

Monetary Policy and Asset Valuation*

Francesco Bianchi
Duke, CEPR, NBER

Martin Lettau
UC Berkeley, CEPR, NBER

Sydney C. Ludvigson
NYU, CEPR, NBER

First draft: May 12, 2016
This draft: March 23, 2021

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 monetary policy stance can generate persistent changes in asset valuations and the equity premium.

JEL: G10, G12, G17

*Bianchi: Department of Economics, Duke, 213 Social Sciences building, Box 90097, Durham, NC 27708-0097; Email: francesco.bianchi@duke.edu, <https://sites.google.com/view/francescobianchi/home>. Lettau: Haas School of Business, University of California at Berkeley, 545 Student Services Bldg. #1900, Berkeley, CA 94720-1900; E-mail: lettau@haas.berkeley.edu; Tel: (510) 643-6349, <https://sites.google.com/view/martinlettau/>. Ludvigson: Department of Economics, New York University, 19 W.4th Street, 6th Floor, New York, NY 10012; Email: sydney.ludvigson@nyu.edu; <https://www.sydneyludvigson.com>. We thank Ian Dew-Becker, Emmanuel Farhi, Simon Gilchrist, Matteo Maggiori, Stefan Nagel (the editor), Emi Nakamura, Monika Piazzesi, Jon Steinsson, Andrea Tamoni, Harald Uhlig, and 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.

1 Introduction

There is growing evidence that the real values of long-term financial assets fluctuate sharply in response to the actions and announcements of central banks. This includes the value of the stock market, a perpetual asset that endures indefinitely. But this creates a puzzle. Asset pricing theories can generally rationalize such large responses only if market participants believe that something related to the conduct of monetary policy will have a long-lasting influence on real variables.¹ Yet the notion that monetary policy shocks could have long-lived effects on real variables is contravened by an agglomeration of foundational New Keynesian macro theories, and empirical evidence appears consistent with this (e.g., Christiano, Eichenbaum, and Evans (2005)).² But if this is so, how does monetary policy influence long-lived assets?

One possibility is that some component of monetary policy does in fact have long-lasting, first-order effects on the aggregate economy, on real interest rates, and on the stock market, even if identified 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 in excess of a widely used measure of the “natural” rate of interest “ r^* ,” a spread referred to hereafter as the *monetary policy spread*, or *mps* for short. Since the Federal Reserve targets the federal funds rate but in theory has no control over the natural rate, a non-zero value for the *mps* may be considered a measure of the stance of monetary policy, i.e., 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, while hawkish regimes coincide with persistently low valuations. 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 of 2006. All other subperiods of the sample are identified as dovish regimes with high valuations and low *mps*.

Our second result is that the dovish, low *mps* regimes coincide with lower equity market return premia, while hawkish, high *mps* regimes coincide with higher return premia. Specifically, in a switch from a hawkish to dovish regime, the estimated present discounted value of future

¹We define a “real variable” here as any non-nominal variable, including risk premia and credit spreads.

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

return premia on the aggregate stock market, as well as that of several equity characteristic portfolios, simultaneously fall to lower levels. Moreover, the return premia of evidently riskier, higher Sharpe ratio portfolios, such as those that go long in value stocks or stocks that have recently appreciated the most, fall more than those of evidently less risky, lower Sharpe ratio portfolios, such as those that go long in growth stocks or stocks that have recently appreciated the least.

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? 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 macro-finance 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 who earns income from investments in two assets: the aggregate stock market and the one-period nominal bond market. This agent may be thought of as a relatively sophisticated investor who typifies the type of wealthy individual or large institution that constitutes a small fraction of the population but owns the vast majority of highly concentrated financial wealth in the U.S. Because the agent is assumed to be vanishingly small relative to the overall population, she takes the macroeconomic dynamics of the economy as given. We refer to this agent interchangeably as the “asset pricing agent” or “investor.”

The second block of the model determines macroeconomic dynamics and is driven by a representative “macro agent” who 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 macro-agent expectations about inflation. Specifically, we allow the evolution of expectations about trend inflation to be potentially influenced by both an adaptive expectations component as well as 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 into data by filtering observations on household inflation expectations from the University of Michigan Survey of Consumers (SOC). Overall perceived trend inflation is then 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 either because the target is imperfectly observed, or because central bank announcements about the target are not viewed as fully credible. Because the weights on the two terms are freely estimated, our approach allows us to directly assess the importance of adaptive expectations and imperfect information about the inflation target.

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 call 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.

This paper does not take a stand on the microfoundations macro block. However, the framework implies that overall macroeconomic dynamics are driven by a central bank who infrequently shifts the conduct of monetary policy, and an “average” household who typifies the vast majority of the population with modest financial assets but whose expectations about inflation and aggregate economic activity preponderate in the general population.

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 they come to view the existing policy stance as the new normal. This aspect of beliefs implies that investors extrapolate too much from the observed continuity in the policy stance, so that the perceived persistence of policy regime shifts overstates their true persistence. The combination of learning plus a fading memory distortion implies that beliefs evolve in a history-dependent manner, with important consequences for how asset valuations adjust in the wake of regime changes in monetary policy.

The full theoretical framework is solved and estimated, with the macro block parameters and latent states estimated using Bayesian methods under flat priors. The estimation uses data on inflation expectations from the SOC, the nominal federal funds rate, output growth, inflation, and a measure of asset valuations. Doing so, we find that the parameters of the monetary policy rule differ markedly across the previously estimated *m*ps regimes. Specifically, we find that the dovish, low *m*ps 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 *m*ps subperiods coincide with a *hawkish policy rule* characterized by a

lower inflation target and greater responsiveness to inflation relative to growth.

With these model estimation results in hand, we identify movements in real variables that are attributable *solely* to the conduct of monetary policy, i.e., to regimes changes in the policy rule. Several results are noteworthy.

First, the estimates imply that changes in the conduct of monetary policy generate large and persistent fluctuations in the short-term real interest rate that last for decades. By contrast, monetary policy *shocks* have far more transitory effects, consistent with prior empirical evidence. Indeed, the estimated model implies that two-thirds of the secular decline in real interest rates observed since the early 1980s is attributable to regime changes in the conduct of monetary policy. This occurs 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:Q2-2008:Q2.

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

Third, dovish policy rules coincide in the model with persistently high asset valuations, a low *mps*, and low equity return premia, while hawkish rules coincide with persistently low valuations, a high *mps*, and high return premia, consistent with our empirical evidence. In addition, the equity return premia estimated from historical data are all strongly positively correlated over our sample with the component of the real interest rate that we estimate is driven by regime changes in monetary policy.

The success of the model in explaining these lower frequency asset pricing phenomena results from the product of two forces: (i) sticky macro-agent 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. Finally, the fading memory component of investor beliefs explains why stock market return premia can be high (low) in hawkish (dovish) subperiods, as documented in the first part of the paper. Intuitively, fading memory of past policy rules means that investors over-react to policy shifts and are always 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 dovish (hawkish).

Because the model economy specifies only an aggregate stock market and a nominal bond market, it cannot speak to one aspect of our empirical evidence, namely the behavior of the cross-section of equity characteristic portfolios across our identified policy regimes. The concluding section of the paper discusses some possibilities for future research to explore the theoretical underpinnings of this finding.

The rest of the paper is organized as follows. The next section discusses related literature. Section 3 discusses the estimation of a joint Markov-switching system for asset valuations and the monetary policy spread and investigates whether the low *mps* regimes are characterized by lower return premia in equity market assets. Section 4 describes the model, explains how it is solved and estimated, and presents results of that estimation. Section 5 concludes. A large amount of additional material, test results, and a detailed data description of each procedure have been placed in an Appendix for online publication.

2 Related Literature

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 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. Yet no matter what the channel, asset pricing 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.⁴ Our work contributes to this literature by providing 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 new 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. Prominent examples in this literature include theories with a “reach-for-yield” motive either in preferences or technologies (e.g., Rajan

³See Hanson and Stein (2015), Gertler and Karadi (2015), Gilchrist, López-Salido, and Zakrajšek (2015), Boyarchenko, Haddad, and Plosser (2016), Jarocinski and Karadi (2019), Cieslak and Schrimpf (2019), and Kekre and Lenel (2019). These studies follow on earlier work finding a link between monetary policy surprises and short-term assets in high frequency data (Cook (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 FOMC-cycle time (Lucca and Moench (2015), Cieslak, Morse, and Vissing-Jorgensen (2015)).

⁴For reviews of frontier asset pricing models, see Cochrane (2005) and Campbell (2017).

(2006); Rajan (2013); Diamond and Rajan (2012); Farhi and Tirole (2012); Drechsler, Savov, and Schnabl (2018); Piazzesi and Schneider (2015); Acharya and Naqvi (2016); Coimbra and Rey (2017); Hanson, Lucca, and Wright (2018)). 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 (2019), driving down return premia. 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. And we provide a new explanation for low return premia in low interest rate regimes based on the idea that investors may extrapolate too much from the observed continuity in the policy stance, thereby creating a wedge between the subjective and objective persistence of policy regime shifts.

A separate body of theoretical work addresses the low and declining interest rates of recent decades with implications for risk premia that are opposite to what would be consistent with our empirical findings. 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 (e.g., Barro and Mollerus (2014); Caballero and Farhi (2014); Hall (2016)). Thus, asset valuations in these theories can only be higher if the decline in the risk-free rate exceeds the rise in risk premia. Our evidence from equity markets implies that low interest rate regimes coincide with lower rather than higher equity market return premia. Our findings for stock market returns in this regard are reminiscent of similar evidence for the Treasury market (e.g., Hanson and Stein (2015)), for U.S. prime money funds (e.g., Di Maggio and Kacperczyk (2015)), and for U.S. corporate bond mutual funds (Choi and Kronlund (2015)). The evidence in these papers pertains to heavily intermediated asset classes. By contrast, our evidence pertains to equity market portfolios, an asset class ostensibly held by retail investors and households, as well as intermediaries.

Finally, our work is related to previous research that has found 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)). Unlike this work, we use a more recent sample and estimate whether there are *joint* regime changes in asset valuations and the *m*_{ps that coincide with regime shifts in the policy rule and risk premia. We also present new evidence, vis-a-vis this literature, of changes in the policy rule parameters toward more dovish monetary policy that occurred at beginning of the 21st century.}

This latter body of work also helps to 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 US 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 being drawn from a mixture of distinct distributions with infrequent transitions between them, rather than a single distribution where a transition occurs each period. The regime switching approach therefore helps us evaluate a key hypothesis of the paper, namely whether infrequent regime changes in the conduct of monetary policy are associated with changes in the real interest rate, valuations, and return premia.

Because the model economy specifies only an aggregate stock market and a nominal bond market, it cannot speak to one aspect of our empirical evidence, namely the behavior of the cross-section of equity characteristic portfolios across our identified policy regimes. The concluding section of the paper discusses some possibilities for future research to explore the theoretical underpinnings of this finding.

The rest of the paper is organized as follows. The next section discusses related literature. Section 3 discusses the estimation of a joint Markov-switching system for asset valuations and the monetary policy spread and investigates whether the low *mps* regimes are characterized by lower return premia in equity market assets. Section 4 describes the model, explains how it is solved and estimated, and presents results of that estimation. Section 5 concludes. A large amount of additional material, test results, and a detailed data description of each procedure have been placed in an Appendix for online publication.

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

3.1 Motivating Evidence

Figure 1 plots the behavior over time of a key instrument of monetary policy, namely the real federal funds rate, 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 sample 1961:Q1-2017:Q3.⁵ The figure shows that there are important lower-frequency fluctuations in the real federal funds rate over the full sample, but little long-term trend. By contrast, the natural rate of interest exhibits a clear downward trend over the entire sample. The right panel plots

⁵The 1961 start date is dictated by the availability of the natural rate of interest measure.

the spread between the real funds rate and the Laubach and Williams (2003) natural rate of interest, a variable we refer to as the *monetary policy spread*. The natural rate of interest measures the component of the real rate whose fluctuations cannot be attributed to monetary policy.⁶ Thus the spread between the real federal funds rate 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 .⁷ According to this measure of the *mps*, monetary policy was accommodative in the sample 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 1 reports the correlations between the real interest rate or the *mps* and different asset valuation metrics. These correlations are reported for the raw series, and for components of the raw series that retain fluctuations with “medium” term cycles, defined to be cycles that take between 8 and 50 years to complete, and “business” cycles, defined to take between 1.5 and 8 years, 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 hereafter), 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. Columns B-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 ratio⁸, Panel C uses the price-dividend ratio for the corporate sector, and Panel D uses the price-earnings ratio for the corporate sector.

Several results in Table 1 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 8-50 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.

⁶Estimates of the natural rate of interest 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 consistent with no change in inflation.

⁷ mps_t is computed as

$$FFR_t - (\text{Expected Inflation})_t - r_t^*$$

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

⁸<http://www.multpl.com/shiller-pe/>

Second, in all cases, the absolute correlation between the valuations and the mps 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 is suggestive that monetary policy as opposed to real rates *per se* play a role in this correlation.

Third, the largest absolute correlation is with $-cay_t$, which has a -0.83 correlation with the real interest rate and a -0.84 correlation with the 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, namely 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.⁹ One reason for this 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 market cash flows appear in the numerator. We observe this mechanism at work in the current data in Panel E of Table 1. 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.

3.2 Regime Changes in cay and mps

This section presents results for a joint Markov-switching model of breaks in the mean of cay and an instrument of monetary policy. As a cleaner indicator of monetary policy, we use the mps rather than the real federal funds rate, in order to purge the later of the trending natural rate component that has nothing to do with monetary policy. We first describe an econometric

⁹See the review of the literature on expected stock market returns in Lettau and Ludvigson (2013). Earlier evidence in the literature has found that cay_t performs better in predicting stock market returns than other predictor variables at all but very long horizons. One reason for this is provided in Lettau and Ludvigson (2005), who show that, at cyclical frequencies, expected returns are high when expected dividend growth is high, a co-movement that diminishes the predictive power of the price-dividend ratio for stock returns, but does not do so for cay_t . This happens because the measures of “fundamentals” in cay_t are c and y , which are only weakly correlated with dividends. Moreover, a review of the literature shows that, even rigorous statistical tests find that cay_t has one of the best track records for predicting stock market returns both in-sample and out-of-sample.

model of regime switches in the mean of cay_t . We then introduce a similar relation for the mps . Finally, we explain how we jointly estimate regime changes in the means of the two variables.

The log valuation variable cay_t is derived from an approximate formula for the log consumption to aggregate (human and non-human) 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.¹⁰ 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

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

where ν is the steady state ratio of human wealth to asset wealth and $r_{a,t}$ is the log return to asset (non-human) 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 above stationary linear combination of c_t , a_t , and y_t may be written

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

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 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 a low and a 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} , and the residuals ϵ_t^{FC} and ϵ_t^c are not the same. Because our procedure jointly recovers the slope coefficients β_a

¹⁰This 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.

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 combine the estimation of changes in the mean of cay_t^{MS} with an isomorphic model for mps_t to estimate a joint Markov-switching model with synchronized regimes. Specifically, we assume that regime changes in the mean of cay_t^{MS} coincide with regime changes in the mean of the mps :

$$mps_t = r_{\xi_t} + \epsilon_t^r, \quad (4)$$

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} . It is worth emphasizing that the same latent state variable, ξ_t , is presumed to follow a two-state Markov-switching process with transition matrix \mathbf{H} , controls both changes in α_{ξ_t} and r_{ξ_t} . Thus the regimes are synchronized across the two means.

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

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

where ξ_t is a latent variable that follows a Markov-switching process with transition matrix \mathbf{H} . Denote the set of parameters to be estimated collectively with the vector

$$\boldsymbol{\theta} = (\alpha_{\xi_t}, \beta_a, \beta_y, r_{\xi_t}, \sigma_{MS}^2, \sigma_r^2, \text{vec}(\mathbf{H}))'.$$

We use Bayesian methods with flat priors to estimate the model parameters in (3) and (4) over the period 1961:Q1-2017:Q3. The sequence $\xi_t = \{\xi_1, \dots, \xi_T\}$ 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 \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 whole sample and knowledge of the parameters $\boldsymbol{\theta}$. We refer to these as the smoothed regime probabilities. We may decompose cay_t^{MS} into two components, a discrete-valued time-varying mean and a continuous-valued random variable:

$$cay_t^{MS} = c_t - (\beta_a a_t + \beta_y y_t) = \bar{\alpha}_t + \epsilon_t^c \quad (5)$$

$$\bar{\alpha}_t = \sum_{i=1}^2 \pi_{t|T}^i \alpha_i. \quad (6)$$

thus $\bar{\alpha}_t$ 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 $\bar{r}_t = \sum_{i=1}^2 \pi_{t|T}^i r_i$.

The posterior distribution of the empirical model (3) and (4) and the corresponding regime probabilities $\pi_{t|t}^i$ and $\pi_{t|T}^i$ 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_t^{MS} and its decomposition into $\bar{\alpha}_t$ and ϵ_t^c , and of mps_t into \bar{r}_t and ϵ_t^r , use the posterior mode of the parameter vector θ 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 Appendix.

The variable cay_t^{MS} 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_t^{MS} as an inverse “wealth” ratio, or equivalently define the (log) *wealth ratio* as $-cay_t^{MS} = -[\epsilon_t^c + \bar{\alpha}_t]$. 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_t^c and ϵ_t^{FC} are mean zero random variables, thus the intercept term $\bar{\alpha}_t$ gives the mean of the inverse wealth ratios.¹¹

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 the high value for r_{ξ_t} as *hawkish* and denote them with an H subscript, and to those with the low value for r_{ξ_t} as *dovish*, and denote them with a D subscript, i.e., $r_H \geq r_D$. Because the regimes in r_{ξ_t} and α_{ξ_t} are synchronized, switches in r_{ξ_t} will by construction coincide with switches in α_{ξ_t} . But the magnitude by which either variable switches, and whether α_{ξ_t} will be high or low when r_{ξ_t} is high are open empirical questions that our estimation is designed to address.

Table 2 reports the parameter estimates, while Figure 2 reports 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 three 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 $\hat{r}_H = 0.0111$ and $\hat{\alpha}_H = -0.7239$. The posterior mode estimates for the low $r_{\xi_t} = r_D$ dovish regime are $\hat{\alpha}_2 = -0.7500$ and $\hat{r}_2 = -0.0252$. 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 ξ^T .¹² Table 3 shows the resulting regime subperiods based on this estimated regime sequence. The hawkish regime prevails for a

¹¹In a finite sample, ϵ_t^{MS} and ϵ_t^{FC} are not necessarily mean zero because of the leads and lags of the first differences included in the DLS regression used to correct for finite sample biases—see the Appendix. In population these variables are mean-zero by definition.

¹²The Appendix describes how the most likely regime sequence is computed from the filtered probabilities.

prolonged period of time from 1978:Q4 to 2001:Q3, during which the smoothed probability that $r = \hat{r}_H$ 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 $r = \hat{r}_H$ is virtually 0. The hawkish regime briefly reappears from 2006:Q1 to 2008:Q2 following a string of 17 target federal funds rate hikes by the Federal Reserve that began on June 30, 2004 and ended with the nominal rate standing at 5.25% on the 29th of June 2006. The target funds rate remained above 4% until January 2008, when it was lowered to 3%.

How different are the hawkish parameter values from the dovish ones? Table 2 reports summary statistics for the differences $\hat{\alpha}_H - \hat{\alpha}_D$ and $\hat{r}_H - \hat{r}_D$, along with percentiles of their posterior distributions. The 90% credible sets for $\hat{\alpha}_H - \hat{\alpha}_D$ and $\hat{r}_H - \hat{r}_D$ are non-zero 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 \mathbf{H} , also reported in Table 2.

To give a visual impression of the properties of these regimes, Figure 3 plots $-cay_t^{MS}$ and the *mps* over time. Also plotted as horizontal lines are the values $-\bar{\alpha}_t$ and \bar{r}_t that arise in each regime over the sample. The figure shows that these estimated values differ by quantitatively large magnitudes across the 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 in the middle, 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 the latter part of 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 the “Greenspan Put,” in Figure 3 after the perceived attempt of Chairman 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

Chairman Bernanke that promised in 2011 to keep interest rates at ultra low levels for an extended period of time, possibly longer than that warranted by a 2% inflation objective. We argue 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.

3.3 Return Premia

The evidence above suggests that a persistently low mps_t is associated with persistently high asset valuations. A natural question is whether persistently low/high real interest rates environments associated with the previously estimated dovish/hawkish mps regimes influence asset valuations only through the riskless real interest rate changes, or whether the return premia also change.¹³ We begin with a loglinear framework to illustrate how real interest rate regimes could be revealed in return premia both in the time series and in a cross-section.

3.3.1 A Loglinear Framework

We consider a log-linearization that constructs earnings from book-market and return data using clean surplus accounting. Let B_t denote book value and M_t denote market value, and let the logarithm of the book-market ratio $\log(B_t/M_t)$ be denoted θ_t . Vuolteenaho (2000) shows that θ_t of an asset or portfolio can be decomposed as:

$$\theta_t = \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t r_{t+1+j} + \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t f_{t+1+j} - \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t e_{t+1+j}^* \quad (7)$$

where $\rho < 1$ is a parameter, and r_{t+1+j} , f_{t+1+j} , and e_{t+1+j}^* denote the log excess return, log risk-free rate, and log earnings, respectively.¹⁴ In other words, the logarithm of the book-market ratio θ_t depends on the present discounted value (PDV) of expected excess returns (return premia), expected risk-free rates, and expected earnings.

One way to assess the possibility that return premia are affected by our estimated policy regimes is to exploit differences across assets. A change in discount rates driven by the riskless rate alone influences all assets in the same way, regardless of their riskiness. By contrast, if shifts between dovish and hawkish mps regimes affect return premia, its reasonable to expect high return premia assets to be more affected than low return premia assets. Specifically, suppose

¹³In what follows we use the terms “risk” premia and “return” premia interchangeably to refer to the expected return on an asset in excess of the risk-free rate. We remain agnostic as to whether the observed premia are attributable to genuine covariance with systematic risk factors or mispricing, or both.

¹⁴Specifically, e^* is the log of 1 plus the earnings-book ratio, adjusted for approximation error. See Vuolteenaho (2000).

we have two portfolios x and y , the spread in their book-market ratios, $\theta_{x,t} - \theta_{y,t}$, is given by:¹⁵

$$\underbrace{\theta_{x,t} - \theta_{y,t}}_{\text{Spread in BM ratios}} = \underbrace{\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (r_{x,t+1+j} - r_{y,t+1+j})}_{\text{PDV of spread in return premia}} - \underbrace{\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (e_{x,t+1+j}^* - e_{y,t+1+j}^*)}_{\text{PDV of spread in expected earnings}}$$

Note that the risk-free rate has no affect on this spread, since all portfolios are affected in the same way by the risk-free rate. Instead only the return premium differential and expected earnings differential affect the book-market spread. We can adjust the book-market spread for the spread in expected earnings to isolate the return premium differential:

$$\underbrace{\theta_{x,t} - \theta_{y,t} + \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (e_{x,t+1+j}^* - e_{y,t+1+j}^*)}_{\text{Spread in BM ratios adjusted for earnings} \equiv \tilde{\theta}_{x,t} - \tilde{\theta}_{y,t}} = \underbrace{\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (r_{x,t+1+j} - r_{y,t+1+j})}_{\text{PDV of the spread in risk premia}} \quad (8)$$

The above expression shows that the spread in book-market ratios adjusted for expected future earnings is identically equal to the PDV of the spread in return premia.

Denote the *adjusted* (for expected earnings) book-market ratio for portfolio x in regime i with a tilde as

$$\tilde{\theta}_{x,t}^i \equiv \theta_{x,t}^i + \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t e_{x,t+1+j}^{*i}.$$

Let x denote a high return premium portfolio while y denotes a low return premium portfolio. To investigate whether return premia fall when policy turns dovish (and rise when it turns hawkish), we can ask whether, in a shift from a hawkish ($i = H$) to a dovish ($i = D$) *mps* regime, the adjusted book-market ratio of x falls more than y , implying $(\tilde{\theta}_{x,t}^H - \tilde{\theta}_{x,t}^D) - (\tilde{\theta}_{y,t}^H - \tilde{\theta}_{y,t}^D) > 0$, or that the *difference-in-difference* of adjusted book-market ratios is positive across regimes:

$$(\tilde{\theta}_{x,t}^H - \tilde{\theta}_{y,t}^H) - (\tilde{\theta}_{x,t}^D - \tilde{\theta}_{y,t}^D) > 0. \quad (9)$$

Thus if return premia fall in dovish subperiods, the *spread* in the *adjusted* book-market ratios between the high premia portfolio and the low premia portfolio should be lower in dovish regimes. Equivalently, in a switch from a hawkish to dovish *mps* regime, the PDV of return premia on the high return premium portfolio should fall more than that of the low return premium portfolio, compressing spreads.

Empirical methods can be used to estimate the PDV of the spread in return premia $\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (r_{x,t+1+j} - r_{y,t+1+j})$ for any two stock portfolios, or for the difference-in-difference (9) in the PDV of return premia, or for the PDV $\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t r_{t+1+j}$ of return premia on the aggregate stock market in excess of a short-term interest rate. We do all of these exercises here.

¹⁵This derivation follows Cohen, Polk, and Vuolteenaho (2003) and imposes the approximation that ρ is constant across portfolios. Cohen, Polk, and Vuolteenaho (2003) find that the approximation error generated by this assumption is small. We further follow these authors 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, $\rho = (0.96)^{.25}$.

3.3.2 Policy Regimes and Equity Return Premia

For this purpose we estimate Markov-switching vector autoregressions (MS-VARs) taking 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 .¹⁶ The MS-VAR coefficients and shock volatilities are not fixed over the sample but instead 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} .

In our empirical application, the MS-VAR parameters A_{ξ_t} and V_{ξ_t} , and \mathbf{H} are freely estimated under flat priors. Our objective in this section is to establish whether return premia differ across the two *previously estimated* regimes. We therefore impose the previously 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 the regimes. All parameters other than the regime sequence are freely estimated and could in principle show no shift across the previously estimated regime subperiods.

To form estimates of the PDV $\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t r_{t+1+j}$ of return premia, we need an estimate of the conditional expectation terms, i.e., $\mathbb{E}_t r_{t+1+j}$. For this we use the MS-VAR to compute econometric, time t expectations of excess returns multiple steps ahead. Given that the excess returns of interest are one element of Z_t , these are all functions of $\mathbb{E}(Z_{t+s} | \mathbb{I}_t)$. This expectation conditions on the information set \mathbb{I}_t available at time t . Note that \mathbb{I}_t includes not only the history of observations Z^t , but also knowledge of the regime in place at time t , which can be observed from the most likely regime sequence, and the VAR parameters for each regime. The time t conditional expectation also takes into account the possibility of *future* regime changes. Intuitively, this is done by computing expectations conditional on every possible future regime path, $\xi_{t+1}, \dots, \xi_{t+s}$, and weighting these expectations by their probabilities:

$$\begin{aligned} \mathbb{E}(Z_{t+s} | \mathbb{I}_t) &= \sum_{\forall \{\xi_{t+1}, \dots, \xi_{t+s}\}} \mathbb{E}(Z_{t+s} | \mathbb{I}_t, \xi_{t+1}, \dots, \xi_{t+s}) \times \Pr(\xi_{t+1}, \dots, \xi_{t+s} | \mathbb{I}_t) \\ &= \sum_{\forall \{\xi_{t+1}, \dots, \xi_{t+s}\}} \left(\prod_{j=1}^s \hat{A}_{\xi_{t+j}} \right) Z_t \times \Pr(\xi_{t+1}, \dots, \xi_{t+s} | \mathbb{I}_t) \end{aligned}$$

where $\Pr(\xi_{t+1}, \dots, \xi_{t+s} | \mathbb{I}_t)$ is the probability of observing a particular path for future regimes $\xi_{t+1}, \dots, \xi_{t+s}$, conditional on the current regime. Since the probability of moving across regimes is controlled by the transition matrix \mathbf{H} , $\Pr(\xi_{t+1}, \dots, \xi_{t+s} | \mathbb{I}_t)$ may be computed using just two pieces of information: the regime in place at t and knowledge of \mathbf{H} . In summary, the expected

¹⁶If the MS-VAR has more than one lag, the companion form can be used to recast the model as illustrated above.

values computed from the MS-VAR are forecasts that take into account the probability of future regime changes.

For portfolio data, we choose a few equity characteristic portfolios that are known to exhibit large cross-sectional variation in average return premia. Specifically, we use the equity return data available from Kenneth French's Dartmouth website on portfolios formed from double-sorting all stocks in the AMEX, NYSE, NASDAQ into categories on the basis of five size categories and five BM categories, and alternatively single-sorting into 10 momentum categories based on recent past return performance.¹⁷ We then use CRSP/Compustat to construct the BM ratios of the corresponding portfolios. It is well known that high BM portfolios earn much higher average returns than low BM portfolios, exhibiting a value-spread, especially in the small size quintiles. Along the momentum dimension, recent past winner stocks earn much higher returns than recent past losers. These statistical facts are evident from Table 4, which reports sample statistics for returns on two value spread portfolios—those long in the extreme value portfolio (highest BM ratio) and short in the extreme growth portfolio (lowest BM ratio) while being in the smallest and second smallest size quintile. The same is reported for a momentum spread portfolio—the portfolio that is long in the extreme winner portfolio (M10) and short in the extreme loser portfolio (M1). The annualized Sharpe ratio for the smallest value spread portfolio is 0.60 with an annualized mean return of 10%. Similarly, the momentum spread portfolio has an annualized Sharpe ratio of 0.63 and annualized mean return of over 15%. Both of these strategies have much higher annualized Sharpe ratios and average return premia than the CRSP value-weighted stock market return in excess of the three-month Treasury bill return, where the corresponding numbers are 0.36 and 6%, respectively.

We estimate a single MS-VAR for the two value spread portfolios and the momentum spread portfolio along with other data that are predictor variables for the returns on these portfolios, chosen on the basis of an Akaike Information Criterion (AIC) selection procedure. The variables included in Z_t are: (a) the momentum return spread, i.e., the difference between the excess return of the extreme winner (M10) portfolio and the excess return of the extreme loser (M1) portfolio; (b) the value return spread (S1), i.e., the difference between the excess return of the small (size 1) high BM portfolio and the excess return of the small (size 1) low BM portfolio; (c) the value return spread (S2), i.e., the difference between the excess return of the size 2 high BM portfolio and the excess return of the small size 2 low BM portfolio; (d) the momentum BM spread: the difference between the logarithm of the BM ratio of the M10 and M1 portfolios; (e) the value BM spread (S1): The difference between the logarithm of the BM ratio of the small (size quintile 1) high book-market portfolio and the logarithm of the BM ratio of the small (size 1) low book-market portfolio; (f) the value BM spread (S2): the difference between the logarithm of the BM ratio of the size quintile 2 high book-market portfolio and the logarithm

¹⁷http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

of the BM ratio of the size 2 low book-market portfolio; (g) the real federal funds rate (nominal federal funds rate minus inflation); (h) the excess return on the small value portfolio.¹⁸ The sample for this estimation is 1964:Q1-2017:Q3.¹⁹

For the aggregate stock market, we follow a similar procedure to form an estimate of $\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t r_{t+1+j}$, where r_{t+1+j} in this instance is a measure of the return on the aggregate stock market in excess of a short-term interest rate. The variables in Z_t for this purpose are predictor variables more relevant for the aggregate stock market.²⁰

We begin by computing the *regime average* values of the adjusted BM *spreads* between the high and low return premia portfolios, $\tilde{\theta}_{xy}^i \equiv \tilde{\theta}_x^i - \tilde{\theta}_y^i$, for each regime i , equivalent to the PDV of the spread in return premia, i.e., $\tilde{\theta}_x^i - \tilde{\theta}_y^i = \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (r_{x,t+1+j} - r_{y,t+1+j})$. The regime average value of $\tilde{\theta}_{xy}^i$ is defined as the expected value of $\tilde{\theta}_{xy,t}$, 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 analogy for the aggregate stock market, is $\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t r_{t+1+j}$, the PDV of the spread between the market return and the short-term interest rate. The Appendix gives formal expressions for the regime average, and explains how they are computed from the MS-VAR parameters.

Table 5 reports the median and 68% credible sets for $\tilde{\theta}_{xy}^i$, computed from each draw of the VAR parameters from the posterior distribution. The high (x) and low (y) return premia portfolios along the BM dimension are always the extreme value (highest BM) and the extreme growth portfolio (lowest BM), respectively, in each size category. Likewise, the high and low return premia portfolios along the momentum dimension are always the extreme winner (M10) and extreme loser portfolio (M1). For the market risk premium, x is the market return and y is the risk-free rate. The third row reports the analogous values for the regime average of the difference-in-difference of the PDV of risk premia between the high and low return premia port-

¹⁸The BM spreads are included because they represent the natural valuation ratios for the portfolio return spreads that we are trying to predict (Cohen, Polk, and Vuolteenaho (2003)). The real FFR and the excess return on the small value portfolio are selected based on the Akaike Information Criterion (AIC) among a set of possible additional regressors. The five Fama/French factors (Fama and French (2015)) are considered as possible additional regressors, but not selected based on the AIC. The Online Appendix provides additional details on the variable selection procedure.

The data on BM ratios for individual portfolios are constructed from CRSP and Compustat exactly as the Fama-French portfolio returns are constructed. We mimic the selection and breakpoints of this construction and compute the book-market ratio of each portfolio.

¹⁹This is shorter than the sample used previously because reliable data for book-market ratios are not available prior to 1964:Q1.

