Evolving Reputation for Commitment: Understanding US Inflation and Inflation Expectations*

Robert G. King† Yang K. Lu‡

December 2024

Abstract

We provide a model to study interaction between inflation expectation and optimal monetary policy. Private agents form forward-looking expectation based on their belief about how likely the central bank can commit – implementing the announced state-contingent policy plan. Their belief is governed by Bayesian learning from observing the history of inflation. The central bank that can commit chooses the optimal announced plan taking into account how it affects inflation expectations. The central bank that cannot commit chooses its policy taking inflation expectations as given. We develop a recursive solution and a quantitative implementation of the model. Our results show that the interaction between expectation and policy is central to understanding the dynamics of US inflation and inflation expectations.

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

JEL classifications: E52, D82, D83.

*We thank the participants in various conferences, seminars, and workshops for valuable feedback. Particularly useful advice was provided by Thorsten Drautzburg, Narayana Kocherlakota, Albert Marcet, Elmar Mertens, Mark Watson, and Thomas Sargent; valuable discussions were made by Donghai Zhang and Shengliang Ou. 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

Managing expectations is central to monetary policy decisions, because inflation expectations are relevant for economic activities (Coibion et al. (2018b)), can help to explain inflation (Beaudry et al. (2024)), and are responsive to central bank policy communication (Blinder et al. (2024)).

In this paper, we develop a theory for the interaction between inflation expectations and policy. Past inflation policies influence how inflation expectations respond to inflation surprises, while these expectations, in turn, shape the formation of future policy. The central link between expectations and policy is reputation, which reflects the private sector's belief in the central bank's commitment to following its announced plans. Reputation determines how inflation expectations respond to the central bank's communication about future policy. Reputation evolves as the private sector updates its belief based on the history of past inflation, leading inflation expectations to adjust in response to past inflation surprises. Importantly, inflation expectations play a crucial role in shaping optimal policy, as they incorporate the effects of past shocks and announced policy measures.

To show how our theory can be used to enhance understanding of the interplay of inflation expectations and inflation policy in the United States, 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 stochastic changes in regime.[1] 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 sector learning. A central quantitative result is that evolving reputation is crucial for matching key aspects of U.S. inflation history.

Our modeling approach builds on a mass literature, pioneered by Lucas, Sargent, Kydland and Prescott, that stresses the importance of policymaker commitment capacity in economic

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

policy. Influential studies by Lucas and Sargent, as summarized in their manifesto (1979), 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 sector learning is important and, indeed, that policymaker commitment capacity is inherently *unobservable*. A substantial body of literature now integrates private sector learning into the theory of economic policy.[2] Yet, an important gap remains as little research features 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. Forward-looking behavior of private agents, coupled with both types of policymakers being purposeful, necessitates our novel theoretical approach. In recursive equilibrium, reputation – defined as private agents' likelihood that the policymaker can commit – emerges as a key endogenous state variable.

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 agents' expectations, which are affected by

[2]See for examples: Barro (1986), Backus and Driffill (1985), Phelan (2006), Dovis and Kirpalani (2022).
[3]See, for example, Clarida et al. (1999)

future committed policy, future opportunistic policy, and reputation (private agents' likelihood that each policy will take place). But 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]

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 quantitative 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 agents' 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 to jointly identify three structural shocks, the three state variables, and regime change events.[8]

[4] A 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 are – to our minds – undesirable.

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

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

[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).

Explaining joint dynamics of expected and actual inflation The nonlinear filter produces estimated reputation that exhibits a big swing, declining throughout 1970s to near zero by the end of 1980 and gradually climbing back afterwards. Using the estimated shocks and states, we compute the model-implied inflation values. Remarkably, these values align closely with the observed U.S. inflation, despite the fact that the observed inflation data is not used in estimating shocks and states. To assess the importance of having optimal committed policy purposefully influence private sector learning, we also conduct a counterfactual exercise in which a naive committed policymaker optimizes but ignores the effect of his policy on reputation evolution. 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. It is because such policymaker naivete results in a narrower policy difference between the committed and opportunistic policymakers, which slows down private sector learning and policymaker reputation building.

Forecast revision regressions validate our theory According to our theory, long-term inflation expectations depend on a reputation state that evolves through Bayesian updating of inflation forecast errors. Moreover, the sensitivity of reputation to forecast errors depends on optimal inflation policies and thus on the time-varying reputation state. We test these implications using regressions of SPF long-term forecast revisions on nowcast forecast errors and find that the estimated coefficient on forecast errors indeed is time-varying and in a pattern consistent with theoretical predictions.

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

[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."

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 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 agents' beliefs* and we use the SPF to extract the evolution of such beliefs. In seeking to recover the evolution of private agents' 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.

Hazell et al. (2022), Carvalho et al. (2023), and Beaudry et al. (2024) emphasizes the importance of inflation expectation in understanding inflation dynamics. Our work complements theirs by highlighting the role played by monetary policy in determining the dynamics of inflation expectation, and how optimal monetary policy could take into account its role in shaping inflation expectations to improve future output-inflation trade-off.

Use of the SPF also links our research to the large and growing literature on survey measures of inflation (Coibion et al. (2018a)). 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. (2018a)) or the model generating the data (Farmer et al. (2021)). Our

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

[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 sector learning include Ball (1995), Erceg and Levin (2003), Orphanides and Williams (2005), Goodfriend and King (2005), Cogley et al. (2015), and Melosi (2016).

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 perform quantitative analysis of U.S. inflation history. Section 6 provides empirical support for our model's prediction on time-varying sensitivity of long-term inflation forecasts to inflation surprises. Section 7 concludes.

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

[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 fiscal links to inflation.

executes it in all subsequent periods until replaced.[14] The committed inflation plan therefore shapes private agents' expected inflation. An *opportunistic* type ($\tau = 0$) makes the same announcements,[15] but chooses intended inflation on a period-by-period basis.

