NBER WORKING PAPER SERIES

A STRUCTURAL APPROACH TO HIGH-FREQUENCY EVENT STUDIES: THE FED AND MARKETS AS CASE HISTORY

Francesco Bianchi
Sydney C. Ludvigson
Sai Ma

Working Paper 30072
http://www.nber.org/papers/w30072

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2022, Revised June 2025

This paper superscedes the old paper “Monetary Based Asset Pricing: A Mixed-Frequency Structural Approach.” We are grateful to Courtney Weigand, Nicholas Zarra, and Steven Zheng for excellent research assistance. Bianchi and Ludvigson received financial support from the National Science Foun- dation under Grant 2116641. We thank Adrien Auclert, Michael Bauer, Anna Cieslak, Andrea Ferrero, Eric Ghysels, Jonas Jensen, Carolin Pflueger, Frank Schorfheide, and seminar participants at the Spring 2023 NBER Asset Pricing meeting, the 2022 American Economic Association annual meetings, the May 2022 Bank of England Macro-Finance Conference, the June 2022 EDHEC PhD Forum, the September 2022 SOFIE webinar, the September 2022 Bank of Finland and CEPR Joint Conference on Monetary Policy in the Post-Pandemic Era, the November 2022 NBER Monetary Economics meeting, Harvard Business School, Notre Dame, HKUST, Bocconi University, McGill, MIT, and the University of Chicago for helpful comments. The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Board, the Federal Reserve System, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Francesco Bianchi, Sydney C. Ludvigson, and Sai Ma. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Structural Approach to High-Frequency Event Studies: The Fed and Markets as Case History
Francesco Bianchi, Sydney C. Ludvigson, and Sai Ma
NBER Working Paper No. 30072
May 2022, Revised June 2025
JEL No. E52, E58, E7, G12

ABSTRACT

We develop a methodology to integrate a high-frequency event study into a macro-finance model and structural estimation. The methodology is applied to Federal Reserve announcements in a model where investor beliefs about the economic state and/or regime change in future policy can jump in response to monetary news. We find that stock market volatility in narrow windows around policy announcements is frequently driven by jumps in beliefs about future policy rules that affect subjective risk premia. Such jumps often generate positive comovement between short rates and the stock market, erroneously suggesting a role for "Fed information shocks."

Francesco Bianchi
Johns Hopkins University
Department of Economics
and CEPR
and also NBER
francesco.bianchi@jhu.edu

Sydney C. Ludvigson
New York University
Department of Economics
and NBER
sydney.ludvigson@nyu.edu

Sai Ma
Federal Reserve Board of Governors
sai.ma@frb.gov

A online appendix is available at http://www.nber.org/data-appendix/w30072

1 Introduction

A growing academic literature has offered a myriad of competing explanations for why financial markets react strongly to the announcements of central banks. A classic New-Keynesian view is that surprise central bank announcements proxy for shocks to a nominal interest rate rule of the type emphasized by Taylor (1993), which influence the real economy through an intertemporal substitution channel. More recently, other theories emphasize the effects of central bank announcements on financial market risk premia (the risk premium channel), the information they impart about the state of the economy (the "Fed information shock" channel), or the role they play in revising the public's understanding of the central bank's reaction function and objectives (the perceived policy rule channel).

As the mushrooming debate over how to interpret this evidence indicates, many questions about the interplay between central bank news and financial markets remain unanswered. In this paper we consider two of them. First, although extant theories typically focus on a single channel of monetary transmission, there is no obvious reason why multiple channels might not be operative simultaneously. We therefore consider whether more than one channel might explain why markets react to central bank news, and quantify their roles. Second, much like other types of news to which markets are strongly attentive, monetary news typically covers a wide-ranging set of themes. In the case of announcements by the U.S. Federal Reserve, these themes often range broadly from interest rate policy, to forward guidance, to quantitative interventions, to the macroeconomic outlook, a plethora offering little reason to expect that the forces driving reactions to one announcement should be similar to another. We therefore consider announcement-specific reactions that differ according to how market participants' perceptions of primitive economic risks have changed as a result of the news.

Addressing these questions requires a high-frequency event study that can allow for distinct theoretical channels of varying degrees of importance across event-time. Unfortunately, existing event study approaches are poorly suited to the task, both because they provide little to no announcement-specific evidence or conceptual framework, and because they offer limited insight into *why* market participants react to specific announcements. When beliefs react to news causing equilibrium prices to change, which primitive shocks did investors perceive and what was their quantitative relation to the resulting price change?

Our contribution to addressing these questions is to integrate a high-frequency event study into a mixed-frequency macro-finance model and structural estimation. The resulting 'mixed-frequency structural approach' to event study analysis allows us to address the question of why–as seen through the lens of a theoretical framework–financial markets reacted to a given piece of news the way they did, effectively identifying state-dependent reaction functions to real-world events. The general approach can be applied in a wide variety of structural and semi-structural settings whenever the research objective calls for a conceptual understanding of financial market responses to any type of news or event.[1]

In this article, we apply the approach to undertake a case study of the stock market's reactions to Federal Reserve (Fed) communications. Our structural estimation uses dozens of forward-looking series ranging from minutely financial market data to biannual survey forecast data. These data, along with the model and structural estimation, allow us to infer jumps in investor beliefs about the latent state of the economy, the perceived sources of primitive economic risk driving that state, and the future conduct of monetary policy, all in response to Fed news. In turn, movements in these beliefs generate endogenous movements in subjective risk premia.

We study a two-agent asset pricing model with New Keynesian style macroeconomic dynamics in which the two agents have heterogeneous beliefs, as in Bianchi, Lettau, and Ludvigson (2022). One agent is a representative "investor" who is forward-looking, reacts swiftly to news, and earns income solely from investments in the stock market and a one-period nominal bond. Macroeconomic dynamics are specified by a set of equations similar to those in New Keynesian models, and can be thought of as driven by a representative "household/worker" that supplies labor, has access to the nominal bond, but holds no stock market wealth. Unlike investors, the household/worker forms expectations in a backward-looking manner using adaptive learning rules, consistent with survey evidence on how households actually arrive at their beliefs (Malmendier and Nagel (2016), Coibion, Gorodnichenko, Kumar, and Pedemonte (2020)).

An important feature of our model is that the conduct of monetary policy is allowed to undergo possible nonrecurrent regime shifts, or "structural breaks," that take the form of shifts in the parameters of a nominal interest rate rule. Such regime changes in what we refer to as the *conduct* of monetary policy give rise to endogenously long-lasting changes in real interest rates that are conceptually distinct from those generated by the monetary policy *shock*, an innovation in the nominal rate that is uncorrelated with inflation, economic growth, and shifts in the policy rule parameters.

With this feature at hand, we explicitly model investor beliefs about future regime change in the conduct of monetary policy. Investors in the model closely monitor central bank communications for information that would lead them to revise their perceived probability of transitioning out of the current policy regime into a perceived "Alternative regime" that they believe will come next. Investors are aware that they may change their minds about the likelihood of near-term monetary regime change given new information,

[1] Bianchi, Ludvigson, and Ma (2024) apply the approach to study over- and under-reaction in the stock market to macroeconomic data announcements, corporate earnings announcements, as well as Fed news.

and take that into account when forming expectations.

The precise nature of relevant new information can take any form, but in this study we focus on Fed announcements. A Fed announcement in our model is a bonafide news shock to which investors may react by revising their nowcasts and forecasts of the current and future economic state, their beliefs about the future conduct of monetary policy, and their perceptions of financial market risk. By allowing for news-driven revisions in expectations about the economic state, we can directly evaluate the extent to which "Fed information shocks" drive market volatility around policy announcements. To ensure that model expectations evolve in a manner that closely aligns with observed expectations, we map data on numerous forward-looking variables, including multiple surveys of expectations as well as financial market indicators from spot and futures markets, onto model-implied equations for investor beliefs and expectations, estimating all parameters and latent states.

Our main empirical results may be summarized as follows. First, the structural estimation implies that investors seldom learn only about conventional monetary policy shocks from central bank announcements. Instead, jumps in financial market variables in response to central bank communications are caused by a mix of factors that are the product of context-specific forces. For each event, the methodology allows us to quantify their relative importance.

For example, on January 3, 2001 the Fed surprised markets by reducing its target for the federal funds rate (FFR) by 50 basis points, resulting in a 4.2% vault in the stock market over the 20 minutes following the announcement. Our estimates imply that the perception of an accommodative policy shock to the short rate played only a small first-order role in the stock market surge. Instead we find that most quantitatively important contributors were a downward revision in investor expectations of financial market liquidity premia, and an upward revision in expectations of corporate earnings as a share of aggregate output. These revisions in beliefs are identified in the data by jumps in credit spreads and in professional survey nowcasts of S&P 500 earnings measured immediately before and after the policy announcement.

This event can be contrasted with the FOMC announcement of April 18, 2001, in which the market increased 2.5% after the Greenspan Fed again surprised with another 50 basis point reduction in the funds rate. In this case, the biggest estimated driver of the stock market jump was a change in investor beliefs about the conduct of future monetary policy. Specifically, we find that this announcement coincided with a shift in investor beliefs that drove markets to price in a higher likelihood that the Fed would soon shift its policy rule toward one that more aggressively stabilized fluctuations in economic growth. This belief is tantamount to predicting a stronger "Fed put," in which the central bank engineers a lower discount rate whenever future cash flows are expected to grow more slowly, causing subjective risk premia to decline and stock prices to rise. These estimates are identified in the data by jumps in both the stock market and in Fed funds futures rates that cannot be explained by rolling forward the current policy rule. The results for this event illustrate an additional channel of monetary transmission to markets, namely the role of Fed communications in altering investor beliefs about future Fed policy to contain economic risks, thereby immediately impacting subjective risk premia.

Our second main finding pertains to the role of "Fed information shocks." Following Jarocinski and Karadi (2020), we define an event with a Fed information shock as one that features a positive comovement between short-term interest rates and the stock market within a narrow window around a policy announcement. We find that the most quantitatively important announcements of this type in our sample (measured by how much the stock market responded) are difficult to reconcile with standard notions of a Fed information shock. For example, on October 29, 2008, the Fed announced that it would lower its target for the FFR by 50 basis points to 1%, just matching its lowest previous level. Despite the decline in rates, the stock market fell 2% in the 20 minutes following the announcement. Why? According to our estimates it was not because investors learned of new negative information about the economic outlook from the announcement. If anything our estimates imply that overall perceptions about the economic state, as seen in e.g., credit spreads and high-frequency nowcasts of corporate earnings, were revised in *favorable* directions around the policy announcement. This suggests that investors believed the rate cut would at least help to support the economy. But these reasons for optimism were more than offset by a newly found pessimism that future monetary policy would soon shift toward a more growth-sensitive policy rule. The net effect on markets pushed stocks downward, culminating in a 2% decline. The findings for this announcement are consistent with reports that some investors were disappointed by the relatively modest size of the cut given the unprecedented magnitude of the global financial crisis.

Finally, our results indicate that shifts in investor beliefs about future policy conduct around policy announcements are important for stock market fluctuations primarily because of their role in shaping the perceived quantity of stock market risk and thus subjective risk premia. By contrast, jumps in beliefs about the future policy rule do not appear to operate primarily through expected future short rates or cash flow growth.

Related Literature The research in this paper connects with a large and growing body of evidence that the values of long-term financial assets and expected return premia respond sharply to the announcements of central banks.[2] A classic assumption of this

[2] See Cochrane and Piazzesi (2002), Piazzesi (2005), Bernanke and Kuttner (2005), Krishnamurthy and Vissing-Jorgensen (2011), Hanson and Stein (2015), Gertler and Karadi (2015), Gilchrist, López-Salido, and Zakrajšek (2015), Brooks, Katz, and Lustig (2018), Kekre and Lenel (2021), and Pflueger

literature is that high-frequency financial market reactions to Fed announcements proxy for conventional monetary policy "shocks," i.e., innovations in a Taylor (1993)-type nominal interest rate rule. By contrast, Jarocinski and Karadi (2020), Cieslak and Schrimpf (2019) and Hillenbrand (2021) argue that some of the fluctuations are likely driven by the revelation of information by the Fed, a "Fed information effect" channel emphasized in earlier work by Romer and Romer (2000), Campbell, Evans, Fisher, and Justiniano (2012), Melosi (2017), and Nakamura and Steinsson (2018). Related, Cieslak and Pang (2021) identify monetary, growth, and risk premium shocks from Fed news using sign-restricted VARs. Bauer and Swanson (2023) instead argue that markets are surprised by the Fed's response to recent economic events, while Bauer, Pflueger, and Sundaram (2022) use survey data to estimate perceived policy rules, finding that they are subject to substantial time-variation. The mixed-frequency structural approach proposed in this paper extends these literatures by allowing for announcement-specific results seen through the lens of a structural asset pricing model, and by providing evidence that expected return premia vary, in part, because the perceived quantity of stock market risk fluctuates around policy announcements due to changing beliefs about future (rather than current) monetary policy conduct.

Our work relates to a theoretical literature focused on the implications of monetary policy for asset prices going back to Piazzesi (2005). Kekre and Lenel (2021) and Pflueger and Rinaldi (2020) develop carefully calibrated theoretical models that imply stock market return premia vary in response to a monetary policy shock. These theories use different mechanisms but are all silent on the possible role of Fed announcement effects or of changing policy rules in driving market fluctuations, features that are at the heart of our analysis. Our structural estimates are consistent with the interpretation of Cieslak and Vissing-Jorgensen (2021) based on textual analyses of FOMC documents and indicate that subjective risk premia vary with beliefs about the strength of a "Fed put" on the stock market.

The two-agent structural model of this paper builds on Bianchi, Lettau, and Ludvigson (2022) (BLL hereafter), who focus on the low frequency implications for asset valuations of changes in the conduct of monetary policy. The mixed-frequency structural approach of this paper offers a significant methodological advancement over BLL and to the best of our knowledge the extant literature, by developing a methodology to provide theoretical insights into why markets respond to news, including estimating revisions in investor beliefs about future monetary policy in the minutes surrounding Fed announcements. We show how forward looking variables, such as survey expectations and asset prices, can be used both to estimate the market's perceived probability of a near-term policy regime change, and to extract beliefs about the nature of future policy regimes.

and Rinaldi (2020).

BLL, and more recently Nagel (2024), underscore the importance of adopting empirical realism in the modeling of household inflation expectations for studying how monetary policy affects financial markets. Extensive evidence from survey microdata shows that households do not form expectations rationally according to the rules assumed by classic New Keynesian models, among others. Instead, survey evidence is consistent with the hypothesis that monetary policy must pursue large and persistent changes in real rates in order to counter inertia in household expectations that–unlike markets–are relatively impervious to Fed announcements. Such persistent monetary non-neutrality rationalizes why markets for ultra long-duration assets like stocks pay so much attention to Fed announcements in the first place.

In contemporaneous work, Caballero and Simsek (2022) study a two-agent, "two-speed" economy with investors and households similar in spirit to BLL. Their purely theoretical investigation studies the relation between Fed policy and a broad-based financial conditions index allowing for two-way feedback between financial conditions and the economy. Our work is an empirical complement that estimates the high-frequency impact of Fed news on the stock market. In future work, we plan to extend our structural estimation to allow for simultaneous feedback between markets, the Fed, and the economy.

Finally, our mixed-frequency structural approach connects with a pre-existing reduced-form forecasting/nowcasting literature using mixed-frequency data in state space models with the objective of augmenting lower frequency prediction models with more timely high-frequency data (e.g., Giannone, Reichlin, and Small (2008), Ghysels and Wright (2009), Schorfheide and Song (2015)). Our use of mixed-frequency data is designed for a very different purpose, namely as way of modeling and measuring the effects of news shocks. We use high-frequency, forward-looking data available *within* the decision interval to infer revisions in the intraperiod beliefs of investors about the economic state to be realized at the *end* of the decision interval. This allows us to treat Fed announcements as bonafide news shocks (as perceived by investors) rather than as ultra high frequency macro shocks that happen to occur around Fed communications.

The rest of this paper is organized as follows. The next section presents preliminary empirical evidence that we use to pin down the timing of monetary regime changes in our sample. Section 3 describes the mixed-frequency structural macro-finance model and equilibrium solution. Section 4 describes the structural estimation, while Section 5 presents our empirical findings from the structural estimation. Section 6 concludes. A large amount of additional material on the model, estimation, and data has been placed in an Online Appendix.

2 Preliminary Evidence

In the structural model of the next section, investors form beliefs about future regime change in the conduct of monetary policy. We therefore begin by presenting preliminary evidence suggestive of infrequent, sizable shifts in the conduct of monetary policy over our the course of our sample. Our evidence on break dates (if any) are dictated by the data and freely estimated.

Figure I panel (a) plots the real federal funds rate, r_t, measured for the purposes of this plot in real terms as the nominal rate minus a four quarter moving average of inflation, while panel (b) plots the difference between this rate and an estimate of the neutral rate of interest, r_t^*, from Laubach and Williams (2003), effectively a low-frequency component of the real funds rate. We refer to the spread between r_t and r_t^* as the *monetary policy spread*, and denote its time t value as mps_t. The mps_t may be considered a crude measure of the stance of monetary policy, i.e., whether monetary policy is accommodative or restrictive.

We allow for the possibility of infrequent regime changes in the means of r_t and mps_t, governed by a discrete valued latent state variable, ξ_t^P that is presumed to follow a N_P-state *nonrecurrent* regime-switching Markov process, i.e., structural breaks. That is, there is no expectation that regime shifts in the means of either variable must move to a new regime that is *identically* equal to one in the past (mathematically a probability zero event), though it could be quite similar. This specification is both more plausible and more flexible than recurrent regime switching, since the estimation is free to choose parameter values across regimes that are arbitrarily close to those that have occurred in the past without restricting them to be identically equal.[3]

Figure I reports the results for the case of two structural breaks ($N_P = 3$) estimated separately for r_t and mps_t, with the estimated regime subperiods reported in the figure notes. Regardless of whether we measure breaks in the mean of r_t or mps_t, the break dates are identical and thus so are the regime subperiods. This shows that the chosen r_t^* measure–while useful to get a sense of the persistence of swings in r_t around that low frequency component–has no influence on the estimated break dates.

We identify decades-long breaks in both r_t and mps_t, consistent with an earlier literature documenting that monetary regime changes are infrequent (Clarida, Gali, and Gertler (2000); Lubik and Schorfheide (2004); Sims and Zha (2006); Bianchi (2013)). The first estimated subperiod spans 1961:Q1 to 1978:Q3, a time period in which r_t was low and mps_t was persistently negative. This "Great Inflation" regime coincides with a run up in inflation and with two oil shocks in the 1970s that were arguably exacerbated by a Fed that failed to react sufficiently proactively. A second, "Great Moderation,"

[3]Details of this procedure are provided in Appendix C of the Online Appendix.

regime begins in 1978:Q4, when a structural break drove upward jumps in both r_t and mps_t. This period of more restrictive monetary policy lasted through 2001:Q3 and covers the Volcker disinflation and moderation in economic volatility that followed. The third, "Post Millennial," regime spans 2001:Q4 to 2020:Q1 and represents a new prolonged period of low real interest rates. The beginning of this regime follows shortly after the inception of public narratives on the "Greenspan Put," the perceived attempt of Chair Greenspan to prop up securities markets in the wake of the IT bust, a recession, and the aftermath of 9/11, by lowering interest rates. This low rate subperiod continues with the explicit forward guidance "low-for-long" policies under Chair Bernanke that repeatedly promised over several years to keep interest rates at ultra low levels for an extended period of time. Below we refer to the Great Inflation, the Great Moderation and the Post Millennial regimes in abbreviated terms as the GI, GM, and PM regimes.

Figure I: Breaks in Monetary Policy

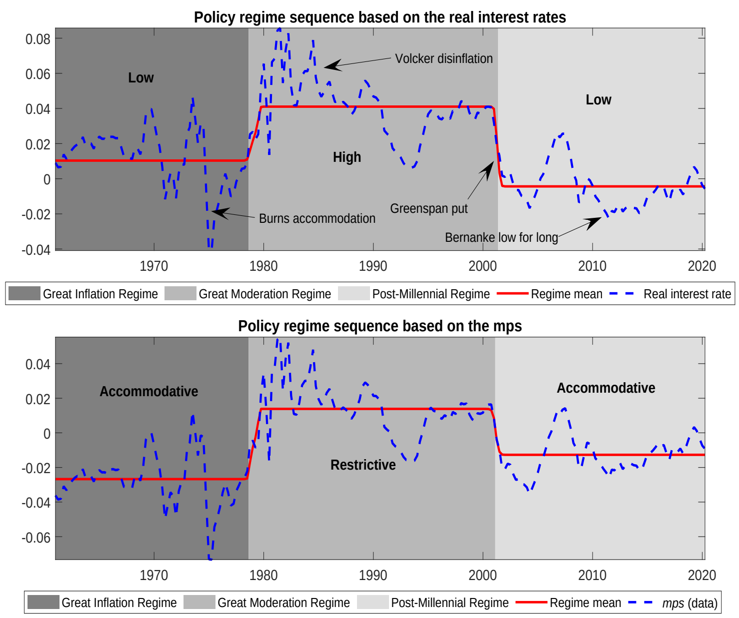

Notes: Monetary policy spread $mps_t \equiv FFR_t - \text{Expected Inflation}_t - r_t^*$. r^* is from Laubach and Williams (2003). The blue (dashed) line represents the data. The red (solid) line is the estimated regime mean of each series. Great Inflation Regime: 1961:Q1-1978:Q3. Great Moderation Regime: 1978:Q4-2001:Q3. Post-Millennial Regime: 2001:Q4-2020:Q1. The sample spans 1961:Q1-2020:Q1.

The preliminary evidence of this section is used to set the *timing* of possible policy regime changes in the structural model.[4] To accomplish this, we set the break dates for regime changes in the model's policy rule to coincide with the regime sequence for ξ_t^P displayed in Figure I. It should be emphasized, however, that all regime-dependent

[4] This procedure allows us to build a structural model to fit large observed breaks in the behavior of the policy instrument, rather than choosing break dates that fit a detailed structural model.

parameters of the policy rule are freely estimated under symmetric priors, thus are treated as equally likely to increase or decrease across the regime subperiods for ξ_t^P, if they change at all. We use Bayesian model comparison of estimated structural models that differ according to the number of policy regimes to decide on the appropriate number and find $N_P = 3$ works well. Our structural estimation therefore spans three different policy regimes across the Great Inflation, the Great Moderation, and the Post Millennial subperiods shown in Figure I.

3 Mixed-Frequency Macro-Finance Model

This section presents a two-agent dynamic asset pricing model of monetary policy transmission. Risky asset prices are determined by the behavior of a forward-looking representative investor who reacts swiftly to news and forms beliefs about future monetary policy. Households/workers supply labor, invest only in the bond, and form expectations using adaptive learning rules that predominate in aggregate inflation and output growth expectations. As in BLL, it is through such heterogeneity in beliefs that regime changes in the conduct of monetary policy have large and prolonged effects on real interest rates, despite the forward-looking, non-inertial nature of market participant expectations. We work with a risk-adjusted loglinear approximation to the model that can be solved analytically, in which all random variables are conditionally lognormally distributed.

Let the "decision" interval t of both agents be monthly and let lowercase variables denote log variables, e.g., $\ln(D_t) = d_t$. For investors, this means that they receive payout and can only observe the economic state S_t at the end of each month. However, as explained below, they nevertheless price assets continuously and update expectations in the wake of Fed announcements.

Asset Pricing Block Assets are priced by a representative investor who consumes per-capita aggregate shareholder payout, D_t, and earns all income from trade in two assets: a one-period nominal risk-free bond and a stock market. The investor's intertemporal marginal rate of substitution in consumption is the stochastic discount factor (SDF) and its logarithm takes the form:

$$m_{t+1} = \ln\left(\beta_p\right) + \vartheta_{pt} - \sigma_p\left(\Delta d_{t+1}\right). \tag{1}$$

where σ_p is a relative risk aversion coefficient and $\ln\left[\beta_p \exp\left(\vartheta_{pt}\right)\right]$ is a subjective time discount factor that varies over time with the patience shifter ϑ_{pt}. Individual investors take ϑ_{pt} as given, driven by the market as a whole.[5] A time-varying specification for the subjective time-discount factor is essential for ensuring that, in equilibrium, investors

[5] This specification for ϑ_{pt} is a generalization of those considered in previous work (e.g., Campbell and Cochrane (1999) and Lettau and Wachter (2007)) where the preference shifter is taken as an exogenous

are willing to hold the nominal bond at the interest rate set by the central bank's policy rule, specified below.

Aggregate payout is a time-varying share K_t of real output Y_t, implying $D_t = K_t Y_t$ or in logs $d_t - \ln(Y_t) = k_t$. Since in the model all earnings are paid out to shareholders, we refer to K_t simply as the *earnings share* hereafter. Variation in k_t, follows an exogenous primitive process:

$$k_t - \overline{k} = (1-\rho_k)\,\lambda_{k,\Delta y}\Delta y_t + \rho_k\left(k_{t-1} - \overline{k}\right) + \sigma_k \varepsilon_{k,t}.$$

Thus k_t varies with economic growth and an independent i.i.d. shock $\varepsilon_{kt} \sim N(0,1)$.

The first-order-condition for optimal holdings of the one-period nominal risk-free bond with a face value equal to one nominal unit is

$$LP_t^{-1} Q_t = \mathbb{E}_t^b\left[M_{t+1}\Pi_{t+1}^{-1}\right], \tag{2}$$

where Q_t is the nominal bond price, $\mathbb{E}_t^b$ denotes the subjective expectations of the investor, and $\Pi_{t+1} = P_{t+1}/P_t$ is the gross rate of general price inflation. Investors' subjective beliefs, indicated with a "b" superscript, play a central role in asset pricing and are discussed in detail below. Investors have a time-varying preference for nominal risk-free assets over equity, accounted for by $LP_t > 1$, implying that Q_t is higher than it would be absent these benefits, i.e., when $LP_t = 1$.

Taking logs of (2) and using the properties of conditional lognormality delivers the real interest rate as perceived by the investor:

$$i_t - \mathbb{E}_t^b\left[\pi_{t+1}\right] = -\mathbb{E}_t^b\left[m_{t+1}\right] - .5\mathbb{V}_t^b\left[m_{t+1} - \pi_{t+1}\right] - lp_t \tag{3}$$

where $i_t = -\ln(Q_t)$, $\pi_{t+1} \equiv \ln(\Pi_{t+1})$ is net inflation, $\mathbb{V}_t^b[\cdot]$ is the conditional variance under the subjective beliefs of the investor, and $lp_t \equiv \ln(LP_t) > 0$. Variation in lp_t follows an AR(1) process

$$lp_t - \overline{lp} = \rho_{lp}\left(lp_{t-1} - \overline{lp}\right) + \sigma_{lp}\varepsilon_{lp,t}$$

subject to an i.i.d. shock $\varepsilon_{lp,t} \sim N(0,1)$.

