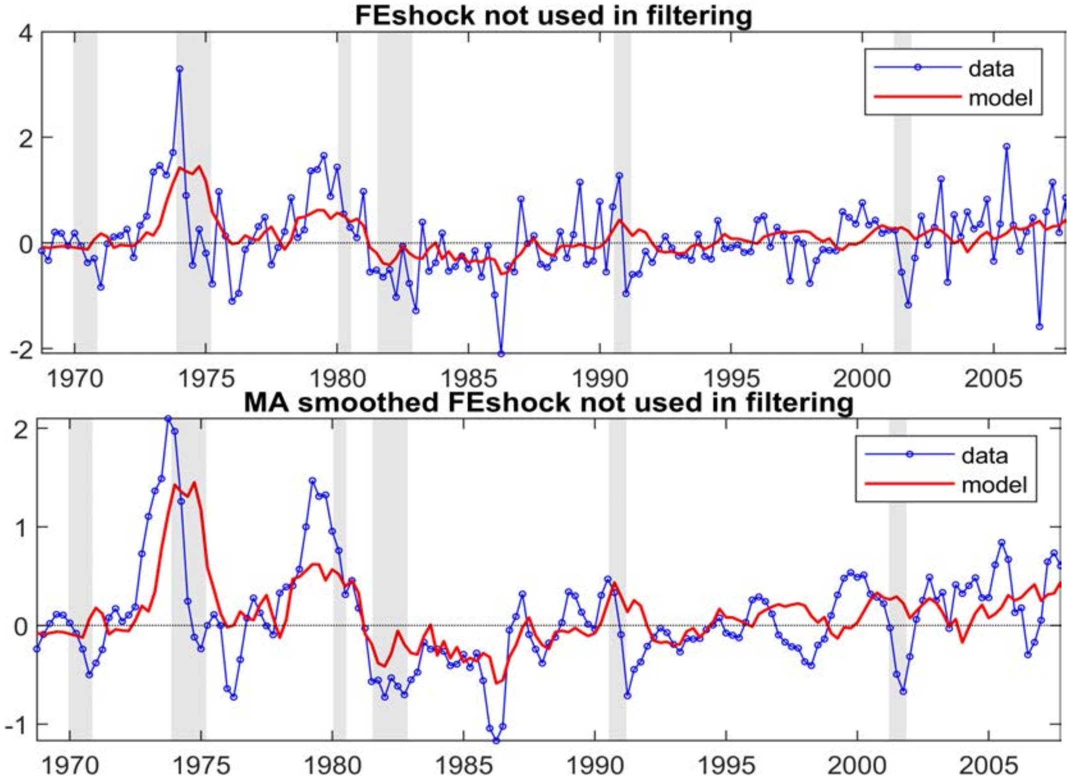

Note: Comparing smoothed estimates of cost-push shock $\widehat{\varsigma}$ to the FEshock – "Food and Energy price shock," constructed as the difference between the growth rate of the overall personal consumption deflator and its counterpart excluding food and energy.