At the start of each period, the policymaker may be replaced through a publicly observed mechanism: the replacement event occurs 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 ζ_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]

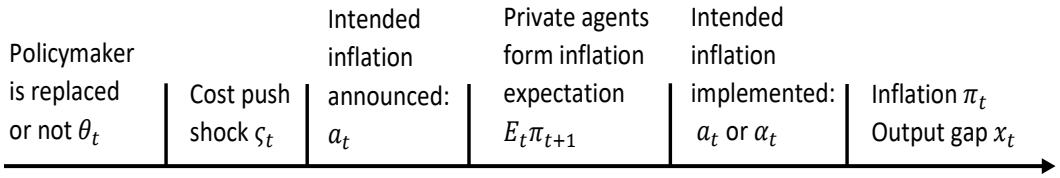
[14] We specify intended inflation rather than intended output for analytical convenience. If policy instead controlled intended real aggregate demand $\underline{x}_{\tau t}$ and $x_{\tau t} = \underline{x}_{\tau t} + \sigma_{x\tau}\varepsilon_t$, the Phillips curve $\pi_t = \kappa x_t + e_t + \varsigma_t$ 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 revelation. In a related fiscal model, Lu (2013) shows that 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 of signalling equilibria.

[16] We interpret random implementation error as a reduced-form representation for all unforeseeable factors that affect inflation beyond 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.

Figure 1: Timing of events within a period



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 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 $g(\pi_t|\alpha_t)$,[20] and an output gap x_t determined by the Phillips curve. New information leads private agents to update their beliefs about policymaker type.

[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.

[19] More specifically, $\rho_0 = \phi_t\rho_t + (1-\phi_t)v_{\rho,t}$, where $\phi_t \sim$ Bernoulli(ζ_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] Recall that π is governed by (2) so that $g(\pi|a)$ is the density function of a normal distribution with mean a and std $\sigma_{v,\pi}$; $g(\pi|\alpha)$ is the density function of a normal distribution with mean α and std $\sigma_{v,\pi}$.

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 $\{e(h_t)\}_{t=0}^{\infty}$.

3.2 Perfect Bayesian Equilibria

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

3.2.1 Consistent beliefs: reputation

Consistency of beliefs requires that private agents' assessments of policymaker type are updated according to Bayes' rule (4) which depends on policymakers' equilibrium strategies and observed inflation π_t. Within a regime, private agents' belief ρ are updated recursively,

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

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.

[21] This restriction is innocuous 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)).

3.2.2 Consistent beliefs: inflation expectations

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

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

Notice that a higher value of ρ makes the nowcast more responsive to a. 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 higher $\rho(h_t)$ making $e(h_t)$ more responsive to future inflation under commitment. Thus, in (5) and (6), higher reputation increases committed policymaker leverage over expectations. Turning to the details of (6),

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

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

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.

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

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 agents' inflation expectations. We consider 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.

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

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

[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.

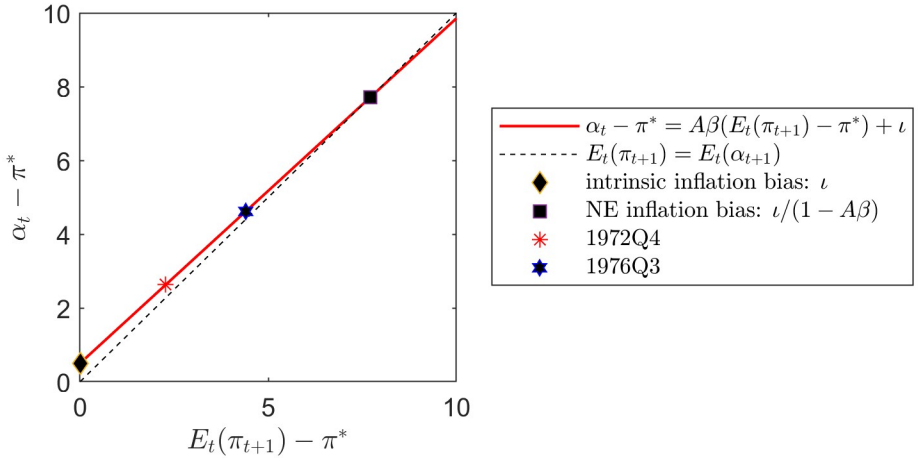
[24] There is a slight abuse of notation here by using summation Σ over history to capture the joint effects of continuous distribution of π and discrete Markov chain distribution of ς.

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 agents' 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 agents expect 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.

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

Figure 2: Optimal Response of Opportunistic Policy to Inflation Expectations



This figure plots how the opportunistic policymaker's optimal inflation bias varies with private agents' expected inflation, together 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. The diamond marker indicates the *intrinsic inflation bias* when private agents expect the inflation to be at target in our imperfect information framework. 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.

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 $\rho(h_t)$, private agents' nowcast of inflation $z(h_t)$ conditional on a replacement satisfies (5);

(ii) given $\alpha(h_t)$, $a(h_t)$, and $z(h_t)$, private agents' beliefs of policymaker type $\rho(h_{t+1})$ are updated according to (4); and their 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 agents' 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 agents' 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 agents' 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 agents' 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 $\lambda(\pi_t, a_t, \alpha_t) \equiv g(\pi_t|\alpha_t)/g(\pi_t|a_t)$ 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]

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

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

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: one role played by γ is that it enters as a shadow price of inflation expectations management (choice of e).[28] The pseudo state variable μ records past committed type promises (contained in ω).[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 reputation state ρ, and the divergent probability beliefs about inflation π held by the committed policymaker and private agents.[31]

[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.

[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.

The evolution of the pseudo state to μ' in (14) and the reputation state to ρ' in (15) identifies two channels through which the committed policymaker manages private agents' inflation expectations since these will influence future policy choices a' and α'.[32]

Expectation anchoring: The next-period pseudo state μ' evolves according to (14): a higher shadow price γ increases μ', making it more costly for the committed type to raise a' in the subsequent period period. That is, 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 anchoring is moderated by private sector skepticism, as the influence of γ on μ' depends on the reputation state ρ. With optimal γ, we write $\mu' = m(s)$ for convenience.

