Predictable Forecast Errors in Full-Information Rational Expectations Models with Regime Shifts*

Ina Hajdini
Federal Reserve Bank of Cleveland

André Kurmann
Drexel University

March 26, 2026

Abstract

This paper shows that regime shifts in Full-Information Rational Expectations (FIRE) models generate predictable, regime-dependent forecast errors. The key mechanism is that under FIRE, agents form expectations as a weighted average of regime-conditional forecasts, so that ex-post forecast errors depend systematically on the realized sequence of regimes relative to agents' expectations. Hence, even in relatively large samples, forecast error predictability alone is neither sufficient to reject FIRE nor informative about alternative expectations theories. We propose a regime-robust test of FIRE for models with regime shifts. Applying the test to an estimated New Keynesian model with regime shifts in monetary policy, we find that the baseline test rejects FIRE for the 1970s and the post-2020 period but not for intermediate subperiods. The rejection for the 1970s disappears when regime transition probabilities are re-estimated for each subperiod using observed SPF inflation expectations, thus allowing for changes in perceived monetary policy credibility. The interpretation is that agents who rationally anticipated an eventual shift to hawkish policy systematically under-predicted inflation during a dovish regime that turned out to be more persistent than expected. The post-2020 rejection persists, providing empirical motivation to consider richer regime structures or alternative theories of expectations formation for this episode.

JEL Classification: C53; E37

Keywords: Full-information Rational Expectations; Markov Regime Shifts; Forecasting Errors; Waves of Over- and Under-Reaction; Survey of Professional Forecasters.

*First version: June 6, 2021. Hajdini: Federal Reserve Bank of Cleveland, 1455 E 6th St, Cleveland, OH 44114 (email: ina.hajdini@clev.frb.org); Kurmann: Drexel University, LeBow College of Business, School of Economics, 3220 Market Street, Philadelphia, PA 19104 (email: kurmann.andre@drexel.edu). We thank Olivier Coibion, Leland Farmer, Andrew Foerster, Zhen Huo, Kyle Jurado, Robert King, Alexandre Kohlhas, Albert Marcet, Fabian Winkler as well as participants at various conferences and seminars for many helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Cleveland or the Federal Reserve System. All errors are our own.

1 Introduction

Much of modern macroeconomic research operates under the assumption that agents perfectly know the current state of the economy and form expectations rationally based on a "model-consistent" calculation of the equilibrium. One of the hallmarks of this full-information rational expectations (FIRE) hypothesis is that forecast errors are unpredictable. Yet, a growing body of research based on survey expectations data shows that ex-post forecast errors are often predictable in systematic and quantitatively important ways. This has been taken as evidence against FIRE and has sparked a burgeoning literature introducing information frictions, departures from rational expectations, or combinations thereof to explain observed forecast error patterns.[1]

In this paper, we study the predictability of ex-post forecast errors in FIRE models in the presence of regime shifts in the parameters governing the data-generating process. Such regime shifts, due, for example, to changes in the economic environment or the stance of fiscal and monetary policy, are well-documented and the focus of a large literature.[2]

The main result of our investigation is that regime shifts in FIRE models lead to predictable, regime-dependent forecast errors. Intuitively, agents under FIRE form expectations as a weighted average of regime-conditional forecasts. The regime that actually realizes ex-post therefore differs from what agents expected, resulting in a forecast error that is systematically related to information available at the time of forecast. Over a given sample, the estimated regression coefficient averages these regime-dependent forecast errors, so that its sign and magnitude reflect how the frequency of realized regimes compares to agents' expected regime persistence. If, for instance, a high-inflation regime persists longer than agents expected, the regression will indicate systematic under-prediction of inflation and under-reaction to news. Because regime shifts are infrequent, convergence of sample regime frequencies to their population counterparts can be very slow, and forecast error predictability can persist even in relatively large samples.

This result has two important implications. First, in the presence of regime shifts, predictability of forecast errors is not a sufficient condition to reject FIRE. A researcher estimating reduced-

[1] See Coibion and Gorodnichenko (2015), Angeletos, Huo, and Sastry (2020), Bordalo et al. (2020), Kohlhas and Walther (2021), and Farmer, Nakamura, and Steinsson (2023), among many others. Also see Coibion, Gorodnichenko, and Kamdar (2018) for a summary of the literature.

[2] Prominent examples include Clarida, Gali, and Gertler (2000), Leeper and Zha (2003), Stock and Watson (2002), Cogley and Sargent (2004), Lubik and Schorfheide (2004), Boivin and Giannoni (2006), Sims and Zha (2006), and Bianchi (2013). Also see Hamilton (2016) for a survey and references therein.

form regressions of ex-post forecast errors on current information may find significant non-zero coefficients even if the data have been generated under FIRE. As described above, the sign of the estimated coefficient depends on the realized regime sequence, producing waves of apparent over- and under-reaction across subperiods. Forecast error predictability vanishes only as the sample grows large enough for the distribution of regime realizations to converge to its population counterpart, but regime shifts may be too infrequent for this convergence to occur in available samples of macroeconomic forecasting data.

Second, regression estimates alone are not informative about alternative theories of expectations formation. This is because ex-post forecast errors by forward-looking agents—whether fully informed rational or not—depend on the full sequence of regime realizations, which is generally unobserved by the researcher. Furthermore, in richer models, forecast errors depend on the full state vector rather than just the predictor variable used in the regression, introducing omitted variable bias. We view this second implication as perhaps most important, since the literature has used the *sign patterns* of regression estimates—under-reaction to news and over-reaction to current realizations—to argue in favor of specific information frictions (e.g., Coibion and Gorodnichenko, 2015; Kohlhas and Walther, 2021).

Our results should be taken as neither an endorsement of FIRE nor a rejection of alternative theories of expectations formation. Indeed, there is much empirical evidence that even sophisticated market participants are subject to imperfect information and make decisions that are hard to square with the assumption of rational expectations.[3] The question, rather, is whether FIRE constitutes an appropriate approximation for average expectations and macroeconomic dynamics, or whether alternative theories of expectations formation provide a better fit with the data. This is important because such alternative theories can lead to policy prescriptions that differ substantially from those under FIRE.[4] Our results imply that to answer this question, researchers should assess FIRE against alternative expectations processes as part of fully specified structural models that incorporate plausible regime shifts. To that end, we propose a simulation-based regime-robust test that computes the probability that the forecast error predictability estimated in the data was generated from a given model under FIRE.

[3]See Tversky and Kahneman (1973), Kahneman and Tversky (1973), De Bondt and Thaler (1985), De Bondt and Thaler (1989), Adam (2007), and Malmendier and Nagel (2016), among many others. Furthermore, there is pervasive heterogeneity in the level and accuracy of forecasts across economic agents, directly contradicting the FIRE hypothesis. See Carroll (2003); Coibion, Gorodnichenko, and Kamdar (2018); Broer et al. (2021); and Weber et al. (2022) for examples.

[4]See for example Ball, Mankiw, and Reis (2005) and Paciello and Wiederholt (2014).

The rest of the paper proceeds as follows. Section 2 sets the stage by reviewing the existing empirical evidence on the predictability of forecast errors with data for US inflation from the Survey of Professional Forecasters (SPF). Beyond well-known results for the full sample that are consistent with the literature, we document large changes in forecast error predictability across subperiods: forecasters under-react to news during periods of high and volatile inflation (the 1970s and early 2020s) and over-react during periods of disinflation (the 1980s and early 1990s). These swings, which are associated with different regimes of inflation and monetary policy, present a challenge for theories of forecast error predictability based on departures from FIRE alone, as such theories imply constant over- or under-reaction.

Motivated by these findings, in Section 3 we illustrate the implications of regime shifts for forecast error predictability in a univariate FIRE model whose coefficients switch according to a Markov process. Agents have perfect information about the current state of the economy, including the realized regime, and form rational expectations about the future based on full knowledge of the environment. The simplicity of the model admits a closed-form solution of ex-post forecast errors as a function of the current state, with the sign of this relationship depending on the future regime realization. We derive the expected forecast error regression coefficient and show that the sign and magnitude of the estimates depend on the sequence of regime realizations over the sample period relative to agents' expectations. Monte Carlo simulations confirm that waves of apparent over- and under-reaction can be sizable and that convergence to the unconditional distribution is slow, exceeding the available time series of survey expectations.

Section 4 generalizes the analysis to any Markov-switching FIRE model with a minimum state variable solution, confirming that predictable forecast errors are a generic feature of any such model. The generalized framework also reveals that forecast errors depend on the full state vector and the entire regime path over the forecast horizon. Since the predictors in forecast error regressions as used in literature do not span this state vector, these regressions are generally subject to omitted variable bias. This is another reason why forecast error regression estimates do not have a structural interpretation and are therefore not informative about the underlying expectations data-generating process.

To move beyond a simple critique of forecast error regressions, Section 5 proposes a regime-robust test of FIRE. The test takes a FIRE model with regime shifts as the null, simulates the distribution of forecast error regression estimates implied by the model, and computes the sig-

nificance level at which the empirical estimates allow one to reject the null. The test is similar in spirit to simulation-based tests by Andolfatto, Hendry, and Moran (2008) and Adam, Marcet, and Beutel (2017), but takes into account not only finite sample uncertainty but also uncertainty about the data-generating process and the regime path.

We illustrate the test with a simple univariate model estimated on US inflation data. The baseline test rejects FIRE for the 1970s, but we show that this rejection reflects the assumption that agents' perceived regime transition probabilities were constant over time. Evidence on long-run inflation expectations suggests that agents during the 1970s assigned substantial probability to a return to low inflation, implying that the prevailing high-inflation regime was perceived as less persistent than full-sample estimates suggest. When transition probabilities are tilted to reflect this, the test no longer rejects FIRE.

Section 6 applies the regime-robust test to a New Keynesian model with Markov regime shifts in both the response coefficients and the inflation target of the monetary policy rule, estimated on US output growth, inflation, and the federal funds rate from 1954 to 2024. In contrast to the univariate model, this structural framework provides a monetary policy interpretation of regime shifts and allows us to investigate the role of changing monetary policy credibility for forecast error predictability. The baseline test rejects FIRE for the 1970s and the post-2020 period but not for intermediate subperiods. When we re-estimate regime transition probabilities for each subperiod using observed SPF inflation expectations—thereby providing a disciplined measure of changing monetary policy credibility—the model generates substantial forecast error predictability across subperiods and the rejection for the 1970s disappears. The post-2020 episode continues to reject, providing empirical motivation to consider richer regime structures or alternative theories of expectations formation.

The paper is related to several literatures. As reviewed in Section 2, the paper contributes to a large literature on the predictability of survey-based forecast errors of macro aggregates. The key insight of our analysis is that in the presence of regime shifts, predictable forecast errors are not a sufficient condition to reject FIRE. As already emphasized, we do not interpret this result as a critique of alternative theories of expectations formation. Our point instead is that reduced-form forecast error regressions on their own are not informative about alternative expectations theories relative to FIRE.

The result shares clear parallels with an earlier asset pricing literature on tests of the efficient

markets hypothesis in the presence of so-called peso problems, i.e., anticipated changes in the probability distribution of asset prices. See, for instance, Rietz (1988); Engel and Hamilton (1990); Cecchetti, Lam, and Mark (1993); Kaminsky (1993); Evans and Lewis (1995a, 1995b); Bekaert, Hodrick, and Marshall (2001); and Barro (2006).[5] The main differences are that we (i) show formally that the mechanism extends to any Markov-switching FIRE model with a minimum state variable solution, (ii) propose a regime-robust test that accounts for uncertainty about both the DGP and the regime path, and (iii) apply the test to an estimated New Keynesian dynamic general equilibrium model.

The paper also contributes to a recent literature that analyzes the extent to which learning in an equilibrium context can explain salient features of survey-based forecast errors of macroeconomic aggregates. Aside from the work by Andolfatto, Hendry, and Moran (2008) and Adam, Marcet, and Beutel (2017) mentioned above, the paper perhaps most closely related to ours is King and Lu (2021), who propose a model with endogenous regime shifts in monetary policy and private sector learning to account for the rise and fall in US inflation and the concomitant dynamics of inflation forecast errors in the SPF. Other related papers are Farmer, Nakamura, and Steinsson (2023), who propose a model of professional forecasters who learn about low-frequency features of the underlying data-generating process to account for various "forecast anomalies," and Andolfatto and Gomme (2003); Davig (2004); Schorfheide (2005); Bullard and Singh (2012); Richter and Throckmorton (2015); and Foerster and Matthes (2022), among others, who introduce imperfect information and learning into otherwise rational expectations DSGE models with Markov regime shifts. The distinguishing feature of our analysis is to show that even with perfectly informed rational agents, regime shifts can generate predictable forecast errors. Our changing credibility interpretation, in which agents' perceived transition probabilities vary across subperiods while FIRE is maintained within each subperiod, shares similarities with learning models but at the same time is distinct because agents still have access to full information up to and including the time of forecast and remain fully rational.

[5]The name peso problem goes back to the empirical puzzle that forward rates on the Mexican peso traded below the dollar exchange rate for much of the early 1970s even though the peso was pegged to the dollar. Then, in 1976, the peso was allowed to float and depreciated by almost 50 percent. Ex-post, the forward-spot rate difference prior to the devaluation looks like a predictable forecast error, but ex-ante it is consistent with rational expectations under the assumption of regime shifts. See Lewis (2008) for a review.

2 Empirical Evidence about Survey-Based Forecast Errors

Starting with Mincer and Zarnowitz (1969), a large literature documents that survey-based forecasts of macroeconomic aggregates are often biased and that ex-post forecast errors – the difference between actual realizations and ex-ante forecasts – are autocorrelated. See, for example, the reviews Croushore (2010) and Coibion, Gorodnichenko, and Kamdar (2018) as well as the references therein.

More recently, several prominent papers have expanded on this empirical evidence by estimating linear regressions of ex-post forecast errors on information available at the time of forecast. The two most commonly estimated regressions in the literature are

$$y_{t+h} - F_t y_{t+h} = \mu + \delta \left(F_t y_{t+h} - F_{t-1} y_{t+h} \right) + e_{t+h}, \tag{1}$$

where $y_{t+h} - F_t y_{t+h}$ denotes the ex-post forecast error about the period $t+h$ realization of some variable of interest, y_{t+h}, relative to its forecast based on information available as of period t; $F_t y_{t+h} - F_{t-1} y_{t+h}$ denotes forecast revisions reflecting news at the time of forecast; and e_{t+h} is an error term; and

$$y_{t+h} - F_t y_{t+h} = \theta + \gamma y_t + e_{t+h}, \tag{2}$$

where y_t is the current realization known at the time of forecast. Regression (1) was first proposed by Coibion and Gorodnichenko (2015) to test whether forecasters under-react or over-react to news, while regression (2) is estimated in Bordalo et al. (2020), Angeletos, Huo, and Sastry (2020), and Kohlhas and Walther (2021) to test whether forecasters under-react or over-react to current realizations.

The slope estimate $\hat{\delta}_T$ of regression (1) is typically found to be positive while the slope estimate $\hat{\gamma}_T$ of regression (2) is often found to be negative, although the significance and even the sign of this latter estimate depends somewhat on data details and the sample period considered.[6] This simultaneous under-reaction to news and over-reaction to current realizations has been widely

[6] Some studies compute forecast errors by averaging forecasts across survey participants, while other studies use individual forecasts and estimate the two regressions with individual fixed effects. The results are typically similar. Bordalo et al. (2020) and others, in turn, estimate regression (1) as a panel using individual forecast errors $y_{t+h} - F_{it} y_{t+h}$ and individual forecast revisions $F_{it} y_{t+h} - F_{it-1} y_{t+h}$. They report negative as opposed to positive estimates of δ. As Angeletos, Huo, and Sastry (2020) and Kohlhas and Walther (2021) point out, however, the sign of this estimate depends on the treatment of outliers in the individual forecast data and the sample period.

interpreted as a rejection of FIRE and has led different authors to propose new theories of expectations formation based on information rigidity (e.g., Coibion and Gorodnichenko, 2015, Angeletos, Huo, and Sastry, 2020) and asymmetric attention (e.g., Kohlhas and Walther, 2021).

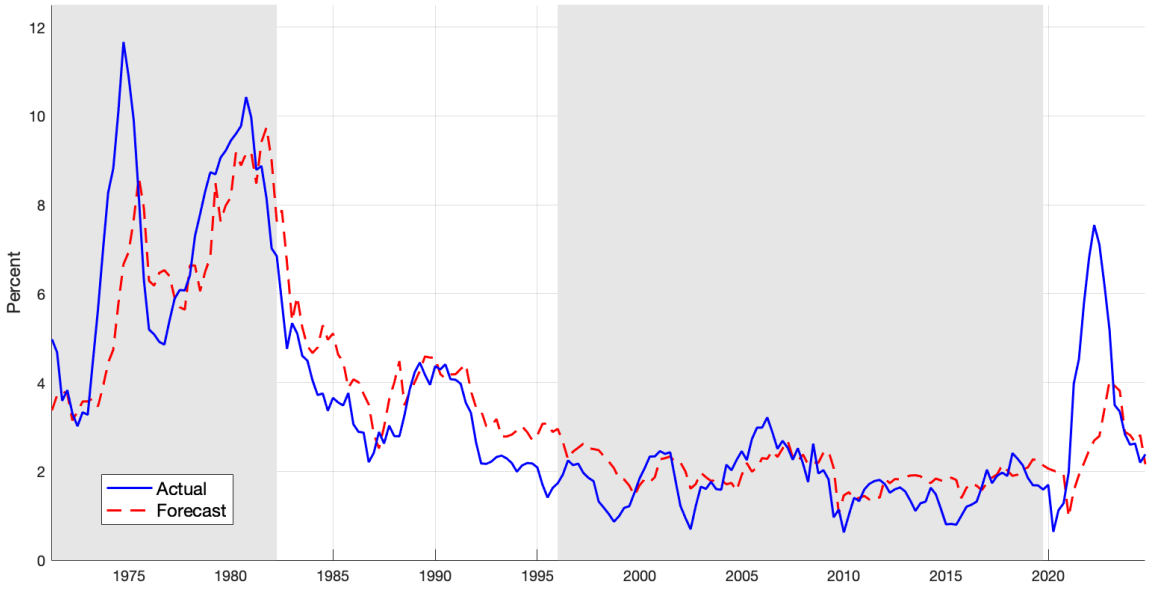
Here, we review some of this evidence and then expand on it by documenting large variations in bias and forecast error predictability across subsamples. Our main focus is on U.S. inflation because inflation expectations have received disproportionate attention in the literature and because U.S. inflation and monetary policy is characterized by distinct regimes that naturally lend themselves to our analysis. In the appendix, we provide evidence for other macro aggregates.

2.1 Data

Following Coibion and Gorodnichenko (2015) and many others, we measure inflation expectations from the Survey of Professional Forecasters (SPF), which is a quarterly survey of a rotating panel of professional forecasters. The inflation forecasts with the longest coverage pertain to the GDP/GNP deflator and are available starting in 1969:4 at horizons ranging from the current quarter to four quarters ahead. To compute current realizations and forecast errors, we use real-time data of the GDP/GNP deflator from the Real-Time Data Research Center of the Philadelphia Fed. Our measure of inflation is therefore consistent with what survey participants forecast and abstracts from reclassifications and revisions not available at the time of forecast (Croushore, 2010).

Figure 1 shows quarterly year-on-year inflation and the corresponding year-on-year mean SPF forecasts from 1971:2 to 2024:4. The shaded / non-shaded areas divide the sample into four subperiods that are motivated by our model estimations below and broadly capture different U.S. inflation and Federal Reserve policy regimes: (i) 1971:2-1982:2 – high and volatile inflation with low monetary policy credibility; (ii) 1982:3-1995:4 – declining inflation with monetary policy credibility building; (iii) 1996:1-2019:4 – low and stable inflation with established monetary policy credibility; (iv) 2020:1-2024:4 – high and volatile inflation with established monetary policy credibility.

Figure 1: U.S. Inflation and Inflation Forecasts



Notes: The graph shows quarterly year-on-year U.S. inflation based on real-time data of the GDP/GNP deflator and the corresponding year-on-year mean SPF forecasts for 1971:2 to 2024:4. The dates on the horizontal axis correspond to the quarter when inflation is realized, not the time of forecast (4 quarters earlier). The shaded / non-shaded areas correspond to the subsamples 1971:2-1982:2, 1982:3-1995:4, 1996:1-2019:4, and 2020:1-2024:4.

2.2 Bias and predictability of forecast errors

To start, we estimate regressions (1) and (2) for different forecast horizons over the full sample period 1970:2-2024:4. Instead of focusing solely on the slope coefficients as the above-cited literature does, we also report results for the constants, which provide evidence on the average bias in forecasts. For forecast horizons $h = 1, ..., 4$, we use annualized quarterly inflation rates. For year-on-year inflation, denoted $h = 1 - 4$, we combine forecasts over horizons one through four. Following the above-mentioned studies, we do not correct for finite sample bias and compute HAC-robust standard errors for inference.[7]

Panel A of Table 1 shows the results for regression (1). The slope estimate $\hat{\delta}_T$ is positive and, at least for short and long horizons, highly significant, confirming the main finding of Coibion and Gorodnichenko (2015) that forecasters under-react to news. The intercept $\hat{\mu}_T$, in turn, is close to

[7]OLS coefficient estimates are biased in finite samples if the regressors are not strictly exogenous. As discussed in Section 3, strict exogeneity is typically violated for the type of regressions in (2) and (1), independent of whether the data-generating process contains regime shifts or not. Other studies such as Adam, Marcet, and Beutel (2017) bias-correct their estimates in a related context. The regime-shift robust test that we propose in Section 5 automatically takes finite sample bias into account.