²⁰The MS-VAR specification for the market return premium includes the following variables: (a) 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; (b) $-cay^{MS}$; (c) following Campbell and Vuolteenaho (2004), the small stock value spread (log-difference in the book to market ratio of the S1 value and S1 growth portfolio); (d) the SMB factor from Fama and French; (e) the HML factor from Fama and French. These variables are included because they improve the AIC criterion.

folios across the two previously estimated regimes, i.e., the difference $(\tilde{\theta}_{x,t}^H - \tilde{\theta}_{y,t}^H) - (\tilde{\theta}_{x,t}^D - \tilde{\theta}_{y,t}^D)$, 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 5 shows that the median values of the adjusted BM spreads $\tilde{\theta}_{xy}^i$ between the high and low return premia portfolios are positive in both the hawkish and dovish *mps* regimes. This is not surprising, since portfolios that have higher return premia should have lower market-to-book values, holding fixed expected earnings. Ostensibly riskier portfolios have a higher PDV of return premium on average, regardless of the regime (see (8)). More significantly, spreads are larger in the hawkish regime than in the dovish, implying that the difference-in-difference across regimes is always positive. Thus, the return premia of evidently riskier/higher return assets decline more in environments with persistently low real interest rates than do less risky/lower return assets.

The third row of Table 5 reports the 68% posterior credible sets in parentheses for the difference-in-difference. Although the 68% posterior credible sets include negative values for the BM spreads and the market risk premium (though not the momentum spread), this does not imply that negative values are likely. The posterior distribution of the diff-in-diff displays substantial negative skewness and, as a consequence, the probability assigned to positive values, i.e., to lower premia in the dovish *mps* subperiods, is quite high in all cases: 81%, 90%, 70%, and 64%, for the market premium, the momentum spread, the S1 BM spread, and the S2 BM spread, respectively. The odds that premia decline in the dovish *mps* regime is over 4 to 1 for the market premium, over 9 to 1 for the momentum spread, over 2.3 to 1 and 1.8 to 1 for the S1 and S2 BM spreads, respectively. 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.

Finally, we estimate the PDV of return premia evolution over the sample using the time t MS-VAR forecasts. Given the posterior distribution of the VAR parameters, these forecasts have their own posterior distribution. Figure 4 reports the median values of the PDVs over our sample as solid (blue) lines, while the regime averages are indicated by the red (dashed) lines. Although the return premia are volatile, it is evident that they fluctuate around distinct means across the regime subperiods. Return premia reach lows or near-lows in the post-millennial period, after shooting up briefly in the aftermath of the financial crisis of 2007-2008. Estimated return premia return to low levels in the post-crisis zero-lower-bound period of our sample.

4 A Macro-Finance Model of Monetary Transmission

To provide an explanation for these findings, this section proposes a new dynamic macro-finance model of monetary policy transmission, with two “blocks” describing the behavior of two different representative agents, “investors” and “households.” In both blocks we work with a loglinear approximation to 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 real macroeconomic variables such as 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 models would imply. At the same time, financial markets evidently react swiftly to central bank communications and actions. This suggests that the expectations of sophisticated financial market participants are subject to little inertia, at least insofar as pertains to beliefs about monetary policy. 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 asset pricing block that determines equilibrium risky asset prices in the model. This block is driven by the optimal behavior of a representative agent who earns all of her income 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 who typifies the type of wealthy individual or large institution that constitutes a small fraction of the population but owns the vast majority of highly concentrated financial wealth in the U.S.²¹ We assume that this agent is small enough relative to the overall economy that she takes the macroeconomic dynamics as given, including the beliefs of the broader public which drive expectations about inflation and output. We refer to this agent interchangeably as the “asset pricing 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. But 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. Taken together, these departures imply that the model can generate persistent (but not permanent)

²¹Only 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 the stock market value and earns a relatively small fraction of income as labor compensation. See GLL for further discussion.

departures from monetary neutrality. The magnitude 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, although we do not explicitly take a stand on the microfoundations of the macro block, macroeconomic dynamics can be thought of as driven by a central bank and an “average” household who 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 asset pricing 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 must learn about its duration. Second, we assume that agents exhibit a form of bounded rationality motivated by evidence in MN and Malmendier and Nagel (2011) that manifests here as “fading memory” of past policy rules. Specifically, once 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 distortion that overstates the true persistence of the regime shifts. As we discuss below, the combination of these two features of investor beliefs (learning plus a fading memory distortion) implies that asset prices in the model respond to monetary policy regime changes by initially under-reacting but eventually over-reacting. 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 under-react to shocks but subsequently over-react (Angeletos, Huo, and Sastry (2020); Bianchi, Ludvigson, and Ma (2020)).

4.1 Model Description

Asset Pricing Block The model allows for a continuum of identical investors indexed by i who derive utility from consumption, $C_{p,t}^i$, at time t . We use the suffix “ p ” to denote variables pertaining to these asset pricing agents. Investors trade in two assets: a nominal bond and a stock market. The agents’ intertemporal marginal rates of substitution in consumption takes the form:

$$M_{t+1} = \underbrace{\delta \exp(\vartheta_t^p)}_{\delta_t} \left(\frac{C_{p,t+1}^i}{C_{p,t}^i} \right)^{-\sigma_p},$$

where δ_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_t^p 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, 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 $D_t = k\mathbb{Y}_t$.²² We therefore drop the i superscript from here on and denote the consumption of the representative investor as $C_{p,t} = D_t = k\mathbb{Y}_t$.

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

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

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

$$\overline{LP}^{-1}Q_t = \mathbb{E}_t^p [M_{t+1}\Pi_{t+1}^{-1}], \quad (10)$$

where Q_t is the nominal bond price, \mathbb{E}_t^p denotes the subjective expectations of the asset pricing 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 across monetary policy regimes is driven solely by endogenous fluctuations in investor beliefs about monetary policy, which is the central focus of this paper.²³

Taking logs of (10) and using the properties of conditional lognormality delivers an expression for the real interest rate as perceived by the investor:

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

²²The assumption of a constant capital share k is made for simplicity in the current model. An extension of the model to allow for a time-varying k could in principle account for evidence that factor shares fluctuations have influenced trends in equity valuations (GLL). Our focus here is to isolate the component attributable to monetary policy so we keep k constant.

²³One could in principle reverse engineer an exogenous process for LP_t to match our evidence on return premia across policy regimes. We instead take the approach of asking how much of this variation can be explained by investor beliefs about policy shifts alone, without proliferating the model’s degrees of freedom to account for the empirical findings of the first part of the paper.

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 agent. This shows that $\vartheta_{p,t}$ is implicitly defined as

$$\vartheta_{p,t} = -[i_t - \mathbb{E}_t^p[\pi_{t+1}]] + \mathbb{E}_t^p[\sigma_p \Delta c_{p,t+1}] - .5\mathbb{V}_t^p[-\sigma_p \Delta c_{p,t+1} - \pi_{t+1}] - \bar{l}p - \ln(\delta). \quad (11)$$

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

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

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

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

where $pd_t \equiv \ln\left(\frac{P_t^D}{D_t}\right)$.

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

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

where $\kappa_1 = \exp(\bar{p}d)/(1 + \exp(\bar{p}d))$, and $\kappa_0 = \log(\exp(\bar{p}d) + 1) - \kappa_1 \bar{p}d$. Combining all of the above, the log equity premium is:

$$\underbrace{\mathbb{E}_t^p [r_{t+1}^D] - (i_t - \mathbb{E}_t^p[\pi_{t+1}])}_{\text{Equity Premium}} = \underbrace{\left[\begin{array}{l} -.5\mathbb{V}_t^p [r_{t+1}^D] - \text{COV}_t^p [m_{t+1}, r_{t+1}^D] \\ +.5\mathbb{V}_t^p [\pi_{t+1}] - \text{COV}_t^p [m_{t+1}, \pi_{t+1}] \end{array} \right]}_{\text{Risk Premium}} + \underbrace{\bar{l}p}_{\text{Liquidity Premium}}, \quad (12)$$

where $\text{COV}_t^p[\cdot]$ 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 non-human) wealth, respectively (see LL). LL show that, if the streams of income accruing to human and non-human wealth are discounted at the same rate, these coefficients are identically equal to the capital and labor *income* shares in models where 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 all 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, i.e., $-cay_t = kpd_t + const.$ ²⁴

Summarizing, the model implies the following asset pricing relations:

1. SDF:

$$m_t = \log(\delta) - \sigma_p \Delta d_t + \vartheta_{p,t-1} \quad (13)$$

2. PD ratio:

$$\begin{aligned} pd_t &= \kappa_0 + \mu + \mathbb{E}_t^p [m_{t+1} + \Delta d_{t+1} + \kappa_1 pd_{t+1}] + \\ &+ .5\mathbb{V}_t^p [m_{t+1} + \Delta d_{t+1} + \kappa_1 pd_{t+1}] \end{aligned} \quad (14)$$

3. Log Euler equation for bonds:

$$i_t - \mathbb{E}_t^p [\pi_{t+1}] = -\mathbb{E}_t^p [m_{t+1}] - .5\mathbb{V}_t^p [m_{t+1} + i_t - \pi_{t+1}] - lp \quad (15)$$

4. Wealth ratio, $-cay$:

$$-cay_t = kpd_t + const. \quad (16)$$

5. Log excess stock market return:

$$er_{t+1}^D = r_{t+1}^D - (i_t - \pi_{t+1}) = \kappa_0 + \kappa_1 pd_{t+1} - pd_t + \Delta d_{t+1} + \mu - (i_t - \pi_{t+1}) \quad (17)$$

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.²⁵ In keeping with New Keynesian models, we assume that real variables grow non-stochastically along a balanced growth path and write all equations in the macro block in terms of detrended real variables. Hereafter, detrended variables are denoted with a tilde, e.g., $\widetilde{\ln(\mathbb{Y}_t)} \equiv \widetilde{y}_t$ 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 who holds small amounts of wealth in the form of nominal bonds and

²⁴To keep the estimation tractible, the model abstracts from one aspect of the data here, namely that $-cay_t$ and pd_t are not perfectly correlated. This is attributable 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.

²⁵Outside of these two distinctive features, macroeconomic dynamics are identical to those that arise from the prototypical New Keynesian model of Galí (2015), Chapter 3.

no equity. In equilibrium, she consumes a fixed fraction $(1 - k)$ of \mathbb{Y}_t , so that log detrended consumption growth of the macro agent is $\Delta \tilde{y}_{t+1}$ and the linearized Euler equation takes the form

$$\tilde{y}_t = \mathbb{E}_t^m (\tilde{y}_{t+1}) - \sigma [i_t - \mathbb{E}_t^m (\pi_{t+1}) - r_{ss}] + f_t \quad (18)$$

where i_t is the short-term nominal interest rate, $\mathbb{E}_t^m (\pi_{t+1})$ 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 that 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:

$$\pi_t - \bar{\pi}_t = \beta \mathbb{E}_t^m [\pi_{t+1} - \bar{\pi}_t] + \kappa [\tilde{y}_{t-1} - \tilde{y}_{t-1}^*] \quad (19)$$

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

The central bank obeys the following nominal interest rate rule:

$$\begin{aligned} i_t - \left(r_{ss} + \pi_{\xi_t}^T \right) &= (1 - \rho_{i, \xi_t}) \left[\psi_{\pi, \xi_t} \left(\pi_t - \pi_{\xi_t}^T \right) + \psi_{\Delta y, \xi_t} \left(\tilde{y}_t - \tilde{y}_{t-1} \right) \right] \\ &+ \rho_{i, \xi_t} \left[i_{t-1} - \left(r_{ss} + \pi_{\xi_t}^T \right) \right] + \sigma_i \varepsilon_i, \quad \varepsilon_i \sim N(0, 1). \end{aligned} \quad (20)$$

Note 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 $\pi_{\xi_t}^T$ and in the activism coefficients ψ_{π, ξ_t} , and $\psi_{\Delta y, \xi_t}$ that govern how strongly the central bank responds to deviation 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 assumed to take on two values, $\xi_t = H$ or $\xi_t = D$, corresponding to either “hawkish” or “dovish” monetary policy. It is important to emphasize that these labels do not imply that we impose any constraints on the estimated values of parameters across the previously estimated regimes. Since we freely estimate the parameters under flat priors, the parameters could in principle show no shift across regimes, or shifts that go in the “wrong” direction with respect to the previously estimated hawkish and dovish regimes.

We interpret equations (18) through (20) as equilibrium dynamics and not a micro-founded structural model. We consider an equilibrium in which bonds are in zero-net-supply in both the macro and asset pricing blocks and thus there is no trade between the asset pricing agent and macro agent.

The macro agent's expectations about inflation are formed using an adaptive algorithm, following evidence in Malmendier and Nagel (2016) (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 .²⁶ Each period, agents form a belief about α , denoted α_t^m , that is updated over time. Updating not only affects beliefs about next period inflation, it also affects beliefs about long-term trend inflation. Define *perceived trend inflation* to be the $\lim_{h \rightarrow \infty} \mathbb{E}_t^m [\pi_{t+h}]$ and denote it by $\bar{\pi}_t$. Given the presumed autoregressive process, the Online Appendix shows that $\bar{\pi}_t = (1 - \phi)^{-1} \alpha_t^m$. This implies that expectations of one step ahead inflation are a weighted average of perceived trend inflation and current inflation:

$$\mathbb{E}_t^m [\pi_{t+1}] = \alpha_t^m + \phi\pi_t = (1 - \phi) \bar{\pi}_t + \phi\pi_t. \quad (21)$$

We allow the evolution of beliefs about α_t^m and $\bar{\pi}_t$ 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 Survey of Consumers (SOC) mean inflation forecast is well described by a constant gain learning algorithm. For the signal component, we assume that beliefs could be partly shaped 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 yields updating rules for α_t^m and $\bar{\pi}_t$ that are a weighted averages of two terms:

$$\alpha_t^m = (1 - \gamma^T) \left[\underbrace{\alpha_{t-1}^m + \gamma (\pi_t - \phi\pi_{t-1} - \alpha_{t-1}^m)}_{\alpha_t^{mCG}} \right] + \gamma^T \left[(1 - \phi) \pi_{\xi_t}^T \right] \quad (22)$$

$$\bar{\pi}_t = (1 - \gamma^T) \left[\underbrace{\bar{\pi}_{t-1} + \gamma (1 - \phi)^{-1} (\pi_t - \phi\pi_{t-1} - (1 - \phi) \bar{\pi}_{t-1})}_{\bar{\pi}_t^{CG}} \right] + \gamma^T \left[\pi_{\xi_t}^T \right] \quad (23)$$

The first terms in square brackets, α_t^{mCG} and $\bar{\pi}_t^{CG}$, 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 $\bar{\pi}_{t-1}$ are updated given new information, π_t . The second term

²⁶In principle one could introduce learning about ϕ as well. We forgo doing this in order 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.

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 either 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:

$$\mathbb{E}_t^m (\tilde{y}_{t+1}) = \varrho \tilde{y}_{t-1}. \quad (24)$$

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 Online Appendix proves that monetary neutrality holds in the long run.

Using equations (21), (23), and (24), we substitute out $\mathbb{E}_t^m [\pi_{t+1}]$, $\bar{\pi}_t$, and $\mathbb{E}_t^m (\tilde{y}_{t+1})$ in the model equations (18), (19), and (20) to obtain the following system of equations that must hold in equilibrium:

1. Real activity

$$\tilde{y}_t = \varrho \tilde{y}_{t-1} - \sigma [i_t - \phi \pi_t - (1 - \phi) \bar{\pi}_t - r_{ss}] + f_t. \quad (25)$$

2. Phillips curve:

$$\pi_t = \bar{\pi}_t + \frac{\kappa}{1 - \beta\phi} [\tilde{y}_{t-1} - \tilde{y}_{t-1}^*]. \quad (26)$$

3. Monetary policy rule with changes in target:

$$\begin{aligned} i_t - (r_{ss} + \pi_{\xi_t}^T) &= (1 - \rho_{i,\xi_t}) \left[\psi_{\pi,\xi_t} (\pi_t - \pi_{\xi_t}^T) + \psi_{\Delta y,\xi_t} (\tilde{y}_t - \tilde{y}_{t-1}) \right] \\ &+ \rho_{i,\xi_t} \left[i_{t-1} - (r_{ss} + \pi_{\xi_t}^T) \right] + \sigma_i \varepsilon_i, \quad \varepsilon_i \sim N(0, 1). \end{aligned} \quad (27)$$

4. Law of motion for f_t :

$$f_t = \rho_f f_{t-1} + \sigma_f \varepsilon_f, \quad \varepsilon_f \sim N(0, 1). \quad (28)$$

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

$$\tilde{y}_t^* = \rho_{y^*} \tilde{y}_{t-1}^* + \sigma_{y^*} \varepsilon_{y^*}, \quad \varepsilon_{y^*} \sim N(0, 1). \quad (29)$$

6. Perceived trend inflation:

$$\bar{\pi}_t = [1 - \gamma^T] [\bar{\pi}_{t-1} + \gamma(1 - \phi)^{-1} (\pi_t - \phi \pi_{t-1} - (1 - \phi) \bar{\pi}_{t-1})] + \gamma^T \pi_{\xi_t}^T. \quad (30)$$

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

For the learning part, 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, central banks have for decades 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 “economic conditions are likely to warrant exceptionally low levels for the federal funds rate *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 that the Committee 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 part, 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 one that is 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 and 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 over-extrapolation implies that the investor will always be surprised whenever there is a switch to a new policy rule after many periods in an 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). All 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 Online 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 a 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:

$$\mathbf{H}^p = \left[\begin{array}{cc|cc} p_{11} & 0 & 0 & p_{14} \\ 0 & p_{22} & p_{23} & p_{24} \\ \hline 0 & p_{32} & p_{33} & 0 \\ p_{41} & p_{42} & 0 & p_{44} \end{array} \right] \quad (31)$$

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, while $\xi_t^p = 4$ 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.²⁷

Since the asset pricing agent knows the structure of the macro block and can observe 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 $\{\xi_t, \xi_{t-1}, \dots\}$

²⁷We 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.

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 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 above restrictions on the perceived transition matrix \mathbf{H}^p will have implications for how beliefs evolve over time. The Online 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 (13)-(17) in the asset pricing block are satisfied, equations (25)-(30) 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 (31), under the parameter restrictions given in the previous subsection.

4.2 Model Solution and Estimation

The model is solved in two steps. First, we solve for the macro dynamics. This returns a MS-VAR in the macro block state vector $S_t = [y_t, y_t^*, \pi_t, i_t, r_t^*, \bar{\pi}_t, f_t]'$. Second, conditional on this

solution and on the probability assigned by the asset pricing 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 asset pricing relations (13)-(17), 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 $\tilde{S}_t = [S_t, m_t, pd_t, \mathbb{E}_t^p(m_{t+1}), \mathbb{E}_t^p(pd_{t+1})]$.

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 federal funds rate (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 asset pricing 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 asset pricing 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 asset pricing 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. We show below 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, $\pi_{\xi_t}^T$, 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 3. Importantly, however, the parameters characterizing the policy regimes as well as the transition matrix are freely estimated.²⁸ 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. Not only are the parameters permitted to show no shift across the previously estimated regimes,

²⁸We use the regime sequence $\hat{\xi}^T = \{\hat{\xi}_1, \dots, \hat{\xi}_T\}$ that is most likely to have occurred, given our estimated posterior mode parameter values for θ . See the Appendix for details.

there is no prior imposed that predisposes the estimates to implying that they are more likely to shift one direction or the other.

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 zero lower bound (ZLB) period. The 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 either with an estimated shadow rate, or with the one-year Treasury bill rate. The reason for this is that the policy rule regime changes we uncover are not mainly tied to the ZLB period.

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

4.3 Model Estimation Results

This section presents results from the model estimation. The first subsection discusses the parameter and latent state estimates. The next two subsections discuss the model implications for how monetary policy regime shifts affect the real interest rate, asset valuations, and return premia on the aggregate stock market.

4.3.1 Parameter and Latent State Estimates

Table 6 reports the prior and posterior distributions for the macro block model parameters. For the policy rule parameter estimates for $\pi_{\xi_t}^T$, 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 3) are associated with quantitatively large changes in the estimated policy rule. Specifically, the hawkish high *mps* regime is characterized by what we will call a *hawkish monetary policy rule* with lower inflation target $\pi_{\xi_t}^T$ 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 1, the threshold generally associated with the “Taylor principle” (Taylor (1993)), which prescribes that the central bank should raise nominal rates by more than one-for-one in response to deviations of inflation from target 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)). But we find here that, starting 2001:Q4, parameters shifted back to values consistent with accommodative monetary

policy. With the exception of a brief interlude from 2006:Q2-2008:Q2, the 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 7 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 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 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 non-policy-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 quite 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 very 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. Most important for the application here, 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 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 $\pi_{\xi_t}^T$ shown in Table 7. Although this parameter plays the role of an “inflation target” in the interest rate rule, unlike traditional New

Keynesian models, $\pi_{\xi_t}^T$ is not a value to which true inflation and inflation expectations in the model necessarily tend in the long-run. This happens 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 the inflation target can deviate substantially from actual inflation and inflation expectations for an extended period of time. For example, consider the value for $\pi_{\xi_t}^T$ under the dovish policy rule ($\pi_{\xi_t}^T = \pi_D^T$), in the post-millennial dovish subperiod. In this case π_D^T 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 or inflation expectations tend toward the estimated 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 (post-millennial) subperiods to share the same policy rule parameter values. Additional results (not reported) 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 $\pi_{\xi_t}^T$.²⁹

To interpret this result, note that the post-millennial 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 by comparison elevated. 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 post-millennial 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 8 reports the parameter values for the asset pricing 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 $\hat{l}_p = 5.8\%$ 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,

²⁹These additional results are based on a re-estimation of the model using the same regime sequence as in the baseline case, but allowing the policy rule parameters to differ freely across the early- and late- (post-millennial) dovish subperiods.

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 declines (increases). This drives up (down) the SDF strongly due to a large change in $\vartheta_{p,t}$, as seen from equation (11). 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.³⁰ The resulting estimate of $\widehat{lp} = 5.8\%$ is admittedly high relative to estimates in e.g., 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 correspondingly reduce \widehat{lp} . We forgo doing so in this paper in order to keep the model simple focusing on fluctuations in the equity premium arising from distorted beliefs about monetary policy.

But note that the model generates a plausible mean and volatility for the real interest rate (Figure 5). Taken together, this shows that the model accounts for the behavior of equity premia and real interest rates reasonably well. Table 8 also reports parameters of the perceived transition matrix. The key implications of these parameters for the model implications come through the affect they have on the evolution of beliefs, discussed below. For now we can observe that the estimated perceived probability of switching out of a long-lasting regime of one type into the short-lasting of the other type is close to unity in both cases, i.e., $p_{23}/(p_{23} + p_{24}) = 0.986$, $p_{41}/(p_{41} + p_{42}) = 0.9999$. 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 To illustrate implications of the perceived transition probabilities, Figure 6 reports the 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 approximately one that she is in the short-lasting dovish regime. The left panel reports the perceived probability at each point in time of being in the long-lasting

³⁰Inflation risk also contributes little to the equity premium since the conditional variance and covariance terms in (12) involving inflation are small due to the high level of inertia in macro-agent inflation expectations.

hawkish regime (blue solid line) and the long-lasting dovish regime (red dashed line). The right panel reports 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.

From the left panel we 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 1. Likewise, as the economy switches into the hawkish policy rule under Volcker, the agent initially places very low probability on the switch enduring. This can be seen from the right panel, which shows that, immediately after the switch, the probability of being in the hawkish regime in one year's time is less than 0.3, and in 20 year's time is effectively zero. The agent only eventually comes to see the hawkish policy rule as a long-lasting feature, after observing years of continuously restrictive policy. This shows that long-lasting expectations are "sticky;" they only change when agents become convinced that monetary policy 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 a sudden but modest jumps in response to the changes in short-term expectations, followed by further changes investors revise the probability of remaining in the new policy framework.

On the other hand, when regime shifts are more frequent, even expectations about the long-term can move quickly. Consider the time after the switch out of the long hawkish subperiod that extended 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.

4.3.2 Conduct of Monetary Policy, the Real Interest Rate, and Asset Valuation

We now investigate the importance of changes in the conduct of monetary policy for 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 the results of a simulation in which the observables and estimated state vector are taken as they were at the beginning of our sample with all Gaussian shocks are shut down. Thus, the only source of variation in the variables plotted in the figure arises from changes in the conduct of monetary policy, i.e., from changes in the policy rule parameters. These isolated movements are the only

ones, other than the policy shock, that the model stipulates can be purely the result of the behavior of the monetary authority. Monetary policy also affects the propagation of the two non-policy shocks, but these effects are not solely the result of changes monetary policy.

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 $\bar{\pi}_t = \pi_{\xi_t}^T$. As explained above, this corresponds to the case where $\gamma^T = 1$. 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 mostly driven 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 drop in inflation that occurred in the early 1980s. Instead, inflation would have kept increasing. What is more relevant and less obvious is the behavior of the real federal funds rate. 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 evident 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 shown in the blue solid line with the black dotted line showing the counterfactual in which the inflation target changes but there are no accompanying changes in the activism coefficients, it is evident 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:Q4-2009: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 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 stance of policy 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

The Secular Decline in Real Interest Rates Figure 8 studies the implications of our model for the contribution of regime changes in the conduct monetary policy in the secular decline of 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 afterwards, 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. Such changes 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 are not associated with a shift in the policy rule parameters, but are instead attributed to a combination of the model’s Gaussian shocks. By contrast, a substantial portion of the secular *trend* downward in real rates since the early 1980s is attributable to regime changes in the conduct of monetary policy. The peak of the real federal funds rate in our sample is 10.22%, which occurs in 1981:Q3. Since that time, the real federal funds rate 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, regime changes in monetary policy generate a peak in the real federal funds rate 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 reports 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 reports 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 identified 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 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 of a pure regime change 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, it is a truism that 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 agent’s 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 again reports impulse responses implied by the model to policy rule regime changes, under different counterfactual simulations. The top row reports responses to a regime change from the dovish to the hawkish regime (curbing inflation), with all Gaussian shocks set to zero. The bottom row reports the analogous responses to a regime change from hawkish to dovish (lifting inflation). The blue (solid) lines in all figures of both rows report the responses in the baseline model. The red (dotted) line shows a counterfactual in which the asset pricing 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 regimes. The magenta (dashed-dotted) line is a counterfactual that combines AP rational expectations with the case where the macro agent’s perceived trend value of inflation, $\bar{\pi}_t$, coincides in every period with the inflation target, $\pi_{\xi_t}^T$.

Note that the asset pricing 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 have a large role in the responses of asset valuations ($-cay_t$), as shown in the last column. In the baseline model the wealth ratio responds to policy regime changes by jumping modestly on impact. 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 first part 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 evident that valuation ratios in the baseline model initially under-react to the policy rule regime shifts. Indeed, under AP rational expectations, the wealth ratio jumps on impact to its maximal response in almost one period. The wealth ratio nonetheless moves smoothly even under rational expectations back toward its steady state value, a reflection of the adaptive learning mechanism in the macro block that drives the persistent behavior of the real interest rate observed in in the third column of Figure 10.

The role of over-extrapolation can be seen by comparing the red dotted (AP rational expectations) line to the black dashed “No AP learning” line, in which we keep the fading memory distortion but shut off learning. Over-extrapolation 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 over-extrapolation, the baseline wealth ratio responds to regime shifts in the policy rule by initially under-reacting but eventually over-reacting vis-a-vis the case of AP rational expectations.

The magenta (dashed-dotted) line of Figure 10 combines AP rational expectations with $\bar{\pi}_t = \pi_{\xi_t}^T$ 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 *increases* the wealth ratio, while a shift to the dovish policy rule slightly *decreases* it. Although policy regime changes in this case have no affect 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 increases asset valuations. The opposite happens in a switch to the dovish rule. This demonstrates that, 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 $m\psi_t$ over our sample. To form a direct comparison with the data, this figure repeats the information from Figure 3, which plots the corresponding series $-cay_t^{MS}$ and $m\psi_t$ from the data, along with horizontal lines that show the conditional means for these series in each regime. The figure also reports the component of the model-implied values for $-cay_t$ and the $m\psi_t$ that we estimate are attributable solely to regime changes in the conduct of monetary policy, shown as black dashed-dotted lines. The magenta dotted lines report the same components under the AP rational expectations counterfactual.

Figure 11 shows that fluctuations in the model-implied $-cay_t$ and $m\psi_t$ attributable to regime changes in monetary policy match well the movements in the conditional means of these variables found in the data. The black dashed lines in both panels fluctuate around the data means across the regime subperiods. These lines 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 $m\psi_t$ that are attributable to shifts in the policy rule, consistent with the Markov-switching specification. For the wealth ratio, however, 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 $m\psi_t$, as it does for example after the switch to the first hawkish subperiod of the sample. Since both beliefs and the interest rate rule are history dependent, this does not happen after every switch, however.

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 gradually reverts toward its steady state value as the real interest rate reverts.

Summarizing the lessons from the previous two figures, we have shown 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 asset pricing agent immediately and fully revises her expectations. Without over-extrapolation, the wealth ratio would not respond to regime changes in monetary policy by overshooting the case where the agent is fully aware of the underlying transition matrix. This pattern has important implications for the PDV of equity return premia. We turn to this next.

4.3.3 Monetary Policy and Equity Return Premia

Figure 12 plots the estimated PDV of forecasted excess returns (return premia) for the portfolios analyzed in Section 3.3 (red dashed line, right axis) together with the component of the real interest rate attributable to regime changes in the monetary policy rule (solid line, left axis). There is discernible positive comovement between the two series, implying that low interest rates associated with dovish monetary policy are also associated with low return premia.

To confirm this visual impression, the second column of Table 9 reports the correlation between the PDV of return premia of the different equity portfolios and the component of the real interest rate driven solely by changes in the policy rule, denoted RIR_t^{MPR} in the table. This correlation is high in all cases, equal to 0.88 for the momentum spread, 0.82 for the market excess return, 0.75 for the value spread in the small size quintile, and 0.70 for the value spread in the second smallest size quintile. Importantly, the correlation of premia with RIR_t^{MPR} is systematically larger than that with the residual component of the real interest rate, $RIR_t - RIR_t^{MPR}$, and thus also larger than the correlation of premia with the real interest rate itself (RIR_t). This shows that shifts in the monetary policy stance play an important role in generating the positive comovement between premia and the real interest rate in the data, but that other movements in the real rate do not share this property. This may be because persistent low- or high-interest rate environments that are the consequence of shifts in the conduct of monetary policy have effects that last for decades, in contrast to movements in real rates driven by other factors.

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 dataset. To ensure that the artificial samples 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 variables include the PDV of return premia on the stock market, which are computed from 20,000 Bayesian estimations of an MS-VAR using the same methodology that produced the PDV of return premia from the data reported in Figure 4. For each t , we report the average (across simulations) of the model-implied PDV in both panels of Figure 13. The left panel of Figure 13 superimposes the average (across simulations) of the model-implied log wealth ratio $-cay_t$, while the right panel reports the same for the model-implied real interest rate. Since we average across sample paths that differ only by the Gaussian shocks, the plotted series reveal fluctuations that are attributable 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 of return premia, while hawkish monetary policy is associated with a low wealth ratio and high PDV, consistent with the data. When the economy switches into the first hawkish subperiod of the sample, coinciding with the Volcker disinflation, the PDV of return premia at first declines slightly before eventually rising to a new, significantly higher level. Meanwhile, 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, as investors take time to learn about the persistence of the shift.

To understand why dovish monetary policy is associated with a low PDV of return premia and vice-versa for hawkish policy, consider what an econometrician armed with our historical sample who computes the PDVs using an MS-VAR taking into account the probability of future regime shifts would find. Because of learning, an econometrician 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 over-extrapolate and eventually come to believe that the regime will persist indefinitely, asset values ultimately over-react 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 *long*-horizon return premia are always high in hawkish regimes, since an econometrician would find a predictable jump upward in returns when the inevitable switch back to dovish policy occurs.

Since the PDV 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 over-extrapolation on long-horizon premia. This

happens 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 the PDV moves monotonically after the subsequent switches in the sample. Regardless of the initial trajectory of the PDV, the model implies that the PDV of return premia is always higher on average in hawkish regimes than in dovish regimes, consistent with the data.

To investigate this further, we use simulated data to compute the model-implied posterior probability that the regime average PDV of return premia is lower in the dovish regime than in the hawkish regime, following the same procedure used to do so in historical data (Table 5).³¹ For the baseline model, this probability is 73%, a magnitude in the ballpark of the 81%, 70%, and 64% probabilities estimated from historical data for the market premium, the S1 BM spread, and the S2 BM spread, respectively. 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.³² If instead we consider a counterfactual that retains the over-extrapolation 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 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.

To summarize, in the early stages of a hawkish (dovish) regime, investor learning shows up as initial under-reaction of asset values and low (high) short-horizon return premia. Over-extrapolation shows up as eventual over-reaction of asset values and high (low) long-horizon return premia. In terms of the model’s implications for the PDV of all future return premia, there is thus a tug-of-war between the learning and over-extrapolation aspects of investor beliefs. The effect of learning is outweighed by the effect of over-extrapolation because the speed of learning is relatively quick compared to the persistence of monetary policy regimes.

5 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 concurrently with equally important shifts in the level of the short-term real interest rate in excess of a widely used measure of

³¹These probabilities are obtained as the fraction of draws from the posterior distribution for which the average PDV is lower in the dovish regime than in the hawkish regime.

³²Even 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.

the “natural” rate of interest, a variable we refer to as the monetary policy spread, 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 of 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 attributable to monetary policy, we solve and estimate a novel macro-finance 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, on equity markets, and on return premia.

The model and its estimates speak to the origins of persistently declining real interest rates over the past 40 years. A striking finding is that two-thirds of the downward trajectory in real rates observed since the early 1980s can be attributed to monetary policy, i.e., to regime changes in the conduct of policy. This occurs 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:Q2-2008:Q2.