Let P_t^D denote total value of market equity, i.e., price per share times shares outstanding. Optimal shareholder consumption obeys the following log Euler equation:

$$\begin{aligned} pd_t &= \kappa_{pd,0} + \mathbb{E}_t^b\left[m_{t+1} + \Delta d_{t+1} + \kappa_{pd,1} pd_{t+1}\right] + \\ &\quad + .5\mathbb{V}_t^b\left[m_{t+1} + \Delta d_{t+1} + \kappa_{pd,1} pd_{t+1}\right], \end{aligned}$$

process that is the same for each shareholder. Combining (1) and (3) below, we see that $\vartheta_{p,t}$ is implicitly defined as

$$\vartheta_t^p = -\left[i_t - \mathbb{E}_t^b\left[\pi_{t+1}\right]\right] + \mathbb{E}_t^b\left[\sigma_p \Delta d_{t+1}\right] - .5\mathbb{V}_t^b\left[-\sigma_p \Delta d_{p,t+1} - \pi_{t+1}\right] - lp_t - \ln\left(\beta_p\right).$$

where $pd_t \equiv \ln\left(P_t^D/D_t\right)$. The log equity return $r_{t+1}^D \equiv \ln\left(P_{t+1}^D + D_{t+1}\right) - \ln\left(P_t^D\right)$ obeys the following approximate identity (Campbell and Shiller (1989)):

$$r_{t+1}^D = \kappa_{pd,0} + \kappa_{pd,1} pd_{t+1} - pd_t + \Delta d_{t+1},$$

where $\kappa_{pd,1} = \exp(\overline{pd})/(1+\exp(\overline{pd}))$, and $\kappa_{pd,0} = \log\left(\exp(\overline{pd}) + 1\right) - \kappa_{pd,1}\overline{pd}$. Combining the above, the log equity premium as perceived by the investor is:

$$\underbrace{\mathbb{E}_t^b\left[r_{t+1}^D\right] - \left(i_t - \mathbb{E}_t^b\left[\pi_{t+1}\right]\right)}_{\text{subj. equity premium}} = \underbrace{\begin{bmatrix} -.5\mathbb{V}_t^b\left[r_{t+1}^D\right] - \mathbb{COV}_t^b\left[m_{t+1}, r_{t+1}^D\right] \\ +.5\mathbb{V}_t^b\left[\pi_{t+1}\right] - \mathbb{COV}_t^b\left[m_{t+1}, \pi_{t+1}\right] \end{bmatrix}}_{\text{subjective risk premium}} + \underbrace{lp_t}_{\text{liquidity Premium}}, \quad (4)$$

where $\mathbb{COV}_t^b[\cdot]$ is the investor's subjective conditional covariance.

The equity premium has two components, a subjective risk premium is attributable to the agent's subjective perception of risk, and a "liquidity premium" lp_t that represents a time-varying preference for risk-free nominal debt over equity. The subjective risk premium varies endogenously in the model with fluctuations in investor beliefs about the conduct of future monetary policy, as explained below. The liquidity premium captures all sources of time-variation in the equity premium other than those attributable to subjective beliefs about the monetary policy rule. These could include variation in the liquidity and safety attributes of nominal risk-free assets (e.g., Krishnamurthy and Vissing-Jorgensen (2012)), variation in risk aversion, flights to quality, or jumps in sentiment.

Macro Dynamics Macroeconomic dynamics are described by a set of equations similar to prototypical New Keynesian models, but with two distinctive features: adaptive learning, and regime changes in the conduct of monetary policy. Strictly speaking, we consider equations (5) through (7) below as equilibrium dynamics and not a micro-founded structural model. We consider an equilibrium in which bonds are in zero-net-supply in both the macro and asset pricing blocks and thus there is no trade between the investor and households.[6]

Let $\ln(A_t/A_{t-1}) \equiv g_t$ represent the stochastic trend growth of the economy, which follows an AR(1) process $g_t = g + \rho_g(g_{t-1} - g) + \sigma_g \varepsilon_{g,t}$, $\varepsilon_{g,t} \sim N(0,1)$. Log of detrended output in the model is defined as $\ln(Y_t/A_t)$. Let variables with tildes, e.g., $\widetilde{y}_t = \ln(Y_t/A_t)$, denote detrended variables. Thus $\widetilde{y}_t > 0$ (< 0) when y_t is above (below)

[6]Models with trade are computationally slow to solve and would present a significant challenge to estimation; hence we leave this to future research. However, an empirically plausible version of our model with trade may not imply appreciably different aggregate dynamics. For example, Chang, Chen, and Schorfheide (2021) provide econometric evidence that spillovers between aggregate and distributional dynamics in heterogeneous agent models are generally small.

potential output, so $\widetilde{y}_t \neq 0$ can be interpreted as a New Keynesian output gap. In keeping with New Keynesian models, we write most equations in the macro block in terms of detrended real variables.

Macroeconomic dynamics satisfy a loglinear Euler or "IS" equation that is a function of household consumption $(1 - K_t)Y_t$:[7]

$$\widetilde{y}_t = \mathbb{E}_t^m(\widetilde{y}_{t+1}) - \sigma\left[i_t - \mathbb{E}_t^m(\pi_{t+1}) - \overline{r}\right] + f_t \tag{5}$$

where $\mathbb{E}_t^m(\cdot)$ is the expectation under the subjective beliefs of the macro agent, $\overline{r}$ is the steady state real interest rate, and f_t is a demand shock that also absorbs any variation in the macro agent's consumption attributable to movements in the labor share, $\ln(1 - K_t)$. The demand shock follows an AR(1) process $f_t = \rho_f f_{t-1} + \sigma_f \varepsilon_f$, $\varepsilon_f \sim N(0,1)$. The coefficient σ in (5) is a positive parameter.

Inflation dynamics are described by the following equation, which takes the form of a New Keynesian Phillips curve:

$$\begin{aligned}\pi_t - \overline{\pi}_t &= \beta\left(1 - \lambda_{\pi,1} - \lambda_{\pi,2}\right)\mathbb{E}_t^m\left[\pi_{t+1} - \overline{\pi}_t\right] + \beta\lambda_{\pi,1}\left[\pi_{t-1} - \overline{\pi}_t\right] \\ &\quad + \beta\lambda_{\pi,2}\left[\pi_{t-2} - \overline{\pi}_t\right] + \kappa_0\widetilde{y}_t + \kappa_1\widetilde{y}_{t-1} + \sigma_\mu\varepsilon_{\mu,t}\end{aligned} \tag{6}$$

where $\overline{\pi}_t$ denotes the household's perceived trend inflation rate (specified below) and $\varepsilon_{\mu,t} \sim N(0,1)$ is a markup shock.[8] Lags beyond the current values of variables are used to capture persistent inflation dynamics. The coefficients β, λ_{π_1}, λ_{π_2}, κ_0, and κ_1 are positive parameters.

The central bank obeys the following nominal interest rate rule subject to nonrecurrent regime changes in its parameters:

$$\begin{aligned}i_t - \left(\overline{r} + \pi^T_{\xi^p_t}\right) &= \left(1 - \rho_{i,\xi^p_t} - \rho_{i_2,\xi^p_t}\right)\left[\psi_{\pi,\xi^p_t}\widehat{\pi}_{t,t-3} + \psi_{\Delta y,\xi^p_t}\left(4\widehat{\Delta y}_{t,t-3}\right)\right] \\ &\quad + \rho_{i_1,\xi^p_t}\left[i_{t-1} - \left(\overline{r} + \pi^T_{\xi^p_t}\right)\right] + \rho_{i_2,\xi^p_t}\left[i_{t-2} - \left(\overline{r} + \pi^T_{\xi^p_t}\right)\right] + \sigma_i\varepsilon_i.\end{aligned} \tag{7}$$

The central bank is presumed to react to quarterly data (at monthly frequency) given that it is unlikely to react to the more volatile monthly variation in growth and inflation. Thus $\widehat{\pi}_{t,t-3} \equiv \sum_{l=0}^{2}\left(\pi_{t-l} - \pi^T_{\xi^p_t}\right)$ is quarterly inflation in deviations from the implicit time t target $\pi^T_{\xi^p_t}$, $4\Delta y_{t,t-3} \equiv 4\sum_{l=0}^{2}(\Delta y_t - g)$ is annualized quarterly output growth in deviations from steady-state growth g, and $\varepsilon_{i,t} \sim N(0,1)$ is an i.i.d. monetary policy

[7] We assume that the Euler equation (5) holds under nonrational expectations. Honkapohja, Mitra, and Evans (2013) provide microfoundations for such Euler equations with nonrational beliefs.

[8] This equation can be micro-founded by assuming that managers of firms are workers who form expectations as households/workers do rather than as shareholders do, consistent with evidence that the discount rates managers use when making investment and employment decisions are different from those observed in financial markets (Gormsen and Huber (2022)), and with evidence that those expectationsn do not appear rational (Gennaioli, Ma, and Shleifer (2016)).

shock. Lags of the left-hand-side variable appear in the rule to capture the observed smoothness in adjustments to the central bank's target interest rate.

The interest rate policy rule allows for nonrecurrent regime changes in the conduct of monetary policy driven by ξ_t^p, which indexes changes in the parameters of (7). The parameter $\pi^T_{\xi_t^P}$ plays the role of an *implicit* time-t inflation target. In particular, this time-varying parameter may deviate from the central bank's stated long-term inflation objective, such as when it is engaging in forward guidance to actively move inflation back toward that objective. In that case, a higher than usual inflation target implies low rates for long and a signal about tolerance for higher inflation. The activism coefficients ψ_{π,ξ_t^P}, and $\psi_{\Delta y,\xi_t^P}$ govern how strongly the central bank responds to deviation from the implicit target and to economic growth and are also subject to regime shifts, as are the autocorrelation coefficients ρ_{i,ξ_t^P} and ρ_{i_2,ξ_t^P}. We treat shifts in the policy rule parameters as exogenous and latent parameters to be estimated.[9] These coefficients vary with ξ_t^P and the identified regime sequence for $r_{\xi_t^P}$ from Figure I. It is important to emphasize, however, that we freely estimate the policy rule parameters under symmetric priors, so they could in principle show no shift across regimes.

We assume that households form expectations about inflation using an adaptive algorithm on the autoregressive process, $\pi_t = \alpha + \phi\pi_{t-1} + \eta_t$, where the agent must learn about α. Each period, agents update a belief α_t^m about α. Define *perceived trend inflation* to be the $\lim_{h\to\infty} \mathbb{E}_t^m [\pi_{t+h}]$ and denote it by $\overline{\pi}_t$. Given the presumed autoregressive process, it can be shown that $\overline{\pi}_t = (1-\phi)^{-1} \alpha_t^m$ and that $\mathbb{E}_t^m [\pi_{t+1}] = (1-\phi)\overline{\pi}_t + \phi\pi_t$.

We allow the evolution of beliefs about α_t^m and $\overline{\pi}_t$ to potentially reflect both adaptive learning as well as a signal about the central bank's inflation target that could reflect the opinion of experts (as in Malmendier and Nagel (2016)) or a credible central bank announcement. For the adaptive learning component, we follow evidence in Malmendier and Nagel (2016) that the University of Michigan Survey of Consumers (SOC) mean inflation forecast is well described by a constant gain learning algorithm. Combining these yields updating rules for α_t^m and $\overline{\pi}_t$:

$$\alpha_t^m = \left(1-\gamma^T\right)\left[\alpha_{t-1}^m + \gamma\left(\pi_t - \phi\pi_{t-1} - \alpha_{t-1}^m\right)\right] + \gamma^T\left[(1-\phi)\,\pi^T_{\xi_t}\right] \quad (8)$$

$$\overline{\pi}_t = \left(1-\gamma^T\right)\left[\overline{\pi}_{t-1} + \gamma(1-\phi)^{-1}\left(\pi_t - \phi\pi_{t-1} - (1-\phi)\,\overline{\pi}_{t-1}\right)\right] + \gamma^T\pi^T_{\xi_t}, \quad (9)$$

where γ is the constant gain parameter that governs how much last period's beliefs α_{t-1}^m and $\overline{\pi}_{t-1}$ are updated given new information, π_t. The second term in square brackets captures the effect of the signal about the implicit inflation target $\pi^T_{\xi_t}$. The parameter

[9]This approach side-steps the need to take a stand on why the Fed changes its policy rule, an important consideration given that the reasons for such changes would be difficult if not impossible to credibly identify as a function of past historical data, due to the degree of discretion the Fed has in interpreting its dual mandate and the likelihood that distinct regimes are the result of a gradual learning process interacting with the bespoke perspectives of different central bank leaders across time.

γ^T controls the informativeness of the inflation target signal. If $\gamma^T = 1$, the signal is perfectly informative and the household's belief about trend inflation is the same as the implicit target. If $\gamma^T = 0$, the signal is completely uninformative and the agent's belief about trend inflation depends only on the adaptive learning algorithm. A weight of $\gamma^T < 1$ could arise either because the target is imperfectly observed, or because central bank announcements about the target are not viewed as fully informative or credible. Small values for γ^T are indicative of slow learning and low central bank credibility, since in that case the macro agent continues to base inflation expectations mostly on a backward looking rule even when there has been a shift in the inflation target.

Finally, expectations about detrended output follow a simple backward looking rule:

$$\mathbb{E}_t^m\left(\widetilde{y}_{t+1}\right) = \varrho_1 \widetilde{y}_{t-1} + \varrho_2 \widetilde{y}_{t-2} + \varrho_3 \widetilde{y}_{t-3}. \tag{10}$$

Investors take the above dynamics into account when forming expectations but they must form beliefs about the future conduct of monetary policy.

Investor Beliefs About Future Monetary Policy To model the uncertainty investors face about the conduct of monetary policy, we assume that they expect the monetary policy rule to be subject to infrequent, nonrecurrent regime changes. We further assume that investors can accurately estimate the policy rule currently in place. This latter assumption, we argue, is a reasonable approximation of the role of beliefs in stock market dynamics in the context of infrequent regime changes, for two reasons. First, plausible uncertainty solely about the current rule–holding fixed beliefs about future monetary policy–is unlikely to be important for long duration assets such as the stock market. This is because what matters for heavily forward-looking assets is not where the policy rule is today, but where it will be for the foreseeable future. In our specification, investors continuously update their beliefs about the probability of moving to a new policy regime as soon as the beginning of the next month, so any reasonable additional uncertainty about where the rule is currently is relatively unimportant. Appendix J of the Online Appendix, shows that results for the stock market shown below are virtually unchanged if investors update their understanding of the current rule after the announcement–even by sizable amounts. This contrasts with the large effects found for changing beliefs about when and where policy will settle for the foreseeable future. Second, these assumptions are consistent with evidence that investors closely monitor Fed communications, combined with the observed practice of the Fed to clearly telegraph an intentional change in the stance of policy but to be comparatively vague about how long that change will last and what will come afterwards. Investors closely scrutinize Fed communications, not because they are most interested in learning what the central bank is doing today, but rather because they understand they must contemplate a future in which the central bank could operate differently from the one today, or any that has come before.

To model these circumstances, we assume that, for each time t policy rule regime indexed by ξ_t^P, investors hold in their minds a perceived "Alternative policy rule" indexed by ξ_t^A that they believe will come next, whenever the current policy regime ends:

$$\begin{aligned} i_t - \left(\overline{r} + \pi^T_{\xi_t^A}\right) &= \left(1 - \rho_{i,\xi_t^A} - \rho_{i_2,\xi_t^A}\right)\left[\psi_{\pi,\xi_t^A}\widehat{\pi}_{t,t-3} + \psi_{\Delta y,\xi_t^A}\left(4\widehat{\Delta y}_{t,t-3}\right)\right] \\ &\quad + \rho_{i_1,\xi_t^A}\left[i_{t-1} - \left(\overline{r} + \pi^T_{\xi_t^A}\right)\right] + \rho_{i_2,\xi_t^A}\left[i_{t-2} - \left(\overline{r} + \pi^T_{\xi_t^A}\right)\right] + \sigma_i\varepsilon_i, \end{aligned} \tag{11}$$

Investors do not have perfect foresight. When the current policy regime ends, the new policy regime that replaces it will never be exactly as previously imagined by the investor. When a regime ends, investors update their understanding of the new current policy rule and proceed to contemplate a new perceived Alternative for the next rule.

Investors in the model form beliefs not only about what the next policy rule ξ_t^A will look like, they also continuously assess the likelihood of switching to ξ_t^A by the beginning of next period. Specifically, for each ξ_t^P, investors have beliefs about the probability of remaining in ξ_t^P versus changing to ξ_t^A next month, but do not consider anything after that. This latter aspect of the specification is a form of bounded rationality that is arguably plausible in the context of infrequent regime changes. In the nonrecurrent regime setup of the model, this implies that the pondered Alternative is treated as an absorbing state as of time t, since the probability of returning to any previous rule must be zero by definition.

We formalize these ideas with a *belief regime* sequence governed by a discrete-valued variable $\xi_t^b \in \{1, 2, ...B, B+1\}$ with $B+1$ states. The regimes $\xi_t^b = 1, 2, ...B$ represent a grid of beliefs taking the form of perceived probabilities that the current policy rule will still be in place next period. The regime $\xi_t^b = B+1$ is a belief regime capturing the perceived probability of staying in the Alternative regime once it is reached. We order these so that belief regime $\xi_t^b = 1$ is the lowest perceived probability that the current policy rule will remain in place and $\xi_t^b = B$ is the highest.

The perceived regimes are modeled with a transition matrix taking the form:

$$\mathbf{H}^b = \begin{bmatrix} p_{b1}p_s & p_{b2}p_{\Delta 1|2} & \cdots & p_{bB}p_{\Delta 1|B} & 0 \\ p_{b1}p_{\Delta 2|1} & p_{b2}p_s & & p_{bB}p_{\Delta 2|B} & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ p_{b1}p_{\Delta B|1} & & & p_{bB}p_s & 0 \\ 1 - p_{b1} & 1 - p_{b2} & \cdots & 1 - p_{bB} & p_{B+1,B+1} = 1 \end{bmatrix}, \tag{12}$$

where $\mathbf{H}^b_{ij} \equiv p\left(\xi_t^b = i | \xi_{t-1}^b = j\right)$ and $\sum_{i \neq j} p_{\Delta i|j} = 1 - p_s$. In the above, p_{b1} is the subjective probability of remaining in the current policy rule under belief 1. For example, $p_{b1} = 0.05$ implies that investors believe there is a 5% chance that the current policy rule will still be in place next period. The non-zero off diagonal elements in the upper left $B \times B$ submatrix allow for the possibility that investors might receive subsequent information that could change their beliefs, and take that into account when forming expectations. The parameter, p_s is the probability that investors will have the same beliefs tomorrow as today. Thus, $1 - p_s$ is the probability that investors will change their minds. The parameter $p_{\Delta i|j}$ is the probability that agents will change to belief i tomorrow as a result of new information, conditional on having belief j today. Thus $p_{bj}p_s$ measures the probability of being in belief j tomorrow, conditional on having belief j today, while $p_{bj}p_{\Delta i|j}$ is the probability of being in belief i tomorrow conditional on having belief j today. Finally, $1-p_{bi}$ is the probability of having belief i today but exiting to the Alternative regime tomorrow. The parameter $p_{B+1,B+1}$ is the perceived probability of remaining in the Alternative regime conditional on having moved there. With perceived nonrecurrent regimes and our bounded rationality assumption, this probability is unity by definition. The model of beliefs therefore takes the form of a reducible Markov chain, implying that investors believe with probability 1 that they will eventually transition out of the current policy rule to the perceived Alternative rule.

Define the *overall policy regime* $\xi_t^{P,A} = \left\{\xi_t^P, \xi_t^A\right\}$ as characterized by the time t policy regime ξ_t^P and its associated time t perceived Alternative policy rule ξ_t^A. Thus with $N_p = 3$ true policy regimes over the course of the sample, there are also 3 perceived Alternative regimes over the same time span.

Equilibrium An equilibrium is defined as a set of prices (bond prices, stock prices), macro quantities (inflation, output growth, inflation expectations), laws of motion, and investor beliefs such that the equations in the asset pricing block are satisfied, the equations in the macro block are satisfied, with investor beliefs about monetary policy characterized by the perceived Alternative policy rule (11) and the perceived belief regime sequence described above with transition matrix (12).

Model Solution To solve the model we use the algorithm of Farmer, Waggoner, and Zha (2011). Appendix I of the Online Appendix explains the approximation used to preserve lognormality of the entire system, following Bansal and Zhou (2002) and Bianchi, Kung, and Tirskikh (2018). The solution of the model takes the form of a Markov-switching vector autoregression (MS-VAR) in the state vector

$$S_t = \left[S_t^M, m_t, pd_t, k_t, lp_t, \mathbb{E}_t^b\left(m_{t+1}\right), \mathbb{E}_t^b\left(pd_{t+1}\right)\right],$$

where $S_t^M \equiv [\widetilde{y}_t, g_t, \pi_t, i_t, \overline{\pi}_t, f_t]'$, with

$$S_t = \underbrace{C\left(\theta_{\xi_t^{P,A}}, \xi_t^b, \mathbf{H}^b\right)}_{\text{level}} + \underbrace{T(\theta_{\xi_t^{P,A}}, \xi_t^b, \mathbf{H}^b)}_{\text{propagation}}S_{t-1} + \underbrace{R(\theta_{\xi_t^{P,A}}, \xi_t^b, \mathbf{H}^b)}_{\text{amplification}}Q\varepsilon_t, \tag{13}$$

where $C(\cdot)$, $T(\cdot)$, and $R(\cdot)$ are matrices whose elements depend on primitive parameters, $\varepsilon_t = (\varepsilon_{f,t}, \varepsilon_{i,t}, \varepsilon_{g,t}, \varepsilon_{\mu,t}, \varepsilon_{k,t}, \varepsilon_{lp,t})$ is the vector of primitive Gaussian shocks, and $\theta_{\xi_t^{P,A}}$ is a vector of parameters that include the time-varying parameters of the current policy regime ξ_t^P, and the time-varying parameters of each associated Alternative regime ξ_t^A.

To solve the model we use the assumption that investors condition on the economic state S_t once it is observed at the end of each month. With this assumption, investor expectations in the presence of nonrecurrent regime switching and the perceived Alternative rule maybe be computed for any variable, as explained in Appendix G of the Online Appendix.

Equation (13) shows that the realized policy regime ξ_t^P, the associated Alternative regime ξ_t^A, and investor beliefs ξ_t^b about the probability of a shift in the policy rule amplify and propagate shocks in three distinct ways. First, there are "level" effects, captured by the coefficients $C\left(\theta_{\xi_t^{P,A}}, \xi_t^b, \mathbf{H}^b\right)$, that affect the economy absent shocks. These are driven by changes in the central bank's objectives such as the inflation target, as well as by the perceived risk of the stock market given by the risk-premium terms in (4). Second, there are "propagation" effects governed by the matrix coefficient $T(\theta_{\xi_t^{P,A}}, \xi_t^b, \mathbf{H}^b)$ that determine how today's economic state is related to tomorrow's. Third, there are "amplification" effects governed by the matrix coefficient $R(\theta_{\xi_t^{P,A}}, \xi_t^b, \mathbf{H}^b)$ that generate endogenous time-varying heteroskedasticity of the primitive Gaussian shocks.

This feature of the model, namely the state-contingent heteroskedasticity of primitive shocks, causes the perceived quantity of risk in the stock market to vary over time. Indeed, it is only through $R(\theta_{\xi_t^{P,A}}, \xi_t^b, \mathbf{H}^b)$ that the subjective risk premium and in particular $\mathbb{COV}_t^b\left[m_{t+1}, r_{t+1}^D\right]$ (equation (4)) varies at all. In turn, $R(\theta_{\xi_t^{P,A}}, \xi_t^b, \mathbf{H}^b)$ varies only with variables: (i) realized regime changes ξ_t^P in the conduct of monetary policy (ii) the Alternative regime ξ_t^A investors perceive will come next, and (iii) time-varying beliefs ξ_t^b about the probability of switching to ξ_t^A. Perceptions about the Alternative rule and the probability of moving there are especially important for the perceived quantity of stock market risk because of the key role played by the Alternative rule parameter $\psi_{\Delta y,\xi_t^A}$. This parameter reflects the investor's view of the degree to which the central bank will react to future fluctuations in economic growth. The larger is $\psi_{\Delta y,\xi_t^A}$, the more agents expect future policy to aggressively stabilize economic growth by lowering interest rates when payout falls. This belief is tantamount to expecting a "Fed put," since it means that the central bank engineers a lower discount rate whenever cash flows are expected to grow more slowly, putting a floor on stock prices. Overall, perceptions that the next rule will feature a higher $\psi_{\Delta y,\xi_t}$ have the effect of raising $\mathbb{COV}_t^b\left[m_{t+1}, r_{t+1}^D\right]$ and thus lower the subjective risk premium. This can be seen from Figure II, which shows the almost linear negative relationship between the subjective risk premium and the parameter $\psi_{\Delta y,\xi_t}$.

Figure II: Subjective Risk Premium and Policy Rule Coefficients

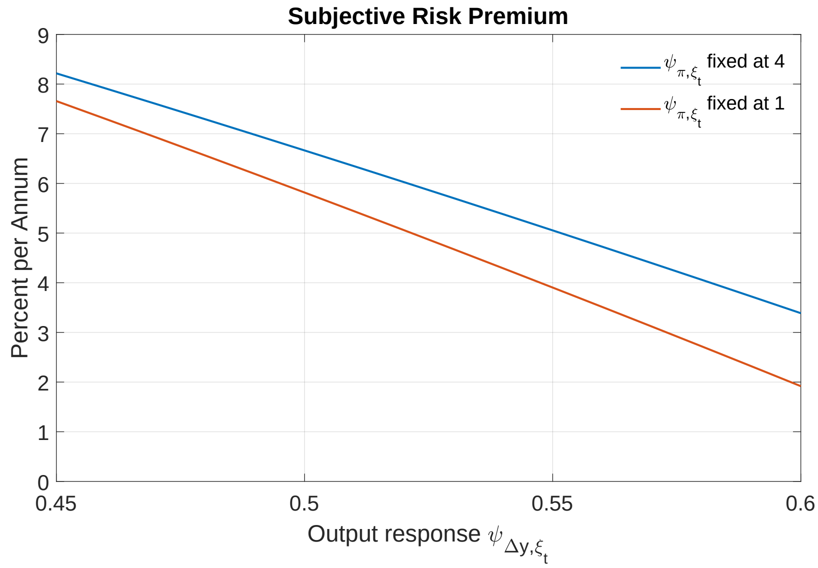

Notes: This figure reports the subjective risk premium as a function of different values for policy rule coefficients.

Investor Information and Updating Let $\mathbb{I}_t$ denote the time t information set of investors, which includes the current policy regime ξ_t^P, their perceived Alternative regime ξ_t^A, their beliefs about monetary policy ξ_t^b, and additional data available at mixed frequencies that we don't explicitly specify. We assume that investors observe S_t and consume payout at a monthly frequency. Investors can observe S_t only at the end of each month, but they price assets continuously by revising expectations in response to news. With S_t observed only at the end of the month, any Fed news that the investor attends to *within* a month results in the updating of a *nowcast* of S_t, which we assume they produce by filtering a potentially extensive database of timely, high-frequency information in $\mathbb{I}_t$. This database is unobserved by the econometrician, but our estimation will not rely on it as explained below.

Investors use $\mathbb{I}_t$ in two ways. First, given a baseline monthly decision interval, they update their previous nowcasts and subjective expectations once S_t is observed at the end of every t. Second, investors allocate attention to updating nowcasts of S_t and beliefs ξ_t^b about future monetary policy at specific times *within* a month when the central bank releases information. This higher-frequency attentiveness to Fed news echoes real-world "Fed watching" and is the mechanism through which the model accommodates swift market reactions to surprise central bank announcements, driving jumps in investor perceptions of stock market risk $\mathbb{COV}_t^b\left[m_{t+1}, r_{t+1}^D\right]$.

4 Structural Estimation

Define $\xi_t \equiv \left(\xi_t^P, \xi_t^A, \xi_t^b\right)$ to be the collection of belief and policy regime indicators. The system of estimable equations may be written in state-space form by combining the state equations (13) with an observation equation taking the form

$$X_t = D_{\xi_t,t} + Z_{\xi_t,t}\left[S_t', \widetilde{y}_{t-1}\right]' + U_t v_t \tag{14}$$

where X_t denotes a vector of data, $v_t \sim N(0, I)$ is a vector of observation errors, U_t is a diagonal matrix with the standard deviations of the observation errors on the main diagonal, and $D_{\xi_t,t}$, and $Z_{\xi_t,t}$ are parameters mapping the model counterparts of X_t into the latent discrete- and continuous-valued state variables ξ_t and S_t, respectively, in the model. The matrices $Z_{\xi_t,t}$, U_t, and the vector $D_{\xi_t,t}$ depend on t independently of ξ_t because some of our observable series are not available at all frequencies and/or over the full sample. As a result, the state-space estimation uses different measurement equations to include series when they are available, and exclude them when they are missing.

