The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXXVII, NO. 2 • APRIL 2022

Monetary Policy and Asset Valuation

FRANCESCO BIANCHI, MARTIN LETTAU, and SYDNEY C. LUDVIGSON

ABSTRACT

We document large, longer term, joint regime shifts in asset valuations and the real federal funds rate-r^* spread. To interpret these findings, we estimate a novel macro-finance model of monetary transmission and find that the documented regimes coincide with shifts in the parameters of a policy rule, with long-term consequences for the real interest rate. Estimates imply that two-thirds of the decline in the real interest rate since the early 1980s is attributable to regime changes in monetary policy. The model explains how infrequent changes in the stance of monetary policy can generate persistent changes in asset valuations and the equity premium.

A GROWING LITERATURE DOCUMENTS THAT the real values of long-term financial assets—including the stock market, a perpetual asset that endures indefinitely—fluctuate sharply in response to the actions and announcements of central banks. But this leads to a puzzle. Asset pricing (AP) theories can generally rationalize such large responses only if market participants believe that something related to monetary policy will have a long-lasting effect on real

---

Francesco Bianchi is with Department of Economics, Duke. Martin Lettau is with Haas School of Business, University of California at Berkeley. Sydney C. Ludvigson is with Department of Economics, New York University. Ludvigson received financial support from the C.V. Starr Center for Applied Economics at NYU, and was a member of the Economic Advisory Panel of the Federal Reserve Bank of New York, which meets twice per year to discuss policy areas of interest to the president of the bank; this position is unpaid. We thank Ian Dew-Becker; Emmanuel Farhi; Simon Gilchrist; Matteo Maggiori; Stefan Nagel (Editor); Emi Nakamura; Monika Piazzesi; Jon Steinsson; Andrea Tamoni; Harald Uhlig; Michael Weber; three anonymous referees; and seminar participants at the Bundesbank-Columbia, HEC Montreal, NYU, UCL, UCLA, the University of Texas at Austin, Minneapolis Fed, Indiana University, Carey – Johns Hopkins University, the National Bank of Belgium, Booth – University of Chicago, the Workshop on Monetary Policy (NY Fed, April 2017), the Financial Markets and Macroeconomic Performance conference (Frankfurt, May 2017), the Developments in Empirical Monetary Economics Conference (FRB May 2017), the SED (June 2017), the CCBS Research Forum on Macrofinance (Bank of England June 2017), the Barcelona GSE Summer Symposium (June 2017), the Computing in Economics and Finance conference (June 2017), Wolfe Research's Global Quantitative and Macro Investment Conference (November 2017), and the 2018 American Economic Association meetings for helpful comments. We thank Paulo Manoel for excellent research assistance. Any errors or omissions are the responsibility of the authors. The authors have read *The Journal of Finance* disclosure policy and have no conflicts of interest to disclose.

Correspondence: Sydney C. Ludvigson, Department of Economics, New York University, 19 W. 4th Street, 6th Floor New York, NY 10012; e-mail: sydney.ludvigson@nyu.edu.

DOI: 10.1111/jofi.13107

variables.[1] Yet the notion that monetary policy shocks could have long-lived effects on real variables is contravened by both foundational New Keynesian macro theories and prior empirical evidence (e.g., Christiano, Eichenbaum, and Evans (2005)).[2] But if monetary policy shocks have at most short-lived effects on real variables, how can monetary policy influence long-lived real assets?

One possibility is that some component of monetary policy has long-lasting, first-order effects on the aggregate economy, on real interest rates, and on the stock market, even if monetary policy shocks do not. In this paper, we present new empirical evidence consistent with this hypothesis, and a new theoretical explanation consistent with the evidence.

We begin by showing that the U.S. economy is characterized by quantitatively large, decades-long regime shifts in asset values relative to macroeconomic fundamentals. These movements coincide with equally important regime shifts in the level of the real federal funds rate (FFR) in excess of a widely used measure of the "natural" rate of interest r^*, a spread that we refer to hereafter as the monetary policy spread (mps). Since the Federal Reserve targets the FFR but in theory has no control over the natural rate, a nonzero value for the mps may be considered a measure of the stance of monetary policy, that is, whether monetary policy is accommodative or restrictive. We refer to accommodative regimes with persistently negative values for the mps as "dovish," and restrictive regimes with persistently positive and high values for the mps as "hawkish."

Dovish regimes in our sample coincide with persistently high asset valuations and low equity market return premia, while hawkish regimes coincide with persistently low valuations and high equity market return premia. The estimation identifies two hawkish subperiods characterized by low valuations and a high mps: 1978:Q4 to 2001:Q3 and 2006:Q2 to 2008:Q2. The first period spans the Volcker disinflation and its aftermath, while the second follows 17 consecutive Federal Reserve rate increases that left the nominal funds rate standing at 5.25% in June 2006. All other subperiods of the sample are identified as dovish regimes with high valuations and low mps.

Taken together, this evidence suggests that low-frequency movements in short-term real interest rates are directly linked to low-frequency regime shifts in asset valuations and equity return premia. But how much—if any—of these findings can plausibly be attributed to monetary policy, even theorestically? After all, the canonical models described above would be wholly inconsistent with this evidence, since monetary policy in those paradigms has only short-lived effects on real variables.

To address this question, we specify and estimate a new macrofinance model of monetary policy transmission with two "blocks." The first block determines risky asset prices and is driven by the optimal behavior of a representative agent that earns income from investments in two assets: the aggregate stock

---

[1] We define a "real variable" here as any nonnominal variable, including risk premia and credit spreads.

[2] For a review of New Keynesian models, see Galí (2015).

market and the one-period nominal bond market. This agent may be thought of as a relatively sophisticated investor such as a wealthy individual or large institution. Although she owns the overwhelming majority of highly concentrated financial wealth in the United States, she is small enough relative to the overall population that she takes macroeconomic dynamics as given. We refer to this agent interchangeably as the AP agent or investor.

The second block of the model determines macroeconomic dynamics and is driven by a representative macro agent that has access to the nominal bond but holds no stock market wealth. This block consists of a set of equations similar to those commonly featured in New Keynesian models. But contrary to standard New Keynesian models, macro dynamics here are influenced by two distinctive features that, taken together, imply that the model can be consistent with long-lasting (but not permanent) departures from monetary neutrality.

The first such feature is sticky macroagent expectations about inflation. Specifically, we allow the evolution of expectations about trend inflation to be influenced by both an adaptive expectations component and a signal about the central bank's inflation target. For the adaptive component, expectations about future inflation are formed using a constant-gain learning algorithm, following the survey evidence established in Malmendier and Nagel (2016, MN). To ensure that model expectations evolve in a manner that closely aligns with observed expectations, we map the learning algorithm to data by filtering observations on household inflation expectations from the University of Michigan Survey of Consumers (SOC).

The second distinctive feature of the macro block is that we allow for regime changes in the conduct of monetary policy. These take the form of shifts in the parameters of a nominal interest rate rule that include both the inflation target and the activism coefficients governing how strongly the monetary authority responds to inflation-target deviations and to economic growth. Such changes in what we refer to as the *conduct* of monetary policy give rise to movements in the nominal interest rate 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.

A key aspect of the model for explaining the stock market behavior documented in the first part of the paper is the evolution of investor beliefs about infrequent shifts in the monetary policy rule. Investors in the model are presumed to closely follow central bank communications, so they observe when shifts in the policy rule occur. However, investors have no way of observing how long any observed shift in policy will last and must learn about its duration. We further assume that, once investors have spent enough time in a particular policy regime, memory of past policy rules fades and investors come to view the existing policy stance as the new normal. The combination of learning plus a fading memory distortion implies that investor beliefs evolve in a history-dependent manner, with important consequences for how asset valuations adjust in the wake of regime changes in monetary policy.

We solve and estimate the full theoretical framework, with the macro block parameters and latent states estimated using Bayesian methods under flat priors. The results imply that the parameters of the monetary policy rule differ markedly across the previously estimated mps regimes. Specifically, we find that the dovish, low-mps subperiods coincide with a dovish policy rule characterized by a comparatively higher inflation target and less responsiveness to inflation relative to growth, while the hawkish, high-mps subperiods coincide with a hawkish policy rule characterized by a lower inflation target and greater responsiveness to inflation relative to growth.

With the model estimation results in hand, we identify movements in real variables that are attributable solely to the conduct of monetary policy, that is to regime changes in the policy rule. Several results are worth noting.

First, the estimates imply that changes in the conduct of monetary policy have generated large and persistent fluctuations in the short-term real interest rate that last for decades. Indeed, the estimated model implies that two-thirds of the secular decline in real interest rates observed since the early 1980s is due to regime changes in the conduct of monetary policy. This result obtains because the policy rule parameters exhibit a decisive shift toward more hawkish values around the time of Volcker's appointment, but then exhibit an equally decisive shift back to more dovish values in the aftermath of the near collapse of Long-Term Capital Management, the tech bust in the stock market, and the 9/11 terrorist attacks. The conduct of monetary policy has remained dovish since, with the exception of a brief interlude from 2006:Q2 to 2008:Q2. These results stand in contrast to those for monetary policy shocks, which are found to have far more transitory effects, consistent with prior empirical evidence.

Second, our estimate of perceived trend inflation closely follows the adaptive learning rule, which plays a crucial role in the results. If perceived trend inflation is counterfactually set equal to the inflation target, regime changes in the conduct of monetary policy have no effect on the real interest rate.

Third, the model parameter estimates imply that dovish policy rules generate persistently high asset valuations, a low mps, and low equity return premia, while hawkish rules generate persistently low valuations, a high mps, and high return premia, consistent with our motivating empirical evidence. In addition, the conditional equity return premium estimated from historical data is strongly positively correlated over our sample with the component of the real interest rate that we find is driven by regime changes in monetary policy.

The success of the model in explaining these lower frequency AP phenomena comes from the product of two forces: (i) sticky macroagent expectations about inflation and (ii) revisions in investor expectations about future monetary policy. Sticky inflation expectations are necessary for monetary policy to generate the persistent movements in the real interest rate that in turn trigger large and persistent fluctuations in asset valuations. Investor learning about the persistence of regime shifts delivers a plausible, gradual adjustment in valuation ratios after regime shift dates, while the fading memory distortion of investor beliefs explains the behavior of stock market return premia across the hawkish and dovish subperiods of our sample. Intuitively, fading memory of past policy rules means that investors extrapolate too much from the observed continuity in the policy stance and are therefore surprised by the inevitable transition out of the existing policy rule. It follows that an econometrician looking back on the historical sample would find that hawkish (dovish) subperiods are predictably followed by a "surprise" (from the perspective of investors) increase (decrease) in excess returns as policy switches back to a dovish (hawkish) stance.

The research in this paper touches on several different strands of literature that connect monetary policy to movements in asset values. Although not focused specifically on announcement effects, our work is related to a growing body of evidence that finds that the values of long-term financial assets respond to the actions and announcements of central banks.[3] Economists have proposed various explanations for these responses, including the revelation of private central bank information and the response of risk premia.[4] Yet no matter the channel, AP models can typically only rationalize such large responses if something associated with the announcement is expected to have a long-lasting influence on real variables or risk premia.[5] Our work contributes to this literature by providing new evidence of regime changes in the conduct of monetary policy that have long-lasting effects on real interest rates, asset valuations, and equity market return premia, and by providing a novel theoretical explanation for these empirical findings.

Our empirical findings also relate to a theoretical literature in which shifts in the risk-free interest rate coincide with shifts in return premia.[6] Our empirical findings contribute to this literature by showing that persistently high asset valuations and persistently low return premia are associated with evidence of a persistently dovish monetary policy stance. In contrast to this literature, we provide a new explanation for low return premia in low interest rate regimes based on the idea that investors may overextrapolate from the observed continuity in the policy stance, thereby creating a wedge between the subjective and objective persistence of policy regimes.

---

[3] See Hanson and Stein (2015), Gertler and Karadi (2015), Gilchrist, López-Salido, and Zakrajšek (2015), Boyarchenko, Haddad, and Plosser (2016), Jarocinski and Karadi (2020), Cieslak and Schrimpf (2019), and Kekre and Lenel (2021). These studies follow on earlier work finding a link between monetary policy surprises and short-term assets in high-frequency data (Cook and Hahn (1989), Bernanke and Kuttner (2005), Gürkaynak, Sack, and Swanson (2005)). A separate literature studies the timing of when premia in the aggregate stock market are earned in weeks related to Federal Open Market Committee (FOMC)-cycle time (Lucca and Moench (2015), Cieslak, Morse, and Vissing-Jorgensen (2019)).

[4] See, for example, Nakamura and Steinsson (2018) on information effects and Gertler and Karadi (2015) on risk premia effects.

[5] For reviews of frontier AP models, see Cochrane (2005) and Campbell (2017).

[6] Prominent examples in this literature include theories with a "reach-for-yield" motive either in preferences or technologies (e.g., Rajan (2006, 2013), Diamond and Rajan (2012), Farhi and Tirole (2012), Coimbra and Rey (2017), Drechsler, Savov, and Schnabl (2018), Acharya and Naqvi (2019), Hanson, Lucca, and Wright (2021), Piazzesi and Schneider (2021)). Alternatively, a decline in real rates driven by monetary policy could increase the fraction of wealth held by more risk-tolerant investors, as in Kekre and Lenel (2021), driving down return premia.

Our findings relate to a body of theoretical work that connects the low and declining real interest rates of recent decades to risk premia.[7] In these theories, declining real rates are the result of shocks that increase the fraction of wealth held by more risk-averse or more pessimistic investors, implying that risk premia rise rather than fall as interest rates decline. Our findings differ in two ways from these papers. First, our estimates imply that two-thirds of the decline in short-term real interest rates since the early 1980s can be attributed to shifts in the monetary policy stance. Second, we find that low interest rate regimes coincide with lower rather than higher return premia. In this regard, our findings for stock market returns are reminiscent of similar evidence for the Treasury market (e.g., Hanson and Stein (2015)), U.S. prime money funds (e.g., Di Maggio and Kacperczyk (2015)), and U.S. corporate bond mutual funds (Choi and Kronlund (2018)).

Finally, our work is related to previous research that finds evidence of infrequent regime changes in the parameters of an estimated monetary policy rule (e.g., Clarida, Gali, and Gertler (2000), Lubik and Schorfheide (2004), Bianchi (2013)). In contrast to this work, we use a more recent sample and estimate whether there are *joint* regime changes in asset valuations and the real FFR-r^* spread, and we show that such changes coincide with regime shifts in the policy rule and equity return premia. We also present new evidence, relative to this literature, of changes in the policy rule parameters that are consistent with a persistently more dovish stance of monetary policy starting in the beginning of the 21st century.

This latter body of work also helps motivate why we use regime switching over alternative procedures such as slowly drifting means to document joint variation in valuations and policy rates. The papers cited in the previous paragraph use regimes to identify different phases of U.S. monetary history. It is quite natural to model changes in the conduct of monetary policy as occurring with discrete regime changes. Different Chairs of the Federal Reserve bring their own views and priorities to the conduct of monetary policy, suggesting that data on policy interest rates are likely to be better described as drawn from a mixture of distinct distributions with infrequent transitions between them, rather than from a single distribution in which a transition occurs each period.

This paper is organized as follows. Section I describes the procedure and results for measuring and modeling regimes in asset valuations, the mps, and return premia. Section II presents our macrofinance model of monetary transmission, details of the structural estimation procedure, and estimation results. Section IV concludes.

I. Regimes in Valuations, Interest Rates, and Equity Return Premia

This section describes how we model and estimate regimes in asset valuations and the mps using a Markov-switching model, and how we evaluate

---

[7] Barro et al. (2014), Caballero and Farhi (2014), and Hall (2016).

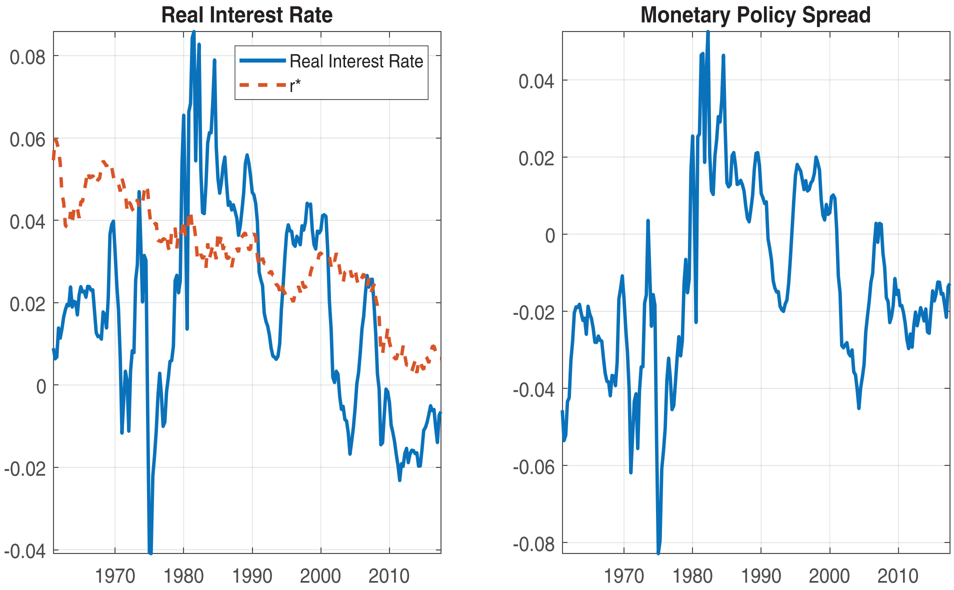

Figure 1. Real interest rate and monetary policy spread (mps). The real interest rate is the difference between the nominal FFR (FFR) and expected inflation, where expected inflation is computed as a four-quarter moving average of inflation. The mps is defined as $mps_t = FFR_t - Expected\ Inflation_t - r_t^*$., where r_t^* is the natural rate of interest from Laubach and Williams (2003). The sample spans the period 1961:Q1 to 2017:Q3. (Color figure can be viewed at wileyonlinelibrary.com)

whether these regimes are associated with movements in return premia. Before discussing the Markov-switching estimation, we begin by presenting preliminary evidence that helps motivate the evidence for long-lived regimes in these variables.

Figure 1 plots the behavior over time of a key monetary policy instrument, namely, the real FFR, measured for the purposes of this plot as the nominal rate minus a four-quarter moving average of inflation. The left panel plots this series along with an estimate of r^* from Laubach and Williams (2003). The data are quarterly and span the period 1961:Q1 to 2017:Q3.[8] The figure shows that there are important lower frequency fluctuations in the real FFR over the full sample, but little long-term trend. By contrast, the natural rate of interest exhibits a clear downward trend over the full sample. The right panel plots the spread between the real funds rate and the Laubach and Williams (2003) natural rate of interest, a variable we refer to as the mps. The natural rate of interest captures the component of the real rate whose fluctuations cannot be attributed to monetary policy.[9] Thus, the spread between the real

---

[8] The 1961 start date is dictated by the availability of the natural interest rate measure.