The model is silent on one aspect of our empirical evidence, namely that spreads in the cross-section of equity characteristic portfolio returns are significantly compressed in dovish subperiods compared to hawkish ones. Future research could explore the possibility that a reach-for-yield in low interest rate environments might be amplified by the type of fading memory belief distortion considered in the model of this paper. For example, as interest rates move from high to low in a switch from hawkish to dovish monetary policy, a reach-for-yield could explain why portfolios with lower market-to-book ratios experience a greater increase in market value than those with higher market-to-book ratios, compressing return differentials. But it is unclear how quantitatively important that mechanism alone could be, if regime shifts

are not perceived to be sufficiently persistent in the first place. One possibility is that reaching-for-yield is amplified by a fading memory distortion in which all investors over-estimate the persistence of policy regime shifts.

References

- ACHARYA, V. V., AND H. NAQVI (2016): “On reaching for yield and the coexistence of bubbles and negative bubbles,” *Available at SSRN 2618973*.
- ANDREWS, D. (1991): “Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation,” *Econometrica*, 59, 817–858.
- ANGELETOS, G.-M., Z. HUO, AND K. A. SASTRY (2020): “Imperfect Macroeconomic Expectations: Evidence and Theory,” Discussion paper, National Bureau of Economic Research.
- BANSAL, R., AND H. ZHOU (2002): “Term structure of interest rates with regime shifts,” *The Journal of Finance*, 57(5), 1997–2043.
- BARRO, R. J., AND A. MOLLERUS (2014): “Safe assets,” National Bureau of Economic Research No. w20652.
- BERNANKE, B. S., AND K. N. KUTTNER (2005): “What Explains the Stock Market’s Reaction to Federal Reserve Policy?,” *Journal of Finance*.
- BIANCHI, F. (2013): “Regime Switches, Agents’ Beliefs, and Post-World War II U.S. Macroeconomic Dynamics,” *The Review of Economic Studies*, 80(2), 463–490.
- (2016): “Methods for measuring expectations and uncertainty in Markov-switching models,” *Journal of Econometrics*, 190(1), 79–99.
- BIANCHI, F., S. LUDVIGSON, AND S. MA (2020): “Belief Distortions and Macroeconomic Fluctuations,” NBER working paper 27406.
- BIANCHI, F., AND L. MELOSI (2016): “Modeling the evolution of expectations and uncertainty in general equilibrium,” *International Economic Review*, 57(2), 717–756.
- BLACK, F. (1995): “Interest rates as options,” *the Journal of Finance*, 50(5), 1371–1376.
- BOYARCHENKO, N., V. HADDAD, AND M. C. PLOSSER (2016): “The Federal Reserve and market confidence,” .
- CABALLERO, R. J., AND E. FARHI (2014): “The Safety Trap,” Discussion paper, National Bureau of Economic Research No. w19927.
- CAMPBELL, J. Y. (2017): *Financial decisions and markets: a course in asset pricing*. Princeton University Press.

- CAMPBELL, J. Y., AND R. J. SHILLER (1989): “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors,” *Review of Financial Studies*, 1(3), 195–228.
- CAMPBELL, J. Y., AND T. VUOLTEENAHO (2004): “Bad Beta, Good Beta,” *American Economic Review*, 94(5), 1249–1275.
- CHOI, J., AND M. KRONLUND (2015): “Reaching for yield or playing it safe? Risk taking by bond mutual funds,” Available at SSRN: <https://ssrn.com/abstract=2669465>.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113(1), 1–45.
- CIESLAK, A., A. MORSE, AND A. VISSING-JORGENSEN (2015): “Stock returns over the FOMC cycle,” Available at SSRN 2687614.
- CIESLAK, A., AND A. SCHRIMPF (2019): “Non-monetary news in central bank communication,” *Journal of International Economics*, 118, 293–315.
- CLARIDA, R., J. GALL, AND M. GERTLER (2000): “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *Quarterly Journal of Economics*, 115(1), 147–180.
- COCHRANE, J. H. (2005): *Asset Pricing, Revised Edition*. Princeton University Press, Princeton, NJ.
- COHEN, R. B., C. POLK, AND T. VUOLTEENAHO (2003): “The value spread,” *The Journal of Finance*, 58(2), 609–641.
- COIMBRA, N., AND H. REY (2017): “Financial cycles with heterogeneous intermediaries,” Discussion paper, National Bureau of Economic Research.
- COOK, T. A. H. (1989): “The effect of changes in the federal funds rate target on market interest rates in the 1970s,” *Journal of Monetary Economics*, 24(3), 331–351.
- COSTA, O., M. FRAGOSO, AND R. MARQUES (2004): *Discrete-Time Markov Jump Linear Systems*. Springer, New York.
- DI MAGGIO, M., AND M. T. KACPERCZYK (2015): “The unintended consequences of the zero lower bound policy,” Columbia Business School Research Paper.
- DIAMOND, D. W., AND R. G. RAJAN (2012): “Illiquid banks, financial stability, and interest rate policy,” *Journal of Political Economy*, 120(3), 552–591.

- DRECHSLER, I., A. SAVOV, AND P. SCHNABL (2018): “A model of monetary policy and risk premia,” *Journal of Finance*, 71(1), 317–373.
- EVANS, G., AND S. HONKAPOHJA (2001): *Learning and Expectations In Macroeconomics*. Princeton University Press, Princeton, NJ.
- FAMA, E. F., AND K. R. FRENCH (2015): “A five-factor asset pricing model,” *Journal of Financial Economics*, 116(1), 1–22.
- FARHI, E., AND J. TIROLE (2012): “Collective moral hazard, maturity mismatch, and systemic bailouts,” *The American Economic Review*, 102(1), 60–93.
- GALÍ, J. (2015): *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.
- GERTLER, M., AND P. KARADI (2015): “Monetary policy surprises, credit costs, and economic activity,” *American Economic Journal: Macroeconomics*, 7(1), 44–76.
- GEWEKE, J. F. (1992): “Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments,” *Bayesian Statistics*, 4, 169–193.
- GILCHRIST, S., D. LÓPEZ-SALIDO, AND E. ZAKRAJŠEK (2015): “Monetary policy and real borrowing costs at the zero lower bound,” *American Economic Journal: Macroeconomics*, 7(1), 77–109.
- GREENWALD, D., M. LETTAU, AND S. C. LUDVIGSON (2019): “How the Wealth Was Won: Factors Shares as Market Fundamentals,” National Bureau of Economic Research Working Paper No. 25769.
- GÜRKAYNAK, R. S., B. SACK, AND E. SWANSON (2005): “The sensitivity of long-term interest rates to economic news: evidence and implications for macroeconomic models,” *The American Economic Review*, 95(1), 425–436.
- HALL, R. E. (2016): “Understanding the Decline in the Safe Real Interest Rate,” National Bureau of Economic Research No. w22196.
- HAMILTON, J. D. (1994): *Time Series Analysis*. Princeton University Press, Princeton, NY.
- HANSON, S. G., D. O. LUCCA, AND J. H. WRIGHT (2018): “Interest Rate Conundrums in the 21st Century,” Discussion paper, unpublished working paper.
- HANSON, S. G., AND J. C. STEIN (2015): “Monetary policy and long-term real rates,” *Journal of Financial Economics*, 115(3), 429–448.

- JAROCINSKI, M., AND P. KARADI (2019): “Deconstructing monetary policy surprises: the role of information shocks,” *American Economic Journal: Macroeconomics* forthcoming.
- KEKRE, R., AND M. LENEL (2019): “Monetary Policy, Redistribution, and Risk Premia,” Unpublished paper, Booth School of Business, University of Chicago.
- KIM, C.-J., AND C. R. NELSON (1999): *State-Space Models with Regime Switching*. The MIT Press, Cambridge, Massachusetts.
- KRISHNAMURTHY, A., AND A. VISSING-JORGENSEN (2012): “The aggregate demand for treasury debt,” *Journal of Political Economy*, 120(2), 233–267.
- LAUBACH, T., AND J. C. WILLIAMS (2003): “Measuring the natural rate of interest,” *Review of Economics and Statistics*, 85(4), 1063–1070.
- LETTAU, M., AND S. C. LUDVIGSON (2001): “Consumption, Aggregate Wealth and Expected Stock Returns,” *Journal of Finance*, 56(3), 815–849.
- (2005): “Expected Returns and Expected Dividend Growth,” *Journal of Financial Economics*, 76, 583–626.
- (2010): “Measuring and Modeling Variation in the Risk-Return Tradeoff,” in *Handbook of Financial Econometrics*, ed. by Y. Ait-Sahalia, and L. P. Hansen, vol. 1, pp. 617–90. Elsevier Science B.V., North Holland, Amsterdam.
- LETTAU, M., AND S. C. LUDVIGSON (2013): “Shocks and Crashes,” in *National Bureau of Economics Research Macroeconomics Annual: 2013*, ed. by J. Parker, and M. Woodford, vol. 28, pp. 293–354. MIT Press, Cambridge and London.
- LUBIK, T. A., AND F. SCHORFHEIDE (2004): “Testing for indeterminacy: an application to US monetary policy,” *American Economic Review*, pp. 190–217.
- LUCCA, D. O., AND E. MOENCH (2015): “The Pre-FOMC Announcement Drift,” *The Journal of Finance*, 70(1), 329–371.
- MALMENDIER, U., AND S. NAGEL (2011): “Depression babies: do macroeconomic experiences affect risk taking?,” *The Quarterly Journal of Economics*, 126(1), 373–416.
- (2016): “Learning from inflation experiences,” *The Quarterly Journal of Economics*, 131(1), 53–87.
- NEWKEY, W. K., AND K. D. WEST (1987): “A Simple, Positive Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55, 703–708.

- PARK, J. Y., AND M. OGAKI (1991): “Inference in Cointegrated Models using VAR prewhitening to estimate Shortrun Dynamics,” Rochester Center for Economic Research Working Paper No. 281, University of Rochester.
- PHILLIPS, P., AND S. OULIARIS (1990): “Asymptotic Properties of Residual Based Tests for Cointegration,” *Econometrica*, 58, 165–193.
- PIAZZESI, M., AND M. SCHNEIDER (2015): “Payments, credit and asset prices,” Unpublished manuscript, Stanford University.
- RAFTERY, A., AND S. LEWIS (1992): “How many iterations in the Gibbs sampler?,” *Bayesian Statistics*, 4, 763–773.
- RAJAN, R. G. (2006): “Has finance made the world riskier?,” *European Financial Management*, 12(4), 499–533.
- (2013): “A step in the dark: unconventional monetary policy after the crisis,” *Andrew Crockett Memorial Lecture, BIS, Basel*, 23.
- SIMS, C. A., AND T. ZHA (2006): “Were There Regime Switches in US Monetary Policy?,” *American Economic Review*, 91(1), 54–81.
- STOCK, J. H., AND M. W. WATSON (1993): “A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems,” *Econometrica*, 61, 783–820.
- TAYLOR, J. B. (1993): “Discretion versus policy rules in practice,” in *Carnegie-Rochester Conference Series on Public Policy*, vol. 39, pp. 195–214.
- VUOLTEENAHO, T. (2000): “Understanding the Aggregate Book-Market Ratio and its Implications to Current Equity-Premium Expectations,” Unpublished paper, Harvard University.
- WU, J. C., AND F. D. XIA (2016): “Measuring the macroeconomic impact of monetary policy at the zero lower bound,” *Journal of Money, Credit and Banking*, 48(2-3), 253–291.

Tables and Figures

	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

Table 1: 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 = $FFR_t - \text{Expected Inflation}_t - r_t^*$, 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:Q1-2017: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

Parameter Estimates: cay^{FC}

α	γ_a	γ_y
-0.8132	0.2048	0.8551
(0.0448)	(0.0186)	(0.0233)

Table 2: 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 parantheses. The two distinct values for the Markov-switch parameters are denoted with H and D subscripts to indicate hawkish or dovish values. The sample is quarterly and spans the period 1961:Q1 to 2017:Q3.

	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)

Table 3: 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.

Annualized Sharpe Ratios and Mean Returns					
Portfolio	SR	Mean	Portfolio	SR	Mean
Market	0.3665	0.0623	V-G (S1)	0.5966	0.0994
W-L	0.6261	0.1517	V-G (S2)	0.3279	0.0545

Table 4: The table reports annualized Sharpe ratios, "SR," and mean returns, "Mean," for the stock market and different portfolios. The Sharpe ratio is defined to be the unconditional mean return divided by the standard deviation of the portfolio return. The long-short portfolios "V-G" are the value-growth portfolios in a given size quintile, S1=smallest, S2= second smallest. long-short portfolios "W-L" are the winner-loser portfolio. For each size category, the return of the V-G portfolio portfolio return is the difference between the return on the extreme value (highest BM ratio) and the return of the extreme growth portfolio (lowest BM ratio). The return of the W-L portfolio return is the difference in returns between the extreme winner (M10) and the extreme loser (M1). All returns are computed at quarterly frequencies but the Sharpe ratios and mean returns are reported in annualized units. The sample spans the period 1964:Q1-2017:Q3.

Breaks in Market Premium and Book-Market Ratio Spreads				
	Market	W-L	Val-Gr (S1)	Val-Gr (S2)
Hawkish Regime	1.5896 (0.8960,2.2558)	4.0432 (2.8912,5.2514)	2.6647 (1.7120,3.6701)	1.4875 (0.6118,2.4214)
Dovish Regime	1.2848 (0.5652,1.9219)	3.4000 (2.2660,4.5196)	2.4460 (1.6679,3.2698)	1.3577 (0.6063,2.1193)
Diff-in-Diff	0.2987 (-0.0367,0.7089)	0.6011 (0.1462,1.2409)	0.1958 (-0.1865,0.6593)	0.1251 (-0.2431,0.5429)
Prob. decline	0.81	0.90	0.70	0.64
Odds ratio	4.26	9.00	2.33	1.78

Table 5: The first two rows report the regime averages for the present discounted value of market expected excess returns and the spread in the present discounted value of portfolio expected excess returns. The columns labeled "Val-Gr" report the spreads for portfolios sorted along the book-market dimension, in a given size category (extreme value minus extreme growth). The columns labeled "W-L" report the spreads for portfolios sorted along the recent past return performance dimension (extreme winner minus extreme loser). The row labeled "Diff-in-Diff" reports the difference between these spreads across the two wealth ratio/interest rate regimes. The numbers in each cell are the median values of the statistic from the posterior distribution while in parentheses we report 68% posterior credible sets. The last row reports the probability that premia decline when moving from the hawkish to dovish regime. These probabilities are obtained by computing 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.

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

Table 6: 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, and B for Beta, 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 mean and standard deviation, respectively. For the uniform distribution they correspond to the lower and upper bound. The last four rows report the standard deviations of the observation errors. The sample spans the period 1961:Q1-2017:Q3.

	<i>Mode</i>	<i>Mean</i>	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

Table 7: 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:Q1-2017: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
eq. premium	5.5%	$p_{23}/(p_{23} + p_{24})$	0.9864
		$p_{41}/(p_{41} + p_{42})$	0.9999

Table 8: 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 eq. premium are annualized log units. The sample is quarterly and spans the period 1961:Q1 to 2017:Q3.

	RIR_t	RIR_t^{MPR}	RIR_t^{RES}
Market Excess Return	0.43	0.82	0.09
Momentum Spread	0.52	0.88	0.18
Value Spread (S1)	0.51	0.75	0.23
Value Spread (S2)	0.46	0.70	0.20

Table 9: Corr. between PDV of market risk premium (rp) or portfolio rp spreads with RIR_t^{MPR} vs. movements in real rates driven by Gaussian shocks (RIR_t^{RES}). Market is the PDV of the market rp, Momentum W-L is the PDV of the difference in the Winner-Loser risk premia; Value Spread is the PDV of the difference in the High-Low book-market ratio portfolios in the smallest (S1) and next to smallest (S2) size quintiles. RIR_t is defined as FFR minus expected inflation (based on the model). The sample is quarterly and spans 1961:Q1 - 2017:Q3.





Figure 1: Real interest rate and Monetary Policy Spread (MPS). The real interest rate is the difference between the nominal federal funds rate (FFR) and expected inflation, where expected inflation is computed as a four quarter moving average of inflation. The monetary policy spread is defined as $MPS_t \equiv FFR_t - Expected\ Inflation_t - r_t^*$, where r_t^* is the natural rate of interest from Laubach and Williams (2003). The sample spans the period 1961:Q1-2017:Q3.

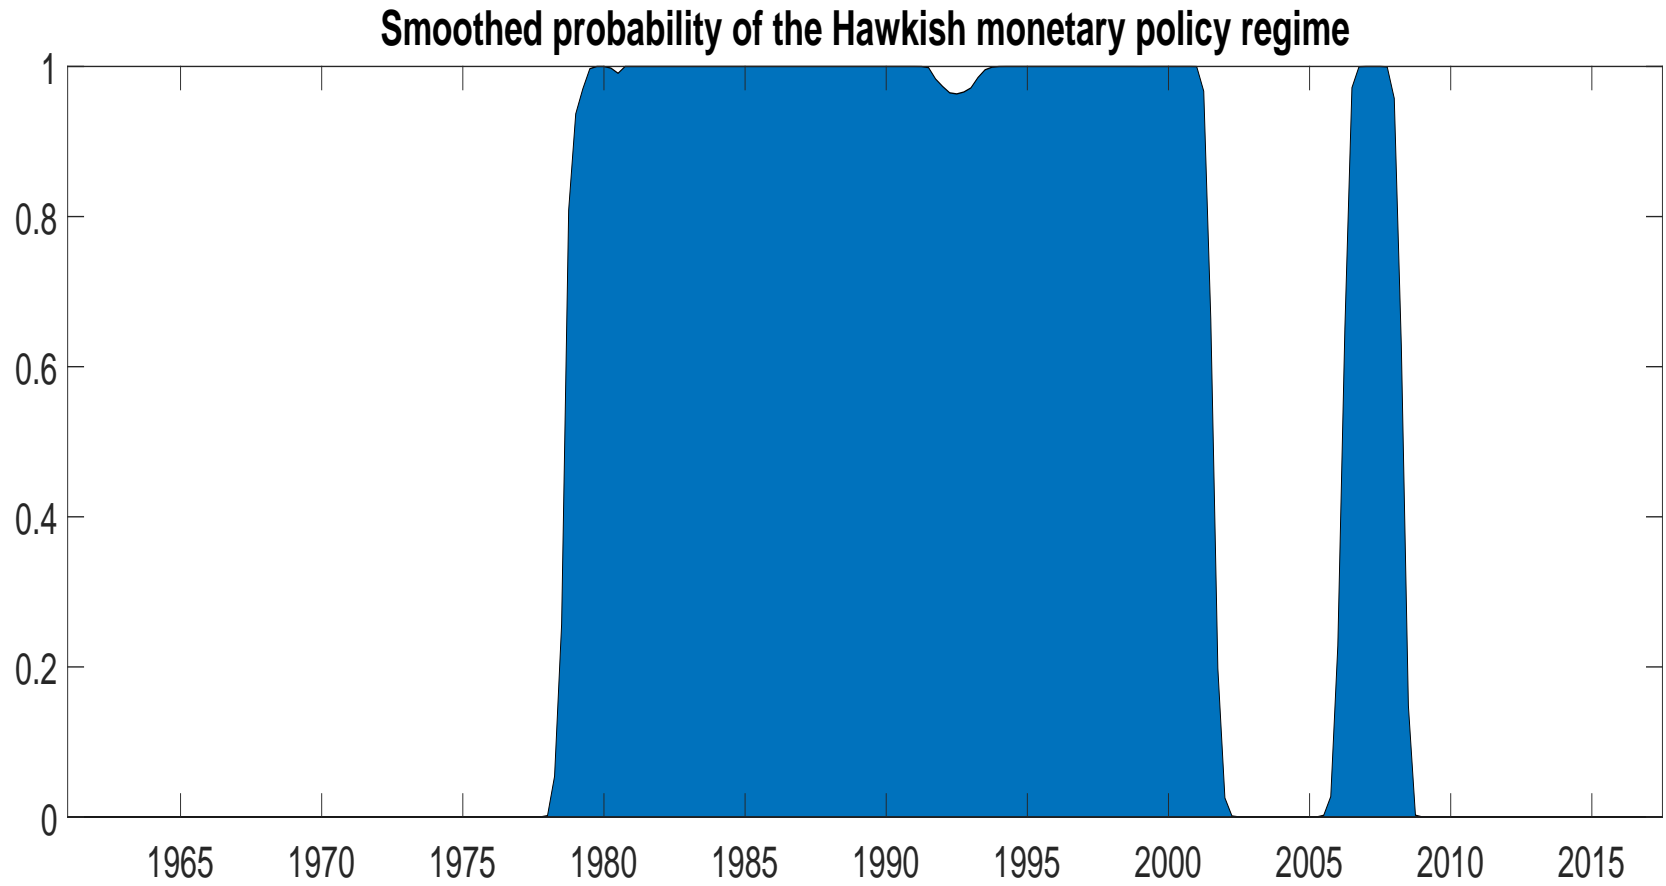


Figure 2: Regime probabilities. Smoothed probabilities of the Hawkish monetary policy regime. The sample is quarterly and spans the period 1961:Q1 to 2017:Q3.

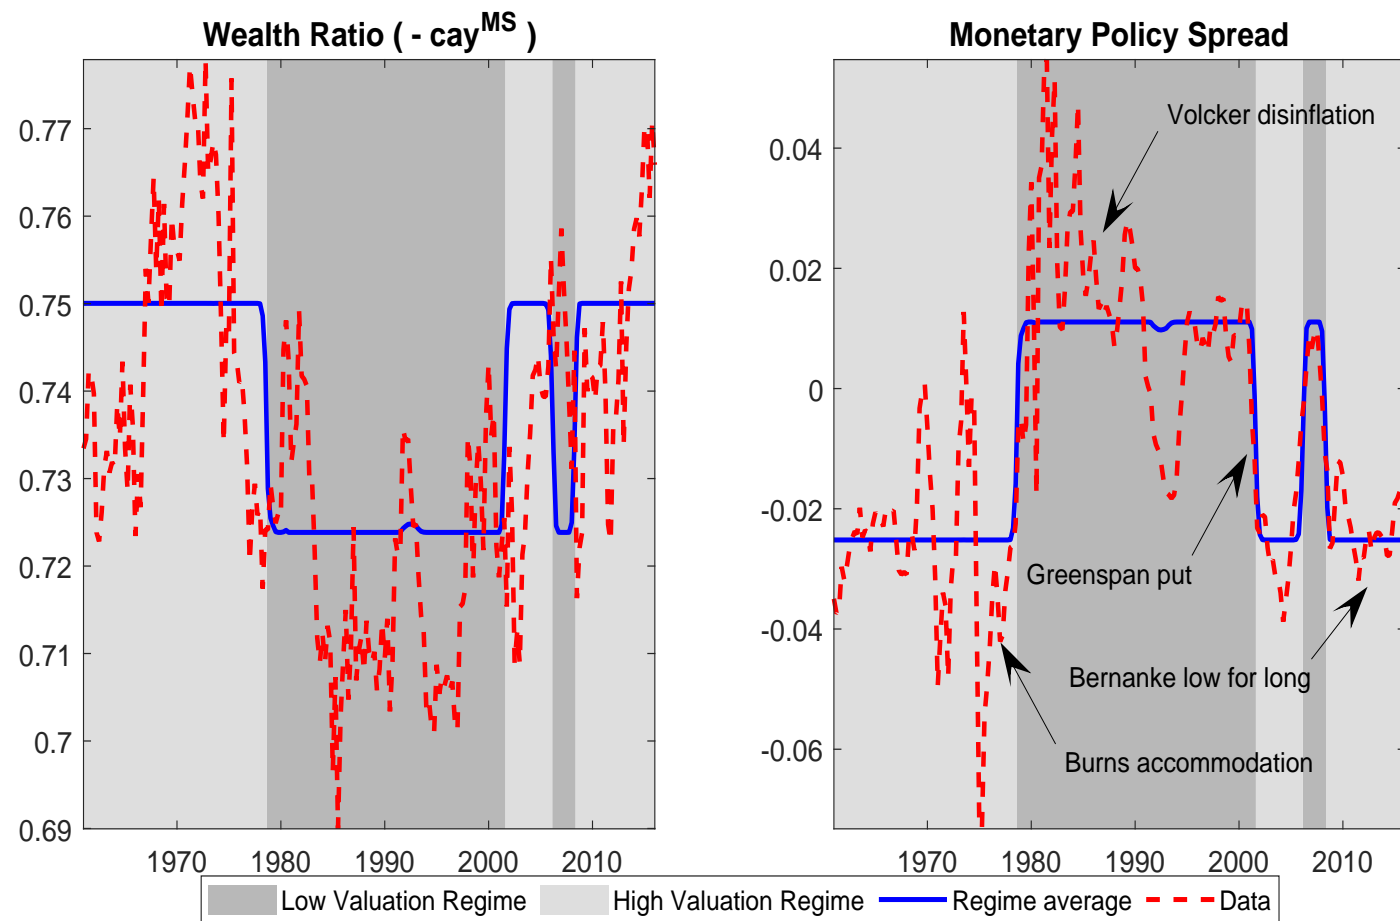


Figure 3: Wealth ratio and MPS in the data. Figure plots the wealth ratio ($-cay^{MS}$) and the monetary policy spread $MPS_t \equiv FFR_t - Expected\ Inflation_t - r_t^*$. The series for r_t^* is from Laubach and Williams (2003). The solid line corresponds to the estimated mean at the posterior mode. The sample spans 1961:Q1-2017:Q3.

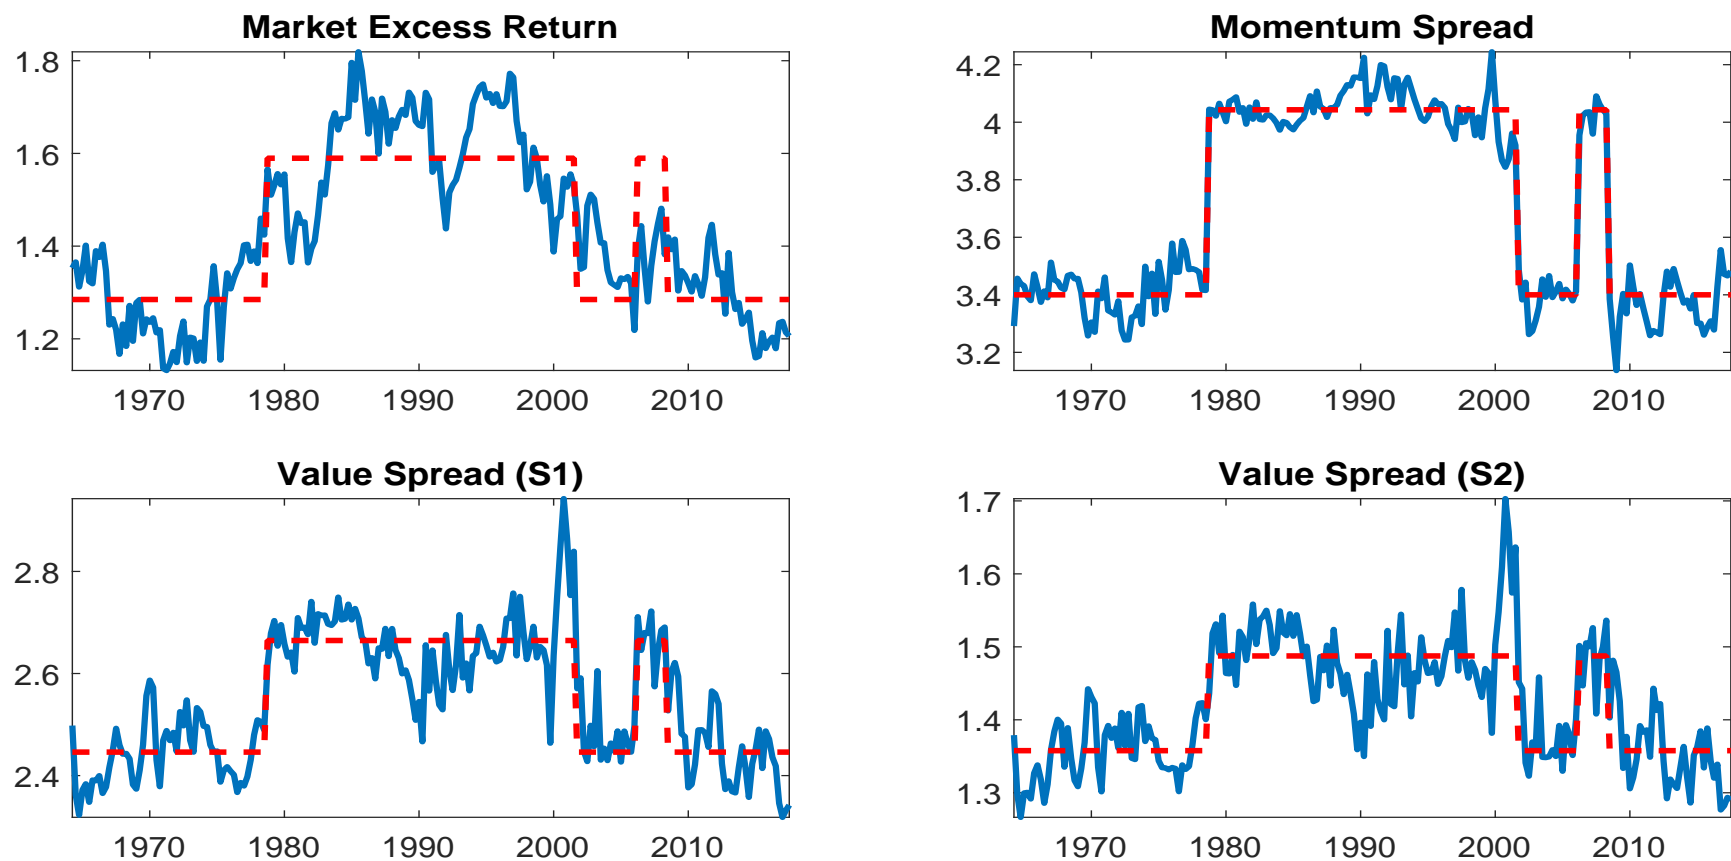


Figure 4: Evolution of Risk Premia in the data. The figure reports the evolution of the PDV of risk premia for the stock market and three different spread portfolios. The blue solid line reports the evolution of the risk premia over time, while the red dashed line corresponds to the conditional steady state of the PDV based on the regime in place. Both are computed by taking into account the possibility of regime changes. The sample spans the period 1964:Q1-2017:Q3.







Figure 5: Macroeconomic series and their filtered counterparts. The figure reports 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 - 2017:Q3.




Figure 6: Evolution of investor beliefs under learning. The left panel reports 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 reports 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 $\text{Pr} \approx 1$ to being in the short-lasting dovish regime. The sample spans 1961:Q1 - 2017:Q3.




Figure 7: The role of changes in the monetary policy rule and adaptive expectations. The blue line corresponds to the 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 the π^T . The hawkish regime has a low π^T and high activism against deviations from π^T . The sample spans 1961:Q1 - 2017:Q3.