We estimate the state-space representation with Bayesian methods using a modified version of Kim's (Kim (1994)) basic filter and approximation to the likelihood for Markov-switching state space models, and a random-walk metropolis Hastings MCMC algorithm to characterize uncertainty. The parameters of the monetary policy rule are estimated under symmetric priors, while the priors on the other parameters are standard and specified to be loosely informative except where there are strong restrictions dictated by theory, e.g., risk aversion must be non-negative. A complete description of the priors is provided in Appendix A of the Online Appendix.

Mixed-Frequency Filtering Algorithm To estimate how an investor in the model would react to news, we undertake a high-frequency event study using narrow (30 minute) windows around a news event. The distinction between our approach and common event study analyses is that we study reactions to specific events as they would be understood through the lens of a structural model.

In the structural model here, investors observe S_t and consume payout at a monthly frequency. At the same time, they price assets continuously by updating their beliefs in response to news. News generates a set of signals about the current economic state S_t with varying degrees of precision. For example, an FOMC statement announcing a change in the target federal funds rate is a signal with infinite precision because it removes all uncertainty about the end-of-month target funds rate. By contrast, announcements that speak to the macroeconomic outlook is a signal with considerable noise because the Fed cannot directly control the outcome. In either case we are interested in measuring the jump in investor expectations attributable to the news *per se.*

To do so, we design a filter to measure the model's implications for high-frequency belief updating attributable to news.

The filtering algorithm described in this section can be used to infer real-time updating of expectations about the current economic state S_t to be revealed at the *end* of the period, in response to news events that occur *within* the period, i.e., before S_t is observed. Since these represent expectations for the current period, we refer to them as nowcasts. To facilitate interpretation, we explain the algorithm for a simple state space representation with a single state variable and no regime switching. The details for the full structural model with belief regimes are provided in Appendix H of the Online Appendix.

Let t denote a month. Let d_h denote the number of time units that have passed within a month when we have reached a particular point in time indexed by h, and let nd denote the total number of time units in the month. Then $0 \leq d_h/nd \leq 1$, and the intramonth time period is denoted $t-1+\delta_h$ with $\delta_h \equiv d_h/nd \in [0,1]$. For example, if the time unit is in days and we are at the beginning of the 11th day in a 31 day month, then $\delta_h \equiv 10/31 = 0.3226$.

Consider a simple state space model in which there is a single state for inflation, π_t, that follows an autoregressive process with no regime changes. The process is specified at monthly frequency:

$$\pi_t = \rho \pi_{t-1} + \sigma_\pi \varepsilon_{\pi,t}. \tag{15}$$

Investors are asked to predict inflation for next month $(t+1)$, even when surveyed intramonth $(t-1+\delta_h < t)$. Agents know the data generating process (15). At the end of month t, we assume that the data are fully revealed and, accordingly, inflation expectations for the next month $(t+1)$ reflect the realized value of inflation at the end of the current month (t), i.e.,

$$\mathbb{E}_t\left[\pi_{t+1}\right] = \rho \pi_t.$$

Now suppose we are within month t at day $t-1+\delta_h$, where $\delta_h < 1$, and the investor receives some new information about what π_t is likely to be. The investor will update their time $t-1$ forecast $\mathbb{E}_{t-1}\left[\pi_{t+1}\right]$ by revising their *nowcast* of π_t. Letting $\pi_{t\backslash t-1+\delta_h}$ denote their revised nowcast, their updated forecast of π_{t+1} is

$$\begin{aligned}
\mathbb{E}_{t-1+\delta_h}\left[\pi_{t+1}\right] &= \rho \pi_{t\backslash t-1+\delta_h} \\
&= \rho[\rho \pi_{t-1} + \sigma_\pi \varepsilon_{\pi,t\backslash t-1+\delta_h}] \\
&= \mathbb{E}_{t-1}\left[\pi_{t+1}\right] + \rho \sigma_\pi \varepsilon_{\pi,t\backslash t-1+\delta_h},
\end{aligned}$$

where we use the symbol "\" to indicate that the conditioning is with respect to the agent's information set. This shows that the revised nowcast $\pi_{t\backslash t-1+\delta_h}$ implicitly depends on the *perceived shock* $\varepsilon_{\pi,t\backslash t-1+\delta_h}$, which is the perception about $\varepsilon_{\pi,t}$ needed to rationalize the update in beliefs about π_t at $t-1+\delta_h$ attributable to the news. Given that investors observe π_t at the end of the month, $\pi_{t\backslash t} \equiv \pi_t$.

The above expressions are a simple theoretical model. To estimate high-frequency jumps in beliefs in such a model, we require data and an econometric approach. For this we assume that the perceived shocks calculated within a month have the same distribution as the realized shocks observed at the end of the month. This allows us to rerun the filter to infer changes in the perceived shocks without introducing additional parameters and structure on the investors updating process, ultimately a latent process. We use superscipt o to denote the observed value of variables in the data, in order to distinguish them from the theoretical values in the model. Suppose that we observe both monthly inflation data π_t^o and a daily measure of inflation expectations such as from a survey or market, where the daily expectation as of time $t-1+\delta_h$ is denoted $\mathbb{E}_{t-1+\delta_h}[\pi_{t+1}]^o$. Combine data and model in a state space representation. With the state (transition) equation given by (15), we have the following observation equations.

1. At the end of month t, both inflation and inflation expectations are observed. The observation equation is:

$$\underbrace{\begin{bmatrix} \pi_t^o \\ \mathbb{E}_t[\pi_{t+1}]^o \end{bmatrix}}_{\equiv X_t} = \underbrace{\begin{bmatrix} 0 \\ 0 \end{bmatrix}}_{\equiv D_t} + \underbrace{\begin{bmatrix} 1 \\ \rho \end{bmatrix}}_{\equiv Z_t} \pi_t + \underbrace{\begin{bmatrix} \sigma_{u,\pi} & 0 \\ 0 & \sigma_{u,\mathbb{E}[\pi]} \end{bmatrix}}_{\equiv U_t} \underbrace{\begin{bmatrix} u_{\pi,t} \\ u_{\mathbb{E}[\pi],t} \end{bmatrix}}_{u_t},$$

   where we allow for observation errors u_t to avoid the stochastic singularity arising because we have two observables (inflation and expected inflation) mapped onto one stochastic process (model inflation).

2. Intramonth, only the daily inflation expectations data are observed. These data can be used to filter out what agents expect inflation will be over the current month:

$$\underbrace{\begin{bmatrix} - \\ \mathbb{E}_{t-1+\delta_h}[\pi_{t+1}]^o \end{bmatrix}}_{\equiv X_{t-1+\delta_h}} = \underbrace{\begin{bmatrix} - \\ 0 \end{bmatrix}}_{\equiv D_{t-1+\delta_h}} + \underbrace{\begin{bmatrix} - \\ \rho \end{bmatrix}}_{\equiv Z_{t-1+\delta_h}} \pi_{t\backslash t-1+\delta} + \underbrace{\begin{bmatrix} - & - \\ - & \sigma_{u,\mathbb{E}[\pi]} \end{bmatrix}}_{\equiv U_{t-1+\delta_h}} \begin{bmatrix} u_{\pi,t-1+\delta_h} \\ u_{\mathbb{E}[\pi],t-1+\delta_h} \end{bmatrix}$$

   where

$$\pi_{t\backslash t-1+\delta_h} = \rho\pi_{t-1} + \sigma_\pi \varepsilon_{\pi,t\backslash t-1+\delta} \tag{16}$$

   This provides a specific interpretation of our filtering results. What we have filtered out is an estimate of the belief $\mathbb{E}_{t-1+\delta_h}[\pi_{t+1}]$ attributable to whatever news occurred at day $t-1+\delta_h$, rather than a high frequency forecast of future inflation driven by an ultra high-frequency inflation shock at $t-1+\delta_h$.[10]

[10]This differs from common reduced-form settings in which high-frequency data are used primarily to

The key insight of this approach is that our high-frequency data is *forward-looking* (e.g., from surveys or financial markets) and thus directly measures updating of beliefs. These data can be filtered to infer updates to investor nowcasts in tight windows around news events, without having to take a stand on the (ultimately unobservable) nowcasting models and information sets of investors, or the ultra high-frequency behavior of π_t. We use the symbol "|" to refer to conditioning in the filter that is with respect to the econometrician's information set to distinguish it from "\", which denotes conditioning of the investor based on their latent (to the econometrician) information set.

Putting this all together, the algorithm has the following steps:

1. The econometrician has information up through month $t-1$. Compute one step-ahead nowcast estimates:

$$\begin{aligned} \pi_{(t \backslash t-1+\delta_h)|t-1} &= \rho \pi_{t-1|t-1} \\ P_{(t \backslash t-1+\delta_h)|t-1} &= \rho^2 P_{t-1|t-1} + \sigma_\pi^2 \end{aligned}$$

2. Compute forecast error given high-frequency information at $t-1+\delta_h$:

$$\begin{aligned} e_{(t \backslash t-1+\delta_h)|t-1+\delta_h,t-1} &= X_{t-1+\delta_h} - D_{t-1+\delta_h} - Z_{t-1+\delta_h} \pi_{(t \backslash t-1+\delta_h)|t-1} \\ f_{(t \backslash t-1+\delta_h)|t-1+\delta_h,t-1} &= Z_{t-1+\delta_h} P_{(t \backslash t-1+\delta_h)|t-1} Z'_{t-1+\delta_h} + U^2_{t-1+\delta_h} \end{aligned}$$

3. Update estimates of nowcast and its variance based on new information at $t-1+\delta_h$

$$\begin{aligned} \pi_{(t \backslash t-1+\delta_h)|t-1+\delta_h} &= \pi_{(t \backslash t-1+\delta_h)|t-1} + P_{(t \backslash t-1+\delta_h)|t-1} Z'_{t-1+\delta_h} \cdot \\ &\quad \left(f_{(t \backslash t-1+\delta_h)|t-1+\delta_h,t-1}\right)^{-1} e_{(t \backslash t-1+\delta_h)|t-1+\delta_h,t-1} \\ P_{(t \backslash t-1+\delta_h)|t-1+\delta_h} &= P_{(t \backslash t-1+\delta_h)|t-1} - P_{(t \backslash t-1+\delta_h)|t-1} Z'_{t-1+\delta_h} \cdot \\ &\quad \left(f_{(t \backslash t-1+\delta_h)|t-1+\delta_h,t-1}\right)^{-1} Z_{t-1+\delta_h} P_{(t \backslash t-1+\delta_h)|t-1} \end{aligned}$$

4. Repeat the above at δ_h corresponding to the desired time units before and again after an announcement to infer revisions in nowcasts due to news. If $t-1+\delta_h = t$, move to the next period by setting $t-1 = t$ and returning to step 1.

The full algorithm applied to the actual structural model refers to the state space equations (13) and (14), where δ_h corresponds the number of time units that have passed when we are at 10 minutes before or 20 after an FOMC event, and requires combining the Kalman and Hamilton filters, as described in Appendix H of the Online Appendix.

Two additional points about this algorithm bear noting. First, the filter can be rerun as frequently as desired within a month, even as transition dynamics are still specified

augment prediction models, an objective typically accomplished by specifying the state/transition equations at the highest frequency of data used. Our mixed-frequency algorithm is designed for a different purpose, namely to measure reactions to news in a structural model. In this context, the state/transition equation is part of the structural model and needs to correspond to the monthly frequency over which investors observe S_t and consume payout.

across months. It is therefore straightforward to handle news events that are spaced non-uniformly over the sampling interval, as when the timing and/or number of FOMC meetings during a month varies over the sample. Second, the entire perceived state vector S_t can be reestimated at any point and over any small interval within a month, provided that a subset of forward-looking data are available at high-frequency. Thus we can infer revisions in perceptions on any macro variable in S_t from the information encoded in more timely financial market observations, even if data on output, earnings, inflation, etc., are only available once per month.

Data and Measurement Our full dataset spans January 1961 through February 2020. The sample of Fed news consists of 220 Federal Open Market Committee (FOMC) press releases covering February 4th, 1994 to January 29th, 2020. Observations on most series are available monthly. For quarterly GDP growth we interpolate to monthly frequency using the method in Stock and Watson (2010). An explicit description of the mapping between our observables and model counterparts and complete description of each data series and sources is given in Appendix H of the Online Appendix.

We use high-frequency pre- and post-FOMC observations on the following variables: daily survey expectations of inflation, GDP growth, and S&P 500 earnings from Bloomberg (BBG), daily observations on the 20-year Baa credit spread with the 20-year Treasury bond rate (Baa spread hereafter), minutely observations on four distinct federal funds futures (FFF) contract rates with different expiries, and minutely observations on the S&P 500 market value. These high-frequency data serve two purposes. First, they allow us to measure the effect of Fed news on investor beliefs and perceptions in tight windows around announcements. Second, the timely information contained in these forward-looking series allow us to account for economic news that pre-dated the FOMC announcement but arrived after the latest observations on stale monthly survey data (Bauer and Swanson (2023)). By conditioning on close-range, pre- and post-announcement observations for inflation, GDP growth, and earnings expectations, as well as credit spreads (at daily frequency), interest rate futures, and the stock market (10 minutes before and 20 minutes after), we ensure that our estimated post-announcement jumps in nowcasts cannot be readily attributed to stale economic news.

At lower frequencies, we use the household-level Survey of Consumers (SOC) from the University of Michigan to discipline household expectations and three additional professional forecaster surveys from Bluechip (BC), Survey of Professional Forecasters (SPF) and Livingston (LIV) to discipline investor expectations. We measure investor expectations at multiple horizons using the four different professional surveys and treat each of these as a noisy signal on the true underlying investor expectations process.

A number of series are used because they have obvious model counterparts. Data for Gross Domestic Product (GDP) growth and inflation are mapped into the model implications for output growth and inflation; data on the current effective federal funds rate (FFR) are mapped into the model's implications for the current nominal interest rate; data on the FFF market and the BC survey measure of the expected FFR 12 months-ahead are mapped into the model's implications for investor expectations of the FFR.[11] The inclusion of data on long-dated FFF contracts and survey forecasts of the funds rate a year or more out are especially helpful for identifying the parameters of the Alternative policy rule, since investors' longer-term forecasts are dominated by where they believe future policy will be and not by the rule currently in place.

We discipline the earnings share of output K_t with observations on the ratio of S&P 500 earnings to GDP. We account for the fact that earnings in the data differs from the payout shareholders actually receive by mapping the theoretical concept for k_t into its respective data series allowing for observation error in the relevant observation equation.

Finally, data on the Baa spread are mapped into the model's implications for the liquidity premium, lp_t, a catchall for many factors outside of the model that could effect the subjective equity premium, including changes in the perceived liquidity and safety attributes of Treasuries, default risk, flights to quality, and/or sentiment. We use the Baa spread as an observable likely to be correlated with many of these factors, but our measurement equation allows for both a constant and a slope coefficient on the Baa spread along with observation error, in order to soak up variation in this latent variable that may not move identically with the spread.

Estimating Beliefs We take the parameters p_{bi} in $\mathbf{H}^b$ from a discretized beta distribution, estimating its mean and variance as parameters of the structural estimation. The parameters $p_{\Delta i|j}$ are specified as $(1-p_s)\left(\rho_b^{|i-j-1|}/\sum_{i\neq j}\rho_b^{|i-j-1|}\right)$, where p_s and $\rho_b < 1$ are estimated parameters and $|i-j-1|$ measures the distance between beliefs j and i, for $i \neq j \in (1, 2, ..., B)$. This creates a decaying function that makes the probability of moving to contiguous beliefs more likely than jumping to very different beliefs.

Let T be the sample size used in the estimation and let X_t be the vector of observations as of time t. Let $\Pr\left(\xi_t^b = i|X_T;\boldsymbol{\theta}\right) \equiv \pi_{t|T}^i$ denote the probability that $\xi_t^b = i$, for $i = 1, 2,...,B+1$, based on information that can be extracted from the whole sample and knowledge of the parameters $\boldsymbol{\theta}$, while $\pi_{t|T}$ is a $(B+1)\times 1$ vector containing the elements $\left\{\pi_{t|T}^i\right\}_{i=1}^{B+1}$. We refer to these as the smoothed regime probabilities. The time

[11] In principle, fed funds futures market rates may contain a risk premium that varies over time. If such variation exists, it is absorbed in the estimation by the observation error for these equations. In practice, risk premia variation in fed funds futures is known to be small when that variation is measured over the short 30-minute windows surrounding FOMC announcements that we analyze (Piazzesi and Swanson (2008)).

t perceived probability of exiting the current policy rule, i.e., of transitioning in the next period to the Alternative policy regime ξ_t^A, is given by $\overline{P}_t^{bE} \equiv \sum_{i=1}^{B} \pi_{t|T}^i (1 - p_{bi})$. The time t perceived probability of exiting the current policy rule and transitioning in h periods to ξ_t^A is $\overline{P}_{t+h,t}^{bE} = \mathbf{1}'_{B+1} \left(\mathbf{H}^b\right)^h \pi_{t|T}$, where $\mathbf{1}'_{B+1}$ is an indicator vector with 1 in the $(B+1)$th position and zeros elsewhere. We use these estimated regime probabilities to compute the most likely belief regime at each point in time and track how it changes around Fed announcements and the whole sample. In the applied estimation, we set $B = 11$.

5 Estimation Results

This section presents results from the structural estimation based on the modal values of the posterior distribution for the parameters. The estimated credible sets indicate that the parameters are tightly identified and we report other moments of the posterior in Table **??** of the Online Appendix. The state space estimation allows for observation errors on theoretical variables for which we use multiple observable series, treating each a distinct noisy signal on the theoretical variable. For example, we have multiple survey expectations to map on to model-implied investor expectations. For variables where we use only a single observable signal (e.g., inflation, GDP, the FFR, and the S&P 500), we set observation error to zero. The estimated model-implied series track their empirical counterparts closely, as shown in Figure A.I of the Online Appendix.

Parameter and Latent State Estimates Table I reports the posterior modes for the policy rule parameters $\pi^T_{\xi_t^P}$, ψ_{π,ξ_t^P}, $\psi_{\Delta y,\xi_t^P}$ and ρ_{i,ξ_t^P}, where we use symmetric priors. The results imply that the regime subperiods reported in Figure I are associated with quantitatively large changes in the estimated policy rule, as well as in the associated Alternative policy rules that we estimate investors perceived would come next. The Great Inflation (GI) regime (1961:Q1-1978:Q3) is characterized by a high implicit

Table I: **Policy Rule Parameters**

	Great Inflation Regime		Great Moderation Regime		Post-Millennial Regime	
	Realized	Alternative	Realized	Alternative	Realized	Alternative
π_ξ^T	12.3763	3.7021	2.2366	0.7033	2.5032	0.0600
ψ_π	1.8617	0.6817	1.9532	2.7249	0.9163	0.8103
ψ_y	1.0214	0.4541	0.1015	0.6472	0.0550	0.5499
ψ_π/ψ_y	1.8227	1.5012	19.2433	4.2103	16.6600	1.4735
$x = \rho_{i,1} + \rho_{i,2}$	0.9955	0.9818	0.9868	0.5680	0.9958	0.8922

Notes: Posterior mode values of the parameters for the current and Alternative policy rules. Great Inflation Regime: 1961:Q1-1978:Q3. Great Moderation Regime: 1978:Q4-2001:Q3. Post-Millennial Regime: 2001:Q4-2020:Q1. The estimation sample spans 1961:Q1-2020:Q1.

inflation target and a moderate level of inflation activism (ψ_{π,ξ_t^P}), consistent with previous research arguing that the Fed accommodated high inflation during this period. The perceived Alternative policy rule for this subperiod has a much lower inflation target, but features less activism against both inflation and output growth, with inflation stabilization perceived as the main objective. The anticipation of a lower inflation target is in fact a defining feature of the subsequent Great Moderation (GM) regime that began in 1978:Q4. The GM also featured a stronger emphasis on inflation stabilization than the GI regime but little activism on economic growth. Moving to the Post-Millennial (PM) regime, we find that policy rule parameters then shifted back toward slightly more accommodative values with a higher implicit inflation target, but with far less activism on inflation and comparably low activism on output growth .

The estimated perceived Alternative policy rules of each regime show how investors expected policy to change in the future. In the GM regime, investors evidently expected the next rule to have an inflation target that was even lower than what was actively in place at the time, along with greater activism in stabilizing both inflation and economic growth. In the PM period investors expected an inflation target that was lower still, but with a greater emphasis on output growth stabilization relative to inflation stabilization compared to the realized rule during the PM period. Thus both the GM and PM periods are characterized by expectations that the next policy rule would be both more hawkish and more active on output growth than the realized rules of those periods. Since more activism on output growth is indicative of more aggressive action to stabilize the real economy, these features of the perceived Alternative rules are closely related to perceived risk in the stock market, as discussed below.

A comment is in order about the estimated magnitudes for $\pi^T_{\xi_t^P}$ shown in Table I. Although this parameter plays the role of an "inflation target" in the interest rate rule, unlike traditional New Keynesian models with a time invariant inflation target, $\pi^T_{\xi_t}$ is appropriately interpreted as an implicit time t target rather than an explicit long-run objective. To understand why, consider the PM period as an example. The structural estimation implies that, to achieve the observed average CPI inflation of roughly 1.96% over this period, $\pi^T_{\xi_t^P}$ needed to be 2.5%, well above what officially became in 2012 the explicitly stated long-run inflation objective of 2%. Forward guidance "low-for-long" interest rate policies and quantitative easing, two tools that were employed at the effective lower bound (ELB), are channels that manifest in the model as a higher values for $\pi^T_{\xi_t^P}$, since with $\gamma^T > 0$ these tools generate higher perceived trend inflation by households even as nominal interest rates remain unchanged at the ELB (equation (9)). Likewise, the high value for $\pi^T_{\xi_t^P}$ in the GI regime represents an implicit central bank accommodation of the high inflation of the 1970s that is difficult to explain otherwise.

Table II presents estimation results for key model parameters other than those of the policy rule.[12] The estimates imply a very high level of inertia in household inflation expectations. The constant gain parameter γ controlling the speed with which households update beliefs about inflation with new information on inflation is estimated to be low ($\gamma = 0.0001$). Furthermore, the parameter γ^T controlling the speed with which households' perceived trend inflation is influenced by shifts in $\pi^T_{\xi^P_t}$ is also estimated to be small, though non-zero ($\gamma^T = 0.006$). Taken together, these findings imply that households revise their beliefs about trend inflation only very slowly over time, both in response to changes in the implicit inflation target and with past inflation realizations. Regarding investors' beliefs, the parameter p_s is estimated to be 0.94, indicating that investors' beliefs are mildly persistent. Finally, we estimate a moderate level of risk aversion for the investor ($\sigma_P = 6.9$).

Table II: **Other Key Parameters**

Parameter	Mode	Parameter	Mode	Parameter	Mode	Parameter	Mode
σ	0.1051	γ^T	0.0056	σ_f	6.4146	σ_{lp}	0.2035
β	0.7552	σ_p	6.8668	σ_i	0.0352	σ_g	1.4568
ϕ	0.7477	β_p	0.9964	σ_μ	0.1302		
γ	0.0001	p_s	0.9410	σ_k	6.3330		

Notes: Posterior mode values of the parameters named in the row. The sample spans 1961:Q1-2020:Q1.

Before leaving this section we report the model implications for basic asset pricing moments. Table III shows that the model based moments for the log stock return, real interest rate, and earnings growth, based on the modal parameter and latent state estimates, match their data counterparts closely.

Table III: **Asset Pricing Moments**

Moments	Model		Data	
	Mean	StD	Mean	StD
Log Excess Return	7.20	14.93	7.42	14.85
Real Interest Rate	1.65	2.48	1.72	2.53
Log Real Earning Growth	2.62	25.06	1.96	17.24

Notes: Annualized monthly statistics (means are multiplied by 12 and standard deviations by $\sqrt{12}$) and reported in units of percent. Excess returns are the log difference in the SP500 market capitalization minus FFR. Real interest rate is FFR minus the average of the average of the one-year ahead forecasts of inflation from the BC, SPF, SOC, and Livingston surveys. SP500 Earnings is deflated using the GDP deflator and divided by population. The sample is 1961:M1 - 2020:M2.

Investor Beliefs About Monetary Policy Over Time Figure III plots the historical times series of estimated perceived probabilities that investors assign to being

[12] The model has a large number of additional auxiliary parameters that are used to map observables into their model counterparts. To conserve space, estimates of these parameters are reported in the Online Appendix.

in a new policy rule regime within one year. Specifically, the figure reports the end-of-the-month value for $\overline{P}^{bE}_{t+12,t} \equiv \pi^{B+1}_{t+h,t|T} = \mathbf{1}'_{B+1}\left(\mathbf{H}^{b}\right)^{12}\pi_{t|T}$, where $\mathbf{1}'_{B+1}$ is an indicator vector with 1 in the $(B+1)$th position and zeros elsewhere and $\pi_{t|T}$ is the vector of smoothed time t belief regime probabilities. The vertical lines mark the timing of the two realized policy regime changes in our sample.

Figure III: Perceived Probability of Monetary Policy Regime Change

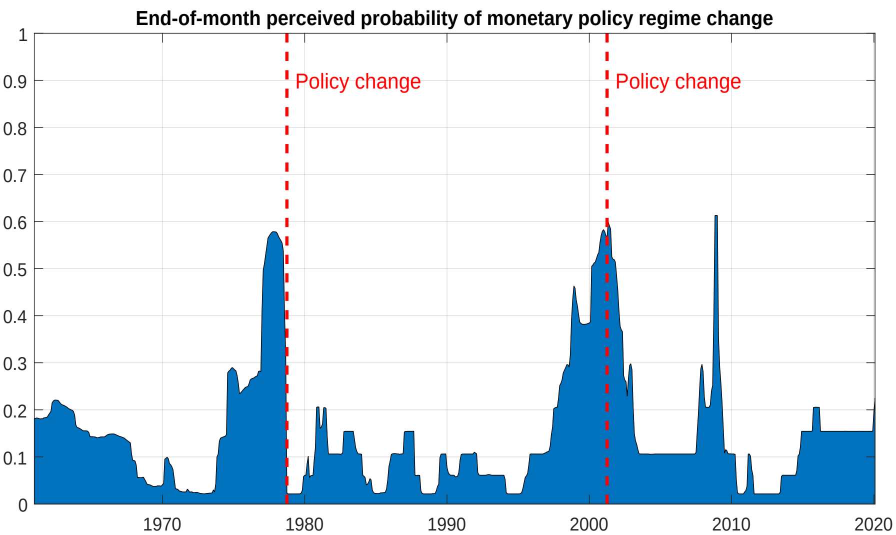

Notes: Estimated end-of-month perceived probability that investors assign to exiting the current monetary policy rule within one year. The sample spans 1961:M1-2020:M2.

Figure III shows that the perceived probability of a policy rule regime change fluctuates strongly over the sample and typically increases before a realized policy change, suggesting that financial markets have some ability to anticipate realized shifts in the conduct of policy even though they cannot perfectly predict what the next policy rule will look like. The perceived probability of a policy rule change also spikes upward sharply in the financial crisis when no actual change occurred subsequently, though this "mistake" is short-lasting.

Figure IV shows the *change* in the estimated perceived probability of a monetary policy regime change within the next year in tight windows around every FOMC announcement in our sample. Naturally, many FOMC announcements carry little news of any kind, consistent with the majority of points lining up along the horizontal line and the idea that significant changes in the policy rule are infrequent. Yet there are clearly some announcements that are associated with sizable changes in the perceived probability of exiting the current policy regime. The largest decline in this perceived probability occurred on October 29, 2008 when the FOMC announced a 50 basis point reduction in the fed funds rate target to 1% on an annual basis. In this case, the perceived probability of a regime change in the next year declined by more than 2% in the 30 minutes surrounding the FOMC press release. The largest increase in the perceived probability of a policy regime change occurs on March 11, 2008 when the Federal Reserve announced an expansion of securities lending in a new facility, the Term Securities Lending Facility (TSLF). In this case, the perceived probability of policy regime change increased more than 2%. We discuss these episodes further below.