Table 1: Bias and Predictability of Inflation Forecast Errors across Forecast Horizons

	Forecast Horizon				
	1	2	3	4	$1-4$
	Panel A: $y_{t+h} - F_t y_{t+h} = \mu + \delta(F_t y_{t+h} - F_{t-1} y_{t+h}) + e_{t+h}$				
Intercept $\hat{\mu}_T$	-0.01	-0.03	-0.04	-0.07	-0.03
	(0.08)	(0.12)	(1.61)	(0.19)	(0.13)
Slope $\hat{\delta}_T$	0.49^{***}	0.48	0.82	1.10^{**}	1.10^{***}
	(0.18)	(0.38)	(0.50)	(0.42)	(0.40)
	Panel B: $y_{t+h} - F_t y_{t+h} = \theta + \gamma y_t + e_{t+h}$				
Intercept $\hat{\theta}_T$	-0.21	-0.28	-0.42	-0.42	-0.16
	(0.13)	(0.19)	(0.22)	(0.25)	(0.22)
Slope $\hat{\gamma}_T$	0.06	0.07	0.11	0.11	0.04
	(0.11)	(0.07)	(0.08)	(0.09)	(0.07)

Notes: The table reports OLS coefficient estimates of the forecast error regressions in (1)-(2) for inflation and the corresponding HAC-robust standard errors for different forecast horizons. Sample: 1970:2–2024:4. See text for details on the data construction. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

zero at all forecast horizons, indicating no discernible bias over the full sample.

Panel B reports the results for regression (2). The slope estimate $\hat{\gamma}_T$ is positive while the intercept $\hat{\theta}_T$ is negative, although neither of them is significantly different from zero. In other words, there is no clear evidence of over- or under-reaction to current realizations over the full sample.

2.3 Large variations across subperiods

As discussed in Kohlhas and Walther (2021), for shorter samples starting in the mid-1980s and inflation measures based on the consumer price index (CPI), the slope estimate of γ can be significantly negative. We now investigate the extent to which the results differ across subsamples by estimating (1) and (2) across the four subperiods described above. Following Coibion and Gorodnichenko (2015) and many others in the literature, we focus on year-on-year forecasts, although the results would be similar for the different quarterly forecasts.

Table 2 provides evidence of large differences in bias and predictability across subperiods. During the periods of high and volatile inflation of the 1970s and the early 2020s, forecasters

systematically expect lower than realized inflation as evidenced by the positive bias, under-react to news and over-react to current information. During the low-inflation periods of the 1980s and early 1990s the estimates flip: forecasters systematically expect higher than realized inflation, over-react to news and no longer over-react to current realizations. Finally, during 1990s through the end of the 2010s, forecasters still expect higher than realized inflation, although to a lesser extent, and there is no longer a significant association of forecast errors with either news or current realizations.

Table 2: Bias and Predictability of Inflation Forecast Errors across Subperiods

	Subsample			
	1970:2-1982:2	1982:3-1995:4	1996:1-2019:4	2020:1-2024:4
	Panel A: $\pi_{t+4} - F_t\pi_{t+4} = \theta + \delta(F_t\pi_{t+4} - F_{t-1}\pi_{t+4}) + e_{t+4}$			
Intercept $\hat{\theta}_T$	0.31	-0.75^{***}	-0.21^{**}	1.39
	(0.33)	(0.13)	(0.10)	(0.82)
Slope $\hat{\delta}_T$	1.01^{**}	-0.40^{**}	0.38	1.30^{**}
	(0.48)	(0.18)	(0.33)	(0.55)
	Panel B: $\pi_{t+4} - F_t\pi_{t+4} = \mu + \gamma y_t + e_{t+4}$			
Intercept $\hat{\mu}_T$	2.50^{**}	-0.65^{**}	-0.19	3.04^{***}
	(1.08)	(0.30)	(0.23)	(0.94)
Slope $\hat{\gamma}_T$	-0.28^{**}	-0.01	-0.02	-0.37^{***}
	(0.13)	(0.10)	(0.14)	(0.11)

Notes: The table reports OLS coefficient estimates of the forecast error regressions in (1)-(2) for inflation and the corresponding HAC-robust standard errors for different subsamples. The forecast horizon is always over 4 quarters. See text for details on the data construction. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

We view these changes in bias and predictability over subperiods as interesting new stylized facts. On the one hand, some of the changes could be due to small sample uncertainty. On the other hand, the association of these changes with different regimes of inflation and Fed policy making suggests that a potentially more structural explanation related to expectation formation under regime switching may be at work. We formally explore this possibility in what follows. Independent of the forces behind these results, we note that large changes in over- and under-reaction to news and current realizations presents a challenge for theories of forecast error predictability based on departures from FIRE alone, as these

theories imply constant over- or under-reaction.

3 Predictable Forecast Errors in a Univariate Model

This section considers predictability of forecast errors in a univariate FIRE model, first without regime shifts and then with regime shifts. While very simple from an empirical standpoint, the model has the advantage that the relationship between forecast errors and current information can be derived analytically and has clear intuition.

3.1 No regime shifts

Consider an endogenous variable of interest y_t governed by expectational Euler equation

$$y_t - \bar{y} = \beta \left(\mathbb{E}_t y_{t+1} - \bar{y} \right) + \psi x_t \tag{3}$$

where $\bar{y}$ is a constant mean and x_t is a stochastic driving variable that evolves according to

$$x_t = \phi x_{t-1} + \varepsilon_t, \tag{4}$$

with $\varepsilon_t \sim i.i.d.(0, \sigma^2)$.

For $|\beta| \leq 1$ and $\phi < 1$, the unique FIRE solution is

$$y_t = \bar{y} + a = \frac{\psi}{1 - \beta\phi} x_t = \bar{y} + a x_t. \tag{5}$$

Given (4) and (5), FIRE implies that for any horizon $h \geq 1$, agents' forecasts of x_{t+h} conditional on information at time t are

$$\mathbb{E}_t y_{t+h} = \bar{y} + a\phi^h x_t, \tag{6}$$

and ex-post forecast errors are

$$y_{t+h} - \mathbb{E}_t y_{t+h} = \bar{y} + a\phi^h x_t + a \sum_{\tau=1}^{h} \phi^{h-\tau} \varepsilon_{t+\tau} - \bar{y} - a\phi^h x_t = a \sum_{\tau=1}^{h} \phi^{h-\tau} \varepsilon_{t+\tau}. \tag{7}$$

In the absence of regime shifts, ex-post forecast errors under FIRE are a linear combination of i.i.d. innovations $\{\varepsilon_{t+\tau}\}_{\tau=1}^{h}$ that are mean zero and unpredictable based on current realizations, y_t, news as captured by ex-ante forecast revisions, $\mathbb{E}_t y_{t+h} - \mathbb{E}_{t-1} y_{t+h} = a\phi^h \varepsilon_t$, or any other information available at the time of forecast.

As reviewed in Section 4, in the absence of regime shifts, the same result of unpredictable forecast errors holds for any linear FIRE model. The result constitutes the starting point for the above-discussed literature documenting that ex-post forecasting errors constructed from survey data are in fact predictable, a finding that is typically taken as evidence against FIRE.

3.2 Markov regime shifts

Suppose instead that both the mean and the response coefficient to x_t of the FIRE solution in (5) switch between two regimes $s_t \in \{1, 2\}$; i.e.,

$$y_t = \bar{y}_{s_t} + a_{s_t} x_t, \tag{8}$$

where

$$\bar{y}_{s_t} = \begin{cases} \bar{y}_1 & \text{if } s_t = 1 \\ \bar{y}_2 & \text{if } s_t = 2 \end{cases} \text{ and } a_{s_t} = \begin{cases} a_1 & \text{if } s_t = 1 \\ a_2 & \text{if } s_t = 2 \end{cases}, \tag{9}$$

and regime switching is governed by an exogenous Markov process with transition matrix

$$P = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix}, \tag{10}$$

where $p_{ij} = Pr(s_t = j \mid s_{t-1} = i)$ with $0 < p_{ij} < 1$ and $\sum_{j=1}^{2} p_{ij} = 1$ for both $i, j = 1, 2$.[8] For the purpose of this illustration, two common regimes for $\bar{y}_{s_t}$ and a_{s_t} is sufficient, though the results easily generalize to separate Markov processes with more than two regimes each. Also, all results carry through if we allow for regime shifts in the persistence parameter ϕ of the exogenous variable x_t. In this section, we abstract from these generalizations to keep the example as simple as possible.

Given (8)-(10), FIRE implies that for any horizon $h \geq 1$, agents' forecasts of y_{t+h} conditional on information at time t (including regime realization s_t) are given by

$$\mathbb{E}_t y_{t+h} = \left(P^{(h)}_{s_t,1} \bar{y}_1 + P^{(h)}_{s_t,2} \bar{y}_2 \right) + \left(P^{(h)}_{s_t,1} a_1 + P^{(h)}_{s_t,2} a_2 \right) \phi^h x_t, \tag{11}$$

where $P^{(h)}_{s_t,s_{t+h}}$ is the (s_t, s_{t+h}) element of P^h. Hence, agents' expectations are a weighted average of regime-conditional forecasts: $\bar{y}_1$ and $a_1 \phi^h x_t$ if the first regime realizes in $t + h$, which occurs with probability $P^{(h)}_{s_t,1}$, respectively $\bar{y}_2$ and $a_2 \phi^h x_t$ if the second regime realizes in $t+h$, which occurs with probability $P^{(h)}_{s_t,2}$.

[8] Returning to the expectational Euler equation in (3), suppose the parameters $\{\beta, \phi\}$ take on different values across the two regimes $s_t \in \{1, 2\}$. Then under conditions described in Davig and Leeper (2007), the FIRE solution takes the form in (9) with $\boldsymbol{a} = \begin{bmatrix} a_1 & a_2 \end{bmatrix}' = (I_2 - \phi \boldsymbol{\beta} P)^{-1} \begin{bmatrix} \psi_1 & \psi_2 \end{bmatrix}'$ and $\boldsymbol{\beta} = \begin{bmatrix} \beta_1 & 0 \\ 0 & \beta_2 \end{bmatrix}$.

Based on (11), we derive the following result.

PROPOSITION 1. *Given the exogenous forcing process (??) and regime-switching model (8)-(10), ex-post forecast errors under FIRE are related to current realizations by*

$$y_{t+h} - \mathbb{E}_t y_{t+h} = \theta^{(h)}_{s_t,s_{t+h}} + \gamma^{(h)}_{s_t,s_{t+h}} y_t + \xi_{t+h}, \tag{12}$$

where $\theta^{(h)}_{s_t,s_{t+h}} \equiv (-1)^{s_{t+h}-1}(\bar{y}_1 - \bar{y}_2)\left(1 - P^{(h)}_{s_t,s_{t+h}}\right)$, $\gamma^{(h)}_{s_t,s_{t+h}} \equiv \frac{(-1)^{s_{t+h}-1}(a_1-a_2)\left(1-P^{(h)}_{s_t,s_{t+h}}\right)\phi^h}{a_{s_t}}$ *and* $\xi_{t+h} \equiv a_{s_{t+h}} \sum_{\tau=1}^{h} \phi^{h-\tau}\varepsilon_{t+\tau}$ *with* $\mathbb{E}_t\xi_{t+h} = 0$.

Proof. See Appendix A.1. □

Proposition 1 establishes that in the presence of Markov regime shifts, ex-post forecasting errors are generally different from zero and systematically related to current realization y_t even though agents have full information and are fully rational. Intuitively, and in contrast to the case without regime shifts, the portion of y_{t+h} that is part of current information ($\bar{y}_{s_{t+h}} + a_{s_{t+h}}\phi^h x_t$) differs from the agents' forecast because, as described in (11), agents' expectations at time t are a weighted average of regime-conditional forecasts.

The signs of $\theta^{(h)}_{s_t,s_{t+h}}$ and $\gamma^{(h)}_{s_t,s_{t+h}}$ change with the regime realized in $t+h$ as described by the following corollary.

COROLLARY 1. *Without loss of generality, assume from hereon that* $\bar{y}_1 > \bar{y}_2$ *and* $a_1 > a_2$. *Then,* $\theta^{(h)}_{s_t,s_{t+h}} > 0$ *and* $\gamma^{(h)}_{s_t,s_{t+h}} > 0$ *if* $s_{t+h} = 1$ *and* $\theta^{(h)}_{s_t,s_{t+h}} < 0$ *and* $\gamma^{(h)}_{s_t,s_{t+h}} < 0$ *if* $s_{t+h} = 2$.

Proof. See Appendix A.2. □

To understand this result, consider an agent forecasting y_{t+h} based on time t information, including the realized regime s_t. Since the agent puts non-zero probability on either the high regime $(\bar{y}_1, a_1)$ or the low regime $(\bar{y}_2, a_2)$ occuring in $t+h$, $\mathbb{E}_t y_{t+h}$ is a weighted average of the two regimes as given by (11). The agent therefore systematically underpredicts high regime realizations in $t+h$.

Similar to Proposition 1, we can also derive that in the presence of Markov regime shifts, ex-post forecast errors under FIRE are systematically predictable by news.

PROPOSITION 2. *Given the environment described in Proposition 1, ex-post forecast errors under FIRE are related to ex-ante forecast revisions by*

$$y_{t+h} - \mathbb{E}_t y_{t+h} = \mu^{(h)}_{s_t,s_{t+h}} + \delta^{(h)}_{s_t,s_{t+h}} \left(\mathbb{E}_t y_{t+h} - \mathbb{E}_{t-1} y_{t+h}\right) + \lambda^{(h+1)}_{s_{t-1},s_{t+h}} y_{t-1} + \xi_{t+h}, \tag{13}$$

where $\mu^{(h)}_{s_t,s_{t+h}} \equiv (-1)^{s_{t+h}-1}(\bar{y}_1 - \bar{y}_2)\left(1 - P^{(h)}_{s_t,s_{t+h}}\right) = \theta^{(h)}_{s_t,s_{t+h}}$, $\delta^{(h)}_{s_t,s_{t+h}} \equiv \frac{(-1)^{s_{t+h}-1}(a_1-a_2)\left(1-P^{(h)}_{s_t,s_{t+h}}\right)}{P^{(h)}_{s_t:}\boldsymbol{a}}$, $\lambda^{(h+1)}_{s_{t-1},s_{t+h}} \equiv \frac{(-1)^{s_{t+h}-1}(a_1-a_2)\left(1-P^{(h)}_{s_t,s_{t+h}}\right)P^{(h+1)}_{s_{t-1}:}\boldsymbol{a}}{\phi a_{s_{t-1}} P^{(h+1)}_{s_{t-1}:}\boldsymbol{a}}$, *and* ξ_{t+h} *is defined as in Proposition 1.*

Proof. See Appendix A.3. □

Following Corollary 1, the sign of $\delta^{(h)}_{s_t,s_{t+h}}$ is the same as the sign of $\gamma^{(h)}_{s_t,s_{t+h}}$ for this univariate regime shift model. As discussed in Section 4, however, this result does not necessarily hold for richer FIRE models with regime shifts. Also note that equation (13) includes a lagged term. As we shall see below, such additional variables can lead to substantial omitted variable bias in reduced-form forecast error regressions.

3.3 Implications for reduced-form forecast error regressions

Propositions 1 and 2 have two key implications for the type of reduced-form forecast error regressions considered in the empirical literature. First, in the presence of regime shifts, the expected regression coefficients from estimating (1) or (2) on data generated by the regime-switching model (8)–(10) under FIRE are generally different from zero and depend on the sequence of regime realizations over the sample period relative to agents' expectations. Second, because the econometrician does not observe the sequence of realized regimes $\{s_t\}_{t=1}^{T+h}$, the sampling distribution of regression coefficient estimates under the null of FIRE is non-standard and generally more disperse than implied by HAC-based standard errors used in the literature.

To provide intuition, consider estimating a demeaned version of regression (2), $y_{t+h} - F_t y_{t+h} = \gamma y_t + e_{t+h}$, on data generated by model (8)–(10) with $\bar{y}_1 = \bar{y}_2 = 0$, which isolates the slope coefficient from the intercept that captures average forecast bias. As described formally in Appendix B, the expected OLS estimate of the slope coefficient $\mathbb{E}\left[\hat{\gamma}_T | \{s_t\}_{t=1}^{T+h}; \Theta\right]$ *conditional on* a given sequence of regime realizations and agents' regime expectations given by transition probabilities $\Theta = [p_{11}, p_{22}]'$ can be decomposed into a bias-corrected component and a finite-sample bias. The bias-corrected component equals a weighted average of the regime-conditional coefficients $\gamma^{(h)}_{s_t,s_{t+h}}$ from Proposition 1, some of which are positive and some negative. Crucially, the sign of this weighted average depends on the conditional sample frequencies of regime realizations $f^{(h)}_{ii} \equiv \frac{1}{T}\sum_{t=1}^{T} \mathbf{1}(s_t = i | s_{t+h} = i)$ relative to agents' expected regime persistence $p^{(h)}_{ii}$. For instance, when the high-volatility regime a_1 is more persistent than expected ($f^{(h)}_{11} > p^{(h)}_{11}$), so that agents assign too high a probability to transitioning out of this regime relative to what materializes, the bias-corrected component is positive in expected value (apparent over-reaction). The opposite case ($f^{(h)}_{11} < p^{(h)}_{11}$) generates apparent under-reaction. In both cases, agents' forecasts are formed under the

belief that regime change is more likely than what the sample delivers, leading to systematic forecast errors in the direction of the less-frequently-realized regime. The finite-sample bias, in turn, is negative whenever the forcing variable is positively autocorrelated, a standard result that is amplified by regime shifts.

Three consequences follow. First, since the distribution of regression coefficients is model-dependent, forecast error regression-based tests of FIRE are necessarily model-dependent. Second, the expected value of the regression coefficient under the null of FIRE is generally different from zero in finite samples. Third, the non-standard distribution implies that researchers using HAC-based critical values are likely to reject FIRE too frequently. Note also that while our derivations assume a data-generating process under FIRE, the same results would obtain under the assumption of imperfect information or departures from rationality as long as agents form expectations over possible future regime realizations. The results therefore imply that forecast error regressions by themselves are not informative about alternative theories of expectations formation. We view this as an important point since the recent empirical literature has used estimates from forecast error regressions to argue in favor of or against specific forms of information frictions and departures from rationality (e.g., Coibion and Gorodnichenko, 2015; Angeletos, Huo, and Sastry, 2020; or Kohlhas and Walther, 2021). As we will show in Section 4, the implications are even more far-reaching in richer FIRE models, where forecast error regressions are additionally subject to omitted variable bias. All of this motivates the simulation-based regime-robust test of FIRE proposed in Section 5.

The discussion raises the natural question of how large the deviations of $\mathbb{E}[\hat{\gamma}_T|\{s_t\}_{t=1}^{T+h};\Theta]$ from zero are for typical sample sizes encountered in the forecast error literature. As shown formally in Appendix B, forecast error predictability vanishes as $T \to \infty$: with a sufficiently long sample, the conditional regime frequencies $f_{ii}^{(h)}$ converge to the transition probabilities $p_{ii}^{(h)}$, and the bias-corrected component converges to zero. However, this convergence can be very slow when regimes are persistent, since long stretches of a single regime are needed before the sample frequencies become representative. To assess the practical relevance of this point, we simulate the univariate regime-shift model 10,000 times over 1,000 periods, keeping the same sequence of regime realizations for each simulation, and compute $\mathbb{E}[\hat{\gamma}_T|\{s_t\}_{t=1}^{T+h};\Theta]$ for rolling windows of $T = 80$ quarters—comparable to the sample lengths considered in the literature. As shown in the appendix, $\mathbb{E}[\hat{\gamma}_T|\{s_t\};\Theta]$ varies substantially across different regime sequences and is on average negative due to the finite-sample bias, generating waves of apparent over- and under-reaction even under the null of FIRE. For samples of this length, the effect of a single persistent regime episode can dominate the regression estimate, making it difficult to distinguish regime-shift-induced forecast error predictability from genuine departures from FIRE.

4 Generalized Model

In this section, we show that ex-post predictability of forecast errors is a generic feature of any regime-shift FIRE model with a minimum state variable (MSV) solution. We also show that forecast errors can no longer be represented by a univariate equation. Instead, the complexity of the forecast error representation increases with that of the underlying DGP, which has several important implications.

4.1 Environment

The MSV solution to any FIRE model with regime shifts can be expressed as

$$X_t = C_{s_t} + A_{s_t} X_{t-1} + B_{s_t} \epsilon_t \tag{14}$$

where X_t is an $n_x \times 1$ vector of model variables; ϵ_t is an $n_\epsilon \times 1$ vector of innovations with $\mathbb{E}(\epsilon_t \epsilon_t') = \Sigma_\epsilon$ with Σ_ϵ being a diagonal matrix and $\mathbb{E}(\epsilon_t \epsilon_{t+h}') = \mathbf{0}_{n_\epsilon \times n_\epsilon}$ for any $h \neq 0$; and C_{s_t}, A_{s_t}, and B_{s_t} are conformable matrices that can take on $s_t \in \{1, 2\}$ different values capturing two possible regime realizations in period t that are governed by a Markov transition matrix P. Note that this formulation allows for regime shifts not only in the dynamics of the different variables but also in the variables' trends (e.g., a shift in the inflation target; output growth trend, etc.). An $n_y \times 1$ vector of observables Y_t is then mapped to the vector of model variables as follows[9]

$$Y_t = \Psi_0 + \Psi_1 X_t \tag{15}$$

Proposition 3 provides an expression for the FIRE forecast of X_{t+h}, for any forecast horizon $h > 0$.

Proposition 3. *Given the MSV solution of the model in (14), the rational expectations forecast of X_{t+h} conditional on the full information set available at time t, including the path of regime realization up to period t, is given by*

$$\mathbb{E}_t X_{t+h} = M_{t,t+h} + Q_{t,t+h} X_t, \tag{16}$$

where matrices $M_{t,t+h}$ and $Q_{t,t+h}$ depend on the regime realized in period t, the transition matrix P, the forecast horizon $h > 0$, as well as matrices A_{s_t}, B_{s_t}, and C_{s_t}.

Proof. See Appendix A.5. □

[9] Adding a vector of measurement errors in (15) would not change any of the results that follow.