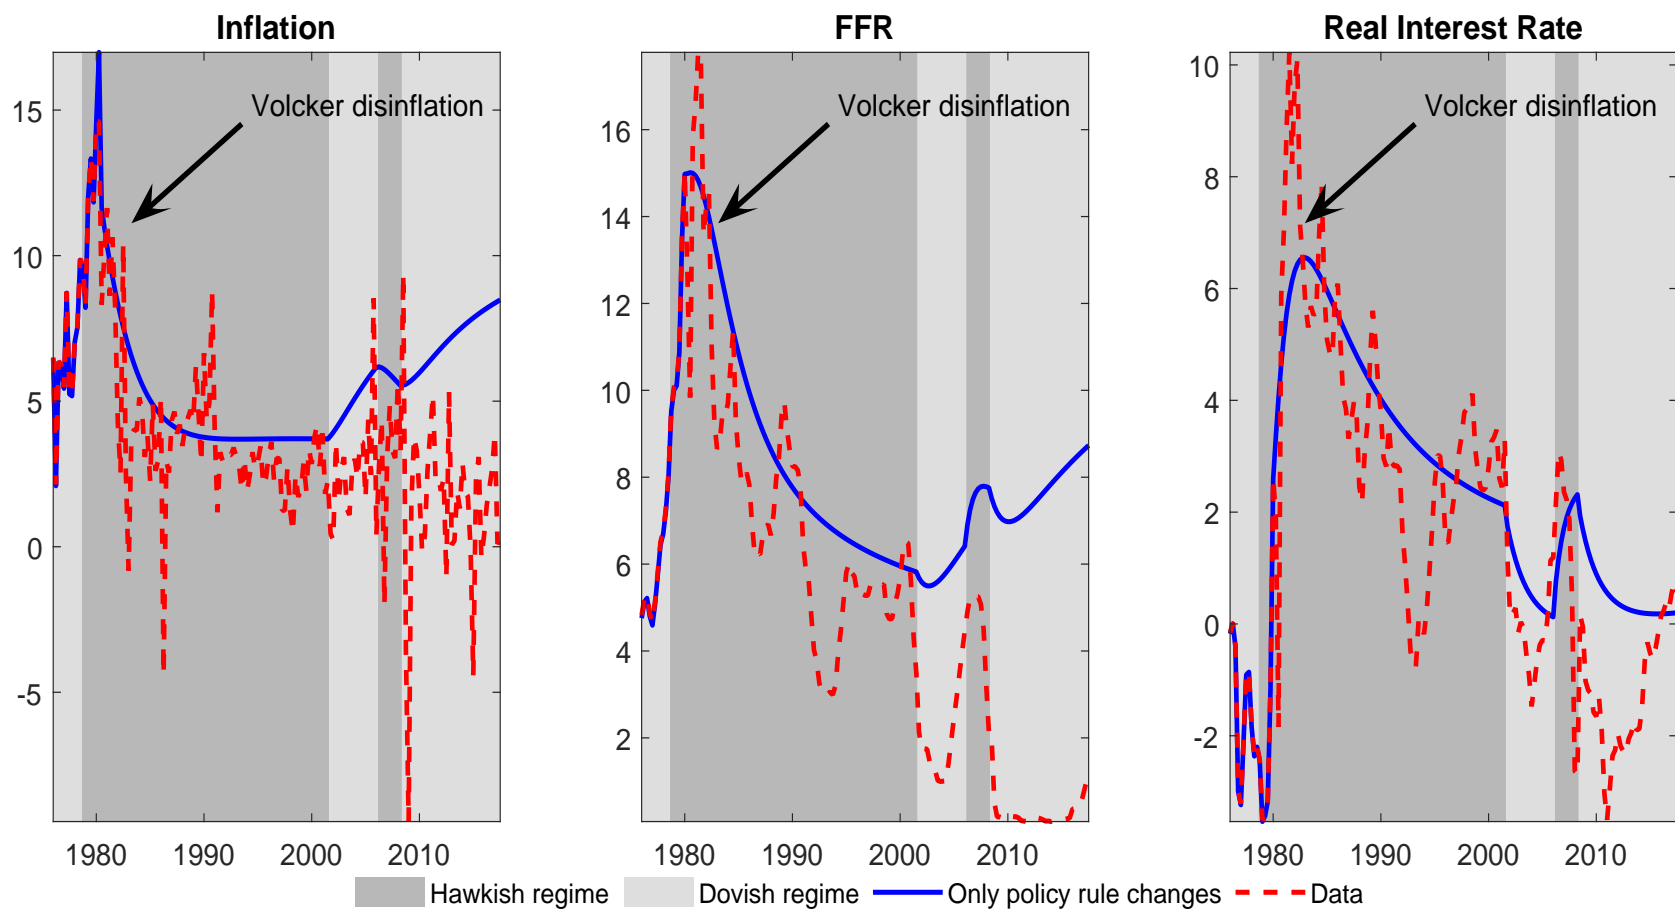


Figure 8: The Volcker disinflation. We start the economy as it was in 1980:Q1 and remove all Gaussian shocks that occurred after that period, but keep the estimated regime sequence. The dashed line corresponds to the data. The real interest rate is computed as the difference between the FFR and expected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

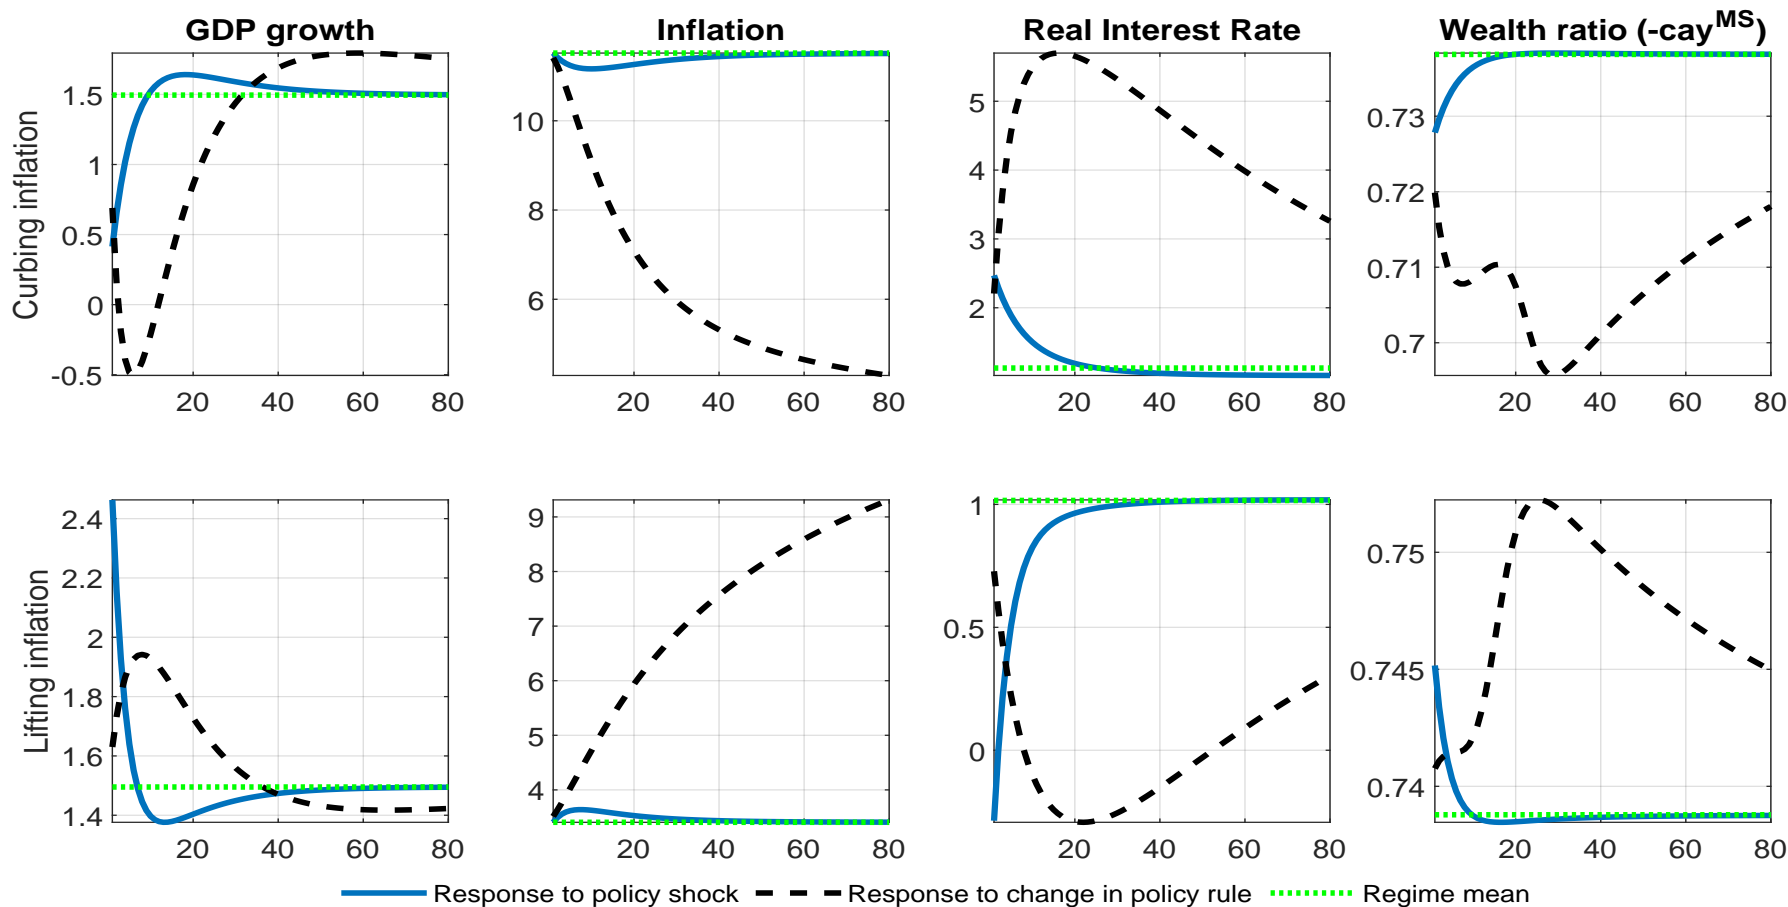


Figure 9: Regime changes versus policy shocks. Top row: curbing inflation. The economy is initially in the dovish regime. The blue solid line presents the evolution of the macro variables and the wealth ratio in response to a two standard deviation contractionary monetary policy shock and no regime change. 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.

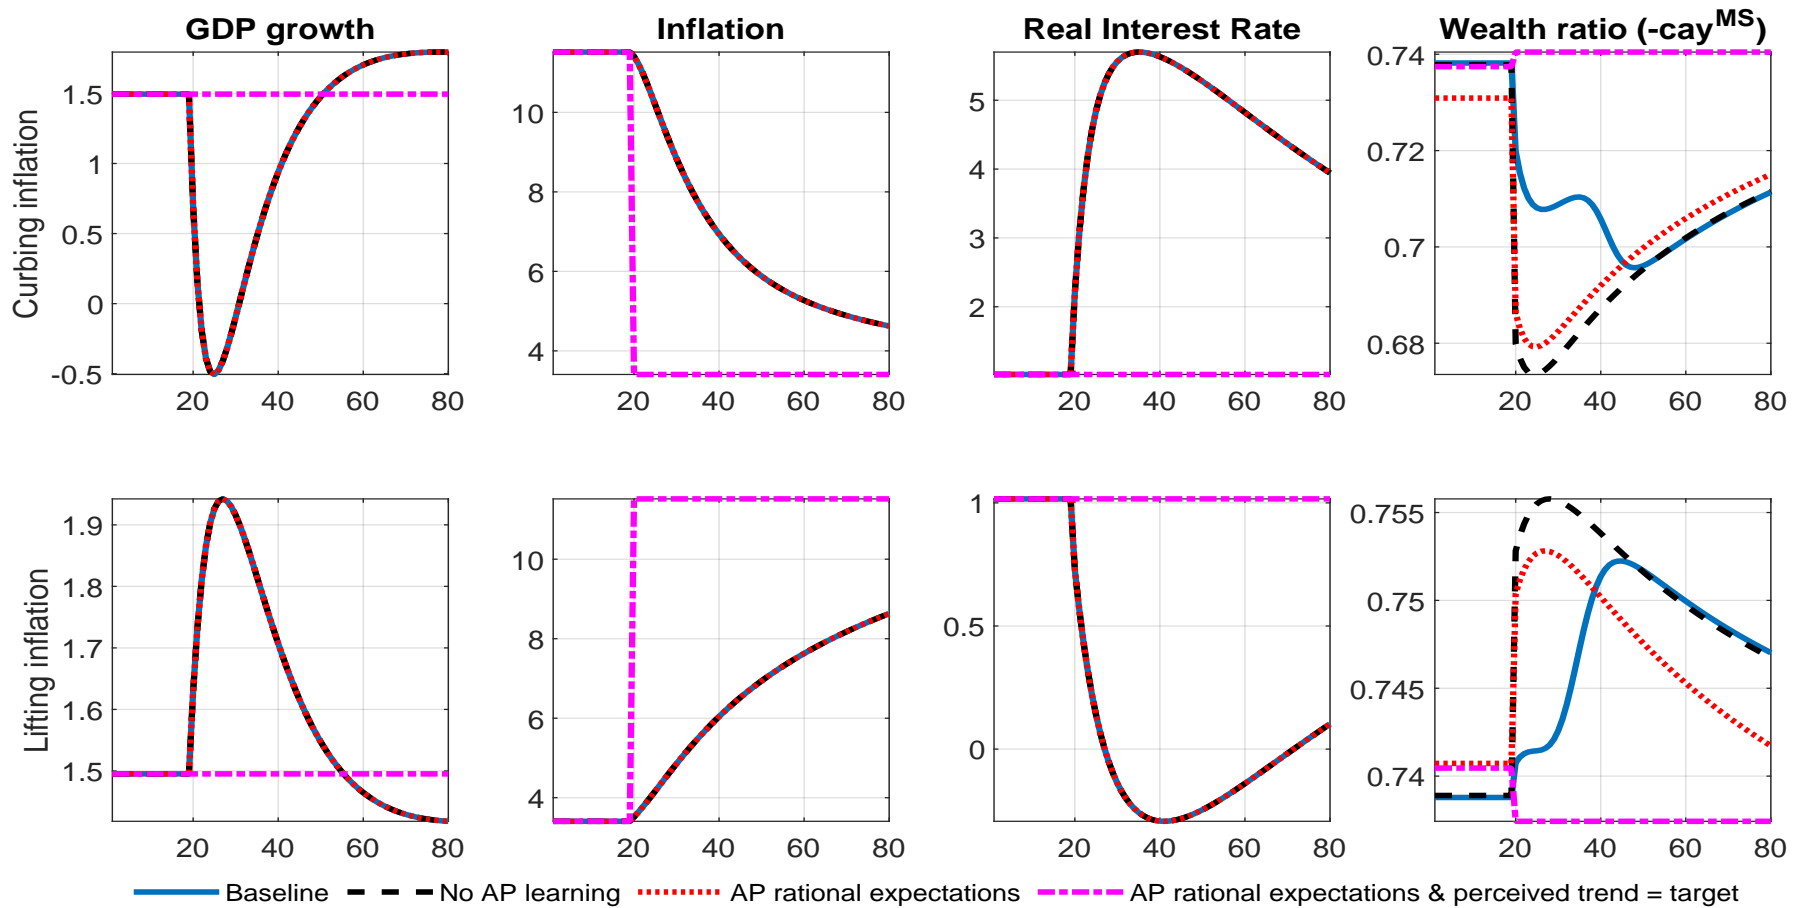


Figure 10: The role of AP learning and of macro stickiness. The blue solid line corresponds to the baseline model with learning of the asset pricing (AP) agent about the probability of moving across regime, over-reaction of the asset pricing agent about the persistence of regime changes, and adaptive expectations of the macro agent; the black dashed line shuts down learning of the asset pricing agent; the red dashed line is the case where the asset pricing 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 asset pricing 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.




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 agent observes the true transition matrix of the Markov-switching process controlling changes in monetary policy. The sample spans 1961:Q1 - 2017:Q3.

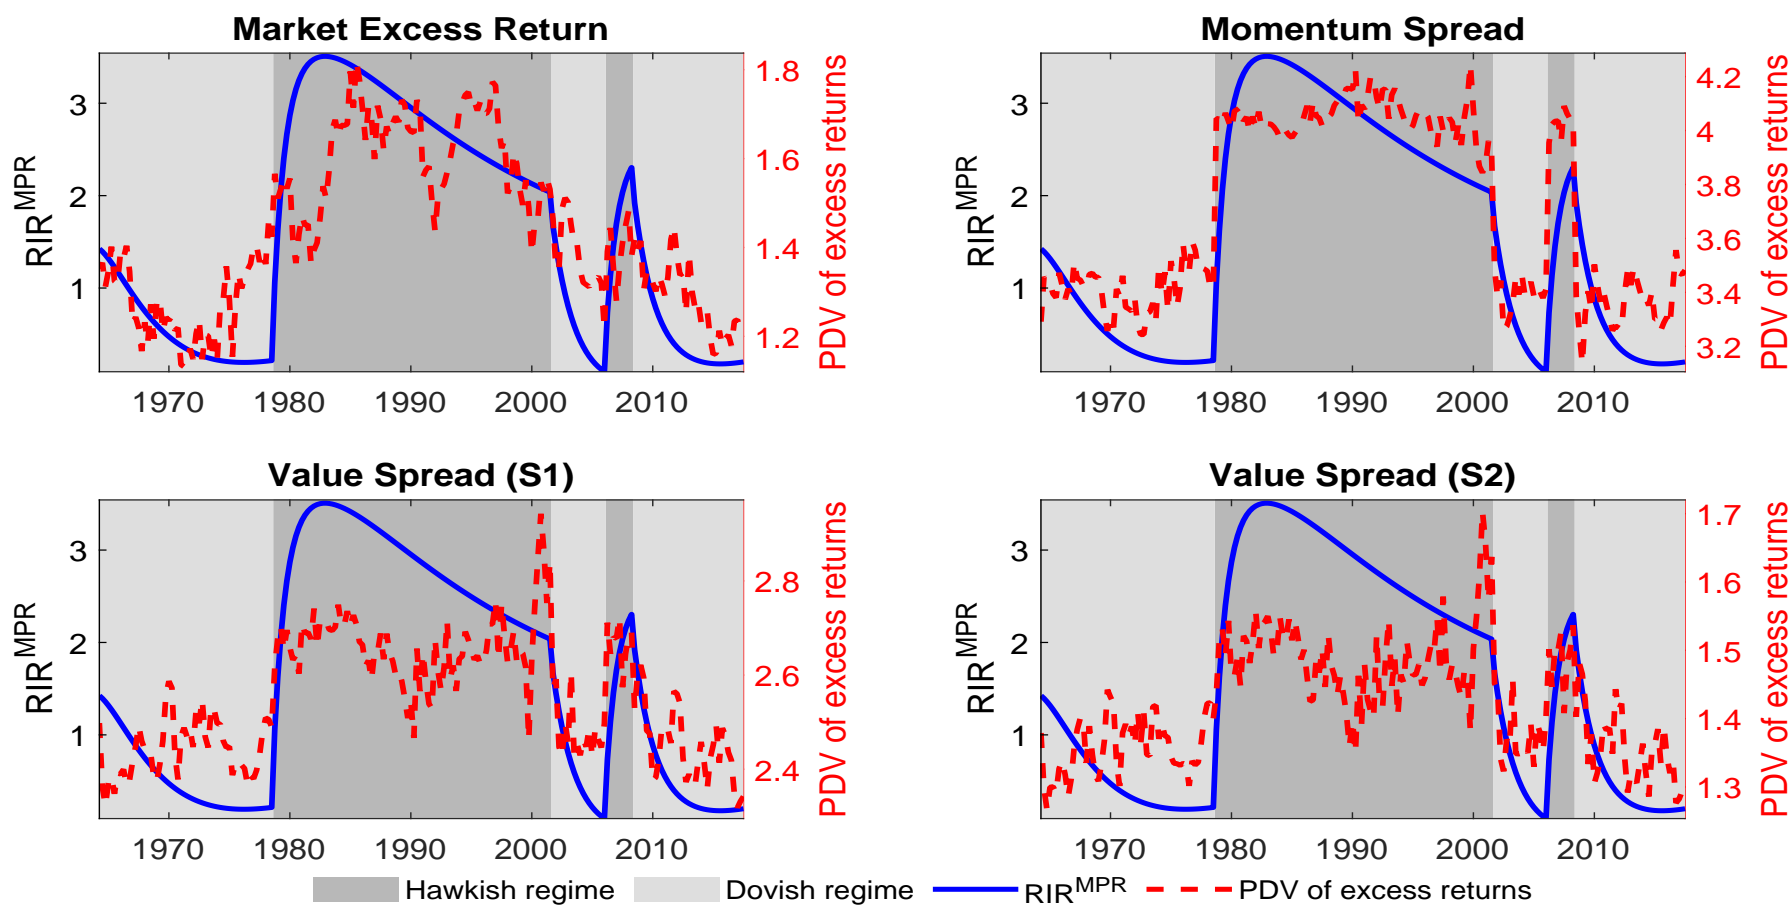


Figure 12: Excess returns and policy rule changes. The figure reports the time series of the PDV of expected excess returns for different portfolios (dashed line, right axis) together with fluctuations of the real interest rate due to changes in the monetary policy rule (solid line, left axis). The sample is quarterly and spans the period 1961:Q1 to 2017:Q3.

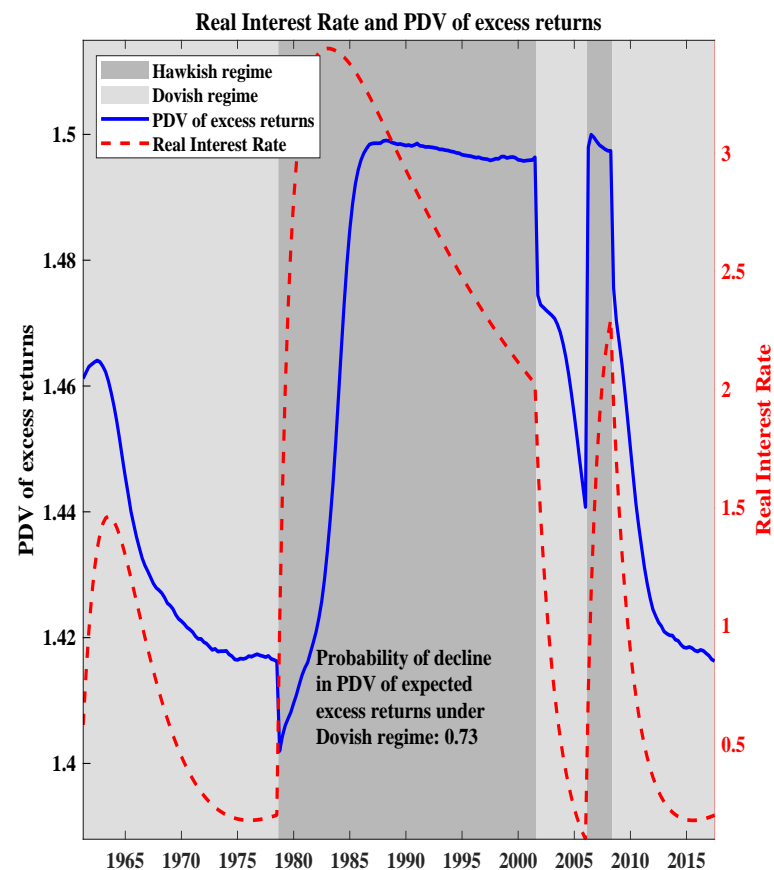



Figure 13: Simulated wealth ratio, real interest, and implied PDV of expected excess returns. This figure plots results from simulating the dynamic macro-finance model at the posterior mode parameter values 20,000 times using a sample length and regime sequence equal to that in our historical data. Using data from each simulated sample, we estimate a 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.

Appendix for Online Publication

Data Appendix

This appendix describes the data used in this study.

CONSUMPTION

Consumption is measured as either total personal consumption expenditure or expenditure on nondurables and services, excluding shoes and clothing. The quarterly data are seasonally adjusted at annual rates, in billions of chain-weighted 2005 dollars. The components are chain-weighted together, and this series is scaled up so that the sample mean matches the sample mean of total personal consumption expenditures. Our source is the U.S. Department of Commerce, Bureau of Economic Analysis.

LABOR INCOME

Labor income is defined as wages and salaries + transfer payments + employer contributions for employee pensions and insurance - employee contributions for social insurance - taxes. Taxes are defined as $[\text{wages and salaries}/(\text{wages and salaries} + \text{proprietors' income with IVA and CCADJ} + \text{rental income} + \text{personal dividends} + \text{personal interest income})]$ times personal current taxes, where IVA is inventory valuation and CCADJ is capital consumption adjustments. The quarterly data are in current dollars. Our source is the Bureau of Economic Analysis.

POPULATION

A measure of population is created by dividing real total disposable income by real per capita disposable income. Our source is the Bureau of Economic Analysis.

WEALTH

Total wealth is household net worth in billions of current dollars, measured at the end of the period. A break-down of net worth into its major components is given in the table below. Stock market wealth includes direct household holdings, mutual fund holdings, holdings of private and public pension plans, personal trusts, and insurance companies. Nonstock wealth includes tangible/real estate wealth, nonstock financial assets (all deposits, open market paper, U.S. Treasuries and Agency securities, municipal securities, corporate and foreign bonds and mortgages), and also includes ownership of privately traded companies in noncorporate equity, and other. Subtracted off are liabilities, including mortgage loans and loans made under home equity lines of credit and secured by junior liens, installment consumer debt and other. Wealth is measured at the end of the period. A timing convention for wealth is needed because the level of consumption is a flow during the quarter rather than a point-in-time estimate as is wealth (consumption data are time-averaged). If we think of a given quarter's consumption data as measuring spending at the beginning of the quarter, then wealth for the quarter should

be measured at the beginning of the period. If we think of the consumption data as measuring spending at the end of the quarter, then wealth for the quarter should be measured at the end of the period. None of our main findings discussed below (estimates of the cointegrating parameters, error-correction specification, or permanent-transitory decomposition) are sensitive to this timing convention. Given our finding that most of the variation in wealth is not associated with consumption, this timing convention is conservative in that the use of end-of-period wealth produces a higher contemporaneous correlation between consumption growth and wealth growth. Our source is the Board of Governors of the Federal Reserve System. A complete description of these data may be found at <http://www.federalreserve.gov/releases/Z1/Current/>.

CRSP PRICE-DIVIDEND RATIO

The stock price is measured using the Center for Research on Securities Pricing (CRSP) value-weighted stock market index covering stocks on the NASDAQ, AMEX, and NYSE. The data are monthly. The stock market price is the price of a portfolio that does not reinvest dividends. The CRSP dataset consists of $vwretx(t) = (P_t/P_{t-1}) - 1$, the return on a portfolio that doesn't pay dividends, and $vwretd_t = (P_t + D_t)/P_t - 1$, the return on a portfolio that does pay dividends. The stock price index we use is the price P_t^x of a portfolio that does not reinvest dividends, which can be computed iteratively as

$$P_{t+1}^x = P_t^x (1 + vwretx_{t+1}),$$

where $P_0^x = 1$. Dividends on this portfolio that does not reinvest are computed as

$$D_t = P_{t-1}^x (vwretd_t - vwretx_t).$$

The above give monthly returns, dividends and prices. The annual log return is the sum of the 12 monthly log returns over the year. We create annual log dividend growth rates by summing the log differences over the 12 months in the year: $d_{t+12} - d_t = d_{t+12} - d_{t+11} + d_{t+11} - d_{t+10} + \dots + d_{t+1} - d_t$. The annual log price-dividend ratio is created by summing dividends in levels over the year to obtain an annual dividend in levels, D_t^A , where t denotes a year here. The annual observation on P_t^x is taken to be the last monthly price observation of the year, P_t^{Ax} . The annual log price-dividend ratio is $\ln(P_t^{Ax}/D_t^A)$. Note that this value for dividend growth is only used to compute the CRSP price-dividend ratio on this hypothetical portfolio. When we investigate the behavior of stock market dividend growth in the MS-VAR, we use actual dividend data from all firms on NYSE, NASDAQ, and AMEX. See the data description for MS-VARs below.

FLOW OF FUNDS EQUITY PAYOUT, DIVIDENDS, PRICE

Flow of Funds payout is measured as “Net dividends plus net repurchases” and is computed using the Flow of Funds Table F.103 (nonfinancial corporate business sector) by subtracting Net Equity Issuance (FA103164103) from Net Dividends (FA106121075). We define net repurchases

to be repurchases net of share issuance, so net repurchases is the negative of net equity issuance. Net dividends consists of payments in cash or other assets, excluding the corporation's own stock, made by corporations located in the United States and abroad to stockholders who are U.S. residents. The payments are netted against dividends received by U.S. corporations, thereby providing a measure of the dividends paid by U.S. corporations to other sectors. The price used for FOF price-dividend and price-payout ratios is "Equity," the flow of funds measure of equities (LM103164103).

PRICE DEFLATOR FOR CONSUMPTION AND ASSET WEALTH

The nominal after-tax labor income and wealth data are deflated by the personal consumption expenditure chain-type deflator (2005=100), seasonally adjusted. In principle, one would like a measure of the price deflator for total flow consumption here. Since this variable is unobservable, we use the total expenditure deflator as a proxy. Our source is the Bureau of Economic Analysis.

DATA FOR MS-VAR TO ESTIMATE RISK PREMIA

The variables included in the MS-VAR for the equity characteristics portfolio data are: (a) the momentum return spread, i.e., the difference between the excess return of the extreme winner (M10) portfolio and the excess return of the extreme loser (M1) portfolio; (b) the value return spread (S1), i.e., the difference between the excess return of the small (size 1) high BM portfolio and the excess return of the small (size 1) low BM portfolio; (c) the value return spread (S2), i.e., the difference between the excess return of the size 2 high BM portfolio and the excess return of the small size 2 low BM portfolio; (d) the momentum BM spread: the difference between the logarithm of the BM ratio of the M10 and M1 portfolios; (e) the value BM spread (S1): The difference between the logarithm of the BM ratio of the small (size quintile 1) high book-market portfolio and the logarithm of the BM ratio of the small (size 1) low book-market portfolio; (f) the value BM spread (S2): the difference between the logarithm of the BM ratio of the size quintile 2 high book-market portfolio and the logarithm of the BM ratio of the size 2 low book-market portfolio; (g) the real FFR (FFR minus inflation); (h) the excess return on the small value portfolio. We then use CRSP/Compustat to construct the BM ratios of the corresponding portfolios.

The MS-VAR specification for the market risk premium includes the following variables: (a) 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; (b) $-cay^{MS}$; (c) the small stock value spread (log-difference in the book to market ratio of the S1 value and S1 growth portfolio); (d) the SMB factor from Fama and French; (e) the HML factor from Fama and French. These variables are obtained from Kenneth French's Dartmouth webpage.

DATA FOR MODEL ESTIMATION

Inflation expectations are taken from the mean inflation forecasts of one year ahead inflation, provided by the University of Michigan Survey of Consumers. Our data sources for output growth are the NIPA tables constructed by the Bureau of Economic Analysis and the St. Louis Fed. Real GDP per capita is obtained by dividing nominal GDP (NIPA 1.1.5, line 1) by the GDP deflator (NIPA 1.1.4, line 1) and population. Population is measured as Civilian Non-institutional Population (CNP16OV) and downloaded from FRED, a website maintained by the Federal Reserve Bank of St. Louis. Inflation is measured as the quarter-to-quarter log-change of CPI. Both the CPI and the FFR series are downloaded from FRED, a website maintained by the Federal Reserve Bank of St. Louis. Expected inflation is the mean of the one-year-ahead expected inflation based on the Michigan survey. All variables are annualized.

Computing cay^{MS}

Let \mathbf{z}_t be a 3×1 vector of data on c_t , a_t , and \tilde{y}_t , and k leads and k lags of Δa_t and Δy_t and let $\mathbf{Z}_t = (\mathbf{z}_t, \mathbf{z}_{t-1}, \dots, \mathbf{z}_1)$ be a vector containing all observations obtained through date t . To estimate the parameters of this stationary linear combination we modify the standard fixed coefficient dynamic least squares regression (DLS–Stock and Watson (1993)) regression to allow for shifts in the intercept α_{ξ_t} :

$$c_t = \alpha_{\xi_t} + \beta_a a_t + \beta_y y_t + \sum_{i=-k}^k b_{a,i} \Delta a_{t+i} + \sum_{i=-k}^k b_{y,i} \Delta y_{t+i} + \sigma^c \varepsilon_t^c \quad (\text{A1})$$

where $\varepsilon_t \sim N(0, 1)$.³³ The parameters of the econometric model include the cointegrating parameters and additional slope coefficients $\beta = (\beta_a, \beta_y, b)'$, where $b = (b_{a,-k}, \dots, b_{a,k}, b_{y,-k}, \dots, b_{y,k})'$, the two intercept values α_1 and α_2 , the standard deviation of the residual σ , and the transition probabilities contained in the matrix \mathbf{H} .

We combine the estimation of changes in the mean of cay_t^{MS} with an isomorphic model for the monetary policy spread. Specifically, we assume that regime changes in the mean of cay_t^{MS} coincide with regime changes in the mean of the mps :

$$mps_t = r_{\xi_t} + \epsilon_t^r, \quad (\text{A2})$$

where $\epsilon_t^r \sim N(0, \sigma_r^2)$. Importantly, unlike cay_t^{MS} , mps_t is an observed variable. Thus, in this case we only need to conduct inference about the MS intercept coefficient r_{ξ_t} . It is worth emphasizing that the same latent state variable, ξ_t , presumed to follow a two-state Markov-switching process with transition matrix \mathbf{H} , controls both changes in α_{ξ_t} and \bar{r}_{ξ_t} .

³³The DLS regression controls for leads and lags of the right-hand-side variables to adjust for the inefficiencies attributable to regressor endogeneity that arise in finite samples.

The model can then be summarized as follows:

$$\begin{aligned}
c_t &= \alpha_{\xi_t} + \beta_a a_t + \beta_y y_t + \sum_{i=-k}^k b_{a,i} \Delta a_{t+i} + \sum_{i=-k}^k b_{y,i} \Delta y_{t+i} + \sigma^c \epsilon_t^c \\
mps_t &= r_{\xi_t} + \sigma^r \epsilon_t^r \\
\epsilon_t^c &\sim N(0, 1), \quad \epsilon_t^r \sim N(0, 1)
\end{aligned}$$

where ξ_t is a hidden variable that follows a Markov-switching process with transition matrix \mathbf{H} . Collect all model parameters into a vector $\boldsymbol{\theta} = (r_{\xi_t}, \sigma^r, \beta, \alpha_{\xi_t}, \sigma^c, \mathbf{H})'$. The model can be thought as a multivariate regression with regime changes in which some of the parameters are restricted to zero.

Our estimate of cay_t^{MS} is based on the posterior mode of the parameter vector $\boldsymbol{\theta}$ and the corresponding regime probabilities. To simplify notation, we denote the vector containing all variables whose coefficients are allowed to vary over time $x_{M,t}$, while $x_{F,t}$ is used to denote the vector containing all the variables whose coefficients are kept constant. We then obtain:

$$\begin{aligned}
c_t &= \alpha_{\xi_t} x_{M,t} + \beta x_{F,t} + \sigma^c \epsilon_t^c \\
mps_t &= r_{\xi_t} x_{M,t} + \sigma^r \epsilon_t^r
\end{aligned}$$

where, in our case, $\beta = [\beta_a, \beta_y, b_{a,-k}, \dots, b_{a,+k}, b_{y,-k}, \dots, b_{y,+k}]$ and the vector $x_{M,t}$ is unidimensional and always equal to 1.