Reputation building: The next-period reputation state ρ' evolves according to Bayes' rule (15). The committed policymaker affects ρ' by choosing a difference between his intended inflation (a) and the intended inflation of an opportunistic type (α). A wider 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 leverage that the committed policymaker's intended inflation has on private agents' expected inflation. With optimal a and α, we write $\rho' = b(s, \pi)$ for convenience.

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]

[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 private agents' 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

Our model has three structural shocks: v_ς to the cost-push shock, to reputation in the event of replacement v_ρ and to inflation v_π. It also features three binary states $(\theta_t, \phi_t, \tau_t)$, indicating the event of policymaker replacement, whether or not there is reputation inheritance, and the type of current policymaker. Our theory identifies three continuous 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, and provides equilibrium functions of the three state variables that map the structural shocks and the binary states to macro variables that can be directly measured using data.

We build the quantitative model in two stages. First, calibrating model parameters, we can compute the equilibrium 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 agents'

[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. Since we are only considering equilibrium behavior from this point on, we also streamline notation by dropping the superscript $*$.

inflation forecasts at other horizons, $f(s_t, j) = E(\pi_{t+j}|s_t)$.

Second, 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 inflation forecasts 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]

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 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

[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.

Table 1: Parameters

π^*	Inflation target	1.5%
β, β_a	Discount factor (private agents, 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
ζ_ρ	prob of reputation inheritance	0.9
$\overline{\rho}$	mean of reputation draw	0.1
σ_ρ	std of reputation draw	0.05
ξ_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	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.

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 $\phi_t \sim$ Bernoulli(ζ_ρ) and $v_{\rho,t} \sim$ 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 $\zeta_\rho = 0.9$, $\overline{\rho} = 0.1$, 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 ξ_ς and innovation volatility $\sigma_{v,\varsigma}$. The transition probability matrix $\varphi(\varsigma'; \varsigma)$ is calibrated to approximate $\varsigma' = \xi_\varsigma \varsigma + v_\varsigma$ 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

[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.

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 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 agents' 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 ($\theta = 1$,

[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.

[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).

$\phi = 1$, $\tau = 1$), 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.

5.3 The state space model with Markov-switching

We now detail the evolution of continuous state variables $S_t = [\varsigma_t, \rho_t, \mu_t, \pi_t]$ resulting from $v_t = [v_{\varsigma,t}, v_{\rho,t}, v_{\pi,t}]$, taking as given the discrete state $\Theta_t \in \{1, 2, ..., 6\}$ defined above.

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

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:

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

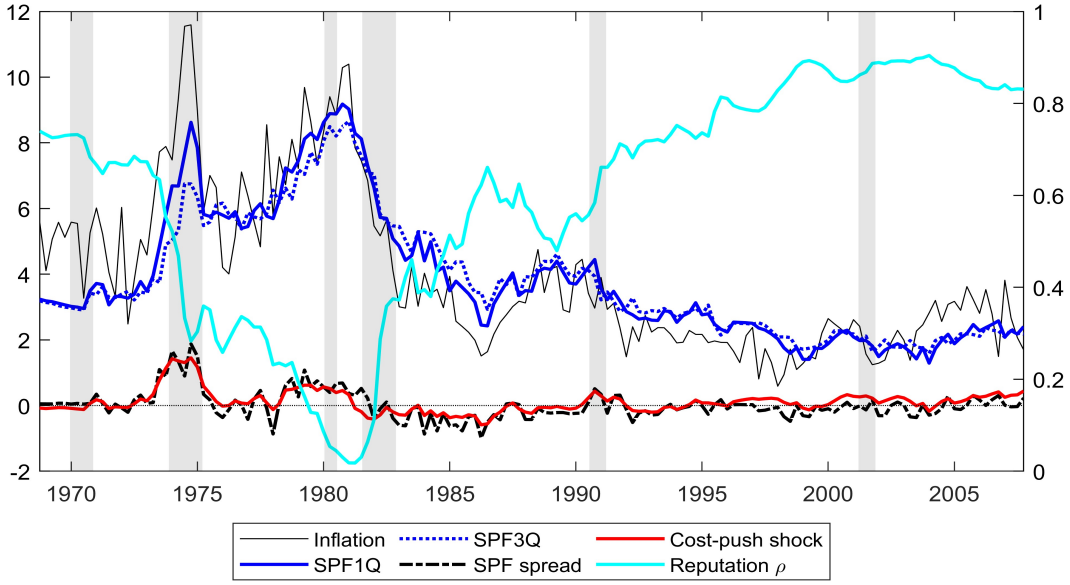
As our model is not linear, we cannot use the standard Kalman filter. We adopt a particular "sigma point" approximation method – the unscented Kalman filter – that has been shown to work well in nonlinear regime-switching models. Appendix C.3 provides details on the algorithm, which also employs the collapsing approach of Kim (1994) and

[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).

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).

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



This figure plots selected inputs and outputs of our nonlinear filter. The inputs are the one-quarter-ahead and three-quarter-ahead SPF inflation forecasts. Appendix C provides details on our SPF constructions. The SPF spread (black dashed line) is the difference between SPF1Q and SPF3Q. It rises during the first (1974-75) and the second (1978-80) inflation surges (black solid line). The outputs are smoothed estimates of reputation (cyan and measured on the right axis) and cost-push shocks (red). Consistent with our strategy of exploiting greater sensitivity of near-term forecasts to transitory shocks, the estimated cost-push shock covaries positively with the SPF spread.

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) are reported in Figure 3.[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.[49]

[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.6 Figure 12 reports how our estimated cost-push shock covaries with the FE shock.