Combining this expression for the FIRE forecast of X_{t+h} with the measurement equation in (15) then yields the FIRE forecast of the observables' vector,

$$\mathbb{E}_t Y_{t+h} = \Psi_0 + \Psi_1 M_{t,t+h} + \Psi_1 Q_{t,t+h} X_t. \tag{17}$$

4.2 Relationship of forecast errors with current information

Proposition 4 describes ex-post forecast errors about Y_{t+h} as a function of the vector of current realizations of the endogenous variables of the model, X_t, as well as a function of ex-ante forecast error revisions of the vector of the endogenous variables of the model, $\mathbb{E}_t X_{t+h} - \mathbb{E}_{t-1} X_{t+h}$.

PROPOSITION 4. *Given state-space representation (14)-(15) and regime sequence* $\{s_t, s_{t+1}, ..., s_{t+h}\}$*, ex-post forecast errors under FIRE for any forecasting horizon* $h \geq 1$ *can be expressed as*

$$Y_{t+h} - \mathbb{E}_t Y_{t+h} = \underbrace{\Theta_{t,t+h}}_{bias} + \underbrace{\Gamma_{t,t+h}}_{predictability} X_t + \xi_{t+h} \tag{18}$$

and

$$Y_{t+h} - E_t Y_{t+h} = \underbrace{\Omega_{t,t+h}}_{bias} + \underbrace{\Delta_{t,t+h}}_{predictability} (\mathbb{E}_t X_{t+h} - \mathbb{E}_{t-1} X_{t+h}) + \underbrace{\Lambda_{t-1,t+h}}_{predictability} X_{t-1} + \xi_{t+h}, \tag{19}$$

where $\Theta_{t,t+h}$*,* $\Omega_{t,t+h}$*,* $\Gamma_{t,t+h}$*, and* $\Lambda_{t-1,t+h}$ *depend on the ex-post realized regime path* $\{s_{t-1}, s_t, s_{t+1}, ..., s_{t+h}\}$*, the transition matrix* P*, forecasting horizon* h*, as well as matrices* A_{s_t}*,* B_{s_t}*, and* C_{s_t}*, while the error term* ξ_{t+h} *is uncorrelated with* X_t*,* $(\mathbb{E}_t X_{t+h} - \mathbb{E}_{t-1} X_{t+h})$*, or* X_{t-1}*.*

Proof. See Appendix A.5. □

Proposition 4 has three important implications. First, it confirms that forecast error bias and ex-post predictability (with respect to current information or ex-ante forecast revisions) are generic features of any FIRE model with regime shifts.

Second, comparison of Proposition 4 with Propositions 1 and 2 makes clear that the relationship of forecast errors with current information in the generalized framework differs in two key aspects from the univariate example: (i) ex-post forecast errors in the generalized framework depend on the *entire* vector X_t of available information at the time of forecast and not just on the realization of the variable that is being forecasted; and (ii) matrices $\Gamma_{t,t+h}$ and $\Delta_{t,t+h}$ linking ex-post forecast errors to current information X_t and ex-ante forecast revisions $\mathbb{E}_t X_{t+h} - \mathbb{E}_{t-1} X_{t+h}$ are contingent on the *entire* sequence of regime realizations between periods t and $t+h$, $\{s_t, s_{t+1}, ..., s_{t+h}\}$, and not only on the regimes realized in periods t and $t+h$.

Third and as alluded to previously in Section 3, Proposition 4 implies that the reduced-form forecast error regressions considered in the literature will generally be subject to omitted variable bias. Specifically, consider forecast error regressions (2) and (1) for the i^{th} variable in Y; i.e.

$$Y_{i,t+h} - F_t Y_{i,t+h} = \theta + \gamma Y_{it} + e_{t+h} \tag{20}$$

$$Y_{i,t+h} - F_t Y_{i,t+h} = \omega + \delta(F_t Y_{i,t+h} - F_{t-1} Y_{i,t+h}) + e_{t+h} \tag{21}$$

Since $\Gamma_{t,t+h}$ and $\Delta_{t,t+h}$ are generally non-diagonal matrices, Y_{it} and $F_t Y_{i,t+h} - F_{t-1} Y_{i,t+h}$ will not span all the information in X_t, respectively in $\mathbb{E}_t X_{t+h} - \mathbb{E}_{t-1} X_{t+h}$ and X_{t-1}, that fully informed rational agents use to forecast $Y_{i,t+h}$. Consequently, Y_{it} and $F_t Y_{i,t+h} - F_{t-1} Y_{i,t+h}$ will be correlated with e_{t+h}, resulting in omitted variable bias.

This last result provides a potential explanation for the instability of estimates of γ and δ when additional regressors are added, as documented in the literature. Moreover, since $\Delta_{t,t+h} \neq \Gamma_{t,t+h}$, the expected signs of the estimates of γ and δ may differ from each other. In other words, a sufficiently rich FIRE model with regime shifts may, depending on the sequence of regime realizations, generate simultaneously negative OLS estimates of γ and positive OLS estimates of δ, as, for example, Coibion and Gorodnichenko (2015), Bordalo et al. (2020), Angeletos, Huo, and Sastry (2020), and Kohlhas and Walther (2021) have found.

Corollary 2 further explores the consequences of regime shifts for three special cases.

Corollary 2. *Proposition 4 nests the following special cases:*

1. *Suppose that there are regime shifts only in the vector of constants; i.e., $C_1 \neq C_2$, but $A_1 = A_2$ and $B_1 = B_2$. Then, $\Gamma_{t,t+h} = \Delta_{t,t+h} = \Lambda_{t-1,t+h} = \mathbf{0}_{n_y \times n_x}$ and $\Theta_{t,t+h} = \Omega_{t,t+h} \neq \mathbf{0}_{n_y \times 1}$.*

2. *Now, suppose there are regime shifts only in the relationship between endogenous variables and innovations; i.e., $C_1 = C_2$ and $A_1 = A_2$, but $B_1 \neq B_2$. Then, $\Gamma_{t,t+h} = \Delta_{t,t+h} = \Lambda_{t-1,t+h} = \mathbf{0}_{n_y \times n_x}$ and $\Theta_{t,t+h} = \Omega_{t,t+h} = \mathbf{0}_{n_y \times 1}$.*

3. *Finally, suppose there are no regime shifts; i.e., $C_1 = C_2$, $A_1 = A_2$ and $B_1 = B_2$. Then, $\Gamma_{t,t+h} = \Delta_{t,t+h} = \Lambda_{t-1,t+h} = \mathbf{0}_{n_y \times n_x}$ and $\Theta_{t,t+h} = \mathbf{0}_{n_y \times 1}$.*

Proof. See Appendix A.6. □

If only the vector of constants is subject to regime shifts, then ex-post forecast errors under FIRE will be biased but not systematically related to current information X_t or to ex-ante forecast revisions $\mathbb{E}_t X_{t+h} - \mathbb{E}_{t-1} X_{t+h}$. If regimes apply only to the relationship between the model's endogenous variables

and innovations, then forecast errors under FIRE will not exhibit any ex-post predictability.[10] Finally, absent regime shifts, forecast errors under FIRE are not predictable regardless of the complexity of the underlying DGP.

The generalized results motivate a test that does not rely on the analytical distribution of regression coefficients â which, as we have shown, is non-standard and model-dependent â but instead simulates this distribution from an estimated DGP.

5 A Regime-Shift Robust Test of FIRE

The results in Sections 3 and 4 imply that in the presence of regime shifts, standard statistical tests of FIRE based on reduced-form regressions are misspecified, both because the null of unpredictability of forecast errors is violated in finite samples, and because the usual standard errors do not take into account the uncertainty implied by regime shifts. Moreover, as shown in Section 4, forecast error regressions are generally subject to omitted variable bias because the variables used as predictors do not span the information set that fully informed rational agents use. Here, we propose a new regime-shift robust test that – given an underlying data-generating process (DGP) – allows one to assess the FIRE hypothesis with the type of reduced-form regressions used in the literature. The test is similar in spirit to simulation-based tests of RE models with imperfect information by Andolfatto, Hendry, and Moran (2008) or Adam, Marcet, and Beutel (2017). Different from these tests, however, our test is applied to FIRE models with regime shifts and takes into account not only finite sample uncertainty but also uncertainty about the DGP and uncertainty about the regime path. We first describe the test procedure in general terms and then illustrate it with the univariate model from Section 3. In Section **??**, we apply the test to the empirically more relevant case of a medium-scale DSGE model.

5.1 Test procedure

Consider a FIRE model with regime shifts whose MSV solution is given by (14), with uncertainty about this DGP characterized by the posterior parameter distribution $P(\Theta|\mathcal{Z})$ estimated based on data $\mathcal{Z}$.[11] To simulate the finite sample distribution of the reduced-form regression estimates $\hat{\gamma}_T$ and $\hat{\delta}_T$ under the null of FIRE, we proceed in three steps:

[10] One can easily show that this result applies if, for instance, one considered Markov regime shifts in the variance of innovations only.

[11] Given that under the null, there exists an invariant DGP, the data $\mathcal{Z}$ used to estimate $P(\Theta|\mathcal{Z})$ may cover a larger period than the sample $\mathcal{Y}_T$ used to estimate the reduced-form regressions. These data may also include different variables than the ones used to estimate the reduced-form regressions.

1. Draw $n = 1, ..., N$ parameter vectors Θ^n from $P(\Theta|\mathcal{Z})$.

2. For each Θ^n, simulate $k = 1, ..., K$ samples of observations $\left\{y_t^{n,k}\right\}_{t=1}^{T+h}$ and $\left\{\mathbb{E}_t y_{t+h}^{n,k}\right\}_{t=1}^{T}$ from the DGP defined by Θ^n.

3. Estimate $\hat{\gamma}_T^{n,k}$ and $\hat{\delta}_T^{n,k}$ for each of these samples.

The resulting distributions can then be used to compute the probability that the simulated $\hat{\gamma}_T^{n,k}$, respectively $\hat{\delta}_T^{n,k}$, are larger in absolute value than the $\hat{\gamma}_T$, respectively $\hat{\delta}_T$, estimated from observed data. This provides a p-value for a t-test of the null of FIRE.

Several comments are in order about this procedure. First, by simulating artificial samples based on different parameter draws of $P(\Theta|\mathcal{Z})$, the procedure incorporates uncertainty about the DGP. Second, the simulation of artificial samples in step (ii) is non-standard because of the need to incorporate the uncertainty about the sequence of realized regimes. We do so by drawing regime realizations $\left\{s_t^{n,k}\right\}_{t=1}^{T+h}$ for each sample k from the smoothed probabilities $\hat{Pr}(s_t \mid \mathcal{Z}_T; \Theta^n)$ implied by the DGP and the data $\mathcal{Z}_T$ over the sample period $t = 1, ...T$ used to estimate the reduced-form regressions.[12] Third, the procedure naturally provides us with the finite sample distribution of regression coefficients $\hat{\gamma}_T$ and $\hat{\delta}_T$. Hence, we do not need to bias-correct the estimates (e.g., through bootstrapping), and we can conduct inference without assuming normality.

5.2 Application to the univariate model

To illustrate the test, we consider the univariate FIRE model with regime shifts from Section 3 for the case of inflation

$$\pi_t = \bar{\pi}_{s_t} + a_{s_t} x_t, \tag{22}$$

with $x_t = \phi x_{t-1} + \sigma\varepsilon_t$. We assume that $\bar{\pi}_{s_t}$ and a_{s_t} each follow and independent 2-state regime switching process with Markov transition probabilities $(p_{12}^{\bar{\pi}}, p_{21}^{\bar{\pi}})$ and (p_{12}^{a}, p_{21}^{a}).

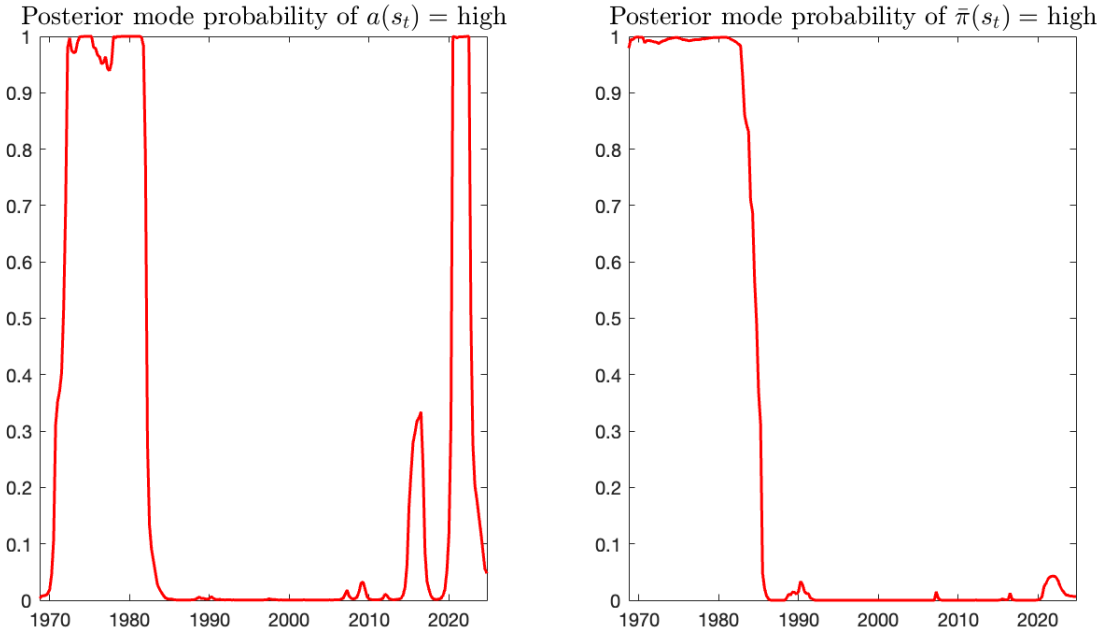
We estimate the parameters and transition probabilities on quarterly U.S. inflation data from 1947 to 2024 using standard Bayesian techniques.[13] The posterior mode estimates identify two distinct trend inflation regimes ($\bar{\pi}_1 = 1.17$, $\bar{\pi}_2 = 0.55$) and two distinct volatility regimes ($a_1 = 5.26$, $a_2 = 1.59$). Both

[12] Alternatively, for each parameter draw Θ^n and simulation run k, one could draw unconditional regime realizations using the Markov transition matrix associated with Θ^n. However, this approach would not take into account that the test we want to implement is conditional on a sequence of regime realizations associated with the sample period $t = 1, ...T$.

[13] The estimation can be implemented either by treating x_t as unobserved or by measuring x_t with an observable. Since we do not want to impose additional assumptions on the nature of x_t, we treat it as unobserved and use the Kalman and Hamilton filters as described in Kim and Nelson (1999) for estimation. See Appendix **??** for details. Note that the data strongly prefer the univariate model with regime shifts over the alternative without regime shifts.

regime processes are highly persistent, with transition probabilities out of each regime $p^a_{12} = 0.03$ and $p^a_{21} = 0.02$, respectively $p^{\bar{\pi}}_{12} = 0.02$ and $p^{\bar{\pi}}_{21} = 0.01$).[14]. As shown in Figure 2, the smoothed regime probabilities implied by these estimates and the data characterize the 1970s through the early 1980s as a high trend inflation / high volatility episode with little estimated probability of transitioning to the low regime. Thereafter, the estimation attributes high probability to low trend inflation / low volatility until 2020 when the process temporarily switches to low trend inflation / high volatility.

Figure 2: Smoothed regime probabilities

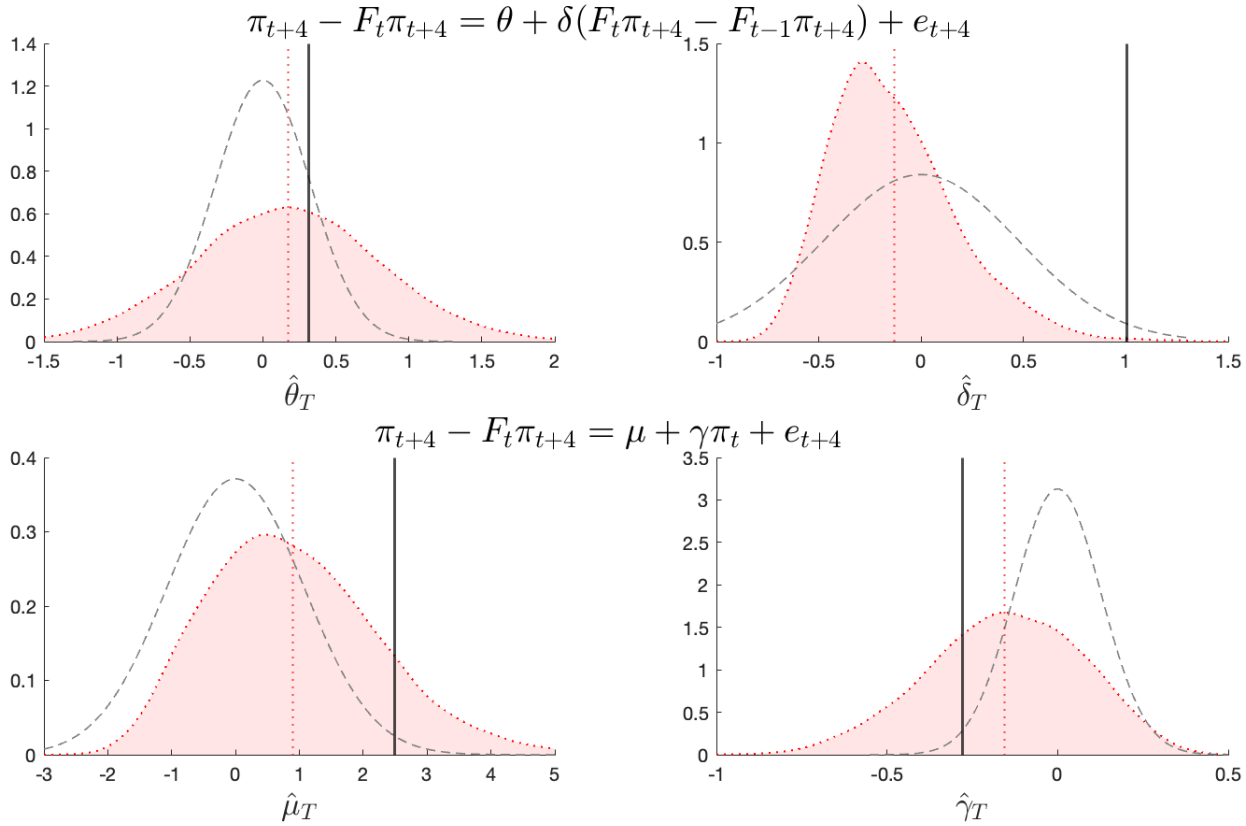


Notes: The two panels show the smoothed regime probabilities implied by the posterior mode estimates of the univariate model for $a_{s_t} = a_1 = 5.26$ (left panel) and $\bar{\pi}_{s_t} = \bar{\pi}_1 = 1.17$ (right panel) for the the sample 1970:2 - 2024:4.

Conditional on this estimated DGP, we perform the test for the forecast error regressions in Section 2, with forecast horizon set to $h = 4$. Figure 3 illustrates the test for the subsample 1970:2-1982:2 by comparing the simulated distributions of the four regression coefficients under the null of the estimated univariate FIRE model with regime shifts (red shaded areas) to the HAC-robust distributions under the standard null of zero (dashed lines). The solid black lines indicate the OLS point estimates from Table 2. Two features stand out. First, the FIRE-RS distributions are generally shifted away from zero, reflecting the non-zero expected regression coefficients implied by regime shifts and finite-sample bias. Second, in three out of four cases ($\hat{\theta}_T$, $\hat{\mu}_T$, and $\hat{\gamma}_T$), the FIRE-RS distributions are more disperse than the HAC-based distributions, illustrating the additional uncertainty that regime shifts impart on regression coefficient estimates.

[14]For the exogenous driving process the posterior mode estimates are $\phi = 0.66$ for the persistence parameter and $\sigma = 0.12$ for the innovation standard deviation.

Figure 3: Regime-shift robust test of FIRE for the 1970:2-1982:2 subperiod



Notes: The different panels show the sample distributions of the OLS estimate of forecast error regressions (1) and (2) under the usual empirical null of $\mathbb{E}[\hat{\gamma}_T] = 0$ and HAC-robust standard errors (dashed black lines) and under the null that the data were generated by the univariate FIRE model with regime shifts (red shaded area). The sample under consideration is 1970:2-1982:2.

Despite the shifts and wider distributions, the OLS estimate $\hat{\delta}_T = 1.01$ remains in the far right tail of the FIRE-RS distribution, yielding a p-value of essentially zero, while $\hat{\gamma}_T = -0.28$ is well within the FIRE-RS distribution with a p-value of 0.45. Hence, the regime-robust test rejects FIRE based on the news regression but not based on the regression on current realizations. Table 3 summarizes these results.