Collect the conditional probabilities $\pi_{t|t}^i = p(\xi_t = i | Y^t; \boldsymbol{\theta})$ for $i = 1, \dots, m$ into an $m \times 1$ vector $\boldsymbol{\pi}_{t|t} = p(\xi_t | Y^t; \boldsymbol{\theta})$. The filtered probabilities reflect the probability of a regime conditional on the data up to time t , $\boldsymbol{\pi}_{t|t} = p(\xi_t | Y^t; \boldsymbol{\theta})$, for $t = 1, \dots, T$, and are part of the output obtained computing the likelihood function associated with the parameter vector $\boldsymbol{\theta} = \{r_{\xi_t}, \sigma^r, \beta, \alpha_{\xi_t}, \sigma, \mathbf{H}\}$. They can be obtained using the following recursive algorithm given by the Hamilton filter:

$$\begin{aligned}
\boldsymbol{\pi}_{t|t} &= \frac{\boldsymbol{\pi}_{t|t-1} \odot \boldsymbol{\eta}_t}{\mathbf{1}' (\boldsymbol{\pi}_{t|t-1} \odot \boldsymbol{\eta}_t)} \\
\boldsymbol{\pi}_{t+1|t} &= \mathbf{H} \boldsymbol{\pi}_{t|t}
\end{aligned} \tag{A3}$$

where $\boldsymbol{\eta}_t$ is a vector whose j -th element contains the conditional density $p(c_t, mps_t | \xi_t = j, x_{M,t}, x_{F,t}; \boldsymbol{\theta})$, i.e.,

$$p(c_t, mps_t | \xi_t = j, x_{M,t}, x_{F,t}; \boldsymbol{\theta}) = \frac{1}{\sqrt{2\pi\sigma^{c,2}}} \frac{1}{\sqrt{2\pi\sigma^{r,2}}} \exp \left[-\frac{[c_t - (\alpha_j x_{M,t} + \beta x_{F,t})]^2}{2\sigma^{c,2}} - \frac{[mps_t - r_j x_{M,t}]^2}{2\sigma^{r,2}} \right], \tag{A4}$$

the symbol \odot denotes element by element multiplication, and $\mathbf{1}$ is a vector with all elements equal to 1. To initialize the recursive calculation we need an assumption on the distribution of ξ_0 . We assume that the two regimes have equal probabilities: $p(\xi_0 = 1) = .5 = p(\xi_0 = 2)$.

The smoothed probabilities reflect all the information that can be extracted from the whole data sample, $\pi_{t|T} = p(\xi_t|Y^T; \boldsymbol{\theta})$. The final term, $\pi_{T|T}$ is returned with the final step of the filtering algorithm. Then, a recursive algorithm can be implemented to derive the other probabilities:

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

where (\div) denotes element by element division.

In using the DLS regression (A1) to estimate cointegrating parameters, we lose 6 leads and 6 lags. For estimates of cay_t^{FC} , we take the estimated coefficients and we apply them to the whole sample. To extend our estimates of cay_t^{MS} over the full sample, we can likewise apply the parameter estimates to the whole sample but we need an estimate of the regime probabilities in the first 6 and last 6 observations of the full sample. For this we run the Hamilton filter from period from -5 to $T + 6$ as follows. When starting at -5 , we assume no lagged values are available and the DLS regression omits all lags, but all leads are included. When at $t = -4$ we assume only one lag is available and the DLS regression includes only one lag, and so on, until we reach $t = 0$ when all lags are included. Proceeding forward when $t = T + 1$ is reached we assume all lags are available and all leads except one are available, when $t = T + 2$, we assume all lags and all leads but two are available, etc. Smoothed probabilities are then computed with standard methods as they only involve the filtered probabilities and the transition matrix \mathbf{H} .

Gibbs Sampling Algorithm

This appendix describes the Bayesian methods used to characterize uncertainty in the regression parameters. To simplify notation, we denote the vector containing all variables whose coefficients are allowed to vary over time $x_{M,t}$, while $x_{F,t}$ is used to denote the vector containing all the variables whose coefficients are kept constant. We then obtain:

$$c_t = \alpha_{\xi_t} x_{M,t} + \beta x_{F,t} + \sigma^c \varepsilon_t^c \quad (\text{A5})$$

$$mps_t = r_{\xi_t} x_{M,t} + \sigma^r \varepsilon_t^r \quad (\text{A6})$$

where, in our case, $\beta = [\beta_a, \beta_y, b_{a,-k}, \dots, b_{a,+k}, b_{y,-k}, \dots, b_{y,+k}]$ and the vector $x_{M,t}$ is unidimensional and always equal to 1.

Suppose the Gibbs sampling algorithm has reached the n -th iteration. We then have draws for $r_{\xi_t,n}, \sigma_n^r, \beta_n, \alpha_{\xi_t,n}, \sigma_n^c, \mathbf{H}_n$, and ξ_n^T , where $\xi_n^T = \{\xi_{1,n}, \xi_{2,n}, \dots, \xi_{T,n}\}$ denotes a draw for the whole regime sequence. The parameters for equations (A5) and (A6) can be drawn separately, while the regime sequence ξ_n^T requires a joint evaluation of the Hamilton filter. Finally, the transition matrix \mathbf{H}_n is drawn conditionally on the regime sequence.

Specifically, the sampling algorithm is described as follows.

1. **Sampling** β_{n+1} : Given $\alpha_{\xi_t, n}$, σ_n^c , and ξ_n^T we transform the data:

$$\tilde{c}_t = \frac{c_t - \alpha_{\xi_t, n} x_{M,t}}{\sigma_n^c} = \beta \frac{x_{F,t}}{\sigma_n^c} + \varepsilon_t = \beta \tilde{x}_t + \varepsilon_t.$$

The above is a regression with fixed coefficients β and standardized residual shocks. Standard Bayesian methods may be used to draw the coefficients of the regression. We assume a Normal conjugate prior $\beta \sim N(B_{\beta,0}, V_{\beta,0})$, so that the conditional (on $\alpha_{\xi_t, n}$, σ_n^c , and ξ_n^T) posterior distribution is given by

$$\beta_{n+1} \sim N(B_{\beta,T}, V_{\beta,T})$$

with $V_{\beta,T} = \left(V_{\beta,0}^{-1} + \tilde{X}'_F \tilde{X}_F \right)^{-1}$ and $B_{\beta,T} = V_{\beta,T} \left[V_{\beta,0}^{-1} B_{\beta,0} + \tilde{X}'_F \tilde{C} \right]$, where $\tilde{C} = (\tilde{c}_1, \dots, \tilde{c}_T)'$ and $\tilde{X}_F = (x_{F,1}, \dots, x_{F,T})'$ and $B_{\beta,0}$ and $V_{\beta,0}^{-1}$ control the priors for the fixed coefficients of the regression. Keeping in mind that the residuals have been normalized to have unit variance, with flat priors, $B_{\beta,0} = 0$ and $V_{\beta,0}^{-1} = 0$ and $B_{\beta,T}$ and $V_{\beta,T}$ coincide with the maximum likelihood estimates, conditional on the other parameters.

2. **Sampling** $\alpha_{i,n+1}$ for $i = 1, 2$: Given β_{n+1} , σ_n^c , and ξ_n^T we transform the data:

$$\tilde{c}_t = \frac{c_t - \beta_{n+1} x_{F,t}}{\sigma_n^c} = \alpha_{\xi_t} \frac{x_{M,t}}{\sigma_n^c} + \varepsilon_t = \alpha_{\xi_t} \tilde{x}_{M,t} + \varepsilon_t.$$

The above regression has standardized shocks and Markov-switching coefficients in the transformed data. Using ξ_n^T we can group all the observations that pertain to the same regime i . Given the prior $\alpha_i \sim N(B_{\alpha_i,0}, V_{\alpha_i,0})$ for $i = 1, 2$ we use standard Bayesian methods to draw α_i from the conditional (on β_{n+1} , σ_n^c , and ξ_n^T) posterior distribution:

$$\alpha_{i,n+1} \sim N(B_{\alpha_i,T}, V_{\alpha_i,T}) \text{ for } i = 1, 2$$

where $V_{\alpha_i,T} = \left(V_{\alpha_i,0}^{-1} + \tilde{X}'_{M,i} \tilde{X}_{M,i} \right)^{-1}$ and $B_{\alpha_i,T} = V_{\alpha_i,T} \left[V_{\alpha_i,0}^{-1} B_{\alpha_i,0} + \tilde{X}'_{M,i} \tilde{C}_i \right]$ where \tilde{C}_i and $\tilde{X}_{M,i}$ collect all the observations for the transformed data for which regime i is in place. The parameters $B_{\alpha_i,0}$ and $V_{\alpha_i,0}^{-1}$ control the priors for the MS coefficients of the regression: $\alpha_i \sim N(B_{\alpha_i,0}, V_{\alpha_i,0})$ for $i = 1, 2$. With flat priors, we have $B_{\alpha_i,0} = 0$ and $V_{\alpha_i,0}^{-1} = 0$ and $B_{\alpha_i,T}$ and $V_{\alpha_i,T}$ coincide with the maximum likelihood estimates, conditional on the other parameters.

3. **Sampling** $r_{i,n+1}$ for $i = 1, 2$: Given σ_n^r and ξ_n^T we transform the data:

$$\widetilde{mps}_t = \frac{mps_t}{\sigma_n^r} = r_{\xi_t} \frac{x_{M,t}}{\sigma_n^r} + \varepsilon_t = \alpha_{\xi_t} \tilde{x}_{M,t} + \varepsilon_t^r.$$

The above regression has standardized shocks and Markov-switching coefficients in the transformed data. Using ξ_n^T we can group all the observations that pertain to the same

regime i . Given the prior $r_i \sim N(B_{r_i,0}, V_{r_i,0})$ for $i = 1, 2$ we use standard Bayesian methods to draw r_i from the conditional $(\sigma_i^r$ and $\xi_n^T)$ posterior distribution:

$$r_{i,n+1} \sim N(B_{r_i,T}, V_{r_i,T}) \text{ for } i = 1, 2$$

where $V_{r_i,T} = \left(V_{r_i,0}^{-1} + \tilde{X}'_{M,i} \tilde{X}_{M,i} \right)^{-1}$ and $B_{r_i,T} = V_{r_i,T} \left[V_{r_i,0}^{-1} B_{r_i,0} + \tilde{X}'_{r,i} \tilde{R}_i \right]$ where \tilde{R}_i and $\tilde{X}_{r,i}$ collect all the observations for the transformed data for which regime i is in place. The parameters $B_{r_i,0}$ and $V_{r_i,0}^{-1}$ control the priors for the MS coefficients of the regression: $r_i \sim N(B_{r_i,0}, V_{r_i,0})$ for $i = 1, 2$. With flat priors, we have $B_{r_i,0} = 0$ and $V_{r_i,0}^{-1} = 0$ and $B_{r_i,T}$ and $V_{r_i,T}$ coincide with the maximum likelihood estimates, conditional on the other parameters.

4. **Sampling** σ_{n+1}^c : Given β_{n+1} , $\alpha_{\xi_t, n+1}$, and ξ_n^T we can compute the residuals of the regression:

$$\tilde{c}_t = c_t - \beta_{n+1} x_{F,t} - \alpha_{\xi_t} x_{M,t} = \sigma^c \varepsilon_t.$$

With the prior that σ^c has an inverse gamma distribution, $\sigma^c \sim IG(Q_0, v_0)$, we use Bayesian methods to draw σ_{n+1}^c from the conditional (on β_{n+1} , $\alpha_{\xi_t, n+1}$, and ξ_n^T) posterior inverse gamma distribution:

$$\sigma_{n+1}^c \sim IG(Q_T^c, v_T), \quad v_T = T + v_0, \quad Q_T^c = Q_0 + E^c E^c$$

where E^c is a vector containing the residuals, T is the sample size, and Q_0 and v_0 control the priors for the standard deviation of the innovations: $\sigma^c \sim IG(Q_0, v_0)$. The mean of a random variable with distribution $\sigma^c \sim IG(Q_T^c, v_T^c)$ is Q_T^c / v_T^c . With flat priors we have $Q_0 = 0$ and $v_0 = 0$, and the mean of σ^c is therefore $(E^c E^c) / T$, which coincides with the standard maximum likelihood (MLE) estimate of σ^c , conditional on the other parameters.

5. **Sampling** σ_{n+1}^r : Given $r_{\xi_t, n+1}$ and ξ_n^T we can compute the residuals of the regression:

$$\widetilde{mps}_t = mps_t - r_{\xi_t} x_{M,t} = \sigma^r \varepsilon_t^r.$$

With the prior that σ^r has an inverse gamma distribution, $\sigma^r \sim IG(Q_0, v_0)$, we use Bayesian methods to draw σ_{n+1}^r from the conditional (on $r_{\xi_t, n+1}$ and ξ_n^T) posterior inverse gamma distribution:

$$\sigma_{n+1}^r \sim IG(Q_T^r, v_T), \quad v_T = T + v_0, \quad Q_T^r = Q_0 + E^r E^r$$

where E is a vector containing the residuals, T is the sample size, and Q_0 and v_0 control the priors for the standard deviation of the innovations: $\sigma^r \sim IG(Q_0, v_0)$. The mean of a random variable with distribution $\sigma^r \sim IG(Q_T^r, v_T^r)$ is Q_T^r / v_T^r . With flat priors we have $Q_0 = 0$ and $v_0 = 0$, and the mean of σ^r is therefore $(E^r E^r) / T$, which coincides with the maximum likelihood (MLE) estimate of σ^r , conditional on the other parameters.

6. **Sampling ξ_{n+1}^T** : Given $r_{\xi_t, n+1}, \sigma_{n+1}^r, \beta_{n+1}, \alpha_{\xi_t, n+1}, \sigma_{n+1}^c$, and \mathbf{H}_n , we can treat equations (A5) and (A6) as a multivariate regression in which some parameters are restricted to zero. This allows to obtain filtered probabilities for the regimes using the filter described in Hamilton (1994). Following Kim and Nelson (1999) we then use a Multi-Move Gibbs sampling to draw a regime sequence ξ_{n+1}^T .
7. **Sampling \mathbf{H}_{n+1}** : Given the draws for the MS state variables ξ_{n+1}^T , the posterior for the transition probabilities does not depend on other parameters of the model and follows a Dirichlet distribution if we assume a prior Dirichlet distribution.³⁴ For each column of \mathbf{H}_{n+1} the posterior distribution is given by:

$$\mathbf{H}_{n+1}(:, i) \sim D(a_{ii} + \eta_{ii, n+1}, a_{ij} + \eta_{ij, n+1})$$

where $\eta_{ij, n+1}$ denotes the number of transitions from state i to state j based on ξ_{n+1}^T , while a_{ii} and a_{ij} the corresponding priors. With flat priors, we have $a_{ii} = 0$ and $a_{ij} = 0$.

8. If $n + 1 < N$, where N is the desired number of draws, go to step 1, otherwise stop.

These steps are repeated until convergence to the posterior distribution is reached. We check convergence by using the Raftery-Lewis Diagnostics for each parameter in the chain. See section below. We use the draws obtained with the Gibbs sampling algorithm to characterize parameter uncertainty in Table 2. The Gibbs sampling algorithm is used to generate a distribution for the difference between the two means in the same manner it is used to generate a distribution for any parameter. For each draw from the joint distribution of the model parameters, we compute the difference and store it. We may then compute means and/or medians, and error bands, as for any other parameter of interest.

Convergence Checks

The 90% credible sets are obtained making 2,000,000 draws from the posterior using the Gibbs sampling algorithm. One in every one thousand draws is retained. We check convergence using the methods suggested by Raftery and Lewis (1992) and Geweke (1992). For Raftery and Lewis (1992) checks, we target 90% credible sets, with a 1% accuracy to be achieved with a 95% minimum probability. We initialize the Gibbs sampling algorithm making a draw around the posterior mode. Sims and Zha (2006) point out that in Markov-switching models it is important to first find the posterior mode and then use it as a starting point for the MCMC algorithm due to the fact that the likelihood can have multiple peaks. The tables below pertain to convergence of the Gibbs sampling algorithm.

³⁴The Dirichlet distribution is a generalization of the beta distribution that allows one to potentially consider more than 2 regimes. See e.g., Sims and Zha (2006).

Most Likely Regime Sequence

In this appendix we explain how to compute the most likely regime sequence. This most likely regime sequence is based on our estimates for the breaks in cay^{MS} and mps , and is taken as given in the portfolio MS-VAR and the model estimation. Specifically, we choose the particular regime sequence $\xi_n^T = \{\hat{\xi}_{1,n}, \dots, \hat{\xi}_{T,n}\}$ that is most likely to have occurred, given our estimated posterior mode parameter values for θ . This sequence is computed as follows. Let $P(\xi_t = i | \mathbf{Z}_{t-1}; \theta) \equiv \pi_{t|t-1}^i$. First, we run Hamilton's filter to get the vector of filtered probabilities $\pi_{t|t}$, $t = 1, 2, \dots, T$. The Hamilton filter can be expressed iteratively as

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

where η_t is a vector whose j -th element contains the conditional density $p(c_t | \xi_t = j, x_{M,t}, x_{F,t}; \theta)$, the symbol \odot denotes element by element multiplication, and $\mathbf{1}$ is a vector with all elements equal to 1. The final term, $\pi_{T|T}$ is returned with the final step of the filtering algorithm. Then, a recursive algorithm can be implemented to derive the other smoothed probabilities:

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

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

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

We first take $\pi_{T|T}$ from the Hamilton filter and choose the regime that is associated with the largest probability, i.e., if $\pi_{T|T} = (.9, .1)$, where the first element corresponds to the probability of regime 1, we select $\hat{\xi}_T = 1$, indicating that we are in regime 1 in period T . We now update $\pi_{T|T} = (1, 0)$ and plug into the right-hand-side above along with the estimated filtered probabilities for $\pi_{T-1|T-1}$, $\pi_{T|T-1}$ and estimated transition matrix \mathbf{H} to get $\pi_{T-1|T}$ on the left-hand-side. Now we repeat the same procedure by choosing the regime for $T - 1$ that has the largest probability at $T - 1$, e.g., if $\pi_{T-1|T} = (.2, .8)$ we select $\hat{\xi}_{T-1} = 2$, indicating that we are in regime 2 in period $T - 1$, we then update to $\pi_{T-1|T} = (0, 1)$, which is used again on the right-hand-side now

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

We proceed in this manner until we have a most likely regime sequence ξ^T for the entire sample $t = 1, 2, \dots, T$. Two aspects of this procedure are worth noting. First, it fails if the updated

probabilities are exactly (.5, .5). Mathematically this is virtually zero. Second, note that this procedure allows us to choose the most likely regime sequence by using the recursive formula above to update the filtered probabilities sequentially from T to time $t = 1$. This allows us to take into account the time dependence in the regime sequence as dictated by the transition probabilities.

Book-to-Market Ratio and PDVs

We use the methods and assumptions of the previous subsection to obtain the present value decomposition of the book to market ratio. Consider an MS-VAR:

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

where Z_t is a column vector containing n variables observable at time t and $\xi_t = 1, \dots, m$, with m the number of regimes, evolves following the transition matrix \mathbf{H} . If the MS-VAR has more than one lag, the companion form can be used to recast the model as illustrated above.

The subappendix “conditional expectations and volatility” below shows how to compute $\mathbb{E}_t(Z_{t+s}) = wq_{t+s|t}$, where

$$\begin{aligned} q_{t+s|t}^i &\equiv \mathbb{E}_t(Z_{t+s} 1_{\xi_t=i}) = \mathbb{E}(Z_{t+s} 1_{\xi_t=i} | \mathbb{I}_t) \\ 1'_x &= [0, \dots, 1, \dots, 0, 0, 0]', \quad mn = m * n \end{aligned}$$

and where \mathbb{I}_t contains all the information that agents have at time t , including knowledge of the regime in place, for the case where there are m regimes.

Now consider the formula from Vuolteenaho (1999):

$$\theta_t = \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t r_{t+1+j} + \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t f_{t+1+j} - \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t e_{t+1+j}^*$$

Given that our goal is to assess if assets with different risk profiles are affected differently by the breaks in the long-term interest rates, we focus on the *difference* between the book-to-market ratios. Specifically, given two portfolios x and y , we are interested in how the difference in their book-to-market ratios, $\theta_{x,t} - \theta_{y,t}$, varies across the two regimes:

$$\underbrace{\theta_{x,t} - \theta_{y,t}}_{\text{Spread in BM ratios}} = \underbrace{\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (r_{x,t+1+j} - r_{y,t+1+j})}_{\text{PDV of the difference in expected excess returns}} - \underbrace{\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (e_{x,t+1+j}^* - e_{y,t+1+j}^*)}_{\text{PDV of the difference in expected earnings}}$$

If then we want to correct the spread in BM ratios by taking into account expected earnings, we have:

$$\underbrace{\theta_{x,t} - \theta_{y,t} + \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (e_{x,t+1+j}^* - e_{y,t+1+j}^*)}_{\text{Spread in BM ratios corrected for earnings}} = \underbrace{\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (r_{x,t+1+j} - r_{y,t+1+j})}_{\text{PDV of the expected spread in excess returns}} \quad (\text{A7})$$

The spread in excess returns, $r_{xy,t} \equiv r_{x,t} - r_{y,t}$. Then the right hand side of (A7) can be computed as:

$$\begin{aligned} \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (r_{xy,t+1+j}) &= \sum_{j=0}^{\infty} \rho^j 1'_{r_{xy}} w q_{t+1+j|t} \\ &= 1'_{r_{xy}} w (I - \rho\Omega)^{-1} [\Omega q_{t|t} + C (I - \rho\mathbf{H})^{-1} \mathbf{H}\pi_{t|t}]. \end{aligned}$$

Therefore, we have:

$$\tilde{\theta}_{xy,t} \equiv \underbrace{\tilde{\theta}_{x,t} - \tilde{\theta}_{y,t} + \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (e_{x,t+1+j}^* - e_{y,t+1+j}^*)}_{\text{Spread in BM ratios corrected for earnings}} = 1'_{r_{xy}} w (I - \rho\Omega)^{-1} [\Omega q_{t|t} + C (I - \rho\mathbf{H})^{-1} \mathbf{H}\pi_{t|t}] \quad (\text{A8})$$

where we have used $\tilde{\theta}_{xy,t}$ to define the spread in BM ratios corrected for earnings.

Similar formulas are used to compute risk premia for the individual portfolios. The premium for a portfolio z coincides with the present discounted value of its excess returns:

$$\underbrace{\text{premia}_{z,t}}_{\text{Premia}} \equiv \underbrace{\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (r_{z,t+1+j})}_{\text{PDV of excess returns}} = 1'_{r_z} w (I - \rho\Omega)^{-1} [\Omega q_{t|t} + C (I - \rho\mathbf{H})^{-1} \mathbf{H}\pi_{t|t}], \quad (\text{A9})$$

where $1'_{r_z}$ is a vector used to extract the PDV of excess returns from a vector containing the PDV of all variables included in the VAR. In our VAR application, we compute $\pi_{t|t}$ to correspond to the most likely regime sequence, as defined in the subsection below. This implies that the vector $\pi_{t|t}$ assumes only one of two values, $(1, 0)'$ or $(0, 1)'$.

Regime Average We also compute the *regime average* value of $\tilde{\theta}_{xy,t}$. The regime average is defined as:

$$\tilde{\theta}_{xy}^i \equiv 1'_{r_{xy}} w (I - \rho\Omega)^{-1} [\Omega \bar{q}_i + C (I - \rho\mathbf{H})^{-1} \mathbf{H}\bar{\pi}_i]$$

where $\bar{\pi}_i = 1_i$ and $\bar{q}_i \equiv [0, \dots, \bar{\mu}_i, \dots, 0]$ is a column vector that contains the conditional steady state of for the mean value of Z_t conditional on being in regime i , i.e., $\mathbb{E}_i(Z_t) = \bar{\mu}_i = (I_n - A_i)^{-1} c_i$, and zero otherwise. Recall that the conditional steady state, $\bar{\mu}_i$, is a vector that contains the expected value of Z_t conditional on being in regime i . Therefore, the vector captures the values to which the variables of the VAR converge if regime i is in place forever. Although none of our regimes are estimated to be absorbing states, this is still a good approximation for regimes that can be expected to persist for prolonged periods of time. Note that $\tilde{\theta}_{xy}^i$ is computed by conditioning on the economy being initially at $Z_t = \bar{\mu}_i$ and in regime i , but taking into account that there might be regime changes in the future. Therefore, we can also think about $\tilde{\theta}_{xy}^i$ as the expected value of $\tilde{\theta}_{xy,t}$, conditional on being in regime i today and on the variables of the VAR being equal to the conditional steady state mean values for regime i . Formally:

$$\tilde{\theta}_{xy}^i = \mathbb{E} \left(\tilde{\theta}_{xy,t} | \xi_t = i, Z_t = \bar{\mu}_i \right). \quad (\text{A10})$$

Similarly, we can compute the *regime average* value of risk premia for an individual portfolio z , $premia_{z,t}$:

$$\overline{premia}_z^i \equiv 1'_{r_z} w (I - \rho\Omega)^{-1} [\Omega\bar{q}_i + C(I - \rho\mathbf{H})^{-1} \mathbf{H}\bar{\pi}_i]. \quad (\text{A11})$$

Formulas (A8), (A9), (A10), and (A11) are used in the paper to produce Figure 4 and Table 5. For each draw of the VAR parameters from the posterior distribution, we can compute the evolution of $\tilde{\theta}_{xy,t}$ and individual portfolio $premia_{z,t}$, by using (A8) and (A9). Thus, we obtain a posterior distribution for $\tilde{\theta}_{xy,t}$ and $premia_{z,t}$. The medians of these posterior distributions are reported as the blue solid lines in Figure 4. Similarly, for each draw of the VAR coefficients, we compute $\tilde{\theta}_{xy}^i$ and the difference $\tilde{\theta}_{xy}^1 - \tilde{\theta}_{xy}^2$. Thus, we obtain a posterior distribution for $\tilde{\theta}_{xy}^i$ and for the difference $\tilde{\theta}_{xy}^1 - \tilde{\theta}_{xy}^2$. The medians of the distribution of $\tilde{\theta}_{xy}^i$ and \overline{premia}_z^i for $i = 1, 2$, are reported in Figure 4 (red dashed line). Table 5 reports the median and the 68% posterior credible sets both for the distribution of $\tilde{\theta}_{xy}^i$, for $i = 1, 2$, and for the difference in these across regimes, $\tilde{\theta}_{xy}^1 - \tilde{\theta}_{xy}^2$. Finally, the last row of Table 5 reports the percentage of draws for which $\tilde{\theta}_{xy}^1 - \tilde{\theta}_{xy}^2 > 0$ and $\overline{premia}_z^1 - \overline{premia}_z^2 > 0$ as the probability that risk premia are lower in the high asset valuation/low interest rate regime than they are in the low asset valuation/high interest rate regime.

Variable Selection for VARs to Compute PDV of Risk Premia

We start with a series of fixed regressors that are relevant for predicting market excess returns or the return of the spread portfolios. To limit the size of the MS-VAR, we then use the Akaike information criterion (AIC) to decide whether to include some additional regressors. Specifically, we compute the AIC for the equation(s) that correspond(s) to the return(s) that we are trying to predict. We then choose the specification that minimizes the AIC.

Here are the details:

1. MS-VAR for the Market excess return:

Fixed regressors (all lagged): Market excess return, inverse valuation ratio based on cay^{MS} . The inverse valuation ratio is included because it represents a measure of asset valuation that can predict future stock market returns. Note that given that we are conditioning to the regime sequence obtained when estimating cay^{MS} , the intercept for the corresponding equation will adjust in a way to reflect the low frequency breaks identified above.

Possible additional variables to be chosen for the estimation based on the AIC: Value (small) spread (log-difference in the book to market ratio of the small value portfolios and the book to market ratio of the small growth portfolios), Real FFR, term yield spread, four of the five Fama and French factors (SMB, HML, RMW, CMA), cay (based

on PCE, available on Martin Lettau's website.) Note that we do not include the market excess return from Fama and French (MKTMINRF) as a possible additional regressor because our dependant variable is the excess market return itself. Therefore, this variable is automatically included in the MS-VAR.

Additional regressors selected based on the AIC: Value Spread, and SMB and HML factors from Fama and French.

2. MS-VAR for (a) Momentum return spread: The difference between the excess return of the extreme winner (M10) portfolio and the excess return of the extreme loser (M1) portfolio; (b) Value return spread (S1): The difference between the excess return of the small (size 1) high BM portfolio and the excess return of the small (size 1) low BM portfolio; (c) Value return spread (S2): The difference between the excess return of the size 2 high BM portfolio and the excess return of the small size 2 low BM portfolio.

Fixed regressors (all lagged): (a) Momentum return spread; (b) Value return spread (S1); (c) Value return spread (S2); (d) Momentum BM spread: The difference between the logarithm of the BM ratio of the extreme winner (M10) portfolio and the logarithm of the BM ratio of the extreme loser (M1) portfolio; (e) Value BM spread (S1): The difference between the logarithm of the BM ratio of the small (size quintile 1) high book-market portfolio and the logarithm of the BM ratio of the small (size 1) low book-market portfolio; (f) Value BM spread (S2): The difference between the logarithm of the BM ratio of the size quintile 2 high book-market portfolio and the logarithm of the BM ratio of the size 2 low book-market portfolio.

Possible additional variables to be chosen for the estimation based on the AIC: Real FFR computed as the difference between FFR and Inflation, excess return of small growth portfolio, excess return of small value portfolio, five Fama-French factors (SMB, HML, RMW, CMA, MKTMINRF.)

Additional regressors selected based on the AIC: Real FFR and excess return of the small value portfolio.

Estimation of the MS-VAR

In this appendix we provide details on the estimation of the MS-VAR. Given that we take the regime sequence as given, we need only estimate the transition matrix and the parameters of the MS-VAR across the two regimes. The model is estimated by using Bayesian methods with flat priors on all parameters. As a first step, we group all the observations that belong to the same regime. Conditional on a regime, we have a fixed coefficients VAR. We can then follow standard procedures to make draws for the VAR parameters as follows.

Rewrite the VAR as

$$\begin{aligned} Y_{T \times n} &= X A_{\xi_t} + \varepsilon_{T \times n}, \quad \xi_t = 1, 2 \\ \varepsilon_t &\sim N(0, \Sigma_{\xi_t}) \end{aligned}$$

where $Y = [Z_1, \dots, Z_T]'$, the t -th row of X is $X_t = [1, Z'_{t-1}, Z'_{t-2}]$, $A_{\xi_t} = [c_{\xi_t}, A_{1,\xi_t}, A_{2,\xi_t}]'$, the t -th row of ε is ε_t , and where $\Sigma_{\xi_t} = V_{\xi_t} V_{\xi_t}'$. If we specify a Normal-Wishart prior for A_{ξ_t} and V_{ξ_t} :

$$\begin{aligned} \Sigma_{\xi_t}^{-1} &\sim W(S_0^{-1}/v_0, v_0) \\ \text{vec}(A_{\xi_t} | \Sigma_{\xi_t}) &\sim N(\text{vec}(B_0), \Sigma_{\xi_t} \otimes N_0^{-1}) \end{aligned}$$

where $E(\Sigma_{\xi_t}^{-1}) = S_0^{-1}$, the posterior distribution is still in the Normal-Wishart family and is given by

$$\begin{aligned} \Sigma_{\xi_t}^{-1} &\sim W(S_T^{-1}/v_T, v_T) \\ \text{vec}(A_{\xi_t} | \Sigma_{\xi_t}) &\sim N(\text{vec}(B_T), \Sigma_{\xi_t} \otimes N_T^{-1}) \end{aligned}$$

Using the estimated regime sequence ξ_n^T we can group all the observations that pertain to the same regime i . Therefore the parameters of the posterior are computed as

$$\begin{aligned} v_T &= T_i + v_0, \quad N_T = X_i' X_i + N_0 \\ B_T &= N_T^{-1} (N_0 B_0 + X_i' X_i \hat{B}_{MLE}) \\ S_T &= \frac{v_0}{v_T} S_0 + \frac{T_i}{v_T} \hat{\Sigma}_{MLE} + \frac{1}{v_T} (\hat{B}_{MLE} - \hat{B}_0)' N_0 N_T^{-1} X_i' X_i (\hat{B}_{MLE} - \hat{B}_0) \\ \hat{B}_{MLE} &= (X_i' X_i)^{-1} (X_i' Y_i), \quad \hat{\Sigma}_{MLE} = \frac{1}{T_i} (Y_i - X_i \hat{B}_{MLE})' (Y_i - X_i \hat{B}_{MLE}), \end{aligned}$$

where T_i, Y_i, X_i denote the number and sample of observations in regime i . We choose flat priors ($v_0 = 0, N_0 = 0$) so the expressions above coincide with the MLE estimates using observations in regime i :

$$v_T = T_i, \quad N_T = X_i' X_i, \quad B_T = \hat{B}_{MLE}, \quad S_T = \hat{\Sigma}_{MLE}.$$

Armed with these parameters in each regime, we can make draws from the posterior distributions for $\Sigma_{\xi_t}^{-1}$ and A_{ξ_t} in regime i to characterize parameter uncertainty about these parameters.

Given that we condition the MS-VAR estimates on the most likely regime sequence, ξ_n^T , for cay^{MS} , it is still of interest to estimate the elements of the transition probability matrix for the MS-VAR parameters, \mathbf{H}^A , conditional on this regime sequence. Because we impose this regime sequence, the posterior of \mathbf{H}^A only depends on ξ_n^T and does not depend on other parameters of the model. The posterior has a Dirichlet distribution if we assume a prior Dirichlet distribution.³⁵ For each column of \mathbf{H}^A the posterior distribution is given by:

$$\mathbf{H}^A(:, i) \sim D(a_{ii} + \eta_{ii,r+1}, a_{ij} + \eta_{ij,r+1})$$

where $\eta_{ij,r+1}$ denotes the number of transitions from regime i to regime j based on ξ_n^T , while a_{ii} and a_{ij} the corresponding priors. With flat priors, we have $a_{ii} = 0$ and $a_{ij} = 0$. Armed with this posterior distribution, we can characterize uncertainty about \mathbf{H}^A . Note that the posterior \mathbf{H}^A will be in general different from the posterior distribution of \mathbf{H} because the former is based on a particular regime sequence ξ_n^T , while the latter reflects the entire posterior distribution for ξ_n^T . The estimated transition matrix \mathbf{H}^A can in turn be used to compute expectations taking into account the possibility of regime change (see the next subsection).