[9] Estimates of the natural interest rate apply theoretical restrictions on the behavior of real interest rates to identify the natural rate component. In Laubach and Williams (2003) these restrictions amount to estimates of the level of the real rate that are consistent with no change in inflation.

FFR and the natural rate is a measure of the stance of monetary policy, with spreads above zero indicative of restrictive monetary policy and those below zero indicative of accommodative monetary policy. Denote the time t value of this spread mps_t.[10] According to this measure of the mps, monetary policy in the sample was accommodative up until about 1980, then sharply restrictive from about 1980 to about 2000, and subsequently mostly accommodative. While there is no secular trend downward in real interest rates over the full sample, there is a noticeable downward trend in both the real interest rate and the mps since about 1980, a point we come back to below.

Next, Table I reports the correlations between the real interest rate or mps_t and different asset valuation metrics. These correlations are reported both for the raw series and for components of the raw series that retain fluctuations with medium-term cycles, defined as cycles that take between 8 and 50 years to complete, and business cycles, defined as cycles that take between 1.5 and 8 years to complete, computed with a bandpass filter. Panel A reports these correlations with $-cay_t$, the negative of the log consumption-wealth variable of Lettau and Ludvigson (2001, LL), one of the broadest asset valuation metrics available. With cay_t, asset values are measured relative to two macroeconomic fundamentals: log consumption c_t and log labor income y_t. The a_t is total household net worth, which is highly correlated with the return on the aggregate stock market. We use $-cay_t$ to put asset values in the numerator, and refer to it simply as a "wealth" ratio. Panels B to D consider alternative valuation ratios each of which has some measure of stock market wealth in the numerator: Panel B uses the Shiller price-earnings ratio[11], Panel C uses the price-dividend ratio for the corporate sector, and Panel D uses the price-earnings ratio for the corporate sector.

Several observations in Table I stand out. First, correlations between the valuation ratios and either the real funds rate or the mps are all negative at medium-term frequencies. Thus, over cycles of 8 to 50 years, persistently high valuations tend to coincide with indicators of monetary policy that are persistently more accommodative. By contrast, the correlations are all positive at business cycle frequencies and generally weaker in absolute terms.

Second, in all cases, the absolute correlation between the valuations and mps_t is greater than that between valuations and the real interest rate itself. Thus, purging the funds rate of the component estimated to be unrelated to monetary policy leads to greater negative comovement, which suggests that monetary policy as opposed to real rates per se plays a role in this correlation.

Third, the largest absolute correlation is with $-cay_t$, which has a correlation of -0.83 with the real interest rate and a correlation of -0.84 with the

---

[10] We compute mps_t as

$$FFR_t - (\text{expected inflation})_t - r_t^*,$$

where FFR is the nominal FFR, expected inflation is a four-quarter moving average of inflation, and r_t^* is the natural interest rate from Laubach and Williams (2003). The quarterly nominal funds rate is the average of monthly values of the effective FFR.

[11] http://www.multpl.com/shiller-pe/.

Table I

Correlation between Valuation Ratios and Fed Funds—r^* Spread

Results under "Medium"use series filtered to retain fluctuations with cycles between 8 and 50 years; "Business"retains cycles x, $1.5 \leq x \leq 8$ years. r^* is from Laubach and Williams (2003). Monetary policy spread $= FFR_t$ − expected inflation$_t$ $- r^*_t$, where expected inflation is a four-period moving average of inflation. Corp. PD ratio is the ratio of market equity (ME) to net dividends for the corporate sector from the flow of funds. Corp. PE ratio is the ratio of ME to after-tax profits of the corporate sector. NVA is net value added for the nonfinancial corporate sector. The sample spans 1961:Q1 to 2017:Q3.

	Overall	Medium	Business
Panel A: Correlations with $-cay_t$			
Real interest rate	−0.41	−0.83	0.25
Monetary policy spread	−0.52	−0.84	0.16
Panel B: Correlations with Shiller PE ratio			
Real interest rate	−0.30	−0.19	0.22
Monetary policy spread	−0.13	−0.30	0.18
Panel C: Correlations with Corp. PD ratio			
Real interest rate	−0.22	−0.49	0.22
Monetary policy spread	−0.25	−0.60	0.19
Panel D: Correlations with Corp. PE ratio			
Real interest rate	−0.28	−0.20	0.39
Monetary policy spread	−0.04	−0.30	0.29
Panel E: Correlations with earnings—NVA ratio			
Real interest rate	−0.54	−0.38	−0.27
Monetary policy spread	−0.35	−0.46	−0.16

mps at medium-term frequencies. This is followed by correlations of −0.49 and −0.60, respectively, with the corporate sector price-dividend ratio, −0.19 and −0.30 with the Shiller price-earnings ratio, and −0.20 and −0.30 with the corporate sector price-earnings ratio. This finding—that lower frequency movements in cay_t are more highly correlated in absolute terms with short-term interest rates than are other valuation ratios—is consistent with prior evidence that cay_t picks up more variation in expected stock market returns than do other stock market valuation ratios, and other stock market predictor variables in general.[12] One reason for this result is that some variation in expected stock market returns appears to be positively correlated with expected growth in stock market cash flows but not with expected growth in c_t or y_t (Lettau and Ludvigson (2005)). These movements in expected returns are therefore obscured in stock market valuation ratios where, unlike cay_t, expected stock

---

[12] See the review of the literature on expected stock market returns in Lettau and Ludvigson (2013).

market cash flows appear in the numerator. We observe this mechanism at work in the current data in Panel E of Table I. At medium-term frequencies, decreases in the real interest rate or mps, which tend to drive stock market valuation ratios up, are simultaneously associated with increases in the earnings share of output, which tend to drive them down. Since cay_t is not as subject to this type of confounding cash flow effect, and since discount rate movements are at the core of what we investigate in this study, we use $-cay_t$ as a measure of valuations in our formal econometric analysis, discussed next.

A. A Joint Regime-Switching Specification

This section presents results for a joint Markov-switching model of breaks in the means of *cay* and the mps.

The log valuation variable cay_t is derived from an approximate formula for the log consumption to aggregate (human and nonhuman) wealth ratio, and its relationship with future growth rates of a_t and/or future growth rates of c_t and y_t can be motivated from an aggregated household budget constraint.[13] An approximate expression linking c_t, a_t, and y_t to expected future returns to asset wealth, consumption growth, and labor income growth may be derived to yield

$$cay_t \equiv c_t - \gamma_a a_t - \gamma_y y_t \approx \alpha + \mathbb{E}_t \sum\nolimits_{i=1}^{\infty} \rho_w^i \big((1-\nu) r_{a,t+i} - \Delta c_{t+i} + \nu \Delta y_{t+1+i}\big), \quad (1)$$

where ν is the steady-state ratio of human wealth to asset wealth and $r_{a,t}$ is the log return to asset (nonhuman) wealth. Theory typically implies that c_t, a_t, and y_t should be cointegrated, or that the linear combination of variables in cay_t should be covariance-stationary.

In the standard estimation without regime shifts in any parameters, the stationary linear combination of c_t, a_t, and y_t may be written as

$$cay_t^{FC} \equiv c_t - \gamma_a a_t - \gamma_y y_t = \alpha + \epsilon_t^{FC}, \quad (2)$$

where the parameters to be estimated are α, γ_a, and γ_y. The residual ϵ_t^{FC} is the mean zero stationary linear combination of these data, referred to as the cointegrating residual. Note that ϵ_t^{FC} is not in general an i.i.d. shock. The superscript "*FC*" stands for "fixed coefficients" to underscore the fact that no parameters in this relation are time-varying.

In this paper, we estimate a Markov-switching version of this variable, analogously written as

$$cay_t^{MS} \equiv c_t - \beta_a a_t - \beta_y y_t = \alpha_{\xi_t} + \epsilon_t^c, \quad (3)$$

where $\epsilon_t^c \sim N(0, \sigma_{MS}^2)$. The intercept term, α_{ξ_t}, is a time-varying mean that depends on the existence of a latent state variable, ξ_t, presumed to follow a

---

[13] This formula is derived under several assumptions described in LL and elaborated on in Lettau and Ludvigson (2010). If labor income is modeled as the dividend paid to human capital, we get the formulation below.

two-state Markov-switching process with transition matrix $\mathbf{H}$. Thus, α_{ξ_t} assumes one of two discrete values, α_1 or α_2. The choice of two regimes is not crucial, but provides a readily interpretable way to organize the data into low- and high-valuation regimes. The residual ϵ_t^c is a stationary, continuous-valued random variable by assumption. The slope coefficients β_a and β_y are analogous to γ_a and γ_y in the fixed coefficient regression (2). They are denoted differently to underscore the point that the coefficients in (2) and (3) are not the same, just as the parameters α and α_{ξ_t}, or the residuals ϵ_t^{FC} and ϵ_t^c are not the same. Because our procedure jointly recovers the slope coefficients β_a and β_y, the timing of regime changes, and, as an implication, the decomposition of cay_t^{MS} into α_{ξ_t} and ϵ_t^c, all three statistical objects can differ.

We assume that regime changes in the mean of cay_t^{MS} coincide with regime changes in the mean of the mps:

$$mps_t = r_{\xi_t} + \epsilon_t^r, \tag{4}$$

where $\epsilon_t^r \sim N(0, \sigma_r^2)$. Unlike cay_t^{MS}, mps_t is an observed variable. Thus, in this case we only need to estimate the Markov-switching intercept coefficient r_{ξ_t}. The same latent state variable, ξ_t, is presumed to follow a two-state Markov-switching process with transition matrix $\mathbf{H}$, and controls changes in both α_{ξ_t} and r_{ξ_t}. The regimes are therefore synchronized across the two means.

The econometric model may be succinctly stated as a joint Markov-switching regression system with synchronized regimes,

$$\begin{aligned} c_t &= \alpha_{\xi_t} + \beta_a a_t + \beta_y y_t + \epsilon_t^c, \quad && \epsilon_t^c \sim N\left(0, \sigma_{MS}^2\right) \\ mps_t &= r_{\xi_t} + \epsilon_t^r, && \epsilon_t^r \sim N\left(0, \sigma_r^2\right), \end{aligned}$$

where ξ_t is a latent variable that follows a Markov-switching process with transition matrix $\mathbf{H}$. Let the vector $\boldsymbol{\theta} = (\alpha_{\xi_t}, \beta_a, \beta_y, r_{\xi_t}, \sigma_{MS}^2, \sigma_r^2, vec(\mathbf{H})')'$ denote the set of parameters to be estimated collectively.

We use Bayesian methods with flat priors to estimate the model parameters in (3) and (4) over the period 1961:Q1 to 2017:Q3. The sequence $\xi_t = \{\xi_1, \ldots, \xi_T\}$ of regimes in place at each point is unobservable and needs to be inferred jointly with the other parameters of the model. Estimates of α_{ξ_t} and r_{ξ_t} are formed by weighting their two estimated values by their state probabilities at each point in time. Let T be the sample size used in the estimation and let the vector of observations as of time t be denoted by $\mathbf{Z}_t$. Let $P(\xi_t = i|\mathbf{Z}_T; \boldsymbol{\theta}) \equiv \pi_{t|T}^i$ denote the probability that $\xi_t = i$, for $i = 1, 2$, based on information that can be extracted from the full sample and knowledge of the parameters $\boldsymbol{\theta}$. We decompose cay_t^{MS} into two components, namely, a discrete-valued time-varying mean and a continuous-valued random variable,

$$cay_t^{MS} = c_t - \left(\beta_a a_t + \beta_y y_t\right) = \overline{\alpha}_t + \epsilon_t^c \tag{5}$$

$$\overline{\alpha}_t = \sum\nolimits_{i=1}^{2} \pi_{t|T}^i \alpha_i, \tag{6}$$

and thus $\overline{\alpha}_t$ is the probability-weighted average of the Markov-switching means. An analogous bifurcation exists for mps_t, where r_{ξ_t} may be computed as $\overline{r}_t = \sum_{i=1}^{2} \pi^i_{t|T} r_i$.

The posterior distribution of the empirical model (3) and (4) and the corresponding regime probabilities $\pi^i_{t|t}$ and $\pi^i_{t|T}$ are obtained by computing the likelihood using the Hamilton filter (Hamilton (1994)) and combining it with priors. Since we use flat priors, the posterior coincides with the likelihood. Our estimate of cay^{MS}_t and its decomposition into $\overline{\alpha}_t$ and ϵ^c_t, and of mps_t into $\overline{r}_t$ and ϵ^r_t, use the posterior mode of the parameter vector $\boldsymbol{\theta}$ and the corresponding regime probabilities. Uncertainty about the parameters, or about any transformation of the model parameters, is characterized using a Gibbs sampling algorithm. The full statement of the procedure and sampling algorithm is given in the Internet Appendix.[14]

The variable cay^{MS}_t may be interpreted as log inverse asset valuation ratios, akin to a log dividend-price ratio as opposed to log price-dividend ratio. For brevity, we refer to cay^{MS}_t as an inverse wealth ratio, or equivalently define the (log) wealth ratio as $-cay^{MS}_t = -[\epsilon^c_t + \overline{\alpha}_t]$. Thus, a high α_i corresponds to a low wealth ratio, since $c_t - \beta_a a_t - \beta_y y_t$ is high whenever a_t is low relative to $c_t - \beta_y y_t$. In population, ϵ^c_t and ϵ^{FC}_t are mean-zero random variables and thus the intercept term $\overline{\alpha}_t$ gives the mean inverse wealth ratio.[15]

Since high values for mps_t are indicative of restrictive monetary policy while low values are indicative of accommodative policy, we refer to regimes with a high value for r_{ξ_t} as hawkish, denoted them with an H subscript, and to those with a low value for r_{ξ_t} as dovish, denoted with a D subscript, that is, $r_H \geq r_D$. Because the regimes in r_{ξ_t} and α_{ξ_t} are synchronized, changes in r_{ξ_t} will by construction coincide with changes in α_{ξ_t}. However, the magnitude by which either variable changes, and whether α_{ξ_t} will be high or low when r_{ξ_t} is high, are open empirical questions.

Table II reports the parameter estimates, while Figure 2 plots the probability of a hawkish regime over time for the Markov-switching intercept r_{ξ_t} based on the posterior mode parameter estimates. The results show that the sample is divided into five subperiods characterized by the two regimes for α and r. The hawkish regime with the high value for $r_{\xi_t} = r_H$ is also a high-α regime with posterior mode point estimates equal to $\widehat{r}_H = 0.0111$ and $\widehat{\alpha}_H = -0.7239$. The posterior mode estimates for the low $r_{\xi_t} = r_D$ dovish regime are $\hat{\alpha}_2 = -0.7500$ and $\widehat{r}_2 = -0.0252$. Since a high α for cay corresponds to a low valuation ratio, this implies that the dovish mps regime coincides with high asset valuations, while the hawkish mps regime coincides with low asset valuations.

The overall sample is divided into estimated regime subperiods using the most likely estimated regime sequence, a T-dimensional vector denoted by

---

[14] The Internet Appendix may be found in the online version of this article.

[15] In a finite sample, ϵ^{MS}_t and ϵ^{FC}_t are not necessarily mean-zero because of the leads and lags of the first differences included in the dynamic least squares regression used to correct for finite-sample biases—see the Internet Appendix. In population, these variables are mean-zero by definition.

Table II

Joint Regime Switching Model: Parameter Estimates

The top panel reports posterior modes, means, and 90% error bands of the parameters of the Markov-switching cointegrating relation. Flat priors are used on all parameters of the model. The lower panel reports parameter estimates for the fixed coefficient cointegrating relation. Standard errors are in parentheses. The two distinct values for the Markov-switching parameters are denoted by H and D subscripts to indicate hawkish or dovish values. The sample is quarterly and spans the period 1961:Q1 to 2017:Q3.

Parameter	Mode	Mean	5%	95%
α_H	−0.7239	−0.7121	−0.7796	−0.6465
α_D	−0.7500	−0.7376	−0.8034	−0.6717
r_H	0.0111	0.0132	0.0097	0.0165
r_D	−0.0252	−0.0244	−0.0266	−0.0222
$\alpha_H - \alpha_D$	0.0262	0.0255	0.0212	0.0296
$r_H - r_D$	0.0363	0.0376	0.0342	0.0411
β_a	0.2762	0.2721	0.2414	0.3014
β_y	0.7619	0.7657	0.7286	0.8042
σ_c	0.0128	0.0143	0.0130	0.0157
σ_r	0.0141	0.0135	0.0123	0.0150
$\mathbf{H}_{HH}$	0.9793	0.9696	0.9306	0.9943
$\mathbf{H}_{DD}$	0.9830	0.9785	0.9539	0.9950

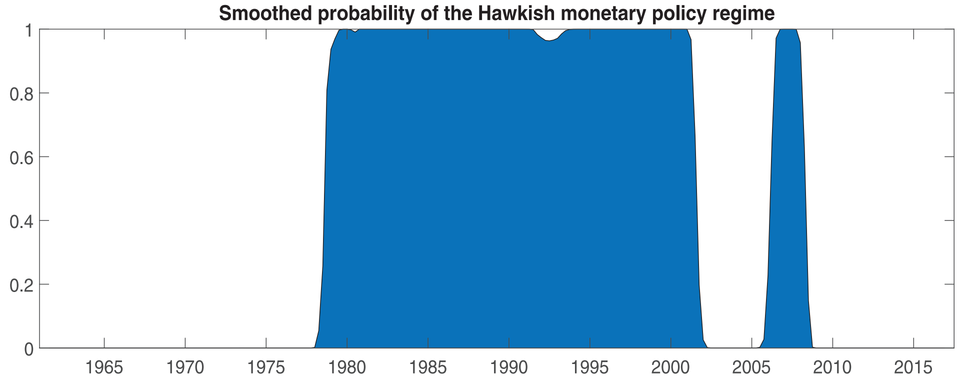

Figure 2. Regime probabilities. Smoothed probabilities of the hawkish monetary policy regime. The sample is quarterly and spans the period 1961:Q1 to 2017:Q3. (Color figure can be viewed at wileyonlinelibrary.com)

ξ^T.[16] Table III shows the resulting regime subperiods based on this estimated regime sequence. The hawkish regime prevails for a prolonged period of time from 1978:Q4 to 2001:Q3, during which the smoothed probability that $r = \widehat{r}_H$ is very close to unity. By contrast, the pre-1978 and most of the post-2001 subsample are dovish subperiods with high asset valuations, where the probability that $r = \widehat{r}_H$ is virtually zero. The hawkish regime briefly reappears from

---

[16] The Internet Appendix describes how the most likely regime sequence is computed from the filtered probabilities.