Monetary policy credibility and regime transition probabilities

The baseline test shows that while the FIRE-RS distributions in Figure 3 are shifted and widened relative to the standard null, they are not shifted far enough to account for the large empirical estimates — in particular, the model generates Mean$[\hat{\delta}_T^{FIRE}] = -0.13$, far from the empirical $\hat{\delta}_T = 1.01$. To understand why, note that the univariate model is estimated over the full sample 1947–2024 using realized inflation data only, under the maintained assumption that transition probabilities p_{12}^a and $p_{12}^{\bar{\pi}}$ are constant. Since both the trend-inflation and volatility regimes were highly persistent ex post, the estimated transition probabilities are small ($p_{12}^a \approx 0.03$, $p_{12}^{\bar{\pi}} \approx 0.02$). When we draw regime sequences from the smoothed

Table 3: Regime-robust test for univariate model: 1970:2-1982:2

$$\pi_{t+4} - F_t\pi_{t+4} = \theta + \delta(F_t\pi_{t+4} - F_{t-1}\pi_{t+4}) + e_{t+4}$$

$\hat{\theta}_T$	Mean$[\hat{\theta}_T^{FIRE}]$	SD$[\hat{\theta}_T^{FIRE}]$	$p(FIRE)$	$\hat{\delta}_T$	Mean$[\hat{\delta}_T^{FIRE}]$	SD$[\hat{\delta}_T^{FIRE}]$	$p(FIRE)$
0.31	0.17	0.63	**0.82**	1.01	−0.13	0.31	**0.00**

$$\pi_{t+4} - F_t\pi_{t+4} = \mu + \gamma\pi_t + e_{t+4}$$

$\hat{\mu}_T$	Mean$[\hat{\mu}_T^{FIRE}]$	SD$[\hat{\mu}_T^{FIRE}]$	$p(FIRE)$	$\hat{\gamma}_T$	Mean$[\hat{\gamma}_T^{FIRE}]$	SD$[\hat{\gamma}_T^{FIRE}]$	$p(FIRE)$
2.49	0.90	1.33	**0.23**	−0.28	−0.16	0.23	**0.60**

Notes: The table reports the results of the regime-robust test of FIRE for the univariate model (22) for inflation over the 1970:2-1982:2 subsample. See text for details.

probabilities for the 1970:2–1982:2 subsample, the simulated conditional frequencies $f_{11}^{(h)}$ and $f_{22}^{(h)}$ are therefore close to the agents' expected persistence $p_{11}^{(h)}$ and $p_{22}^{(h)}$, leaving little room for forecast error predictability under the null. This inference, however, rests on the assumption of constant transition probabilities — an assumption that deserves scrutiny.

In Markov regime-switching models, transition probabilities play a dual role: they govern the stochastic process for regime realizations, and they describe agents' expectations about future regime transitions based on current information. Agents' long-run inflation forecast, for instance, equals the ergodic mean $\mathbb{E}_t\pi_{t+\infty} = \frac{p_{21}}{p_{12}+p_{21}}\bar{\pi}_1 + \frac{p_{12}}{p_{12}+p_{21}}\bar{\pi}_2$, which is a function of p_{12} and p_{21} alone. Under the standard assumption that these probabilities are constant, agents' long-run expectations are therefore invariant — the only thing that changes over time are regime *realizations*. This is clearly counterfactual. Long-run inflation expectations varied substantially over the postwar period, as documented by, among others, Kozicki and Tinsley (2001) and Groen and Middeldorp (2013). In the regime-switching framework, time-varying long-run expectations imply that agents' perceived transition probabilities changed as they obtained new information about the conduct and credibility of monetary policy.[15]

Consider the 1970s specifically. Agents observing high-$\bar{\pi}$ / high-a regime realizations likely expected the economy to return to a low-inflation regime with substantial probability. Evidence from Adrian, Crump, and Moench (2013) and Groen and Middeldorp (2013) shows that long-run inflation expectations and inflation term premia remained relatively low throughout most of the 1970s, well below realized inflation — consistent with agents assigning substantial weight to an eventual return to low inflation. In

[15]The idea that agents' regime expectations evolve with credibility is related to the learning literature on regime change (e.g., Andolfatto and Gomme, 2003; Foerster and Matthes, 2022). Our approach differs in that we allow transition probabilities to vary across subperiods but maintain FIRE within each subperiod.

terms of the model, this suggests that agents' perceived $p_{12}^{\bar{\pi}}$ and p_{12}^{a} during the 1970s were considerably higher than the full-sample posterior estimates of 0.02 and 0.03. Ex post, however, agents' expectations of a return to low inflation turned out to be wrong: the economy remained in the high-$\bar{\pi}$ / high-a regime throughout the 1970s and into the early 1980s. The realized conditional frequencies therefore satisfy $f_{12}^{\bar{\pi}} < p_{12}^{\bar{\pi}}$ and $f_{12}^{a} < p_{12}^{a}$, or equivalently, the high-inflation regime was more persistent than agents expected. By the analytical results in Section 3 (see Appendix **??**), this gap between perceived transition probabilities and realized regime frequencies generates under-prediction of average inflation and under-reaction to news — precisely the patterns observed in the data for this subperiod.

To illustrate the quantitative impact of this credibility channel, we re-run the regime-robust test for the 1970:2–1982:2 subsample with tilted transition probabilities $p_{12}^{a} = p_{12}^{\bar{\pi}} = 0.33$, keeping all other parameters at their baseline posterior mode estimates. The choice of $p_{12}^{a} = p_{12}^{\bar{\pi}} = 0.33$ is admittedly arbitrary. In Section 6, we provide a more disciplined approach based on a structural New Keynesian model where SPF long-run inflation expectations data are used to re-estimate the transition probabilities for each subperiod.

Table 4 reports the results. The model now generates substantially stronger forecast error predictability: Mean$[\hat{\delta}_T^{FIRE}]$ rises from -0.13 to 0.29, and the p-value of the news regression increases from essentially 0 to 0.18. The model also generates positive bias (Mean$[\hat{\theta}_T^{FIRE}] = 1.43$) and negative Mean$[\hat{\gamma}_T^{FIRE}] = -0.07$, both consistent with under-prediction of inflation and over-reaction to current realizations during this subperiod. The test now fails to reject the null of FIRE for any of the regression coefficients.

Table 4: Regime-robust test with tilted transition probabilities: 1970:2-1982:2

$$\pi_{t+4} - F_t\pi_{t+4} = \theta + \delta(F_t\pi_{t+4} - F_{t-1}\pi_{t+4}) + e_{t+4}$$

$\hat{\theta}_T$	Mean$[\hat{\theta}_T^{FIRE}]$	SD$[\hat{\theta}_T^{FIRE}]$	$p(FIRE)$	$\hat{\delta}_T$	Mean$[\hat{\delta}_T^{FIRE}]$	SD$[\hat{\delta}_T^{FIRE}]$	$p(FIRE)$
0.31	1.43	0.72	**0.12**	1.01	0.29	0.56	**0.18**

$$\pi_{t+4} - F_t\pi_{t+4} = \mu + \gamma\pi_t + e_{t+4}$$

$\hat{\mu}_T$	Mean$[\hat{\mu}_T^{FIRE}]$	SD$[\hat{\mu}_T^{FIRE}]$	$p(FIRE)$	$\hat{\gamma}_T$	Mean$[\hat{\gamma}_T^{FIRE}]$	SD$[\hat{\gamma}_T^{FIRE}]$	$p(FIRE)$
2.494	1.84	1.39	**0.65**	-0.28	-0.07	0.24	**0.38**

Notes: The table reports the results of the regime-robust test of FIRE for the 1970:2-1982:2 subsample with tilted transition probabilities $p_{12}^{a} = p_{12}^{\bar{\pi}} = 0.33$. All other parameters are set at their baseline posterior mode estimates. See text for details.

6 Application with a New Keynesian FIRE-RS model

We finish by assessing the extent to which a New Keynesian model with regime shifts that imposes FIRE and fits macroeconomic dynamics reasonably well is quantitatively consistent with the reduced-form evidence on the predictability of forecast errors discussed in Section 2. In contrast to the univariate regime-shift model for inflation in (22), which lacks structural interpretation, the model considered here provides a structural monetary policy interpretation of regime shifts and allows us to investigate the role of changing monetary policy credibility for forecast error predictability.

The model we consider is based on Ireland (2007), augmented with Markov regime shifts in the central bank's inflation target and the interest rate response coefficients to current economic conditions. We estimate the model with US macroeconomic aggregates using Bayesian likelihood-based techniques and then perform the regime-shift robust test of FIRE proposed in Section 5.

6.1 Model

The economy is populated by a representative household, labor unions, intermediate firms, a final goods producer, and a monetary policy authority. The household maximizes expected lifetime utility over consumption and labor effort, subject to external habit formation in consumption and a standard intertemporal budget constraint. Firms produce differentiated goods using labor and face quadratic price adjustment costs à la **?**, with partial indexation to both lagged inflation and the central bank's inflation target. Technology is labor-only, with productivity following a unit root process. The economy is subject to four structural shocks: a preference shock, a cost-push/markup shock, a technology shock, and a monetary policy shock. We refer the reader to Appendix C for the full set of model equations and instead focus here on the monetary policy rule, which is the key ingredient for our analysis.

The central bank sets the nominal interest rate R_t according to a partial-adjustment rule with a time-varying inflation target:

$$R_t = \phi^r_{s^r_t} R_{t-1} + (1 - \phi^r_{s^r_t}) \left[\phi^\pi (\pi_t - \bar{\pi}_{s^\pi_t}) + \phi^{\Delta y}_{s^r_t} \Delta y_t \right] + v_t, \tag{23}$$

where R_t is the nominal short-term interest rate, π_t is inflation, $\bar{\pi}_{s^\pi_t}$ is the regime-dependent inflation target, Δy_t is real output growth, and v_t is an exogenous monetary policy shock.

Two features of this rule are worth highlighting. First, the interest rate smoothing coefficient $\phi^r_{s^r_t}$ and the response to output growth $\phi^{\Delta y}_{s^r_t}$ are governed by a two-state Markov process $s^r_t \in \{1, 2\}$ with transition matrix P^r, while the inflation target $\bar{\pi}_{s^\pi_t}$ is governed by a *separate* two-state Markov process $s^\pi_t \in \{1, 2\}$

with transition matrix P^{π}. Second, to obtain identification in the estimation, the response coefficient on the inflation gap ϕ^{π} is *constant* across regimes. All regime shifts in monetary policy thus operate through the degree of interest rate smoothing, the responsiveness to output growth, and the level of the inflation target—but not through the central bank's response to the inflation gap itself.

6.2 Model solution and estimation

We solve the model under FIRE and estimate the parameters, including the Markov transition probabilities, with Bayesian likelihood-based techniques using quarterly US data for output growth, inflation, and the federal funds rate from 1954:3 to 2024:4.[16] See Appendix C for details on the mapping from model variables to observables and the estimation procedure.

Table 5 reports the posterior mode estimates for the monetary policy parameters and regime transition probabilities.

Table 5: Posterior mode estimates of monetary policy rule

Dovish / high target regime	ϕ_1^r	$\phi_1^{\Delta y}$	ϕ^{π}	$\bar{\pi}_1$	$p_{12}^{\phi^r,\phi^{\Delta y}}$	$p_{12}^{\bar{\pi}}$
	0.75	0.60	1.56	1.77	0.02	0.05
Hawkish / low target regime	ϕ_2^r	$\phi_2^{\Delta y}$	ϕ^{π}	$\bar{\pi}_2$	$p_{21}^{\phi^r,\phi^{\Delta y}}$	$p_{21}^{\bar{\pi}}$
	0.48	0.10	1.56	0.72	0.03	0.02

Notes: The table reports posterior mode estimates of the monetary policy rule parameters in (23). The interest rate smoothing ϕ_i^r and output growth response $\phi_i^{\Delta y}$ are governed by a Markov process with transition probability p_{12}^r (from regime i to the other), while the inflation target $\bar{\pi}_i$ is governed by a separate Markov process with transition probability p_{12}^{π}. The inflation gap response ϕ^{π} is constant across regimes.

The estimates identify two distinct monetary policy regimes. The "dovish" regime (regime 1) features a high inflation target ($\bar{\pi}_1 = 1.77$), substantial interest rate smoothing ($\phi_1^r = 0.75$), and a large response to output growth ($\phi_1^{\Delta y} = 0.60$). Because of the high degree of smoothing, the *effective* response of the interest rate to the inflation gap is only $(1-\phi_1^r)\phi^{\pi} = 0.25 \times 1.56 = 0.39$, while the effective response to output growth is $(1-\phi_1^r)\phi_1^{\Delta y} = 0.25 \times 0.60 = 0.15$. The "hawkish" regime (regime 2), by contrast, has a low inflation target ($\bar{\pi}_2 = 0.72$), less interest rate smoothing ($\phi_2^r = 0.48$), and a small response to output growth ($\phi_2^{\Delta y} = 0.10$). As a result, the effective response to the inflation gap more than doubles to $(1-\phi_2^r)\phi^{\pi} = 0.52 \times 1.56 = 0.81$, while the effective response to output growth falls to $(1-\phi_2^r)\phi_2^{\Delta y} = 0.52 \times 0.10 = 0.05$. The dovish regime thus implies considerably more volatile equilibrium inflation dynamics than the hawkish regime.

[16] For now, we refrain from using SPF or other forecast data in the estimation. However, we will return to incorporation this information below.

Both regime processes are estimated to be highly persistent, with transition probabilities of $p^r_{12} = 0.02$ and $p^\pi_{12} = 0.05$ for the dovish-to-hawkish transition, and $p^r_{21} = 0.03$ and $p^\pi_{21} = 0.02$ for the reverse.

6.3 Regime-shift robust test of FIRE

We apply the regime-shift robust test of FIRE from Section 5 to inflation forecast errors using the estimated New Keynesian model as the data-generating process. Specifically, we take the estimated model and draw $n = 1, ..., 1000$ parameter vectors Θ^n from the estimated posterior distribution $P(\Theta|\mathcal{Z})$. Then for each Θ^n, we simulate $k = 1, ..., 200$ samples of observations $\left\{\pi_t^{n,k}\right\}_{t=1}^{T+h}$ and $\left\{\mathbb{E}_t\pi_{t+h}^{n,k}\right\}_{t=1}^{T}$ conditional on realized regimes drawn from smoothed probabilities $\hat{Pr}(s_t \mid \Theta^n, \mathcal{Z}_T)$, and estimate the forecast error regression coefficients for each of these samples. The resulting distributions are used to compute p-values for the null that the empirical estimates were generated by the model under FIRE.

Table 6: Test of New Keynesian FIRE-RS model

$$\pi_{t+4} - F_t\pi_{t+4} = \theta + \delta(F_t\pi_{t+4} - F_{t-1}\pi_{t+4}) + e_{t+4}$$

	$\hat{\theta}_T$	Mean$[\hat{\theta}_T^{FIRE}]$	SD$[\hat{\theta}_T^{FIRE}]$	$p(FIRE)$	$\hat{\delta}_T$	Mean$[\hat{\delta}_T^{FIRE}]$	SD$[\hat{\delta}_T^{FIRE}]$	$p(FIRE)$
1970:2-1982:2	0.31	-0.04	0.35	**0.32**	1.01	-0.24	0.15	**0.00**
1982:3-1995:4	-0.75	-0.19	0.39	**0.15**	-0.40	-0.15	0.46	**0.57**
1996:1-2019:4	-0.21	-0.19	0.21	**0.90**	0.38	-0.09	0.29	**0.10**
2020:1-2024:4	1.39	-0.20	0.72	**0.03**	1.30	-0.39	0.75	**0.03**

$$\pi_{t+4} - F_t\pi_{t+4} = \mu + \gamma\pi_t + e_{t+4}$$

	$\hat{\mu}_T$	Mean$[\hat{\mu}_T^{FIRE}]$	SD$[\hat{\mu}_T^{FIRE}]$	$p(FIRE)$	$\hat{\gamma}_T$	Mean$[\hat{\gamma}_T^{FIRE}]$	SD$[\hat{\gamma}_T^{FIRE}]$	$p(FIRE)$
1970:2-1982:2	2.49	1.73	1.01	**0.44**	-0.28	-0.44	0.21	**0.45**
1982:3-1995:4	-0.65	0.25	0.83	**0.27**	-0.01	-0.14	0.22	**0.55**
1996:1-2019:4	-0.19	0.10	0.57	**0.62**	-0.02	-0.09	0.17	**0.69**
2020:1-2024:4	3.04	0.99	1.71	**0.20**	-0.36	-0.40	0.44	**0.94**

Notes: The table reports the empirical OLS estimates of the forecast error regressions in (1)–(2) for inflation across four subperiods, together with the mean and standard deviation of the model-implied distribution of regression coefficients under FIRE, and the associated p-value $p(FIRE)$. The model-implied distribution is based on 200,000 simulated samples from the estimated New Keynesian model evaluated at the posterior mode. See Section 5.1 for details.

Table 6 reports the results for each of the four subperiods from Section 2. For the forecast error regression on forecast revisions (1), the baseline test rejects FIRE strongly in the 1970:2–1982:2 subperiod: $\hat{\delta}_T = 1.01$ falls far outside the model-implied distribution, which has $\mathbb{E}[\hat{\delta}_T^{FIRE}] = -0.24$ and $p(FIRE) = 0.00$. The test also rejects for the 2020:1–2024:4 subperiod, where $\hat{\delta}_T = 1.30$ yields $p(FIRE) = 0.03$. The

two intermediate subperiods, by contrast, are broadly consistent with the model. For the forecast error regression on current inflation (2), the results are more favorable: the model cannot be rejected for any of the subperiods.

The rejection for the forecast error regression on forecast revisions (1) echoes the finding for the univariate model in Section 5.2 and has the same root cause. Because the full-sample estimated regime transition probabilities are very small (Table 5), the model implies that agents expected both the dovish reaction-coefficient regime and the high inflation-target regime to be highly persistent. Ex post, this expectation turns out to be approximately correct: regimes did persist. But agents' expectations about regime persistence are identified from the *full* historical sample, including the eventual regime change in the early 1980s. Based on the information available *during* the 1970s, agents—including the central bank—may well have perceived regime change as considerably more likely, consistent with the relatively low long-run inflation expectations and inflation risk premia that prevailed throughout most of that decade.

To formalize this idea, we re-estimate the regime transition probabilities for each subperiod using observed SPF inflation expectations, keeping all other model parameters fixed at their full-sample posterior mode estimates. This yields substantially different transition probabilities for the 1970:2–1982:2 subperiod: the probability of transitioning from the dovish to the hawkish regime rises from $p_{12}^{\phi^r,\phi^{\Delta y}} = 0.02$ to 0.34 for the reaction-coefficient process and from $p_{12}^{\bar{\pi}} = 0.05$ to 0.37 for the inflation-target process, while the reverse transition probabilities remain small at $p_{21}^{\phi^r,\phi^{\Delta y}} = 0.02$ and $p_{21}^{\bar{\pi}} = 0.01$. In other words, during the 1970s, agents perceived the dovish regime as much less persistent than what the full-sample estimates imply—consistent with substantial credibility of the central bank's commitment to eventually bring down inflation.

Table 7 reports the test results with these re-estimated transition probabilities. The improvement for the 1970:2–1982:2 subperiod is striking: for the forecast revision regression, $\mathbb{E}[\hat{\delta}_T^{FIRE}]$ shifts from -0.24 to 0.52 and $p(FIRE)$ rises from 0.00 to 0.24, so that the test no longer rejects FIRE. Intuitively, agents who expected the dovish regime to end soon would have systematically under-predicted inflation and under-reacted to news during the 1970s—precisely the pattern observed in the data—not because they violated FIRE but because the regime turned out to be more persistent than they rationally anticipated. Results for the other subperiods are broadly similar to the baseline, confirming that the credibility channel primarily matters for the 1970s. The 2020:1–2024:4 subperiod continues to reject, consistent with the possibility that the post-pandemic inflation surge involved mechanisms beyond what our model captures.

The application with the New Keynesian FIRE-RS model yields two main lessons. First, the baseline regime-shift robust test rejects FIRE for subperiods characterized by major regime transitions—the

Table 7: Test of New Keynesian FIRE-RS model with re-estimated probabilities

$$\pi_{t+4} - F_t\pi_{t+4} = \theta + \delta(F_t\pi_{t+4} - F_{t-1}\pi_{t+4}) + e_{t+4}$$

	$\hat{\omega}_T$	Mean$[\hat{\omega}_T^{FIRE}]$	SD$[\hat{\omega}_T^{FIRE}]$	$p(FIRE)$	$\hat{\delta}_T$	Mean$[\hat{\delta}_T^{FIRE}]$	SD$[\hat{\delta}_T^{FIRE}]$	$p(FIRE)$
1970:2-1982:2	0.31	0.64	0.36	**0.37**	1.01	0.52	0.42	**0.24**
1982:3-1995:4	-0.75	-0.09	0.40	**0.10**	-0.40	-0.12	0.49	**0.56**
1996:1-2019:4	-0.21	-0.09	0.28	**0.68**	0.38	-0.10	0.37	**0.18**
2020:1-2024:4	1.39	-0.09	0.73	**0.04**	1.30	-0.39	0.79	**0.04**

$$\pi_{t+4} - F_t\pi_{t+4} = \theta + \gamma\pi_t + e_{t+4}$$

	$\hat{\theta}_T$	Mean$[\hat{\theta}_T^{FIRE}]$	SD$[\hat{\theta}_T^{FIRE}]$	$p(FIRE)$	$\hat{\gamma}_T$	Mean$[\hat{\gamma}_T^{FIRE}]$	SD$[\hat{\gamma}_T^{FIRE}]$	$p(FIRE)$
1970:2-1982:2	2.49	1.35	0.84	**0.16**	-0.28	-0.15	0.17	**0.45**
1982:3-1995:4	-0.65	0.33	0.84	**0.23**	-0.01	-0.14	0.22	**0.57**
1996:1-2019:4	-0.19	0.16	0.60	**0.56**	-0.02	-0.09	0.17	**0.71**
2020:1-2024:4	3.04	1.00	1.70	**0.19**	-0.36	-0.37	0.44	**0.99**

Notes: Same as Table 6, except that the regime transition probabilities are re-estimated for each subperiod using observed SPF inflation expectations, keeping all other model parameters fixed at their full-sample posterior mode. See text for details.

1970:2–1982:2 and 2020:1–2024:4 episodes—but not for the intermediate subperiods. Second, accounting for changing monetary policy credibility by re-estimating regime transition probabilities with SPF inflation expectations substantially improves the model's ability to match the forecast error evidence, especially for the 1970s. This suggests that much of what appears as systematic forecast error predictability during this period is consistent with rational agents who correctly anticipated an eventual shift to hawkish policy but experienced a dovish regime that turned out to be more persistent than expected. The post-2020 episode, by contrast, remains difficult for the model to match, providing empirical motivation to consider richer regime structures—for example, alternative theories of expectations formation.