Conditional Expectations and Volatility

In this appendix we explain how expectations and economic uncertainty are computed for variables in the MS-VAR. More details can be found in Bianchi (2016). Consider the following first-order MS-VAR:

$$Z_t = c_{\xi_t} + A_{\xi_t} Z_{t-1} + V_{\xi_t} \varepsilon_t, \varepsilon_t \sim N(0, I) \tag{A12}$$

and suppose that we are interested in $\mathbb{E}_0(Z_t) = \mathbb{E}(Z_t | \mathbb{I}_0)$ with \mathbb{I}_0 being the information set available at time 0. The first-order VAR is not restrictive because any VAR with $l > 1$ lags can be rewritten as above by using the first-order companion form, and the methods below applied to the companion form.

Let n be the number of variables in the VAR of the previous Appendix section. Let m be the number of Markov-switching states. Define the $mn \times 1$ column vector q_t as:

$$q_t = [q_t^1, \dots, q_t^m]'$$

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

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

³⁵The Dirichlet distribution is a generalization of the beta distribution that allows one to potentially consider more than 2 regimes. See e.g., Sims and Zha (2006).

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

$$\mu_t = \mathbb{E}_0(Z_t) = \sum_{i=1}^m q_t^i = wq_t$$

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

Proposition 1 *Consider a Markov-switching model whose law of motion can be described by (A12) and define $q_t^i = \mathbb{E}_0(Z_t 1_{\xi_t=i})$ for $i = 1 \dots m$. Then $q_t^j = c_j \pi_t^j + \sum_{i=1}^m A_j q_{t-1}^i h_{ji}$.*

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

$$q_{t+s|t} = [q_{t+s|t}^{1'}, \dots, q_{t+s|t}^{m'}]'$$

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

$$\mu_{t+s|t} = \mathbb{E}_t(Z_{t+s}) = wq_{t+s|t}, \quad (\text{A13})$$

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

$$q_{t+s|t} = C\pi_{t+s|t} + \Omega q_{t+s-1|t} \quad (\text{A14})$$

$$\pi_{t+s|t} = \mathbf{H}\pi_{t+s-1|t} \quad (\text{A15})$$

with $\pi_{t+s|t} = [\pi_{t+s|t}^{1'}, \dots, \pi_{t+s|t}^{m'}]'$, $\Omega = bdiag(A_1, \dots, A_m)(\mathbf{H} \otimes I_n)$, and $C = bdiag(c_1, \dots, c_m)_{mn \times m}$, where e.g., c_1 is the $n \times 1$ vector of constants in regime 1, \otimes represents the Kronecker product and $bdiag$ is a matrix operator that takes a sequence of matrices and use them to construct a block diagonal matrix.

Similar formulas hold for the second moments. Before proceeding, let us define the vectorization operator $\varphi(X)$ that takes the matrix X as an input and returns a column vector stacking the columns of the matrix X on top of one another. We will also make use of the following result: $\varphi(X_1 X_2 X_3) = (X_3' \otimes X_1) \varphi(X_2)$.

Define the vector $n^2 m \times 1$ column vector Q_t as:

$$Q_t = [Q_t^{1'}, \dots, Q_t^{m'}]'$$

where the $n^2 \times 1$ vector Q_t^i is given by $Q_t^i = \varphi[\mathbb{E}_0(Z_t Z_t' 1_{\xi_t=i})]$. This implies that we can compute the vectorized matrix of second moments $M_t = \varphi[\mathbb{E}_0(Z_t Z_t')]$ as:

$$M_t = \varphi[\mathbb{E}_0(Z_t Z_t')] = \sum_{i=1}^m Q_t^i = WQ_t$$

where the matrix $W = [I_{n^2}, \dots, I_{n^2}]$ is obtained placing side by side m n^2 -dimensional identity matrices. We can then state the following proposition:

Proposition 2 Consider a Markov-switching model whose law of motion can be described by (A12) and define $Q_t^i = \varphi [\mathbb{E}_0 (Z_t Z_t' 1_{\xi_t=i})]$, $q_t^i = \mathbb{E}_0 [Z_t 1_{\xi_t=i}]$, and $\pi_t^i = P_0 (\xi_t = i)$, for $i = 1 \dots m$. Then $Q_t^j = [\widehat{c}c_j + \widehat{V}V_j \varphi [I_k]] \pi_t^j + \sum_{i=1}^m [\widehat{A}A_j Q_{t-1}^i + \widehat{D}A C_j q_{t-1}^i] h_{ji}$, where $\widehat{c}c_j = (c_j \otimes c_j)$, $\widehat{V}V_j = (V_j \otimes V_j)$, $\widehat{A}A_j = (A_j \otimes A_j)$, and $\widehat{D}A C_j = (A_j \otimes c_j) + (c_j \otimes A_j)$.

It is then straightforward to compute the evolution of second moments conditional on the information available at a particular point in time. Suppose we are interested in $\mathbb{E}_t (Z_{t+s} Z_{t+s}')$, i.e. the second moment of the vector Z_{t+s} conditional on the information available at time t . If we define:

$$Q_{t+s|t} = [Q_{t+s|t}^{1'}, \dots, Q_{t+s|t}^{m'}]'$$

where $Q_{t+s|t}^i = \varphi (\mathbb{E}_t (Z_{t+s} Z_{t+s}' 1_{\xi_t=i})) = \varphi (\mathbb{E}_t (Z_{t+s} Z_{t+s}' | \xi_t = i)) \pi_{t+s|t}^i$, we obtain $\varphi (\mathbb{E}_t (Z_{t+s} Z_{t+s}')) = W Q_{t+s|t}$. Using matrix algebra we obtain:

$$Q_{t+s|t} = \Xi Q_{t+s-1|t} + \widehat{D}A C q_{t+s-1|t} + \widehat{V}c \pi_{t+s|t} \quad (\text{A16})$$

$$q_{t+s|t} = C \pi_{t+s|t} + \Omega q_{t+s-1|t}, \quad \pi_{t+s|t} = \mathbf{H} \pi_{t+s-1|t}. \quad (\text{A17})$$

where

$$\begin{aligned} \Xi &= bdiag(\widehat{A}A_1, \dots, \widehat{A}A_m)(\mathbf{H} \otimes I_{n^2}), \quad \widehat{V}c = [\widehat{V}V + \widehat{c}c], \quad \widehat{c}c = bdiag(\widehat{c}c_1, \dots, \widehat{c}c_m), \\ \widehat{V}V &= bdiag(\widehat{V}V_1 \varphi [I_k], \dots, \widehat{V}V_m \varphi [I_k]), \quad \widehat{D}A C = bdiag(\widehat{D}A C_1, \dots, \widehat{D}A C_m)(\mathbf{H} \otimes I_n). \end{aligned}$$

With the first and second moments at hand, it is then possible to compute the variance s periods ahead conditional on the information available at time t :

$$\varphi [\mathbb{V}_t (Z_{t+s})] = M_{t+s|t} - \varphi [\mu_{t+s|t} \mu_{t+s|t}'], \quad (\text{A18})$$

where $M_{t+s|t} = \varphi (\mathbb{E}_t (Z_{t+s} Z_{t+s}')) = \sum_{i=1}^m Q_{t+s|t}^i = W Q_{t+s|t}$.

To report estimates of (A13) and (A18) we proceed as follows. Note that $\mu_{t+s|t} = \mathbb{E}_t (Z_{t+s}) = w q_{t+s|t}$ and $M_{t+s|t}$ depend only on $q_{t+s|t}$ and $Q_{t+s|t}$. Furthermore we can express (A14)-(A15) and (A16)-(A17) in a compact form as

$$\tilde{Q}_{t+s|t} = \tilde{\Xi}^s \tilde{Q}_{t|t} \quad \text{where} \quad \tilde{\Xi} = \left[\begin{array}{c|cc} \Xi & \widehat{D}A C & \widehat{V}c \mathbf{H} \\ \hline & \Omega & C \mathbf{H} \\ & & \mathbf{H} \end{array} \right], \quad (\text{A19})$$

where $\tilde{Q}_{t+s|t} = [Q'_{t+s|t}, q'_{t+s|t}, \pi'_{t+s|t}]'$. Armed with starting values $\tilde{Q}_{t|t} = [Q'_{t|t}, q'_{t|t}, \pi'_{t|t}]'$ we can then compute (A13) and (A18) using (A19). To obtain $\pi'_{t|t}$ recall that we assume that \mathbb{I}_t includes knowledge of the regime in place at time t , the data up to time t , Z^t , and the VAR parameters for each regime. Given that we assume knowledge of the current regime, $\pi_{t|t}^i \equiv P (\xi_t = i | \mathbb{I}_t)$ can only assume two values, 0 or 1. As a result $\pi'_{t|t}$ will be (1, 0) or (0, 1). As

a result, and given $Z_t \in \mathbb{I}_t$, $q'_{t|t} = [q_{t|t}^1, q_{t|t}^2]'$ with $q_{t|t}^i \equiv \mathbb{E}_t(Z_t | \xi_t = i) \pi_{t|t}^i$, will be $[Z_t' \cdot 1, Z_t' \cdot 0]'$ or $[Z_t' \cdot 0, Z_t' \cdot 1]'$. Analogously, $Q'_{t|t} = [Q_{t|t}^1, Q_{t|t}^2]'$ with $Q_{t|t}^i \equiv \varphi(\mathbb{E}_t(Z_t Z_t' | \xi_t = i)) \pi_{t|t}^i$ will be $[\varphi(Z_t Z_t' \cdot 1)', \varphi(Z_t Z_t' \cdot 0)']'$ or $[\varphi(Z_t Z_t' \cdot 0)', \varphi(Z_t Z_t' \cdot 1)']'$.

Mean Square Stability

We consider the following MS-VAR model with n variables and $m = 2$ regimes:

$$Z_t = c_{\xi_t} + A_{1,\xi_t} Z_{t-1} + A_{2,\xi_t} Z_{t-2} + V_{\xi_t} \varepsilon_t, \quad \varepsilon_t \sim N(0, I) \quad (\text{A20})$$

where Z_t is an $n \times 1$ vector of variables, c_{ξ_t} is an $n \times 1$ vector of constants, A_{l,ξ_t} for $l = 1, 2$ is an $n \times n$ matrix of coefficients, $V_{\xi_t} V_{\xi_t}'$ is an $n \times n$ covariance matrix for the $n \times 1$ vector of shocks ε_t . The process ξ_t controls the regime that is in place at time t and evolves based on the transition matrix \mathbf{H} .

When estimating the MS-VAR we require the model to be mean square stable. Mean square stability is defined as follows:

Definition 1 *An n -dimensional process Z_t is mean square stable if and only if there exists an n -vector $\bar{\mu}$ and an n^2 -vector \bar{M} such that:*

- 1) $\lim_{t \rightarrow \infty} \mathbb{E}_0[Z_t] = \bar{\mu}$
- 2) $\lim_{t \rightarrow \infty} \mathbb{E}_0[Z_t Z_t'] = \bar{M}$

for any initial Z_0 and ξ_0 .

Mean-square-stability requires that first and second moments converge as the time horizon goes to ∞ . Under the assumptions that the Markov-switching process ξ_t is ergodic and that the innovation process ε_t is asymptotically covariance stationary, Costa, Fragoso, and Marques (2004) show that a multivariate Markov-switching model as the one described by (A20) is mean-square stable if and only if it is asymptotically covariance stationary. Both conditions hold for the models studied in this paper and are usually verified in economic models.

Costa, Fragoso, and Marques (2004) show that in order to establish MSS of a process such as the one described by (A20), it is enough to check MSS stability of the correspondent homogeneous process: $Z_t = A_{\xi_t} Z_{t-1}$. In this case, formulas for the evolution of first and second moments simplify substantially: $q_t = \Omega q_{t-1}$ and $Q_t = \Xi Q_{t-1}$. Let $r_\sigma(X)$ be the operator that given a square matrix X computes its largest eigenvalue. We then have:

Proposition 3 *A Markov-switching process whose law of motion can be described by (A20) is mean square stable if and only if $r_\sigma(\Xi) < 1$.*

Mean square stability allows us to compute finite measures of uncertainty as the time horizon goes to infinity. Mean square stability also implies that shocks do not have permanent effects on the variables included in the MSVAR.

Conditional Steady State

Consider a MS-VAR:

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

where Z_t is a column vector containing n variables observable at time t and $\xi_t = 1, \dots, m$, with m the number of regimes, evolves following the transition matrix \mathbf{H} . If the MS-VAR has more than one lag, the companion form can be used to recast the model as illustrated above.

The conditional steady state for the mean corresponds to the expected value for the vector Z_t conditional on being in a particular regime. This is computed by imposing that a certain regime is in place forever:

$$\mathbb{E}_i(Z_t) = \bar{\mu}_i = (I_n - A_i)^{-1} c_i \quad (\text{A21})$$

where I_n is an identity matrix with the appropriate size. Note that unless the VAR coefficients imply very slow moving dynamics, after a switch from regime j to regime i , the variables of the VAR will converge (in expectation) to $\mathbb{E}_i(Z_t)$ over a finite horizon. If there are no further switches, we can then expect the variables to fluctuate around $\mathbb{E}_i(Z_t)$. Therefore, the conditional steady states for the mean can also be thought as the values to which the variables converge if regime i is in place for a long enough period of time.

The conditional steady state for the standard deviation corresponds to the standard deviation for the vector Z_t conditional on being in a particular regime. The conditional standard deviations for the elements in Z_t are computed by taking the square root of the main diagonal elements of the covariance matrix $\mathbb{V}_i(Z_t)$ obtained imposing that a certain regime is in place forever:

$$\varphi(\mathbb{V}_i(Z_t)) = (I_{n^2} - A_i \otimes A_i)^{-1} \varphi(V_{\xi_t} V_{\xi_t}') \quad (\text{A22})$$

where I_{n^2} is an identity matrix with the appropriate size, \otimes denotes the Kronecker product, and the vectorization operator $\varphi(X)$ takes a matrix X as an input and returns a column vector stacking the columns of the matrix X on top of one another.

Dynamic Macro-Finance Model: Macro block

This section reports technical details about the MS-DSGE model.

Constant Gain Adaptive Learning

Suppose the representative macro agent believes that inflation evolves according to an AR(1) process:

$$\pi_t = \alpha + \phi \pi_{t-1} + \eta_t. \quad (\text{A23})$$

Macro agents undertake an adaptive learning process whereby they estimate $b \equiv (\alpha, \phi)'$ from past data following

$$\begin{aligned} R_t &= R_{t-1} + \gamma_t (x_{t-1} x'_{t-1} - R_{t-1}) \\ b_t &= b_{t-1} + \gamma_t R_t^{-1} x_{t-1} (\pi_t - b'_{t-1} x_{t-1}) \end{aligned} \quad (\text{A24})$$

where $x_t = (1, \pi_t)'$. Assume that the recursion is started at some point in the distant past. The sequence of gains $0 < \gamma_t < 1$ determines the speed of updating when faced with an inflation surprise at time t . For $\gamma_t = 1/t$ the algorithm represents a recursive formulation of an ordinary least squares estimation that uses all available data until time t with equal weights (see Evans and Honkapohja (2001)). By contrast, for constant $\gamma_t = \gamma$, it represents a constant-gain learning algorithm with exponentially decaying weights on past observations. This implies that the agent gives more weight to the more recent observations, possibly to guard against parameter instability, as in this model. This specification simplifies if we assume that agents are only uncertain about the long term value of inflation, but not its persistence. If agents only learn about α and the recursion has started in the distant past we have:

$$R_t = 1 \text{ if } R_{t-1} = 1 \quad (\text{A25})$$

$$\alpha_t^m = \alpha_{t-1}^m + \gamma_t (\pi_t - \phi \pi_{t-1} - \alpha_{t-1}^m) \quad (\text{A26})$$

To see the above, note that if ϕ were known, the agent would estimate α by running a regression of $\pi_t - \phi \pi_{t-1}$ on a constant, or a vector of ones. So $x_t = 1$ in every period and $R_t = R_{t-1} + \gamma_t (x_{t-1} x'_{t-1} - R_{t-1}) = R_t = R_{t-1} + \gamma_t (1 - R_{t-1})$. Starting value for $R = R_0 \Rightarrow R_1 = R_0 + \gamma(1 - R_0)$. Continuing to iterate, this converges to 1 no matter what R_0 as long as $0 < \gamma_t < 1$. Set $x_t = R_1 = 1$ in (A24) to get (A26).

With constant gain learning, the variable γ_t is a constant parameter that we denote γ . This implies:

$$\alpha_t^m = \alpha_{t-1}^m + \gamma (\pi_t - \phi \pi_{t-1} - \alpha_{t-1}^m). \quad (\text{A27})$$

Hereafter we assume that expectations are formed using a constant gain adaptive rule.

Perceived trend inflation, $\bar{\pi}_t$, is defined as the $\lim_{h \rightarrow \infty} \mathbb{E}_t^m (\pi_{t+h})$. Observe that, since expectations obey the constant gain adaptive rule, $\bar{\pi}_t$ is not constant but varies with information at time t . This can be seen by taking expectations on both sides of equation A23 to find,

$$\begin{aligned} \bar{\pi}_t &= \lim_{h \rightarrow \infty} \mathbb{E}_t^m (\pi_{t+h}) \\ &= \lim_{h \rightarrow \infty} \mathbb{E}_t^m (\alpha_t^m + \phi \pi_{t+h-1}) \\ &= \lim_{h \rightarrow \infty} \mathbb{E}_t^m (\alpha_t^m + \phi \alpha_t^m + \phi^2 \alpha_t^m + \dots \phi^{h-1} \alpha_t^m + \phi^h \pi_t) \\ &= \alpha_t^m / (1 - \phi), \end{aligned}$$

where we plug in the value of α_t^m that agents perceive at t as the last step. In the above we have used the standard notion of “anticipated utility,” whereby beliefs at time t about α_t^m are perceived by the agent to hold forever in the future; i.e., the agent does not recognize that she will update her estimate of α_t^m in future periods. With this, the AR(1) process implies a one-to-one mapping between the perceived constant α_t^m and perceived trend inflation $\bar{\pi}_t$. Using the relation between $\bar{\pi}_t$ and α_t^m , we get:

$$\begin{aligned}\alpha_t^m &= (1 - \phi) \bar{\pi}_t \Rightarrow \\ (1 - \phi) \bar{\pi}_t &= (1 - \phi) \bar{\pi}_{t-1} + \gamma (\pi_t - \phi \pi_{t-1} - (1 - \phi) \bar{\pi}_{t-1}) \Rightarrow \\ \bar{\pi}_t &= \bar{\pi}_{t-1} + \gamma (1 - \phi)^{-1} (\pi_t - \phi \pi_{t-1} - (1 - \phi) \bar{\pi}_{t-1}),\end{aligned}\tag{A28}$$

where the second equation above follows from (A27).

Finally, the unconditional mean of inflation as perceived by the agent is estimated recursively under the constant gain adaptive rule and so depends on the sample of data she uses at time t to estimate α . Denote this information \mathbb{I}_t . Taking perceived unconditional means on both sides of (A23), we find that the unconditional mean of inflation as perceived by the agent at time t is the same as perceived trend inflation:

$$\mathbb{E}^m (\pi_t | \mathbb{I}_t) = \alpha + \phi \mathbb{E}^m (\pi_t | \mathbb{I}_t) \Rightarrow \mathbb{E}_{T_t^m}^m (\pi_t) = \bar{\pi}_t = \alpha_t^m / (1 - \phi).$$

Signal About the Inflation Target In our model, we combine the constant gain learning algorithm described above with a signal about the central bank’s inflation target, thereby allowing beliefs to be partly shaped by additional information the agent receives about the target. This signal could reflect the opinion of experts (as in MN), or a credible central bank announcement. If we use α_t^{mCG} and $\bar{\pi}_t^{CG}$ to denote the beliefs implied by the constant gain learning described above, we obtain modified updating rules for α_t^m and $\bar{\pi}_t$ that are a weighted averages of two terms:

$$\begin{aligned}\alpha_t^m &= (1 - \gamma^T) \left[\underbrace{\alpha_{t-1}^m + \gamma (\pi_t - \phi \pi_{t-1} - \alpha_{t-1}^m)}_{\alpha_t^{mCG}} \right] + \gamma^T \left[(1 - \phi) \pi_{\xi_t}^T \right]. \\ \bar{\pi}_t &= (1 - \gamma^T) \left[\underbrace{\bar{\pi}_{t-1} + \gamma (1 - \phi)^{-1} (\pi_t - \phi \pi_{t-1} - (1 - \phi) \bar{\pi}_{t-1})}_{\bar{\pi}_t^{CG}} \right] + \gamma^T \left[\pi_{\xi_t}^T \right].\end{aligned}$$

The first terms in square brackets, α_t^{mCG} and $\bar{\pi}_t^{CG}$, are the recursive updating rules implied by constant gain learning as in (A27) and (A28). These terms are combined with two terms that involve the central bank’s current inflation target $\pi_{\xi_t}^T$. Note that, since $\alpha_t^m = (1 - \phi) \bar{\pi}_t$

under the autoregressive model, the term $(1 - \phi) \pi_{\xi_t}^T$ is simply the value of α_t^m that would arise if $\bar{\pi}_t = \pi_{\xi_t}^T$. 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 learning algorithm. Thus, the resulting laws of motion for the beliefs are a weighted average of what would arise under constant gain learning and a term reflecting information about the current inflation target.

Expected inflation

Expected inflation from the point of view of the agents in the model is formed based on equation A23 and their beliefs about the constant α , i.e., α_t^m . Specifically, we have

$$\begin{aligned}\mathbb{E}_t^m [\pi_{t+1}] &= \alpha_t^m + \phi \pi_t \\ \mathbb{E}_t^m [\pi_{t+2}] &= \alpha_t^m + \phi \alpha_t^m + \phi^2 \pi_t \\ \mathbb{E}_t^m [\pi_{t+3}] &= \alpha_t^m + \phi \alpha_t^m + \phi^2 \alpha_t^m + \phi^3 \pi_t \\ \mathbb{E}_t^m [\pi_{t+4}] &= \alpha_t^m + \phi \alpha_t^m + \phi^2 \alpha_t^m + \phi^3 \alpha_t^m + \phi^4 \pi_t\end{aligned}$$

where, in line with the learning literature, we have assumed that agents do not take into account that their beliefs might change in the future (i.e., they do not have anticipated utility).

Cumulative inflation over the next year is:

$$\begin{aligned}\mathbb{E}_t^m [\pi_{t,t+4}] &= [4 + 3\phi + 2\phi^2 + \phi^3] \alpha_t^m + [\phi + \phi^2 + \phi^3 + \phi^4] \pi_t \\ &= [4 + 3\phi + 2\phi^2 + \phi^3] (1 - \phi) \bar{\pi}_t + [\phi + \phi^2 + \phi^3 + \phi^4] \pi_t\end{aligned}$$

where in the second row we have used the fact that $\bar{\pi}_t = \alpha_t^m / (1 - \phi)$. The general formulas are:

$$\begin{aligned}\mathbb{E}_t^m [\pi_{t+h}] &= \alpha_t^m + \phi \alpha_t^m + \dots + \phi^{h-1} \alpha_t^m + \phi^h \pi_t \\ \mathbb{E}_t^m [\pi_{t,t+h}] &= (1/h) \sum_{i=1}^h \mathbb{E}_t^m [\pi_{t+i}]\end{aligned}$$

Using matrix algebra, we can express the *perceived* law of motion for inflation as:

$$\begin{bmatrix} \alpha_t^m \\ \pi_{t+1} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 1 & \phi \end{bmatrix} \begin{bmatrix} \alpha_t^m \\ \pi_t \end{bmatrix} + \begin{bmatrix} 0 \\ \eta_{t+1} \end{bmatrix}$$

This is equivalent to:

$$\begin{bmatrix} \bar{\pi}_t \\ \pi_{t+1} \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 0 \\ 1 - \phi & \phi \end{bmatrix}}_{\Omega} \underbrace{\begin{bmatrix} \bar{\pi}_t \\ \pi_t \end{bmatrix}}_{e_{\pi\bar{\pi}} S_t} + \begin{bmatrix} 0 \\ \eta_{t+1} \end{bmatrix}$$

where once again we have used $\bar{\pi}_t = \alpha_t^m / (1 - \phi)$ and the matrix $e_{\pi\bar{\pi}}$ is used to extract both inflation π_t and the perceived long term inflation $\bar{\pi}_t$ from the state vector S_t . The latter

formulation is used for the solution of the model since it is $\bar{\pi}_t$ rather than α_t^m that appears in the state space representation of the model. It follows that

$$\mathbb{E}_t^m [\pi_{t,t+h}] = e_\pi \Omega (I - \Omega)^{-1} (I - \Omega^4) (e_{\pi\bar{\pi}} S_t)$$

where the vector e_π is used to extract inflation.

Long-run Monetary Neutrality

Suppose the central bank were to permanently change the inflation target. Would this have a long-run influence on real activity? In a model with rational expectations, the relation between inflation and the output gap is controlled by a New-Keynesian Phillips curve:

$$\pi_t - \bar{\pi}_t = \beta \mathbb{E}_t [\pi_{t+1} - \bar{\pi}_t] + \kappa [y_{t-1} - y_{t-1}^*]$$

where $\bar{\pi}_t$ denotes the long term value of inflation that coincides, under rational expectations, with the central bank's inflation target $\pi_{\xi_t}^T$. Taking unconditional expectation on both sides, we have:

$$\begin{aligned} \mathbb{E} [\pi_t - \bar{\pi}_t] &= \beta \mathbb{E} [\pi_{t+1} - \bar{\pi}_t] + \kappa \mathbb{E} [y_{t-1} - y_{t-1}^*] \\ \mathbb{E} [\pi_t] - \mathbb{E} [\bar{\pi}_t] &= \beta \mathbb{E} [\bar{\pi}_t] - \beta \mathbb{E} [\bar{\pi}_t] + \kappa \mathbb{E} [y_{t-1} - y_{t-1}^*] \\ 0 &= \kappa \mathbb{E} [y_{t-1} - y_{t-1}^*] \end{aligned}$$

where we have used the fact that $\bar{\pi}_t = \mathbb{E} [\pi_t] = \pi_{\xi_t}^T$. Therefore, we have: $\mathbb{E} [y_{t-1}^*] = \mathbb{E} [y_{t-1}] = 0$. Thus in the long-run, real output is expected to equal the natural rate, and monetary policy is neutral.

With sticky expectations, long-term neutrality still holds. In a rational expectations model, the econometrician's beliefs and the agent's beliefs about trend inflation are always aligned, even in the short-run. These beliefs in turn align with the central bank's target inflation. In the constant gain adaptive world, the agent's beliefs about long-term inflation, $\bar{\pi}_t$, align with the econometrician's beliefs and with the central bank's inflation target only in the long-run. But, even with sticky expectations, if the central bank permanently changes the target, we still have $\lim_{h \rightarrow \infty} \mathbb{E}_t [\bar{\pi}_{t+h}] = \mathbb{E} [\bar{\pi}_t] = \pi_{\xi_t}^T = \mathbb{E} [\pi_t]$, where $\mathbb{E}_t [\cdot]$ denotes the expectations of the econometrician. Then,

$$\begin{aligned} \pi_t - \bar{\pi}_t &= \beta \phi [\pi_t - \bar{\pi}_t] + \kappa [y_{t-1} - y_{t-1}^*] \\ \mathbb{E} [\pi_t] &= \mathbb{E} [\bar{\pi}_t] + \frac{\kappa}{1 - \beta \phi} \mathbb{E} [y_{t-1} - y_{t-1}^*] \\ 0 &= \kappa \mathbb{E} [y_{t-1} - y_{t-1}^*]. \end{aligned}$$

Therefore, we again have: $\mathbb{E} [y_{t-1}^*] = \mathbb{E} [y_{t-1}] = 0$.

Solution and Estimation of the Macro Block

We can rewrite the system of equations as:

$$\tilde{y}_t = \varrho \tilde{y}_{t-1} - \sigma [i_t - \phi \pi_t - (1 - \phi) \bar{\pi}_t - r] + f_t \quad (\text{A29})$$

$$\pi_t = \bar{\pi}_t + \frac{\kappa}{1 - \beta \phi} [y_{t-1} - y_{t-1}^*] \quad (\text{A30})$$

$$i_t - \left(r + \pi_{\xi_t}^T \right) = (1 - \rho_{i, \xi_t}) \left[\psi_{\pi, \xi_t} \left(\pi_t - \pi_{\xi_t}^T \right) + \psi_{\Delta y, \xi_t} (\tilde{y}_t - \tilde{y}_{t-1}) \right] \quad (\text{A31})$$

$$+ \rho_{i, \xi_t} \left[i_{t-1} - \left(r + \pi_{\xi_t}^T \right) \right] + \sigma_i \varepsilon_{i,t} \quad (\text{A32})$$

$$\tilde{y}_t^* = \rho_{y^*} \tilde{y}_{t-1}^* + \sigma_{y^*} \varepsilon_{y^*,t} \quad (\text{A33})$$

$$\bar{\pi}_t = [1 - \gamma^T] \left[\bar{\pi}_{t-1} + \gamma (1 - \phi)^{-1} (\pi_t - \phi \pi_{t-1} - (1 - \phi) \bar{\pi}_{t-1}) \right] + \gamma^T \pi_{\xi_t}^T \quad (\text{A34})$$

$$f_t = \rho_f f_{t-1} + \sigma_f \varepsilon_{f,t} \quad (\text{A35})$$

State space and parameter vectors

Define the parameter vectors θ_{ξ_t} and $\theta_{\xi_t}^c$ as

$$\theta_{\xi_t} = [\varrho, \sigma, \beta, \kappa, \psi_{\pi, \xi_t}, \psi_{\Delta y, \xi_t}, \rho_{i, \xi_t}, \rho_{y^*}, \gamma^T, \gamma, \phi, \rho_f]'$$

$$\theta_{\xi_t}^c = [\pi_{\xi_t}^T, r]'$$

and the state vector S_t and the vector of Gaussian shocks ε_t as

$$S_t = [\tilde{y}_t, \tilde{y}_t^*, \pi_t, i_t, \bar{\pi}_t, f_t]'$$

$$\varepsilon_t = [\varepsilon_{i,t}, \varepsilon_{y_t^*}, \varepsilon_{f,t}]', \quad \varepsilon_t \sim N(0, I)$$

Let the matrix $Q = \text{diag}(\sigma_i, \sigma_{y^*}, \sigma_d)$ be a square matrix with the shock standard deviations on the main diagonal. Conditional on each regime, the system of equations can be rewritten using matrix notation:

$$\Gamma_0(\theta_{\xi_t}) S_t = \Gamma_c(\theta_{\xi_t}^c) + \Gamma_1(\theta_{\xi_t}) S_{t-1} + Q \varepsilon_t$$

Note that the vector $\Gamma_c(\theta_{\xi_t}^c)$ includes the inflation target for the corresponding regime.

Inverting the matrix $\Gamma_0(\theta_{\xi_t})$, we obtain the solution of the model as MS-VAR:

$$S_t = C(\theta_{\xi_t}^c, \theta_{\xi_t}) + T(\theta_{\xi_t}) S_{t-1} + R(\theta_{\xi_t}) Q \varepsilon_t$$

where $C(\theta_{\xi_t}^c, \theta_{\xi_t}) = \Gamma_0^{-1}(\theta_{\xi_t}) \Gamma_c(\theta_{\xi_t}^c)$, $T(\theta_{\xi_t}) = \Gamma_0^{-1}(\theta_{\xi_t}) \Gamma_1(\theta_{\xi_t})$, and $R(\theta_{\xi_t}) = \Gamma_0^{-1}(\theta_{\xi_t})$.

The solution of the model can be combined with an observation equation to estimate the model. Given that we know the regime sequence, we can estimate the model with a standard

Kalman filter algorithm. The only caveat is that the associated transition equation (A37), below, varies over time. We thus have the following state space representation:

$$X_t = D + Z [S'_t, y_{t-1}]' + Uv_t \quad (\text{A36})$$

$$S_t = C \left(\theta_{\xi_t}^c, \theta_{\xi_t} \right) + T(\theta_{\xi_t})S_{t-1} + R(\theta_{\xi_t})Q\varepsilon_t \quad (\text{A37})$$

$$v_t \sim N(0, I) \quad (\text{A38})$$

where v_t is a vector of observation errors and U is a diagonal matrix with the standard deviations of the observation errors on the main diagonal. As said before, we condition on a regime sequence ξ_t , so the transition equation (A37) in at each point in time is known.

In our estimation, we use four observables: Real GDP per capita growth, Inflation, Federal Funds rate, and the mean of the Michigan survey one-year-ahead inflation forecasts. All variables are annualized. We have observation errors on all variables because we have 3 shocks for four observables.