Figure 13: Untargeted inflation: filtered result

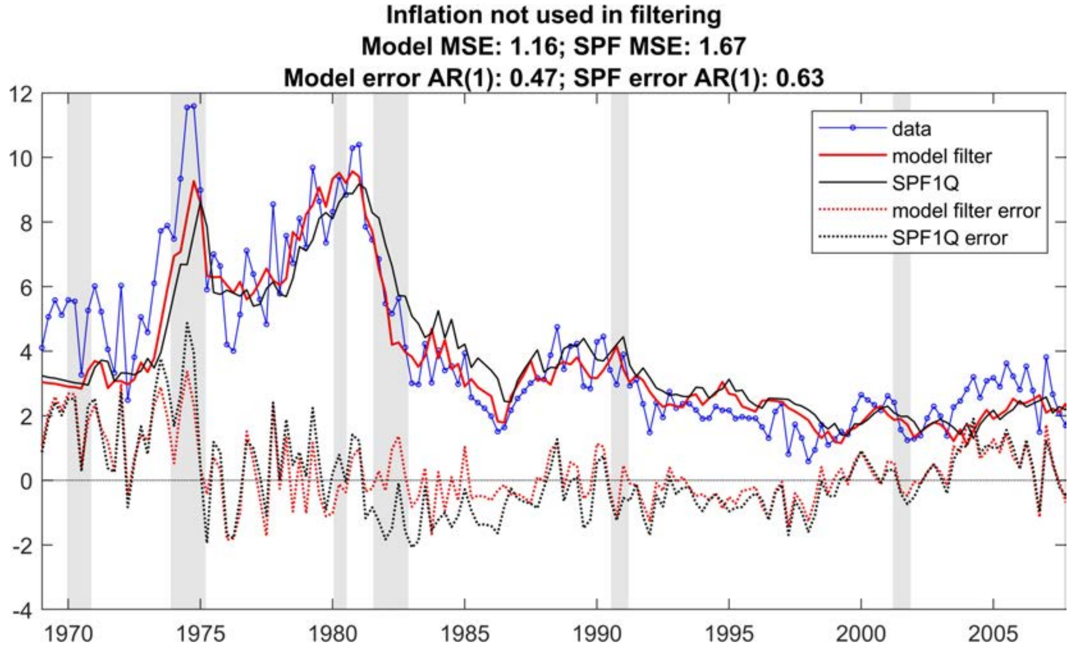

Figure 14: Model-implied and SPF forecasts of inflation

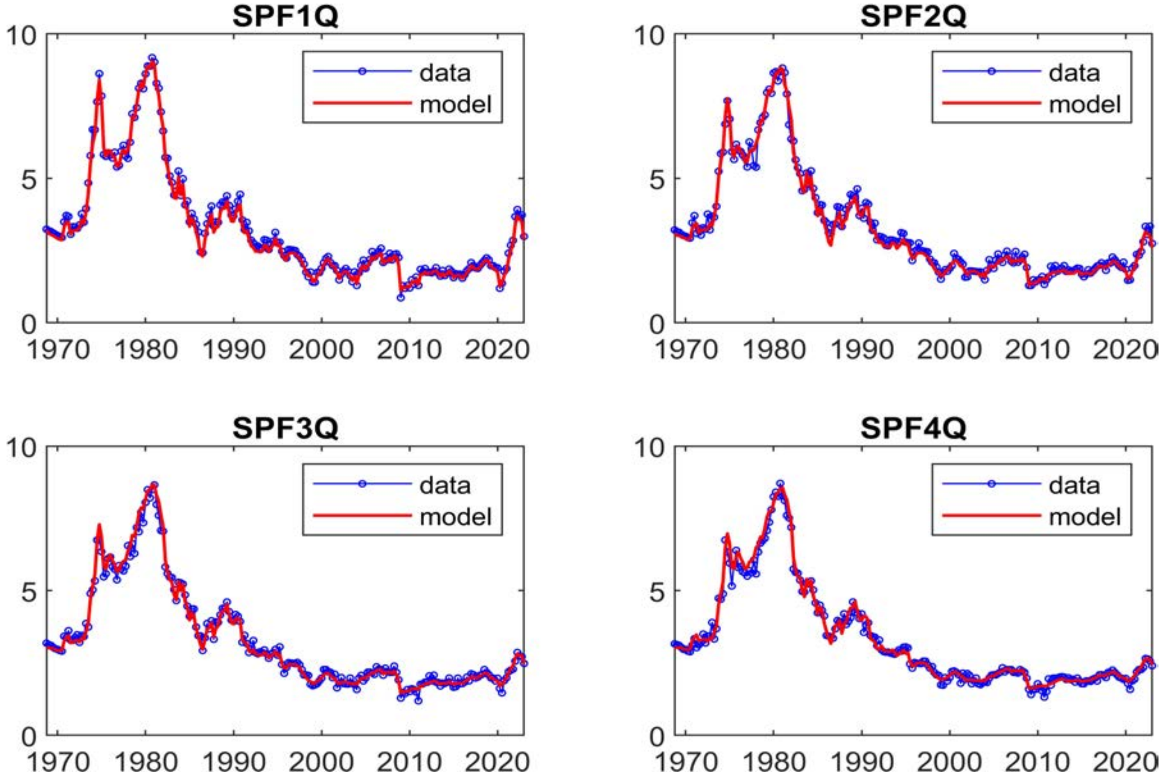

Figure 15: Inflation history and model-implied policies