Table III
Estimated Regime Sequence

The table reports the most likely regime sequence based on the posterior mode estimates. Dovish refers to the low monetary policy spread regime and hawkish refers to the high regime.

	1961:Q1–1978:Q3	1978:Q4–2001:Q3	2001:Q4–2006:Q1	2006:Q2–2008:Q2	2008:Q3–2017:Q3
Regime	Dovish (2)	Hawkish (1)	Dovish (2)	Hawkish (1)	Dovish (2)

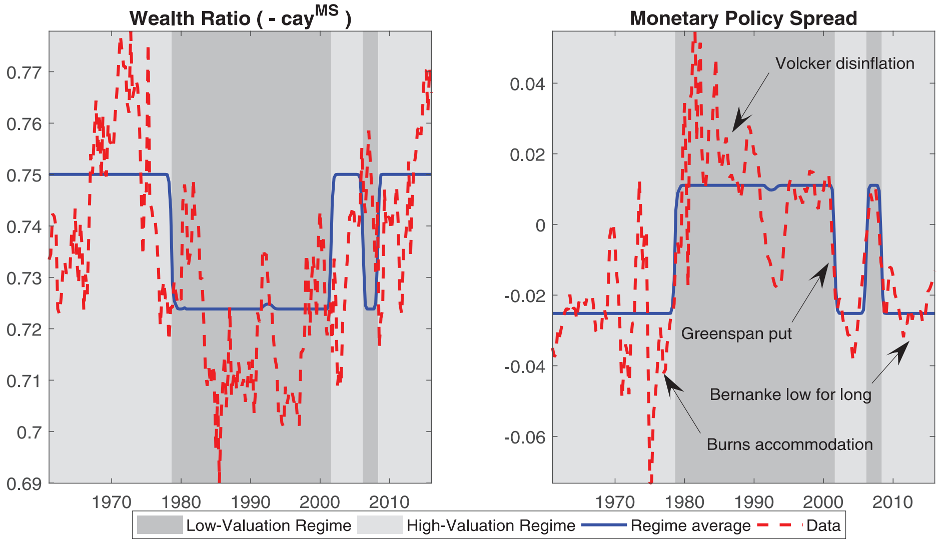

Figure 3. Wealth ratio and mps in the data. This figure plots the wealth ratio ($-cay^{MS}$) and the monetary policy spread $mps_t \equiv FFR_t - Expected\ Inflation_t - r_t^*$. The series for r_t^* comes from Laubach and Williams (2003). The solid line corresponds to the estimated mean at the posterior mode. The sample spans 1961:Q1 to 2017:Q3. (Color figure can be viewed at wileyonlinelibrary.com)

2006:Q1 to 2008:Q2 following a string of 17 target FFR hikes by the Federal Reserve that began on June 30, 2004, and ended with the nominal rate standing at 5.25% on June 29, 2006. The target funds rate remained above 4% until January 2008, when it was lowered to 3%.

The 90% credible sets for $\hat{\alpha}_H - \hat{\alpha}_D$ and $\widehat{r}_H - \widehat{r}_D$ are nonzero and positive, indicating that the data strongly favor changes in the mean of the log wealth ratio and the mps across the estimated regime subperiods. The two regimes are stationary but persistent, as indicated by the estimated diagonal elements of the transition matrix **H**, also reported in Table II.

Figure 3 plots $-cay_t^{MS}$ and the mps over time, along with the values $-\overline{\alpha}_t$ and $\overline{r}_t$ that arise in each regime over the sample. The figure shows that these estimated values differ by quantitatively large magnitudes across regime subperiods. The wealth ratio $-cay_t^{MS}$ fluctuates around two distinct means in five separate periods of the sample: a high mean in the early part of the sample, a low mean from 1978:Q4 to 2001:Q3, a high mean from 2001:Q4 to 2005:Q4, a low mean in the shorter subperiod from 2006:Q1 to 2008:Q2, and a high mean again at the end of the sample. The mps is a mirror image, fluctuating around a low mean in the early part of the sample, a high mean in the middle, and, with the exception of 2006:Q1 to 2008:Q2, a low mean everywhere else in the sample.

Several narrative "events" in monetary history are labeled in the mps panel of Figure 3. The first occurrence of the high-asset valuation/low-mps regime from 1961:Q1 to 1978:Q3 coincides with the run-up of inflation in the 1960s and 1970s and low real interest rates. Researchers have concluded that monetary policy failed to react aggressively to inflation during those years (Clarida, Gali, and Gertler (2000), Lubik and Schorfheide (2004), Sims and Zha (2006), Bianchi (2013)). This is labeled the "Burns Accommodation," after Arthur Burns who chaired the Federal Reserve Board over much of this subperiod. Real interest rates increased significantly during the "Volcker disinflation" and remained high for a prolonged period of time, coinciding with low valuations and high mps. The beginning of second occurrence of the high asset valuation/low *mps* regime is labeled "Greenspan Put" in Figure 3, after 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 and (allegedly) resulting in a perception of put protection on asset prices. The high-valuation/low-mps subperiod at the end of the sample overlaps with the explicit forward guidance "low-for-long" policies under Chair Bernanke, who promised in 2011 to keep interest rates at ultra low levels for an extended period of time, possibly longer than warranted by a 2% inflation objective. We argue that narrative events such as these are likely to coincide with infrequent shifts in the stance of monetary policy, shifts that are well captured by a Markov-switching specification.

B. Return Premia

The evidence above suggests that a persistently low mps_t is associated with persistently high asset valuations. A natural question that arises is whether the estimated dovish/hawkish mps regimes influence asset valuations only through the real interest rate, or whether return premia also change. To shed light on this question, we estimate a Markov-switching vector autoregression (MS-VAR) that takes the form

$$Z_t = A_{\xi_t} Z_{t-1} + V_{\xi_t} \varepsilon_t,$$

where Z_t is a column vector containing n demeaned variables observable at time t.[17] The variables in Z_t are a measure of the return on the U.S. stock market in excess of a short-term real interest rate along with predictor variables

---

[17] If the MS-VAR has more than one lag, the companion form can be used to recast the model as illustrated above.

relevant for the excess return.[18] The MS-VAR coefficients and shock volatilities vary with the discrete-valued random variable ξ_t, which evolves in our application according to a two-state Markov-switching process with transition matrix $\mathbf{H}$. To establish whether return premia differ across the two previously estimated regimes, we impose the estimated mps regime sequence ξ^T on the MS-VAR. Note that there is no implication from this procedure that return premia must necessarily show evidence of structural change across regimes. All parameters other than the regime sequence (i.e., A_{ξ_t}, V_{ξ_t}, $\mathbf{H}$) are freely estimated under flat priors and could in principle show no shift across the previously estimated regime subperiods.

It is instructive to consider how regime-switching can affect estimates of return premia over time. We are interested in the behavior of the present discounted value of all future return premia, which we denote $pdv_t \equiv \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t r_{t+1+j}$, where r_t is a measure of the return on the market less a proxy for the riskless real interest rate.[19] To estimate the conditional expectation terms, that is, $\mathbb{E}_t r_{t+1+j}$, we use the MS-VAR to compute econometric, time t forecasts of excess returns multiple steps ahead. Since excess returns are one element of Z_t, these are all functions of $\mathbb{E}(Z_{t+s}|\mathbb{I}_t)$. This expectation conditions on $\mathbb{I}_t$, which includes the history of observations Z^t, knowledge of the regime in place at time t, and the VAR parameters for each regime. With regime-switching, the time t conditional forecast also takes into account the likelihood of *future* regime changes. Intuitively, this is done by computing multistep-ahead VAR forecasts of returns at $t+s$, conditional on every possible future regime path, $\xi_{t+1}, \ldots, \xi_{t+s}$, and weighting these forecasts by the probability $\Pr(\xi_{t+1}, \ldots, \xi_{t+s}|\mathbb{I}_t)$ of each path. (The Internet Appendix gives the precise formulas.) The probability $\Pr(\xi_{t+1}, \ldots, \xi_{t+s}|\mathbb{I}_t)$ can be computed using just two pieces of information, namely, the regime in place at t and knowledge of $\mathbf{H}$. This shows that, as the horizon s grows, econometric forecasts reflect the increasing likelihood of a future regime change and its consequence for the variables in Z_{t+s}.

Table IV reports the regime-average values of pdv_t^i for each regime i, denoted $\overline{pdv}_t^i$. The regime average $\overline{pdv}_t^i$ is defined as the expected value of pdv_t conditional on being in regime i today and on the variables of the MS-VAR being equal to their conditional steady-state mean values for regime i. The Internet Appendix gives formal expressions for the regime average, and shows how they are computed from the MS-VAR parameters. Table IV reports the median and 68% credible sets for $\overline{pdv}_t^i$, computed from each draw of the VAR

---

[18] The MS-VAR includes (i) the market excess return, computed as the difference in the CRSP value-weighted stock market return (including dividend redistributions) and the three-month Treasury bill rate; (ii) $-cay^{MS}$; (iii) the small stock value spread (log-difference in the book-to-market ratio of the S1 value and S1 growth portfolio) following Campbell and Vuolteenaho (2004); (iv) the SMB factor from Fama and French (1993); and (v) the HML factor from Fama and French (1993). These variables are included because they improve the Akaike information criterion.

[19] We follow Cohen, Polk, and Vuolteenaho (2003) and set $\rho = 0.9898$ at a quarterly rate, or the annual rate used in Cohen, Polk, and Vuolteenaho (2003) raised to the power 0.25, $\rho = (0.96)^{0.25}$.

Table IV

Breaks in Market Premium

The first two rows report the regime averages of the present discounted value (PDV) of market expected excess returns. The row labeled "Difference" reports the difference between the PDVs of the hawkish and dovish regimes. The numbers in each cell are the median values of the statistic from the posterior distribution, with 68% posterior credible sets in parentheses. The last row reports the probability that premia decline when moving from the hawkish to the dovish regime, computed as the fraction of draws from the posterior distribution for which the premia under the dovish regime are lower than the premia under the hawkish regime.

	Market
Hawkish regime	1.5896
	(0.8960,2.2558)
Dovish regime	1.2848
	(0.5652,1.9219)
Difference	0.2987
	(−0.0367,0.7089)
Prob. decline	0.81
Odds ratio	4.26

parameters from the posterior distribution. The third row reports the difference between the $\overline{pdv}_t^i$ in the hawkish and dovish regimes as implied by the MS-VAR estimates. The last row reports the posterior probability that return premia decline in the dovish, low-mps regime, computed as the percentage of draws from the posterior distribution of regime-averages for which return premia are lower in the dovish regime than the hawkish regime.

Table IV shows that the median value of $\overline{pdv}_t^i$ is larger in the hawkish regime than in the dovish regime, implying that the difference between the two is always positive. This implies in turn that the estimated equity return premium is on average lower in environments with persistently low real interest rates. Although the 68% posterior credible sets include negative values for the difference in equity premium, this does not imply that negative values are likely. The posterior distribution of the difference displays substantial negative skewness and hence the probability assigned to a lower equity premium in the dovish mps subperiods is 81%. The odds that the equity premium is lower in the dovish mps regime is over four to one. In short, the mass of probability overwhelmingly favors one particular interpretation, namely, that return premia are lower in dovish mps subperiods than they are in hawkish subperiods.

Figure 4 depicts the median value of pdv_t over our sample as solid (blue) lines, with the regime average values indicated by the red (dashed) lines. Although the equity premium is volatile, it fluctuates around distinct means across the regime subperiods. The equity premium reaches lows or near-lows in the postmillennial period, after rising modestly in the aftermath of the financial crisis of 2007 to 2008. The equity premium then returns to historically low levels in the postcrisis zero-lower-bound (ZLB) period of our sample.

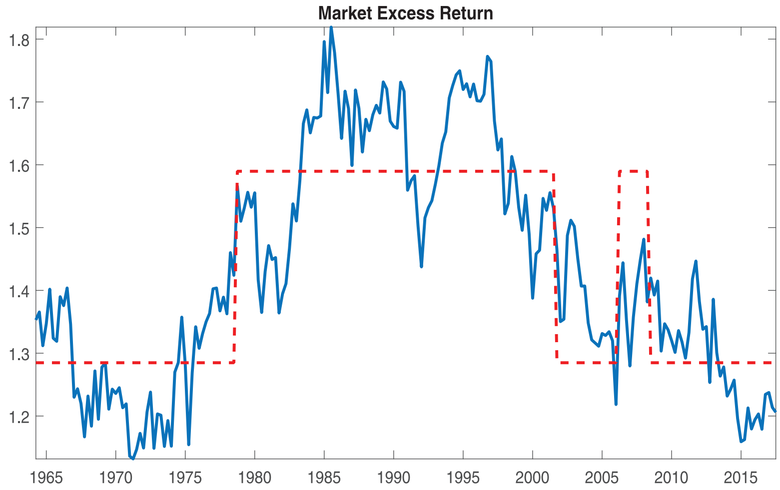

Figure 4. Evolution of return premia in the data. The figure reports the evolution of the present discounted value of risk premia for the stock market and three different spread portfolios. The blue solid line reports the evolution of the risk premia over time, while the red dashed line corresponds to the conditional steady state of the present discounted value based on the regime in place. Both are computed by taking into account the possibility of regime changes. The sample spans the period 1964:Q1 to 2017:Q3. (Color figure can be viewed at wileyonlinelibrary.com)

II. A Macrofinance Model of Monetary Transmission

To interpret the evidence in the previous sections, we propose a new dynamic macrofinance model of monetary policy transmission with two "blocks." These blocks describe the behavior of two different representative agents, "investors" and "households." In both blocks we work with a loglinear approximation of the model that can be solved analytically in which all random variables are conditionally lognormally distributed.

To understand the impetus for modeling two types of agents, note that the motivating evidence of the previous sections suggests that monetary policy has large and persistent effects on the real interest rate. Such persistent real effects are inconsistent with canonical New Keynesian models because agents' rational expectations quickly adapt to changes in monetary policy. This suggests that macro expectations may be subject to more inertia than what rational expectations would imply. However, financial markets react swiftly to central bank communications and actions. This suggests that the expectations of financial market participants are subject to little inertia, at least as pertains to beliefs about central bank actions. The framework below reconciles these seemingly contradictory observations by considering two types of agents with different beliefs. We now describe the two blocks of the model.

The first block is an AP block that determines equilibrium risky asset prices in the model. This block is driven by the optimal behavior of a representative agent whose income comes from investments in two assets: the stock market and the one-period nominal bond market. This agent may be thought of as a relatively sophisticated investor that typifies a wealthy individual or large institution, that is, an investor that constitutes a small fraction of the population but owns the vast majority of highly concentrated financial wealth in the United States.[20] We assume that this agent is small enough as a fraction of the general population that she takes macroeconomic dynamics as given, including the beliefs of households, which are assumed to be the key drivers of expectations about macro variables. We refer to this agent interchangeably as the AP agent or investor.

The second block of the model determines macroeconomic dynamics. This block is driven by a set of reduced-form equations similar to those standard in New Keynesian models. However, contrary to standard New Keynesian models, macroeconomic dynamics here are influenced by two distinctive features: sticky expectations about inflation of the type documented in Malmendier and Nagel (2016) (MN), and regime changes in the conduct of monetary policy. The purpose of these departures is to arrive at a model that can generate persistent (though not permanent) departures from monetary neutrality, as suggested by the evidence above. The extent of sticky expectations in inflation is disciplined by forcing the model to match data on household inflation expectations from the University of Michigan's SOC. Thus, macroeconomic dynamics can be thought of as driven by a central bank and an "average" household that typifies the vast majority of the population with comparatively negligible financial assets but whose expectations about inflation and aggregate economic activity preponderate in the general population.

An important aspect of the AP block of the model is the evolution of investor beliefs about infrequent shifts in the monetary policy rule. These beliefs are central to how shifts in the stance of monetary policy affect asset valuations and return premia. Investors in the model are presumed to closely follow central bank communications, so they observe when shifts in the monetary policy rule occur. However, we make two departures from the standard rational expectations assumption that the agent can observe the true transition matrix for monetary policy regime shifts. First, we assume that agents are uncertain about how long any observed policy shift will last and hence must learn about its duration. Second, we assume that agents exhibit fading memory of past policy rules, similar to evidence uncovered in surveys of inflation by MN and Malmendier and Nagel (2011). Specifically, if agents spend enough time in a particular policy regime, memory of past policy rules fades and they come to believe that the existing policy stance will persist indefinitely, a form of overextrapolation that overstates the true persistence of the regime shifts. As we discuss below, the combination of these two features of investor beliefs (learning plus

---

[20] Only about half of households report owning stocks either directly or indirectly in 2016 according to the Survey of Consumer Finances (SCF). More importantly, even among those households that own equity, most own very little: the top 5% of the stock wealth distribution owns 76% of stock market value and earns a relatively small fraction of income as labor compensation. See Greenwald, Lettau, and Ludvigson (2019) for further discussion.

a fading memory distortion) implies that asset prices in the model respond to monetary policy regime changes by initially underreacting but eventually overreacting. These features of beliefs imply that the model is qualitatively consistent with independent empirical evidence showing that survey expectations—including those of professional forecasters—initially underreact to shocks but subsequently overreact (Angeletos, Huo, and Sastry (2020), Bianchi, Ludvigson, and Ma (2020)).

A. Model Description

AP Block. The model allows for a continuum of identical investors indexed by i who derive utility from consumption, $C^i_{p,t}$, at time t. We use the suffix "p" to denote variables pertaining to these AP agents. Investors trade in two assets: a nominal bond and equity. The agents' intertemporal marginal rates of substitution in consumption takes the form

$$M_{t+1} = \delta_t \left(C^i_{p,t+1}/C^i_{p,t}\right)^{-\sigma_p},$$

where $\delta_t \equiv \delta \exp(\vartheta^p_t)$ is a time-varying subjective time discount factor. The time discount factor is subject to an externality in the form of a patience shifter ϑ^p_t that individual investors take as given, driven by the market as a whole. A time-varying specification for the subjective time discount factor is essential for ensuring that, in equilibrium, investors are willing to hold the nominal bond at the interest rate set by the central bank's policy rule as specified below.

We assume that investors derive income only from asset holdings and that the nominal bond is in zero net supply. It follows that, in equilibrium, assets are priced by a representative investor who consumes per-capita aggregate equity payout, D_t. We further assume that aggregate payout is derived from a constant capital share k of aggregate output $\mathbb{Y}_t$, implying $D_t = k\mathbb{Y}_t$.[21] We therefore drop the i superscript here forward and denote the consumption of the representative investor as $C_{p,t} = D_t = k\mathbb{Y}_t$.