7 Conclusion

The present paper shows that regime shifts in FIRE models lead to predictable, regime-dependent ex-post forecast errors. In general, in the presence of regime shifts, forecast errors become a complicated function of the current state of the economy and the sequence of realized regimes over the entire forecast horizon. This implies that reduced-form forecast error regressions by themselves are not informative about alternative theories of expectations formation. Furthermore, regime shifts imply that expectations exhibit waves of over- and under-reaction to current information across different subperiods. Using survey-based

forecast data of inflation and output growth constructed from the SPF, we confirm the existence of such waves.

Based on these insights, we propose a regime-shift robust test of FIRE and apply it to an estimated New Keynesian model with regime shifts in the monetary policy interest rate rule and the inflation target, estimated on US output growth, inflation, and interest rate data from 1954 to 2024. The baseline test rejects FIRE for the 1970:2–1982:2 subperiod, during which the dovish monetary policy regime persisted longer than agents—given the full-sample estimated transition probabilities—would have expected. We show, however, that this rejection is driven by the assumption that agents' perceived regime transition probabilities remained constant over the entire sample. When we re-estimate these probabilities for each subperiod using observed SPF inflation expectations—thereby allowing for changes in perceived monetary policy credibility—the test no longer rejects FIRE for the 1970s. The intuition is that agents who rationally anticipated an eventual shift to hawkish policy would have systematically under-predicted inflation during a dovish regime that turned out to be more persistent than expected, generating precisely the pattern of forecast error predictability observed in the data.

This should be taken as neither an endorsement of FIRE nor a dismissal of alternative theories of expectations formation. The test continues to reject for the 2020:1–2024:4 subperiod, and the credibility interpretation we develop is specific to regime shifts in monetary policy. We view our results as providing a new empirical motivation to consider richer data-generating processes as well as alternative theories of expectations formation that may complement regime shifts in accounting for the dynamics of forecast errors across the full postwar sample.

References

Adam, Klaus. 2007. "Experimental Evidence on the Persistence of Output and Inflation." *The Economic Journal* 117 (520):603–636. URL https://www.jstor.org/stable/4625509. 3

Adam, Klaus, Albert Marcet, and Johannes Beutel. 2017. "Stock Price Booms and Expected Capital Gains." *American Economic Review* 107 (8):2352–2408. URL https://doi.org/10.1257/aer.20140205. 1, 7, 5, 21

Adrian, Tobias, Richard K. Crump, and Emanuel Moench. 2013. "Pricing the Term Structure with Linear Regressions." *Journal of Financial Economics* 110 (1):110–138. 5.2

Andolfatto, David and Paul Gomme. 2003. "Monetary Policy Regimes and Beliefs." *International Economic Review* 44 (1):1–30. URL https://doi.org/10.1111/1468-2354.t01-1-00061. 1, 15

Andolfatto, David, Scott Hendry, and Kevin Moran. 2008. "Are inflation expectations rational?" *Journal of Monetary Economics* 55 (2):406–422. URL https://doi.org/10.1016/j.jmoneco.2007.07.004. 1, 5

Angeletos, George-Marios, Zhen Huo, and Karthik A. Sastry. 2020. "Imperfect Macroeconomic Expectations: Evidence and Theory." *NBER Macroeconomics Annual* 35:1–86. URL https://doi.org/10.1086/712313. 1, 2, 6, 3.3, 4.2

Ball, Laurence, N. Gregory Mankiw, and Ricardo Reis. 2005. "Monetary Policy for Inattentive Economies." *Journal of Monetary Economics* 52 (4):703–25. URL https://doi.org/10.1016/j.jmoneco.2005.03.002. 4

Barro, Robert J. 2006. "Rare disasters and asset markets in the 20th century." *Quarterly Journal of Economics* 121 (3):823–866. URL https://doi.org/10.1162/qjec.121.3.823. 1

Bekaert, Geert, Robert J. Hodrick, and David A. Marshall. 2001. "Peso problem explanations for term structure anomalies." *Journal of Monetary Economics* 48 (2):241 – 270. URL https://doi.org/10.1016/S0304-3932(01)00075-7. 1

Bianchi, Francesco. 2013. "Regime Switches, Agents' Beliefs, and Post-World War II US Macroeconomic Dynamics." *The Review of Economic Studies* 80 (2):463–490. URL https://doi.org/10.1093/restud/rds032. 2

Boivin, Jean and Marc P. Giannoni. 2006. "Has monetary Policy Become More Effective?" *Review of Economics and Statistics* 88 (3):445–462. URL https://doi.org/10.1162/rest.88.3.445. 2

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. 2020. "Overreaction in Macroeconomic Expectations." *American Economic Review* 110 (9):2748–2782. URL https://doi.org/10.1257/aer.20181219. 1, 2, 6, 4.2

Broer, Tobias, Alexandre Kohlhas, Kurt Mitman, and Kathrin Schlafmann. 2021. "Information and Wealth Heterogeneity in the Macroeconomy." Tech. Rep. 15934, CEPR Press. URL https://cepr.org/publications/dp15934. 3

Bullard, James and Aarti Singh. 2012. "Learning and the Great Moderation." *International Economic Review* 53 (2):375–397. URL https://doi.org/10.1111/j.1468-2354.2012.00685.x. 1

Carroll, Christopher D. 2003. "Macroeconomic Expectations of Households and Professional Forecasters." *Quarterly Journal of Economics* 118 (1):269–298. URL https://doi.org/10.1162/00335530360535207. 3

Cecchetti, Stephen G., Pok-sang Lam, and Nelson C. Mark. 1993. "The equity premium and the risk-free rate: Matching the moments." *Journal of Monetary Economics* 31 (1):21–45. URL https://doi.org/10.1016/0304-3932(93)90015-8. 1

Clarida, Richard, Jordi Gali, and Mark Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics* 115 (1):147–180. URL https://www.jstor.org/stable/2586937. 2

Cogley, Timothy and Thomas J. Sargent. 2004. "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII US." *Review of Economic Dynamics* 8 (2):262–302. URL https://doi.org/10.1016/j.red.2004.10.009. 2

Coibion, Olivier and Yuriy Gorodnichenko. 2015. "Information Rigidity and the Expectations Formation Process: A simple Framework and New Facts." *American Economic Review* 105 (8):2644–2678. URL https://doi.org/10.1257/aer.20110306. 1, 1, 2, 2.1, 2.2, 2.3, 3.3, 4.2

Coibion, Olivier, Yuriy Gorodnichenko, and Rupal Kamdar. 2018. "The Formation of Expectations, Inflation, and the Phillips Curve." *Journal of Economic Literature* 56(4) (4):1447–1491. URL https://doi.org/10.1257/jel.20171300. 1, 3, 2

Croushore, Dean. 2010. "An Evaluation of Inflation Forecasts from Surveys Using Real-Time Data." *B. E. Journal of Macroeconomics* 10 (1):1–32. URL https://doi.org/10.2202/1935-1690.1677. 2, 2.1

Davig, Troy and Eric M. Leeper. 2007. "Generalizing the Taylor Principle." *American Economic Review* 97 (3):607–635. URL https://doi.org/10.1257/aer.97.3.607. 8

Davig, Troy A. 2004. "Regime-switching debt and taxation." *Journal of Monetary Economics* 51 (4):837–859. URL https://api.semanticscholar.org/CorpusID:2582376. 1

De Bondt, Werner F. M. and Richard H. Thaler. 1985. "Does the Stock Market Overreact?" *Journal of Finance* 40 (3):793–805. URL https://doi.org/10.2307/2327804. 3

———. 1989. "Anomalies: A Mean-Reverting Walk Down Wall Street." *Journal of Economic Perspectives* 3 (1):189–202. URL https://doi.org/10.1257/jep.3.1.189. 3

Engel, Charles and James D. Hamilton. 1990. "Long Swings in the Dollar: Are They in the Data and Do Markets Know It?" *American Economic Review* 80 (4):689–713. URL https://www.jstor.org/stable/2006703. 1

Evans, Martin D. D. and Karen K. Lewis. 1995a. "Do expected shifts in inflation affect estimates of the long-run Fisher relation?" *Journal of Finance* 50 (1):225–253. URL https://doi.org/10.1111/j.1540-6261.1995.tb05172.x. 1

———. 1995b. "Do Long-Term Swings in the Dollar Affect Estimates of the Risk Premia?" *Review of Financial Studies* 8:225–253. URL https://doi.org/10.1093/rfs/8.3.709. 1

Farmer, Leland, Emi Nakamura, and Jón Steinsson. 2023. "Learning About the Long Run." *Journal of Political Economy* URL https://doi.org/10.1086/730207. 1, 1

Foerster, Andrew and Christian Matthes. 2022. "Learning about Regime Change." *International Economic Review* 63 (4):1829–1859. URL https://doi.org/10.1111/iere.12585. 1, 15

Groen, Jan J. J. and Menno Middeldorp. 2013. "Creating a History of U.S. Inflation Expectations." Liberty street economics, Federal Reserve Bank of New York. August 21. 5.2

Hamilton, James D. 1994. *Time Series Analysis*. Princeton University Press. B.1, 20, B.1, B.1

———. 2016. "Chapter 3 - Macroeconomic Regimes and Regime Shifts." Elsevier, 163–201. URL https://doi.org/10.1016/bs.hesmac.2016.03.004. 2

Ireland, Peter N. 2007. "Changes in the Federal Reserve's Inflation Target: Causes and Consequences." *Journal of Money, Credit and Banking* 39 (8):1851–1882. 6, C.1, C.1, C.2, C.4, C.5

Kahneman, Daniel and Amos Tversky. 1973. "On the psychology of prediction." *Psychological Review* 80:237–251. URL https://psycnet.apa.org/doi/10.1037/h0034747. 3

Kaminsky, Graciela. 1993. "Is there a peso problem? Evidence from the dollar/pound exchange rate: 1976-1987." *American Economic Review* 83 (3):450–472. URL https://www.jstor.org/stable/2117528. 1

Kim, Chang-Jin and Charles R. Nelson. 1999. *State-Space Models with Regime Switching*. The MIT Press. URL https://doi.org/10.7551/mitpress/6444.001.0001. 13

King, Robert G and Yang K Lu. 2021. "The Rise, Fall and Stabilization of U.S. Inflation: Shifting Regimes and Evolving Reputation." Working Paper 29585, National Bureau of Economic Research. URL http://www.nber.org/papers/w29585. 1

Kohlhas, Alexandre N. and Ansgar Walther. 2021. "Asymmetric Attention." *American Economic Review* 111 (9):2879–2925. URL https://doi.org/10.1257/aer.20191432. 1, 1, 2, 6, 2.3, 3.3, 4.2

Kozicki, Sharon and P. A. Tinsley. 2001. "Shifting Endpoints in the Term Structure of Interest Rates." *Journal of Monetary Economics* 47 (3):613–652. 5.2

Leeper, Eric M. and Tao Zha. 2003. "Modest Policy Interventions." *Journal of Monetary Economics* 50 (8):2383–2424. URL https://doi.org/10.1016/j.jmoneco.2003.01.002. 2

Lewis, Karen K. 2008. "Peso Problem." In *The New Palgrave Dictionary of Economics*. Palgrave Macmillan. URL https://karenkaylewis.com/wp-content/uploads/2022/08/PesoProblem_Palgrave.pdf. 5

Lubik, Thomas A. and Frank Schorfheide. 2004. "Testing for Indeterminacy: An Application to US Monetary Policy." *American Economic Review* 94 (1):190–217. URL https://doi.org/10.1257/000282804322970760. 2

Malmendier, Ulrike and Stefan Nagel. 2016. "Learning from Inflation Experiences." *The Quarterly Journal of Economics* 131 (1):53–87. URL https://doi.org/10.1093/qje/qjv037. 3

Mincer, Jacob A. and Victor Zarnowitz. 1969. "The Evaluation of Economic Forecasts." In *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*, edited by Jacob A. Mincer. NBER, 3–46. URL https://www.nber.org/system/files/chapters/c1214/c1214. 2

Paciello, Luigi and Mirko Wiederholt. 2014. "Exogenous Information, Endogenous Information, and Optimal Monetary Policy." *Review of Economic Studies* 81 (1):356–88. URL https://doi.org/10.1093/restud/rdt024. 4

Richter, Alexander W. and Nathaniel A. Throckmorton. 2015. "The Consequences of an Unknown Debt Target." *European Economic Review* 78:76–96. URL https://doi.org/10.1016/j.euroecorev.2015.05.002. 1

Rietz, Thomas A. 1988. "The equity risk premium a solution." *Journal of Monetary Economics* 22 (1):117–131. URL https://doi.org/10.1016/0304-3932(88)90172-9. 1

Schorfheide, Frank. 2005. "Learning and Monetary Policy Shifts." *Review of Economic Dynamics* 8 (2):392–419. URL https://doi.org/10.1016/j.red.2005.01.001. 1

Sims, Christopher A. and Tao Zha. 2006. "Were There Regime Switches in US Monetary Policy?" *American Economic Review* 96 (1):54–81. URL https://www.aeaweb.org/articles?id=10.1257/000282806776157678. 2

Stock, James H. and Mark W. Watson. 2002. "Has the Business Cycle Changed and Why?" *NBER Macroeconomics Annual* 17:159–218. URL https://doi.org/10.1086/ma.17.3585284. 2

Tversky, Amos and Daniel Kahneman. 1973. "Availability: A heuristic for judging frequency and probability." *Cognitive Psychology* 5:207–232. URL https://doi.org/10.1016/0010-0285(73)90033-9. 3

Weber, Michael, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion. 2022. "The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications." *Journal of Economic Perspectives* 36 (3):157–184. URL https://doi.org/10.1257/jep.36.3.157. 3

Appendix

A Proofs

A.1 Proof of Proposition 1

The forecasting errors about y_{t+h}, for any $h > 0$, are given by

$$\begin{aligned}
FE_{t,t+h} &= y_{t+h} - \mathbb{E}_t y_{t+h} \\
&= a_{s_{t+h}}\phi^h x_t - P^{(h)}_{s_t:}\boldsymbol{a}\phi^h x_t + a_{s_{t+h}}\sum_{\tau=1}^{h}\phi^{h-\tau}\varepsilon_{t+\tau} \\
&= \left(a_{s_{t+h}} - P^{(h)}_{s_t:}\boldsymbol{a}\right)\phi^h x_t + a_{s_{t+h}}\sum_{\tau=1}^{h}\phi^{h-\tau}\varepsilon_{t+\tau} \\
&= (a_{s_{t+h}} - P^{(h)}_{s_t,1}a_1 - P^{(h)}_{s_t,2}a_2)\phi^h x_t + a_{s_{t+h}}\sum_{\tau=1}^{h}\phi^{h-\tau}\varepsilon_{t+\tau} \\
&= \begin{cases} (a_1 - a_2)(1 - P^{(h)}_{s_t,1})\phi^h x_t + a_1\sum_{\tau=1}^{h}\phi^{h-\tau}\varepsilon_{t+\tau} & \text{if } s_{t+h} = 1 \\ -(a_1 - a_2)(1 - P^{(h)}_{s_t,2})\phi^h x_t + a_2\sum_{\tau=1}^{h}\phi^{h-\tau}\varepsilon_{t+\tau} & \text{if } s_{t+h} = 2 \end{cases}
\end{aligned} \tag{A.1}$$

where the last equality follows from $P^{(h)}_{s_t,1} + P^{(h)}_{s_t,2} = 1$. We show this through proof by induction. Clearly, $P^{(1)}_{s_t,1} + P^{(1)}_{s_t,2} \equiv p_{s_t,1} + p_{s_t,2} = 1$ for any $s_t \in \{1, 2\}$. Suppose that $P^{(h)}_{s_t,1} + P^{(h)}_{s_t,2} = 1$. We should prove that $P^{(h+1)}_{s_t,1} + P^{(h+1)}_{s_t,2} = 1$. We have that

$$P^{(h+1)} = PP^{(h)} = \begin{bmatrix} p_{11}P^{(h)}_{11} + p_{12}P^{(h)}_{21} & p_{11}P^{(h)}_{12} + p_{12}P^{(h)}_{22} \\ p_{21}P^{(h)}_{11} + p_{22}P^{(h)}_{21} & p_{21}P^{(h)}_{12} + p_{22}P^{(h)}_{22} \end{bmatrix} \tag{A.2}$$

Then, $p_{11}P^{(h)}_{11} + p_{12}P^{(h)}_{21} + p_{11}P^{(h)}_{12} + p_{12}P^{(h)}_{22} = p_{11} + p_{12} = 1$ and $p_{21}P^{(h)}_{11} + p_{22}P^{(h)}_{21} + p_{21}P^{(h)}_{12} + p_{22}P^{(h)}_{22} = p_{21} + p_{22} = 1$, which establishes that $P^{(h+1)}_{s_t,1} + P^{(h+1)}_{s_t,2} = 1$. Finally, we use $x_t = \frac{y_t}{a_{s_t}}$ to summarize the above expression as

$$FE_{t,t+h} = \frac{(-1)^{s_{t+h}-1}(a_1 - a_2)(1 - P^{(h)}_{s_t,s_{t+h}})\phi^h}{a_{s_t}}x_t + a_{s_{t+h}}\sum_{\tau=1}^{h}\phi^{h-\tau}\varepsilon_{t+\tau}, \tag{A.3}$$

which is 1 in the main text.

A.2 Proof of Corollary 1

Consider

$$a_{s_{t+h}} - P^{(h)}_{s_t:}\boldsymbol{a} = \begin{cases} (a_1 - a_2)(1 - P^{(h)}_{s_t,1}) & \text{if } s_{t+h} = 1 \\ -(a_1 - a_2)(1 - P^{(h)}_{s_t,2}) & \text{if } s_{t+h} = 2 \end{cases}$$

One can show via proof by induction that for $0 < p_{11}, p_{22} < 1$, it follows that $0 < P^{(h)}_{s_t,s_{t+h}} < 1$ for any $s_t, s_{t+h} \in \{1, 2\}$. Clearly, $1 - p_{s_t,s_{t+h}} > 0$ for any $s_t \in \{1, 2\}$. Suppose that $1 - P^{(h)}_{s_t,s_{t+h}} > 0$. We should prove that $1 - P^{(h+1)}_{s_t,s_{t+h}} > 0$. We have that

$$P^{(h+1)} = PP^{(h)} = \begin{bmatrix} p_{11}P^{(h)}_{11} + p_{12}P^{(h)}_{21} & p_{11}P^{(h)}_{12} + p_{12}P^{(h)}_{22} \\ p_{21}P^{(h)}_{11} + p_{22}P^{(h)}_{21} & p_{21}P^{(h)}_{12} + p_{22}P^{(h)}_{22} \end{bmatrix} \tag{A.4}$$

Then,

$$1-p_{11}P^{(h)}_{11}-p_{12}P^{(h)}_{21} = 1-p_{11}(P^{(h)}_{11}-P^{(h)}_{21})-P^{(h)}_{21} = P^{(h)}_{22}-p_{11}(P^{(h)}_{11}+P^{(h)}_{22}-1) = P^{(h)}_{22}(1-p_{11})+p_{11}(1-P^{(h)}_{11}) > 0$$

$$1 - p_{11}P^{(h)}_{12} - p_{12}P^{(h)}_{22} = 1 - p_{11}(1 - P^{(h)}_{11} - P^{(h)}_{22}) - P^{(h)}_{22} = P^{(h)}_{21}(1 - p_{11}) + p_{11}P^{(h)}_{11} > 0$$

$$1 - p_{21}P^{(h)}_{11} - p_{22}P^{(h)}_{21} = 1 - p_{21}(P^{(h)}_{11} - 1 + P^{(h)}_{22}) - P^{(h)}_{21} = P^{(h)}_{22}p_{22} + p_{21}(1 - P^{(h)}_{11}) > 0$$

$$1 - p_{21}P^{(h)}_{12} - p_{22}P^{(h)}_{22} = 1 - p_{21}(P^{(h)}_{12} - 1 + P^{(h)}_{21}) - P^{(h)}_{22} = P^{(h)}_{21}p_{11} + p_{12}P^{(h)}_{11} > 0$$

Hence,

$$sign(a_{s_{t+h}} - P^{(h)}_{s_t:}\boldsymbol{a}) = \begin{cases} sign(a_1 - a_2) & \text{if } s_{t+h} = 1 \\ -sign(a_1 - a_2) & \text{if } s_{t+h} = 2 \end{cases},$$

which is Corollary 1 in the main text.

A.3 Proof of Proposition 2

Forecast revisions can be expressed as

$$\mathbb{E}_t y_{t+h} - \mathbb{E}_{t-1} y_{t+h} = \phi^h \frac{P^{(h)}_{s_t:}\boldsymbol{a}}{a_{s_t}} y_t - \phi^{h+1} \frac{P^{(h+1)}_{s_{t-1}:}\boldsymbol{a}}{a_{s_{t-1}}} y_{t-1}. \tag{A.5}$$

To obtain (**??**), we isolate y_t from (A.5) and use it to substitute y_t in (12).

A.4 Proof of Proposition 3

Full-information rational expectations about X_{t+h} are given by:

$$\begin{aligned}\mathbb{E}_t X_{t+h} &= \mathbb{E}_t \left(C_{t+h} + A_{t+h} X_{t+h-1} + B_{t+h}\epsilon_{t+h}\right) \\ &= \mathbb{E}_t \left(C_{t+h} + \sum_{\tau=1}^{h-1} \left(\prod_{l=0}^{\tau} A_{t+h-\tau}\right) C_{t+h-\tau} + \prod_{\tau=0}^{h-1} A_{t+h-\tau} X_t\right)\end{aligned}$$

We will show through a proof by induction that

$$\begin{aligned}\mathbb{E}_t X_{t+h} &= \left(\begin{bmatrix} C_1 & C_2 \end{bmatrix} (P')^h + \begin{bmatrix} A_1 & A_2 \end{bmatrix} (P' \otimes I_{n_x}) \sum_{\tau=1}^{h-1} \left(\widetilde{A}(P' \otimes I_{n_x})\right)^{\tau-1} \widetilde{C}(P')^{h-\tau}\right) I_{:s_t} \\ &+ \begin{bmatrix} A_1 & A_2 \end{bmatrix} (P' \otimes I_{n_x}) \left(\widetilde{A}(P' \otimes I_{n_x})\right)^{h-1} \bar{\iota}_{s_t} X_t\end{aligned}$$

where $\widetilde{C} = \begin{bmatrix} C_1 & 0_{n_x \times 1} \\ 0_{n_x \times 1} & C_2 \end{bmatrix}$, $\widetilde{A} = \begin{bmatrix} A_1 & 0_{n_x \times n_x} \\ 0_{n_x \times n_x} & A_2 \end{bmatrix}$, $\bar{\iota}_{s_t}$ is a $2n_x \times n_x$ size matrix whose s_t^{th} block of n_x rows together with the columns form an identity matrix and the rest of the elements are 0, and $I_{:s_t}$ is the s_t^{th} column of a 2×2 identity matrix.