Thus, the vector of data X_t is defined as:

$$\begin{bmatrix} \Delta GDP \\ Inflation \\ FFR \\ E(Inflation) \end{bmatrix} = \begin{bmatrix} \overline{\Delta GDP} \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 4y_t - 4y_{t-1} \\ 4\pi_t \\ 4i_t \\ [4 + 3\phi + 2\phi^2 + \phi^3] (1 - \phi) \bar{\pi}_t + [\phi + \phi^2 + \phi^3 + \phi^4] \pi_t \end{bmatrix} + \begin{bmatrix} v_t^y \\ v_t^\pi \\ v_t^f \\ v_t^e \end{bmatrix}$$

where in the last row we have used the fact that expectations for an agent in the model is:

$$\begin{aligned} \mathbb{E}_t^m [\pi_{t,t+4}] &= [4 + 3\phi + 2\phi^2 + \phi^3] \alpha_t^m + [\phi + \phi^2 + \phi^3 + \phi^4] \pi_t \\ &= [4 + 3\phi + 2\phi^2 + \phi^3] (1 - \phi) \bar{\pi}_t + [\phi + \phi^2 + \phi^3 + \phi^4] \pi_t \end{aligned}$$

The mapping from the variables of the model to the observables can be written using matrix algebra. The vector D is then:

$$D = \begin{bmatrix} \overline{\Delta GDP} \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

The matrix Z is then:

$$Z = \begin{bmatrix} 4 & 0 & 0 & 0 & 0 & 0 & 0 & -4 \\ 0 & 0 & 4 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 4 & 0 & 0 & 0 & 0 \\ 0 & 0 & [\phi + \phi^2 + \phi^3 + \phi^4] & 0 & 0 & 0 & [4 + 3\phi + 2\phi^2 + \phi^3] (1 - \phi) & 0 \end{bmatrix}$$

Note that the matrix Z loads detrended output (y_t) and lagged detrended output (y_{t-1}).

The likelihood is computed with the Kalman filter and then combined with a prior distribution for the parameters to obtain the posterior. As a first step, a block algorithm is used to find the posterior mode, while a Metropolis-Hastings algorithm is used to draw from the posterior distribution.

Draws from the posterior are obtained using a standard Metropolis-Hastings algorithm initialized around the posterior mode. When working with models whose posterior distribution is very complicated in shape it is very important to find the posterior mode. Here are the key steps of the Metropolis-Hastings algorithm:

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

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

The only aspect of the estimation that it is not traditional is that the transition equation (A37) varies over time. However, given that we estimate the model fixing the regime sequence, we can easily modify the standard Kalman filter to handle this change. Specifically, the modified Kalman filter is described as follows.

Given a sequence of regimes $\xi^T = \xi_1 \dots \xi_T$, the Kalman filter involves the following steps for each $t = 1 \dots T$:

1. Prediction:

$$S_{t|t-1} = C(\theta_{\xi_t}^c, \theta_{\xi_t}) + T(\theta_{\xi_t})S_{t-1|t-1} \quad (\text{A39})$$

$$P_{t|t-1} = T(\theta_{\xi_t})P_{t-1|t-1}T(\theta_{\xi_t})' + R(\theta_{\xi_t})Q^2R(\theta_{\xi_t})' \quad (\text{A40})$$

$$\eta_{t|t-1} = X_t - X_{t|t-1} = X_t - D - Z * S_{t|t-1} \quad (\text{A41})$$

$$f_{t|t-1} = ZP_{t|t-1}Z' + U^2 \quad (\text{A42})$$

2. Updating

$$S_{t|t} = S_{t|t-1} + K_t \eta_{t|t-1} \quad (\text{A43})$$

$$P_{t|t} = P_{t|t-1} - K_t Z P_{t|t-1} \quad (\text{A44})$$

where $K_t = P_{t|t-1} Z' f_{t|t-1}^{-1}$ is the Kalman gain.

The log-likelihood $\ln L$ is then obtained as:

$$\ln L = -.5 \sum_{t=1}^T \ln (2\pi f_{t|t-1}) - .5 \sum_{t=1}^T \eta'_{t|t-1} f_{t|t-1}^{-1} \eta_{t|t-1}.$$

Details about the solution. The matrices used to write the model in state space form are described below.

Equations:

$$\begin{aligned} y_t &= \varrho y_{t-1} - \sigma [i_t - \phi \pi_t - (1 - \phi) \bar{\pi}_t - r] + f_t \\ \pi_t &= \bar{\pi}_t + \frac{\kappa}{1 - \beta \phi} [y_{t-1} - y_{t-1}^*] \end{aligned} \quad (\text{A45})$$

$$i_t - \left(r + \pi_{\xi_t}^T \right) = (1 - \rho_{i,\xi_t}) \left[\psi_{\bar{\pi},\xi_t} \left(\bar{\pi}_t - \pi_{\xi_t}^T \right) + \psi_{\pi,\xi_t} \left(\pi_t - \pi_{\xi_t}^T \right) + \psi_{\Delta y,\xi_t} (y_t - y_{t-1}) \right] \quad (\text{A46})$$

$$+ \rho_{i,\xi_t} \left[i_{t-1} - \left(r + \pi_{\xi_t}^T \right) \right] + \sigma_i \varepsilon_{i,t} \quad (\text{A47})$$

$$y_t^* = \rho_{y^*} y_{t-1}^* + \sigma_{y^*} \varepsilon_{y^*,t} \quad (\text{A48})$$

$$\bar{\pi}_t = [1 - \gamma^T] [\bar{\pi}_{t-1} + \gamma (1 - \phi)^{-1} (\pi_t - \phi \pi_{t-1} - (1 - \phi) \bar{\pi}_{t-1})] + \gamma^T \pi_{\xi_t}^T + \sigma_{\bar{\pi}} \varepsilon_{\bar{\pi},t} \quad (\text{A49})$$

$$f_t = \rho_f f_{t-1} + \sigma_f \varepsilon_{f,t} \quad (\text{A50})$$

We get:

$$\begin{aligned}
y_t &= \delta y_{t-1} - \sigma i_t + \sigma \phi \pi_t + \sigma(1 - \phi)\bar{\pi}_t + \sigma r + f_t \\
\pi_t &= \bar{\pi}_t + \frac{\kappa}{1 - \beta\phi} y_{t-1} - \frac{\kappa}{1 - \beta\phi} y_{t-1}^* \\
i_t - (r + \pi_{\xi_t}^T) &= (1 - \rho_{i,\xi_t}) \psi_{\bar{\pi},\xi_t} \bar{\pi}_t - (1 - \rho_{i,\xi_t}) \psi_{\bar{\pi},\xi_t} \pi_{\xi_t}^T \\
&\quad + (1 - \rho_{i,\xi_t}) \psi_{\pi,\xi_t} \pi_t - (1 - \rho_{i,\xi_t}) \psi_{\pi,\xi_t} \pi_{\xi_t}^T \\
&\quad + (1 - \rho_{i,\xi_t}) \psi_{\Delta y,\xi_t} y_t - (1 - \rho_{i,\xi_t}) \psi_{\Delta y,\xi_t} y_{t-1} \\
&\quad + \rho_{i,\xi_t} i_{t-1} - \rho_{i,\xi_t} (r + \pi_{\xi_t}^T) \\
&\quad + \sigma_i \varepsilon_{i,t} \\
y_t^* &= \rho_y^* y_{t-1}^* + \sigma_{y^*} \varepsilon_{y^*,t} \\
\bar{\pi}_t &= (1 - \gamma^T) \bar{\pi}_{t-1} + (1 - \gamma^T) \gamma (1 - \phi)^{-1} \pi_t \\
&\quad - (1 - \gamma^T) \gamma (1 - \phi)^{-1} \phi \pi_{t-1} \\
&\quad - (1 - \gamma^T) \gamma (1 - \phi)^{-1} (1 - \phi) \bar{\pi}_{t-1} \\
&\quad + \gamma^T \pi_{\xi_t}^T + \sigma_{\bar{\pi}} \varepsilon_{\bar{\pi},t} \\
f_t &= \rho_f f_{t-1} + \sigma_f \varepsilon_{f,t}
\end{aligned}$$

Equations with state variables at t on the LHS, everything else on the RHS, and re-ordered to match the state variable vector:

$$\begin{aligned}
y_t + \sigma i_t - \sigma \phi \pi_t - \sigma(1 - \phi)\bar{\pi}_t - f_t &= \delta y_{t-1} + \sigma r \\
y_t^* &= \rho_y^* y_{t-1}^* + \sigma_{y^*} \varepsilon_{y^*,t} \\
\pi_t - \bar{\pi}_t &= + \frac{\kappa}{1 - \beta\phi} y_{t-1} - \frac{\kappa}{1 - \beta\phi} y_{t-1}^* \\
i_t - (1 - \rho_{i,\xi_t}) \psi_{\bar{\pi},\xi_t} \bar{\pi}_t - (1 - \rho_{i,\xi_t}) \psi_{\pi,\xi_t} \pi_t - (1 - \rho_{i,\xi_t}) \psi_{\Delta y,\xi_t} y_t &= -(1 - \rho_{i,\xi_t}) \psi_{\bar{\pi},\xi_t} \pi_{\xi_t}^T \\
&\quad - (1 - \rho_{i,\xi_t}) \psi_{\pi,\xi_t} \pi_{\xi_t}^T \\
&\quad - (1 - \rho_{i,\xi_t}) \psi_{\Delta y,\xi_t} y_{t-1} \\
&\quad + \rho_{i,\xi_t} i_{t-1} - \rho_{i,\xi_t} (r + \pi_{\xi_t}^T) \\
&\quad + \sigma_i \varepsilon_{i,t} + (r + \pi_{\xi_t}^T) \\
\bar{\pi}_t - (1 - \gamma^T) \gamma (1 - \phi)^{-1} \pi_t &= (1 - \gamma^T) \bar{\pi}_{t-1} \\
&\quad - (1 - \gamma^T) \gamma (1 - \phi)^{-1} \phi \pi_{t-1} \\
&\quad - (1 - \gamma^T) \gamma (1 - \phi)^{-1} (1 - \phi) \bar{\pi}_{t-1} \\
&\quad + \gamma^T \pi_{\xi_t}^T + \sigma_{\bar{\pi}} \varepsilon_{\bar{\pi},t} \\
f_t &= \rho_f f_{t-1} + \sigma_f \varepsilon_{f,t}
\end{aligned}$$

Goal: matrix form with $\Gamma_0 S_t = \Gamma_C + \Gamma_1 S_{t-1} + \Psi Q \varepsilon_t$. Γ_0 and Γ_1 are 6×6 matrices. Γ_C is 6×1 . Ψ is 6×4 .

State variables: $S_t = [y_t, y_t^*, \pi_t, i_t, \bar{\pi}_t, f_t]'$.

Stochastic variables: $Q = \text{diag}(\sigma_i, \sigma_{y^*}, \sigma_{\bar{\pi}}, \sigma_d)$.

First, Γ_0 , which is for the time t state variables on the LHS. Empty cells are zero.

$$\Gamma_0 = \begin{bmatrix} & y_t & y_t^* & \pi_t & i_t & \bar{\pi}_t & f_t \\ y_t & 1 & & -\sigma\phi & \sigma & -\sigma(1-\phi) & -1 \\ y_t^* & & 1 & & & & \\ \pi_t & & & 1 & & -1 & \\ i_t & -(1-\rho_{i,\xi_t})\psi_{\Delta y,\xi_t} & & -(1-\rho_{i,\xi_t})\psi_{\pi,\xi_t} & 1 & -(1-\rho_{i,\xi_t})\psi_{\bar{\pi},\xi_t} & \\ \bar{\pi}_t & & & -(1-\gamma^T)\gamma(1-\phi)^{-1} & & 1 & \\ f_t & & & & & & 1 \end{bmatrix}$$

Next, Γ_1 , which is for the time $t-1$ state variables on the RHS. Empty cells are zero.

$$\Gamma_1 = \begin{bmatrix} & y_{t-1} & y_{t-1}^* & \pi_{t-1} & i_{t-1} & \bar{\pi}_{t-1} & f_{t-1} \\ y_t & \varrho & & & & & \\ y_t^* & & \rho_y^* & & & & \\ \pi_t & \frac{\kappa}{1-\beta\phi} & -\frac{\kappa}{1-\beta\phi} & & & & \\ i_t & -(1-\rho_{i,\xi_t})\psi_{\Delta y,\xi_t} & & & \rho_{i,\xi_t} & & \\ \bar{\pi}_t & & & -(1-\gamma^T)\gamma(1-\phi)^{-1}\phi & & (1-\gamma^T)(1-\gamma) & \\ f_t & & & & & & \rho_f \end{bmatrix}$$

Ψ inserts the stochastic processes into each of the equations. Empty cells are zero.

$$\Psi = \begin{bmatrix} & \varepsilon_{i,t} & \varepsilon_{y^*,t} & \varepsilon_{\bar{\pi},t} & \varepsilon_{f,t} \\ y_t & & & & \\ y_t^* & & \sigma_{y^*} & & \\ \pi_t & & & & \\ i_t & \sigma_i & & & \\ \bar{\pi}_t & & & \sigma_{\bar{\pi}} & \\ f_t & & & & \sigma_f \end{bmatrix}$$

Finally, Γ_C collects all of the leftover constant terms on the RHS.

$$\Gamma_C = \begin{array}{c|c} \hline y_t & \sigma r \\ \hline y_t^* & 0 \\ \hline \pi_t & 0 \\ \hline i_t & (1 - \rho_{i,\xi_t})[r + \pi_{\xi_t}^T (1 - \psi_{\pi,\xi_t} - \psi_{\bar{\pi},\xi_t})] \\ \hline \bar{\pi}_t & \gamma^T \pi_{\xi_t}^T \\ \hline f_t & 0 \\ \hline \end{array}$$

Dynamic Macro-Finance Model: Asset prices

In this section, we provide details on how to solve for asset prices in the baseline model with learning on the side of the AP agent. Note that learning on the side of the AP agent does not affect the dynamics of the macro block, only the beliefs of the AP agent about the future evolution of monetary policy. These beliefs affect forecasts of the AP agent about all macro variables in the model and current asset prices.

The results on the evolution of the AP agent's beliefs that we present below build on Bianchi and Melosi (2016). Bianchi and Melosi (2016) develop methods to solve general equilibrium models in which forward-looking agents are uncertain about the statistical properties of the regime changes that they observe. For example, when observing Hawkish monetary policy agents might be uncertain whether such policy rule will persist for a long time or not.

Agents in the model are fully rational, conduct Bayesian learning, and they know that they do not know. Therefore, when forming expectations, agents take into account that their beliefs will evolve according to what they will observe in the future. A maintained assumption of Bianchi and Melosi (2016) is that agents know the transition matrix governing regime changes. However, some regimes only differ in terms of their persistence and the probability of moving to different regimes. Thus, agents engage in Bayesian learning to uncover what kind of policy regime they are currently facing (short-lasting or long-lasting). This implies that agents are still rational, but not perfectly informed.

In this paper, we depart from the assumption that the transition matrix guiding the Bayesian learning process coincides with the DGP transition matrix. This allows us to capture a series of behavioral features that help in explaining the response of asset valuation to structural changes in the conduct of monetary policy. First, while asset pricing agents might always be aware of what the central bank is *currently* doing, they might be uncertain about what this implies for its future behavior. Second, if agents have spent a long time in one policy regime, they might experience excess-extrapolation about what this implies for future monetary policy and memories of previous regimes might fade away. Finally, consistently with the previous assumption, when encountering a policy change after a prolonged period under the same policy regime, agents might initially consider the policy change as temporary and expect to revert to the old regime. Only after spending enough time in the new regime, they might come to

consider the regime change as a structural one.

Beliefs: Overview

The policy rule follows two regime $\xi_t = H$ for hawkish and $\xi_t = D$ for dovish. We assume that the asset pricing agent observes all variables of the economy in the current time period t . If agents can also observe the regime in place ξ_t and know the transition matrix H governing the probability of moving across regimes, we have the full information rational expectations model.

Define the augmented state space $\tilde{S}_t = [S_t, m_t, pd_t, \mathbb{E}_t^p(m_{t+1}), \mathbb{E}_t^p(pd_{t+1})]'$. Suppose first that agents can observe the monetary policy regime in place and they form expectations based on the transition matrix \mathbf{H} of the true data generating process (DGP) transitions across the two policy regimes. In this case, the model can be expressed in the following form:

$$\Gamma_{0,\xi_t} \tilde{S}_t = \Gamma_{c,\xi_t} + \Gamma_{1,\xi_t} \tilde{S}_{t-1} + \Psi_{\xi_t} \varepsilon_t + \Pi \eta_t \quad (\text{A51})$$

where η_t is a vector containing the endogenous expectation errors, and the random vector ε_t contains the familiar Gaussian shocks. The variable ξ_t controls the parameter values in place at time, $\theta(\xi_t)$, assumes discrete values $\xi_t \in \{1, 2\}$, and evolves according to a Markov-switching process with transition matrix \mathbf{H} . Denote the true DGP transition probabilities

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

in which the probability of switching to regime j given that we are in regime i is denoted by p_{ij} , where $j = H, D$. The model can then be solved with any of the solution algorithms developed for Markov Switching Rational Expectations (MS-RE) models.

Now suppose agents have a distorted transition matrix \mathbf{H}^p that differs from \mathbf{H} . The model can be solved in the same way, replacing \mathbf{H} with the perceived transition matrix \mathbf{H}^p . This gives us the “no learning” distorted beliefs case reported in the text, in which agents correctly observe the monetary policy regime in place today, but overstate the probability of remaining in the current regime.

Finally, under the baseline model, we combine distorted beliefs with learning about the persistence of policy regimes. In this case, when a monetary policy regime change occurs, the asset pricing agent initially perceives the shift as a transitory deviation from the old regime, effectively underestimating the true persistence of the regime change. However, as she spends more time in the new regime the agent comes to believe that a structural change has occurred, effectively overstating the true persistence of the regime change. Thus the probabilities that the agent assigns to future monetary policy regimes changes over time. To capture this idea, we introduce the perceived regime sequence $\xi_t^p \in \{1, 2, 3, 4\}$. Some of these perceived regimes are assumed to bring about the same macro block model parameters, $\theta(\xi_t^p)$. Specifically, two of

the perceived regimes are characterized by Hawkish monetary policy, while two of the perceived regimes are characterized by Dovish monetary policy. Without loss of generality, we assume that regimes $\xi_t^p = 1$ and $\xi_t^p = 2$ belong to a block 1: $b_1 = \{\xi_t^p \in \{1, 2\} : \theta(\xi_t) = \theta_{b_1}\}$, characterized by Hawkish monetary policy ($\xi_t = 1$), while regimes $\xi_t^p = 3$ and $\xi_t^p = 4$ belong to a block 2: $b_2 = \{\xi_t^p \in \{3, 4\} : \theta(\xi_t) = \theta_{b_2}\}$, characterized by Dovish monetary policy ($\xi_t = 2$). The regime $\xi_t^p = 1$ is perceived as a short-lasting Hawkish regime, while $\xi_t^p = 2$ is perceived as a long-lasting Hawkish regime. The perceived regime $\xi_t^p = 3$ is assumed to be a short-lasting Dovish regime, while the perceived regime $\xi_t^p = 4$ is assumed to be long-lasting Dovish regime.

Given that agents know the structure of the model and can observe the endogenous variables and the shocks, they can also determine which set of parameters is in place at each point in time. In other words, they can tell whether monetary policy is dovish or hawkish and can back out the history of policy regimes. This allows them to determine ξ_t and the block b_j in place at time t . However, while this is enough for agents to establish the history of blocks, agents cannot exactly infer the realized regime ξ_t^p , because the regimes within each block share the same parameter values. It is important to emphasize that regimes that belong to the same block are not identical in all respects, as they differ in their perceived persistence and therefore the probability of switching to other perceived regimes.

The perceived probabilities of moving across regimes are summarized by the transition matrix:

$$\mathbf{H}^p = \left[\begin{array}{cc|cc} p_{11} & 0 & 0 & p_{14} \\ 0 & p_{22} & p_{23} & p_{24} \\ \hline 0 & p_{32} & p_{33} & 0 \\ p_{41} & p_{42} & 0 & p_{44} \end{array} \right] \quad (\text{A52})$$

in which the probability of switching to regime j given that we are in regime i is denoted by p_{ij} . Since $\xi_t^p = 1$ is the perceived short-lasting Hawkish regime, while $\xi_t^p = 2$ is the perceived long-lasting Hawkish regime, it must be that $p_{22} > p_{11}$. Analogously, since $\xi_t^p = 3$ is the perceived short-lasting Dovish regime, while $\xi_t^p = 4$ is the perceived long-lasting Dovish regime we have $p_{44} > p_{33}$. We set $p_{44} = p_{22} = 0.999$ to capture the idea that, as agents spend more and more time in a regime, they become convinced that this regime will persist indefinitely.³⁶

Suppose that the economy is initially in a state where the agent's posterior probability that she is in the long-lasting Hawkish regime $\xi_t^p = 2$ is unity. If policymakers then start conducting Dovish monetary policy, we further assume that agents will initially believe that this likely represents just a temporary deviation from $\xi_t^p = 2$ regime. This idea is captured by the conditions $p_{23} > p_{24}$, $p_{32} > 0$, $p_{31} = 0$. That is, the probability that she has switched from long-lasting Hawkish to short-lasting Dovish is greater than the probability of switching

³⁶We 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.

from long-lasting Hawkish to long-lasting Dovish, and given that she is in short-lasting Dovish, she can only switch back to long-lasting Hawkish. However, because $p_{44} > p_{33}$, if policymakers remain in the Dovish regime for long enough, the agents perceived posterior probability that they are in a long-lasting Dovish regime goes to unity. There are symmetric restrictions in the second block, corresponding to $p_{41} > p_{42}$, $p_{14} > 0$, $p_{13} = 0$. 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 distorted beliefs component of the baseline model implies that $p_{22} > p_{HH}$ and $p_{44} > p_{DD}$, where recall that the latter transition probabilities p_{HH} equals the true probability of remaining in a Hawkish regime, and p_{DD} equals the true probability of remaining in a Dovish regime.

More generally given arbitrary initial beliefs, the above restrictions on the perceived transition matrix \mathbf{H}^p will have implications for how beliefs evolve over time. Given the model of belief formation described below, if a regime change occurs after many periods of the same monetary policy rule, agents will be almost certain that the deviation is temporary. By contrast, if regime changes are frequent, agents will be uncertain about their nature and beliefs could change more abruptly.

To solve the model, we first need to establish how agents' beliefs about the perceived regimes evolve over time. This will allow us to characterize the evolution of beliefs about future monetary policy, i.e. beliefs about the persistence of the current monetary policy regime ξ_t . We will then define an expanded set of regimes that keep track at the same time of the policy rule in place (ξ_t) and agents' beliefs about future monetary policy (captured by the probabilities assigned by agents to the regimes ξ_t^p belonging to the same block).

We will now proceed in two steps. First, we will characterize the evolution of agents' beliefs *within* a block for given prior beliefs. This will allow us to track the evolution of beliefs as agents observe more and more periods of the same policy rule regime. Second, we will explain how agents' beliefs are pinned down once the economy moves *across* blocks. This will allow us to characterize agents' beliefs when agents observe a change in the conduct of monetary policy. All results are based on the Bayes' theorem. Finally, for each of these cases, we will describe how to recast the model with information frictions as a perfect information rational expectations model obtained by expanding the number of regimes to keep track of agents' beliefs.

Evolution of Beliefs *Within* a Block

In what follows, we will derive the law of motion of agents' beliefs conditional on being in a specific block, i.e. conditional on observing a certain policy rule. The formulas derived below will provide a recursive law of motion for agents' beliefs based on Bayes' theorem. Such recursion applies *for any starting values for agents' beliefs*. These initial values will be determined by agents' beliefs at the moment the system enters the new block, i.e. the moment agents observe a policy regime that is different from the one observed in the previous period. We will characterize these initial beliefs in the next subsection.

As we have noticed in the previous section, agents can infer the history of the blocks (i.e. the history of the policy rule in place, ξ^T). Therefore, at each point in time, agents know the number of consecutive periods spent in the current block since the last switch. Let us denote the number of consecutive realizations of Block i at time t as τ_t^i , $i \in \{1, 2\}$. To fix ideas, suppose that the system is in block 1 (hawkish monetary policy) at time t , implying that $\tau_t^1 > 0$ and $\tau_t^2 = 0$. Then, there are only two possible outcomes for the next period. The economy can spend an additional period in block 1 (hawkish monetary policy), implying that $\tau_{t+1}^1 = \tau_t^1 + 1$ and $\tau_{t+1}^2 = 0$, or it can move to block 2 (dovish monetary policy), implying $\tau_{t+1}^1 = 0$ and $\tau_{t+1}^2 = 1$. In this subsection, we restrict our attention to the first case.

Using Bayes' theorem and the fact that $prob(\xi_{t-1}^p = 2 | \tau_{t-1}^1) = 1 - prob(\xi_{t-1}^p = 1 | \tau_{t-1}^1)$, the probability of being in regime 1 given that we have observed τ_t^1 consecutive realizations of block 1, $prob(\xi_t^p = 1 | \tau_t^1)$, is given by:

$$\begin{aligned} prob(\xi_t^p = 1 | \tau_t^1) &= \frac{prob(\xi_{t-1}^p = 1 | \tau_{t-1}^1) p_{11}}{prob(\xi_{t-1}^p = 1 | \tau_{t-1}^1) p_{11} + prob(\xi_{t-1}^p = 2 | \tau_{t-1}^1) p_{22}} \\ &= \frac{prob(\xi_{t-1}^p = 1 | \tau_{t-1}^1) p_{11}}{prob(\xi_{t-1}^p = 1 | \tau_{t-1}^1) (p_{11} - p_{22}) + p_{22}} \end{aligned} \quad (\text{A53})$$

where $\tau_t^1 = \tau_{t-1}^1 + 1$ and for $\tau_t^1 > 1$. Notice that for $\tau_t^1 = 1$, $prob(\xi_t^p = 1 | \tau_t^1)$ denotes the initial beliefs that will be discussed in Subsection 5. Equation (A53) is a rational first-order difference equation that allows us to *recursively* characterize the evolution of agents' beliefs about being in Regime 1 while the system is in Block 1. As agents observe more and more periods of Block 1 (Hawkish monetary policy), the probability that they assign to the short-lasting Hawkish Regime 1 declines. Once agents have spent enough time under Hawkish monetary policy, they conclude that the probability of a short lasting regime is zero.

Similarly, the probability of being in Regime 3 given that we have observed τ_t^2 consecutive

realizations of Block 2, $prob(\xi_t = 3|\tau_t^2)$, can be analogously derived:

$$\begin{aligned} prob(\xi_t^p = 3|\tau_t^2) &= \frac{prob(\xi_{t-1}^p = 3|\tau_{t-1}^2) p_{33}}{prob(\xi_{t-1}^p = 3|\tau_{t-1}^2) p_{33} + prob(\xi_{t-1}^p = 4|\tau_{t-1}^2) p_{44}} \\ &= \frac{prob(\xi_{t-1}^p = 3|\tau_{t-1}^2) p_{33}}{prob(\xi_{t-1}^p = 3|\tau_{t-1}^2) (p_{33} - p_{44}) + p_{44}}. \end{aligned} \quad (\text{A54})$$

where $\tau_t^2 = \tau_{t-1}^2 + 1$ and for $\tau_t^2 > 1$.

The recursive equations (A53) and (A54) characterize the dynamics of agents' beliefs in both blocks for a given set of prior beliefs. Bianchi and Melosi (IER) show that under our assumptions for the transition matrix, these recursive equations converge as τ_t^1 and τ_t^2 grow. Once these parameters reach sufficiently high values, denoted τ^1 and τ^2 , there is no further significant change to the probabilities assigned to the short- and long-lasting regimes. In particular, the probability assigned to the short-lasting regimes converges toward zero. In what follows, we denote the converging probabilities for the short lasting regimes $prob(\xi_t^p = 1|\tau^1)$ and $prob(\xi_t^p = 3|\tau^2)$ as $\tilde{\lambda}_{b_1}$ and $\tilde{\lambda}_{b_2}$, respectively. The converging probabilities for the respective long-lasting regimes are then $1 - \tilde{\lambda}_{b_1}$ and $1 - \tilde{\lambda}_{b_2}$, respectively. This convergence result will be key to being able to recast the model with learning in terms of a finite dimensional set of regimes indexed with respect to agents' beliefs.

Evolution of Beliefs Across Blocks

In the previous subsection, we characterized the evolution of agents' beliefs conditional on being in a specific block, i.e. conditional on observing additional realizations of the same policy rule. The formulas derived above apply to *any set of initial beliefs*. In this subsection, we will pin down agents' beliefs at the moment the economy moves *across* blocks, i.e. for the alternative case in which the policy regime observed at time t differs from the policy regime in place at time $t - 1$. These beliefs will serve as starting points for the recursions (A53) and (A54) governing the evolution of beliefs within a block.

Suppose for a moment that before switching to the new block, agents are convinced of being in one of the two regimes of the current block (in other words, they believe that they know which ξ_t^p is in place). Notice that in this case the transition matrix conveys all the information necessary to pin down agents' starting beliefs about the regime in place within the new block. Specifically, we have that if the economy moves from block 2 (Dovish) to block 1 (Hawkish), the probability of being in regime 1 (short-lasting Hawkish) is given by

$$prob(\xi_t^p = 1|\xi_{t-1}^p = 3, \tau_t^1 = 1) = \frac{p_{31}}{p_{31} + p_{32}} = 0, \quad (\text{A55})$$

if the economy was under regime 3 (short-lasting Dove) in the previous period, or by

$$prob(\xi_t^p = 1|\xi_{t-1}^p = 4, \tau_t^1 = 1) = \frac{p_{41}}{p_{41} + p_{42}} = 1 \quad (\text{A56})$$

if the economy was under regime 4 (long-lasting Dove) in the previous period. Symmetrically, the initial probability of being in regime 3 (short-lasting Dovish) given that the economy just moved to block 2 (Hawkish monetary policy) is given by

$$prob(\xi_t^p = 3 | \xi_{t-1}^p = 1, \tau_t^2 = 1) = \frac{p_{13}}{p_{13} + p_{14}} = 0, \quad (\text{A57})$$

if the economy was under regime 1 (short-lasting Hawk) in the previous period, or by

$$prob(\xi_t^p = 3 | \xi_{t-1}^p = 2, \tau_t^2 = 1) = \frac{p_{23}}{p_{23} + p_{24}} = 1 \quad (\text{A58})$$

if the economy was previously under regime 2 (long-lasting Hawk).

However, in the model, agents are generally not sure about the nature (short-lasting v.s. long-lasting) of the observed monetary policy regime that is in place. Therefore, their beliefs at the moment the economy moves from one block to the other will be a weighted average of the probabilities outlined above. The weights in general, in turn, will depend on agents' beliefs right before the moment of the switch. Specifically, agents' starting beliefs in a new block 1 upon the shift from block 2 are given by

$$prob\{\xi_t^p = 1 | \mathcal{I}_t\} = \frac{(1 - prob\{\xi_{t-1}^p = 3 | \mathcal{I}_{t-1}\}) p_{41}}{prob\{\xi_{t-1}^p = 3 | \mathcal{I}_{t-1}\} p_{32} + (1 - prob\{\xi_{t-1}^p = 3 | \mathcal{I}_{t-1}\}) (p_{41} + p_{42})} \quad (\text{A59})$$

while if the system just entered block 2, starting beliefs read

$$prob\{\xi_t^p = 3 | \mathcal{I}_t\} = \frac{(1 - prob\{\xi_{t-1}^p = 1 | \mathcal{I}_{t-1}\}) p_{23}}{prob\{\xi_{t-1}^p = 1 | \mathcal{I}_{t-1}\} p_{14} + (1 - prob\{\xi_{t-1}^p = 1 | \mathcal{I}_{t-1}\}) (p_{23} + p_{24})} \quad (\text{A60})$$

where \mathcal{I}_t includes the history of policy regimes (blocks) up to time t . Because the above are recursive formulations, we observe that the only information in \mathcal{I}_t that is relevant for knowing the starting beliefs upon switching to a new block is the agent's beliefs last period and the perceived transition matrix \mathbf{H}^p .

To summarize, taking together movements within and across blocks, two variables pin down the dynamics of beliefs over time: how many consecutive periods the system has spent in the current block *and* the initial beliefs agents had when the system entered the current block.

Tracking beliefs

To solve the model under learning, we need to keep track of the evolution of beliefs. An approximation is required, since beliefs are continuous variables with an infinite number of possible values. To keep the problem tractable, we map beliefs into a grid of possible values. As the number grid points approaches infinity, the approximation becomes arbitrarily accurate. Note that for each point in the grid (A59) and (A60) tell us how beliefs will evolve if we observe

a change in the conduct of monetary policy, while (A53) and (A54) tell us how beliefs will evolve if an additional period of the same policy regime is observed. In other words, these two pairs of equations tell us how beliefs evolve across every possible scenario. This allows us to compute the probability of moving to any point in the grid from any other point, and can be represented by an expanded transition matrix that keeps track at the same time of the evolution of policymakers' behavior and agents' beliefs. Once we have the expanded transition matrix, we can combine it with the model equations to solve the model. When agents form expectations, the expanded transition matrix will determine the evolution of their beliefs about future monetary policy. Importantly, *agents know that they do not know*: They understand that their beliefs will change based on what they will observe in the future. In what follows, we provide the details.