Let lowercase letters denote log variables, for example, $c_{p,t} = \ln(C_{p,t})$. The marginal rate of substitution M_{t+1} is the stochastic discount factor (SDF), with log SDF given by

$$m_t = \log(\delta) - \sigma_p\left(c_{p,t} - c_{p,t-1}\right) + \vartheta_{p,t-1}.$$

The representative investor chooses consumption and optimal nominal bond holdings to maximize the expected present discounted value of a stream of utility derived from consumption and convenience benefits from the nominal bond due to their liquidity and safety (Krishnamurthy and Vissing-Jorgensen (2012)). The resulting first-order condition for optimal holdings of a one-period

---

[21] The assumption of a constant capital share k is made in the current model for simplicity. An extension of the model to allow for time-varying k could in principle account for evidence that factor share fluctuations have influenced trends in equity valuations (Greenwald, Lettau, and Ludvigson (2019)). Our focus here is to isolate the component due to monetary policy, so we keep k constant.

zero-coupon bond with a face value equal to one nominal unit is

$$\overline{LP}^{-1} Q_t = \mathbb{E}_t^p \big[M_{t+1} \Pi_{t+1}^{-1} \big], \tag{7}$$

where Q_t is the nominal bond price, $\mathbb{E}_t^p$ denotes the subjective expectation of the AP agent (discussed below), $\Pi_{t+1} = P_{t+1}/P_t$ is the gross rate of general price inflation, and $\overline{LP} > 1$ is the convenience premium. We make the simplifying assumption that this premium is constant over time, which helps the model match a sizable *average* equity premium while ensuring that any time-variation in the equity premium is driven solely by endogenous fluctuations in investor beliefs about monetary policy, which is the central focus of this paper.

Taking logs of (7) and using the properties of conditional lognormality delivers an expression for the real interest rate as perceived by the investor:

$$i_t - \mathbb{E}_t^p [\pi_{t+1}] = -\mathbb{E}_t^p [m_{t+1}] - .5 \mathbb{V}_t^p [m_{t+1} - \pi_{t+1}] - \overline{lp},$$

where the nominal interest rate $i_t = -\ln(Q_t)$, $\pi_{t+1} \equiv \ln(\Pi_{t+1})$ is net inflation, and $\mathbb{V}_t^p[\cdot]$ is the conditional variance under the subjective beliefs of the investor. This expression shows that $\vartheta_{p,t}$ is implicitly defined as

$$\vartheta_{p,t} = -\big[i_t - \mathbb{E}_t^p [\pi_{t+1}] \big] + \mathbb{E}_t^p \big[\sigma_p \Delta c_{p,t+1} \big] - .5 \mathbb{V}_t^p \big[-\sigma_p \Delta c_{p,t+1} - \pi_{t+1} \big] - \overline{lp} - \ln(\delta). \tag{8}$$

Let P_t^D denote the total value of market equity, that is, the price per share times shares outstanding. Then with D_t equal to the total equity payout, the first-order condition for optimal shareholder consumption implies the Euler equation

$$P_t^D = \mathbb{E}_t^p \big[M_{t+1} \big(P_{t+1}^D + D_{t+1} \big) \big],$$

$$\frac{P_t^D}{D_t} = \mathbb{E}_t^p \left[M_{t+1} \frac{D_{t+1}}{D_t} \frac{P_{t+1}^D + D_{t+1}}{D_{t+1}} \right].$$

Taking logs on both sides and using the properties of conditional lognormality, we obtain an expression for the log price-payout ratio:

$$\begin{aligned} pd_t = {} & \kappa_0 + \mathbb{E}_t^p \big[m_{t+1} + \Delta d_{t+1} + \kappa_1 pd_{t+1} \big] \\ & + 0.5 \mathbb{V}_t^p \big[m_{t+1} + \Delta d_{t+1} + \kappa_1 pd_{t+1} \big], \end{aligned}$$

where $pd_t \equiv \ln(P_t^D / D_t)$.

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

$$r_{t+1}^D = \kappa_0 + \kappa_1 pd_{t+1} - pd_t + \Delta d_{t+1},$$

where $\kappa_1 = \exp(\overline{pd})/(1+\exp(\overline{pd}))$ and $\kappa_0 = \log(\exp(\overline{pd})+1) - \kappa_1 \overline{pd}$. Combining the above, the log equity premium is

$$\underbrace{\mathbb{E}_t^p\left[r_{t+1}^D\right] - \left(i_t - \mathbb{E}_t^p[\pi_{t+1}]\right)}_{\text{Equity Premium}} = \underbrace{\begin{bmatrix} -.5\mathbb{V}_t^p\left[r_{t+1}^D\right] - \mathbb{COV}_t^p\left[m_{t+1}, r_{t+1}^D\right] \\ +.5\mathbb{V}_t^p[\pi_{t+1}] - \mathbb{COV}_t^p[m_{t+1}, \pi_{t+1}] \end{bmatrix}}_{\text{Risk Premium}} + \underbrace{\overline{lp}}_{\text{Liquidity Premium}}, \tag{9}$$

where $\mathbb{COV}_t^p[\cdot]$ is the conditional covariance under the subjective beliefs of the agent.

Finally, we derive cay_t as implied by the model. Let C_t denote aggregate consumption, and let $c_t = \ln(C_t)$. To derive the model-implied cay_t, note that the coefficients γ_a and γ_y in (1), or β_a and β_y in (3), are approximately equal to the shares of asset wealth and human capital in aggregate (human plus nonhuman) wealth, respectively (see LL). LL show that if the streams of income accruing to human and nonhuman wealth are discounted at the same rate, these coefficients are identically equal to the capital and labor income shares in models in which such shares are constant, an assumption we maintain here. Recalling that k is the presumed constant capital share of aggregate income $\mathbb{Y}_t$, it follows that $(1-k)\mathbb{Y}_t = (1-k)C_t$ is implied labor income in the model, where the last equality uses the fact that $\mathbb{Y}_t = C_t$. Since payout is $D_t = k\mathbb{Y}_t$, we have $\Delta d_{t+1} = \Delta c_{t+1} = \Delta \ln(\mathbb{Y}_{t+1})$. Putting this together, the model-implied value for the wealth ratio $-cay_t$ can be shown to be proportional to the log price-payout ratio, pd_t, plus a constant, that is, $-cay_t = kpd_t + const$.[22]

Summarizing, the model implies the following AP relations:

1. SDF:

$$m_t = \log(\delta) - \sigma_p \Delta d_t + \vartheta_{p,t-1}. \tag{10}$$

2. Price-payout ratio:

$$\begin{aligned} pd_t = \kappa_0 + \mu + \mathbb{E}_t^p\left[m_{t+1} + \Delta d_{t+1} + \kappa_1 pd_{t+1}\right] \\ + .5\mathbb{V}_t^p\left[m_{t+1} + \Delta d_{t+1} + \kappa_1 pd_{t+1}\right]. \end{aligned} \tag{11}$$

3. Log Euler equation for bonds:

$$i_t - \mathbb{E}_t^p[\pi_{t+1}] = -\mathbb{E}_t^p[m_{t+1}] - .5\mathbb{V}_t^p[m_{t+1} - \pi_{t+1}] - lp. \tag{12}$$

4. Wealth ratio, $-cay$:

$$-cay_t = kpd_t + const. \tag{13}$$

---

[22] To keep the estimation tractable, the model abstracts from one aspect of the data here, namely, that $-cay_t$ and pd_t are not perfectly correlated. This is due to the simplifying assumption that the "capital" share of $\mathbb{Y}_t$ is a constant k. Future work could extend the analysis to allow the capital share to be time-varying along the lines of Greenwald, Lettau, and Ludvigson (2019), thereby breaking the perfect correlation.

5. Log excess stock market return:

$$er^D_{t+1} = r^D_{t+1} - (i_t - \pi_{t+1}) = \kappa_0 + \kappa_1 pd_{t+1} - pd_t + \Delta d_{t+1} + \mu - (i_t - \pi_{t+1}). \tag{14}$$

Macro Dynamics. Macroeconomic dynamics are driven by a set of equations similar to those commonly featured in New Keynesian models, but with two distinctive features: sticky expectations about inflation and output and regime changes in the conduct of monetary policy.[23] In keeping with New Keynesian models, we assume that real variables grow nonstochastically along a balanced growth path and we write all equations in the macro block in terms of detrended real variables. Hereafter, detrended variables are denoted with a tilde, for example, $\widetilde{\ln(\mathbb{Y}_t)} \equiv \widetilde{y}_t$ denotes detrended log real output.

As in prototypical New Keynesian models, macroeconomic dynamics satisfy a loglinear Euler equation. In our setting this Euler equation is driven by the behavior of an average household referred to as the "macro agent." The macro agent can be considered typical of a household in the general population that holds small amounts of wealth in the form of nominal bonds and no equity. She consumes a fixed fraction $(1-k)$ of $\mathbb{Y}_t$, so that log detrended consumption growth of the macro agent is $\Delta\widetilde{y}_{t+1}$ and the linearized Euler equation takes the form

$$\widetilde{y}_t = \mathbb{E}^m_t(\widetilde{y}_{t+1}) - \sigma\left[i_t - \mathbb{E}^m_t(\pi_{t+1}) - r_{ss}\right] + f_t, \tag{15}$$

where i_t is the short-term nominal interest rate, $\mathbb{E}^m_t(\pi_{t+1})$ is the subjective expected inflation of the macro agent, r_{ss} is the steady-state real interest rate, and f_t is a demand shock, which 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 is a positive parameter.

We introduce two equations for inflation and the nominal interest rate rule. Inflation dynamics are described by the following equation, which takes the form of a New Keynesian Phillips curve:

$$\pi_t - \overline{\pi}_t = \beta\mathbb{E}^m_t[\pi_{t+1} - \overline{\pi}_t] + \kappa\left[\widetilde{y}_{t-1} - \widetilde{y}^*_{t-1}\right], \tag{16}$$

where $\overline{\pi}_t$ denotes the perceived long-term value of inflation that depends on the agent's information $\mathbb{I}_t$. We discuss how macro expectations are formed below. The coefficients β and κ are positive parameters and the variable $\widetilde{y}^*_t$ denotes the natural level of detrended output. Thus, $\widetilde{y}_{t-1} - \widetilde{y}^*_{t-1}$ is the output gap at time $t-1$. We assume an AR(1) process for $\widetilde{y}^*_t = \rho_{y^*}\widetilde{y}^*_{t-1} + \sigma_{y^*}\varepsilon_{y^*}$, $\varepsilon_{y^*} \sim N(0,1)$.

The central bank obeys the following nominal interest rate rule:

$$\begin{aligned} i_t - \left(r_{ss} + \pi^T_{\xi_t}\right) = {} & \left(1 - \rho_{i,\xi_t}\right)\left[\psi_{\pi,\xi_t}\left(\pi_t - \pi^T_{\xi_t}\right) + \psi_{\Delta y,\xi_t}(\widetilde{y}_t - \widetilde{y}_{t-1})\right] \\ & + \rho_{i,\xi_t}\left[i_{t-1} - \left(r_{ss} + \pi^T_{\xi_t}\right)\right] + \sigma_i\varepsilon_i,\ \varepsilon_i \sim N(0,1). \end{aligned} \tag{17}$$

---

[23] Outside of these two distinctive features, macroeconomic dynamics are identical to those that arise from the prototypical New Keynesian model of Galí (2015), Chapter 3.

Note that the interest rate rule is written in deviations from the steady state conditional on being in a particular regime dictated by ξ_t. This means that once inflation reaches the desired target, the economy stabilizes around it, absent shocks.

An important feature of this interest rate policy rule, and a departure from the prototypical model, is that it allows for regime changes in the conduct of monetary policy. These manifest as regime shifts in the inflation target $\pi^T_{\xi_t}$ and in the activism coefficients ψ_{π,ξ_t} and $\psi_{\Delta y,\xi_t}$ that govern how strongly the central bank responds to deviations from the target and to economic growth. The rule also allows for potential regime shifts in the autocorrelation coefficient ρ_{i,ξ_t}. These coefficients are modeled with a Markov-switching process governed by the discrete random variable ξ_t, which is assumed to take one of two values, $\xi_t = H$ or $\xi_t = D$, corresponding to hawkish or dovish monetary policy. These shifts are modeled as exogenous and random. However, it is important to recognize that the labels hawkish or dovish do not imply that the model imposes constraints on the estimated values of parameters across the previously estimated regimes. Since we freely estimate the regime parameters under flat priors, the parameters could in principle show no shift across regimes or even shifts that go in the "wrong" direction with respect to the previously estimated hawkish and dovish regimes.

We interpret equations (15) through (17) as equilibrium dynamics and not a microfounded structural model. We consider an equilibrium in which bonds are in zero net supply in both the macro and AP blocks and thus there is no trade between the AP agent and the macro agent.

The macro agent's expectations about inflation are formed using an adaptive algorithm, following MN. The representative macro agent forms expectations about inflation using an autoregressive process, $\pi_t = \alpha + \phi\pi_{t-1} + \eta_t$, but must learn about the parameter α.[24] Each period, agents form a belief about α, denoted by α^m_t, that is updated over time. Updating affects not only beliefs about next period's inflation, but also beliefs about long-term trend inflation. Define perceived trend inflation to be $\lim_{h\to\infty} \mathbb{E}^m_t[\pi_{t+h}]$, which we denote by $\overline{\pi}_t$. Given the presumed autoregressive process, the Internet Appendix shows that $\overline{\pi}_t = (1-\phi)^{-1}\alpha^m_t$. This implies that expectations of one-step-ahead inflation are a weighted average of perceived trend inflation and current inflation,

$$\mathbb{E}^m_t[\pi_{t+1}] = \alpha^m_t + \phi\pi_t = (1-\phi)\overline{\pi}_t + \phi\pi_t. \tag{18}$$

We allow the evolution of beliefs about α^m_t and $\overline{\pi}_t$ to potentially reflect both an adaptive learning component as well as a signal about the central bank's inflation target. For the adaptive learning component, we follow evidence in MN that the University of Michigan SOC mean inflation forecast is well described

---

[24] In principle one could introduce learning about ϕ as well. We forgo doing this to keep the estimation tractable, since the most important learning aspects in the model involve those parameters such as α that bear most closely on trend inflation.

by a constant gain learning algorithm. For the signal component, we assume that beliefs could be shaped in part by additional information the agent receives about the current inflation target. This signal could reflect the opinion of experts (as in MN) or a credible central bank announcement. Combining these two components yields updating rules for α_t^m and $\overline{\pi}_t$ that are weighted averages of two terms:

$$\alpha_t^m = (1-\gamma^T)\underbrace{\left[\alpha_{t-1}^m + \gamma\left(\pi_t - \phi\pi_{t-1} - \alpha_{t-1}^m\right)\right]}_{\alpha_t^{mCG}} + \gamma^T\left[(1-\phi)\pi_{\xi_t}^T\right], \tag{19}$$

$$\overline{\pi}_t = \left(1-\gamma^T\right)\underbrace{\left[\overline{\pi}_{t-1} + \gamma(1-\phi)^{-1}(\pi_t - \phi\pi_{t-1} - (1-\phi)\overline{\pi}_{t-1})\right]}_{\overline{\pi}_t^{CG}} + \gamma^T\left[\pi_{\xi_t}^T\right]. \tag{20}$$

The first terms in square brackets, α_t^{mCG} and $\overline{\pi}_t^{CG}$, are the recursive updating rules implied by constant-gain learning, 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 current inflation target $\pi_{\xi_t}^T$. If $\gamma^T = 1$, the signal is completely informative and the agent's belief about trend inflation is the same as the inflation 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. Overall, perceived trend inflation is a weighted average of the trend implied by the constant-gain learning rule and the central bank's inflation target. A weight of less than one on the target could arise because the target is imperfectly observed or because central bank announcements about the target are not viewed as fully credible. Note that the parameter γ^T is closely related to the speed with which the agent learns about a new inflation target. Since γ^T is freely estimated, we can empirically assess the magnitude of this speed and its role in macroeconomic fluctuations.

Agents form expectations about detrended output using a simple backward looking rule:

$$\mathbb{E}_t^m(\widetilde{y}_{t+1}) = \varrho\widetilde{y}_{t-1}. \tag{21}$$

Unlike inflation, agents do not perceive a moving mean for detrended output. This assumption is consistent with the equilibrium of the model implying that the central bank cannot have a permanent effect on real activity. The Internet Appendix proves that monetary neutrality holds in the long run.

Using equations (18), (20), and (21), we substitute out $\mathbb{E}_t^m[\pi_{t+1}]$, $\overline{\pi}_t$, and $\mathbb{E}_t^m(\widetilde{y}_{t+1})$ in equations (15), (16), and (17) to obtain the following system of equations that must hold in equilibrium:

1. Real activity:

$$\widetilde{y}_t = \varrho\widetilde{y}_{t-1} - \sigma\left[i_t - \phi\pi_t - (1-\phi)\overline{\pi}_t - r_{ss}\right] + f_t. \tag{22}$$

2. Phillips curve:

$$\pi_t = \overline{\pi}_t + \frac{\kappa}{1-\beta\phi}\left[\widetilde{y}_{t-1} - \widetilde{y}^*_{t-1}\right]. \tag{23}$$

3. Monetary policy rule with changes in target:

$$\begin{aligned} i_t - \left(r_{ss} + \pi^T_{\xi_t}\right) = {} & \left(1 - \rho_{i,\xi_t}\right)\left[\psi_{\pi,\xi_t}\left(\pi_t - \pi^T_{\xi_t}\right) + \psi_{\Delta y,\xi_t}(\widetilde{y}_t - \widetilde{y}_{t-1})\right] \\ & + \rho_{i,\xi_t}\left[i_{t-1} - \left(r_{ss} + \pi^T_{\xi_t}\right)\right] + \sigma_i \varepsilon_i, \ \ \varepsilon_i \sim N(0,1). \end{aligned} \tag{24}$$

4. Law of motion for f_t:

$$f_t = \rho_f f_{t-1} + \sigma_f \varepsilon_f, \ \ \varepsilon_f \sim N(0,1). \tag{25}$$

5. Law of motion for $\widetilde{y}^*_t$:

$$\widetilde{y}^*_t = \rho_{y^*}\widetilde{y}^*_{t-1} + \sigma_{y^*}\varepsilon_{y^*}, \ \ \varepsilon_{y^*} \sim N(0,1). \tag{26}$$

6. Perceived trend inflation:

$$\overline{\pi}_t = \left[1 - \gamma^T\right]\Big[\overline{\pi}_{t-1} + \gamma(1-\phi)^{-1}(\pi_t - \phi\pi_{t-1} - (1-\phi)\overline{\pi}_{t-1})\Big] + \gamma^T \pi^T_{\xi_t}. \tag{27}$$

Investors understand the macro block, can observe equations (22) to (27) and take those dynamics as given. But investors form beliefs about the persistence of any observed regime shift in the monetary policy rule (24).

Investor Beliefs. We now describe how investor beliefs in the model evolve over time. This evolution is influenced by both a learning component and a "fading memory" component.