Figure IV: Change in beliefs about a policy switch around FOMC announcements

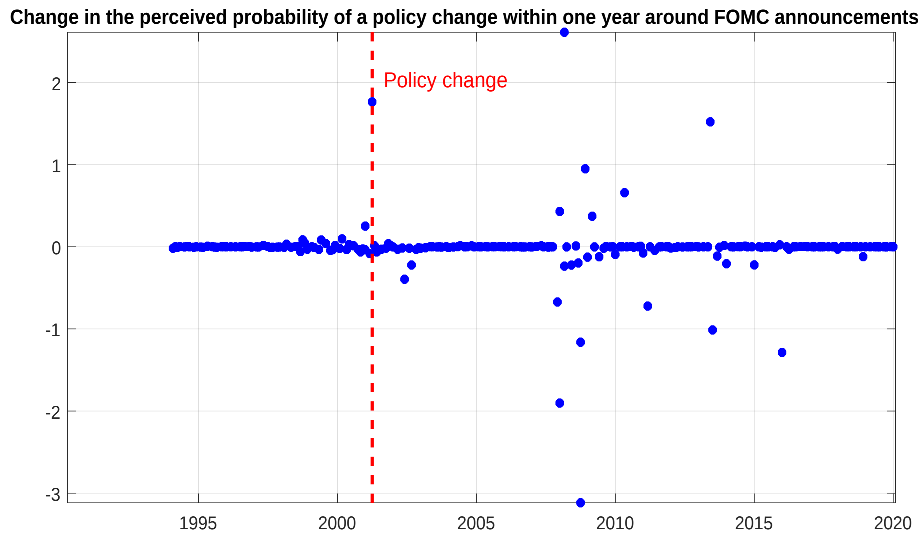

Notes: Pre-/post- FOMC announcement log changes (10 minutes before/20 minutes after) in the probability that financial markets assign to a switch in the monetary policy rule occurring within one year. The full sample has 220 announcements spanning February 4th, 1994 to February 28th, 2020. The sample reported in the figure is 1993:M1-2020:M2.

High-Frequency Analysis We now move on to study specific announcements, investigating market reactions in tight windows around individual FOMC press releases. In our analysis, the pre-FOMC value is always either 10 minutes before or the day before the FOMC press release time, depending on data availability (daily versus minutely), and the post-FOMC value is either 20 minutes after or the day after the release. Figure V displays the log change in pre-/post- FOMC announcement values of variables we measure at high frequency, for each FOMC announcement in our sample. Some announcements are associated with declines in the stock market within 30 minutes surrounding the FOMC press release that exceed 2% in absolute terms or increases above 4%. Many announcements also produce large jumps in other financial market variables such as FFF rates and the Baa spread.

The mixed-frequency structural approach developed in this paper allows us to investigate a variety of possible explanations for these large market reactions. Consider an FOMC announcement in month t. As above, let $\delta_h \in (0,1)$ represent the number of time units that have passed during month t up to some particular point $t-1+\delta_h$. Let $S^i_{(t \backslash t-1+\delta_h)|t-1+\delta_h}$ denote a filtered estimate of investors' perceived time t economic

Figure V: HF Changes in Prices and Expectations

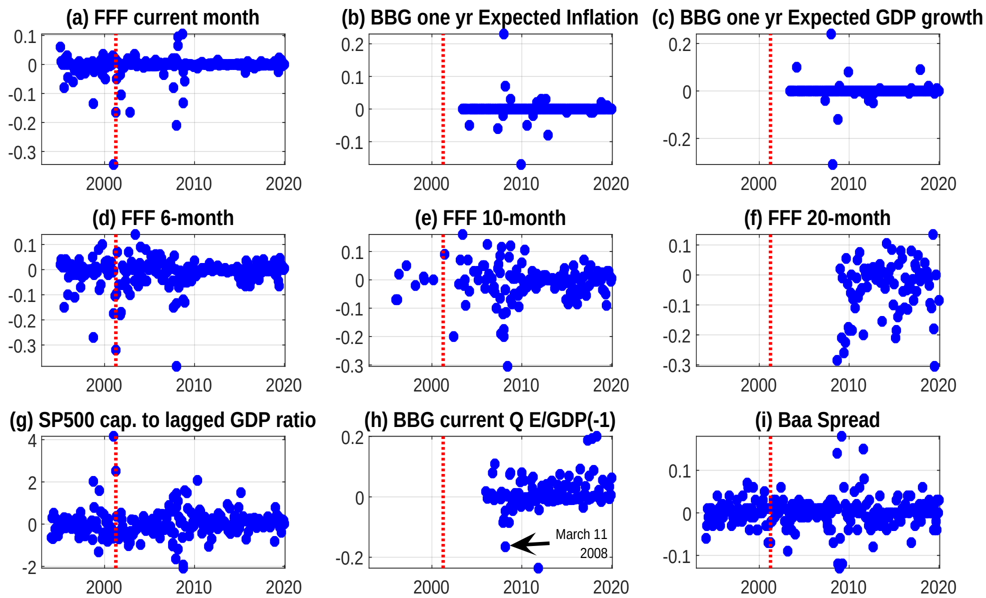

Notes: Log change in the observed variables in a short time-window around FOMC meetings. For all but panels (b) and (c), this corresponds to a change measured from 10 minutes before to 20 minutes after an FOMC statement is released. For panels (b) and (c), this corresponds to one day before to one day after the FOMC statement is released. The full sample has 220 FOMC announcements spanning February 4th, 1994 to February 28th, 2020. The sample reported in the figure is 1993:M1-2020:M2.

state based on their information up to time $t-1+\delta_h$, conditional on $\xi_t^b = i$. We use the filtering algorithm described above along with high-frequency, forward-looking data on investor expectations and financial markets to obtain estimates of the pre- and post-FOMC announcement nowcasts $S^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$, and the associated filtered belief regime probabilities $\pi^i_{t|t-1+\delta_h} \equiv \Pr\left(\xi_t^b = i | X_{t-1+\delta_h}, X^{t-1}\right)$, where δ_h assumes distinct values d_{pre} and d_{post} that denote the times right before and right after the FOMC meeting. These *pre* and *post* differences represent our estimates of the market's revised nowcasts for S and beliefs about the future conduct of monetary policy that are attributable to the FOMC announcement.

Figure VI displays the percent changes in pre-/post- announcement nowcasts of different elements of S_t for every FOMC announcement in our sample. The figure shows that some FOMC announcements led to frequent and large changes in investor perceptions about trend growth g_t, detrended output, $\widetilde{y}_t$, inflation, current demand f_t, the earnings share k_t, and the liquidity premium lp_t. This implies that some announcements cause investors to significantly revise their beliefs about the state of the economy and its core driving forces.

To make further progress of our understanding of what markets learn from FOMC announcements, we use estimates of $S^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$ and the belief regimes $\pi^i_{t|t-1+\delta_h}$ in

Figure VI: HF Changes in State Variables

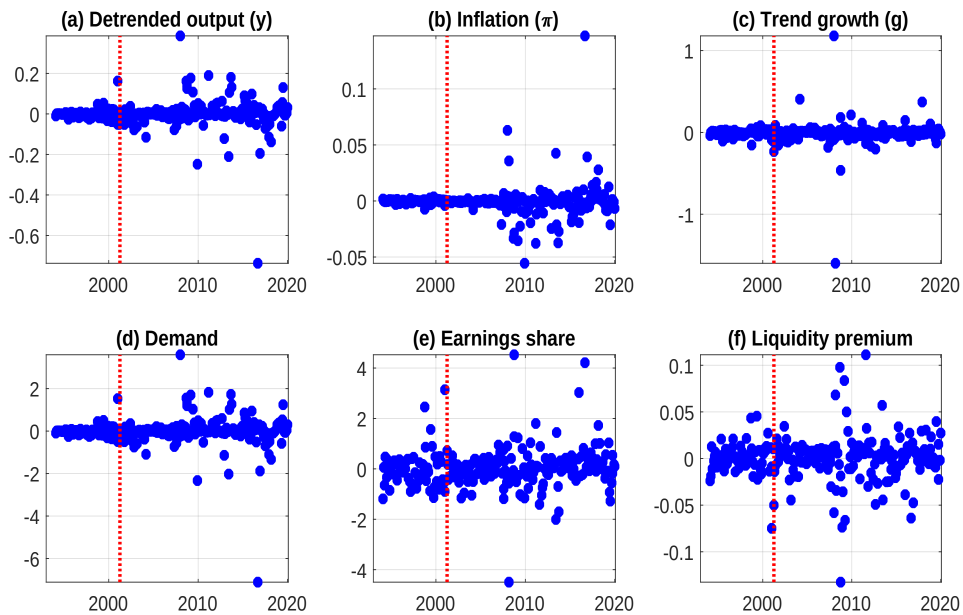

Notes: Estimated changes in the perceived state of the economy from 10 minutes before to 20 minutes after an FOMC press release. The full sample has 220 FOMC announcements spanning February 4th, 1994 to February 28th, 2020. The sample reported in the figure is 1993:M1-2020:M2.

the minutes and days surrounding an FOMC meeting to observe changes in the *perceived shocks* $\varepsilon^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$ that investors must have discerned in order to explain revisions in $S^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$ and $\pi^i_{t|t-1+\delta_h}$. To do so consider the model solution applied to the intramonth nowcasting updates:

$$\begin{aligned} S^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h} &= C\left(\theta_{\xi^{P,A}_t}, \xi^b_t = i, \mathbf{H}^b\right) + T(\theta_{\xi^{P,A}_t}, \xi^b_t = i, \mathbf{H}^b) S^j_{t-1} \\ +R(\theta_{\xi^{P,A}_t}, \xi^b_t &= i, \mathbf{H}^b) Q \varepsilon^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h}, \end{aligned} \quad (17)$$

where $\varepsilon^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$ denotes the perceived Gaussian shocks estimated on the basis of data available at time $t-1+\delta_h$, conditional on being in belief regime $\xi^b_t = i$. Given estimates of $S^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$, $C(\cdot)$, $T(\cdot)$, $R(\cdot)$, Q, and S^j_{t-1} using the most likely belief regime j at $t-1$, invert (17) to solve for $\varepsilon^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$. The contribution of one particular perceived shock k is to variation in $S^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$ is given by:

$$S^{\cdot,k}_{(t\backslash t-1+\delta_h)|t-1+\delta_h} = \sum_{i=1}^{B} \pi^i_{t|t-1+\delta_h} R(\theta_{\xi^{P,A}_t}, \xi^b_t = i, \mathbf{H}^b) Q \varepsilon^{i,k}_{(t\backslash t-1+\delta_h)|t-1+\delta_h} \quad (18)$$

where $\varepsilon^{i,k}_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$ is a vector constructed by setting each element of $\varepsilon^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$ to zero other than the kth. The contribution of the belief regime is the remaining part:

$$S^{\cdot,b}_{(t\backslash t-1+\delta_h)|t-1+\delta_h} = \sum_{i=1}^{B} \pi^i_{t|t-1+\delta_h} \left[C\left(\theta_{\xi^{P,A}_t}, \xi^b_t = i, \mathbf{H}^b\right) + T(\theta_{\xi^{P,A}_t}, \xi^b_t = i, \mathbf{H}^b) S^j_{t-1}\right]. \quad (19)$$

Finally, the contribution of revisions in perceived shocks and belief regimes to jumps in observed variables X_t is computed by taking the difference between the post- and pre-announcement values of $S^{\cdot,k}_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$ and $S^{\cdot,b}_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$ and linking them back to X_t using the mapping (14). We refer to these as shock decompositions.

Figure VII reports shock decompositions for the FOMC announcements that generated the 10 most quantitatively important absolute changes in the stock market, displayed as a series of bar charts. For each shock decomposition figure, we report the jump in a high-frequency series (e.g., the S&P 500) around the policy announcement with a black dot, and the jump implied by the estimated structural model with a red triangle. The difference between the two is due to observation error. For the S&P 500, the figure always shows black dots that lie on top of the red triangles because observation errors are set to zero.

Of the 10 announcements displayed in Figure VII, the one associated with the largest change in the stock market occurred on January 3, 2001 when the Fed met off-cycle to lower the target funds rate by 50 basis points, driving the S&P 500 surge 4.2% over the 30 minutes surrounding the news. The second and third most important FOMC events for the stock market in our sample occurred on April 18, 2001 and October 29, 2008, respectively, when the market increased 2.5% and declined 2%, respectively, in the 30 minutes surrounding those press releases. For these and several subsequent events, jumps in investor beliefs about regime change in the conduct of future monetary policy played the largest quantitative role in the market's jump.

To understand these results, it is important to note that our linearized structural model captures the first-order, direct, effects of each channel. Second- and higher-order effects, though not explicitly modeled, can still be accounted for in the empirical procedure by observed movements in data. The event of Jan. 3, 2001 provides an example. This action–widely interpreted as an unexpected insurance cut in response to the dot-com bust–leads us to estimate that the announcement led to investors to perceive a large accommodative monetary policy shock as can be seen from the shock decompositions for the 6-month FFF in panel (c) of Figure VII for the January 3, 2001 event. This shock explains most of the movement in the FFF rate. Yet panel (d) of Figure VII shows that this same shock was not the main driver of the stock market's spike upward. Instead, the estimates imply that the market surged primarily for two reasons: investors revised up their nowcast of the corporate earnings share, and revised down their nowcast of the liquidity premium. Since the market's spike cannot be largely attributed to the direct impact of the monetary policy shock or to changing beliefs about future policy regimes, such revisions must be related to estimated higher-order effects of the surprise cut in rates.

What in the data identifies these revisions in nowcasts for the earnings share and

Figure VII: Top Ten FOMC: SP500

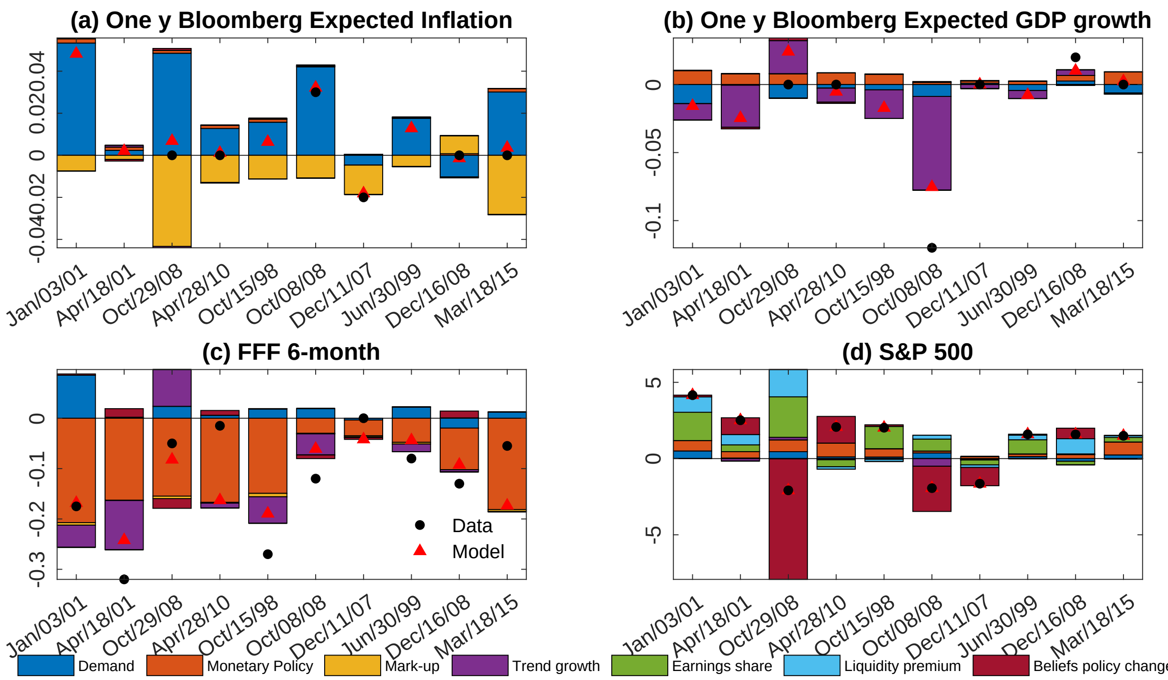

Notes: The figure reports shock decomposition for the 10 most relevant FOMC announcements based on changes in the SP500-lagged GDP ratio. The sample is 1961:M1-2020:M2.

liquidity premium? Figure VIII, panels (b) and (c), shows that they are identified by jumps in two high-frequency series in the data: First, the daily BBG survey expectation for the S&P 500 earnings nowcast for current quarter earnings jumped up around the policy announcement. Second, the BAA spread, which we map into the liquidity premium component of the subjective equity premium, jumped down surrounding the news. We refer to the announcement of January 3, 2001 as a "non-belief" event, because beliefs about the conduct of monetary policy play little role. Two other non-belief events are shown in figure VIII with similar patterns for these high-frequency series. Panel (a) shows the shock decompositions for these events, while panels (b) and (c) show jumps in the BBG and BAA data around the policy announcements.

An alternative interpretation of the January, 2001 announcement is that investors updated their assessment of the parameters of the current rather than future policy rule. This is not a feature of our baseline model, but we allow for it in an alternative version of the model and report the results in Appendix J of the Online Appendix, which shows that the results for the stock market shown in Figure VII are virtually unchanged if investors update their understanding of the current rule after the announcement–even by sizable amounts. The reason for this is that what matters for the long-duration stock market is not where the policy rule is today, but where it is likely to settle for the foreseeable future. Beliefs about the latter are already captured in our baseline

Figure VIII: Biggest Non-Belief Events

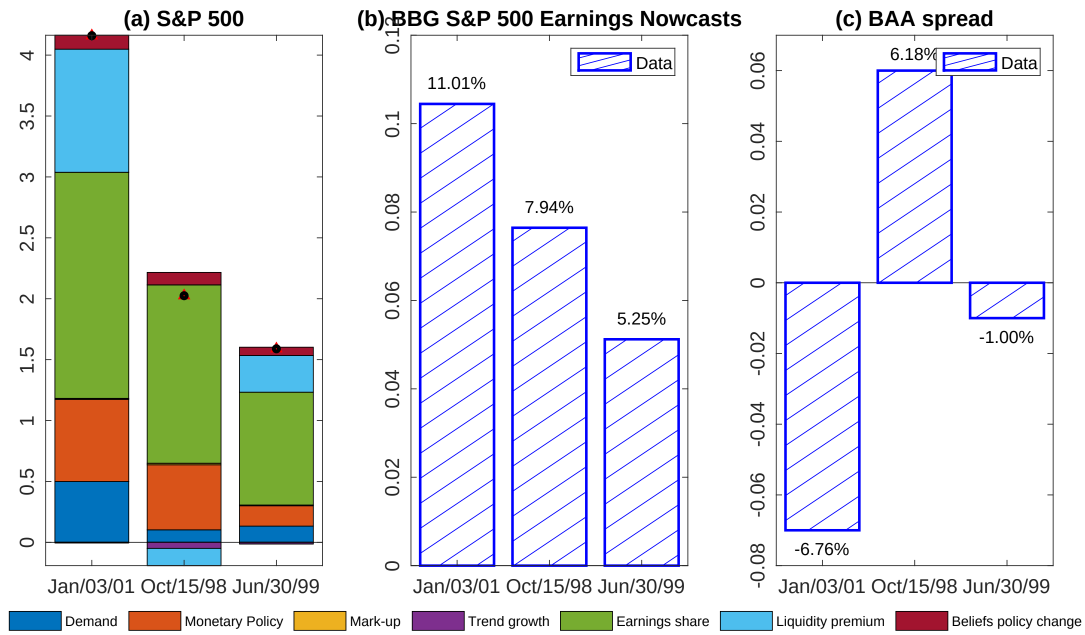

Notes: The figure reports shock decomposition for the 3 most relevant FOMC announcements based on changes in the SP500-lagged GDP ratio in which beliefs about the conduct of future policy play little role. The sample is 1961:M1-2020:M2.

model by the investor's continuous belief updating about the probability of moving to the Alternative rule, so allowing for additional updating about the parameters of the current rule changes little about the long-term outlook relevant for the stock market.

The second largest jump in the stock market around an FOMC announcement in our sample occurs on April 18, 2001 when the market leapt 2.5% after the Greenspan Fed again surprised with another 50 basis point reduction in the funds rate. Figure VII, shows that the big driver of the stock market in this event was a jump in investor beliefs about future monetary policy. Specifically, the estimates imply that investors revised upward the perceived probability that the conduct of monetary policy going forward would more aggressively protect against the downside risks that affect stocks. These estimates are identified in the data by jumps in both the stock market and in Fed funds futures rates that cannot be explained by simply rolling forward the current policy rule at time t inflation and output growth expectations. The results for this event are new to the literature and illustrate an important channel of monetary transmission to markets, namely the role of Fed communications in altering investor beliefs about future Fed policy to contain economic risks, thereby immediately impacting subjective risk premia. The next section provides additional evidence that these movements in beliefs about the future conduct of monetary policy are linked to movements in subjective risk premia.

Fed "Information Shocks" It is instructive to explore the question of Fed "information shocks" as a possible driver of the stock market's reaction to some Fed announcements. To do so, we define an event with a Fed information shock following Jarocinski and Karadi (2020) as one that generates a positive comovement between short-term interest rates and the stock market within a narrow window around a policy announcement. The idea here is that, if the market learns new negative information about the state of the economy from a Fed announcement, the information shock causes markets to fall *despite* a cut in rates. Figure IX shows several panels for the three largest such "information" events in our sample, as measured by how much the stock market responded to the policy announcement in absolute terms. These three events are the FOMC announcements of October 29, October 8, and March 11, of 2008, in which the stock market jumped -2.08%, -1.93%, and 1.09% percent, respectively, in the 30 minutes the policy announcement. Panel (a) shows the stock market shock decompositions for these three announcements. Panel (b) shows the interest rate jump around the policy announcement, as measured by the 1-month FFF rate. Panel (c) shows the resulting change in perceived probability of transitioning to the Alternative policy rule in the next year, due to each announcement. We label these "Quasi-Information" events in the Figure, since we do not find that they can be readily attributed to true Fed information shocks, as discussed next.

The first point to note is that, for all three events, there is a clear positive comovement between the stock market jump and the interest rate, as seen by the common directional change of the black dots in panel (a) for the stock market and panel (b) for the FFF rate. Thus, the three events qualify as Fed information shock events, according to the definition of Jarocinski and Karadi (2020). The second point to note is that, for all events, the positive comovement is entirely attributable to estimated shifts around the policy announcement in investor perceptions about the future conduct of monetary policy, rather than to updates in investor beliefs about the economic state. This can be observed in panel (a) of Figure IX by noting that, in each event, changing beliefs about the future conduct of monetary policy (maroon bars) explains more than 100% of the directional change in the stock market, thereby more than offsetting all other countervailing forces attributable to perceived shocks to the economic state.

Consider the FOMC announcement of October 29, 2008. The Fed announced that it was lowering its target for the FFR by 50 basis points to 1%, while the stock market declined 2% over the 30 minute window around the announcement. Despite the announcement's statements about the slowing 'pace of economic activity,' the shock decompositions in panel (a) of Figure IX imply that perceptions about the economic state were, if anything, revised in a *favorable* direction around the policy announcement, suggesting that investors believed the rate cut would at least help to support the

Figure IX: Biggest Quasi-Information Events

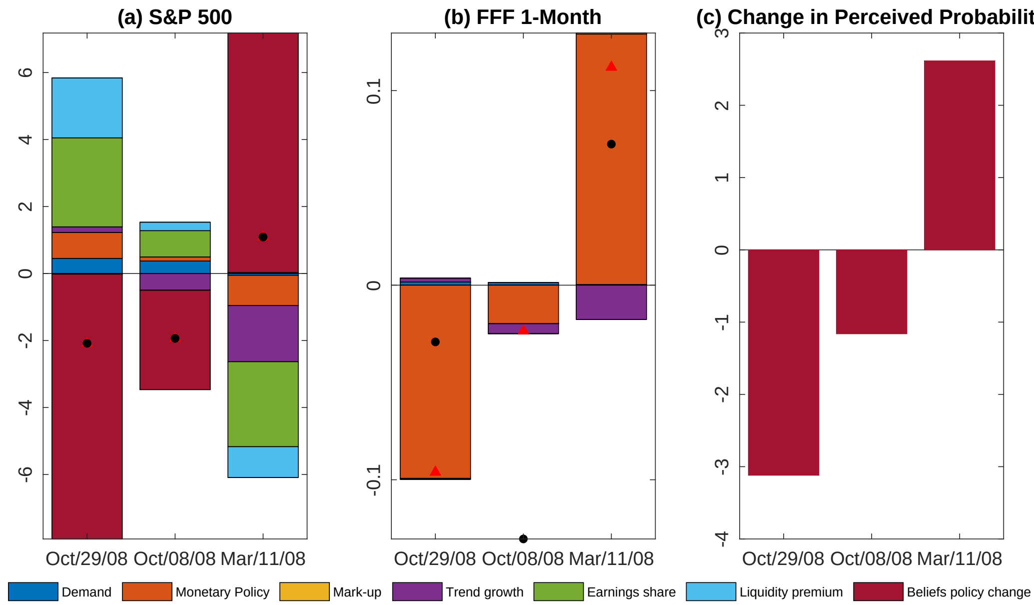

Notes: The figure reports shock decomposition for the three largest "Quasi-Information" effect FOMC events, as measured by the absolute change in the stock market. A "Quasi-Information" effect event is defined as one with a positive comovement between the stock market and the 1-month FFF rate in the 30-minute window surrounding the annoucement. We restrict the sample so as to consider only events in which the FFF rate moves at least one standard deviation from its historic mean. The sample is 1961:M1-2020:M2.

economy. Indeed, for this FOMC event, the announcement-driven upward revision in investor nowcasts for the earnings share and downward revision in the nowcast for the liquidity premium (lowering risk premia) were both positive impulses for the market that by themselves would have lead to a rise in the market. This implies that the market declined for other reasons, and Figure IX, panel (a) shows that the other reason is that investors changed their beliefs about the future conduct of monetary policy. Specifically, they lowered their expectations for a shift in the parameters of the policy rule to more aggressively stabilize economic growth. This change in beliefs was a large negative impulse for the market that more than offset the positive contributions coming from revisions in investor beliefs about the earnings share and liquidity premium. The next two biggest Quasi-Information events show a very similar pattern.

We can ask why beliefs about the future conduct of monetary policy might have predominated in the overall market movements. The answers invariably depend on context-specific forces. For the October 29th, 2008 announcement, one possibility is that the federal funds rate was still left standing at 1%, despite coming in the midst of the worst financial crisis in a century. This 1% rate merely matched the lowest previous level for the funds rate achieved a few years earlier during a non-crisis period. The negative stock market reaction that ensued is consistent with the idea that the announcement fueled worries about the 1% historical nadir becoming a barrier to further action, leaving the Fed with limited capacity to respond to future problems.