One can check that the expression above holds for $h \in \{1, 2\}$. Suppose it also holds for $h = \bar{h}$; does it also apply for $h = \bar{h} + 1$?

$$\begin{aligned}\mathbb{E}_t X_{t+\bar{h}+1} &= \mathbb{E}_t \left(C_{t+\bar{h}+1} + A_{t+\bar{h}+1} X_{t+\bar{h}}\right) = \mathbb{E}_t \left[\mathbb{E}_{t+1} \left(C_{t+\bar{h}+1} + A_{t+\bar{h}+1} X_{t+\bar{h}}\right)\right] \\ &= \mathbb{E}_t \left(\begin{bmatrix} C_1 & C_2 \end{bmatrix} (P')^{\bar{h}} + \begin{bmatrix} A_1 & A_2 \end{bmatrix} (P' \otimes I_{n_x}) \sum_{\tau=1}^{\bar{h}-1} \left(\widetilde{A}(P' \otimes I_{n_x})\right)^{\tau-1} \widetilde{C}(P')^{\bar{h}-\tau}\right) I_{:s_{t+1}} \\ &+ \mathbb{E}_t \left(\begin{bmatrix} A_1 & A_2 \end{bmatrix} (P' \otimes I_{n_x}) \left(\widetilde{A}(P' \otimes I_{n_x})\right)^{\bar{h}-1} \bar{\iota}_{s_{t+1}} X_{t+1}\right) \\ &= \left(\begin{bmatrix} C_1 & C_2 \end{bmatrix} (P')^{\bar{h}} + \begin{bmatrix} A_1 & A_2 \end{bmatrix} (P' \otimes I_{n_x}) \sum_{\tau=1}^{\bar{h}-1} \left(\widetilde{A}(P' \otimes I_{n_x})\right)^{\tau-1} \widetilde{C}(P')^{\bar{h}-\tau}\right) \mathbb{E}_t I_{:s_{t+1}} \\ &+ \begin{bmatrix} A_1 & A_2 \end{bmatrix} (P' \otimes I_{n_x}) \left(\widetilde{A}(P' \otimes I_{n_x})\right)^{\bar{h}-1} \mathbb{E}_t \left(\bar{\iota}_{s_{t+1}} A_{t+1} X_t + \bar{\iota}_{s_{t+1}} C_{t+1}\right) \\ &= \begin{bmatrix} C_1 & C_2 \end{bmatrix} (P')^{\bar{h}+1} I_{:s_t} + \begin{bmatrix} A_1 & A_2 \end{bmatrix} (P' \otimes I_{n_x}) \sum_{\tau=1}^{\bar{h}} \left(\widetilde{A}(P' \otimes I_{n_x})\right)^{\tau-1} \widetilde{C}(P')^{\bar{h}+1-\tau} I_{:s_t} \\ &+ \underbrace{\begin{bmatrix} A_1 & A_2 \end{bmatrix} (P' \otimes I_{n_x}) \left(\widetilde{A}(P' \otimes I_{n_x})\right)^{\bar{h}} \bar{\iota}_{s_t}}_{\text{response to info at time t}} X_t\end{aligned}$$

where the equality in the 7^{th} row follows from $\mathbb{E}_t I_{:s_{t+1}} = P' I_{:s_t}$. Therefore,

$$M_{t,t+h} = \begin{bmatrix} C_1 & C_2 \end{bmatrix} (P')^h I_{:s_t} + \begin{bmatrix} A_1 & A_2 \end{bmatrix} (P' \otimes I_{n_x}) \sum_{\tau=1}^{h-1} \left(\widetilde{A}(P' \otimes I_{n_x}) \right)^{\tau-1} \widetilde{C} (P')^{h-\tau} I_{:s_t} \tag{A.6}$$

$$Q_{t,t+h} = \begin{bmatrix} A_1 & A_2 \end{bmatrix} (P' \otimes I_{n_x}) \left(\widetilde{A}(P' \otimes I_{n_x}) \right)^{h-1} \bar{\iota}_{s_t} \tag{A.7}$$

Finally, we note that $M_{t,t+1} = \begin{bmatrix} C_1 & C_2 \end{bmatrix} P' I_{:s_t}$.

A.5 Proof of Proposition 4

Consider the ex-post forecasting error about vector Y_{t+h}, where $h \geq 1$:

$$\begin{aligned}
FE_{t,t+h} &= Y_{t+h} - \mathbb{E}_t Y_{t+h} \\
&= \Psi_1 \left(C_{s_{t+h}} + A_{s_{t+h}} X_{t+h-1} + B_{t+h} \epsilon_{t+h} - M_{t,t+h} - Q_{t,t+h} X_t \right) \\
&= \Psi_1 C_{s_{t+h}} + \Psi_1 A_{s_{t+h}} \left(C_{s_{t+h-1}} + A_{s+h-1} X_{t+h-1} + B_{s_{t+h-1}} \epsilon_{t+h-1} - M_{t,t+h} - Q_{t,t+h} X_t \right) + \Psi_1 B_{s_{t+h}} \epsilon_{t+h} \\
&= ... \\
&= \underbrace{\Psi_1 \left(C_{s_{t+h}} + \sum_{\tau=1}^{h-1} \left(\prod_{l=0}^{\tau} A_{s_{t+h-l}} \right) C_{s_{t+h-\tau}} - M_{t,t+h} \right)}_{\Theta_{t,t+h} \equiv \text{bias}} + \underbrace{\Psi_1 \left(\prod_{\tau=1}^{h} A_{s_{t+\tau}} - Q_{t,t+h} \right)}_{\Gamma_{t,t+h} \equiv \text{ex-post predictability}} X_t + error_{t+h}
\end{aligned} \tag{A.8}$$

where $error_{t+h} = \Psi_1 \sum_{\tau=1}^{h-1} (\prod_{l=0}^{\tau} A_{s_{t+h-l}}) B_{s_{t+h-\tau}} \epsilon_{t+h-\tau} + \Psi_1 B_{s_{t+h}} \epsilon_{t+h}$. We now turn to expressing ex-post forecast errors as a function of ex-ante forecast revisions. The FIRE forecasts about the endogenous variables vector X_{t+h} in periods t and $(t-1)$ are given by, respectively,

$$\mathbb{E}_t X_{t+h} = M_{t,t+h} + Q_{t,t+h} X_t \tag{A.9}$$

$$\mathbb{E}_{t-1} X_{t+h} = M_{t-1,t+h} + Q_{t-1,t+h} X_{t-1} \tag{A.10}$$

Hence, the ex-ante forecast revision about X_{t+h} is given by

$$FR_{t,t+h} = \mathbb{E}_t X_{t+h} - \mathbb{E}_{t-1} X_{t+h} = M_{t,t+h} - M_{t-1,t+h} + Q_{t,t+h} X_t - Q_{t-1,t+h} X_{t-1} \tag{A.11}$$

If $Q_{t,t+h}$ is invertible, then from (A.11), $X_t = Q_{t,t+h}^{-1} (\mathbb{E}_t X_{t+h} - \mathbb{E}_{t-1} X_{t+h}) + Q_{t,t+h}^{-1} Q_{t-1,t+h} X_{t-1} - Q_{t,t+h}^{-1} (M_{t,t+h} - M_{t-1,t+h})$. Substituting for X_t into (A.8), we can rewrite ex-post forecast errors as a

function of ex-ante forecast revisions.

$$Y_{t+h} - \mathbb{E}_t Y_{t+h} = \underbrace{\left(\Theta_{t,t+h} - \Gamma_{t,t+h} Q_{t,t+h}^{-1}(M_{t,t+h} - M_{t-1,t+h})\right)}_{=\Omega_{t,t+h}} + \underbrace{\Gamma_{t,t+h} Q_{t,t+h}^{-1}}_{=\Delta_{t,t+h}}(\mathbb{E}_t X_{t+h} - \mathbb{E}_{t-1} X_{t+h})$$

$$+ \underbrace{\Gamma_{t,t+h} Q_{t,t+h}^{-1} Q_{t-1,t+h}}_{=\Lambda_{t-1,t+h}} X_{t-1} + error_{t+h} \tag{A.12}$$

If $Q_{t,t+h}$ is non-invertible, we proceed as follows. Let q_{ij} and $\widetilde{q}_{ij}$ denote the element located in row i and column j in matrices $Q_{t,t+h}$ and $\widetilde{Q}_{t-1,t+h}$, respectively. Furthermore, let m_i be the element located in row i in matrix $(M_{t,t+h} - M_{t-1,t+h})$. The ex-ante forecast revision of any variable X_i in X can be written as:

$$FR_{i,t,t+h} = \mathbb{E}_t X_{i,t+h} - \mathbb{E}_{t-1} X_{i,t+h} = m_i + \sum_{j=1}^{n_x} q_{ij} X_{j,t} - \sum_{j=1}^{n_x} \widetilde{q}_{ij} X_{j,t-1} \tag{A.13}$$

Then, any variable X_{kt} in X_t can be written as a function of the ex-ante forecast revision about variable $X_{i,t+h}$, where i is chosen such that $q_{ik} \neq 0$, as well as X_{-kt}, X_{t-1}, and a constant:

$$\begin{aligned} X_{kt} &= \frac{FR_{i,t,t+h} - \sum_{j \neq k} q_{ij} X_{jt} + \sum_j \widetilde{q}_{ij} X_{j,t-1} - m_i}{q_{ik}} \\ &= \underbrace{\begin{bmatrix} 0 & 0 & \dots & \frac{1}{q_{ik}} & \dots & 0 & 0 \end{bmatrix}}_{Q^-(k,:)} FR_{t,t+h} - \underbrace{\begin{bmatrix} \frac{q_{i1}}{q_{ik}} & \dots & \frac{q_{i,k-1}}{q_{ik}} & 0 & \frac{q_{i,k+1}}{q_{ik}} & \dots & \frac{q_{in_x}}{q_{ik}} \end{bmatrix}}_{Q_Q(k,:)} X_t + \underbrace{\begin{bmatrix} \frac{\widetilde{q}_{i1}}{q_{ik}} & \dots & \frac{\widetilde{q}_{in_x}}{q_{ik}} \end{bmatrix}}_{\widetilde{Q}_Q(k,:)} X_{t-1} \\ &- \underbrace{\frac{m_i}{q_{ik}}}_{M_{Q(k,:)}} \end{aligned} \tag{A.14}$$

It follows that vector X_t can be written as a function of ex-ante forecast revisions as described below:

$$X_t = Q^- FR_{t,t+h} - Q_Q X_t + \widetilde{Q}_Q X_{t-1} - M_Q$$

From here, we have that $X_t = (I_{nx} + Q_Q)^{-1}(Q^- FR_{t,t+h} + \widetilde{Q}_Q X_{t-1} - M_Q)$, and that

$$
\begin{aligned}
Y_{t+h} - \mathbb{E}_t Y_{t+h} &= \underbrace{\Theta_{t,t+h} - \Gamma_{t,t+h}(I_{nx} + Q_Q)^{-1} M_Q}_{=\Omega_{t,t+h}} + \underbrace{\Gamma_{t,t+h}(I_{nx} + Q_Q)^{-1} Q^-}_{=\Delta_{t,t+h}} FR_{t,t+h} \qquad \text{(A.15)} \\
&+ \underbrace{\Gamma_{t,t+h}(I_{nx} + Q_Q)^{-1} \widetilde{Q}_Q}_{=\Lambda_{t-1,t+h}} X_{t-1} + error_{t+h}
\end{aligned}
$$

A.6 Proof of Corollary 2

1. Consider the case when $C_1 \neq C_2$, while $A_1 = A_2 = A$ and $B_1 = B_2 = B$. Note that, for $A_1 = A_2 = A$, the following is true:

$$
Q_{t,t+h} = \Psi_1 \begin{bmatrix} A & A \end{bmatrix} (P' \otimes I_{n_x}) \left(\widetilde{A}(P' \otimes I_{n_x}) \right)^{h-1} \bar{\iota}_{s_t} = \Psi_1 A^h \qquad \text{(A.16)}
$$

   Therefore, $\Gamma_{t,t+h} = \Psi_1(A^h - A^h) = \mathbf{0}_{n_y \times n_x}$, and, given that $B_1 = B_2 = B$, we have that $error_{t+h} = \Psi_1 \sum_{\tau=0}^{h-1} A^\tau B \epsilon_{t+h-\tau}$. Furthermore,

$$
M_{t,t+h} = \Psi_1 \sum_{\tau=0}^{h-1} A^\tau \widetilde{C} (P')^{h-\tau} I_{:s_t} \neq \mathbf{0}_{n_y \times 1} \qquad \text{(A.17)}
$$

   In this case, $\Theta_{t,t+h} = \Psi_1 \sum_{\tau=0}^{h-1} A^\tau (C_{s_{t+h-\tau}} - \widetilde{C}(P')^{h-\tau} I_{:s_t}) \neq \mathbf{0}_{n_y \times 1}$,therefore, ex-post forecast errors will be biased, but they will not respond to information embedded in X_t.

2. Now, suppose that $C_1 = C_2 = C$ and $A_1 = A_2 = A$, while $B_1 \neq B_2$. Note that, given that $C_1 = C_2 = C$, the following is true:

$$
M_{t,t+h} = \Psi_1 \sum_{\tau=0}^{h-1} A^\tau C \qquad \text{(A.18)}
$$

   So, $\Theta_{t,t+h} = \mathbf{0}_{n_y \times n_x}$, implying that ex-post forecast error are not biased. Furthermore, $A_1 = A_2 = A$ implies that $\Gamma_{t,t+h} = \mathbf{0}_{n_y \times n_x}$. For $B_1 \neq B_2$, the error term is as defined in Section A.5. Consequently, when the regime shifts affect only the relationship between the endogenous variables and innovations, ex-post forecast errors are just accumulated noise, similar to the case of no regime shifts discussed below.

3. Finally, shutting down all regime shifts in the model implies that $\Theta_{t,t+h} = \mathbf{0}_{n_y \times 1}$, $\Gamma_{t,t+h} = \mathbf{0}_{n_y \times n_x}$, and $error_{t+h} = \Omega \sum_{\tau=0}^{h-1} A^\tau B \epsilon_{t+h-\tau}$. Hence, in this case, forecast errors are accumulated noise similar to the second case.

B Univariate regime-shift model

B.1 Illustration of forecast error predictability

Consider estimating a demeaned version of regression (2),

$$y_{t+h} - F_t y_{t+h} = \gamma y_t + e_{t+h}, \tag{B.1}$$

from sample $\{y_{t+h} - F_t y_{t+h}, y_t\}_{t=1}^T$ generated by regime-switching model (8)-(10) with constant mean set to zero (i.e., $\bar{y}_1 = \bar{y}_2 = 0$).[17] The insights gained from this special case apply equally to the more general case with regime-switching in the mean, but the analytics would be not be as straightforward. Similarly, we could consider estimating a demeaned version of regression (1), $y_{t+h} - F_t y_{t+h} = \delta\left(F_t y_{t+h} - F_{t-1} y_{t+h}\right) + e_{t+h}$. Different from regression (B.1), however, no closed-form solution of the type derived below exists because forecast revisions do not span the information set that agents use to forecast y_{t+h}. We return to this point in Section **??** where we consider a generalized regime-switching model.

Under the assumption that forecasters are fully informed rational expectations agents (i.e., $F_t = \mathbb{E}_t$), Proposition 1 implies that the relation between forecast error $y_{t+h} - F_t y_{t+h}$ and current realization y_t is a function of agents' future regime expectations and actual regime realizations. Hence, the ordinary least squares (OLS) estimate $\hat{\gamma}_T$ of regression (B.1) based on sample $\{y_{t+h} - F_t y_{t+h}, y_t\}_{t=1}^T$ depends on agents' transition probabilities $\Theta = [p_{11}, p_{22}]'$ and the sequence of regime realizations $\{s_t\}_{t=1}^{T+h}$, both of which are unobserved to the econometrician and have to be inferred from available information $\mathcal{I}$.[18] Formally, $\hat{\gamma}_T$ is a continuous mixture of random variables with probability density function

$$f\left(\hat{\gamma}_T|\mathcal{I}\right) = \int ... \int f\left(\hat{\gamma}_T|\left\{s_t\right\}_{t=1}^{T+h};\Theta\right) f\left(\left\{s_t\right\}_{t=1}^{T+h}|\Theta;\mathcal{I}\right) f\left(\Theta|\mathcal{I}\right) d\Theta ds_{T+h} ... ds_1, \tag{B.2}$$

where $f\left(\left\{s_t\right\}_{t=1}^{T+h}|\mathcal{I};\Theta\right)$ is the joint density of the sequence of regime realizations $\{s_t\}_{t=1}^{T+h}$ conditional on agents' regime expectations Θ and the econometrician's information set $\mathcal{I}$, and $f\left(\Theta|\mathcal{I}\right)$ is the density of agents' expectations given the econometrician's information set $\mathcal{I}$.[19] To provide further insights about implications 2 and 3, consider $\mathbb{E}\left[\hat{\gamma}_T|\left\{s_t\right\}_{t=1}^{T+h};\Theta\right]$, the expected OLS estimate of the slope coefficient γ

[17] Using demeaned data is consistent with the FIRE solution of the data-generating process in (8). This type of equation generally arises as the result of log-linearizing optimality conditions of dynamic stochastic problems, which by definition refer to deviations from the mean. Section **??** considers the generalized framework that includes regime shifts in constant terms.

[18] The econometrician's information set $\mathcal{I}$ includes $\mathcal{Y}^{T+h} \equiv (y_{T+h}, y_t, ...y_1)$ as well as potentially further leads and lags of y_t and other observables.

[19] Of course, the econometrician may also be uncertain not just about agents' regime expectations but also about other aspects of the data-generating process. For simplicity, we abstract from this type of model uncertainty here but will integrate in the regime-robust test of FIRE in Section 5.

conditional on a given sequence of regime realizations and agents' regime expectations. By Proposition 1, this expectation can be expressed as

$$\mathbb{E}\left[\hat{\gamma}_T|\left\{s_t\right\}_{t=1}^{T+h};\Theta\right]=\underbrace{\mathbb{E}\left[\sum_{t=1}^{T}\left(\frac{a_{s_t}^2x_t^2}{\sum_{t=1}^{T}a_{s_t}^2x_t^2}\right)\gamma_{s_t,s_{t+h}}^{(h)}\right]}_{\mathbb{E}\left[\hat{\gamma}_T^c|\{s_t\}_{t=1}^{T+h};\Theta\right]}+\underbrace{\mathbb{E}\left[\sum_{t=1}^{T}\left(\frac{a_{s_t}x_t}{\sum_{t=1}^{T}a_{s_t}^2x_t^2}\right)\xi_{t+h}\right]}_{\mathbb{E}\left[\Delta_T|\{s_t\}_{t=1}^{T+h};\Theta\right]}. \tag{B.3}$$

The first part, $\mathbb{E}\left[\hat{\gamma}_T^c|\left\{s_t\right\}_{t=1}^{T+h};\Theta\right]$, is the expected bias-corrected estimate and equals a weighted average of the different values that $\gamma_{s_t,s_{t+h}}^{(h)}$ can take, with the weights depending on the sequence of regime realizations $\{s_t\}_{t=1}^{T+h}$. The second part, $\mathbb{E}\left[\Delta_T|\left\{s_t\right\}_{t=1}^{T+h};\Theta\right]$, is the finite sample bias arising from the fact that the regressor $y_t=a_{s_t}x_t$ is not strictly exogenous with respect to forecast error innovations $\xi_{t+h}\equiv a_{s_{t+h}}\sum_{\tau=1}^{h}\phi^{h-\tau}\varepsilon_{t+\tau}$.

Applying standard time-series analysis (e.g., Hamilton, 1994), the sign of these two expectations can be described as follows.