For the learning component, we assume that investors closely follow central bank communications and are therefore capable of observing when important shifts in the policy rule parameters have occurred. They are uncertain about how long any shift will last, however, and must therefore learn about its duration. This assumption may be motivated by observing that sophisticated financial market participants in the real world expend significant resources on "Fed watching." Moreover, for decades central banks have clearly telegraphed their intentions when they seek to change the stance of monetary policy, but have been comparatively vague about the length of time such a change will last. The Federal Reserve's FOMC statement of August 9, 2011, for example, announced that "economic conditions are likely to warrant exceptionally low levels for the FFR *at least* [emphasis added] through mid-2013." Similarly, the FOMC press release of September 16, 2020, stated "the committee will aim to achieve inflation moderately above 2 percent for *some time* [emphasis added] …" and expects to maintain "an accommodative stance" until "inflation expectations remain *well anchored* [emphasis added] at 2 percent." The emphasized words in these sentences are murky and explicitly convey uncertainty about the length of time such policy changes will last.

For the fading memory component, we assume that expectations are shaped most strongly by recently experienced data, motivated by evidence in MN and Malmendier and Nagel (2011).

To model these two aspects of investor beliefs, we combine Bayesian learning about the persistence of regime changes with distorted beliefs. The key elements of this specification are twofold. First, if a regime change occurs after many periods in a previous regime, the investor will at first be almost certain that the deviation is temporary. However, as she observes more and more periods in a row in which the new regime holds, she gradually updates her beliefs and increasingly views the deviation as likely to persist. Second, once the agent spends enough time in a particular regime, memory of past policy rules fades and she comes to believe that the existing policy stance is the new normal that will persist indefinitely. Since the true policy regime transition matrix is persistent but transitory, fading memory about past policy rules represents a distortion in beliefs whereby agents extrapolate too much from recent continuity in the policy stance. This overextrapolation implies that the investor will always be surprised whenever there is a switch to a new policy rule after many periods in a previous regime.

In the rest of this subsection, we provide the basic idea for how this is modeled. The methodology is an extension of Bianchi and Melosi (2016). Technical details on the evolution of beliefs within and across policy regimes, and on how the model is solved under these beliefs, are provided in the Internet Appendix.

First consider the true data-generating process (DGP) for the monetary policy rule, which we presume follows a two-state Markov-switching process controlled by the variable $\xi_t \in \{H, D\}$ with transition matrix $\mathbf{H}$. Let $\xi_t = H$ be the state characterized by hawkish policy parameters, and $\xi_t = D$ be the state characterized by dovish policy parameters. Denote the true DGP transition probability matrix $\mathbf{H}$ as

$$\mathbf{H} = \begin{bmatrix} p_{HH} & p_{HD} \\ p_{DH} & p_{DD} \end{bmatrix},$$

where p_{ij}, $i, j \in \{H, D\}$, is the probability of switching to regime j given that the state is currently in regime i.

To model the idea that agents must learn about the persistence of regime changes, we assume that agents believe regime shifts can be either long- or short-lasting. This can be accommodated by introducing the notion of the perceived regime process $\xi_t^p \in \{1, 2, 3, 4\}$, with four states. Specifically, two of the perceived regimes are characterized by hawkish monetary policy ($\xi_t = H$), while two of the perceived regimes are characterized by dovish monetary policy ($\xi_t = D$). Without loss of generality, we assume that regimes $\xi_t^p = 1$ and $\xi_t^p = 2$ belong to a hawkish block 1 associated with $\xi_t = H$, while regimes $\xi_t^p = 3$ and $\xi_t^p = 4$ belong to a dovish block 2 associated with $\xi_t = D$. In the hawkish block, $\xi_t^p = 1$ is perceived as a short-lasting hawkish regime, while $\xi_t^p = 2$ is perceived as a long-lasting hawkish regime. In the dovish block, $\xi_t^p = 3$ is perceived as a short-lasting dovish regime, while $\xi_t^p = 4$ is perceived as a long-lasting dovish regime. The perceived probabilities of moving across these regimes are summarized by the transition matrix

$$\mathbf{H}^p = \left[\begin{array}{cc|cc} p_{11} & 0 & 0 & p_{14} \\ 0 & p_{22} & p_{23} & p_{24} \\ \hline 0 & p_{32} & p_{33} & 0 \\ p_{41} & p_{42} & 0 & p_{44} \end{array}\right], \tag{28}$$

where p_{ij} denotes the probability of switching to regime j given that we are in regime i. Since $\xi_t^p = 1$ is the perceived short-lasting hawkish regime, while $\xi_t^p = 2$ is the perceived long-lasting hawkish regime, we have $p_{22} > p_{11}$ by definition. Analogously, since $\xi_t^p = 3$ is the perceived short-lasting dovish regime and $\xi_t^p = 4$ is the perceived long-lasting dovish regime, we have $p_{44} > p_{33}$. To capture the idea that agents eventually "forget" about previous policy regimes once they spend enough time in a regime, we set $p_{44} = p_{22} = 0.999$. This implies that once agents believe they are in a long-lasting regime of either type, they come to view that regime as persisting almost indefinitely.[25]

Since the AP agent knows the structure of the macro block and can observe both the endogenous variables and the shocks at time t, she can also determine which set of policy parameters is in place at each point in time. That is, she can back out the history $\{\xi_t, \xi_{t-1}, \ldots\}$ of policy regimes and the block (dovish or hawkish) in place at time t. However, agents cannot exactly infer the realized perceived regime ξ_t^p, because the regimes within each block share the same policy rule parameter values governed by ξ_t. Thus, after a switch to a new policy regime, agents must learn about which element (short- or long-lasting) of the block they are actually in.

Suppose that the economy is initially in a state where the agent's perceived probability that she is in the long-lasting hawkish regime $\xi_t^p = 2$ is unity. If policymakers then start conducting dovish monetary policy ($\xi = D$), investors initially believe that this likely represents a temporary deviation from the $\xi_t^p = 2$ regime. This idea is captured by the conditions $p_{23} > p_{24}$, $p_{32} > 0$. However, because $p_{44} > p_{33}$, if the dovish regime persists long enough, the agent's perceived posterior probability that she is in a long-lasting dovish regime goes to unity. There are symmetric restrictions in the second block, corresponding to $p_{41} > p_{42}, p_{14} > 0$. Note that the purpose of the perceived short-lasting regimes is merely to model the idea that once investors perceive they are in a long-lasting regime of one type (hawkish or dovish), deviations from that policy rule might initially be viewed as transitory. Thus, we rule out transitions from a perceived short-lasting regime of one type to a short-lasting regime of the opposite type ($p_{31} = p_{13} = 0$) and transitions from a long-lasting regime of one type to a short-lasting regime of the same type ($p_{21} = p_{43} = 0$).

The fading memory distortion is captured by specifying $p_{22} > p_{HH}$ and $p_{44} > p_{DD}$. That is, once the agent spends enough time in a regime, she believes

---

[25] We rule out setting this probability to unity, since without further assumptions it would not be obvious how to model the evolution of investor beliefs when a shift out of the perceived long-lasting regime inevitably occurs.

that the regime will continue virtually indefinitely even though in reality it is persistent but transitory, so any switch out of a perceived long-lasting regime will be a surprise. This distortion leads the agent to eventually overstate the true persistence of policy regimes.

More generally given arbitrary initial beliefs, the restrictions above on the perceived transition matrix $\mathbf{H}^p$ will have implications for how beliefs evolve over time. The Internet Appendix gives recursive formulas for the perceived state probabilities that are history dependent.

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 equations (10) to (14) in the AP block are satisfied, equations (22) to (27) in the macro block are satisfied, and investors beliefs about the persistence of policy regimes are characterized by Bayesian updating about a perceived Markov-switching process with transition matrix (28), under the parameter restrictions given in the previous subsection.

B. Model Solution and Estimation

The model is solved in two steps. First, we solve for the macro dynamics. This returns an MS-VAR in the macro block state vector $S_t = [y_t, y_t^*, \pi_t, i_t, r_t^*, \overline{\pi}_t, f_t]'$. Second, conditional on this solution and on the probability assigned by the AP agent to moving across regimes, we derive the evolution of asset prices. This second step takes the MS-VAR law of motion for the macroeconomy as an input and combines it with the equilibrium AP relations (10) to (14), conditional on the law of motion for agents' beliefs outlined above. The final solution for all variables (macro and asset block) takes the form of MS-VAR in the augmented state space $\widetilde{S}_t = [S_t, m_t, pd_t, \mathbb{E}_t^p(m_{t+1}), \mathbb{E}_t^p(pd_{t+1})]$.

To estimate the model, we exploit the block structure of the solution to take a two-step approach. First, we use Bayesian methods to estimate the macro block by combining the MS-VAR solution for S_t with an observation equation. As data, we use four observable series: real per-capita gross domestic product (GDP) growth, inflation, the nominal FFR, and the mean of inflation expectations from the SOC. Since we have only three shocks to match four observable variables, we allow for observation errors on all variables. Second, conditional on the estimated parameter values from the macro block, the AP block parameters are chosen to minimize the sum (over t) of squared deviations between the model-implied cay_t and the observed series, cay_t^{MS}. Using an objective function penalty, we also require the AP block parameters to return a sizable equity premium. This two-step approach keeps the estimation tractable in the face of both regime shifts in monetary policy and history-dependent beliefs that are part of the AP block.

By using SOC data on inflation expectations, we ask the model to generate realistic behavior for inflation expectations. Specifically, we map the perceived law of motion of inflation into the Michigan survey. Below we show that the model-implied inflation expectations track their empirical counterparts well.

Parameter uncertainty is characterized using a random walk Metropolis–Hastings algorithm. The parameters of the policy rule, $\pi^T_{\xi_t}$, ψ_{π,ξ_t}, $\psi_{\Delta y,\xi_t}$, and ρ_{i,ξ_t}, are permitted to switch between two regimes according to a Markov-switching process. Since we are interested in understanding the connection between the previously estimated dovish/hawkish regimes, short-term interest rates, asset valuations, and return premia, we force the regime sequence for the policy rule parameters to correspond to the estimated sequence for α_{ξ_t} and r_{ξ_t} reported in Table III. Importantly, however, the parameters characterizing the policy regimes as well as the transition matrix are freely estimated.[26] Thus, there is no implication from this procedure that the parameters of the policy rule must necessarily show evidence of structural change. Moreover, since we freely estimate the parameters of the policy regime under flat priors, there is nothing in the model estimation that restricts the low- (high-) mps subperiods to coincide with parameters of the interest rate rule that would imply relatively accommodative (restrictive) monetary policy.

The sample spans the period 1961:Q1 to 2017:Q3, in line with our estimates for the regimes in the means of *cay* and the *mps*. We use the full sample of data, including observations from the ZLB period. The Internet Appendix shows that our findings on the long-lasting real effects of changes in the conduct of monetary policy are robust to replacing the FFR with either an estimated shadow rate or the one-year Treasury bill rate. The reason is that the policy rule regime changes that we uncover are not mainly tied to the ZLB period.

The Internet Appendix provides a detailed description of the data, model solution, and estimation.

C. Model Estimation Results

C.1. Parameter and Latent State Estimates

Table V reports the prior and posterior distributions for the macro block model parameters. For the policy rule parameter estimates for $\pi^T_{\xi_t}$, ψ_{π,ξ_t}, $\psi_{\Delta y,\xi_t}$, and ρ_{i,ξ_t}, where we use flat priors, a key finding is that the previously estimated regime subperiods (given in Table III) are associated with quantitatively large changes in the estimated policy rule. Specifically, the hawkish high-mps regime is characterized by what we refer to as a hawkish monetary policy rule with lower inflation target $\pi^T_{\xi_t}$ and strong activism ψ_{π,ξ_t} against deviations of inflation from the target relative to activism $\psi_{\Delta y,\xi_t}$ on growth. The dovish, low-mps regime is characterized by a dovish monetary policy rule with an inflation target that is comparatively higher and an activism against inflation that is significantly lower. In fact, for the dovish, low-mps regime, the 90% credible set for ψ_{π,ξ_t} includes one, the threshold generally associated with the Taylor (1993) principle, which prescribes that the central bank should raise nominal rates more than one-for-one in response to deviations of inflation from target,

[26] We use the regime sequence $\hat{\xi}^T = \{\hat{\xi}_1, \ldots, \hat{\xi}_T\}$ that is most likely to have occurred, given our estimated posterior mode parameter values for $\boldsymbol{\theta}$. See the Internet Appendix for details.

Table V

Macrofinance Model: Parameter Estimates and Priors

This table reports the posterior mode, mean, and 90% credible sets for the model parameters of the model. Prior distributions are denoted as follows: N stands Normal, G for Gaussian, B for Beta, and U for Uniform, where $Para_1$ and $Para_2$ refer to hyperparameters of the prior. For the Beta, Normal, and Gaussian distributions, the first parameter and second parameter correspond to the mean and standard deviation, respectively. For the Uniform distribution they correspond to the lower and upper bounds. The last four rows report the standard deviations of the observation errors. The sample spans the period 1961:Q1 to 2017:Q3.

	Mode	Mean	5%	95%	Type	Para_1	Para_2
π_H^T	0.8516	0.8411	0.7038	0.9642	U	0	10
$\psi_{\pi,H}$	2.3164	2.8146	1.9184	4.2444	U	0	10
$\rho_{i,H}$	0.8913	0.9080	0.8597	0.9494	B	0.5	0.2
$\psi_{\Delta y,H}$	2.6387	3.6663	1.9532	6.2615	U	0	10
π_D^T	2.8794	2.9040	2.7017	3.1626	U	0	10
$\psi_{\pi,D}$	1.1089	1.1146	0.8266	1.4120	U	0	10
$\rho_{i,D}$	0.8978	0.9264	0.8579	0.9804	B	0.5	0.2
$\psi_{\Delta y,D}$	1.2320	2.6661	0.8990	6.5637	U	0	10
γ	0.0017	0.0019	0.0008	0.0032	B	0.05	0.02
γ^T	0.0132	0.0131	0.0110	0.0152	B	0.2	0.1
σ	0.7970	1.1462	0.5406	2.0439	G	2	1
ϱ	0.9062	0.9008	0.8048	0.9696	B	0.9	0.05
β	0.7696	0.7156	0.5270	0.8909	B	0.8	0.1
κ	0.0343	0.0317	0.0143	0.0520	G	0.4	0.2
ρ_d	0.7589	0.8208	0.6731	0.9368	B	0.5	0.2
ρ_{y^*}	0.9457	0.9177	0.8474	0.9695	B	0.5	0.2
ϕ	0.8057	0.8041	0.7924	0.8149	B	0.5	0.2
r_{ss}	0.2540	0.3002	0.1210	0.5581	G	0.5	0.25
Δy_{ss}	0.3738	0.4126	0.3427	0.4912	G	0.5	0.2
σ_d	0.6033	0.6569	0.5431	0.7987	U	0	10
σ_i	0.1865	0.1950	0.1782	0.2153	U	0	10
σ_{y^*}	2.7349	3.7875	2.1213	7.4099	U	0	10
$\sigma_{oe,\Delta GDP}$	0.2897	0.2857	0.2339	0.3366	U	0	10
$\sigma_{oe,INFL}$	1.2294	1.2514	1.1557	1.3547	U	0	10
$\sigma_{oe,FFR}$	0.0000	0.0006	0.0000	0.0015	U	0	10
$\sigma_{oe,EXP}$	0.0686	0.0696	0.0517	0.0853	U	0	10

thereby raising the real rate and reducing inflationary pressure. The activism coefficient $\psi_{\Delta y,\xi_t}$ for output growth and the autoregressive parameter ρ_{i,ξ_t} are more similar across the two regimes.

These findings indicate that the policy rule parameters shifted to values consistent with restrictive monetary policy in 1978:Q4 around the time of Volcker's appointment, consistent with an older empirical literature (e.g., Clarida, Gali, and Gertler (2000)). However, the results here show that, starting in 2001:Q4, parameters shifted back to values consistent with accommodative monetary policy. With the exception of a brief interlude from 2006:Q2 to 2008:Q2, the

Table VI

Regime Shifts in the Policy Rule

This table reports the posterior mode, mean, and 90% credible sets for the difference between the monetary policy rule parameters across the two regimes, defined as hawkish (H) and dovish (D). The sample spans the period 1961:Q1 to 2017:Q3.

	Mode	Mean	5%	95%
$\pi_H^T - \pi_D^T$	−2.0278	−2.0629	−2.4647	−1.7533
$\psi_{\pi,H} - \psi_{\pi,D}$	1.2074	1.7001	0.6990	3.1387
$\rho_{i,H} - \rho_{i,D}$	−0.0065	−0.0184	−0.0911	0.0576
$\psi_{\Delta y,H} - \psi_{\Delta y,D}$	1.4067	1.0001	−3.3289	4.4279

relatively dovish policy rule has remained in place since, to the end of our sample in 2017:Q3.

Shifts in the policy rule parameters across the two regimes are large in magnitude. Table VI reports the posterior distribution for the differences in the parameters across regimes. The mode of the distribution of the difference in the quarterly $\pi_{\xi_t}^T$ across the two regimes is around 2%. This large value implies a difference in the annualized inflation target across regimes of almost 8%. The 90% credible set also indicates strong statistical evidence in favor of a quantitatively large difference in the inflation target across the two regimes. Similarly, the posterior distribution for the difference in the inflation activism coefficient ψ_{π,ξ_t} is centered on 1.2 with posterior credible sets that are bounded well away from zero, confirming evidence of a change in the degree of activism aimed at stabilizing inflation around the desired target. Finally, the posterior distributions for the difference in activism $\psi_{\Delta y,\xi_t}$ on growth and in the autoregressive parameter ρ_{i,ξ_t} show only weak evidence of change in these parameters. To summarize, there is strong evidence of sizable shifts across the previously estimated regimes in the relative importance of inflation and economic growth in the policy rule and a large shift in the tolerable level of inflation.

For the nonpolicy-rule parameters, it is worth emphasizing that the estimates imply a very high level of inertia in inflation expectations. The constant gain parameter γ, controlling the speed with which beliefs about long-term inflation are updated with new information on inflation, is estimated to be low. Furthermore, the parameter γ^T, controlling the extent to which perceived trend inflation is influenced by the central bank target, is estimated to be very low. Taken together, these findings imply that agents revise their beliefs about long term inflation only slowly over time and mostly based on past realizations of inflation rather than on the inflation target itself.