[49] For the estimates of discrete states, please refe to Appendix C.5.

5.4 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.

5.4.1 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. Appendix C.6 Figure 10 shows that these additional forecasts lie almost entirely on top of the untargeted SPF2Q and SPF4Q, providing support for our state extraction approach.

5.4.2 Observed and estimated inflation

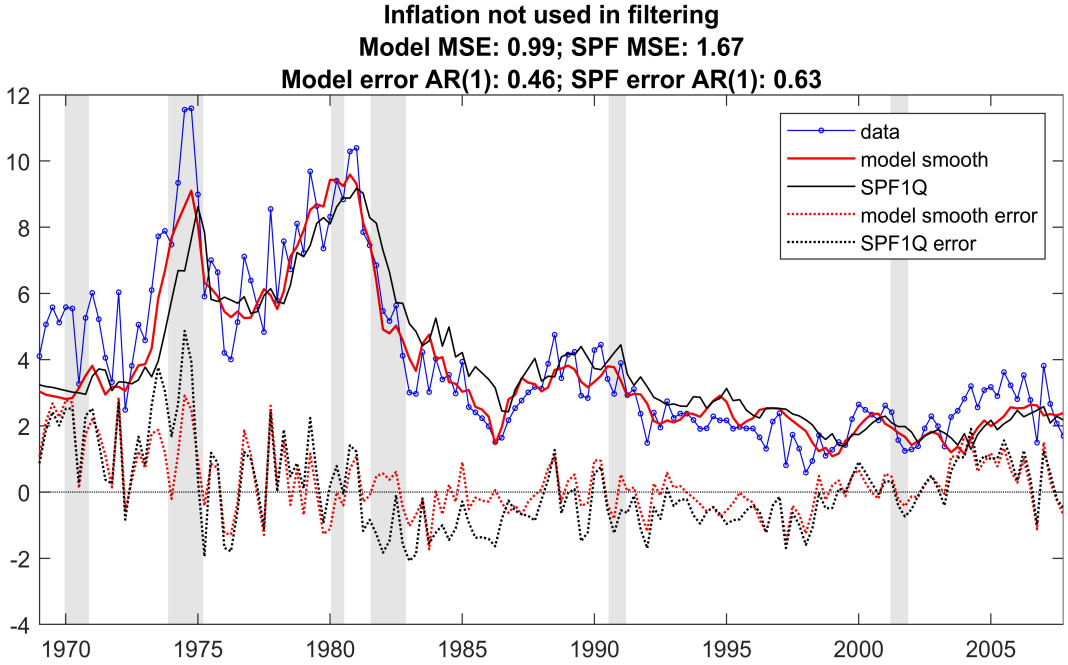
Our state-space model treats inflation as a latent state variable and therefore produces filtered and smoothed estimates for π_t. Because our extraction method only uses SPF data to obtain states, comparing the smoothed estimates $\widehat{\pi}_t$ with the observed inflation data serves as a model validation. To assess how well our estimates correspond to observed inflation data, we use the SPF1Q as a benchmark.[50]

Figure 4 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 concern that our smoothed measure of model inflation is based on the full sample, while the SPF is prepared in real time. In Appendix C.6 Figure 13, we therefore provide a version of Figure 4 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.

[50] Since our extraction is based on SPF1Q and SPF3Q, one may argue that it must work well because the SPF also tracks observed inflation. Using SPF1Q as the performance benchmark speaks to this skepticism.

Figure 4: Inflation history and model-implied inflation



This figure shows that the smoothed estimates for inflation (red) produced by our nonlinear filter align very well with the U.S. inflation data (blue), even though the nonlinear filter only uses SPF data. The one-quarter-ahead SPF inflation forecasts (SPF1Q) from the prior quarter (black) are also plotted as a comparison. The gap between inflation and the smoothed estimate (model smooth error) is plotted using a dashed red line; the gap between inflation and SPF1Q from the prior quarter (SPF1Q error) is plotted using a dashed black line. The comparison reveals that our smoothed estimate fits the inflation data better than the SPF1Q by two statistical measures: its squared error has a lower mean and its error is less persistent.

5.5 Counterfactual with naive committed policy

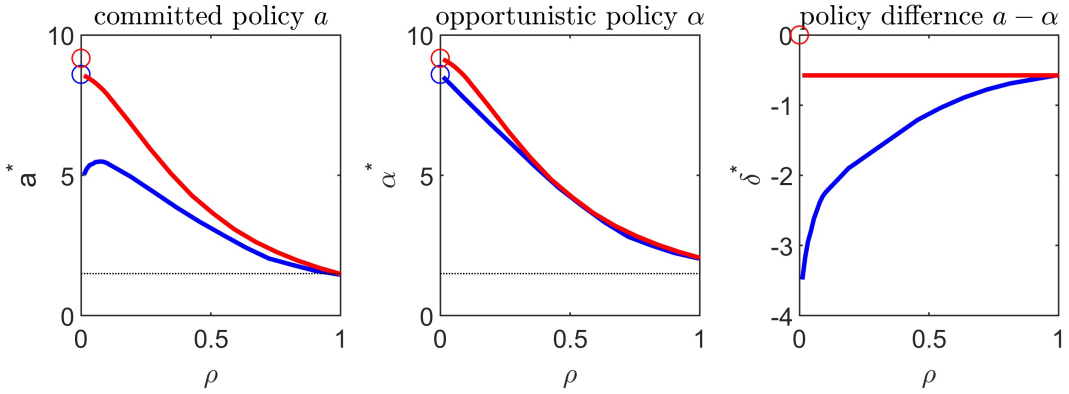
While previous studies have examined the evolution of private agents' beliefs about the policymaker,[51] they generally abstract from the strategic actions policymakers may take to shape those beliefs. In contrast, our paper highlights this intentional influence as a defining feature of optimal committed policy.