Discount Rate or Cash Flow News? In principle, the actions and announcements of central banks can affect financial markets through either discount rate or cash flow effects, or both. To study these different channels, we decompose the price-output ratio into components of the representative investor's subjective expectations:

$$\frac{P_t^D}{Y_{t-1}} = \frac{P_t^D}{D_t}\frac{D_t}{Y_t}\frac{Y_t}{Y_{t-1}}$$

or in logs

$$pgdp_t = pd_t + k_t + \Delta y_t, \tag{20}$$

where $pgdp_t \equiv \ln\left(P_t^D/Y_{t-1}\right)$ and $pd_t \equiv \ln\left(P_t^D/D_t\right)$. Let r_t^{ex} denote the log return r_t^D in excess of the log real interest rate, rir_t. Decompose pd_t as in Campbell and Shiller (1989) into the sum of three forward-looking terms:

$$pd_t = \frac{\kappa_{pd,0}}{1-\kappa_{pd,1}} + pdv_t\left(\Delta d\right) - pdv_t\left(r^{ex}\right) - pdv_t\left(rir\right) \tag{21}$$

where $pdv_t\left(x\right) \equiv \sum_{h=0}^{\infty}\beta_p^h\mathbb{E}_t^b\left[x_{t+1+h}\right]$, $rir_{t+1} \equiv \left(i_{t+1} - \mathbb{E}_t^b\left[\pi_{t+1}\right]\right)$ are computed under the subjective expectations of the investor $\mathbb{E}_t^b\left[\cdot\right]$. Subjectively expected return premia $pdv_t\left(r^{ex}\right)$ are driven in the model by three factors: (i), realized regime change in monetary policy ξ_t^P, (ii) changing investor beliefs about the probability of future regime change ξ_t^b, and (iii) the liquidity premium lp_t. Subjectively expected real interest rates $pdv_t\left(rir\right)$ depend these factors, as well as on expectations about inflation and output growth that enter the monetary policy rule. Substituting (21) into (20), we can decompose $pgdp_t$ as:

$$pgdp_t = \underbrace{ey_t}_{\text{earning share}} + \underbrace{pdv_t\left(\Delta d\right)}_{\text{earnings}} - \underbrace{pdv_t\left(r^{ex}\right)}_{\text{premia}} - \underbrace{pdv_t\left(rir\right)}_{\text{real int rate}}, \tag{22}$$

where $ey_t \equiv \frac{\kappa_{pd,0}}{1-\kappa_{pd,1}} + k_t + \Delta y_t$ is the earnings-to-lagged output ratio, or "earnings share."

Figure X examines cash flow versus discount rate effects at high frequency around FOMC announcements. The figure decomposes the announcement-related jumps in pd_t into fluctuations driven by the $pdv_t\left(\cdot\right)$ components on the right-hand-side of (21) for the 5 most relevant FOMC announcements sorted on the basis of jumps in the estimated perceived probability of a regime change in the conduct of monetary policy over the next year. Panel (a) of Figure X shows how the perceived probabilities of a regime change shifted in the 30 minute windows surrounding each FOMC announcement, while panel (b) shows the decomposition of the jump in pd_t into its $pdv_t\left(\cdot\right)$ components.

Figure X: Jumps in risk perceptions, short rates, and earnings expectations

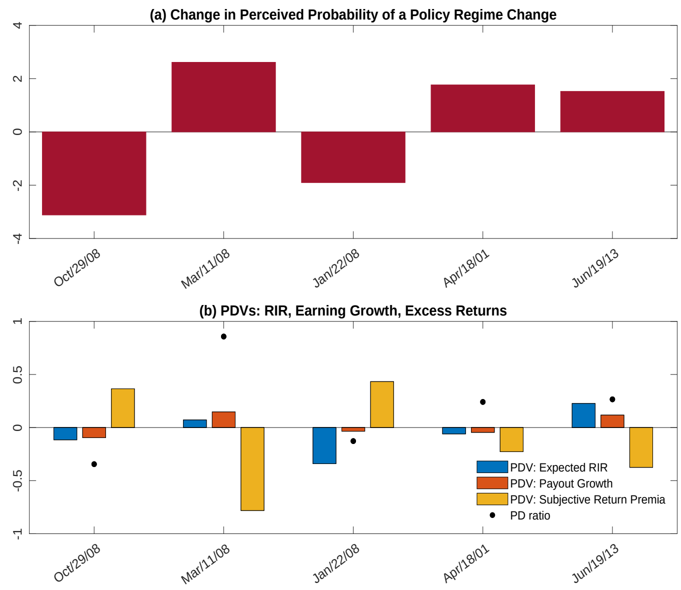

Notes: Panel (a) shows the pre-/post-FOMC announcement change (10 minutes before/20 minutes after) in the perceived probability of a monetary policy regime change occurring within one year. Panel (b) decomposes the jump in the log price-payout ratio $pd = pdv_t(\Delta d) - pdv_t(r^{ex}) - pdv_t(rir)$ into movements in the subjective equity risk premia $pdv_t(r^{ex})$ (yellow bar), subjective expected real interest rates $pdv_t(RIR)$ (blue bar), and subjective expected payout growth $pdv_t(\Delta d)$ (red bar). PD ratio is $pdv_t(\Delta d) - pdv_t(r^{ex}) - pdv_t(rir)$. The sample is 1961:M1-2020:M2.

The FOMC announcement of October 29, 2008, is associated with the largest absolute decline in the perceived probability of monetary regime change, as shown in panel (a). In this case, the perceived probability that the central bank would soon transition to an Alternative policy rule featuring more aggressive stabilization of the real economy falls sharply, resulting in a large jump up in subjective risk premia and jump down in pd. Such a belief would dash hopes for a more aggressive "Fed put" policy rule, since it means that the central bank is less likely to engineer a lower discount rate whenever cash flows are expected to grow more slowly. This event suggests that the announcement fueled worries the Fed might soon re-approach the ELB with limited capacity for lower rates.

Conversely, the March 11, 2008 announcement that the Federal Reserve announced an expansion of securities lending via the creation of a new facility, the Term Securities Lending Facility (TSLF) is the event associated with largest *in*crease in the perceived probability of exiting the policy rule over the next 12 months, and is the mirror image of the October 29, 2008 event. The stock market rose surrounding this announcement, and Panel (b) shows that the most important contributor to this rise was a decline in subjective return premia driven by a lower perceived quantity of risk in the stock market. Evidently, the announcement of new lending facilities to support the economy amid a financial crisis triggered an increase in the perceived probability of shifting to a new monetary policy regime characterized by more aggressive stabilization of economic growth. The "Fed put" flavor of this event shows how Fed news can move markets by altering beliefs about future policies to limit downside risk, immediately changing risk premia.

6 Conclusion

What do financial markets learn from news? We develop a methodology to address this question using the stock market's reaction to Fed announcements as a case study. To do so, we integrate a high-frequency monetary event study into a mixed-frequency macro-finance model and structural estimation. The approach allows for jumps at Fed announcements in investor beliefs, providing granular detail on why markets react to central bank announcements. We also provide a methodology for modeling expectations in the presence of structural breaks, and show how forward-looking data can be used to infer what agents expect from the next policy regime. The overall approach can be used in a variety of other settings to provide a richer understanding of the role of news shocks in driving financial market volatility.

The heightened responsiveness of financial markets to central bank communications raises an important question: What are the underlying drivers of this phenomenon? We find that the reasons involve a mix of factors, including revisions in investor beliefs about the latent state of the economy, uncertainty over the future conduct of monetary policy, and subjective reassessments of risk in the stock market. These dynamics stem from two primary sources. First, beliefs about the conduct of future policy react to Fed news even if current policy is unchanged, affecting the perceived quantity of risk in the stock market. Second, some announcements are associated with sizable shifts in investor perceptions of certain elements of the economic state, including the corporate earnings share of output, and credit spreads. The methodology permits a precise quantification of the relevance of these channels across announcements in event-time.

References

Bansal, R., and H. Zhou (2002): "Term structure of interest rates with regime shifts," *The Journal of Finance*, 57(5), 1997–2043.

Bauer, M. D., C. Pflueger, and A. Sundaram (2022): "Perceptions About Monetary Policy," Unpublished manuscript, University of Chicago.

Bauer, M. D., and E. T. Swanson (2023): "An Alternative Explanation for the "Fed Information Effect"," *American Economic Review*, 113(3), 664–700.

Bernanke, B. S., and K. N. Kuttner (2005): "What explains the stock market's reaction to rederal reserve policy?," *Journal of Finance*, 60, 1221–1257.

Bianchi, F. (2013): "Regime switches, agents' beliefs, and post-World war II U.S. macroeconomic dynamics," *Review of Economic Studies*, 80(2), 463–490.

——— (2016): "Methods for measuring expectations and uncertainty in Markov-switching models," *Journal of Econometrics*, 190(1), 79–99.

Bianchi, F., R. Gómez-Cram, T. Kind, and H. Kung (2023): "Threats to central bank independence: High-frequency identification with twitter," *Journal of Monetary Economics*, 135, 37–54.

Bianchi, F., H. Kung, and M. Tirskikh (2018): "The origins and effects of macroeconomic uncertainty," Discussion paper, National Bureau of Economic Research.

Bianchi, F., M. Lettau, and S. C. Ludvigson (2022): "Monetary Policy and Asset Valuaton," *The Journal of Finance*, 77(2), 967–1017.

Bianchi, F., S. C. Ludvigson, and S. Ma (2024): "What Hundreds of Economic News Events Say About Belief Overreaction in the Stock Market," Discussion paper, National Bureau of Economic Research, No. w32301.

Brooks, J., M. Katz, and H. Lustig (2018): "Post-FOMC announcement drift in US bond markets," Discussion paper, National Bureau of Economic Research.

Caballero, R. J., and A. Simsek (2022): "A Monetary Policy Asset Pricing Model," *Available at SSRN 4113332*.

Campbell, J. R., C. L. Evans, J. D. Fisher, and A. Justiniano (2012): "Macroeconomic effects of federal reserve forward guidance," *Brookings papers on economic activity*, pp. 1–80.

Campbell, J. Y., and J. H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.

Campbell, J. Y., and R. J. Shiller (1989): "The dividend-price ratio and expectations of future dividends and discount factors," *Review of Financial Studies*, 1(3), 195–228.

Chang, M., X. Chen, and F. Schorfheide (2021): "Heterogeneity and aggregate fluctuations," Discussion paper, National Bureau of Economic Research No. w28853.

Cieslak, A., and H. Pang (2021): "Common shocks in stocks and bonds," *Journal of Financial Economics*, 142, 880–904.

Cieslak, A., and A. Schrimpf (2019): "Non-monetary news in central bank communication," *Journal of International Economics*, 118, 293–315.

Cieslak, A., and A. Vissing-Jorgensen (2021): "The economics of the Fed put," *The Review of Financial Studies*, 34(9), 4045–4089.

Clarida, R., J. Gali, and M. Gertler (2000): "Monetary policy rules and macroeconomic stability: evidence and some theory," *Quarterly Journal of Economics*, 115(1), 147–180.

Cochrane, J. H., and M. Piazzesi (2002): "The fed and interest rates-a high-frequency identification," *AEA Papers and Proceedings*, 92(2), 90–95.

Coibion, O., Y. Gorodnichenko, S. Kumar, and M. Pedemonte (2020): "Inflation expectations as a policy tool?," *Journal of International Economics*, 124, 103297.

Farmer, R. E., D. F. Waggoner, and T. Zha (2011): "Minimal state variable solutions to Markov-switching rational expectations models," *Journal of Economic Dynamics and Control*, 35(12), 2150–2166.

Gennaioli, N., Y. Ma, and A. Shleifer (2016): "Expectations and investment," *NBER Macroeconomics Annual*, 30(1), 379–431.

Gertler, M., and P. Karadi (2015): "Monetary policy surprises, credit costs, and economic activity," *American Economic Journal: Macroeconomics*, 7(1), 44–76.

Ghysels, E., and J. H. Wright (2009): "Forecasting professional forecasters," *Journal of Business & Economic Statistics*, 27(4), 504–516.

Giannone, D., L. Reichlin, and D. Small (2008): "Nowcasting: The real-time informational content of macroeconomic data," *Journal of monetary economics*, 55(4), 665–676.

Gilchrist, S., D. López-Salido, and E. Zakrajšek (2015): "Monetary policy and real borrowing costs at the zero lower bound," *American Economic Journal: Macroeconomics*, 7(1), 77–109.

Gormsen, N. J., and K. Huber (2022): "Corporate Discount Rates," *Available at SSRN.*

Hamilton, J. D. (1994): *Time Series Analysis.* Princeton University Press, Princeton, NY.

Hanson, S. G., and J. C. Stein (2015): "Monetary policy and long-term real rates," *Journal of Financial Economics*, 115(3), 429–448.

Hillenbrand, S. (2021): "The Fed and the Secular Decline in Interest Rates," *Available at SSRN 3550593.*

Honkapohja, S., K. Mitra, and G. W. Evans (2013): "Notes on Agents' Behavioural Rules under Adaptive Learning and Studies of Monetary Policy," in *Macroeconomics at the Service of Public Policy*, ed. by T. J. Sargent, and J. Vilmunen, pp. 63–79. Oxford University Press, Great Clarendon Street, Oxford, OX2 6DP.

Jarocinski, M., and P. Karadi (2020): "Deconstructing monetary policy surprises: the role of information shocks," *American Economic Journal: Macroeconomics*, 12, 1–43.

Kekre, R., and M. Lenel (2021): "Monetary Policy, Redistribution, and Risk Premia," NBER Working paper No. 28869.

Kim, C.-J. (1994): "Dynamic Linear Models with Markov-Switching," *Journal of Econometrics*, 60, 1–22.

Krishnamurthy, A., and A. Vissing-Jorgensen (2011): "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy," *Brookings Papers on Economic Activity*, 2011(2), 215–287.

——— (2012): "The aggregate demand for Treasury debt," *Journal of Political Economy*, 120(2), 233–267.

Laubach, T., and J. C. Williams (2003): "Measuring the natural rate of interest," *Review of Economics and Statistics*, 85(4), 1063–1070.

Lettau, M., and J. Wachter (2007): "The Term Structures of Equity and Interest Rates," Unpublished paper.

Lubik, T. A., and F. Schorfheide (2004): "Testing for indeterminacy: An application to US monetary policy," *American Economic Review*, 94, 190–217.

Malmendier, U., and S. Nagel (2016): "Learning from inflation experiences," *Quarterly Journal of Economics*, 131(1), 53–87.

Melosi, L. (2017): "Signalling effects of monetary policy," *The Review of Economic Studies*, 84(2), 853–884.

Nagel, S. (2024): "Leaning against inflation experiences," *Unpublished manuscript, University of Chicago.*

Nakamura, E., and J. Steinsson (2018): "High-frequency identification of monetary non-neutrality: The information effect," *Quarterly Journal of Economics*, 133(3), 1283–1330.

Pflueger, C., and G. Rinaldi (2020): "Why Does the Fed Move Markets so Much? A Model of Monetary Policy and Time-Varying Risk Aversion," Discussion paper, National Bureau of Economic Research.

Piazzesi, M. (2005): "Bond yields and the Federal Reserve," *Journal of Political Economy*, 113(2), 311–344.

Piazzesi, M., and E. Swanson (2008): "Futures Prices as Risk-Adjusted Forecasts of Monetary Policy," *Journal of Monetary Economics*, 55, 677–691.

Romer, C. D., and D. H. Romer (2000): "Federal Reserve information and the behavior of interest rates," *American economic review*, 90(3), 429–457.

Schorfheide, F., and D. Song (2015): "Real-time forecasting with a mixed-frequency VAR," *Journal of Business & Economic Statistics*, 33(3), 366–380.

Sims, C. A., and T. Zha (2006): "Were there regime switches in US monetary policy?," *American Economic Review*, 91(1), 54–81.

Stock, J. H., and M. W. Watson (2010): "Research Memorandum," `https://www.princeton.edu/~mwatson/mgdp_gdi/Monthly_GDP_GDI_Sept20.pdf`.

Taylor, J. B. (1993): "Discretion versus policy rules in practice," in *Carnegie-Rochester conference series on public policy*, vol. 39, pp. 195–214. Elsevier.

Online Appendix

A Priors, Posterior, and Smoothed Series

Table A.I describes the posterior (left-hand-side of the table) and prior (right-hand-side of the table) distributions for the parameters of the model. In the column "Type," N stands for Normal, G stands for Gamma, IG stands for Inverse Gamma, and B stands for Beta distribution, respectively. For all prior distributions, we report the mean and the standard deviation. The priors for all parameters are diffuse and centered around values typically found in the literature. We choose symmetric priors for the parameters of the realized and alternative policy rules. For the posterior, we report the mode and 90% credible sets.

Figure A.I: Smoothed Series

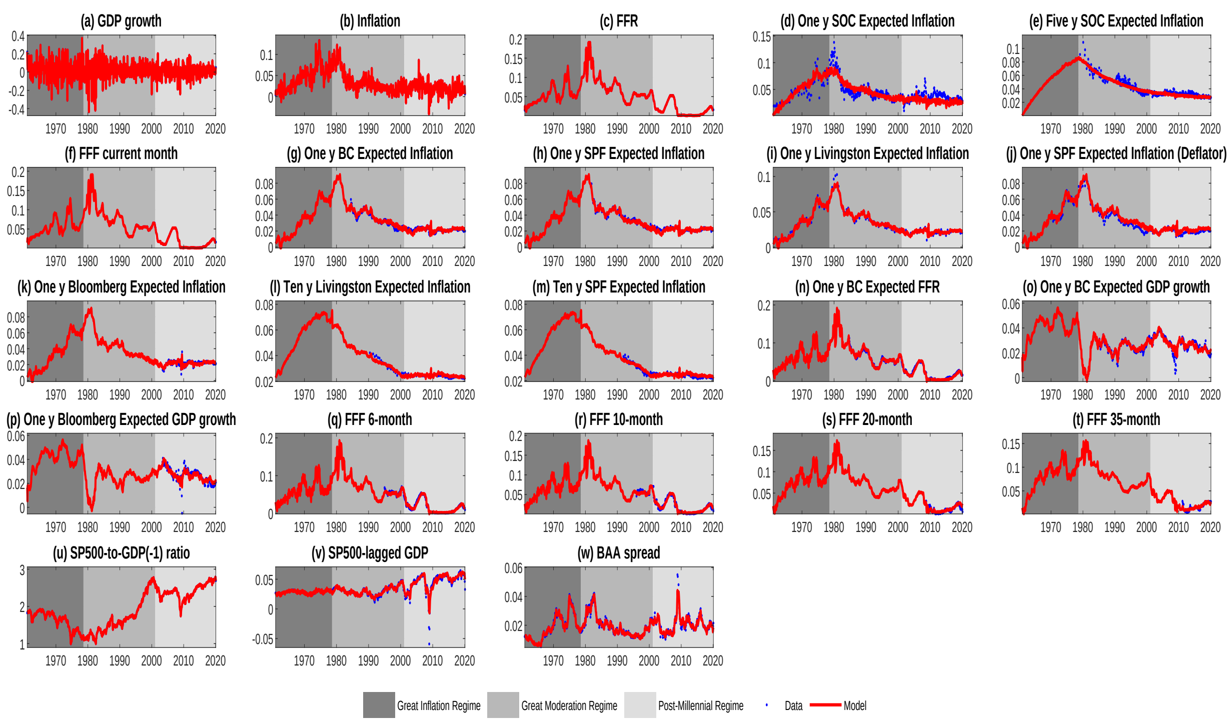

Notes: The figure displays the model-implied series (red, solid line) and the actual series (blue dotted line). The model-implied series are based on smoothed estimates $S_{t|T}$ of S_t, and exploit the mapping to observables in (14) using the modal parameter estimates. The difference between the model-implied series and the observed counterpart is attributable to observation error. We allow for observation errors on all variables except for GDP growth, inflation, the FFR, and the SP500 capitalization to GDP ratio. Great Inflation Regime: 1961:Q1-1978:Q3. Great Moderation Regime: 1978:Q4-2001:Q3. Post-Millennial Regime: 2001:Q4-2020:Q1. The sample is 1961:M1-2020:M2.

B Data

Real GDP

The real Gross Domestic Product is obtained from the US Bureau of Economic Analysis. It is in billions of chained 2012 dollars, quarterly frequency, seasonally adjusted, and at annual rate. The source is from Bureau of Economic Analysis (BEA code: A191RX). The sample spans 1959:Q1 to 2021:Q2. The series was interpolated to monthly frequency using the method in Stock and Watson (2010). The quarterly series was downloaded on August 20th, 2021.

GDP price deflator

The Gross Domestic Product: implicit price deflator is obtained from the US Bureau of Economic Analysis. Index base is 2012=100, quarterly frequency, and seasonally adjusted. The source is from Bureau of Economic Analysis (BEA code: A191RD). The sample spans 1959:Q1 to 2021:Q2. The series was interpolated to monthly frequency using the method in Stock and Watson (2010). The quarterly series was downloaded on August 20th, 2021.

Federal funds rate (FFR)

The Effective Federal Funds Rate is obtained from the Board of Governors of the Federal Reserve System. It is in percentage points, quarterly frequency, and not seasonally adjusted. The sample spans 1960:02 to 2021:06. The series was downloaded on August 20th, 2021

SP500 and SP500 futures

For our high-frequency analysis, we use tick-by-tick data on SP500 index obtained from tickdata.com. The series was downloaded on September 22th, 2021 from `https://www.tickdata.com/`.We create the minutely data using the close price within each minute. Within trading hours, we construct minutely S&P 500 market capitalization by multiplying the S&P 500 index by the previous month's S&P 500 Divisor. (The index is the market capitalization of the 500 companies covered by the index divided by the S&P 500 divisor, roughly the number of shares outstanding across all companies.) The S&P 500 divisor is available at the URL: `https://ycharts.com/indicators/sp_500_divisor`. We supplement SP500 index using SP500 futures for events that occur in off-market hours. We use the current-quarter contract futures. We purchased the SP500 futures from CME group at URL: https://datamine.cmegroup.com/. Our sample spans January 2nd 1986 to September 17th, 2021. The SP500 futures data were downloaded on October 6, 2021.

SP500 Earnings and Market Capitalization

We use S&P 500 earnings divided by GDP as a noisy signal on K_t in the structural estimation. To map into a monthly estimation, we ideally would use monthly earnings data. Instead, we have quarterly S&P 500 earnings per share (EPS) data that starts in 1988:Q2 from S&P Global `https://www.spglobal.com/spdji/en/documents/additional-material/sp-500-eps-est.xlsx`, which we linearly interpolate (and divide it by 3) to monthly observations. To extend our sample backward, we use monthly (interpolated) EPS data on the S&P 500 from Robert Shiller's data depository at URL:`http://www.econ.yale.edu/~shiller/data/ie_data.xls`. These are monthly EPS data computed from the S&P four-quarter trailing totals (and divide it by 12) and we use data for this series that span 1959:01-1988:03. (There is no quarterly EPS data prior to 1988:03. Observation error in the structural estimation absorbs fluctuations attributable to splicing together an annual moving average with an annualized quarterly sequence of observations.) Splicing the two series together results in a monthly earnings per share series spanning 1959:01 to 2021:06. In the S&P Global dataset there is one observation in 2008:Q4 with a negative EPS, which appears in 2008:12 after interpolation. To deal with this in earnings growth rate calculations we replace the 2008:12 observation with the Shiller four quarter total trailing EPS observation in 2008:12. To convert EPS to total earnings, we multiply EPS by the monthly S&P 500 Divisor available at URL:`https://ycharts.com/indicators/sp_500_divisor`. For a monthly stock market value series, we use the S&P 500 market capitalization, obtained as the end-of-month series from Ycharts.com available at `https://ycharts.com/indicators/sp\_500\_market\_cap`. All series span the periods 1959:01 to 2021:06 and were downloaded on December 22nd, 2021.

Baa Spread, 20-yr T-bond, Long-term US government securities

We obtained daily Moody's Baa Corporate Bond Yield from FRED (series ID: DBAA) at URL: `https://fred.stlouisfed.org/series/BAA`, US Treasury securities at 20-year constant maturity from FRED (series ID: DGS20) at URL: `https://fred.stlouisfed.org/series/DGS20`, and long-term US government securities from FRED (series ID: LTGOVTBD) at URL: `https://fred.stlouisfed.org/series/LTGOVTBD`. The sample for Baa spans the periods 1986:01 to 2021:06. To construct the long term bond yields, we use LTGOVTBD before 2000 (1959:01 to 1999:12) and use DGS20 after 2000 (2000:01 to 2021:06). The Baa spread is the difference between the Moody's Corporate bond yield and the 20-year US government yield. The excess bond premium is obtained at URL: `https://www.federalreserve.gov/econres/notes/feds-notes/ebp_csv.csv`. All series were downloaded on February 21, 2022.

Bloomberg Consensus Inflation, GDP forecasts, and Earnings Nowcasts

We obtain the Bloomberg (BBG) US GDP (id: ECGDUS) and inflation (id: ECPIUS) consensus mean forecast from the Bloomberg Terminal available on a daily basis up to a few days before the release of GDP and inflation data. The Bloomberg (BBG) US consensus forecasts are updated daily (except for weekends and holidays) and reports daily quarter-over-quarter real GDP growth and CPI forecasts from 2003:Q1 to 2021Q2. These forecasts provide more high-frequency information on the professional outlook for economic indicators. Both forecast series were downloaded on October 21, 2021.

We obtain daily nowcasts of S&P 500 earnings per share (EPS) from BBG. The survey respondents are equity strategists that are asked to provide nowcasts of earnings per share (EPS) for the constituents of the S&P 500. For each S&P 500 constituent, BBG provides the mean nowcast across survey respondents as well as a bottom-up aggregate nowcast of EPS for the S&P 500 by aggregating the EPS nowcasts across the S&P 500 constituents. We construct a mean respondent nowcast for the level of S&P 500 earnings by multiplying this aggregate with the S&P 500 index divisor. (The index is the market capitalization of the 500 companies covered by the index divided by the S&P 500 divisor, roughly the number of shares outsanding across all companies.) The S&P 500 divisor is available at the URL: `https://ycharts.com/indicators/sp_500_divisor`. These nowcasts are available daily for the current standard financial quarter (Jan-Mar, Apr-June, Jul-Sep, Oct-Dec). For example, the observation for July 10, 2024 (which falls under the standard financial quarter of 2024:Q3) would contain nowcasts for 2024:Q3 EPS for the S&P 500. Bloomberg does not require respondents to submit their nowcasts on a specific timeline or frequency. Instead, respondents voluntarily decide how often to update their nowcasts. To ensure that consensus nowcasts are not heavily influenced by outdated information, Bloomberg excludes stale nowcasts submitted before the most recent earnings announcement date. The data was downloaded from the Bloomberg terminal on July 24, 2024, using the Earnings & Estimates (EE) function on the S&P 500 index (SPX Index). The aggregated consensus nowcasts are available daily, except weekends and holidays, spanning the period from Jan 2, 2006 to July 24, 2024. The divisor series series span the periods 1959:01 to 2021:06 and was downloaded on December 22nd, 2021.

Livingston Survey Inflation Forecast

We obtained the Livingston Survey mean 1-year and 10-year CPI inflation forecast from the Federal Reserve Bank of Philadelphia, URL: `https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/livingston-historical-data`. Our sample spans 1947:06 to 2021:06. The forecast series were downloaded on September 20, 2021.

Michigan Survey of Consumers Inflation Forecasts

We construct MS forecasts of annual inflation of respondents answering at time t. Each month, the SOC contains approximately 50 core questions, and a minimum of 500 interviews are conducted by telephone over the course of the entire month, each month. We use two questions from the monthly survey for which the time series begins in January 1978.

1. Annual CPI inflation: To get a point forecast, we combine the information in the survey responses to questions A12 and A12b.

    - Question A12 asks (emphasis in original): *During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?*
    - A12b asks (emphasis in original): *By about what percent do you expect prices to go (up/down) on the average, during the next 12 months?*

2. Long-run CPI inflation: To get a point forecast, we combine the information in the survey responses to questions A13 and A13b.

    - Question A13 asks (emphasis in original): *What about the outlook for prices over the next 5 to 10 years? Do you think prices will be higher, about the same, or lower, 5 to 10 years from now?*
    - A13b asks (emphasis in original): *By about what percent per year do you expect prices to go (up/down) on the average, during the next 5 to 10 years?*

All series were downloaded on September 17th, 2021.