Figure 5 shows that the model-implied series track their empirical counterparts quite well. In general, observation errors play little to no role in the dynamics of the model-implied series. The model generates a plausible mean and volatility for the real interest rate. The dynamics of the model-implied series for one-step-ahead inflation expectations tracks the SOC series virtually without error. This is relevant since inflation expectations play a key role

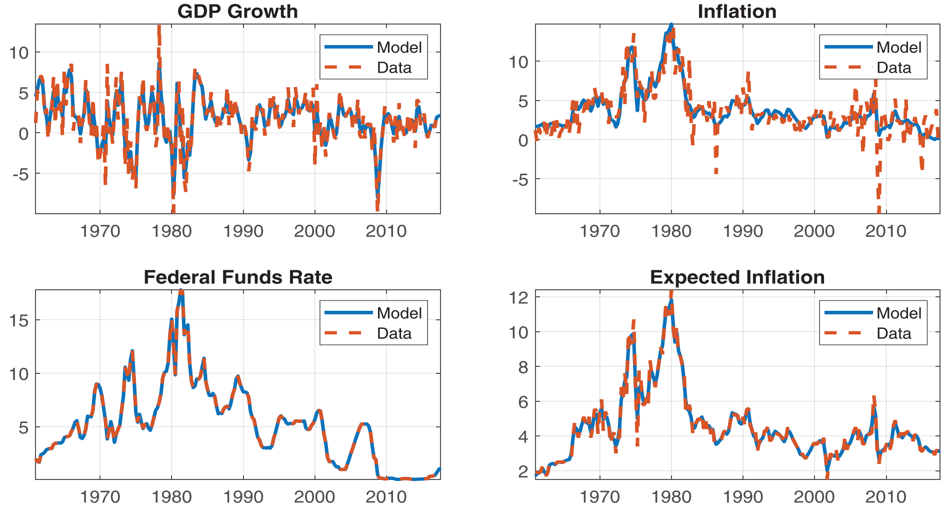

Figure 5. Macroeconomic series and their filtered counterparts. This figure plots the model-implied series and the corresponding observed series. Expected inflation comes from the Michigan Survey of Consumers. The difference is due to observation errors. The sample spans 1961:Q1 to 2017:Q3. (Color figure can be viewed at wileyonlinelibrary.com)

in the model's predictions, as we show below. Figure 5 underscores the extent to which those predictions are predicated on empirically relevant inflation expectations. The other model-implied series also track their empirical counterparts fairly closely. In particular, since the model fits the FFR and inflation expectations well, it also fits the real rate as measured by the difference in the two. For inflation, there are a handful of high-frequency spikes that the model is not well positioned to capture. A richer model could account for these spikes, but since the scope of our investigation is a study of lower frequency shifts in the policy rule, we do not view this as an important drawback of the framework.

A comment is in order about the estimated values for $\pi^T_{\xi_t}$ shown in Table VI. Although this parameter plays the role of an "inflation target" in the interest rate rule, unlike traditional New Keynesian models, $\pi^T_{\xi_t}$ is not a value to which true inflation and inflation expectations in the model necessarily tend in the long run. This is because the model here differs in two ways from the traditional New Keynesian models: macro expectations are strongly backward-looking and the policy rule parameters are not constant but instead vary over time. This combination implies that the inflation target can deviate substantially from actual inflation and inflation expectations for an extended period of time. For example, consider the value for $\pi^T_{\xi_t}$ under the dovish policy rule ($\pi^T_{\xi_t} = \pi^T_D$), in the postmillennial dovish subperiod. In this case π^T_D is quite high, yet the model matches the observed low values for both inflation and inflation expectations over the extended subperiod well (see Figure 5), and neither the model-implied inflation nor inflation expectations tend toward the estimated

Table VII

Parameters of the Asset Pricing Block

The parameters are chosen to minimize the distance between the fluctuations in *cay* implied by the model as a result of regime changes and the actual cay^{MS}. The values for lp and equity premium are annualized log units. The sample is quarterly and spans the period 1961:Q1 to 2017:Q3.

Parameter	Value	Parameter	Value
k	0.4506	p_{11}	0.6750
σ_p (fixed)	3	p_{22} (fixed)	0.9990
δ	0.9329	p_{33}	0.7331
lp	5.8%	p_{44} (fixed)	0.9990
equity premium	5.5%	$p_{23}/(p_{23}+p_{24})$	0.9864
		$p_{41}/(p_{41}+p_{42})$	0.9999

value for π_D^T, which is 2.9% at a quarterly rate. This result is not attributable to the two-state Markov-switching specification, which forces the early-dovish (1960s and early 1970s) and late-dovish (postmillennial) subperiods to share the same policy rule parameter values. Additional results given in the Internet Appendix indicate that the early-dovish and late-dovish subperiods both rationalize a high value for $\pi_{\xi_t}^T$, but for different reasons. In the early subperiod, both observed inflation and inflation expectations were high, which the model rationalizes with a high value for $\pi_{\xi_t}^T$. In the late subperiod, observed inflation is much lower and trending down, but expected inflation remains relatively elevated, causing a gap to open up between the two. This gap is also rationalized in the model by a high value for $\pi_{\xi_t}^T$.

To interpret this result, note that the postmillennial subperiod is characterized by negative demand shocks in the model (to account for the two sharp recessions), subsequent sluggish economic growth, and sustained periods of low and even negative inflation. Yet at the same time, data on inflation *expectations* remain elevated by comparison. The model reconciles this set of facts by indicating that monetary policy was extremely dovish, as exhibited by a high value for π_D^T. In the real world, central banks have additional policy tools for implementing accommodative monetary policy, such as forward guidance and quantitative easing, two tools that were employed in the postmillennial subperiod of our sample. These additional channels are absent from the stylized model, but manifest as a high value for the inflation target policy parameter $\pi_{\xi_t}^T$.

Table VII reports the parameter values for the AP block. The procedure implies a modest relative risk aversion coefficient of $\hat{\sigma}_p = 3$. The equity premium implied by the model parameters is 5.5% at an annual rate, which is slightly lower than the liquidity premium component $\widehat{lp} = 5.8\%$ at an annual rate, implying that the risk premium component of the equity premium is slightly negative. To understand why a (small) negative risk premium arises, first note that the overall risk premium is a weighted average of the risk premia in the dovish and hawkish regimes, where the weights are pinned down by the ergodic regime probabilities. Premia are small in absolute terms in both regimes. However, in the hawkish regime, the central bank is very “active,” so when the economy enters a recession (boom), it cuts (raises) real interest rates aggressively because inflation decreases (increases). This drives the SDF up (down) strongly due to a large change in $\vartheta_{p,t}$, as seen from equation (8). Although recessions (booms) also cause expected future cash flows to fall (rise), the aggressive manipulation of the real interest rate in the hawkish regime implies that these cash flow effects are outweighed by discount rate effects that move stock returns in the opposite direction. Thus, in the hawkish regime, central bank actions effectively mean that stocks provide insurance against cash flow fluctuations, leading to a negative risk premium (but quantitatively small in absolute terms). In the dovish regime, by contrast, the Fed is less active and the standard cash flow effects on stock prices outweigh the discount rate effects, leading to a small positive risk premium.[27] Because the estimation chooses parameters to target the observed equity premium, the resulting estimate of $\widehat{lp} = 5.8\%$ is high relative to estimates in, for example, Krishnamurthy and Vissing-Jorgensen (2012). This could be addressed in the existing setup by modeling payouts as levered, which would increase the risk premium component of the equity premium and reduce $\widehat{lp}$. We choose not to do so in this paper to keep the model simple, focusing on fluctuations in the equity premium arising from distorted beliefs about monetary policy.

Table VII also reports parameters of the perceived transition matrix. The estimated perceived probability of switching out of a long-lasting regime of one type into a short-lasting regime of the other type is close to unity in both cases, that is, $p_{23}/(p_{23} + p_{24}) = 0.986$ and $p_{41}/(p_{41} + p_{42}) = 0.9999$. This implies that any switch to a dovish (hawkish) policy rule when the agent had previously believed she was in a long-lasting hawkish (dovish) regime is initially perceived as a temporary deviation from the old rule.

Evolution of Beliefs. Figure 6 shows the estimated model’s implications for the evolution of investor beliefs about future monetary policy over our sample, under the assumption that the agent begins the sample believing with probably one that she is in the short-lasting dovish regime. The left panel plots the perceived probability at each point in time of being in the long-lasting hawkish regime (blue solid line) and the long-lasting dovish regime (red dashed line). The right panel plots the perceived probability of being in a hawkish regime (either short- or long-lasting) at some future horizon $t + h$, where $h = 1, 4$, or 80 quarters in the future.

In the left panel we can see that, from the beginning of the sample onward, it takes several years of continuously observing dovish monetary policy before the perceived probability of being in a long-lasting dovish regime is close to one. Likewise, as the economy switches into the hawkish policy rule under Volcker, the agent initially places very low probability on the change persisting. This can be seen in the right panel, where immediately after the change the

---

[27] Inflation risk also contributes little to the equity premium since the conditional variance and covariance terms in (9) involving inflation are small due to the high level of inertia in macroagent inflation expectations.

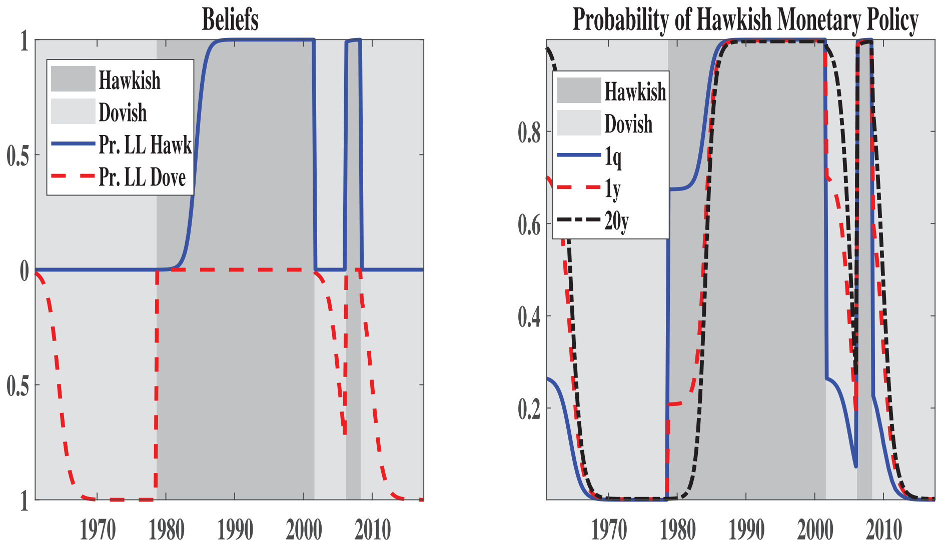

Figure 6. Evolution of investor beliefs under learning. The left panel plots the perceived probability of currently being in the long-lasting hawkish regime (blue solid line in the top scale) or the long-lasting dovish regime (red dashed line in the lower scale). The right panel plots the perceived probability of being in either hawkish (long- or short-lasting) regime at $t+h$, where $h = 1, 4$, or 80 quarters in the future. We initialize the asset pricing agent's beliefs in 1960:Q1 assuming that she assigns $\Pr \approx 1$ to being in the short-lasting dovish regime. The sample spans 1961:Q1 to 2017:Q3. (Color figure can be viewed at wileyonlinelibrary.com)

perceived probability of being in the hawkish regime in one year's (20 year's) time is less than 0.3 (effectively zero). The agent only eventually comes to see the hawkish policy rule as a long-lasting feature after observing several years of continuously restrictive policy. Beliefs about the long term are therefore "sticky,"they change only when agents become convinced that monetary policy has experienced a structural break. By contrast, short-term expectations about future monetary policy can change quickly, as agents take into account the possibility of a temporary deviation from the current policy framework. This implies that asset valuations can experience sudden but modest jumps in response to changes in short-term expectations, followed by further changes as investors revise the probability of remaining in the new policy framework.

At the same time, when regime shifts are observed more frequently, even expectations about the long term can move quickly, as occurs right after the switch out of the long hawkish subperiod from 1978:Q4 to 2001:Q3. When the policy rule switches back to hawkish less than five years later in 2006:Q2, the perceived probability of being in a hawkish regime 20 years later jumps to unity almost immediately. Due to the history dependence in the evolution of beliefs, investors in 2006:Q2 still have a strong recent memory of the previously lengthy hawkish regime and quickly perceive its return.

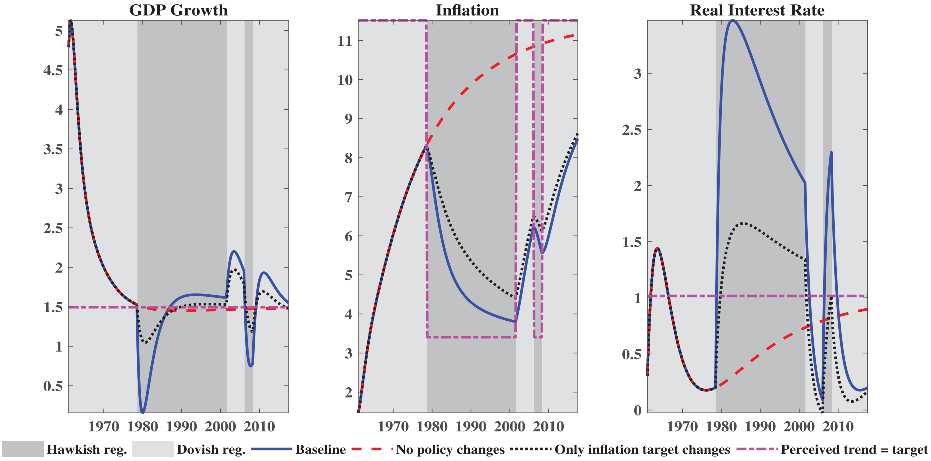

Figure 7. The role of changes in the monetary policy rule and adaptive expectations. The blue line corresponds to fluctuations generated by changes in both the target and the slope coefficients of the policy rule. The red dashed line assumes that monetary policy starts under the dovish regime and no regime changes occur. The black dotted line assumes that changes in the target occurred, but that the slope coefficients in the interest rate rule remain fixed as in the dovish regime. Finally, the magenta dashed-dotted line shows a counterfactual in which the policy rule shifts but the macro agent's perceived trend inflation equals the central bank's target. The dovish regime is defined by a high target π^T and low activism against deviations from π^T. The hawkish regime has a low π^T and high activism against deviations from π^T. The sample spans 1961:Q1 to 2017:Q3. (Color figure can be viewed at wileyonlinelibrary.com)

C.2. Conduct of Monetary Policy, the Real Interest Rate, and Asset Valuation

In this section, we investigates how changes in the conduct of monetary policy affect the real interest rate and asset valuations over our sample. To do so, we consider a number of simulations that isolate the effects of regime changes in the conduct of monetary policy. All figures present the values of the variables at the estimated posterior mode parameter values.

Monetary Policy and Macroeconomy Over the Sample. Figure 7 shows results from a simulation in which the observables and estimated state vector are set to their values at the beginning of our sample with all Gaussian shocks shut down. Thus, the only source of variation in the variables plotted in the figure arises from changes in the conduct of monetary policy, that is, from changes in the policy rule parameters.

For the baseline model, the portion of movements in output growth, inflation, and the real interest rate over our sample that can be directly associated with changes in the policy rule are shown in blue (solid) lines in Figure 7. The figure also considers three counterfactual simulations. The orange (dashed) line assumes that monetary policy starts under the dovish regime and that no subsequent regime change occurs. The black (dotted) line assumes that changes in the target occurred, but that the slope coefficients in the policy rule always remain as they are in the dovish rule. The magenta (dashed-dotted) line assumes that the macro agent's perceived trend value for inflation coincides in every period with the inflation target $\overline{\pi}_t = \pi^T_{\xi_t}$, corresponding to the case in which $\gamma^T = 1$. This value is highly counterfactual, since the estimated value $\hat{\gamma}^T = 0.013$ implies that expectations of trend inflation as implied by the SOC data place virtually no weight on the inflation target and instead are driven mostly by the constant gain adaptive expectations rule.

A series of noteworthy results emerge from Figure 7 . First, if instead of switching to a hawkish stance under Volcker the central bank had maintained the dovish policy rule throughout our sample, the economy would not have experienced the decrease in inflation that occurred in the early 1980s. Instead, inflation would have continued to increase. What is more relevant and less obvious is the behavior of the real FFR. The right panel of Figure 7 shows that changes in the conduct of monetary policy generate fluctuations in the real interest rate that last for decades. Comparing the estimated case with the orange dashed line that counterfactually assumes no policy rule changes in our sample, it is clear that the real FFR would have been substantially more stable had there been no changes in the monetary policy stance.

Second, Figure 7 shows that large, persistent swings in the real interest rate attributable to changes in the conduct of monetary policy were not solely the result of shifts in the inflation target-shifts in the activism coefficients also play a role. Comparing the baseline estimation (blue solid line) with the the counterfactual in which the inflation target changes but there are no accompanying changes in the activism coefficients (black dotted line), it is clear that the sharp increases in the real rate associated with Volcker would have been far smaller had the activism coefficients remained constant. A similar result holds in the short hawkish regime that precedes the Great Recession (2007:Q4 to 2009:Q2). Intuitively, since the hawkish regime exhibits both a lower inflation target and increased activism against deviations from the target, the real interest rate increases much more than it would have if only the inflation target had changed. The combination of the two contributed to sharp contractions in output growth during the recessions of 1980 and 1981 and during the Great Recession, as observed in the first panel. Without the concomitant shifts in the activism coefficients, both inflation and inflation expectations would have remained higher over the entire post-Volcker sample, as observed in the middle panel.

Third, the magenta (dashed-dotted) line of Figure 7 shows that the macro agent's highly adaptive expectations are crucial to understanding the long-lasting effects of regime changes in monetary policy. In the counterfactual economy where the perceived trend value for inflation coincides in every period with the inflation target, inflation jumps immediately to the new target whenever the policy stance changes, with no effect on the real interest rate. Inflation jumps in the counterfactual case because the central bank does not have to "convince" agents about the new inflation target. It is the interaction between changes in the anti-inflationary stance of the central bank and sticky macro expectations that generates long-lasting fluctuations in the real interest rate.