PROPOSITION 5. *Given the environment described in Proposition 1 and assuming* $a_1>a_2$ *and* $0<\phi<1$*, then under the null hypothesis of FIRE:*

1. *For a finite sequence of regime realizations* $\{s_t\}_{t=1}^{T+h}$ *characterized by conditional sample frequencies* $f_{ij}^{(h)}\equiv\frac{1}{T}\sum_{t=1}^{T}\mathbf{1}(s_t=i|s_{t+h}=j)$*,*

$$\mathbb{E}\left[\hat{\gamma}_T^c|\left\{s_t\right\}_{t=1}^{T+h};\Theta\right]\gtrless 0\ \textit{if}\ f_{11}^{(h)}\gtrless g(f_{22}^{(h)})\equiv\frac{a_1p_{11}^{(h)}(1-f_{22}^{(h)})+a_2(f_{22}^{(h)}-p_{22}^{(h)})}{a_1(1-f_{22}^{(h)})+a_2(f_{22}^{(h)}-p_{22}^{(h)})};$$

and

$$\mathbb{E}\left[\Delta_T|\left\{s_t\right\}_{t=1}^{T+h};\Theta\right]<0.$$

2. *For* $T\to\infty$*,* $f_{ii}^{(h)}\to p_{ii}^{(h)}$ *for* $i\in\{1,2\}$ *and therefore* $f_{11}^{(h)}\to g(f_{22}^{(h)})$*; and* $\mathbb{E}\left[\Delta_T|\left\{s_t\right\}_{t=1}^{T+h};\Theta\right]\to 0$*. Hence*

$$\underset{T\to\infty}{plim}\mathbb{E}\left[\hat{\gamma}_T|\left\{s_t\right\}_{t=1}^{T+h};\Theta\right]=0.$$

Proof. The OLS coefficient estimate of regression $y_{t+h}-F_ty_{t+h}=\gamma y_t+e_{t+h}$ is

$$\hat{\gamma}_T=\left(\sum_{t=1}^{T}y_t^2\right)^{-1}\left(\sum_{t=1}^{T}y_t(y_{t+h}-F_ty_{t+h})\right).$$

Assuming that the data are generated by the univariate regime-switching model under FIRE in (8)-(10),

we can use Proposition 1 to express this estimate as

$$\hat{\gamma}_T = \left(\sum_{t=1}^{T} a_{s_t}^2 x_t^2\right)^{-1} \left(\sum_{t=1}^{T} \gamma_{s_t,s_{t+h}}^{(h)} a_{s_t}^2 x_t^2\right) + \left(\sum_{t=1}^{T} a_{s_t}^2 x_t^2\right)^{-1} \left(\sum_{t=1}^{T} a_{s_t} x_t \xi_{t+h}\right),$$

where $\gamma_{s_t,s_{t+h}}^{(h)} \equiv \frac{(-1)^{s_{t+h}-1}(a_1-a_2)\left(1-P_{s_t,s_{t+h}}^{(h)}\right)\phi^h}{a_{s_t}}$ and $\xi_{t+h} \equiv a_{s_{t+h}} \sum_{\tau=1}^{h} \phi^{h-\tau}\varepsilon_{t+\tau}$. Hence, the expected OLS coefficient given regime sequence $\{s_t\}_{t=1}^{T+h}$ is

$$\begin{aligned}\mathbb{E}\left[\hat{\gamma}_T | \{s_t\}_{t=1}^{T+h}\right] = \mathbb{E}&\left[\left(\sum_{t=1}^{T} a_{s_t}^2 x_t^2\right)^{-1} \left(\sum_{t=1}^{T} \gamma_{s_t,s_{t+h}}^{(h)} a_{s_t}^2 x_t^2\right) | \{s_t\}_{t=1}^{T+h}\right] \\ &+ \mathbb{E}\left[\left(\sum_{t=1}^{T} a_{s_t}^2 x_t^2\right)^{-1} \left(\sum_{t=1}^{T} a_{s_t} x_t \xi_{t+h}\right) | \{s_t\}_{t=1}^{T+h}\right].\end{aligned}$$

The first term is the expected estimate implied by the data-generating process, while the second term is the bias arising from the fact that in finite samples, innovations $\varepsilon_{t+1}, \varepsilon_{t+2}, ..., \varepsilon_{t+h}$ affect not only ξ_{t+h} but also $\sum_{t=1}^{T} a_{s_t}^2 x_t^2$. In particular, for positive values of a_1, a_2 and ϕ, realizations of ξ_{t+h} are positively correlated with $\sum_{t=1}^{T} a_{s_t}^2 x_t^2$ in finite samples and thus, the bias is negative.

To make progress on the first term, note that we can rewrite it as

$$\mathbb{E}\left[\hat{\gamma}_T^c | \{s_t\}_{t=1}^{T+h}\right] = \mathbb{E}\left[\left(\frac{1}{T}\sum_{t=1}^{T} a_{s_t}^2 x_t^2\right)^{-1} \left(\frac{1}{T}\sum_{t=1}^{T} \gamma_{s_t,s_{t+h}}^{(h)} a_{s_t}^2 x_t^2\right) | \{s_t\}_{t=1}^{T+h}\right],$$

where $\hat{\gamma}_T^c$ denotes the fact that this is effectively the bias-corrected estimate. Now, imagine repeating the OLS regression for K samples of size $T+h$, with each sample being conditional on the *same* sample regime sequence $\{s_t\}_{t=1}^{T+h}$. Then, since the stochastic process $\{\varepsilon_t\}$ and therefore $\{x_t\}$ are covariance-stationary and ergodic for second moments, the sample moments will converge to the population moments as K goes to infinity; i.e.,

$$\frac{1}{KT}\sum_{k=1}^{K}\sum_{t=1}^{T} a_{s_t}^2 x_{k,t}^2 | \{s_t\}_{t=1}^{T} \xrightarrow{p} \mathbb{E}\left[\frac{1}{KT}\sum_{k=1}^{K}\sum_{t=1}^{T} a_{s_t}^2 x_{k,t}^2 | \{s_t\}_{t=1}^{T}\right] = \sum_{t=1}^{T} a_{s_t}^2 \mathbb{E}\left[x_{k,t}^2\right],$$

and thus, equivalently[20]

$$\left(\frac{1}{KT}\sum_{k=1}^{K}\sum_{t=1}^{T} a_{s_t}^2 x_{k,t}^2 | \{s_t\}_{t=1}^{T}\right)^{-1} \xrightarrow{p} \mathbb{E}\left[\frac{1}{KT}\sum_{k=1}^{K}\sum_{t=1}^{T} a_{s_t}^2 x_{k,t}^2 | \{s_t\}_{t=1}^{T}\right]^{-1} = \left(\sum_{t=1}^{T} a_{s_t}^2 \mathbb{E}\left[x_t^2\right]\right)^{-1},$$

[20] See, for example, Proposition 7.1 in Hamilton (1994).

and

$$\frac{1}{KT}\sum_{k=1}^{K}\sum_{t=1}^{T}\gamma^{(h)}_{s_t,s_{t+h}}a^2_{s_t}x^2_{k,t}|\{s_t\}_{t=1}^{T} \xrightarrow{p} \mathbb{E}\left[\frac{1}{KT}\sum_{k=1}^{K}\sum_{t=1}^{T}\gamma^{(h)}_{s_t,s_{t+h}}a^2_{s_t}x^2_{k,t}|\{s_t\}_{t=1}^{T}\right] = \sum_{t=1}^{T}\gamma^{(h)}_{s_t,s_{t+h}}a^2_{s_t}\mathbb{E}\left[x^2_t\right].$$

Hence, the OLS estimate converges to the linear projection coefficient (see e.g., Chapter 4.2 in Hamilton, 1994), and we can express the expected bias-corrected estimate $\hat{\gamma}^c_T$ conditional on regime sequence $\{s_t\}_{t=1}^{T+h}$ as

$$\mathbb{E}\left[\hat{\gamma}^c_T|\,\{s_t\}_{t=1}^{T+h}\right] = \frac{\sum_{t=1}^{T}\gamma^{(h)}_{s_t,s_{t+h}}a^2_{s_t}\mathbb{E}\left[x^2_t\right]}{\sum_{t=1}^{T}a^2_{s_t}\mathbb{E}\left[x^2_t\right]} = \frac{\sum_{j=1}^{2}\sum_{i=1}^{2}a^2_i\gamma^{(h)}_{ij}\mathcal{F}^{(h)}_T(i,j)}{\sum_{i=1}^{2}a^2_i\mathcal{F}^{(h)}_T(i)},$$

where the second equality makes use of the fact that under the assumption of two regimes, $\gamma^{(h)}_{s_t,s_{t+h}}a^2_{s_t}$ takes on one of four values (each with joint sample frequency $\mathcal{F}^{(h)}_T(i,j)$), while $a^2_{s_t}$ takes on one of two values (each with sample frequency $\mathcal{F}^{(h)}_T(i)$).

To analyze the properties of $\mathbb{E}\left[\hat{\gamma}^c_T|\,\{s_t\}_{t=1}^{T+h}\right]$, define the following conditional sample transition probabilities:

$$f^{(h)}_{ji} = \frac{\sum_{t=1}^{T}\mathbf{1}(s_t=i,s_{t+h}=j)}{\sum_{t=1}^{T}\mathbf{1}(s_{t+h}=j)} = \frac{\mathcal{F}^{(h)}_T(i,j)}{\mathcal{F}^{(h)}_T(j)}$$

where $\sum_{i=1}^{2}f^{(h)}_{ji}=1$. One can show that

$$\mathcal{F}^{(h)}_T(j) = \frac{1}{T}\sum_{t=1}^{T}\mathbf{1}(s_{t+h}=j) \approx \begin{cases} \frac{1-f^{(h)}_{22}}{2-f^{(h)}_{11}-f^{(h)}_{22}} & \text{if } j=1 \\ \\ \frac{1-f^{(h)}_{11}}{2-f^{(h)}_{11}-f^{(h)}_{22}} & \text{if } j=2 \end{cases}$$

Hence, $\mathcal{F}^{(h)}_T(i,j) = f^{(h)}_{ji}\mathcal{F}^{(h)}_T(j)$, and $\mathcal{F}^{(h)}_T(j)$ depends on $f^{(h)}_{11}$and $f^{(h)}_{22}$only. Substituting these expressions together with $\gamma^{(h)}_{s_t,s_{t+h}} \equiv \frac{(-1)^{s_{t+h}-1}(a_1-a_2)\left(1-P^{(h)}_{s_t,s_{t+h}}\right)\phi^h}{a_{s_t}}$ in the above expression $\mathbb{E}\left[\hat{\gamma}^c_T|\,\{s_t\}_{t=1}^{T+h}\right]$, we obtain

$$\mathbb{E}\left[\hat{\gamma}^c_T|\,\{s_t\}_{t=1}^{T+h}\right] = \underbrace{\frac{\phi^h(a_1-a_2)}{a_1^2(1-f^{(h)}_{22})+a_2^2(1-f^{(h)}_{11})}}_{(+)}\underbrace{[a_1(1-f^{(h)}_{22})(f^{(h)}_{11}-p^{(h)}_{11})-a_2(1-f^{(h)}_{11})(f^{(h)}_{22}-p^{(h)}_{22})]}_{g(f^{(h)}_{11},f^{(h)}_{22})} \tag{B.4}$$

Clearly, the sign of $\mathbb{E}[\gamma|\{s_t\}_{t=1}^{T+h}]$ depends on the sign of $g(f^{(h)}_{11},f^{(h)}_{22})$. The frontier in the $(f^{(h)}_{11},f^{(h)}_{22})$ plane for which $\mathbb{E}[\gamma|\{s_t\}_{t=1}^{T+h}]=0$ is given by

$$f_{11}^{(h)} = g(f_{22}^{(h)}) = \frac{a_1 p_{11}^{(h)}(1 - f_{22}^{(h)}) - a_2(p_{22}^{(h)} - f_{22}^{(h)})}{a_1(1 - f_{22}^{(h)}) - a_2(p_{22}^{(h)} - f_{22}^{(h)})}$$

where $a_1(1 - f_{22}^{(h)}) > a_2(p_{22}^{(h)} - f_{22}^{(h)})$ for any $0 \leq f_{22}^{(h)} \leq 1$, given that $a_1 > a_2$. Moreover, $g(p_{22}^{(h)}) = p_{11}^{(h)}$ and $g(1) = 1$. Then, $\mathbb{E}[\gamma|\{s_t\}_{t=1}^{T+h}] \lesseqgtr 0 \iff f_{11}^{(h)} \lesseqgtr g(f_{22}^{(h)})$.

□

Part 1 establishes that in the presence of regime shifts, the null of unpredictable forecast errors is a knife-edge case. Generally, $\mathbb{E}\left[\hat{\gamma}_T|\{s_t\}_{t=1}^{T+h};\Theta\right] \neq 0$ either because $\mathbb{E}\left[\hat{\gamma}_T^c|\{s_t\}_{t=1}^{T+h};\Theta\right] < 0$ or because $\mathbb{E}\left[\hat{\gamma}_T^c|\{s_t\}_{t=1}^{T+h};\Theta\right] > 0$ but is not exactly offset by $\mathbb{E}\left[\Delta_T|\{s_t\}_{t=1}^{T+h};\Theta\right] < 0$. Intuitively, the bias-corrected expectation $\mathbb{E}\left[\hat{\gamma}_T^c|\{s_t\}_{t=1}^{T+h}\right]$ is a weighted average of the four possible values of γ_{ij}, two of which are positive (γ_{11} and γ_{21}) and two of which are negative (γ_{12} and γ_{22}) . The weights and therefore the sign of $\mathbb{E}\left[\hat{\gamma}_T^c|\{s_t\}_{t=1}^{T+h};\Theta\right]$ depend on the conditional sample frequencies $f_{11}^{(h)}$ and $f_{22}^{(h)}$ implied by regime realizations $\{s_t\}_{t=1}^{T+h}$ relative to what agents expect ($p_{11}^{(h)}$ and $p_{22}^{(h)}$) as well as on the values of a_1 and a_2. Appendix **??** provides further details, including a numerical illustration of the different cases. In turn, the finite-sample bias $\mathbb{E}\left[\Delta_T|\{s_t\}_{t=1}^{T+h};\Theta\right]$ is negative as long as the forcing variable x_t is positively autocorrelated because positive ε_t result in both higher forecast error innovations $\xi_{t+h} \equiv a_{s_{t+h}} \sum_{\tau=1}^{h} \phi^{h-\tau}\varepsilon_{t+\tau}$ and higher $\sum_{t=1}^{T} a_{s_t}^2 x_t^2$ and thus $COV\left(\frac{a_{s_t}x_t}{\sum_{t=1}^{T} a_{s_t}^2 x_t^2}, a_{s_{t+h}} \sum_{\tau=1}^{h} \phi^{h-\tau}\varepsilon_{t+\tau}\right) < 0$. Such a negative finite sample bias occurs for any positively autocorrelated time series process even without regime shifts (see e.g., Hamilton, 1994).[21] But the bias here is potentially larger because both the regressor (current realizations y_t) and forecast error innovations are a function of regime shifts.

Part 2 establishes that as T increases, the conditional sample frequencies of regime realizations converge to their population counterparts and therefore agents' expectations. Hence, periods of over- and under-reaction tend to cancel each other out in the limit. Similarly, the finite sample bias vanishes in the limit. The result makes clear that forecast error predictability remains a finite sample phenomenon.

Proposition 5 raises the natural question of how large and variable deviations of $\mathbb{E}\left[\hat{\gamma}_T|\{s_t\}_{t=1}^{T+h};\Theta\right]$ from zero are for typical sample sizes of forecast error data. To answer this question, we simulate the univariate regime model (8)-(10) with constant mean for 10,000 times over 1000 periods each, keeping the same sequence of regime realizations for each of the simulations. [22] We then compute $\mathbb{E}\left[\hat{\gamma}_T|\{s_t\}_{t=1}^{T+h}\right]$ as well as bias-corrected $\mathbb{E}\left[\hat{\gamma}_T^c|\{s_t\}_{t=1}^{T+h}\right]$ for each rolling window sample of $T = 80$ periods by Monte Carlo

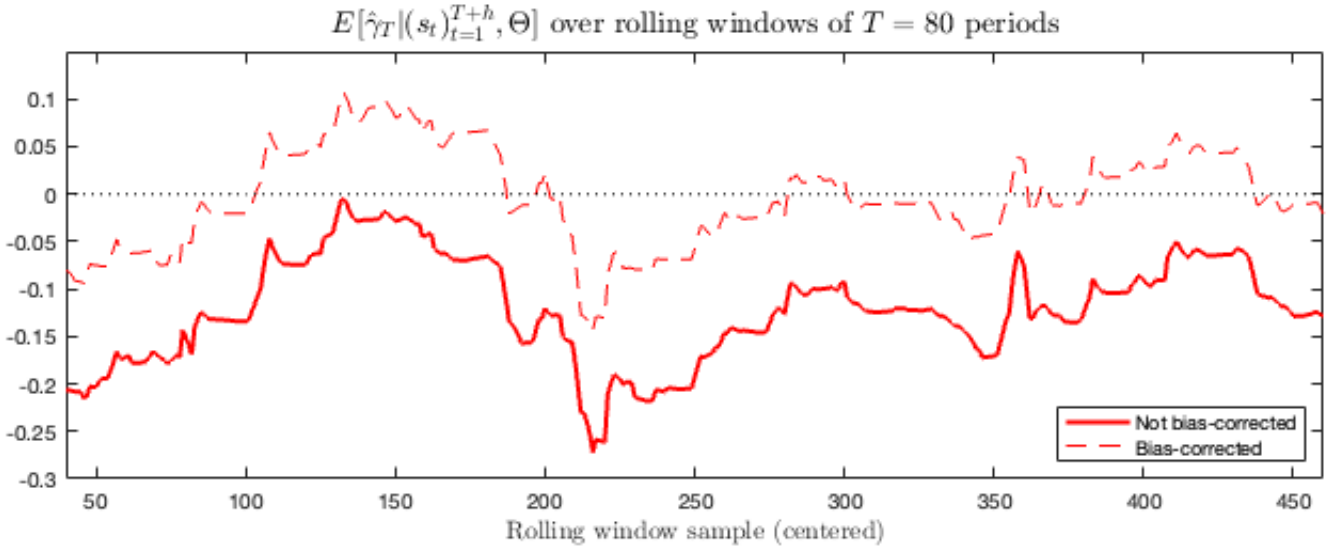
[21] As discussed in Section 2, the empirical literature on forecast error regressions often ignores the resulting finite-sample bias. One exception is Adam, Marcet, and Beutel (2017) who, like us, adopt a simulation-based approach to correct for finite-sample bias. See below for further discussion.

[22] For the simulation without regime shifts, we set $a = 1.25$ and choose $\phi = 0.9$, $\sigma_\varepsilon = 0.01$. For the simulation with regime shifts, we keep the exogenous forcing process unchanged and set $a_1 = 2$, $a_2 = 0.5$ with Markov transition probabilities $p_{11} = p_{22} = 0.7$. Similar results would obtain for other parameterizations.

integration.[23]

Figure 4 reports the results. Across the different samples, $\mathbb{E}\left[\hat{\gamma}_T|\left\{s_t\right\}_{t=1}^{T+h};\Theta\right]$ varies substantially and is on average smaller than zero. As shown by the difference between $\mathbb{E}\left[\hat{\gamma}_T|\left\{s_t\right\}_{t=1}^{T+h};\Theta\right]$ and $\mathbb{E}\left[\hat{\gamma}_T^c|\left\{s_t\right\}_{t=1}^{T+h};\Theta\right]$, the negative average value is due in large part by the finite sample bias, which amounts to about -0.10 for this simulation. But even with bias-correction, there are sizable swings across samples, indicating that even for reasonably large samples, the sequence of realized regimes can differ substantially from expected regime realizations.[24]

Figure 4: Forecast error predictability across samples



Notes: The figure shows average OLS estimates of γ across simulations for rolling windows of 80 periods with data generated from the univariate regime-switching model (8)-(10) with constant mean set to zero. The solid line shows the rolling window average estimates without bias-correction. The dashed line shows the rolling window average estimates with bias-correction.

The illustration yields two important insights. First, in the presence of regime shifts, forecast error

[23]***ADAPT***To correct for finite sample bias, we simulate $i = 1, ..., 10,000$ new samples of 500 periods, preserving the original sequence of regime realizations for each of the samples. For each 40-period rolling window, we then compute the model-implied finite sample bias $fsb_T^i = \frac{\sum_{\tau=t}^{t+39}(y_\tau^i - \bar{y}_{t:t+39}^i)(\xi_{\tau+1}^i - \bar{\xi}_{t+1:t+40}^i)}{\sum_{\tau=t}^{t+39}(y_{\tau,n} - \bar{y}_{t:t+39,n})^2}$ across the different samples i, and subtract the average bias $fsb_T = \sum_{i=1}^{10,000} fsb_T^i/10,000$ from the OLS estimate. Using the bias-corrected OLS point estimates, we compute the fitted values of the regression in (**??**) as well as the standard deviation of the regression error terms for each rolling sample of simulated data. From a normal distribution with that standard deviation and mean 0, we generate $N = 1000$ i.d.d. innovations, and add those disturbances to the fitted values to generate 1000 new datasets of the dependent variable. Preserving the regressor, we re-estimate (**??**) and bias-correct the point estimates for finite sample bias using the model-implied bias averaged across the 1000 simulations. Finally, for each rolling sample we isolate the bottom and top 5 percent bias-free estimates to compute the 90 percent coverage bands.

[24]As shown in Appendix A.X, the same simulation for the univariate model *without* regime shifts would yield a very different result: the bias-corrected $\mathbb{E}\left[\hat{\gamma}_T^c|\left\{s_t\right\}_{t=1}^{T+h}\right]$ would be zero by definition, and $\mathbb{E}\left[\hat{\gamma}_T|\left\{s_t\right\}_{t=1}^{T+h}\right]$ would be negative but essentially constant. Appendix A.X also shows that while finite sample bias becomes negligible for samples of size $T > 150$, deviations of $\mathbb{E}\left[\hat{\gamma}_T|\left\{s_t\right\}_{t=1}^{T+h}\right]$ from zero can remain sizable even for much larger samples.

predictability is not a sufficient condition to reject FIRE. Second, since the econometrician observes neither agents' regime expectations nor the sequence of realized regimes (i.e., p_{11}, p_{22} and f_{11}, f_{22} in the simple example above), the conditional null under FIRE (i.e., $\mathbb{E}\left[\hat{\gamma}_T | \{s_t\}_{t=1}^{T+h}; \Theta\right]$) is itself an uncertain object. The distribution of $\hat{\gamma}_T$ to be used for statistical inference is therefore non-standard and likely to be more disperse than the usual HAC-based standard errors used in the literature.

C New Keynesian FIRE-RS model

C.1 Model

The model, apart from the regime switching part, is closely related to Ireland (2007). The economy is populated by a representative household, a continuum of monopolistically competitive intermediate goods firms, a final goods producer, and a central bank that sets monetary policy according to an interest rate rule with a time-varying inflation target.