In this subsection, we show that this intentional policy influence on private agents' belief is central to understanding the Volcker disinflation. To this end, 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 factor governing the leverage that his

[51] See footnote 12 for examples.

policy a has over inflation expectations.[52] Using these policy functions, we construct time series of inflation policies based on the estimated cost-push shocks from the benchmark model, the endogenous time-varying reputation shaped by the naive policymakers' past responses to shocks, and the model-consistent pseudo-state for the naive committed policymaker.

Figure 5: Equilibrium policy functions



Equilibrium policies as functions of reputation ρ. Each panel compares two policy functions: one in our benchmark model in which the committed policymaker actively manages his reputation (in blue), and the other in a naive policy model in which a committed policymaker takes reputation as an exogenous factor (in red). The policy functions are conditional on the values of two other states: μ and ς. We set μ at its steady state level when $\rho = 1$. The cost push shock ς is set at 0. When the committed policymaker actively manages his reputation, the optimal policy difference widens with poorer reputation. When the committed policymaker takes his reputation as an exogenous factor, the optimal policy difference is narrower and insensitive to reputation.

Policy Functions Figure 5 displays how committed policy a (left panel), opportunistic policy α (middle panel), and their policy difference $\delta = a - \alpha$ (right panel) vary with reputation ρ.[53] Each panel compares two policy functions: one where a committed policymaker acts naively (in red), and our benchmark model in which the committed policymaker actively manages his reputation (in blue).[54]

[52] Appendix E 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.

[53] The policy difference $\delta = a - \alpha$ is always negative because the committed type sees his intended inflation a as affecting expected inflation, in contrast to the opportunistic type's view.

[54] The policy functions are conditional on the values of two other states: μ and ς. Relative to ρ, μ and ς are less important in determining the shape of the policy functions. In this figure, we set μ at its steady state level when $\rho = 1$, and ς at zero.

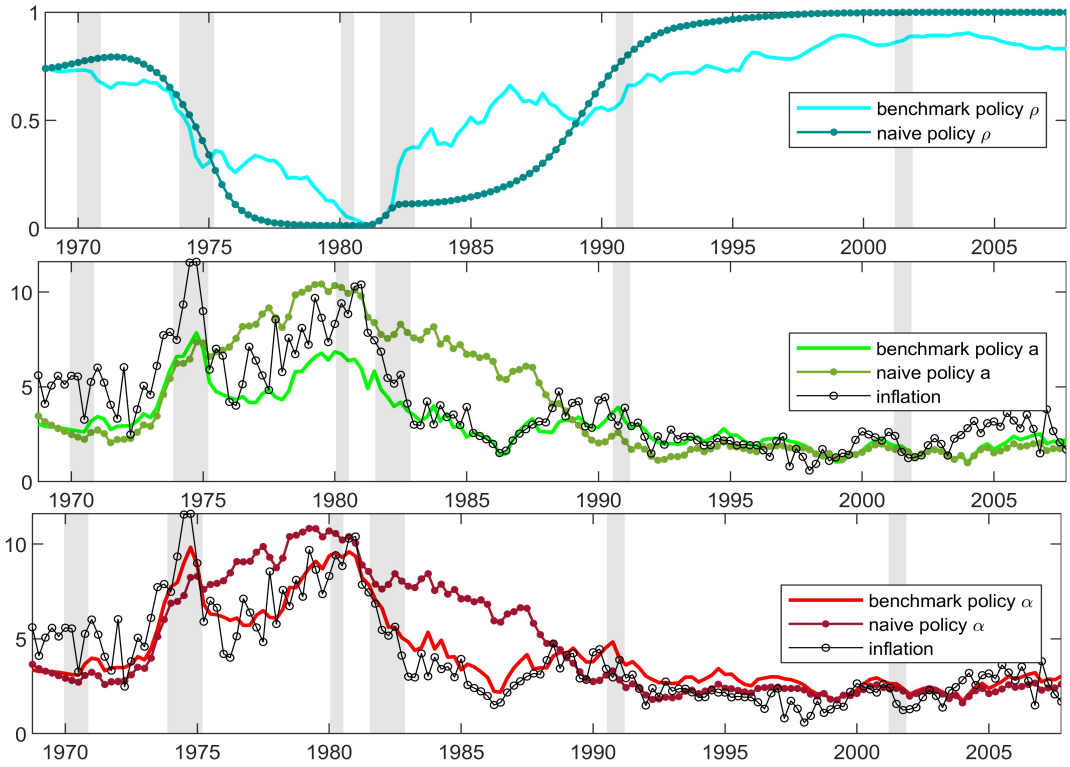
The gap between naive and benchmark policy behavior is most evident in the the policy difference functions (top panels). When reputation gets lower, the policy difference widens under benchmark policy but shrinks under naive policy. As a result, the policy difference is always narrower in the naive policy model than it is in the benchmark policy model, but especially so 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 narrower policy difference at lower reputation. Only when the committed policymaker understands that a wider 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.

Time Series Figure 6 contrasts reputation and intended inflation policies in the counterfactual case of a naive committed policymaker with the benchmark case where the committed policymaker actively manages its reputation.[55]

The period of interest is post-1981 when a committed policy regime is more plausible but reputation starts at a low level. In the benchmark model, the committed intended inflation is nearly three percentage points lower than the opportunistic intended inflation. This substantial and persistent policy difference drives rapid reputation-building during the 1982–83 recession. In contrast, when the committed policymaker behaves naively, reputation remains low for an extended period after 1981, even without implementation errors. This occurs because the naive committed policymaker treats time-varying reputation as an exogenous factor, resulting in a committed policy that closely resembles its opportunistic counterpart. Low reputation leads to high expected inflation, worsening the trade-off between inflation and output. Consequently, the naive committed policymaker takes significantly longer to disinflate the economy, diverging from the rapid disinflation observed in the post-1981 U.S. inflation history.

[55]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 11). 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.