Bluechip Inflation and GDP Forecasts

We obtain Blue Chip expectation data from Blue Chip Financial Forecasts from Wolters Kluwer. The surveys are conducted each month by sending out surveys to forecasters in around 50 financial firms such as Bank of America, Goldman Sachs & Co., Swiss Re, Loomis, Sayles & Company, and J.P. Morgan Chase. The participants are surveyed around the 25th of each month and the results published a few days later on the 1st of the following month. The forecasters are asked to forecast the average of the level of U.S. interest rates over a particular calendar quarter, e.g. the federal funds rate and the set of H.15 Constant Maturity Treasuries (CMT) of the following maturities: 3-month, 6-month, 1-year, 2-year, 5-year and 10-year, and the quarter over quarter percentage changes in Real GDP, the GDP Price Index and the Consumer Price Index, beginning with the current quarter and extending 4 to 5 quarters into the future.

In this study, we look at a subset of the forecasted variables. Specifically, we use the Blue Chip micro data on individual forecasts of the quarter-over-quarter (Q/Q) percentage change in the Real GDP, the GDP Price Index and the CPI, and convert to quarterly observations as explained below.

1. CPI inflation: We use quarter-over-quarter percentage change in the consumer price index, which is defined as

   "Forecasts for the quarter-over-quarter percentage change in the CPI (consumer prices for all urban consumers). Seasonally adjusted, annual rate."

   Quarterly and annual CPI inflation are constructed the same way as for PGDP inflation, except CPI replaces PGDP.

2. For real GDP growth, We use quarter-over-quarter percentage change in the Real GDP, which is defined as

   "Forecasts for the quarter-over-quarter percentage change in the level of chain-weighted real GDP. Seasonally adjusted, annual rate. Prior to 1992, Q/Q % change (SAAR) in real GNP."

The surveys are conducted right before the publication of the newsletter. Each issue is always dated the 1st of the month and the actual survey conducted over a two-day period almost always between 24th and 28th of the month. The major exception is the January issue when the survey is conducted a few days earlier to avoid conflict with the Christmas holiday. Therefore, we assume that the end of the last month (equivalently beginning of current month) is when the forecast is made. For example, for the report in 2008 Feb, we assume that the forecast is made on Feb 1, 2008. We obtained Blue Chip Financial Forecasts from Wolters Kluver in several stages starting in 2017 and with the last update purchased in June of 2022 and received on June 22, 2022. URL:https://law-store.wolterskluwer.com/s/product/blue-chip-financial-forecast-print/ 01tG000000LuDUCIA3.

Survey of Professional Forecasters (SPF)

The SPF is conducted each quarter by sending out surveys to professional forecasters, defined as forecasters. The number of surveys sent varies over time, but recent waves sent around 50 surveys each quarter according to officials at the Federal Reserve Bank of Philadelphia. Only forecasters with sufficient academic training and experience as macroeconomic forecasters are eligible to participate. Over the course of our sample, the number of respondents ranges from a minimum of 9, to a maximum of 83, and the mean number of respondents is 37. The surveys are sent out at the end of the first month of each quarter, and they are collected in the second or third week of the middle month of each quarter. Each survey asks respondents to provide nowcasts and quarterly forecasts from one to four quarters ahead for a variety of variables. Specifically, we use the SPF micro data on individual forecasts of the price level, long-run inflation, and real GDP.[1] Below we provide the exact definitions of these variables as well as our method for constructing nowcasts and forecasts of quarterly and annual inflation for each respondent.[2] Source: Federal Reserve Bank of Philadelphia. All series were downloaded on September 17th, 2021.

The following variables are used on either the right- or left-hand-sides of forecasting models:

1. Quarterly and annual inflation (1968:Q4 - present): We use survey responses for the level of the GDP price index (PGDP), defined as

   "Forecasts for the quarterly and annual level of the chain-weighted GDP price index. Seasonally adjusted, index, base year varies. 1992-1995, GDP implicit deflator. Prior to 1992, GNP implicit deflator. Annual forecasts are for the annual average of the quarterly levels."

   Since advance BEA estimates of these variables for the current quarter are unavailable at the time SPF respondents turn in their forecasts, four quarter-ahead inflation and GDP growth forecasts are constructed by dividing the forecasted level by the survey respondent-type's nowcast. Let $\mathbb{F}_t^{(i)}[P_{t+h}]$ be forecaster i's prediction of PGDP h quarters ahead and $\mathbb{N}_t^{(i)}[P_t]$ be forecaster i's nowcast of PGDP for the current quarter. Annualized inflation forecasts for forecaster i are

   $$\mathbb{F}_t^{(i)}[\pi_{t+h,t}] = (400/h) \times \ln\left(\frac{\mathbb{F}_t^{(i)}[P_{t+h}]}{\mathbb{N}_t^{(i)}[P_t]}\right),$$

   where $h = 1$ for quarterly inflation and $h = 4$ for annual inflation. Similarly, we construct quarterly and annual nowcasts of inflation as

   $$\mathbb{N}_t^{(i)}[\pi_{t,t-h}] = (400/h) \times \ln\left(\frac{\mathbb{N}_t^{(i)}[P_t]}{P_{t-h}}\right),$$

   where $h = 1$ for quarterly inflation and $h = 4$ for annual inflation, and where P_{t-1} is the BEA's advance estimate of PGDP in the previous quarter observed by the respondent in time t, and P_{t-4} is the BEA's most accurate estimate of PGDP four quarters back. After computing inflation for each survey respondent, we calculate

[1] Individual forecasts for all variables can be downloaded at https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/historical-data/individual-forecasts.

[2] The SPF documentation file can be found at https://www.philadelphiafed.org/-/media/research-and-data/real-time-center/survey-of-professional-forecasters/spf-documentation.pdf?la=en.

the 5th through the 95th percentiles as well as the average, variance, and skewness of inflation forecasts across respondents.

2. Long-run inflation (1991:Q4 - present): We use survey responses for 10-year-ahead CPI inflation (CPI10), which is defined as

   "Forecasts for the annual average rate of headline CPI inflation over the next 10 years. Seasonally adjusted, annualized percentage points. The "next 10 years" includes the year in which we conducted the survey and the following nine years. Conceptually, the calculation of inflation is one that runs from the fourth quarter of the year before the survey to the fourth quarter of the year that is ten years beyond the survey year, representing a total of 40 quarters or 10 years. The fourth-quarter level is the quarterly average of the underlying monthly levels."

   Only the median response is provided for CPI10, and it is already reported as an inflation rate, so we do not make any adjustments and cannot compute other moments or percentiles.

3. Real GDP growth (1968:Q4 - present): We use the level of real GDP (RGDP), which is defined as

   *"Forecasts for the quarterly and annual level of chain-weighted real GDP. Seasonally adjusted, annual rate, base year varies. 1992-1995, fixed-weighted real GDP. Prior to 1992, fixed-weighted real GNP. Annual forecasts are for the annual average of the quarterly levels. Prior to 1981:Q3, RGDP is computed by using the formula NGDP / PGDP * 100."*

Quarterly growth rates are constructed the same way as for inflation, except RGDP replaces PGDP.

Fed Funds Futures

We use tick-by-tick data on Fed funds futures (FFF) and Eurodollar futures obtained from the CME Group. Our sample spans January 3, 1995 to June 2, 2020. FFF contracts settle based on the average federal funds rate that prevails over a given calendar month. Fed funds futures are priced at $100 - f_t^{(n)}$, where $f_t^{(n)}$ is the time-t contracted federal funds futures market rate that investors lock in. Contracts are monthly and expire at month-end, with maturities ranging up to 60 months. For the buyer of the futures contract, the amount of $\left(f_t^{(n)} - r_{t+n}\right) \times \$\text{D}$, where r_{t+n} is the ex post realized value of the federal funds rate for month $t+n$ calculated as the average of the daily Fed funds rates in month $t+n$, and \$D is a dollar "deposit", represents the payoff of a zero-cost portfolio.

Contracts are cleaned following communication with the CME Group. First, trades with effectivevolume, which indicate a canceled order, are excluded. Floor trades, which do not require a volume on record, are included. Next, trades with a recorded expiry (in YYMM format) of 9900 indicate bad data and are excluded (Only 1390 trades, or less than 0.01% of the raw Fed funds data, have contract delivery dates of 9900). For trades time stamped to the same second, we and keep the trade with the lowest sequence number, corresponding to the first trade that second.

Fed funds futures trade prices were quoted in different units prior to August 2008. To standardize units across our sample, we start by noting that Fed funds futures are priced to the average effective Fed funds rate realized in the contract month. And in our sample, we expect a reasonable effective Fed funds rate to correspond to prices in the 90 to 100 range. As such, we rescale prices to be less than 100 in the pre-August 2008 subsample.[3] After rescaling, a small number of trades still appear to have prices that are far away from the effective Fed funds rates at both trade day and contract expiry, along with trades in the immediate transactions. The CME Group could not explain this data issue, so following Bianchi, Gómez-Cram, Kind, and Kung (2023) and others in the high frequency equity literature, we apply an additional filter to exclude trades with such non-sensible prices. Specifically, for each maturity contract, we only keep trades where

$$|p_t - \overline{p}_t(k, \delta)| < 3\sigma_t(k, \delta) + \gamma,$$

where p_t denotes the trade price (where t corresponds to a second), and $\overline{p}_t(k, \delta)$ and $\sigma_t(k, \delta)$ denote the average price and standard deviation, respectively, centered with $k/2$ observations on each side of t excluding $\delta k/2$ trades with highest price and excluding $\delta k/2$ trades with lowest price. Finally, γ is a positive constant to account for the cases where prices are constant within the window. Our main specification uses $k = 30$, $\delta = 0.05$ and $\gamma = 0.4$, and alternative parameters produce similar results.

C Structural Breaks as Nonrecurrent Regime-Switching

Let T be the sample size used in the estimation and let the vector of observations as of time t be denoted $z_{r,t}$. The sequence $\xi_t^P = \{\xi_1^P, ..., \xi_T^P\}$ of regimes in place at each point is unobservable and needs to be inferred jointly with the other parameters of the model. We use the Hamilton filter (Hamilton (1994)) to estimate the smoothed regime probabilities $P\left(\xi_t^P = i|z_{r,T}; \boldsymbol{\theta}_r\right)$, where $i = 1, ..., N_P$. We then use these regime probabilities to estimate the most likely historical regime sequence ξ_t^P over our sample as described in the next subsection.

[3]For trades with prices significantly greater than 100, we repeatedly divide by 10 until prices are in the range of 90 to 100. We exclude all trades otherwise.

The specifications to be estimated are

$$z_{r,t} = r_{\xi_t^P} + \epsilon_t^r, \quad z_{r,t} = \{r_t, mps_t\}$$

where $\epsilon_t^r \sim N(0, \sigma_r^2)$, and $r_{\xi_t^P}$ is a time-varying intercept governed by a discrete valued latent state variable, ξ_t^P, that follows a N_P-state nonrecurrent regime-switching Markov with transition matrix $\mathbf{H}$. Bayesian methods with flat priors are used estimate the parameters $\boldsymbol{\theta}_r = \left(r_{\xi_t^P}, \sigma_r^2, vec(\mathbf{H})'\right)'$ over the period 1961:Q1-2020:Q1 and to estimate the most likely historical regime sequence for ξ_t^P over that sample.

To capture the phenomenon of nonrecurrent regimes, we suppose that ξ_t^P follows a Markov-switching process in which new regimes can arise but do not repeat exactly as before. This is modeled by specifying the transition matrix over nonrecurrent states, or "structural breaks." If the historical sample has N_P nonrecurrent regimes (implying $N_P - 1$ structural breaks), the transition matrix for the Markov process takes the form

$$\mathbf{H} = \begin{bmatrix} p_{11} & 0 & \cdots & \cdots & \cdots & \cdots & 0 \\ 1-p_{11} & p_{22} & 0 & \cdots & \cdots & \cdots & 0 \\ 0 & 1-p_{22} & p_{33} & 0 & \cdots & \cdots & \vdots \\ \vdots & 0 & 1-p_{33} & \ddots & & & \\ \vdots & \vdots & 0 & \vdots & \ddots & & \vdots \\ \vdots & \vdots & \vdots & \vdots & \ddots & p_{N_P,N_P} & 0 \\ 0 & \cdots & \cdots & \cdots & 0 & 1-p_{N_P,N_P} & 1 \end{bmatrix}, \qquad \text{(A.1)}$$

where $\mathbf{H}_{ij} \equiv p\left(\xi_t^P = i | \xi_{t-1}^P = j\right)$. For example, if there were $N_P = 2$ nonrecurrent regimes in the sample, we would have

$$\mathbf{H} = \begin{bmatrix} p_{11} & 0 \\ 1-p_{11} & 1 \end{bmatrix}.$$

The above process implies that, if you are currently in regime 1, you will remain there next period with probability p_{11} or exit to regime 2 with probability $1 - p_{11}$. Upon exiting to regime 2, since there are only two regimes in the sample and the probability p_{12} of returning exactly to the previous regime 1 is zero, $p_{22} = 1$.

D Most Likely Regime Sequence

For regime switches in the mean of mps_t where the specification that is estimated is

$$mps_t = r_{\xi_t^P} + \epsilon_t^r,$$

$\epsilon_t^r \sim N(0, \sigma_r^2)$, and $r_{\xi_t^P}$ is an intercept governed by a discrete valued latent state variable, ξ_t^P, that is presumed to follow a N_P-state nonrecurrent regime-switching Markov with transition matrix $\mathbf{H}$. The vector $\boldsymbol{\theta}_r = \left(r_{\xi_t^P}, \sigma_r^2, vec(\mathbf{H})'\right)'$ denotes the set of parameters to be estimated. The most likely regime sequence is the regime sequence

$\xi^{P,T} = \{\hat{\xi}_1^P, ..., \hat{\xi}_T^P\}$ that is most likely to have occurred, given the estimated posterior mode parameter values for $\boldsymbol{\theta}_r$. This sequence is computed as follows.

Let $P\left(\xi_t^P = i | z_{t-1}; \boldsymbol{\theta}_r\right) \equiv \pi_{t|t-1}^i$. First, run Hamilton's filter to get the vector of filtered regime probabilities $\pi_{t|t}$, $t = 1, 2, ..., T$. The Hamilton filter can be expressed iteratively as

$$\begin{aligned} \pi_{t|t} &= \frac{\pi_{t|t-1} \odot \eta_t}{\mathbf{1}'\left(\pi_{t|t-1} \odot \eta_t\right)} \\ \pi_{t+1|t} &= \mathbf{H}\pi_{t|t} \end{aligned}$$

where the symbol $\odot$ denotes element by element multiplication, η_t is a vector whose *j-th* element contains the conditional density $p(mps_t|\xi_t^P = j; \boldsymbol{\theta}_r)$, i.e.,

$$\eta_{j,t} = \frac{1}{\sqrt{2\pi}\sigma_r} \exp\left\{\frac{-\left(mps_t - r_j\right)^2}{2\sigma_r^2}\right\},$$

and where $\mathbf{1}$ is a vector with all elements equal to 1. The final term, $\pi_{T|T}$ is returned with the final step of the filtering algorithm. Then, a recursive algorithm can be implemented to derive the other smoothed probabilities:

$$\pi_{t|T} = \pi_{t|t} \odot \left[\mathbf{H}'\left(\pi_{t+1|T} \left(\div\right) \pi_{t+1|t}\right)\right]$$

where $(\div)$ denotes element by element division. To choose the regime sequence most likely to have occurred given our parameter estimates, consider the recursion in the next to last period $t = T - 1$:

$$\pi_{T-1|T} = \pi_{T-1|T-1} \odot \left[\mathbf{H}'\left(\pi_{T|T} \left(\div\right) \pi_{T|T-1}\right)\right].$$

Suppose we have $N_p = 3$ regimes. We first take $\pi_{T|T}$ from the Hamilton filter and choose the regime that is associated with the largest probability, i.e., if $\pi_{T|T} = (.8, .1, .1)$, where the first element corresponds to the probability of regime 1, we select $\hat{\xi}_T^P = 1$, indicating that we are in regime 1 in period T. We now update $\pi_{T|T} = (1, 0, 0)$ and plug into the right-hand-side above along with the estimated filtered probabilities for $\pi_{T-1|T-1}$, $\pi_{T|T-1}$ and estimated transition matrix $\mathbf{H}$ to get $\pi_{T-1|T}$ on the left-hand-side. Now we repeat the same procedure by choosing the regime for $T-1$ that has the largest probability at $T - 1$, e.g., if $\pi_{T-1|T} = (.2, .7, .1)$ we select $\hat{\xi}_{T-1} = 2$, indicating that we are in regime 2 in period $T-1$, we then update to $\pi_{T-1|T} = (0, 1, 0)$, which is used again on the right-hand-side now

$$\pi_{T-2|T} = \pi_{T-2|T-2} \odot \left[\mathbf{H}'\left(\pi_{T-1|T} \left(\div\right) \pi_{T-1|T-2}\right)\right].$$

We proceed in this manner until we have a most likely regime sequence $\xi^{P,T}$ for the entire sample $t = 1, 2, ..., T$. Two aspects of this procedure are worth noting. First, it fails if the updated probabilities are exactly $(.333, .333, .333)$. Mathematically this is virtually a effectiveprobability event. Second, note that this procedure allows us to choose the most likely regime sequence by using the recursive formula above to update the filtered probabilities sequentially working backwards from $t = T$ to $t = 1$. This allows us to take into account the time dependence in the regime sequence as dictated by the transition probabilities.

Follow the same procedure to obtain the most likely belief regime sequence ξ_t^b, where the structural model is described by B^2 conditional densities

$$f\left(X_{t-1+\delta_h}|\xi_{t-1}^b = j, \xi_t^b = i, X^{t-1}\right) = (2\pi)^{-N_X/2} |f_{t|t-1+\delta_h}^{(i,j)}|^{-1/2} \exp\left\{-\frac{1}{2} e_{t|t-1+\delta_h,t-1}^{(i,j)\prime} f_{t|t-1+\delta_h,t-1}^{(i,j)} e_{t|t-1+\delta_h,t-1}^{(i,j)}\right\}.$$

Define ξ_t^* describe a B^2-state Markov chain incorporating all the (i, j) combinations above and recast $f(\cdot)$ as B^2 densities $\eta_t = f\left(X_{t-1+\delta_h}|\xi_t^* = i, X^{t-1}\right)$ to use in the computation of $\pi_{t|t}$.

E Price-Output Decompositions

Mapping from price to output (measured as GDP_t) is

$$\begin{aligned}
\frac{P_t}{GDP_{t-1}} &= \frac{P_t}{D_t}\frac{D_t}{GDP_t}\frac{GDP_t}{GDP_{t-1}} \\
pgdp_t &= pd_t + k_t + \widetilde{y}_t + g_t - \widetilde{y}_{t-1}
\end{aligned}$$

Below we decompose pd_t to write:

$$\begin{aligned}
pgdp_t &= \underbrace{\frac{\kappa_{pd,0}}{1-\kappa_{pd,1}} + k_t + y_t + g_t - \widetilde{y}_{t-1}}_{\text{earning share component}} + \underbrace{pdv_t(\Delta d)}_{\text{earnings}} - \underbrace{pdv_t(r^{ex})}_{\text{premia}} - \underbrace{pdv_t(rir)}_{\text{RIR}} \\
pgdp_{r^{ex},t} &= \underbrace{\frac{\kappa_{pd,0}}{1-\kappa_{pd,1}} + k_t + \widetilde{y}_t + g_t - \widetilde{y}_{t-1}}_{\text{earning share component}} - \underbrace{pdv_t(r^{ex})}_{\text{premia}} \\
pgdp_{rir,t} &= \underbrace{\frac{\kappa_{pd,0}}{1-\kappa_{pd,1}} + k_t + \widetilde{y}_t + g_t - \widetilde{y}_{t-1}}_{\text{earning share component}} - \underbrace{pdv_t(rir)}_{\text{RIR}} \\
pgdp_{\Delta d,t} &= \underbrace{\frac{\kappa_{pd,0}}{1-\kappa_{pd,1}} + k_t + \widetilde{y}_t + g_t - \widetilde{y}_{t-1}}_{\text{earning share component}} + \underbrace{pdv_t(\Delta d)}_{\text{earnings}}
\end{aligned}$$

where

$$\begin{aligned}
pd_t &= \kappa_{pd,0} + \mathbb{E}_t^b\left[m_{t+1} + \Delta d_{t+1} + \kappa_{pd,1} pd_{t+1}\right] + \\
&\quad +.5\mathbb{V}_t^b\left[m_{t+1} + \Delta d_{t+1} + \kappa_{pd,1} pd_{t+1}\right].
\end{aligned}$$

The solution approximates around the balanced growth path with $\frac{D_{t+1}}{D_t} = G$, where G is the gross growth rate of the economy. The Euler equation under the balanced growth path is

$$\begin{aligned} 1 &= \left[M_{t+1} \left(\frac{P_{t+1}/D_{t+1}+1}{P_t/D_t} \right) \frac{D_{t+1}}{D_t} \right] \\ &= \left[\beta_p \left(\frac{D_{t+1}}{D_t} \right)^{-\sigma_p} \left(\frac{P_{t+1}/D_{t+1}+1}{P_t/D_t} \right) \frac{D_{t+1}}{D_t} \right] \\ &= \left[\underbrace{\beta_p G^{1-\sigma_p}}_{\widetilde{\beta}_p} \left(\frac{P/D+1}{P/D} \right) \right] => \\ \frac{1}{\widetilde{\beta}_p} &= \left(\frac{P/D+1}{P/D} \right) => \\ P/D &= \frac{\widetilde{\beta}_p}{1-\widetilde{\beta}_p}. \end{aligned}$$

Denote the log steady state price-payout ratio as $\ln(P/D) = \overline{pd}$, thus we have

$$\overline{pd} = \ln \left(\frac{\widetilde{\beta}_p}{1-\widetilde{\beta}_p} \right).$$

$$\begin{aligned} \kappa_{pd,1} &= \exp(\overline{pd})/(1+\exp(\overline{pd})) = \frac{\widetilde{\beta}_p}{1-\widetilde{\beta}_p} \left[1 + \frac{\widetilde{\beta}_p}{1-\widetilde{\beta}_p} \right]^{-1} = \widetilde{\beta}_p \\ \kappa_{pd,0} &= \ln(\exp(\overline{pd})+1) - \kappa_{pd,1} pd = \ln \left(\frac{1}{1-\widetilde{\beta}_p} \right) - \widetilde{\beta}_p \ln \frac{\widetilde{\beta}_p}{1-\widetilde{\beta}_p} \\ &= -\widetilde{\beta}_p \ln \widetilde{\beta}_p - \left(1-\widetilde{\beta}_p\right) \ln \left(1-\widetilde{\beta}_p\right) \end{aligned}$$

The log return obeys the following approximate identity (Campbell and Shiller (1989)):

$$r^D_{t+1} = \kappa_{pd,0} + \kappa_{pd,1} pd_{t+1} - pd_t + \Delta d_{t+1},$$

where $\kappa_{pd,1} = \exp(\overline{pd})/(1+\exp(\overline{pd}))$, and $\kappa_{pd,0} = \log\left(\exp(\overline{pd})+1\right) - \kappa_{pd,1}\overline{pd}$. Combining all of the above, the log equity premium is

$$\underbrace{\mathbb{E}^b_t \left[r^D_{t+1} \right] - \left(i_t - \mathbb{E}^b_t \left[\pi_{t+1} \right] \right)}_{\text{Equity Premium}} = \underbrace{\left[\begin{array}{c} -.5\mathbb{V}^b_t \left[r^D_{t+1} \right] - \mathbb{COV}^b_t \left[m_{t+1}, r^D_{t+1} \right] \\ +.5\mathbb{V}^b_t \left[\pi_{t+1} \right] - \mathbb{COV}^b_t \left[m_{t+1}, \pi_{t+1} \right] \end{array} \right]}_{\text{Risk Premium}} + \underbrace{\overline{lp}_t}_{\text{Liquidity Premium}},$$

Then

$$\begin{aligned} pd_t &= \kappa_{pd,0} + \mathbb{E}^b_t \left[\Delta d_{t+1} - r^D_{t+1} + \kappa_{pd,1} pd_{t+1} \right] \\ pd_t &= \kappa_{pd,0} + \mathbb{E}^b_t \left[\Delta d_{t+1} - \left(r^{ex}_{t+1} - rir_{t+1} \right) + \kappa_{pd,1} pd_{t+1} \right] \end{aligned}$$

where $\mathbb{E}_t^b\left[r_{t+1}^{ex}\right] = \mathbb{E}_t^b\left[r_{t+1}^{D}\right] - rir_{t+1}$, where $rir_{t+1} \equiv \left(i_{t+1} - \mathbb{E}_t^b\left[\pi_{t+1}\right]\right)$.

Solving forward:

$$\begin{aligned} pd_t &= \kappa_{pd,0} + \mathbb{E}_t^b\left[\Delta d_{t+1} - r_{t+1}^{ex} - rir_{t+1}\right] + \\ &\quad + \kappa_{pd,1}\mathbb{E}_t^b\left[\kappa_{pd,0} + \mathbb{E}_t^b\left[\Delta d_{t+2} - r_{t+2}^{ex} - rir_{t+1} + \kappa_{pd,1} pd_{t+2}\right]\right] \end{aligned}$$

Thus:

$$pd_t = \frac{\kappa_{pd,0}}{1-\kappa_{pd,1}} + \left(1_{\Delta d} - 1_{\mathbb{E}(r^{ex})} - 1_{rir}\right)\sum_{h=0}^{\infty}\kappa_{pd,1}^h \mathbb{E}_t^b\left[S_{t+1+h}\right]$$

where 1_x is a vector of all zeros except for a 1 in the xth position. This can be written as:

$$pd_t = \frac{\kappa_{pd,0}}{1-\kappa_{pd,1}} + pdv_t\left(\Delta d\right) - pdv_t\left(r^{ex}\right) - pdv_t\left(rir\right)$$

Using the solution:

$$pd_t = \frac{\kappa_{pd,0}}{1-\kappa_{pd,1}} + \left(1_{\Delta d} - 1_{\mathbb{E}(r^{ex})} - 1_{r^b}\right)\left(\mathbf{I} - \kappa_{pd,1}T_{\xi_t}\right)^{-1}\left[T_{\xi_t}S_t + \left(\mathbf{I} - \kappa_{pd,1}\right)^{-1}C_{\xi_t}\right].$$

Thus, we can decompose movements in the pd_t into those attributable to expected dividends, equity premia, and expected real interest rates:

$$pgdp_t = \underbrace{\frac{\kappa_{pd,0}}{1-\kappa_{pd,1}} + k_t + y_t + g_t - y_{t-1}}_{\text{earning share component}} + \underbrace{pdv_t\left(\Delta d\right)}_{\text{earnings}} - \underbrace{pdv_t\left(r^{ex}\right)}_{\text{premia}} - \underbrace{pdv_t\left(rir\right)}_{\text{RIR}}.$$

F Solution and Estimation Details

This appendix presents details on the solution and estimation. An overview of the steps are as follows.

1. We first solve the macro block set of equations involving a set of macro state variables $S_t^M \equiv [\widetilde{y}_t, g_t, \pi_t, i_t, \overline{\pi}_t, f_t]'$. The MS-VAR solution consists of a system of equations taking the form

$$S_t^M = C_M\left(\theta_{\xi_t^P}\right) + T_M(\theta_{\xi_t^P})S_{t-1}^M + R_M(\theta_{\xi_t^P})Q_M\varepsilon_t^M,$$

   where $\varepsilon_t^M = (\varepsilon_{f,t}, \varepsilon_{i,t}, \varepsilon_{g,t}, \varepsilon_{\mu,t})$. Since this block involves no forward-looking variables and only depends on the pre-determined policy regimes, this block can be solved analytically. See Bianchi, Lettau, and Ludvigson (2022).

2. Use the solution for S_t^M based on the current realized policy regime ξ_t^P and then resolve the model based on the Alternative regime, i.e., obtain

$$S_t^M = C_M\left(\theta_{\xi_t^A}\right) + T_M(\theta_{\xi_t^A})S_{t-1}^M + R_M(\theta_{\xi_t^A})Q_M\varepsilon_t^M.$$

   Store the two solutions. S_t^M under ξ_t^P is mapped into the observed current macro variables in our observation equation.