Household sector

The representative household maximizes expected lifetime utility:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} a_t \left[\log(C_t - \gamma C_{t-1}) - N_t\right] \tag{C.1}$$

subject to the intertemporal budget constraint

$$P_t C_t + B_t / R_t \leq W_t N_t + B_{t-1} + D_t, \tag{C.2}$$

where C_t is consumption, N_t labor supply, B_t nominal bond holdings, R_t the gross nominal interest rate between t and $t+1$, W_t the nominal wage, P_t is the price level, D_t profits from firms, $\beta \in (0,1)$ the discount factor, γ the degree of external habit formation, and a_t a preference shock following

$$\log(a_t) = \rho_a \log(a_{t-1}) + \sigma_a \varepsilon_{at}. \tag{C.3}$$

The first-order conditions are:

$$\text{Consumption:} \quad \Lambda_t = \frac{a_t}{C_t - \gamma C_{t-1}} - \beta\gamma \mathbb{E}_t \left(\frac{a_{t+1}}{C_{t+1} - \gamma C_t} \right) \tag{C.4}$$

$$\text{Labor supply:} \quad a_t = \Lambda_t \frac{W_t}{P_t} \tag{C.5}$$

$$\text{Bonds:} \quad \Lambda_t = \beta \mathbb{E}_t \left[\Lambda_{t+1} \frac{R_t}{\Pi_{t+1}} \right] \tag{C.6}$$

where Λ_t denotes the Lagrange multiplier on the nominal budget constraint for period t.

Firm sector

There is a continuum of monopolistically competitive firms, each producing a differentiated good with technology

$$Y_t(i) = Z_t N_t(i) \tag{C.7}$$

where Z_t is productivity and $N_t(i)$ is labor used by firm i. Productivity follows a random walk with drift:

$$\ln(Z_t) = \ln(z) + \ln(Z_{t-1}) + \sigma_z \varepsilon_{zt} \tag{C.8}$$

Final good producers combine intermediate goods according to:

$$Y_t = \left(\int_0^1 Y_t(i)^{\frac{\theta_t - 1}{\theta_t}} di \right)^{\frac{\theta_t}{\theta_t - 1}} \tag{C.9}$$

where $\theta_t > 1$ is the elasticity of substitution between goods, following

$$\ln(\theta_t) = (1 - \rho_\theta) \ln(\theta) + \rho_\theta \ln(\theta_{t-1}) + \sigma_\theta \varepsilon_{\theta t}. \tag{C.10}$$

As in Rotemberg (1982), intermediate goods firms face a quadratic cost of adjusting prices:

$$\Phi(P_t(i), P_{t-1}(i)) = \frac{\phi}{2} \left[\frac{P_t(i)}{\Pi_{t-1}^{\alpha} (\Pi_t^*)^{1-\alpha} P_{t-1}(i)} - 1 \right]^2 Y_t, \tag{C.11}$$

where $\phi \geq 0$ governs the magnitude of the adjustment cost, $\Pi_t = P_t / P_{t-1}$ denotes the gross inflation rate, Π_t^* is the central bank's inflation target, and the parameter $1 \geq \alpha \geq 0$ denotes the degree of inflation indexation in the adjustment cost.

Optimal price setting by intermediate goods, after imposing symmetry across producers, yields:

$$\begin{aligned}\theta_t - 1 = \theta_t \left(\frac{a_t}{\Lambda_t Z_t}\right) - \phi \left[\frac{\Pi_t}{\Pi_{t-1}^{\alpha}(\Pi_t^*)^{1-\alpha}} - 1\right]\left[\frac{\Pi_t}{\Pi_{t-1}^{\alpha}(\Pi_t^*)^{1-\alpha}}\right] \\ +\beta\phi\mathbb{E}_t\left\{\left(\frac{\Lambda_{t+1}}{\Lambda_t}\right)\left[\frac{\Pi_{t+1}}{\Pi_t^{\alpha}(\Pi_{t+1}^*)^{1-\alpha}} - 1\right]\left[\frac{\Pi_{t+1}}{\Pi_t^{\alpha}(\Pi_{t+1}^*)^{1-\alpha}}\right]\left(\frac{Y_{t+1}}{Y_t}\right)\right\}\end{aligned} \tag{C.12}$$

Central bank

The central bank follows an interest rate rule with a time-varying inflation target:

$$\log(R_t) - \log(R_{t-1}) = \rho_\pi \log\left(\frac{\Pi_t}{\Pi_t^*}\right) + \rho_{gy}\log\left(\frac{g_t^Y}{g^Y}\right) + \log(v_t) \tag{C.13}$$

where $g_t^Y = Y_t/Y_{t-1}$ is output growth. In Ireland (2007), the inflation target follows a unit root process:

$$\log(\Pi_t^*) = \log(\Pi_{t-1}^*) - \delta_\theta \varepsilon_{\theta t} - \delta_z \varepsilon_{zt} + \sigma_\pi \varepsilon_{\pi t}. \tag{C.14}$$

In our case, the inflation target instead evolves according to a Markov process:

$$\log(\Pi_t^*) = \begin{cases} \log(\Pi_1^*) & \text{if } s_t^{\bar{\pi}} = 1 \\ \log(\Pi_2^*) & \text{if } s_t^{\bar{\pi}} = 2 \end{cases} \tag{C.15}$$

with transition matrix $P^{\bar{\pi}}$. In addition, we allow for regime switching in the monetary policy response coefficients ρ_{π,s_t^r} and ρ_{gy,s_t^r}, governed by a separate Markov process $s_t^r \in \{1,2\}$ with transition matrix P^r.

Symmetric equilibrium and stationary variables

In a symmetric equilibrium, all intermediate-goods-producing firms make identical decisions, so that $Y_t(i) = Y_t$, $N_t(i) = N_t$, and $P_t(i) = P_t$ for all $i \in [0,1]$. Market clearing for bonds requires $B_t = B_{t-1} = 0$, which implies $Y_t = C_t$ and $N_t = \int_0^1 N_t(i)di$.

The real variables Y_t, C_t, and Λ_t inherit unit roots from the process for technology. However, the transformed variables $y_t = Y_t/Z_t$, $c_t = C_t/Z_t$, $\lambda_t = Z_t\Lambda_t$, and $z_t = Z_t/Z_{t-1}$ are all stationary.

Similarly, the nominal variables Π_t and R_t inherit non-stationarity from the inflation target Π_t^*. The transformed variables $\Pi_t/\Pi_t^* = \pi_t$, $R_t/\Pi_t^* = r_t$, and $\Pi_t^*/\Pi_{t-1}^* = \pi_t^*$ remain stationary. For the Markov process for the inflation target, $\Delta\log(\Pi_t^*) = \log(\pi_t^*) = \mu_{\pi^*}(s_t^{\bar{\pi}}, s_{t-1}^{\bar{\pi}})$ where $\mu_{\pi^*}(\cdot)$ denotes the conditional growth rate of the inflation target.

C.2 Derivation of the NKPC

We loglinearize the normalized system of equations around the steady state, analogous to Ireland (2007). In steady state, $\pi = 1$ (inflation equals its target on average), $\pi^* = 1$ (average inflation target growth equals 1), $z_t = z$ (constant economic growth), and y, λ, a are constants. The steady state implies $z/\beta = \bar{R}/\bar{\pi}$, $\bar{w}\bar{\lambda} = 1$, and $\theta = \theta(\bar{a}/\bar{\lambda}) + 1$.

We derive the NKPC from the firm's optimal pricing condition. The FOC, after imposing symmetry, is:

$$\begin{aligned} \theta_t - 1 = \theta_t \left(\frac{a_t}{\Lambda_t Z_t}\right) - \phi \left[\frac{\Pi_t}{\Pi_{t-1}^{\alpha}(\Pi_t^*)^{1-\alpha}} - 1\right]\left[\frac{\Pi_t}{\Pi_{t-1}^{\alpha}(\Pi_t^*)^{1-\alpha}}\right] \\ +\beta\phi\mathbb{E}_t\left\{\left(\frac{\Lambda_{t+1}}{\Lambda_t}\right)\left[\frac{\Pi_{t+1}}{\Pi_t^{\alpha}(\Pi_{t+1}^*)^{1-\alpha}} - 1\right]\left[\frac{\Pi_{t+1}}{\Pi_t^{\alpha}(\Pi_{t+1}^*)^{1-\alpha}}\right]\left(\frac{Y_{t+1}}{Y_t}\right)\right\} \end{aligned} \tag{C.16}$$

Rewriting in terms of normalized variables ($y_t = Y_t/Z_t$, $\lambda_t = Z_t\Lambda_t$, $\pi_t = \Pi_t/\Pi_t^*$, $\pi_t^* = \Pi_t^*/\Pi_{t-1}^*$):

$$\begin{aligned} \theta_t - 1 = \theta_t \left(\frac{a_t}{\lambda_t}\right) - \phi \left[\pi_t\left(\frac{\pi_t^*}{\pi_{t-1}}\right)^{\alpha} - 1\right]\left[\pi_t\left(\frac{\pi_t^*}{\pi_{t-1}}\right)^{\alpha}\right] \\ +\beta\phi\mathbb{E}_t\left\{\left(\frac{\lambda_{t+1}}{\lambda_t g_t^z}\right)\left[\pi_{t+1}\left(\frac{\pi_{t+1}^*}{\pi_t}\right)^{\alpha} - 1\right]\left[\pi_{t+1}\left(\frac{\pi_{t+1}^*}{\pi_t}\right)^{\alpha}\right]\left(\frac{y_{t+1}g_t^z}{y_t}\right)\right\} \end{aligned} \tag{C.17}$$

In steady state, $\pi = 1$, $\pi^* = 1$, $g_t^z = g^z$, and y, λ, a are constants, which gives:

$$\theta - 1 = \theta\left(\frac{a}{\lambda}\right). \tag{C.18}$$

Loglinearizing (with hats denoting log deviations from steady state), the left-hand side gives $\theta - 1 + \theta\hat{\theta}_t$. The first term on the right gives $(\theta - 1)(1 + \hat{\theta}_t + \hat{a}_t - \hat{\lambda}_t)$. For the Rotemberg adjustment cost terms, the key approximations are:

$$\left[\pi_t\left(\frac{\pi_t^*}{\pi_{t-1}}\right)^{\alpha} - 1\right]\left[\pi_t\left(\frac{\pi_t^*}{\pi_{t-1}}\right)^{\alpha}\right] \approx \hat{\pi}_t + \alpha(\hat{\pi}_t^* - \hat{\pi}_{t-1}) \tag{C.19}$$

and

$$\mathbb{E}_t\left[\pi_{t+1}\left(\frac{\pi_{t+1}^*}{\pi_t}\right)^{\alpha} - 1\right] \approx \mathbb{E}_t\left[\hat{\pi}_{t+1} + \alpha(\hat{\pi}_{t+1}^* - \hat{\pi}_t)\right] \tag{C.20}$$

where we use the fact that in steady state $\pi = \pi^* = 1$ so that $[\pi(\pi^*/\pi)^{\alpha} - 1] = 0$ and $[\pi(\pi^*/\pi)^{\alpha}] = 1$.

Putting it together:

$$\begin{aligned} \theta\hat{\theta}_t = (\theta - 1)(\hat{\theta}_t + \hat{a}_t - \hat{\lambda}_t) - \phi\left(\hat{\pi}_t + \alpha(\hat{\pi}_t^* - \hat{\pi}_{t-1})\right) \\ + \beta\phi\mathbb{E}_t\left\{\hat{\pi}_{t+1} + \alpha(\hat{\pi}_{t+1}^* - \hat{\pi}_t)\right\} \end{aligned} \tag{C.21}$$

Dividing by ϕ and rearranging:

$$\begin{aligned}\frac{1}{\phi}\hat{\theta}_t = \frac{\theta-1}{\phi}(\hat{a}_t - \hat{\lambda}_t) - (1+\beta\alpha)\hat{\pi}_t - \alpha\hat{\pi}_t^* + \alpha\hat{\pi}_{t-1} \\ + \beta\mathbb{E}_t\left\{\hat{\pi}_{t+1} + \alpha\hat{\pi}_{t+1}^*\right\}\end{aligned} \tag{C.22}$$

Defining $\psi = (\theta-1)/\phi$, normalizing the cost push shock as $\hat{e}_t = (1/\phi)\hat{\theta}_t$, and rearranging to isolate $\hat{\pi}_t$:

$$(1+\beta\alpha)\hat{\pi}_t = \alpha\hat{\pi}_{t-1} + \beta\mathbb{E}_t\hat{\pi}_{t+1} + \psi\hat{mc}_t - \hat{e}_t - \alpha\left(\hat{\pi}_t^* - \beta\mathbb{E}_t\hat{\pi}_{t+1}^*\right) \tag{C.23}$$

where $\hat{mc}_t = \hat{a}_t - \hat{\lambda}_t$ is real marginal cost.

Note that if $\pi^* \neq 1$ in steady state, several terms in this loglinearization would not cancel out and inflation dynamics implied by the NKPC would be considerably more complicated.

C.3 NKPC and long-run inflation expectations

Consider the special case of $\alpha = 0$ such that the NKPC (C.34) reduces to

$$\hat{\pi}_t = \beta\mathbb{E}_t\hat{\pi}_{t+1} + \psi\hat{mc}_t - \hat{e}_t \tag{C.24}$$

and use the fact that $\hat{\pi}_t = \log\Pi_t - \log\Pi_t^*$ to expand

$$\log\Pi_t = \log\Pi_t^* + \beta\mathbb{E}_t\hat{\pi}_{t+1} + \psi\hat{mc}_t - \hat{e}_t. \tag{C.25}$$

Solving forward:

$$\log\Pi_t = \sum_{j=0}^{\infty}\beta^j\mathbb{E}_t\log\Pi_{t+j}^* + \psi\sum_{j=0}^{\infty}\beta^j\mathbb{E}_t\hat{mc}_{t+j} + \omega_t \tag{C.26}$$

where $\omega_t = -\sum_{j=0}^{\infty}\beta^j\mathbb{E}_t e_{t+j}$.

Under Ireland's random-walk assumption for the inflation target, $\mathbb{E}_t\log\Pi_{t+j}^* = \log\Pi_t^*$ for any j and thus:

$$\log\Pi_t = (1-\beta)^{-1}\log\Pi_t^* + \psi\sum_{j=0}^{\infty}\beta^j\mathbb{E}_t\hat{mc}_{t+j} + \omega_t \tag{C.27}$$

where $\log\Pi_t^* = \mathbb{E}_t\log\Pi_{t+\infty}^*$ are agents' long-run inflation expectations.

Under the Markov process assumption for the inflation target, $\mathbb{E}_t\left[\log\Pi_{t+j}^*|s_t\right] = \left[P^j\widetilde{\Pi}^*\right]_{s_t.}$ where $\widetilde{\Pi}^* = [\log\Pi_1^* \ \log\Pi_2^*]'$. Thus:

$$\log\Pi_t = \overline{\Pi}^*(s_t) + \psi\sum_{j=0}^{\infty}\beta^j\mathbb{E}_t\hat{mc}_{t+j} + \omega_t \tag{C.28}$$

where

$$\overline{\Pi}^*(s_t) = \sum_{j=0}^{\infty} \beta^j \mathbb{E}_t \left[\log \Pi^*_{t+j} | s_t\right] = \left(\left[I - \beta P\right]^{-1} \widetilde{\Pi}^* \right)_{s_t}. \tag{C.29}$$

is the *regime-dependent* present-value of agents' future inflation target expectations.

If $\hat{mc}_t$ followed an AR(1) with persistence φ, we would get

$$\log \Pi_t = \overline{\Pi}^*(s_t) + a(s_t) x_t \tag{C.30}$$

with $a(s_t)x_t = \psi \left[1 - \beta\varphi\right]^{-1} \hat{mc}_t + \omega_t$, which is the univariate type of model used in Section 3 (with added regime shifts in the different parameters).

Long-run inflation expectations under the two formulations differ importantly. Under (C.27), $\mathbb{E}_t(\log \Pi_{t+\infty}) = (1-\beta)^{-1} \log \Pi^*_t$: long-run inflation expectations equal whatever the central bank's inflation target is today. This is a direct implication of the unit-root assumption and abstracts from the notion that what matters for long-run inflation expectations is the central bank's credibility going forward (**?**). Under (C.28), by contrast:

$$\begin{aligned} \mathbb{E}_t(\log \Pi_{t+\infty}) &= \mathbb{E}_t \left(\overline{\Pi}^*(s_{t+\infty}) \right) \\ &= \frac{p_{21}}{p_{12} + p_{21}} \log(\Pi^*_1) + \frac{p_{12}}{p_{12} + p_{21}} \log(\Pi^*_2). \end{aligned} \tag{C.31}$$

Long-run inflation expectations are explicitly about agents' expectations of central bank credibility as given by the transition probabilities p_{12} and p_{21}. While these probabilities are constant in standard formulations—and thus long-run inflation expectations are constant, which is clearly counterfactual—we can relax this by adjusting p_{12} and p_{21} to fit observed long-run inflation expectations, as described in Section 6.

C.4 Loglinearized system

The loglinearized equilibrium conditions are:

$$\text{Lagrange multiplier: } (z-\gamma)(z-\beta\gamma)\hat{\lambda}_t = \gamma z\hat{y}_{t-1} - (z^2+\beta\gamma^2)\hat{y}_t + \beta\gamma z\mathbb{E}_t\hat{y}_{t+1} + (z-\gamma)(z-\beta\gamma\rho_a)\hat{a}_t - \gamma z\hat{z}_t \tag{C.32}$$

$$\text{Euler equation: } \hat{\lambda}_t = \mathbb{E}_t\left(\hat{\lambda}_{t+1} + \hat{r}_t - \hat{\pi}^*_{t+1} - \hat{\pi}_{t+1}\right) \tag{C.33}$$

$$\text{NKPC: } (1+\beta\alpha)\hat{\pi}_t = \alpha\hat{\pi}_{t-1} + \beta\mathbb{E}_t\hat{\pi}_{t+1} + \psi(\hat{a}_t - \hat{\lambda}_t) - \hat{e}_t - \alpha\left(\hat{\pi}^*_t - \beta\mathbb{E}_t\hat{\pi}^*_{t+1}\right) \tag{C.34}$$

$$\text{Monetary policy: } \hat{r}_t - \hat{r}_{t-1} = \rho_{\pi,s^r_t}\hat{\pi}_t + \rho_{gy,s^r_t}\hat{g}^Y_t - \hat{\pi}^*_t + \hat{v}_t \tag{C.35}$$

$$\text{Output growth: } \hat{g}^Y_t = \hat{y}_t - \hat{y}_{t-1} + \hat{z}_t \tag{C.36}$$

$$\text{Real interest rate: } \hat{r}^{r\pi}_t = \hat{r}_t - \hat{\pi}_t \tag{C.37}$$

$$\text{Inflation target change: } \hat{\pi}^*_t = \log(\Pi^*_t) - \log(\Pi^*_{t-1}) \tag{C.38}$$

where $\psi = (\theta-1)/\phi$ is proportional to the slope of the Phillips curve, $\hat{mc}_t = \hat{a}_t - \hat{\lambda}_t$ is the real marginal cost, and the inflation target $\log(\Pi^*_t)$ is governed by the Markov process $s^{\bar{\pi}}_t$. Note that the NKPC contains the term $-\alpha(\hat{\pi}^*_t - \beta\mathbb{E}_t\hat{\pi}^*_{t+1})$, which differs from Ireland (2007) where $\mathbb{E}_t\hat{\pi}^*_{t+1} = 0$ under the random walk assumption for the inflation target. Under our Markov regime-shift assumption, $\mathbb{E}_t\hat{\pi}^*_{t+1} \neq 0$ in general.

The structural shocks follow exogenous processes:

$$\text{Preference shock: } \hat{a}_t = \rho_a\hat{a}_{t-1} + \sigma_a\varepsilon_{at} \tag{C.39}$$

$$\text{Cost push shock: } \hat{e}_t = \rho_e\hat{e}_{t-1} + \sigma_e\varepsilon_{et} \tag{C.40}$$

$$\text{Technology shock: } \hat{z}_t = \sigma_z\varepsilon_{zt} \tag{C.41}$$

$$\text{Monetary policy shock: } \hat{v}_t = \rho_v\hat{v}_{t-1} + \sigma_v\varepsilon_{vt} \tag{C.42}$$

C.5 Mapping the model to the data

The model is estimated on three observable series: output growth, inflation, and the nominal interest rate. The measurement equations linking observables to model variables are:

$$\text{Output growth: } \log(Y_t) - \log(Y_{t-1}) = \hat{g}_t^Y + \log(z) \tag{C.43}$$

$$\text{Inflation: } \log(\Pi_t) = \hat{\pi}_t + \log(\Pi_t^*) \tag{C.44}$$

$$\text{Nominal interest rate: } \log(R_t) = \hat{r} + \log(r) + \log(\Pi_t^*) \tag{C.45}$$

where, contrary to Ireland (2007), we do not impose that $r = z/\beta$ since the additional 20 years of data in our sample implies a substantially lower average real risk-free rate of about 0.36%, which in turn would require $\beta > 1$ since output growth is about 0.5%.

C.6 Solution and estimation

The regime-dependent MSV solution of the model under FIRE is given by

$$X_t = A_{s_t} X_{t-1} + B_{s_t} \epsilon_t \tag{C.46}$$

where s_t now captures the joint Markov state across both the inflation target and monetary policy response coefficient processes. For the estimation, we map a vector of observable variables, Y_t, with the endogenous variables vector X_t:

$$Y_t = \Psi_0(s_t) + \Psi_1 X_t \tag{C.47}$$

where Y_t contains data on output growth, inflation, and the federal funds rate. Vector $\Psi_0(s_t)$ is given by

$$\Psi_0(s_t) = \begin{bmatrix} \log(z) & \log(\Pi^*_{s_t^{\bar{\pi}}}) & \hat{r} + \log(r) + \log(\Pi^*_{s_t^{\bar{\pi}}}) \end{bmatrix}' \tag{C.48}$$

where $\log(z)$ is the average output growth rate, $\log(\Pi^*_{s_t^{\bar{\pi}}})$ is the regime-dependent inflation target, and $\hat{r} + \log(r)$ captures the steady-state nominal interest rate components. The sample covers U.S. quarterly data from 1954:3 to 2024:4.