Figure 6: Counterfactual



This figure plots the time series of reputation ρ (top panel), committed policy a (middle panel), and opportunistic policy α (bottom panel) produced by the benchmark model and by the naive policy model, using the same cost-push shocks and probabilities of states as estimated in our benchmark model. 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. The observed inflation data (black circled line) are plotted as a reference. In the benchmark model, the committed policy after 1981 is much lower than the opportunistic policy, resulting in a rapid increase in reputation. By contrast, in the naive policy model where the incentives to manage reputation are missing, it takes much longer for the committed policymaker to disinflate the economy after 1981, and reputation remains low for an extended period even without any implementation error.

6 Understanding long-term inflation expectations

We have seen that our quantitative model yields time series that align closely with US inflation history over 1968 to 2007, matching both targeted short-term expected inflation and untargeted observed inflation. We now use a reduced-form approach to examine the implications of our theory on long-term inflation expectations, in particular, how longer-

term inflation forecast revision responds to inflation shocks.

6.1 Long-term inflation expectation proportional to reputation

According to our theory, private sector's inflation forecasts depend on both reputation and cost push shocks. Since cost-push shocks are stationary, their impact on inflation forecast diminishes over forecasting horizons, ultimately becoming irrelevant for long-term inflation forecast. Reputation is therefore the sole determinant for long-term inflation forecasts:

$$f_{\infty|t} = (1-q)[\rho_t \pi^* + (1-\rho_t)\pi^{NE}] + q[\rho_t z(\varsigma = 0, \rho = 1) + (1-\rho_t) z(\varsigma = 0, \rho = 0)] \tag{17}$$

The term in the first squared bracket is the expected long-term inflation conditional on staying in the same regime (with probability $1-q$). In this case, policymaker's type will be ultimately revealed and the inflation will converge to either the inflation target π^* or the Nash-equilibrium inflation bias π^{NE}, depending on the type of policymaker in place. The term in the second squared bracket is the expected "startup inflation" in the new regime when there is a replacement (with probability q).

6.2 Dynamics of reputation

Since the long-term inflation forecast is proportional to reputation, it inherits the time series properties of reputation, which we now examine.

Under our theory, reputation is a function of π and evolves according to

$$\rho_{t+1} = b(\pi_t; a_t, \alpha_t, \rho_t) = \frac{\rho_t g(\pi_t|a_t)}{\rho_t g(\pi_t|a_t) + (1-\rho_t) g(\pi_t|\alpha_t)}$$

We take a linear approximation of the change in reputation around $[\rho_t a_t + (1-\rho_t)\alpha_t]$, the now-cast of inflation made by private agents at the start of the period:[56]

$$\rho_{t+1} - \rho_t \approx -k\{\rho_t(1-\rho_t)[a_t - \alpha_t]\}_1\{\pi_t - [\rho_t a_t + (1-\rho_t)\alpha_t]\}_2 \tag{18}$$

where $\pi_t - [\rho_t a_t + (1-\rho_t)\alpha_t]$ is a measure of inflation surprise, and k depends inversely on the volatility of implementation errors.

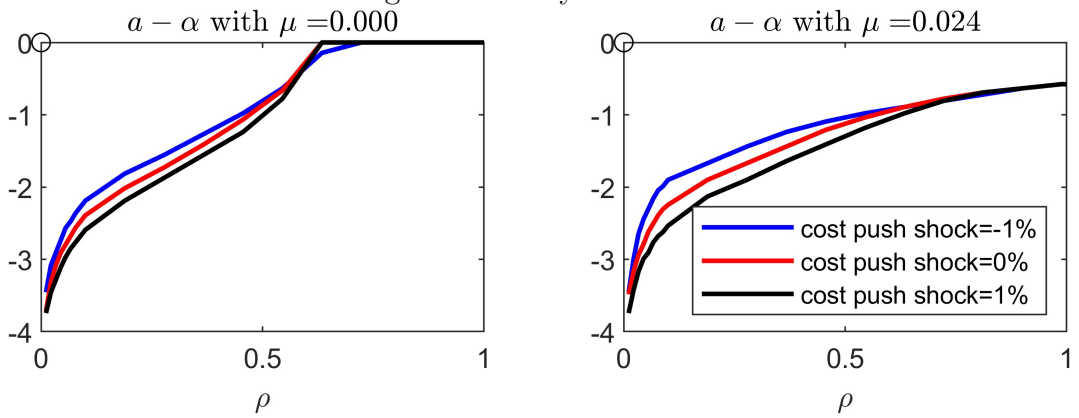
[56] The approximation will be exact in a continuous-time version of imperfect public monitoring games, for example, Faingold and Sannikov (2011). Appendix XX maps the following expression to Proposition 1 in Faingold and Sannikov (2011).

The first term in braces indicates that a time-varying influence of inflation surprises on reputation. The influence is larger when (i) reputation is away from its boundaries – the reputation weight $\rho(1-\rho)$ is maximized at 1/4 when ρ is exactly in the middle of its range – and (ii) the policy difference is greater in absolute value.

The policy difference itself varies with reputation. Figure 7 plot the equilibrium policy difference, $a - \alpha$, as a function of ρ at various levels of cost-push shock and μ. Lower reputation implies larger policy difference in absolute value.

Taking together, the influence of inflation surprises on reputation changes, and in turn on revisions in long-term inflation forecast, is increasing in $[\rho(1-\rho)](1-\rho)$ where the second $(1-\rho)$ approximates the effect of ρ on the policy difference.

Figure 7: Policy Difference



Difference between equilibrium committed policy and opportunistic policy is larger at lower levels of reputation. This property holds for various levels of cost push shock and pseudo state μ. The two values of μ are representative as they correspond to the steady state μ when $\rho = 0$ and $\rho = 1$, respectively.

6.3 Time variation in sensitivity to inflation news

We now study empirically how revisions in long-term inflation forecasts respond to news about inflation, using the SPF CPI inflation forecast data which includes 10-year-ahead inflation forecast available starting in 1991Q4. our sample covers up to 2024Q2.