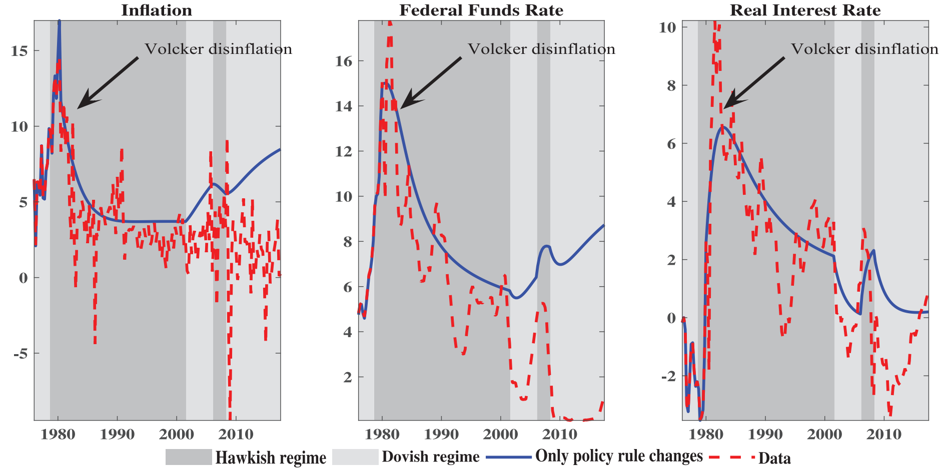

Figure 8. The Volcker disinflation. We start the economy as it was in 1980:Q1 and remove all Gaussian shocks that occurred after that period but keep the estimated regime sequence. The dashed line corresponds to the data. The real interest rate is computed as the difference between the federal funds rate (FFR) and expected inflation. Expected inflation is obtained based on the model solution. (Color figure can be viewed at wileyonlinelibrary.com)

The Secular Decline in Real Interest Rates. Figure 8 studies the implications of our model for how regime changes in the conduct of monetary policy have contributed to the secular decline in real interest rates observed since the early 1980s. For this purpose, we begin a simulation with the economy as it was in 1980:Q1, at the beginning of the Volcker disinflation, when inflation had reached its peak in our sample but before the peak in the real interest rate reached in 1981:Q3. To isolate the effects of changes in the monetary policy rule on the real interest rate under the Volcker disinflation and thereafter, we set all Gaussian shocks after 1981:Q1 to zero. These movements are shown in blue (solid) lines, and the actual values for each series are shown in red (dashed) lines.

The right panel of Figure 8 shows that the sharp run-up in real rates in the 1980s, and much of its decline since that time, can be attributed to changes in the conduct of monetary policy. Changes in the conduct of monetary policy do not track the higher frequency fluctuations in the real rate. For example, there is a sharp decline in the real rate that lasts for several years after the Great Recession. These fluctuations in the real rate are not associated with a shift in the policy rule parameters, but are instead due to a combination of the model's Gaussian shocks. By contrast, a substantial portion of the downward secular *trend* in real rates since the early 1980s is due to regime changes in the conduct of monetary policy. The peak of the real FFR in our sample is 10.22%, which occurs in 1981:Q3. Since that time, the real FFR has gradually trended downward, with the last observation in our sample equal to 0.56% in 2017:Q3. This represents a decline of 9.67%. According to our estimated model,

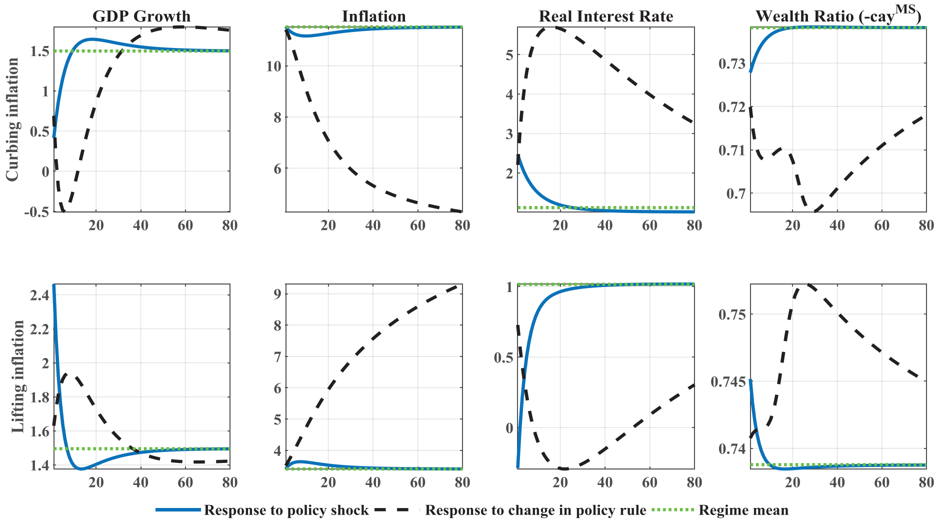

Figure 9. Regime changes versus policy shocks. Top row: curbing inflation. The economy is initially in the dovish regime. The blue solid line presents the evolution of the macro variables and the wealth ratio in response to a two-standard-deviation contractionary monetary policy shock and no regime change. The black dashed line presents the evolution of the macro variables and the wealth ratio in response to a regime change from the dovish regime to the hawkish regime. Bottom row: lifting inflation. The economy is initially in the hawkish regime. The blue solid line presents the evolution of the macro variables and the wealth ratio in response to a two-standard-deviation expansionary monetary policy shock and no regime change. The black dashed line presents the evolution of the macro variables and the wealth ratio in response to a regime change from the hawkish regime to the dovish regime. (Color figure can be viewed at wileyonlinelibrary.com)

regime changes in monetary policy generate a peak in the real FFR of 6.55% in 1983:Q1 and an end-of-sample value of 0.20% in 2017:Q3. This translates into a decline of 6.35%, or roughly two-thirds of the observed secular decline.

Regime Changes versus Policy Shocks. Figure 9 shows the implications of our model for monetary policy *shocks* versus monetary policy regime changes, using two sets of estimated impulse response functions. In the top row, we assume that the economy is initially in the dovish regime and consider the case of the monetary authority attempting to curb inflation. The blue solid line in the top row shows responses to a two-standard-deviation *contractionary* (i.e., positive) monetary policy shock and no policy rule regime change. The black dashed line in the top row shows responses to a regime change from the dovish to the hawkish regime, with all Gaussian shocks (including the monetary policy shock) set to zero. The figure shows the model's implications for the response of GDP growth, inflation, the real interest rate, and the log wealth ratio ($-cay_t$) to policy regime changes versus policy shocks. It is immediately evident that the effects of a regime change in the policy rule parameters are long-lived and last for decades, while those of monetary policy shocks are relatively short-lived, consistent with empirical evidence using observed monetary policy shocks (e.g., Christiano, Eichenbaum, and Evans (2005)). In response to a regime shift to hawkish policy, asset valuations, as measured by the log wealth ratio, fall and remain low for many years, while a contractionary monetary policy shock has negligible effects on valuations.

Because the model is nonlinear, the duration of these effects can differ depending on whether we begin in a dovish or a hawkish regime. In the lower row of Figure 9 we assume that the economy is initially in the hawkish regime and consider the case of the monetary authority attempting to lift inflation. The blue solid line shows the impulse responses to a two-standard-deviation *expansionary* monetary policy shock and no regime change in the conduct of monetary policy. The black dashed line shows responses to a regime shift from the hawkish to the dovish regime, with all Gaussian shocks set to zero. The effects on the real interest rate of a policy rule regime change in this case are even more long-lived than in the curbing-inflation case. The reason is that, under the dovish policy rule, the central bank responds less aggressively to fluctuations in inflation and output, as indicated by the smaller estimated activism coefficients ψ_{π,ξ_t} and $\psi_{\Delta y,\xi_t}$. Thus, when the central bank seeks to lift inflation as opposed to curb it, it does so more gradually, so the real interest rate and the log wealth ratio remain perturbed from their steady-state values for a longer period of time.

In either the lifting- or curbing-inflation case, the effects on the real rate attributable solely to regime changes in the conduct of monetary policy are extremely long lived, lasting more than 90 years in both cases, in sharp contrast to a monetary policy *shock*. Monetary policy shocks have short-lived effects because they cause inflation to move away from target and are always quickly stabilized, even in the dovish regime. By contrast, there is no reason for the central bank to stabilize an intentional change in the stance of monetary policy, so the extent to which regime changes in monetary policy persist in their real effects depends only on how quickly agents adapt their expectations about long-term inflation. Since our parameter estimates imply that agents' expectations adapt very slowly over time, changes in the conduct of monetary policy have effects that last for decades.

The Role of Investor Beliefs. What is the role of investor beliefs in the response of asset valuations to policy rule changes? To illustrate their role, Figure 10 plots impulse responses implied by the model to policy rule regime changes under different counterfactual simulations. The top row reports responses to a change from the dovish to the hawkish regime (curbing inflation), with all Gaussian shocks set to zero. The bottom row shows analogous responses to a change from hawkish to dovish (lifting inflation). The blue (solid) lines in all figures of both rows plot the responses in the baseline model. The red (dotted) line shows a counterfactual in which the AP agent knows the true policy rule transition matrix $\mathbf{H}$, a case we label "AP rational expectations." In this case there is no learning about the persistence of regime shifts and no fading memory distortion. The black (dashed) line labeled "No AP learning" shows a counterfactual that retains the fading memory distortion–implying that investors act as if persistent shifts in the policy rule will continue indefinitely—but we shut off learning about the persistence of

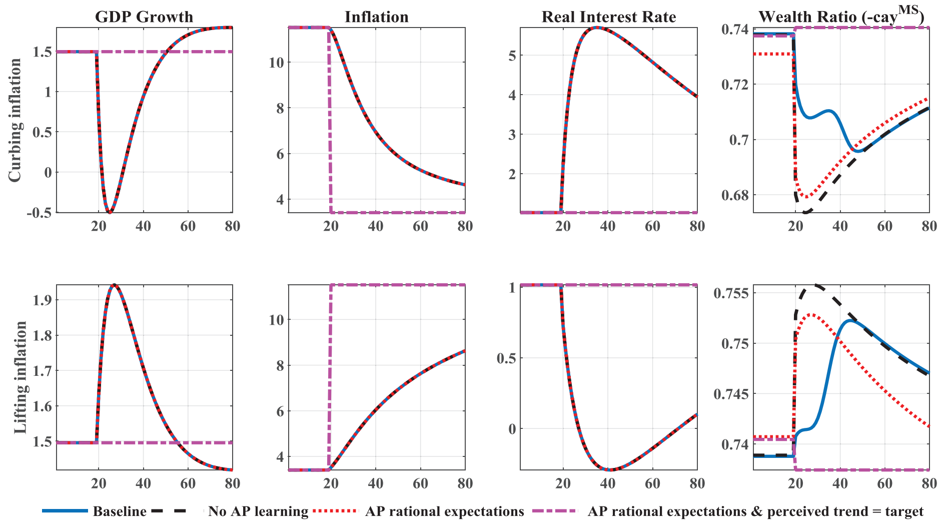

Figure 10. The role of AP learning and of macro stickiness. The blue solid line corresponds to the baseline model with the asset pricing (AP) agent learning about the probability of moving across regimes, overreaction of the AP agent about the persistence of regime changes, and adaptive expectations of the macro agent; the black dashed line shuts down learning of the AP agent; the red dashed line is the case in which the AP agent observes the true transition matrix of the Markov-switching process controlling policy rule regimes; the dotted-dashed magenta line shuts down both learning of the AP agent and adaptive expectations of the macro agent. Top row: curbing inflation. The economy is initially in the dovish regime and in period 20 moves to the hawkish regime. Lower row: lifting inflation. The economy is initially in the hawkish regime and in period 20 moves to the dovish regime. (Color figure can be viewed at wileyonlinelibrary.com)

regimes. The magenta (dashed-dotted) line is a counterfactual that combines AP rational expectations with the case in which the macro agent's perceived trend value of inflation, $\overline{\pi}_t$, coincides in every period with the inflation target, $\pi^T_{\xi_t}$.

The AP agent's beliefs play no role in the macro dynamics. Thus, the blue (solid), black (dashed), and red (dotted) responses for GDP growth, inflation, and the real interest rate in Figure 10 all lie on top of each other. By contrast, investor beliefs play a large role in the responses of asset valuations ($-cay_t$), as shown in the last column. A switch to hawkish (dovish) policy drives the wealth ratio down (up) as the real interest rate rises (falls). Because of learning, the initial jump is only the start of a gradual response and is followed by further changes in the wealth ratio as agents revise upward the probability of remaining in the new policy framework. Comparing the blue line to the red dotted line that corresponds to AP rational expectations, it is clear that valuation ratios in the baseline model initially underreact to the policy rule regime shifts. Under AP rational expectations, the wealth ratio jumps on impact to its maximal response in almost one period. The wealth ratio nonetheless moves smoothly back toward its steady-state value even under rational expectations, a reflection of the adaptive learning mechanism in the macro block that drives the persistent behavior of the real interest rate observed in the third column of Figure 10.

The role of overextrapolation can be seen by comparing the red dotted "AP rational expectations"line to the black dashed "No AP learning" line. Overextrapolation amplifies the response of the wealth ratio to regime changes in monetary policy, but it does not create gradualism in the response. Since the baseline model has both learning and overextrapolation, the baseline wealth ratio responds to regime shifts in the policy rule by initially underreacting but eventually overreacting vis-a-vis the case of AP rational expectations.

The magenta (dashed-dotted) line of Figure 10 combines AP rational expectations with $\overline{\pi}_t = \pi^T_{\xi_t}$ for all t. When expectations of the macro agent are not adaptive, regime changes in the policy rule have no effect on the real interest rate or real GDP growth, as noted above. By contrast, a shift to the hawkish policy rule slightly *in*creases the wealth ratio, while a shift to the dovish policy rule slightly *de*creases it. This result obtains because, although policy regime changes in this case have no effect on the first moments of real variables, they do affect second moments. A switch to the hawkish policy rule implies that the central bank more aggressively stabilizes real activity, which reduces the risk premium on equity and increases asset valuations. The opposite occurs in a switch to the dovish rule. Thus, without adaptive macro expectations about long-term inflation, the model cannot generate the right comovement of valuation ratios with the monetary policy regime sequence observed in the data, either qualitatively or quantitatively.

Monetary Policy and Asset Valuation over the Sample. Figure 11 shows the implications of the model for $-cay_t$ and mps over our sample. To facilitate a direct comparison with the data, this figure repeats the information from Figure 3, which plots the corresponding series $-cay^{MS}_t$ and mps_t from the data, along with horizontal lines that show the regime average values for these series. The figure also shows the component of the model-implied values for $-cay_t$ and the mps that we estimate are due solely to regime changes in the conduct of monetary policy, shown as black dashed-dotted lines. The magenta dotted lines correspond to the same components under the AP rational expectations counterfactual.

Figure 11 shows that fluctuations in the model-implied $-cay_t$ and mps due to regime changes in monetary policy fluctuate closely around the data regime average values for these series across the regime subperiods. Because the AP agent understands the macro block, she knows that persistent regime shifts in the conduct of monetary policy will generate persistent movements in the real interest rate. This leads to large swings in the price of long-duration assets as discount rates vary. The model-implied series show that learning about the persistence of regime changes can coexist with jumps at regime shift dates in the components of the wealth ratio and mps that are due to shifts in the policy rule, consistent with the Markov-switching specification. For the wealth ratio, the initial jump is smaller than its ultimate change due to learning. This implies that the full change in the wealth ratio after a regime switch can sometimes lag the full change in the mps, as it does, for example, after the

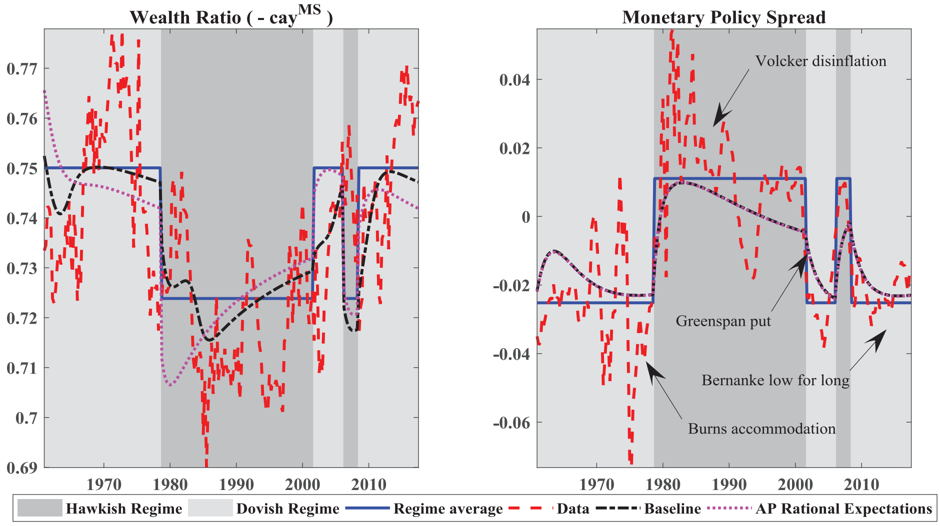

Figure 11. Wealth ratio and mps: Data and model. The figure reports the time series of the log wealth ratio and the monetary policy spread. The red dashed lines represent the data, the blue solid line represent the regime means, the black dashed-dotted lines represent the fluctuations that can be explained by regime changes in monetary policy under the baseline model, and the magenta dotted lines represent the fluctuations that can be explained by regime changes in monetary policy assuming that the asset pricing (AP) agent observes the true transition matrix of the Markov-switching process controlling changes in monetary policy. The sample spans 1961:Q1 to 2017:Q3. (Color figure can be viewed at wileyonlinelibrary.com)

switch to the first hawkish subperiod of the sample. Since both beliefs and the interest rate rule are history dependent, this does not necessarily happen after every switch.

Under the AP rational expectations counterfactual, the model-implied $-cay_t$ jumps after a regime switch to its final destination in almost one period, driven by the revision in expected real interest rates. Investors in this case realize that a prolonged period of high or low real interest rates will follow as the central bank tries to alter inflation in the face of highly adaptive macro expectations. Eventually, inflation adjusts and the wealth ratio reverts toward its steady-state value as the real interest rate reverts.

Summarizing the lessons from the previous two figures, we show that the large movements in the wealth ratio following monetary policy regime changes are the result of the interaction between two forces: (i) sticky macro agent expectations about inflation and (ii) revisions in investor expectations about future monetary policy. Without stickiness in inflation expectations, the model cannot generate persistent movements in the real interest rate that in turn trigger large fluctuations in the wealth ratio. Without investor learning about the persistence of regime shifts, the model produces implausibly large jumps in valuation ratios at regime shift dates as the AP agent immediately and fully revises her expectations. Without overextrapolation, the wealth ratio would

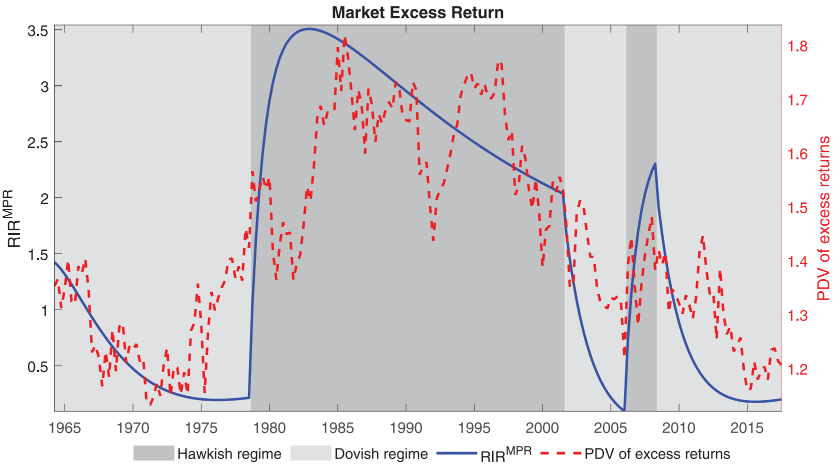

Figure 12. Return premia and policy rule changes. The figure reports the time series of the present discounted value (PDV) of expected stock market excess returns (dashed line, right axis) together with fluctuations of the real interest rate (RER) due to changes in the monetary policy rule (solid line, left axis). The sample is quarterly and spans the period 1961:Q1 to 2017:Q3. (Color figure can be viewed at wileyonlinelibrary.com)

not respond to regime changes by overshooting the case in which the agent is fully aware of the underlying transition matrix. This latter element of investor beliefs has important implications for the present discounted value of equity return premia. We turn to this next.