3. To identify the parameters of the Alternative policy rule, the perceived transition matrix $\mathbf{H}^b$ and belief regime probabilities governing moving to the Alternative rule, we use:

   (a) Measures of expectations from professional forecast surveys and futures markets. Given the perceived transition matrix of the investor $\mathbf{H}^b$, use it to compute investor expectations for future macro variables that take into account the perceived probability of transitioning to the Alternative rule in the future. See the section below on "Computing Expectations with Regime Switching and Alternative Policy Rule." These give us investor expectations of the macro block variables used in our observation equation.

   (b) Stock prices. The asset pricing block of equations involves conditional subjective variance terms that are affected by Markov-switching random variables in the model. The subsection "Risk Adjustment with Lognormal Approximation," below explains the approximation used to preserve lognormality of the entire system. This part uses the approach in Bianchi, Kung, and Tirskikh (2018) who in turn build on Bansal and Zhou (2002) and is combined with the algorithm of Farmer, Waggoner, and Zha (2011) to solve the overall system of model equations, where investors form expectations taking into account the probability of regime change in the future. The state variables for the full system are

$$S_t = \left[S_t^M, m_t, pd_t, k_t, lp_t, \mathbb{E}_t^b\left(m_{t+1}\right), \mathbb{E}_t^b\left(pd_{t+1}\right)\right].$$

   This leaves us with the MS-VAR solution consists of a system of equations taking the form

$$S_t = C\left(\theta_{\xi_t^P}, \xi_t^b, \mathbf{H}^b\right) + T(\theta_{\xi_t^P}, \xi_t^b, \mathbf{H}^b)S_{t-1} + R(\theta_{\xi_t^P}, \xi_t^b, \mathbf{H}^b)Q\varepsilon_t,$$

   where $\varepsilon_t = (\varepsilon_{f,t}, \varepsilon_{i,t}, \varepsilon_{g,t}, \varepsilon_{\mu,t}, \varepsilon_{k,t}, \varepsilon_{lp,t})$. Since pd_t depends the risk adjustment and $\mathbb{E}_t^b\left(pd_{t+1}\right)$, its value is also informative about the parameters of the Alternative rule, $\mathbf{H}^b$ and belief regime probabilities. Unlike the formulas that are required to relate data on expectations to future macro variables in step (a), the formulas governing these relationships are solved numerically using the solution algorithm described above.

4. We estimate the model by combining the solution above with an observation equation that includes macro, asset pricing, and survey expectation variables. See the subsection "Estimation" below.

G Computing Expectations with Regime Switching and Alternative Policy Rule

In what follows, we explain how to use expectations to infer what alternative regimes agents have in mind. Expectations about inflation, FFR, and GDP growth depend on the regime currently in place, the alternative regime, and the probability of moving to such regime. This note is based on "Methods for measuring expectations and uncertainty" in Bianchi (2016). That paper explains how to computed expected values in presence of regime changes. In the models described above, for each policy rule in place, agents would have different beliefs about alternative future policy rules. This would lead to changes in expected values for the endogenous variables of the model.

Consider a MS model:

$$S_t = C_{\xi_t} + T_{\xi_t} S_{t-1} + R_{\xi_t} Q \varepsilon_t \tag{A.2}$$

where $\xi_t = \left\{\xi_t^P, \xi_t^b\right\}$ controls the policy regime ξ_t^P controls the policy rule currently in place and the alternative policy rule, while the belief regime ξ_t^b controls agents' beliefs about the possibility of moving to the alternative policy rule.

Let n be the number of variables in S_t. Let $m = B + 1$ be the number of Markov-switching states and define

$$\xi_t = i \equiv \left\{\xi_t^P, \xi_t^b = i\right\}, \quad i = 1, ..., B + 1.$$

Define the $mn \times 1$ column vector q_t as:

$$\underset{mn \times 1}{q_t} = \left[q_t^{1\prime}, ..., q_t^{m\prime}\right]'$$

where the individual $n \times 1$ vectors $q_t^i = \mathbb{E}_0\left(S_t 1_{\xi_t = i}\right) \equiv \mathbb{E}\left(S_t 1_{\xi_t = i} | \mathbb{I}_0\right)$ and $1_{\xi_t = i}$ is an indicator variable that is one when belief regime i is in place and effectiveotherwise. Note that:

$$q_t^i = \mathbb{E}_0\left(S_t 1_{\xi_t = i}\right) = \mathbb{E}_0\left(S_t | \xi_t = i\right)\pi_t^i$$

where $\pi_t^i = P_0\left(\xi_t = i\right) = P\left(\xi_t = i | \mathbb{I}_0\right)$. Therefore we can express $\mu_t = \mathbb{E}_0\left(S_t\right)$ as:

$$\mu_t = \mathbb{E}_0\left(S_t\right) = \textstyle\sum_{i=1}^m q_t^i = w q_t$$

where the matrix $\underset{n \times mn}{w} = [I_n, ..., I_n]$ is obtained placing side by side m n-dimensional identity matrices. Then the following proposition holds:

PROPOSITION 1: *Consider a Markov-switching model whose law of motion can be described by (A.2) and define* $q_t^i = \mathbb{E}_0\left(S_t 1_{\xi_t = i}\right)$ *for* $i = 1...m$. *Then* $q_t^j = C_j \pi_t^j + \sum_{i=1}^m T_j q_{t-1}^i p_{ji}$.

It is then straightforward to compute expectations conditional on the information available at a particular point in time. Suppose we are interested in $\mu_{t+s|t} \equiv \mathbb{E}_t^b\left(S_{t+s}\right)$, i.e. the expected value for the vector S_{t+s} conditional on the information set available at time t. If we define:

$$q_{t+s|t} = \left[q_{t+s|t}^{1\prime}, ..., q_{t+s|t}^{m\prime}\right]'$$

where $q_{t+s|t}^{i} = \mathbb{E}_t^b \left(S_{t+s}1_{\xi_t=i}\right) = \mathbb{E}_t^b \left(S_{t+s}|\xi_t = i\right) \pi_{t+s|t}^{i}$, where $\pi_{t+s|t}^{i} \equiv P\left(\xi_{t+s} = i|\mathbb{I}_t\right)$, we have

$$\mu_{t+s|t} = \mathbb{E}_t^b \left(S_{t+s}\right) = wq_{t+s|t}, \tag{A.3}$$

where for $s \geq 1$, $q_{t+s|t}$ evolves as:

$$q_{t+s|t} = C\pi_{t+s|t} + \Omega q_{t+s-1|t} \tag{A.4}$$

$$\pi_{t+s|t} = \mathbf{H}^b \pi_{t+s-1|t} \tag{A.5}$$

with $\pi_{t+s|t} = \left[\pi_{t+s|t}^1, ..., \pi_{t+s|t}^m\right]'$, $\Omega = bdiag\left(T_1, ..., T_m\right)\left(\mathbf{H}^b \otimes I_n\right)$, and $\underset{mn\times m}{C} = bdiag\left(C_1, ..., C_m\right)$, where e.g., C_1 is the $n \times 1$ vector of constants in regime 1, $\otimes$ represents the Kronecker product and *bdiag* is a matrix operator that takes a sequence of matrices and use them to construct a block diagonal matrix.

The formulas above are used to compute expectations conditional on each belief regime ξ_t^b and policy rule regime ξ_t^P. For each composite regime $\xi_t = \left\{\xi_t^P, \xi_t^b\right\}$, we can obtain a forecast for each of the variables of the model. For example, conditional on ξ_t^P and $\xi_t^b = j$ in place we have

$$q_{t,\xi_t=j} = e_j \otimes S_t$$

where e_j is a variable that has elements equal to effectiveexcept for the one in position ξ_t^b. For example, with $B = 5$ belief regimes and $\xi_t^b = 3$ we have

$$q_{t,\xi_t=3} = \left[\mathbf{0}', \mathbf{0}', S_t', \mathbf{0}', \mathbf{0}', \mathbf{0}'\right]'.$$

where 0 and S_t are column vectors with n rows. We have $B + 1$ subvectors in $q_{t,\xi_t=j}$ to take into account the alternative policy mix. The fact that all subvectors are effectiveexcept for the one corresponding to the belief regime $b = 3$ reflects the assumption that agents can observe the current state S_t and, by definition, their own beliefs (while the econometrician cannot observe any of the two and she uses macro data and survey expectations to estimate both S_t and agents' beliefs).

Thus, suppose we want to compute the expected value for a variable x over the next year under the assumption that agents' beliefs are $\xi_t^b = j$. With monthly data, we have:

$$\begin{aligned}
\mathbb{E}_t^b \left(x_{t,t+s}|\xi_t = j\right) &= \textstyle\sum_{s=1}^{12} \mathbb{E}_t^b \left(x_{t+s}|\xi_t = j\right) \\
&= e_x \textstyle\sum_{s=1}^{12} \mu_{t+s|t,\xi_t=j} \\
&= e_x w \textstyle\sum_{s=1}^{12} q_{t+s|t,\xi_t=j}
\end{aligned}$$

where for $s \geq 1$, $q_{t+s|t}$ evolves as:

$$q_{t+s|t,\xi_t=j} = C\pi_{t+s|t} + \Omega q_{t+s-1|t,\xi_t=j} \tag{A.6}$$

$$\pi_{t+s|t,\xi_t=j} = \mathbf{H}^b \pi_{t+s-1|t,\xi_t=j} \tag{A.7}$$

with $\pi_{t+s|t} = \left[\pi^1_{t+s|t}, ..., \pi^m_{t+s|t}\right]'$, $\Omega = bdiag\left(T_1, ..., T_m\right)\left(\mathbf{H}^b \otimes I_n\right)$, and $\underset{mn\times m}{C} = bdiag\left(C_1, ..., C_m\right)$, where e.g., C_1 is the $n \times 1$ vector of constants in regime 1, $\otimes$ represents the Kronecker product and $bdiag$ is a matrix operator that takes a sequence of matrices and use them to construct a block diagonal matrix. The recursive algorithm is initialized with $\pi_{t|t,\xi_t=j} = 1_{\xi_t=j}$ and $q_{t,\xi_t=j} = e_j \otimes S_t$.

The formulas (A.6) and (A.7) can be written in a more compact form. If we define $\widetilde{q}_{t|t} = [q'_{t|t}, \pi'_{t|t}]'$, with $\pi_{t|t}$ a vector with elements $\pi^i_{t|t} \equiv P\left(\xi_t = i|\mathbb{I}_t\right)$ we can compute the conditional expectations in one step:

$$\mu_{t+s|t} = \mathbb{E}^b_t\left(S_{t+s}\right) = \widetilde{w}\widetilde{\Omega}^s\widetilde{q}_{t|t} \tag{A.8}$$

where $\widetilde{w} = \left[w, 0_{n\times m}\right]$. The formula above can be used to compute the expected value from the point of view of the agent of the model with beliefs $\xi_t = j$:

$$\mathbb{E}^b_t\left(x_{t+s}|\xi_t = j\right) = e_x\mu_{t+s|t,\xi_t=j} = e_x\widetilde{w}\widetilde{\Omega}^s\widetilde{q}_{t|t,\xi^b_t=j} = \underbrace{e_x w\widetilde{\Omega}^s_{\{1,nm\},\{n(j-1)+1,nj\}}}_{Z_{\xi_t,x_{t+s}}}\underbrace{S_t}_{(n\times1)} + \underbrace{e_x w\widetilde{\Omega}^s_{\{1,nm\},nm+j}}_{D_{\xi_t,x_{t+s}}} \tag{A.9}$$

where $D_{\xi_t,x_{t+s}}$ is a scalar, $Z_{\xi_t,x_{t+s}}$ is an $(1 \times n)$ vector, $\widetilde{\Omega}^s_{\{1,nm\},\{n(j-1)+1,nj\}}$ is the submatrix obtained taking the first nm rows and the columns from $n(j-1)+1$ to nj of $\widetilde{\Omega}^s$, while $\widetilde{\Omega}^s_{\{1,nm\},nm+j}$ is the submatrix obtained taking the first nm rows and the $nm+j$ column of $\widetilde{\Omega}^s$. Thus, we have that conditional on one belief regime and a policy rule regime, we can map the current state of the economy S_t into the expected value reported in the survey. The matrix algebra in (A.9) returns the same results of the recursion in (A.6) and (A.7).

To see what the formulas above do, consider a simple example with $B = 2$ and we are currently in belief regime $b = 2$:

$$\begin{aligned}
\mathbb{E}_t^b\left(x_{t+s}|\xi_t=2\right) &= e_x\widetilde{w}\widetilde{\Omega}^s\widetilde{q}_{t|t,\xi_t=2} = e_x\widetilde{w}\widetilde{\Omega}^s\begin{bmatrix}\underset{n\times1}{\mathbf{0}}\\ \underset{n\times1}{S_t}\\ \underset{n\times1}{\mathbf{0}}\\ 0\\ 1\\ 0\end{bmatrix}\\
&= e_x\widetilde{w}\begin{bmatrix}\widetilde{\Omega}_{11}^s & \widetilde{\Omega}_{12}^s & \widetilde{\Omega}_{13}^s & \widetilde{\Omega}_{14}^s & \widetilde{\Omega}_{15}^s & \widetilde{\Omega}_{16}^s\\ \widetilde{\Omega}_{21}^s & \widetilde{\Omega}_{22}^s & \widetilde{\Omega}_{23}^s & \widetilde{\Omega}_{24}^s & \widetilde{\Omega}_{25}^s & \widetilde{\Omega}_{26}^s\\ \widetilde{\Omega}_{31}^s & \widetilde{\Omega}_{32}^s & \widetilde{\Omega}_{33}^s & \widetilde{\Omega}_{34}^s & \widetilde{\Omega}_{35}^s & \widetilde{\Omega}_{36}^s\\ & & & \widetilde{\Omega}_{44}^s & \widetilde{\Omega}_{45}^s & \widetilde{\Omega}_{46}^s\\ & & & \widetilde{\Omega}_{54}^s & \widetilde{\Omega}_{55}^s & \widetilde{\Omega}_{56}^s\\ & & & \widetilde{\Omega}_{64}^s & \widetilde{\Omega}_{65}^s & \widetilde{\Omega}_{66}^s\end{bmatrix}\begin{bmatrix}\underset{n\times1}{\mathbf{0}}\\ \underset{n\times1}{S_t}\\ \underset{n\times1}{\mathbf{0}}\\ 0\\ 1\\ 0\end{bmatrix}\\
&= e_x\widetilde{w}\begin{bmatrix}\widetilde{\Omega}_{12}^sS_t+\widetilde{\Omega}_{15}^s\\ \widetilde{\Omega}_{22}^sS_t+\widetilde{\Omega}_{25}^s\\ \widetilde{\Omega}_{32}^sS_t+\widetilde{\Omega}_{35}^s\\ \widetilde{\Omega}_{44}^s\\ \widetilde{\Omega}_{54}^s\\ \widetilde{\Omega}_{64}^s\end{bmatrix}\\
&= e_x\left(\widetilde{\Omega}_{12}^s+\widetilde{\Omega}_{22}^s+\widetilde{\Omega}_{32}^s\right)S_t+e_x\left(\widetilde{\Omega}_{15}^s+\widetilde{\Omega}_{25}^s+\widetilde{\Omega}_{35}^s\right)
\end{aligned}$$

Finally, suppose we are interested in the forecast $E_t^b\left(x_{t,t+s}|\xi_t^b=j,\xi_t^p\right)$:

$$\mathbb{E}_t^b\left(x_{t,t+s}|\xi_t=j\right)=\underbrace{\left[e_x\sum_{s=1}^{12}w\widetilde{\Omega}^s_{\{1,nm\},\{n(j-1)+1,nj\}}\right]}_{Z_{\xi_t,x_{t,t+s}}}\underbrace{S_t}_{(n\times1)}+\underbrace{e_x\sum_{s=1}^{12}w\widetilde{\Omega}^s_{\{1,nm\},nm+j}}_{D_{\xi_t,x_{t,t+s}}} \tag{A.10}$$

Thus, we can include $Z_{\xi_t,x_{t,t+s}}$ as a row in Z_{ξ_t} and $D_{\xi_t,x_{t,t+s}}$ as a row in D_{ξ_t} in the mapping from the model to the observables described in (A.11). Note that the matrix Z and vector D are now regime dependent.

For GDP growth, we are interested in the average growth over a certain horizon. Our state vector contains $\widetilde{y}_t$. Thus, we can use the following approach:

$$\begin{aligned}
\mathbb{E}_t^b\left[\left(gdp_{t+h}-gdp_t\right)h^{-1}|\xi_t=j\right] &= \mathbb{E}_t^b\left[\left(\widetilde{y}_{t+h}-\widetilde{y}_t+\sum_{s=1}^h\widehat{g}_t+hg\right)h^{-1}|\xi_t=j\right]\\
&= h^{-1}\mathbb{E}_t^b\left[\widetilde{y}_{t+h}|\xi_t=j\right]-h^{-1}\widetilde{y}_t+g+h^{-1}\sum_{s=1}^h\widehat{g}_t
\end{aligned}$$

where g is the average growth rate in the economy and $\widetilde{y}_t$ is GDP in deviations from the trend. With deterministic growth we have $gdp_{t+h} - gdp_t - hg \equiv \widetilde{y}_{t+h} - \widetilde{y}_t$. We then have

$$\begin{aligned}
& \mathbb{E}_t^b \left[\left(gdp_{t+h} - gdp_t \right) h^{-1} | \xi_t = j \right] \\
= \; & h^{-1} \mathbb{E}_t^b \left[\widetilde{y}_{t+h} | \xi_t = j \right] - h^{-1} \widetilde{y}_t + g \\
= \; & h^{-1} \left[\underbrace{e_{\widetilde{y}} w \widetilde{\Omega}^s_{\{1,nm\},\{n(j-1)+1,nj\}}}_{Z_{\xi_t,\widetilde{y}_{t+s}}} \underbrace{S_t}_{(n\times1)} + \underbrace{e_{\widetilde{y}} w \widetilde{\Omega}^s_{\{1,nm\},nm+j}}_{D_{\xi_t,\widetilde{y}_{t+s}}} - e_{\widetilde{y}} S_t \right] + g \\
= \; & h^{-1} \underbrace{\left[e_{\widetilde{y}} w \widetilde{\Omega}^s_{\{1,nm\},\{n(j-1)+1,nj\}} - e_{\widetilde{y}} \right]}_{Z_{\xi_t,\widetilde{y}_{t+s}-\widetilde{y}_t}} \underbrace{S_t}_{(n\times1)} + h^{-1} \underbrace{e_{\widetilde{y}} w \widetilde{\Omega}^s_{\{1,nm\},nm+j}}_{D_{\xi_t,\widetilde{y}_{t+s}}} + g
\end{aligned}$$

The expected values for the endogenous variables depend on the perceived transition matrix H^b and the properties of the alternative regime. The latter can be seen by recalling that the regime $\xi_t = B + 1$ applies to the perceived Alternative regime. Thus, data from survey expectations and futures markets provide information about the perceived probability of moving across belief regimes as well as the parameters of the Alternative regime.

H Estimation and Filtering

The solution of the model takes the form of a Markov-switching vector autoregression (MS-VAR) in the state vector $S_t = \left[S_t^M, m_t, pd_t, k_t, lp_t, \mathbb{E}_t^b \left(m_{t+1} \right), \mathbb{E}_t^b \left(pd_{t+1} \right) \right]$. Here, S_t^M is a vector of macro block state variables given by $S_t^M \equiv [\widetilde{y}_t, g_t, \pi_t, i_t, \overline{\pi}_t, f_t]'$. The asset pricing block of equations involves conditional subjective variance terms that are affected by Markov-switching random variables in the model. The subsection "Risk Adjustment with Lognormal Approximation," below, explains the approximation used to preserve lognormality of the entire system.

The model solution in state space form is

$$\begin{aligned}
X_t &= D_{\xi_t,t} + Z_{\xi_t,t} \left[S_t', \widetilde{y}_{t-1} \right]' + U_t v_t \\
S_t &= C\left(\theta_{\xi_t^P}, \xi_t^b, \mathbf{H}^b \right) + T(\theta_{\xi_t^P}, \xi_t^b, \mathbf{H}^b) S_{t-1} + R(\theta_{\xi_t^P}, \xi_t^b, \mathbf{H}^b) Q \varepsilon_t \\
Q &= diag \left(\sigma_{\varepsilon_1}, ..., \sigma_{\varepsilon_G} \right), \; \varepsilon_t \sim N(0, I) \\
U &= diag \left(\sigma_1, ..., \sigma_X \right), \; v_t \sim N(0, I) \\
\xi_t^P &= 1...N_P, \; \xi_t^b = 1, ...B+1, \mathbf{H}^b_{ij} = p \left(\xi_t^b = i | \xi_{t-1}^b = j \right).
\end{aligned}$$

where X_t is a $N_X \times 1$ vector of data, v_t are a vector of observation errors, U_t is a diagonal matrix with the standard deviations of the observation errors on the main diagonal, and $D_{\xi_t,t}$, and $Z_{\xi_t,t}$ are parameters mapping the model counterparts of X_t into the latent discrete- and continuous-valued state variables ξ_t and S_t, respectively, in the model.

The vector X_t of observables is explained below. Note that the parameters $D_{\xi_t,t}$, $Z_{\xi_t,t}$, and U_t vary with t independently of ξ_t because not all variables are observed at each data sampling period. To reduce computation time, we calibrate rather than estimate the parameters in $U = diag\left(\sigma_1, ..., \sigma_X\right)$ such that the variance of the observation error is 0.05 times the sample variance of the corresponding variable in X. In addition, some of the parameters in the system are dependent on the current policy rule and the associated Alternative rule, ξ_t^P, and the unobserved, discrete-valued $(B+1)$-state Markov-switching variable ξ_t^b $\left(\xi_t^b = 1, 2, ..., B+1\right)$ with perceived transition probabilities

$$\mathbf{H}^b = \begin{bmatrix} p_{b1}p_s & p_{b2}p_{\Delta 1|2} & \cdots & p_{bB}p_{\Delta 1|B} & 0 \\ p_{b1}p_{\Delta 2|1} & p_{b2}p_s & & p_{bB}p_{\Delta 2|B} & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ p_{b1}p_{\Delta B|1} & & & p_{bB}p_s & 0 \\ 1-p_{b1} & 1-p_{b2} & \cdots & 1-p_{bB} & p_{B+1,B+1}=1 \end{bmatrix},$$

where $H_{ij}^b \equiv p\left(\xi_t^b = i|\xi_{t-1}^b = j\right)$, and $\sum_{i\neq j} p_{\Delta i|j} = 1 - p_s$. We take the parameters p_{bi} from a discretized estimated beta distribution, where the mean and variance of the beta distribution are estimated. We specify the probability of transitioning to belief i tomorrow, conditional on having belief j today, while remaining in the same policy regime, as $p_{\Delta i|j} \equiv (1-p_s)\left(\rho_b^{|i-j-1|}/\sum_{i\neq j}\rho_b^{|i-j-1|}\right)$, where p_s and $\rho_b < 1$ are parameters to be estimated and $|i-j-1|$ measures the distance between beliefs j and i, for $i \neq j \in (1, 2, ..., B)$. This creates a decaying function that makes the probability of moving to contiguous beliefs more likely than jumping to very different beliefs. For computational reasons, we also eliminate very unlikely transitions ($p_{\Delta i|j} < 0.0001$) by setting their probabilities to zero.

We use the following notation:

$$\begin{aligned} C_{\xi_t^P,i} &= C\left(\theta_{\xi_t^P}, \xi_t^b = i\right),\ T_{\xi_t^P,j} = T\left(\theta_{\xi_t^P}, \xi_t^b = i\right),\ R_{\xi_t^P,j} = R\left(\theta_{\xi_t^P}, \xi_t^b = i\right) \\ D_{i,t} &= D_{\xi_t|\xi_t^b=i}, \quad Z_{i,t} = Z_{\xi_t|\xi_t^b=i}. \end{aligned}$$

Kim's Approximation to the Likelihood and Filtering We use Kim's (Kim (1994)) basic filter and approximation to the likelihood.

First note that, from the econometricians viewpoint, investors are only ever observed in the first B regimes, since the perceived Alternative is never actually realized. For this reason the filtering algorithm for the latent belief regimes involves only the upper $B \times B$ submatrix of H^b, rescaled so that the elements sum to unity. Even though the filtering loops over just B states rather than $B+1$, this is done conditional on the parameters for the full $(B+1) \times (B+1)$ transition matrix, which is estimated from all the data by combining the likelihood with the priors, as described below.

The sample is divided into N_P policy regime subperiods indexed by ξ_t^P. Denote the last observation of each regime subperiod of the sample $T_1, ..., T_{N_P}$. The algorithm for the basic filter is described as follows.

Initiate values $\widetilde{S}_{0|0}$, $P_{0|0}$, for the Kalman filter and $\Pr\left(\xi_0^b\right) = \pi_0$ for the Hamilton filter and initialize $L(\theta) = 0$. Denote $X^{t-1} \equiv \{X_1, ..., X_{t-1}\}$ and $\xi^{PT} = \left\{\xi_1^P, ..., \xi_T^P\right\}$.

As explained for the simplified example above, in the mixed-frequency estimation we use high frequency, forward-looking intramonth data to infer updates to investor nowcasts of the state space that will be revealed at the end of the month. Our "final" estimates of the state space are obtained using a more complete set of data available at the end of each month. Let t denote a month. Let d_h denote the number of time units that have passed within a month when we have reached a particular point in time, and let nd denote the total number of time units in the month. Then $0 \leq d_h/nd \leq 1$, and the intramonth time period is denoted $t-1+\delta_h$ with $\delta_h \equiv d_h/nd \in [0,1]$. For example, δ_{100} could denote the point within the month that is exactly 10 minutes before an FOMC meeting during the month, while δ_{130} could denote point in the month 20 minutes after the same meeting. Intra-month observations used just prior to an FOMC meeting will typically include the daily BBG consensus forecasts and Baa credit spread from the day before the meeting, and the 10-minutes before FFF, ED and stock market data. Intermonth observations for the point of the month right after the FOMC meeting will typically include the daily BBG consensus forecasts and Baa spread from the day after the meeting, and the 20-minutes after FFF, ED and stock market data.

Suppose we are within month t at day $t-1+\delta_h$, where $\delta_h < 1$. Investors are presumed to have a belief or *nowcast* of what S will be at the end of t. We assume that such nowcasts are formed from extensive information sets that are by definition unobserved by the econometrician. The filtering algorithm below is designed to allow the econometrician to infer investor nowcasts at any time $t-1+\delta_h$ by using the structural model combined with high frequency forward-looking data, without taking a stand on the investor's unobservable information sets and nowcasting model. We use the suffix $(t\backslash t-1+\delta_h)$ to denote filtered objects related to investor nowcasts which are implicitly based on the agent's latent information set, i.e., $S_{(t\backslash t-1+\delta_h)}$ refers to the investor's nowcast of S_t based on information through $t-1+\delta_h$. Given that investors observe S_t at the end of t, $S_{t\backslash t} \equiv S_t$. We use the symbol "$|$" to refer to conditioning in the filter that is with respect to the econometrician's information set.