To construct the "inflation surprise", we utilize the "backcast" estimate of the prior quarter inflation and a "nowcast" estimate of the current quarter inflation reported by SPF participants in each quarter.

$$v_t^i = f_{t-1|t}^i - f_{t-1|t-1}^i$$

The "backcast" $f^i_{t-1|t}$ is the realized CPI inflation last quarter as perceived by the forecaster i. $f^i_{t-1|t-1}$ is the same forecaster's last quarter's "nowcast" of inflation of the same quarter. For each forecaster, we can also construct 10-year forecast revision as

$$v^i_{40,t} = f^i_{t+40|t} - f^i_{t+39|t-1}$$

We aggregate the individual variables by considering only forecasters present in both t and $t-1$, then taking the mean as the consensus measure.

The empirical counterpart of (18) is:

$$v_{40,t} = \gamma_t v_t + \epsilon_t \tag{19}$$

where the effect of inflation surprise on long-term forecast revision γ_t is time-varying and increasing in $\rho_t(1-\rho_t)^2$.

To capture the dependence of γ_t on ρ_t, we employ the relation between reputation and long-term inflation forecast in (17) to construct a reduced-form measure of reputation, $\widehat{\rho}_t$, with the SPF 10-year CPI inflation forecast. We then use $\widehat{\rho}_t$ to construct weights to inflation surprise and examine if the weighted inflation surprises outperform the unweighted ones in explaining revisions of long-term inflation forecast. We compare three regression models:

$$\text{Model 1: } v_{40,t} = \gamma v_t + \epsilon_t$$

$$\text{Model 2: } v_{40,t} = \gamma \widehat{\rho}_t(1-\widehat{\rho}_t) v_t + \epsilon_t$$

$$\text{Model 3: } v_{40,t} = \gamma \widehat{\rho}_t(1-\widehat{\rho}_t)^2 v_t + \epsilon_t$$

Model 1 uses the unweighted inflation surprise as the regressor. Model 2 weights the inflation surprise by $\widehat{\rho}_t(1-\widehat{\rho}_t)$, reflecting the idea that long-term forecast revisions are proportional to the private sector's belief updating about the policymaker's type —- belief updating is more responsive to shocks when uncertainty about the policymaker's type is higher. Model 3 modifies the weight to $\widehat{\rho}_t(1-\widehat{\rho}_t)^2$, capturing the interaction between policy and reputation: lower reputation increases policy differences and accelerates belief updating by the private sector.

Table 2 reports the performance of the three models in terms of their adjusted R squared and root mean squired error. Compared to unweighted inflation surprise, the weighted surprises show significantly better explanation power for the revision of long-term inflation

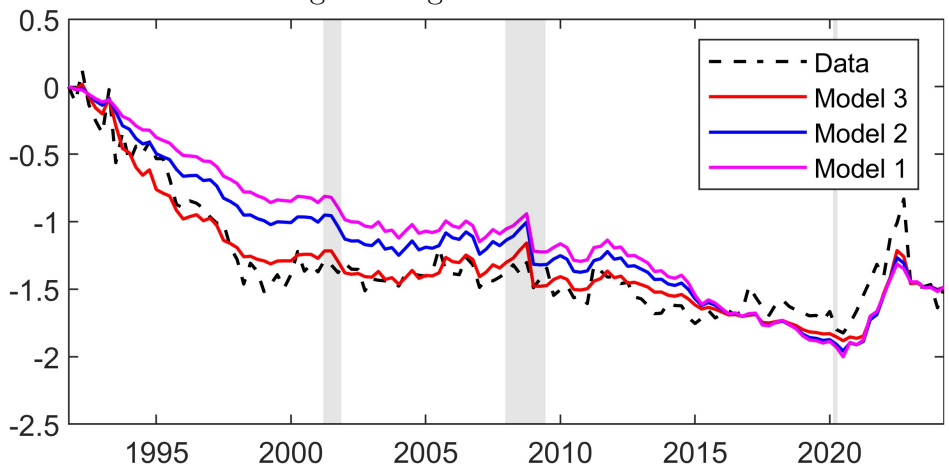
forecast. The adjusted R squared increasing from 0.163 in Model 1 to 0.193 in Model 2 and further to 0.248 in Model 3. The corresponding root mean squared error decreases from 0.117 in Model 1 to 0.114 in Model 2 and to 0.11 in Model 3.

Table 2: Comparison of three regression models

	γ	p value	Adjusted R^2	RMSE	N
Model 1	0.038	1.13E-06	0.163	0.117	130
Model 2	0.384	1.04E-07	0.193	0.114	130
Model 3	3.078	1.02E-09	0.248	0.11	130

OLS regressions of long-term CPI inflation forecast revision on inflation surprise. Sample period is from 1991Q4 to 2024Q2. Model 1 uses the unweighted inflation surprise. Models 2 and 3 weight inflation surprise using $\widehat{\rho}_t(1-\widehat{\rho}_t)$ and $\widehat{\rho}_t(1-\widehat{\rho}_t)^2$, respectively, where $\widehat{\rho}_t$ is constructed using the lagged long-term inflation forecast according to (17). RMSE stands for Root Mean Squared Error. All regressions include a constant, but its coefficients are statistically insignificant.

Figure 8: Cumulative change in long-term CPI inflation forecast since 1991Q4



The cumulative change in the long-term CPI inflation forecast from 1991Q4 to 2024Q2 against its fitted counterparts produced by three models. The black dash line is the data, the magenta, blue, and red lines correspond to the fitted line using Models 1, 2, and 3, respectively. Model 1 uses the unweighted inflation surprise. Models 2 and 3 weight inflation surprise using $\widehat{\rho}_t(1-\widehat{\rho}_t)$ and $\widehat{\rho}_t(1-\widehat{\rho}_t)^2$, respectively, where $\widehat{\rho}_t$ is constructed using the lagged long-term inflation forecast according to (17).