C.3. Monetary Policy and the Equity Premium Over Time

Figure 12 plots the estimated pdv_t of forecasted excess returns (return premia) for the stock market analyzed in Section II.B (red dashed line, right axis) together with the estimated component of the real interest rate due to regime changes in the monetary policy rule (solid line, left axis). Denote this component RIR_t^{MPR}. There is discernible positive comovement between pdv_t and RIR_t^{MPR}, reinforcing the result that low interest rates associated with dovish monetary policy are also associated with low return premia. The correlation between the two, 0.82, is systematically larger than the correlation of 0.09 between pdv_t and the residual component of the real interest rate, $RIR_t - RIR_t^{MPR}$, and thus is also larger than the correlation of 0.43 between premia and the real interest rate itself (RIR_t). This result shows that shifts in the monetary policy stance play an important role in generating the positive comovement between return premia and the real interest rate in the data, while other movements in the real interest rate do not share this property. This may be because persistent low- or high-interest rate environments that are a consequence of shifts in the conduct of monetary policy have effects that last for decades, in contrast to movements in real rates driven by more transitory factors.

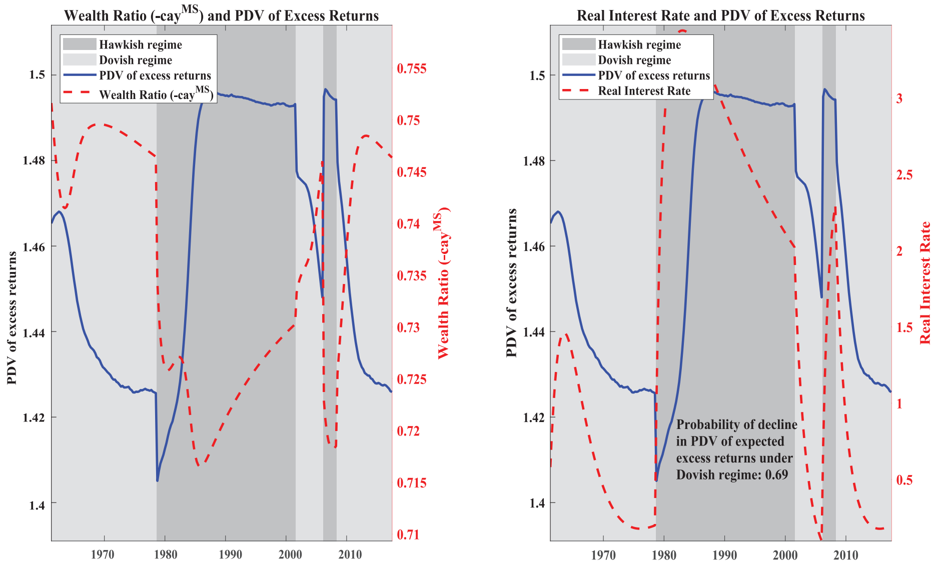

Figure 13. Simulated wealth ratio, real interest, and implied present discounted value (PDV) of expected excess returns. This figure plots results from simulating the dynamic macro-finance model at the posterior mode parameter values 20,000 times using a sample length and regime sequence equal to that in our historical data. Using data from each simulated sample, we estimate an MS-VAR and use it to compute the PDV of expected (i.e., forecasted) future excess returns. At each point in time we compute the average (across simulations) of the PDV of excess returns (reported in both panels), the wealth ratio $-cay_t$ (left panel), and the real interest rate (right panel). Since we average across sample paths, the observed movements in $-cay_t$ and the real interest rate are attributable to changes in the policy rule. (Color figure can be viewed at wileyonlinelibrary.com)

To evaluate the model implications for these comovements, Figure 13 plots output from 20,000 model simulations of length equal to that in our historical data set. To ensure that the artificial samples that we generate have a regime sequence commensurate with that observed in the historical sample, we fix the regime sequence across the simulations, drawing repeatedly from the model's Gaussian shocks. With each artificial sample, we construct a time series of the model-implied values of several variables. These model-implied variables include the present discounted value of return premia on the stock market, pdv_t, which are computed from 20,000 Bayesian estimations of an MS-VAR using the same methodology that produced the empirical pdv_t reported in Figure 4. For each t, we report the average (across simulations) of the model-implied pdv_t in both panels of Figure 13. The left panel superimposes the average (across simulations) of the model-implied log wealth ratio $-cay_t$, while the right panel superimposes the average model-implied real interest rate. Since we average across sample paths that differ only in the Gaussian shocks, the plotted series reveal fluctuations that are due solely to regime changes in the monetary policy rule.

Figure 13 shows that dovish monetary policy is associated with a high wealth ratio and low pdv_t, while hawkish monetary policy is associated with a low wealth ratio and high pdv_t, consistent with the data. When the economy moves to the first hawkish subperiod of the sample, coinciding with the Volcker disinflation, pdv_t at first declines slightly before eventually rising to a new, significantly higher level. The wealth ratio jumps down at the regime shift date, but not all the way to its final destination (left panel). Instead, it gradually adjusts downward for several more periods before reaching its nadir.

To understand this result, consider what an econometrician armed with historical data generated by the model would find. Because of learning, she would find that asset prices initially decline predictably after a switch to a hawkish regime, as investors gradually update their expectation that the regime will last. This implies that, immediately after the switch, *short*-horizon return premia are low rather than high. But because investors also overextrapolate and eventually come to believe that the regime will persist indefinitely, asset values ultimately overreact and fall by too much relative to what would be warranted by the true persistence of the regime change. This means that investors are inevitably surprised by the end of the existing regime. It follows that an econometrician would find that *long*-horizon return premia are always high in hawkish regimes, since returns jump predictably upward when the inevitable switch back to dovish policy occurs.

Because pdv_t is a weighted sum of return premia spanning short to long horizons, it can initially drift in a direction opposite to its longer run trajectory if the effect of learning on short-horizon premia outweighs the effect of overextrapolation on long-horizon premia. This occurs after the switch into the first hawkish subperiod. But because beliefs evolve in a history-dependent manner, this need not happen after all switches. Figure 13 shows that pdv_t moves monotonically after the subsequent switches in the sample. Regardless of the initial trajectory of pdv_t, the model implies that it is always higher on average in hawkish regimes than in dovish regimes, consistent with the data.

The model-implied posterior probability that the regime average of pdv_t is lower in the dovish regime than in the hawkish regime is 73%.[28] By contrast, under the AP rational expectations counterfactual of the model, this same posterior probability is 56%, providing only weak evidence of any change in premia across the regimes. Intuitively, since policy rule regime shifts under rational expectations affect asset valuations primarily by changing the real interest rate, they leave return *premia* largely unaffected.[29] If instead we consider a counterfactual that retains the overextrapolation in investor beliefs but eliminates learning, we find that the posterior probability rises to 75%, slightly higher than the baseline probability of 73%. This shows that learning, which

---

[28] These probabilities are obtained as the fraction of draws from the posterior distribution for which the average present discounted value is lower in the dovish regime than in the hawkish regime.

[29] Even under AP rational expectations the central bank's policy rule has a small effect on return premia due to the implications of the policy rule for macroeconomic stability.

is crucial for explaining a gradual adjustment of valuation ratios after regime shifts, works against the model's ability to explain the behavior of return premia. The effect of learning on return premia is nonetheless small because, at the estimated parameter values, the speed of learning is relatively quick compared to the persistence of policy regimes.

III. Conclusion

We show that the U.S. economy is characterized by large, longer term regime shifts in asset values relative to macroeconomic fundamentals that arise concurrent with equally important shifts in the level of the short-term real interest rate in excess of a widely used measure of the "natural" rate of interest, a variable we refer to as the mps. Our results identify two "hawkish" subperiods of the sample characterized by a high *mps* and low asset valuations: 1978:Q4 to 2001:Q3 and 2006:Q2 to 2008:Q2. The first subperiod spans the Volcker disinflation and its aftermath, while the second subperiod follows 17 consecutive Federal Reserve rate increases that left the nominal funds rate standing at 5.25% in June 2006. All other subperiods through the end of our sample in 2017:Q3 are identified as "dovish" regimes with low mps and high asset valuations. We further document that the dovish subperiods are associated with lower equity market return premia.

To investigate what part of these findings could be due to monetary policy, we solve and estimate a novel macrofinance model of monetary transmission. Estimates of this model imply that the conduct of monetary policy differed markedly across the previously estimated dovish and hawkish subperiods. Specifically, the dovish, low-mps subperiods are characterized by an estimated interest rate rule that is consistent with accommodative monetary policy, while the hawkish, high-mps subperiods are characterized by a rule consistent with restrictive policy. In both the model and the data, subperiods characterized by dovish policy rules are also characterized by persistently low values for the mps, persistently high stock market valuations, and persistently low equity market return premia, while subperiods characterized by hawkish policy rules exhibit the opposite pattern. The model therefore provides a rationale for how monetary policy can have long-lasting effects on real variables, equity markets, and return premia.

The model and its estimates speak to the origins of persistently declining real interest rates over 40 years. A striking finding is that two-thirds of the downward trajectory in short-term real rates observed since the early 1980s can be attributed to monetary policy, that is, to regime changes in the conduct of policy. This result obtains because the policy rule parameters exhibit a decisive shift toward hawkish values around the time of Volcker's appointment to the Federal Reserve, but then exhibit an equally decisive shift back to dovish values in the aftermath of 9/11. The estimated policy rule has remained dovish since, with the exception of a brief interlude from 2006:Q2 to 2008:Q2.

The model fit for other data that were not a target of our estimation is less tight. For example, the model implies that the AP agent's five-year-ahead inflation expectations averaged about 4% at an annual rate from 2000 to the end of our sample in 2017:Q3. By contrast, Blue Chip surveys of corporate executives show five-year inflation expectations fluctuating in narrow bands around values just above 2%. Future work could explore whether allowing for a more general monetary policy reaction function that explicitly incorporates roles for unconventional monetary policy could improve the models' implications along these lines.

Initial submission: June 15, 2018; Accepted: March 9, 2021
Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

REFERENCES

Acharya, Viral V., and Hassan Naqvi, 2019, On reaching for yield and the coexistence of bubbles and negative bubbles, *Journal of Financial Intermediation* 38, 1–10.

Angeletos, George-Marios, Zhen Huo, and Karthik A. Sastry, 2020, Imperfect macroeconomic expectations: evidence and theory, NBER Working paper No. 27308.

Barro, Robert J., Jesús Fernández-Villaverde, Oren Levintal, and Andrew Mollerus, 2014, Safe assets, NBER Working paper No. w20652.

Bernanke, Ben S., and Kenneth N. Kuttner, 2005, What explains the stock market's reaction to Federal Reserve Policy?, *Journal of Finance* 60, 1221–1257.

Bianchi, Francesco, 2013, Regime switches, agents' beliefs, and post-World War II U.S. macroeconomic dynamics, *Review of Economic Studies* 80, 463–490.

Bianchi, Francesco, Sydney Ludvigson, and Sai Ma, 2020, Belief distortions and macroeconomic fluctuations, NBER Working paper No. 27406.

Bianchi, Francesco, and Leonardo Melosi, 2016, Modeling the evolution of expectations and uncertainty in general equilibrium, *International Economic Review* 57, 717–756.

Boyarchenko, Nina, Valentin Haddad, and Matthew C. Plosser, 2016, The federal reserve and market confidence, Staff Reports 773, Federal Reserve Bank of New York.

Caballero, Ricardo J., and Emmanuel Farhi, 2014, The safety trap, NBER Working paper No. 19927.

Campbell, John Y., 2017, *Financial Decisions and Markets: A Course in Asset Pricing* (Princeton University Press, Princeton, NJ).

Campbell, John Y., and Robert J. Shiller, 1989, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.

Campbell, John Y., and Tuomo Vuolteenaho, 2004, Bad beta, good beta, *American Economic Review* 94, 1249–1275.

Choi, Jaewon, and Mathias Kronlund, 2018, Reaching for yield in corporate bond mutual funds, *Review of Financial Studies* 31, 1930–1965.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, 2005, Nominal rigidities and the dynamic effects of a shock to monetary policy, *Journal of Political Economy* 113, 1–45.

Cieslak, Anna, Adair Morse, and Annette Vissing-Jorgensen, 2019, Stock returns over the FOMC cycle, *Journal of Finance* 74, 2201–2248.

Cieslak, Anna, and Andreas Schrimpf, 2019, Non-monetary news in Central Bank communication, *Journal of International Economics* 118, 293–315.

Clarida, Richard, Jordi Gali, and Mark Gertler, 2000, Monetary policy rules and macroeconomic stability: Evidence and some theory, *Quarterly Journal of Economics* 115, 147–180.

Cochrane, John H., 2005, *Asset Pricing: Revised Edition* (Princeton University Press, Princeton, NJ).

Cohen, Randolph B., Christopher Polk, and Tuomo Vuolteenaho, 2003, The value spread, *Journal of Finance* 58, 609–641.

Coimbra, Nuno, and Hélene Rey, 2017, Financial cycles with heterogeneous intermediaries, NBER Working paper No. 23245.

Cook, Timothy, and Thomas Hahn, 1989, The effect of changes in the federal funds rate target on market interest rates in the 1970s, *Journal of Monetary Economics* 24, 331–351.

Di Maggio, Marco, and Marcin T. Kacperczyk, 2015, The unintended consequences of the zero lower bound policy, Working paper, Columbia Business School.

Diamond, Douglas W., and Raghuram G. Rajan, 2012, Illiquid banks, financial stability, and interest rate policy, *Journal of Political Economy* 120, 552–591.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl, 2018, A model of monetary policy and risk premia, *Journal of Finance* 71, 317–373.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Farhi, Emmanuel, and Jean Tirole, 2012, Collective moral hazard, maturity mismatch, and systemic bailouts, *American Economic Review* 102, 60–93.

Galí, Jordi, 2015, Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications (Princeton University Press, Princeton, NJ).

Gertler, Mark, and Peter Karadi, 2015, Monetary policy surprises, credit costs, and economic activity, *American Economic Journal: Macroeconomics* 7, 44–76.

Gilchrist, Simon, David López-Salido, and Egon Zakrajšek, 2015, Monetary policy and real borrowing costs at the zero lower bound, *American Economic Journal: Macroeconomics* 7, 77–109.

Greenwald, Daniel, Martin Lettau, and Sydney C. Ludvigson, 2019, How the wealth was won: Factor shares as market fundamentals, NBER Working paper No. 25769.

Gürkaynak, Refet S., Brian Sack, and Eric Swanson, 2005, The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models, *American Economic Review* 95, 425–436.

Hall, Robert E., 2016, Understanding the decline in the safe real interest rate, NBER Working paper No. w22196.

Hamilton, James D., 1994, Time Series Analysis (Princeton University Press, Princeton, NJ).

Hanson, Samuel G., David O. Lucca, and Jonathan H. Wright, 2021, Rate-amplifying demand and the excess sensitivity of long-term rates, *Quarterly Journal of Economics* 136, 1719–1781.

Hanson, Samuel G., and Jeremy C. Stein, 2015, Monetary policy and long-term real rates, *Journal of Financial Economics* 115, 429–448.

Jarocinski, Marek, and Peter Karadi, 2020, Deconstructing monetary policy surprises: The role of information shocks, *American Economic Journal: Macroeconomics* 12, 1–43.

Kekre, Rohan, and Mortiz Lenel, 2021, Monetary policy, redistribution, and risk premia, NBER Working paper No. 28869.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, 2012, The aggregate demand for treasury debt, *Journal of Political Economy* 120, 233–267.

Laubach, Thomas, and John C. Williams, 2003, Measuring the natural rate of interest, *Review of Economics and Statistics* 85, 1063–1070.

Lettau, Martin, and Sydney C. Ludvigson, 2001, Consumption, aggregate wealth and expected stock returns, *Journal of Finance* 56, 815–849.

Lettau, Martin, and Sydney C. Ludvigson, 2005, Expected returns and expected dividend growth, *Journal of Financial Economics* 76, 583–626.

Lettau, Martin, and Sydney C. Ludvigson, 2010, Measuring and modeling variation in the risk-return tradeoff, in Yacine Ait-Sahalia, and Lars Peter Hansen, eds., Handbook of Financial Econometrics (Elsevier Science B.V., North Holland, Amsterdam).

Lettau, Martin, and Sydney C. Ludvigson, 2013, Shocks and crashes, in Jonathan Parker, and Michael Woodford, eds., National Bureau of Economics Research Macroeconomics Annual: 2013 (MIT Press, Cambridge and London).

Lubik, Thomas A., and Frank Schorfheide, 2004, Testing for indeterminacy: An application to U.S. monetary policy, *American Economic Review* 94, 190–217.

Lucca, David O., and Emanuel Moench, 2015, The pre-FOMC announcement drift, *Journal of Finance* 70, 329–371.

Malmendier, Ulrike, and Stefan Nagel, 2011, Depression babies: Do macroeconomic experiences affect risk taking?, *Quarterly Journal of Economics* 126, 373–416.

Malmendier, Ulrike, and Stefan Nagel, 2016, Learning from inflation experiences, *Quarterly Journal of Economics* 131, 53–87.

Nakamura, Emi, and Jón Steinsson, 2018, High-frequency identification of monetary non-neutrality: The information effect, *Quarterly Journal of Economics* 133, 1283–1330.

Piazzesi, Monika, and Martin Schneider, 2021, Payments, credit and asset prices, Working paper, Stanford University.

Rajan, Raghuram G., 2006, Has finance made the world riskier?, *European Financial Management* 12, 499–533.

Rajan, Raghuram G., 2013, A step in the dark: Unconventional monetary policy after the crisis, *Andrew Crockett Memorial Lecture*, BIS, Basel 23.

Sims, Christopher A., and Tao Zha, 2006, Were there regime switches in U.S. monetary policy?, *American Economic Review* 91, 54–81.

Taylor, John B., 1993, Discretion versus policy rules in practice, in *Carnegie-Rochester Conference Series on Public Policy* 39, 195–214.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.
Replication Code.