- For $t = 1$ to T_1 and $\theta_{\xi_t^P}$ relevant when $\xi_t^P = 1$:

1. Suppose the econometrician has information up through month $t-1$ and new high frequency data arrives at $t-1+\delta_h$. Conditional on $\xi_{t-1}^b = j$ and $\xi_t^b = i$ run the Kalman filter given below for $i, j = 1, 2, ..., B$ to update estimates of the latent

state:

$$\begin{aligned}
S^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1} &= C_{\xi^P_t,i} + T_{\xi^P_t,i} S^j_{t-1|t-1} \\
P^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1} &= T_{\xi^P_t,i} P^j_{t-1|t-1} T'_{\xi^P_t,i} + R_{\xi^P_t,i} Q^2 R'_{\xi^P_t,i} \text{ with } Q^2 \equiv QQ' \\
e^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1+\delta_h,t-1} &= X_{t-1+\delta_h} - D_{i,t-1+\delta_h} - Z_{i,t-1+\delta_h}\left[S^{(i,j)\prime}_{(t\backslash t-1+\delta_h)|t-1}, \widetilde{y}_{t-1}\right] \\
f^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1+\delta_h,t-1} &= Z_{i,t-1+\delta_h} P^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1} Z'_{i,t-1+\delta_h} + U^2_{t-1+\delta_h} \quad \text{with } U^2_{t-1+\delta_h} \equiv U_{t-1+\delta_h} U'_{t-1+\delta_h} \\
S^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1+\delta_h} &= S^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1} + P^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1} Z'_{i,t-1+\delta_h}\left(f^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1+\delta_h,t-1}\right)^{-1} e^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1+\delta_h,t-1} \\
P^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1+\delta_h} &= P^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1} - P^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1} Z'_{i,t-1+\delta_h}\left(f^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1+\delta_h,t-1}\right)^{-1} Z_{i,t-1+\delta_h} P^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1}
\end{aligned}$$

2. Run the Hamilton filter to calculate new regime probabilities $\Pr\left(\xi^b_t, \xi^b_{t-1}|X_{t-1+\delta_h}, X^{t-1}\right)$ and $\Pr\left(\xi^b_t|X_{t-1+\delta_h}, X^{t-1}\right)$, for $i, j = 1, 2, ..., B$

$$\begin{aligned}
\Pr\left(\xi^b_t, \xi^b_{t-1}|X^{t-1}\right) &= \Pr\left(\xi^b_t|\xi^b_{t-1}\right)\Pr\left(\xi^b_{t-1}|X^{t-1}\right) \\
\ell\left(X_{t-1+\delta_h}|X^{t-1}\right) &= \textstyle\sum_{j=1}^B \sum_{i=1}^B f\left(X_{t-1+\delta_h}|\xi^b_{t-1} = j, \xi^b_t = i, X^{t-1}\right) \\
&\quad \Pr\left[\xi^b_{t-1} = j, \xi^b_t = i|X^{t-1}\right] \\
f\left(X_{t-1+\delta_h}|\xi^b_{t-1} = j, \xi^b_t = i, X^{t-1}\right) &= (2\pi)^{-N_X/2}|f^{(i,j)}_{t|t-1+\delta_h}|^{-1/2} \\
&\quad \exp\left\{-\frac{1}{2} e^{(i,j)\prime}_{(t\backslash t-1+\delta_h)|t-1+\delta_h,t-1} f^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1+\delta_h,t-1} e^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1+\delta_h,t-1}\right\} \\
\mathcal{L}(\theta) &= \mathcal{L}(\theta) + \ln\left(\ell\left(X_{t-1+\delta_h}|X^{t-1}\right)\right) \\
\Pr\left(\xi^b_t, \xi^b_{t-1}|X_{t-1+\delta_h}, X^{t-1}\right) &= \frac{f\left(X_{t-1+\delta_h}|\xi^b_t, \xi^b_{t-1}, X^{t-1}\right)\Pr\left(\xi^b_t, \xi^b_{t-1}|X^{t-1}\right)}{\ell\left(X_{t-1+\delta_h}|X^{t-1}\right)} \\
\Pr\left(\xi^b_t|X_{t-1+\delta_h}, X^{t-1}\right) &= \sum_{j=1}^B \Pr\left(\xi^b_t, \xi^b_{t-1} = j|X_{t-1+\delta_h}, X^{t-1}\right)
\end{aligned}$$

3. Using $\Pr\left(\xi^b_t, \xi^b_{t-1}|X_{t-1+\delta_h}, X^{t-1}\right)$ and $\Pr\left(\xi^b_t|X_{t-1+\delta_h}, X^{t-1}\right)$, collapse the $B \times B$ values of $S^{(i,j)}_{t|t-1+\delta_h}$ and $P^{(i,j)}_{t|t-1+\delta_h}$ into B values represented by $S^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$ and $P^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$:

$$S^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h} = \frac{\sum_{j=1}^B \Pr\left[\xi^b_{t-1} = j, \xi^b_t = i|X_{t-1+\delta_h}, X^{t-1}\right] S^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1+\delta_h}}{\Pr\left[\xi^b_t = i|X_{t-1+\delta_h}, X^{t-1}\right]}$$

$$P^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h} = \frac{\sum_{j=1}^B \Pr\left[\xi^b_{t-1} = j, \xi^b_t = i|X_{t-1+\delta_h}, X^{t-1}\right]\left(\begin{array}{c} P^{(i,j)}_{t|t-1+\delta_h} + \left(S^i_{(t\backslash t-1+\delta_h)t-1+\delta_h} - S^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1+\delta_h}\right) \\ \left(S^i_{(t\backslash t-1+\delta_h)|t-1+\delta_h} - S^{(i,j)}_{(t\backslash t-1+\delta_h)|t-1+\delta_h}\right)' \end{array}\right)}{\Pr\left[\xi^b_t = i|X_{t-1+\delta_h}, X^{t-1}\right]}$$

4. If $t - 1 + \delta_h = t$, move to the next period by setting $t - 1 = t$ and returning to step 1

5. else, store the updated $S^{(j)}_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$, $P^{(j)}_{(t\backslash t-1+\delta_h)|t-1+\delta_h}$, $\Pr\left(\xi^b_t, \xi^b_{t-1}|X_{t-1+\delta_h}, X^{t-1}\right)$, and $\Pr\left(\xi^b_t|X_{t-1+\delta_h}, X^{t-1}\right)$, move to the next intramonth time unit $\delta_k > \delta_h$, and repeat steps 1-5 keeping $t - 1$ fixed.

- At $t = T_1 + 1$ use $\theta_{\xi_t^P}$ relevant when $\xi_t^P = 2$, set $t - 1 = t$, and repeat steps 1-5
- At $t = T_2 + 1$ use $\theta_{\xi_t^P}$ relevant when $\xi_t^P = 3$, set $t - 1 = t$, and repeat steps 1-5
- $\vdots$
- At $t = T_{N_P - 1} + 1$ use $\theta_{\xi_t^P}$ relevant when $\xi_t^P = N_P$, set $t - 1 = t$ and repeat steps 1-5
- At $t = T_N = T$ stop. Obtain $L(\theta) = \sum_{t=1}^{T} \sum_{\delta_h \in (0,1)} \ln\left(\ell\left(X_{t-1+\delta_h} | X^{t-1}\right)\right)$.

The algorithm above is described in general terms; in principle the intramonth loop could be repeated at every instant within a month for which we have new data. Since we have only a subset of data intramonth, we vary the dimension of the vector of observables $X_{t-1+\delta_h}$ as a function of time $t - 1 + \delta_h$. In application, we repeat steps 1-5 only at certain minutes or days pre- and post-FOMC meeting. We initialize the algorithm with guesses for the Markov-switching parameters that vary across regime subperiods (only the policy rule parameters), while the fixed-coefficient parameters have guessed values that are identical across regime subperiods. These guesses are used to evaluate the posterior by combining the likelihood $L(\theta)$ with the priors. We continue guessing parameters and evaluating the posterior in this manner, until we find parameter values that maximize the posterior. With the posterior mode in hand, we evaluate the entire posterior distribution, as described below.

Observation Equation The mapping from the variables of the model to the observables in the data can be written using matrix algebra to obtain the observation equation $X_t = D_{\xi_t,t} + Z_{\xi_t,t}\left[S_t', \widetilde{y}_{t-1}\right]' + U_t v_t$. Denote $\widehat{g}_t \equiv g_t - g$, and $\widehat{lp}_t = lp_t - lp$. Using the definition of stochastically detrended output, we have $\widetilde{y}_t = \ln(Y_t/A_t)$, $\Delta\ln(A_t) \equiv g_t = g + \rho_g (g_{t-1} - g) + \sigma_g \varepsilon_{g,t} \Rightarrow \widetilde{y}_t - \widetilde{y}_{t-1} = \Delta\ln(Y_t) - g_t \Rightarrow \Delta\ln(Y_t) = \widetilde{y}_t - \widetilde{y}_{t-1} + g_t = \widetilde{y}_t - \widetilde{y}_{t-1} + \widehat{g}_t + g$. Annualizing the monthly growth rates to get annualized GDP growth we have $\Delta\ln(GDP_t) \equiv 12\Delta\ln(Y_t) = 12g + 12\left(\widetilde{y}_t + \widehat{g}_t - \widetilde{y}_{t-1}\right)$. For quarterly GDP growth we interpolate to monthly frequency using the method in Stock and Watson (2010). For our other quarterly variables (SPF survey measures) and our biannual Liv survey, we drop these from the observation vector in the months for which they aren't available. The observation equation when all variables in X_t are available takes the form:

$$
\begin{bmatrix} \Delta \ln (GDP_t) \\ Inflation \\ FFR \\ SOC\,(Inflation)_{12m} \\ SOC\,(Inflation)_{60m} \\ f_t^{(0)} \\ BC\,(Inflation)_{12m} \\ SPF\,(Inflation)_{12m} \\ Liv\,(Inflation)_{12m} \\ SPF\,(GDPDInfl)_{12m} \\ BBG\,(Inflation)_{12m} \\ Liv\,(Inflation)_{120m} \\ SPF\,(Inflation)_{120m} \\ BC\,(FFR)_{12m} \\ BC\,(\Delta GDP)_{12m} \\ BBG\,(\Delta GDP)_{12m} \\ SPF\,(\Delta GDP)_{12m} \\ f_t^{(n)} \\ Baa_t \\ pgdp_t \\ EGDP_{t,t-1} \end{bmatrix} = \begin{bmatrix} 12g \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ D_{\xi_t,\pi_{t,t+12}} \\ D_{\xi_t,\pi_{t,t+12}} \\ D_{\xi_t,\pi_{t,t+12}} \\ D_{\xi_t,\pi_{t,t+12}} \\ D_{\xi_t,\pi_{t,t+12}} \\ D_{\xi_t,\pi_{t,t+120}} \\ D_{\xi_t,\pi_{t,t+120}} \\ D_{\xi_t,i_{t,t+12}} \\ D_{\xi_t,y_{t+s}} \\ D_{\xi_t,y_{t+s}} \\ D_{\xi_t,y_{t+s}} \\ D_{\xi_t,i_{t+n}} \\ C_{Baa} \\ \ln(K)+g \\ K\exp(g) \end{bmatrix} + \begin{bmatrix} 12\left(\widetilde{y}_t + \widehat{g}_t - \widetilde{y}_{t-1}\right) \\ 12\pi_t \\ 12 i_t \\ \left[h + (h-1)\phi + (h-2)\phi^2 + ... + \phi^{11}\right](1-\phi)\overline{\pi}_t + \Sigma_{j=1}^{12}\phi^j \pi_t \\ \left[h + (h-1)\phi + (h-2)\phi^2 + ... + \phi^{59}\right](1-\phi)\overline{\pi}_t + \Sigma_{j=1}^{60}\phi^j \pi_t \\ 12 i_t \\ Z_{\xi_t,\pi_{t,t+12}} S_t \\ Z_{\xi_t,\pi_{t,t+12}} S_t \\ Z_{\xi_t,\pi_{t,t+12}} S_t \\ Z_{\xi_t,\pi_{t,t+12}} S_t \\ Z_{\xi_t,\pi_{t,t+12}} S_t \\ Z_{\xi_t,\pi_{t,t+120}} S_t \\ Z_{\xi_t,\pi_{t,t+120}} S_t \\ Z_{\xi_t,i_{t,t+12},} S_t \\ Z_{\xi_t,y_{t+s}-y_t} S_t \\ Z_{\xi_t,y_{t+s}-y_t} S_t \\ Z_{\xi_t,y_{t+s}-y_t} S_t \\ Z_{\xi_t,i_{t+n}} S_t \\ B\widehat{lp}_t \\ k_t - k + pd_t + \widehat{g}_t + \widetilde{y}_t - \widetilde{y}_{t-1} \\ K\exp(g)\left[k_t - k + \widehat{g}_t + \widetilde{y}_t - \widetilde{y}_{t-1}\right] \end{bmatrix} + U_t v_t \tag{A.11}
$$

where we have used the fact that expectations for the macro agent in the model is:

$$
\begin{aligned}
\mathbb{E}_t^m\left[\pi_{t,t+h}\right] &= \left[h + (h-1)\phi + (h-2)\phi^2 + ... + \phi^{h-1}\right]\alpha_t^m + \left[\phi + \phi^2 + ... + \phi^h\right]\pi_t \\
&= \left[h + (h-1)\phi + (h-2)\phi^2 + ... + \phi^{h-1}\right](1-\phi)\overline{\pi}_t + \left[\phi + \phi^2 + ... + \phi^h\right]\pi_t
\end{aligned}
$$

The term $Inflation$ in the above stands for CPI inflation; $GDPDInfl$ refers to GDP deflator inflation. The variable $f_t^{(n)}$ refers to the time-t contracted federal funds futures market rate. Here we use $n = \{6, 10, 20, 35\}$. The variable $pgdp$ is the log of the SP500 capitalization-to-lagged GDP ratio, i.e., $\ln(P_t/GDP_{t-1})$; $EGDP_{t,t-1}$ is the level of the SP500 earnings-to-lagged GDP ratio; taking a first order Taylor approximation of $EGDP_{t,t-1}$ around the log earnings-output ratio, we have $EGDP_{t,t-1} \approx K\exp(g)\left[1 + k_t - k + \widehat{g}_t + \widetilde{y}_t - \widetilde{y}_{t-1}\right]$, where K is the steady state level of $EGDP_t = \exp(k)$. To obtain high-frequency information on $EGDP_{t,t-1}$, we use the BBG earnings nowcasts divided by one-month lagged real-time GDP. For announcements that occur during the time frame for which these nowcasts are available (Jan 2, 2006 to July 24, 2024), we use the pre- and post- announcement BBG earnings nowcast-to-lagged GDP ratio. For announcements that pre-date this time frame, we use the end-of-month earnings-lagged GDP ratio allowing for observation error as a noisy proxy for these nowcasts. Baa_t is the Baa spread described above, where C_{Baa} and B are parameters. To allow for the fact that the true liquidity premium and convenience yield is only a function of Baa_t, we add a constant C_{Baa} to our model-implied convenience yield lp_t and scale it by the parameter B to be estimated. Unless otherwise indicated, all survey expectations are 12 month-ahead forecasts in annualized units.

The above uses multiple measures of observables on a single variable, e.g., investor expectations of inflation 12 months ahead are measured by four different surveys (BC, SPF, LIV, and BBG). In the filtering algorithm above, these provide four noisy signals on the same latent variable.

Computing the Posterior

The likelihood is computed with the Kim's approximation to the likelihood, as explained above, and then combined with a prior distribution for the parameters to obtain the posterior. A block algorithm is used to find the posterior mode as a first step. Draws from the posterior are obtained using a standard Metropolis-Hastings algorithm initialized around the posterior mode. Here are the key steps of the Metropolis-Hastings algorithm:

- Step 1: Draw a new set of parameters from the proposal distribution: $\vartheta \sim N\left(\theta_{n-1}, c\overline{\Sigma}\right)$
- Step 2: Compute $\alpha\left(\theta^{m};\vartheta\right) = \min\left\{p\left(\vartheta\right)/p\left(\theta^{m-1}\right), 1\right\}$ where $p\left(\theta\right)$ is the posterior evaluated at θ.
- Step 3: Accept the new parameter and set $\theta^{m} = \vartheta$ if $u < \alpha\left(\theta^{m};\vartheta\right)$ where $u \sim U\left([0,1]\right)$, otherwise set $\theta^{m} = \theta^{m-1}$
- Step 4: If $m \geq n^{sim}$, stop. Otherwise, go back to step 1

The matrix $\overline{\Sigma}$ corresponds to the inverse of the Hessian computed at the posterior mode $\overline{\theta}$. The parameter c is set to obtain an acceptance rate of around 30%. We use four chains of 540,000 draws each (1 of every 200 draws is saved). The four chains combined are used to form an estimate of the posterior distribution from which we make draws. Convergence is checked by using the Brooks-Gelman-Rubin potential reduction scale factor using within and between variances based on the four multiple chains used in the paper.

I Risk Adjustment with Lognormal Approximation

The asset pricing block of equations involves conditional subjective variance terms that are affected by Markov-switching random variables in the model. We extend the approach in Bansal and Zhou (2002) of approximating a model with Markov-switching random variables using a risk-adjustment while maintaining conditional log-normality. Consider the forward looking relation for the price-payout ratio:

$$\begin{aligned} P_t^D &= \mathbb{E}_t^b\left[M_{t+1}\left(P_{t+1}^D + D_{t+1}\right)\right] \\ \frac{P_t^D}{D_t} &= \mathbb{E}_t^b\left[M_{t+1}\frac{D_{t+1}}{D_t}\left(\frac{P_{t+1}^D}{D_{t+1}} + 1\right)\right]. \end{aligned}$$

Taking logs on both sides, we get:

$$pd_t = \log \left[\mathbb{E}_t^b \left[\exp \left(m_{t+1} + \Delta d_{t+1} + \kappa_{pd,0} + \kappa_{pd,1} pd_{t+1} \right) \right] \right].$$

Applying the approximation implied by conditional log-normality we have:

$$\begin{aligned} pd_t &= \kappa_0 + \mathbb{E}_t^b \left[m_{t+1} + \Delta d_{t+1} + \kappa_{pd,1} pd_{t+1} \right] + \\ & +.5 \mathbb{V}_t^b \left[m_{t+1} + \Delta d_{t+1} + \kappa_{pd,1} pd_{t+1} \right]. \end{aligned}$$

To implement the solution, we follow Bansal and Zhou (2002) and approximate the conditional variance as the weighted average of the objective variance across regimes, conditional on ξ_t. Using the simpler notation of the state equation,

$$S_t = C_{\xi_t} + T_{\xi_t} S_{t-1} + R_{\xi_t} Q \varepsilon_t,$$

the approximation takes the form

$$\mathbb{V}_t^b \left[x_{t+1} \right] \approx e_x' \mathbb{E}_t^b \left[R_{\xi_{t+1}} Q Q' R_{\xi_{t+1}}' \right] e_x \tag{A.12}$$

where e_x is a vector used to extract the desired linear combination of the variables in S_t. This approximation maintains conditional log-normality of the entire system. In the solution, C_{ξ_t} depends on the risk adjustment term $V_t^b \left[m_{t+1} + \Delta d_{t+1} + \kappa_{pd,1} pd_{t+1} \right]$ which depends on R_{ξ_t}. Conditional on the risk adjustment term, the numerical solution delivers the appropriate coefficients, C_{ξ_t}, T_{ξ_t}, and R_{ξ_t}. To solve this fixed point problem, we employ the iterative approach of Bianchi, Kung, and Tirskikh (2018). Specifically, we solve the model and get S_t for an initial guess on the risk adjustment V_t^b, denoted $V_t^{b(0)}$. Given the approximation (A.12) the term $V_t^b \left[m_{t+1} + \Delta d_{t+1} + \kappa_{pd,1} pd_{t+1} \right]$ only depends on ξ_t. For each policy regime ξ_t^P our initial guess $V_t^{b(0)}$ is therefore one value of V_t^b for each of the belief regimes ξ_t^b. The initial solution based on the initial guess $V_t^{b(0)}$ gives an initial value for R_{ξ_t}, denoted $R_{\xi_t}^{(0)}$. So far we have not used (A.12). Then, given $R_{\xi_t}^{(0)}$, we use (A.12) to get an updated $V_t^{b(1)} \approx e_x E_t^b \left[R_{\xi_{t+1}}^{(0)} Q Q' R_{\xi_{t+1}}^{(0)\prime} \right] e_x$. Given the updated risk adjustment $V_t^{b(1)}$ we resolve the model for S_t one more time, and verify that the new $R_{\xi_{t+1}}$ is the same as the one obtained before, i.e., the same as $R_{\xi_{t+1}}^{(0)}$. Note that, although $V_t^b \left[x_{t+1} \right]$ depends on $R_{\xi_{t+1}}$ only (it does not depend on C_{ξ_t} due to the approximation (A.12)), $R_{\xi_{t+1}}$ does not depend on V_t^b. Thus, we can stop here.

J Allowing for Belief Uncertainty About the Current Policy Rule

This Appendix shows what would happen in our results if beliefs about the current rule–holding constant beliefs/uncertainty about future policy regimes–had differed from the true estimated rule, as in our baseline estimates. First, we redo the shock decompositions for the stock market assuming that investors change their belief about the current rule after a Fed announcement. Figure A.II shows a decompsition that is analogous to that in Figure VII for the top four FOMC announcements in terms of absolute jumps in the market. Each panel plots the shock decomposition for one announcement. The first bar labeled "base" shows the results for our baseline model, which repeats information from Figure (VII). The next three bars show what would happen if investors had– in constrast to our baseline model–updated beliefs about the current policy rule in the wake of the announcement. Specifically, we assume they update their belief about the activism coefficient on output growth, ψ_y as a result of the announcement. In these cases investors' pre-announcement belief is equal to the true ψ_y, but the post-announcement belief changes to some $\psi_y^* \neq \psi_y$. The plot shows different cases where $\psi_y^* = \left\{0.5\psi_y, 1.5\psi_y, 2\psi_y\right\}$. The red triangles show the jump in the market implied by each specification, while the black dot shows the jump in the data. For the baseline model results shown in the first bars, the black dots and red triangles coincide. For the other specifications, the differences show the result of allowing investors to update their assessment of the current policy rule as a result of the announcement. We can see that the differences are negligble: the difference between the baseline model (black dot) and the red triangles in each case are small. Moreover, the relative contribution of different perceived shocks is virtually unchanged from the baseline case. In particular, changing perceptions about the economic state and/or beliefs about future regime change remain important contributors to the market's jump in all cases.

Next, we show how the model's implication for historical variation in the stock market would have changed if we allow for different beliefs about the parameters of the current policy rule. Panel (a) of A.III repeats the results for the historical variation in the log stock market-lagged GDP ratio implied by our baseline model. The blue line shows $pgdp_t$, (the log stock market-to-last months GDP ratio) from our baseline model which coincides with true data series. The red/dashed (purple/dashed-dotted) line in each panel plots a counterfactual in which the belief regime with the highest (lowest) perceived probability of exiting the policy rule was always in place. Repeating the information from **??**, we see that investor beliefs about the conduct of *future* monetary policy play an outsized role in stock market fluctuations, as can be observed from the quantitatively large gap between the red and purple lines in panel (a). Panel (b) of A.III plots two different counterfactual series, in which investors believed that the activism coefficient ψ_y on output growth in the policy rule had been double (half) the true estimated value over the post-millennial period. Panel (b) shows that a substantial range of different beliefs about the *current* policy rule has negligle effects on the historical variation in the stock market. By contrast, panel (a) shows that differing beliefs about the probability of switching to a new policy rule that is likely to be in place for an extended period of time has large effects.

Figure A.II: Effects of Post-FOMC Updates to Beliefs About Current Rule

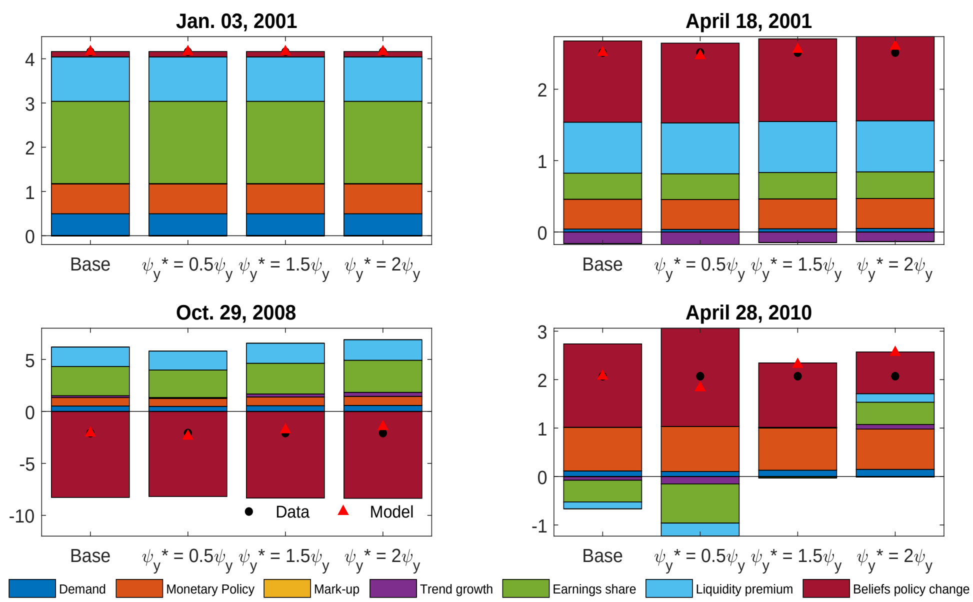

Notes: The figure reports a decomposition for the 4 most relevant FOMC announcements based on changes in the SP500-lagged GDP ratio. The first bar on the left gives our baseline estimate where investors know the true reaction coefficient on output growth in the current regime is ψ_y. The next three bars show what would have happened if instead investors had updated their belief about ψ_y to ψ_y^* *after* the FOMC meeting. The sample is 1961:M1-2020:M2.

Figure A.III: Effects of Beliefs About Future Rule vs Beliefs About Current Rule

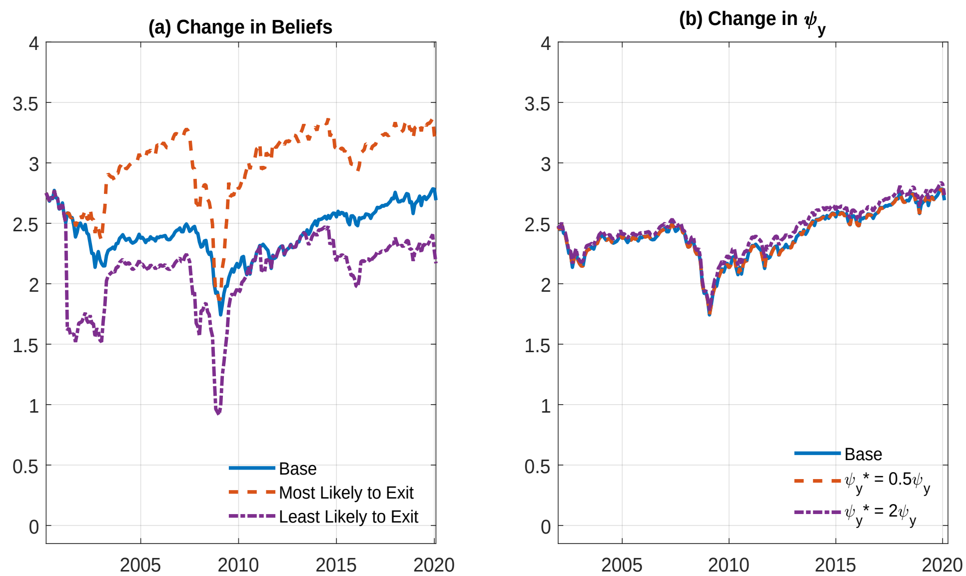

Notes: The blue line in each panel plots the log of the S&P 500 market capitalization-lagged GDP ratio. In panel (a) this coincides with the historical variation implied by the baseline model. The red/dashed (purple/dashed-dotted) line in panel (a) plots a counterfactual S&P 500 to GDP ratio in which the belief regime with the highest (lowest) perceived probability of exiting the policy rule was always in place. The red/dashed (purple/dashed-dotted) line in panel (b) plots a counterfactual S&P 500 to GDP ratio in which the investor had counterfactually believed the activism coefficient on output growth in the policy rule had been double (half) the true estimated value. The sample for the counterfactual spans 2000:M3 to 2020:M2.