Figure 8 visualizes the model comparison by plotting the cumulative change in the long-term inflation forecast against its fitted counterparts produced by three models. The black

dash line is the data, the magenta, blue, and red lines correspond to the fitted line using Models 1, 2, and 3, respectively. It is clear that Model 3 outperforms the other two models in fitting the actual forecast revision. The improvement is most visible before 2000 when there are significant downward revisions of long-term forecast, i.e., an improvement of reputation.

6.4 Stochastic inflation trends

A common element of time series econometric studies of inflation and survey inflation expectations is the assumption that there is a stochastic trend, which is a random walk or martingale. In the univariate analysis of Stock and Watson (2007), for example, this stochastic trend explains much of U.S. inflation but only if the posited unit root process has changes with time-varying volatility. Multivariate analyses that employ inflation and survey data, such as Kozicki and Tinsley (2012) and Mertens (2016), develop various stochastic trend specifications and exploit cointegration-type restrictions between actual and expected inflation. Such studies frequently suggest that the evolution of long-horizon inflation expectations is related to the credibility of the commitment of monetary authorities to low inflation.

Under our theory, reputation is a martingale within a regime from the standpoint of the private sector as the likelihood of observing an inflation realization of π_t is $\rho_t g(\pi_t|a_t) + (1-\rho_t)g(\pi_t|\alpha_t)$ so that it is direct that $\int b(\pi; a_t, \alpha_t, \rho_t)[\rho_t g(\pi|a_t) + (1-\rho_t)g(\pi|\alpha_t)]d\pi = \rho_t$.

Since reputation is a martingale, long-horizon forecasts will inherit this property: our theory thus provides an explanation of the stochastic trend in empirical studies and, in line with the conjectures in some of these studies, it is linked to changes in the private sector beliefs about policy commitment to low inflation.[57]

Further, our theory predicts that a stochastic inflation trend should have "stochastic volatility" since $f_{\infty|t+1} - f_{\infty|t}$ is proportional to $\rho_{t+1} - \rho_t$. As shown in (18) above, there is a time-varying link between changes in reputation and inflation surprises – relative to private sector nowcasts – captured by $\rho_t(1-\rho_t)[a_t - \alpha_t]$.[58] On this basis, volatility should be larger when (i) reputation is away from its boundaries and (ii) the policy difference is greater in absolute value.

[57] In contrast to some other explanations of inflation trends – such as the shifting long-run inflation target in Ireland (1999), Smets and Wouters (2003) and Erceg and Levin (2003), it arises endogenously in macroeconomic equilibrium with a purposeful monetary authority.

[58] Specifically, the trend volatility is time-varying even when the distribution of policy implementation errors v_π is time invariant.

7 Summary, Conclusions and Final Remarks

We characterize the recursive equilibrium of a dynamic 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. In the game, 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.

Harnessing the insights of modern contract theory, we develop a computable recursive equilibrium, where the equilibrium policies of both policymaker types and the rational expectations of private agents are shown to be functions of only three state variables, including an important reputation state that captures the evolution of private agents' beliefs about the commitment capacity of current policymaker. Using SPF short-term inflation forecasts data as observables, we extract latent states of the model via a nonlinear filter based on our theoretical model's dynamic system. The model-implied inflation 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 results reveal that evolving reputation is central to understanding the interplay between inflation expectation and policy. In particular, without the incentives of rebuilding reputation, a switch to committed policy regime in 1981 cannot account for the observed post-1981 Volcker disinflation. Moreover, long-term inflation forecasts depend on reputation that evolves through Bayesian updating of inflation forecast errors. Our model implies a nonlinear relation between reputation and the sensitivity of long-term inflation forecast to forecast errors, which is supported by empirical evidence from regressions of SPF long-term forecast revision on nowcast forecast errors.

Our model is deliberately stark. But it yields results that have surprised us and others. We believe its success in matching U.S. experience of inflation and inflation expectations indicates great promise to further research on leveraging this theory framework to guide how to manage expectations in optimal policy.

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.

Beaudry, Paul, Chenyu Hou, and Franck Portier, "The Dominant Role of Expectations and Broad-Based Supply Shocks in Driving Inflation," in "NBER Macroeconomics Annual 2024, volume 39," University of Chicago Press, 2024.

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 S., Michael Ehrmann, Jakob de Haan, and David-Jan Jansen, "Central Bank Communication with the General Public: Promise or False Hope?," *Journal of Economic Literature*, June 2024, *62* (2), 425–457.

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," 2023, *15* (1), 1–47.

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.

_ , _ , and Saten Kumar, "How Do Firms Form Their Expectations? New Survey Evidence," *American Economic Review*, September 2018, *108* (9), 2671–2713.

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.

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.

Faingold, Eduardo and Yuliy Sannikov, "Reputation in Continuous-Time Games," *Econometrica*, May 2011, *79* (3), 773–876.

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.

Ireland, Peter N., "Does the time-consistency problem explain the behavior of inflation in the United States?," *Journal of Monetary Economics*, 1999, *44* (2), 279–291.

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.

Kozicki, Sharon and P.A. Tinsley, "Effective Use of Survey Information in Estimating the Evolution of Expected Inflation," *Journal of Money, Credit and Banking*, 2012, *44* (1), 145–169.

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, Federal Reserve Bank of Minneapolis*, 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, "Measuring the Level and Uncertainty of Trend Inflation," *The Review of Economics and Statistics*, 2016, *98* (5), 950–967. Publisher: The MIT Press.

_ **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.

Smets, Frank and Raf Wouters, "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area," *Journal of the European Economic Association*, September 2003, *1* (5), 1123–1175.

Stock, James H. and Mark W. Watson, "Why Has U.S. Inflation Become Harder to Forecast?," *Journal of Money, Credit and Banking*, February 2007, *39* (s1), 3–33.

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.