Denote the grid for beliefs $prob \{\xi_t^p = 1 | \mathcal{I}_t\}$ as $\mathcal{G}_{b_1} = \{\mathcal{G}_1, \dots, \mathcal{G}_{g_1}\}$ and for beliefs $prob \{\xi_t^p = 3 | \mathcal{I}_t\}$ as $\mathcal{G}_{b_2} = \{\mathcal{G}_{g_1+1}, \dots, \mathcal{G}_{g_1+g_2}\}$ where $0 \leq \mathcal{G}_i \leq 1$, all $1 \leq i \leq g = g_1 + g_2$. Furthermore, we denote the whole grid as $\mathcal{G} = \mathcal{G}_{b_1} \cup \mathcal{G}_{b_2}$. Endowed with such a grid, we can keep track of agents' beliefs and policymakers' behavior by introducing a new set of regimes ζ_t^p . The new regime ζ_t^p characterizes the policy regime in place and the knot of the grid \mathcal{G} that best approximates agents' beliefs; that is, in our notation $prob \{\xi_t^p = 1 | \mathcal{I}_t\}$ when the system is in block 1 and $prob \{\xi_t^p = 3 | \mathcal{I}_t\}$ when the system is in block 2. Thus, each regime ζ_t^p is associated with a pair $\{\xi_t = 1, prob \{\xi_t^p = 1\} = \mathcal{G}_{b_1}\}$ or $\{\xi_t = 2, prob \{\xi_t^p = 3\} = \mathcal{G}_{b_2}\}$. For example, the regime $\zeta_t^p = g_1 + i$ is associated with the pair $\{\xi_t = 2, prob \{\xi_t^p = 3\} = \mathcal{G}_{g_1+i}\}$ and corresponds to monetary policy being Dovish ($\xi_t = 2$) and agents thinking that the probability of being in the short lasting Dovish regime is \mathcal{G}_{g_1+i} .

The expanded transition probability matrix for these new regimes can be pinned down using the recursions (A53) and (A54) and the initial conditions (A59) and (A60). Denote this expanded perceived transition matrix $\hat{\mathbf{H}}^p$. The algorithm below illustrates how exactly to perform this task.

Algorithm Initialize the transition matrix $\hat{\mathbf{H}}^p$ for the new regimes ζ_t^p , setting $\hat{\mathbf{H}}^p = \mathbf{0}_{g \times g}$.

Step 1 For each of the two blocks, do the following steps (without loss of generality we describe the steps for Block 1):

Step 1.1 For any grid point $\mathcal{G}_i \in \mathcal{G}_{b_1}$, $1 \leq i \leq g_1$, compute

$$\hat{\mathbf{H}}^p(i, j) = prob \{\xi_{t-1}^p = 1 | \mathcal{I}_{t-1}\} p_{11} + (1 - prob \{\xi_{t-1}^p = 1 | \mathcal{I}_{t-1}\}) p_{22}$$

where $prob \{\xi_{t-1}^p = 1 | \mathcal{I}_{t-1}\} = \mathcal{G}_i$ and $j \leq g_1$ is set so as to $\min |prob \{\xi_t^p = 1 | \mathcal{I}_t\} - \mathcal{G}_j|$, where $prob \{\xi_t^p = 1 | \mathcal{I}_t\}$ is computed using the recursive equation (A53) by approximating $prob \{\xi_{t-1}^p = 1 | \mathcal{I}_{t-1}\} = \mathcal{G}_i$. To ensure the convergence of beliefs, we correct

j as follows: if $j = i$ and $\mathcal{G}_i \neq \tilde{\lambda}_{b_1}$, then set $j = \min(j + 1, g_1)$ if $\mathcal{G}_i < \tilde{\lambda}_{b_1}$ or $j = \max(1, j - 1)$ if $\mathcal{G}_i > \tilde{\lambda}_{b_1}$.

Step 1.2 For any grid point $\mathcal{G}_i \in \mathcal{G}_{b_1}$, $1 \leq i \leq g_1$, compute $\hat{\mathbf{H}}^p(i, l) = 1 - \hat{\mathbf{H}}^p(i, j)$ with $l > g_1$ satisfying

$$\min \left| \frac{(1 - \text{prob}\{\xi_{t-1}^p = 1 | \mathcal{I}_{t-1}\}) p_{23}}{\text{prob}\{\xi_{t-1}^p = 1 | \mathcal{I}_{t-1}\} p_{14} + (1 - \text{prob}\{\xi_{t-1}^p = 1 | \mathcal{I}_{t-1}\}) (p_{23} + p_{24})} - \mathcal{G}_l \right|$$

where $\text{prob}\{\xi_{t-1}^p = 1 | \mathcal{I}_{t-1}\} = \mathcal{G}_i$.

Step 2 If no column of $\hat{\mathbf{H}}^p$ has all zero elements, stop. Otherwise, go to Step 3.

Step 3 Construct the matrix T as follows. Set $j = 1$ and $l = 1$. While $j \leq g$, if $\sum_{i=1}^g \hat{\mathbf{H}}^p(i, j) = 0$ set $j = j + 1$. Otherwise, if $\sum_{i=1}^g \hat{\mathbf{H}}^p(i, j) \neq 0$: (1) set $T(j, l) = 1$, (2) set $T(j, v) = 0$ for any $1 \leq v \leq g$ and $v \neq l$, (3) set $l = l + 1$ and $j = j + 1$.

Step 4 Write the transition equation as $\tilde{H}^p = T \cdot \hat{\mathbf{H}}^p \cdot T'$. If no column of \tilde{H}^p has all zero elements, set $\hat{\mathbf{H}}^p = \tilde{H}^p$ and stop. Otherwise, go to step 3.

Step 1.1 determines the regime j the system will go to if it stays in block 1 next period and fills up the appropriate element (i, j) of the transition matrix $\hat{\mathbf{H}}^p$ with the probability of moving to Regime j . Step 1.2 computes the regime l the system will go to if it leaves block 1 and fills up the appropriate element (i, l) of matrix \tilde{H}^p . Steps 2-4 are not necessary but help to keep the dimension of the grid small, getting rid of regimes that will never be reached. For computational convenience, we always add the convergence points for the two blocks (i.e., $\tilde{\lambda}_{b_1}$ in the case of block 1) to the grid \mathcal{G} . On many occasions, it is a good idea to make the grid near the convergence knot very fine to improve the precision of the approximation.

At the end of this algorithm we end up with a transition matrix for the expanded regime space with elements taking the form

$$\hat{H}_{ij} = \Pr(\zeta_{t+1} = j | \zeta_t = i). \quad (\text{A61})$$

Solving the Dynamic Macro-Finance Model

The model can be solved in two steps. First, we solve for the macro dynamics. This returns a MS-VAR in the state vector S_t defined above. Then, conditional on this solution and the probability assigned by the asset pricing agent to moving across perceived regimes as captured by the expanded $(g \times g)$ transition matrix $\hat{\mathbf{H}}^p$, we can solve for the evolution of asset prices.

In equations, the first step returns a MS-VAR in the macro state vector:

$$S_t = \mathbf{C}_{\xi_t} + \mathbf{T}_{\xi_t} S_{t-1} + \mathbf{R}_{\xi_t} \mathbf{Q} \varepsilon_t.$$

The second step takes this regime specific law of motion for the macroeconomy as an input and combines it with the equilibrium asset pricing relations, conditional on the law of motion for agents' beliefs as captured by the transition matrix $\hat{\mathbf{H}}^p$. All variables that enter the asset pricing system of equations are linear transformation of the variables entering the macro block for which we have a solution. For example, the log SDF can be expressed as a function of the macro state vector S_t : $m_t = e_m S_t$, where e_m is a vector that extracts the desired linear combination of the variables contained in S_t . We have:

$$m_t = \log(\delta) - \sigma_p e_c (S_t - S_{t-1}) + \vartheta_{p,t-1} \quad (\text{A62})$$

$$e_i S_t - \mathbb{E}_t^p [e_\pi S_{t+1}] = -\mathbb{E}_t^p [m_{t+1}] \underbrace{- .5 \mathbb{V}_t^p [m_{t+1} + e_i S_t - e_\pi S_{t+1}]}_{\text{Regime dependent risk adjustment}} - lp \quad (\text{A63})$$

$$pd_t = \kappa_0 + \mu + \underbrace{[.5 \mathbb{V}_t^p [m_{t+1} + e_c (S_{t+1} - S_t) + \kappa_1 pd_{t+1}]]}_{\text{Regime dependent risk adjustment}} \quad (\text{A64})$$

$$+ \mathbb{E}_t^p [m_{t+1} + e_c (S_{t+1} - S_t) + \kappa_1 pd_{t+1}]$$

$$\eta_t^{pd} = pd_t - \mathbb{E}_{t-1}^p (pd_t) \quad (\text{A65})$$

$$\eta_t^m = m_t - \mathbb{E}_{t-1}^p (m_t) \quad (\text{A66})$$

$$S_t = \mathbf{C}_{\xi_t} + \mathbf{T}_{\xi_t} S_{t-1} + \mathbf{R}_{\xi_{t+1}} \mathbf{Q} \varepsilon_t \quad (\text{A67})$$

where e_x is a vector that extracts the desired linear combination of the variables contained in S_t : $x_t = e_x S_t$.

Notice that the solution of the macro block implies heteroskedasticity for the endogenous variables, the Markov-switching coefficients in the equation for S_t . To keep the framework one that is conditionally lognormal with a risk adjustment, we follow Bansal and Zhou (2002) and compute the one-step-ahead conditional variance as the weighted average of the conditional variances across regimes resulting from the Gaussian shocks. This implies that lognormality is assumed, conditional on ξ_{t+1} . (The section ‘‘Solving the model with a risk adjustment,’’ below provides details on lognormality in a setting with regime shifts.) Define the augmented state space as $\tilde{S}_t = [S_t, m_t, pd_t, \mathbb{E}_t^p (m_{t+1}), \mathbb{E}_t^p (pd_{t+1})]$. For any variable x_t in the asset pricing block, conditional lognormality assumption implies:

$$\mathbb{V}_t^p [x_{t+1}] \approx \mathbb{E}_t^p (\mathbb{V}_t [x_{t+1} | \xi_{t+1}]) = e_x \mathbb{E}_t^p [\mathbf{R}_{\xi_{t+1}} \mathbf{Q} \mathbf{Q}' \mathbf{R}'_{\xi_{t+1}}] e_x \quad (\text{A68})$$

where e_x is a vector used to extract the desired linear combination of the variables in S_t . This assumption maintains conditional log-normality of the entire system and guarantees the algorithm above converges in one step. Notice that $\mathbb{V}_t [\cdot | \xi_{t+1}]$ without a ‘‘p’’ superscript is the conditional variance under the objective measure given the specification of the lognormal shocks in the model.

The second step consists of expanding the number of regimes to reflect the evolution of beliefs. To do so, we recast the model in terms of the new set of regimes ζ_t that keep track

both of the behavior of monetary authority (as captured by ξ_t) and agents' beliefs about the nature of these regime changes (i.e. beliefs about ξ_t^p). Furthermore, given the approximation (A68), the one-step-ahead variance $\mathbb{V}_t^p[x_{t+1}]$ is only a function of the expanded regimes at time t , ζ_t . This leaves us with a new system to be solved, given by:

$$\Gamma_0(\zeta_t)\tilde{S}_t = \Gamma_c(\zeta_t) + \Gamma_1(\zeta_t)\tilde{S}_{t-1} + \Psi(\zeta_t)\mathbf{Q}\varepsilon_t + \Pi\boldsymbol{\eta}_t \quad (\text{A69})$$

where the regime $\zeta_t \in \{1, \dots, g_1 + g_2\}$ follows the transition matrix $\hat{\mathbf{H}}^p$ and the terms $\Gamma_c(\zeta_t)$ now also contain the regime-specific risk adjustment terms $\mathbb{V}_t^p[x_{t+1}]$ that are part of the asset pricing block. Note that $\Gamma_c(\zeta_t)$ depends on $\mathbb{V}_t^p[x_{t+1}]$ as given in (A68). For variables in the system (A62)-(A67) expressed in recursive form, like pd_t , the vector e_x is not known until we solve for \tilde{S}_t . We therefore employ an iterative procedure. First, we guess a value for e_x . We can then use solution methods available for dynamic macro models with Markov-switching random variables. The resulting solution takes the form once again as a MS-VAR:

$$\tilde{S}_t = \tilde{\mathbf{C}}(\zeta_t, \hat{\mathbf{H}}^p) + \tilde{\mathbf{T}}(\zeta_t, \hat{\mathbf{H}}^p)S_{t-1} + \tilde{\mathbf{R}}(\zeta_t, \hat{\mathbf{H}}^p)\mathbf{Q}\varepsilon_t \quad (\text{A70})$$

We use the solution to update e_x , then solve the model again. The iteration converges in one step due to linear system and the fact that the risk corrections only affect $\Gamma_c(\zeta_t)$. The desired observables can then be reconstructed starting from the augmented state vector.

Armed with \tilde{S}_t , any vector of endogenous variables Y_t in the model has a solution taking the form

$$Y_t = D + Z\tilde{S}_t,$$

where D is a constant vector and Z is a constant matrix.

Let the model solution for the price-dividend ratio be denoted $pd_t = pd(\tilde{S}_t)$, where $pd(\cdot)$ is a linear transformation. The solution satisfies the recursion below. To see how agents' beliefs matter for asset prices, consider the recursive formulation for the price-dividend ratio:

$$pd(\tilde{S}_t) = \kappa_0 + \mu + \underbrace{.5\mathbb{V}_t^p[m_{t+1} + e_c(S_{t+1} - S_t) + \kappa_1 pd_{t+1}]}_{\text{Regime dependent risk adjustment}} + \mathbb{E}_t^p[m_{t+1} + e_c(S_{t+1} - S_t) + \kappa_1 pd(\tilde{S}_{t+1})]$$

As explained above, the regime dependent risk adjustment only depends on the regime in place at time t . Then:

$$pd(\tilde{S}_t) = \kappa_0 + \mu + \underbrace{\left[.5e_x \mathbb{E}_t^p \left[\tilde{\mathbf{R}}(\zeta_t, \hat{\mathbf{H}}^p) \mathbf{Q} \mathbf{Q}' \tilde{\mathbf{R}}(\zeta_t, \hat{\mathbf{H}}^p)' \right] e_x \right]}_{\text{Regime dependent risk adjustment}} + \mathbb{E}_t^p[m_{t+1} + e_c(S_{t+1} - S_t) + \kappa_1 pd(\tilde{S}_{t+1})]$$

where we use e_x to denote a vector that extracts the desired linear combination from the one

step ahead covariance matrix. We then have:

$$pd\left(\tilde{S}_t\right) = \kappa_0 + \mu + \mathbb{E}_t^p \left[.5e_x \tilde{\mathbf{R}}_{\zeta_{t+1}} \mathbf{Q}\mathbf{Q}\tilde{\mathbf{R}}'_{\zeta_{t+1}} e_x + m_{t+1} + e_c (S_{t+1} - S_t) + \kappa_1 pd\left(\tilde{S}_{t+1}\right) \right]$$

$$pd\left(\tilde{S}_t\right) = \kappa_0 + \mu + \sum_{j=1}^{g_1+g_2} P \left\{ \zeta_{t+1} = j | \zeta_t = i \right\} \mathbb{E}_t^p \left[\begin{array}{l} .5e_x \tilde{\mathbf{R}}_{\zeta_{t+1}} \mathbf{Q}\mathbf{Q}\tilde{\mathbf{R}}'_{\zeta_{t+1}} e_x + m_{t+1} \\ + e_c (S_{t+1} - S_t) + \kappa_1 pd\left(\tilde{S}_{t+1}\right) \end{array} \right]$$

where we have used the output in (A61) from the algorithm discussed above to obtain the $P \left\{ \zeta_{t+1} = j | \zeta_t = i \right\}$ that are elements of the expanded transition matrix $\hat{\mathbf{H}}^p$.

Solving a model with risk adjustment

This section provides more details about solving the model with a risk adjustment. As explained in the main text, our approach is quite common in the asset pricing and macro-finance literatures. This appendix provides the following points:

1. The method can be characterized as a guess-and-verify approach. This is because once the model is log-linearized and solved, with or without a risk-adjustment, the variables of the model follow a linear process in logs and are therefore log-normal in levels. The method exploits this property of the solution when log-linearizing the model and implements a risk-adjusted log-linearization. This affects only the equilibrium conditions in which an expectational term appears. Note that log-normality does not affect the rest of the log-linearized equations. When introducing regime changes, the process becomes conditionally log-normal, conditional on the regimes.
2. To understand why the solution without risk adjustment already implies lognormality, it is important to notice that all shocks are specified as shocks to log variables. Thus, when taking a lognormal approximation, the solution of the model implies a linear process in logs with Gaussian innovations.
3. The solution with risk-adjustment allows us to take into account the effects of risk on asset prices.

Conditional log-normality

Suppose that a variable Z_{t+1} has a log-normal distribution such that $z_{t+1} = \log(Z_{t+1})$ follows the process:

$$z_{t+1} = c + az_t + \sigma\varepsilon_{t+1}$$

Then:

$$\ln(\mathbb{E}_t[Z_{t+1}]) = \mathbb{E}_t[z_{t+1}] + .5\mathbb{V}_t[z_{t+1}] = c + az_t + .5\sigma^2 \tag{A71}$$

Now, suppose $z_{t+1} = \log(Z_{t+1})$ follows a Markov-switching process:

$$z_{t+1} = c_{\xi_{t+1}} + a_{\xi_{t+1}} z_t + \sigma_{\xi_{t+1}} \varepsilon_{t+1} \quad (\text{A72})$$

where ξ_{t+1} denotes the regime at time $t + 1$. The solution of the model, presented in the main text, has this form. When we log-linearize the system of model equations, we are facing log-linearization equations of the following form:

$$\mathbb{E}_t[e^{z_{t+1}}]. \quad (\text{A73})$$

We extend the approach in Bansal and Zhou (2002), who utilize conditional log-normality of the process in equation (A72). Conditioning on the regime in the next period, log-normality holds:

$$\begin{aligned} \mathbb{E}_t[e^{z_{t+1}}|\xi_{t+1}] &= e^{\mathbb{E}_t[z_{t+1}|\xi_{t+1}] + 0.5\mathbb{V}_t[z_{t+1}|\xi_{t+1}]}, \\ \ln(\mathbb{E}_t[e^{z_{t+1}}|\xi_{t+1}]) &= \mathbb{E}_t[z_{t+1}|\xi_{t+1}] + 0.5\mathbb{V}_t[z_{t+1}|\xi_{t+1}]. \end{aligned}$$

Therefore, using the law of iterated expectations:

$$\begin{aligned} \mathbb{E}_t[e^{z_{t+1}}] &= \mathbb{E}_t[\mathbb{E}_t[e^{z_{t+1}}|\xi_{t+1}]] = \mathbb{E}_t[e^{\mathbb{E}_t[z_{t+1}|\xi_{t+1}] + 0.5\mathbb{V}_t[z_{t+1}|\xi_{t+1}]}] = \\ &= \mathbb{E}_t[e^{c_{\xi_{t+1}} + a_{\xi_{t+1}} z_t + \sigma_{\xi_{t+1}} \varepsilon_{t+1}}]. \end{aligned}$$

To proceed, we follow Bansal and Zhou (2002) and use the approximation: $e^{c_{\xi_{t+1}} + a_{\xi_{t+1}} z_t + 0.5\sigma_{\xi_{t+1}}^2} \approx 1 + c_{\xi_{t+1}} + a_{\xi_{t+1}} z_t + 0.5\sigma_{\xi_{t+1}}^2$. With this approximation, we have

$$\mathbb{E}_t[e^{z_{t+1}}] = \mathbb{E}_t[\mathbb{E}_t[e^{z_{t+1}}|\xi_{t+1}]] \approx \mathbb{E}_t[1 + c_{\xi_{t+1}} + a_{\xi_{t+1}} z_t + 0.5\sigma_{\xi_{t+1}}^2] = \quad (\text{A74})$$

$$= 1 + \mathbb{E}_t[c_{\xi_{t+1}} + a_{\xi_{t+1}} z_t] + 0.5\mathbb{E}_t[\sigma_{\xi_{t+1}}^2] \quad (\text{A75})$$

Thus, we obtain:

$$\ln(\mathbb{E}_t[Z_{t+1}]) \approx \mathbb{E}_t[c_{\xi_{t+1}} + a_{\xi_{t+1}} z_t] + 0.5\mathbb{E}_t[\sigma_{\xi_{t+1}}^2] \quad (\text{A76})$$

again using the approximation $\ln(1 + x) \approx x$, for x small.

Above we have made use of the fact that $z_{t+1} = c_{\xi_{t+1}} + a_{\xi_{t+1}} z_t + \sigma_{\xi_{t+1}} \varepsilon_{t+1}$ is close to zero. But the solution is always approximating around the steady state values. The same approximation holds even if z_{t+1} is not close zero. To see this, suppose z is the steady state of z_{t+1} and $\tilde{z}_{t+1} \equiv z_{t+1} - z$ is the log-deviation of Z_{t+1} from its mean. Then, we have:

$$e^z \mathbb{E}_t[e^{\tilde{z}_{t+1}}] = e^z \mathbb{E}_t[\mathbb{E}_t[e^{\tilde{z}_{t+1}}|\xi_{t+1}]] \quad (\text{A77})$$

$$= e^z \mathbb{E}_t[e^{\mathbb{E}_t[\tilde{z}_{t+1}|\xi_{t+1}] + 0.5\mathbb{V}_t[\tilde{z}_{t+1}|\xi_{t+1}]}] \quad (\text{A78})$$

$$\approx e^z \mathbb{E}_t[1 + \mathbb{E}_t[\tilde{z}_{t+1}|\xi_{t+1}] + 0.5\mathbb{V}_t[\tilde{z}_{t+1}|\xi_{t+1}]] \quad (\text{A79})$$

$$= e^z [1 + \mathbb{E}_t[\tilde{z}_{t+1}] + 0.5\mathbb{E}_t[\sigma_{\xi_{t+1}}^2]] \quad (\text{A80})$$

where we have used the fact that $\mathbb{V}_t[\tilde{z}_{t+1}|\xi_{t+1}] = \mathbb{V}_t[z_{t+1}|\xi_{t+1}]$. Then:

$$\log(\mathbb{E}_t[Z_{t+1}]) = \log(e^z \mathbb{E}_t[e^{\tilde{z}_{t+1}}]) \quad (\text{A81})$$

$$\approx z + \mathbb{E}_t[\tilde{z}_{t+1}] + 0.5\mathbb{E}_t[\sigma_{\xi_{t+1}}^2] \quad (\text{A82})$$

$$= \mathbb{E}_t[z_{t+1}] + 0.5\mathbb{E}_t[\sigma_{\xi_{t+1}}^2] \quad (\text{A83})$$

$$= \mathbb{E}_t[c_{\xi_{t+1}} + a_{\xi_{t+1}}z_t] + 0.5\mathbb{E}_t[\sigma_{\xi_{t+1}}^2]. \quad (\text{A84})$$

To see how the method works in our model, note that the above approximations hold both under the objective probability distribution in the model as well as under the distorted beliefs $\mathbb{E}_t^p[\cdot]$, since in both cases the random variables are conditionally lognormal. Consider the forward looking relation for the price-payout ratio:

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

Taking logs on both sides, we get:

$$pd_t = \log [\mathbb{E}_t^p [\exp (m_{t+1} + \Delta d_{t+1} + \kappa_0 + \kappa_1 pd_{t+1})]].$$

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

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

where under the conditional lognormality approximation we have:

$$\mathbb{V}_t^p [m_{t+1} + \Delta d_{t+1} + \kappa_1 pd_{t+1}] \approx \mathbb{E}_t^p [\mathbb{V}_t [m_{t+1} + \Delta d_{t+1} + \kappa_1 pd_{t+1}|\xi_{t+1}]].$$

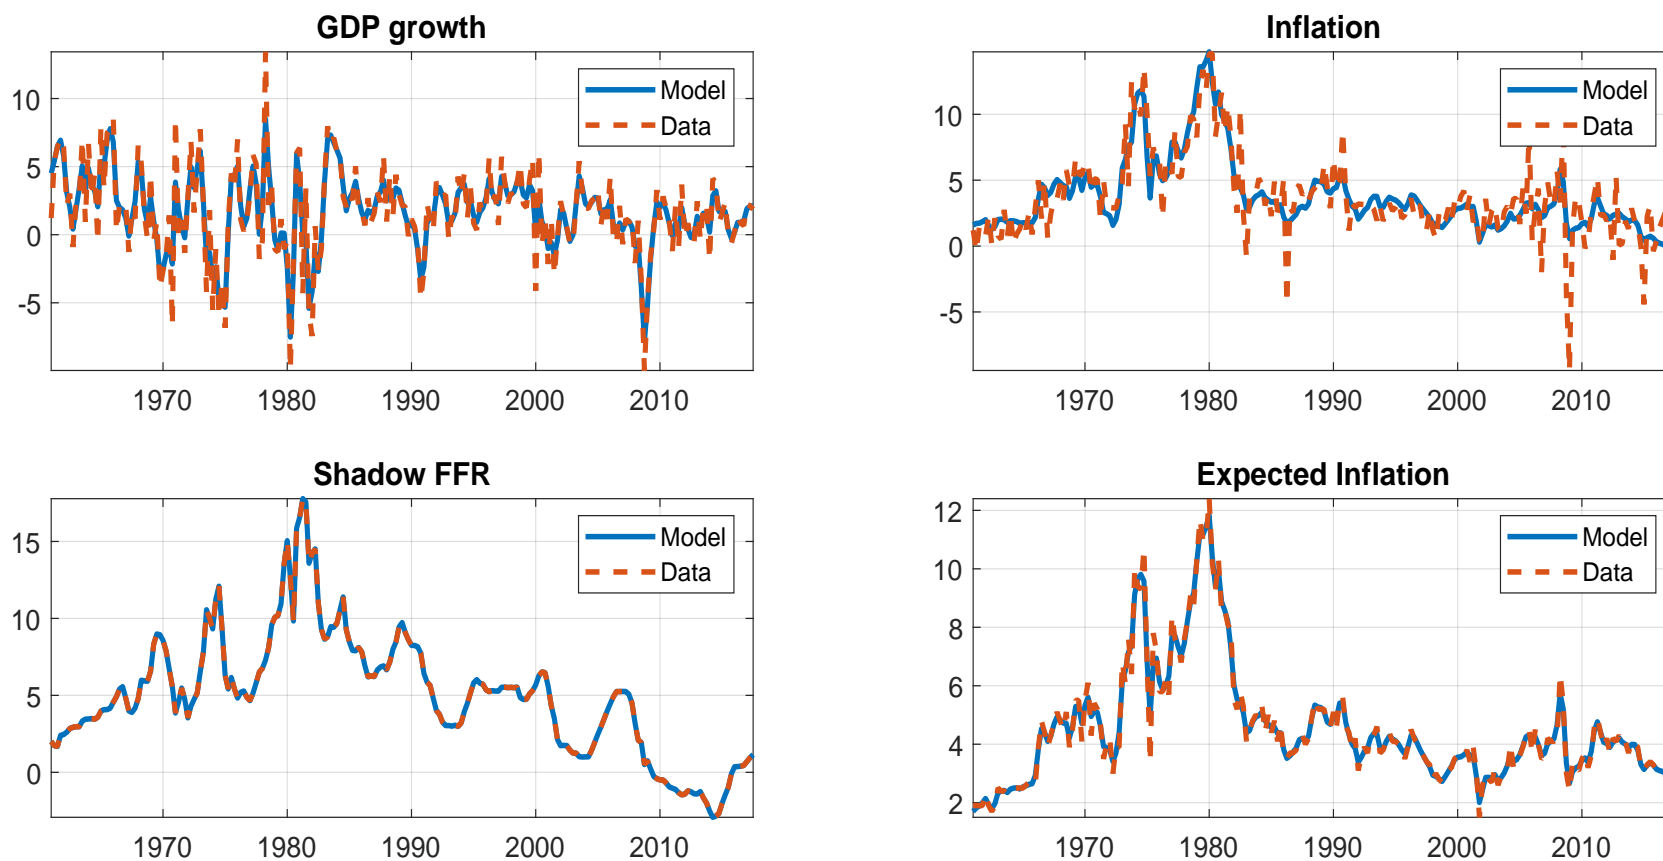


Figure A.1: Shadow rate estimates of the macro-finance model. The figure reports the model implied series and the corresponding observed series. Expected inflation comes from the Michigan Survey of Consumers. The difference is due to observation errors. The sample spans 1961:Q1 - 2017:Q3. Results are based on estimates obtained replacing the FFR with the Wu-Xia shadow rate when the ZLB is binding.




Figure A.2: Shadow rate estimates of the macro-finance model. The blue line corresponds to the fluctuations generated by changes in both the target and the slope coefficients. The orange line assumes that monetary policy starts under the Dovish regime and no regime change occurs. Finally, the black dotted line assumes that changes in the target occurred, but that the slope coefficients in the Taylor rule coefficients always remain as in the Dovish-high target regime. Results are based on estimates obtained replacing the FFR with the Wu-Xia shadow rate when the ZLB is binding.

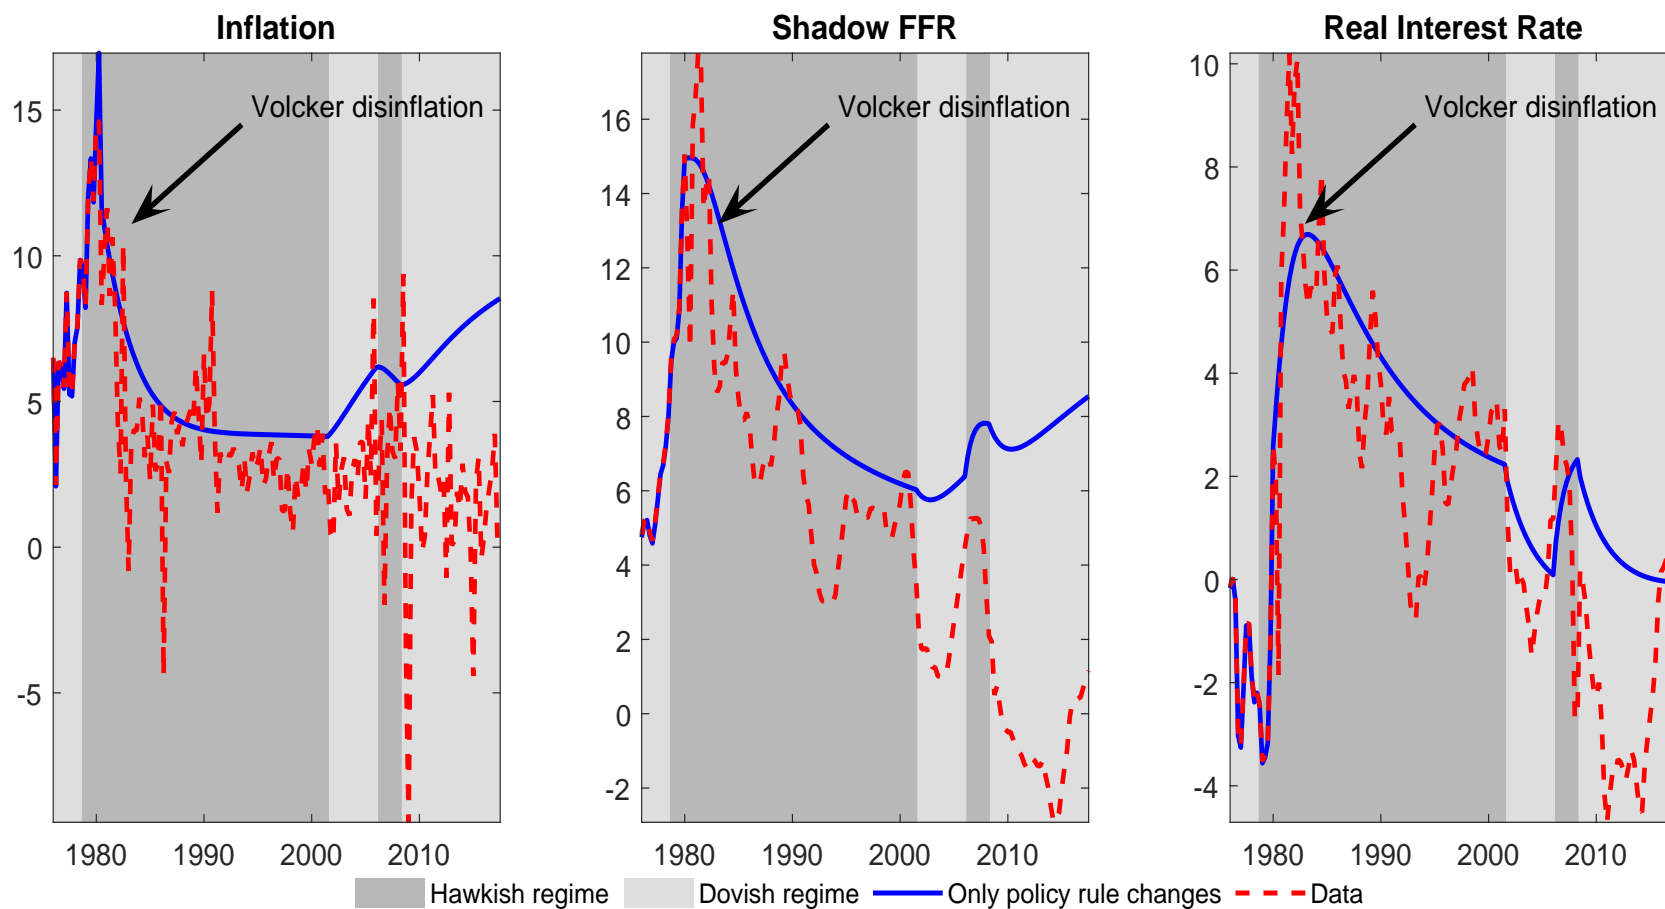


Figure A.3: Shadow rate estimates of the macro-finance model. The Volcker disinflation. We start the economy as it was in 1980:Q1 and remove all Gaussian shocks that occurred after that period, but keep the estimated regime sequence. The dashed line corresponds to the data. The real interest rate is computed as the difference between the FFR and expected inflation. Expected inflation is obtained based on the model solution. Results are based on estimates obtained replacing the FFR with the Wu-Xia shadow rate when the ZLB is binding.

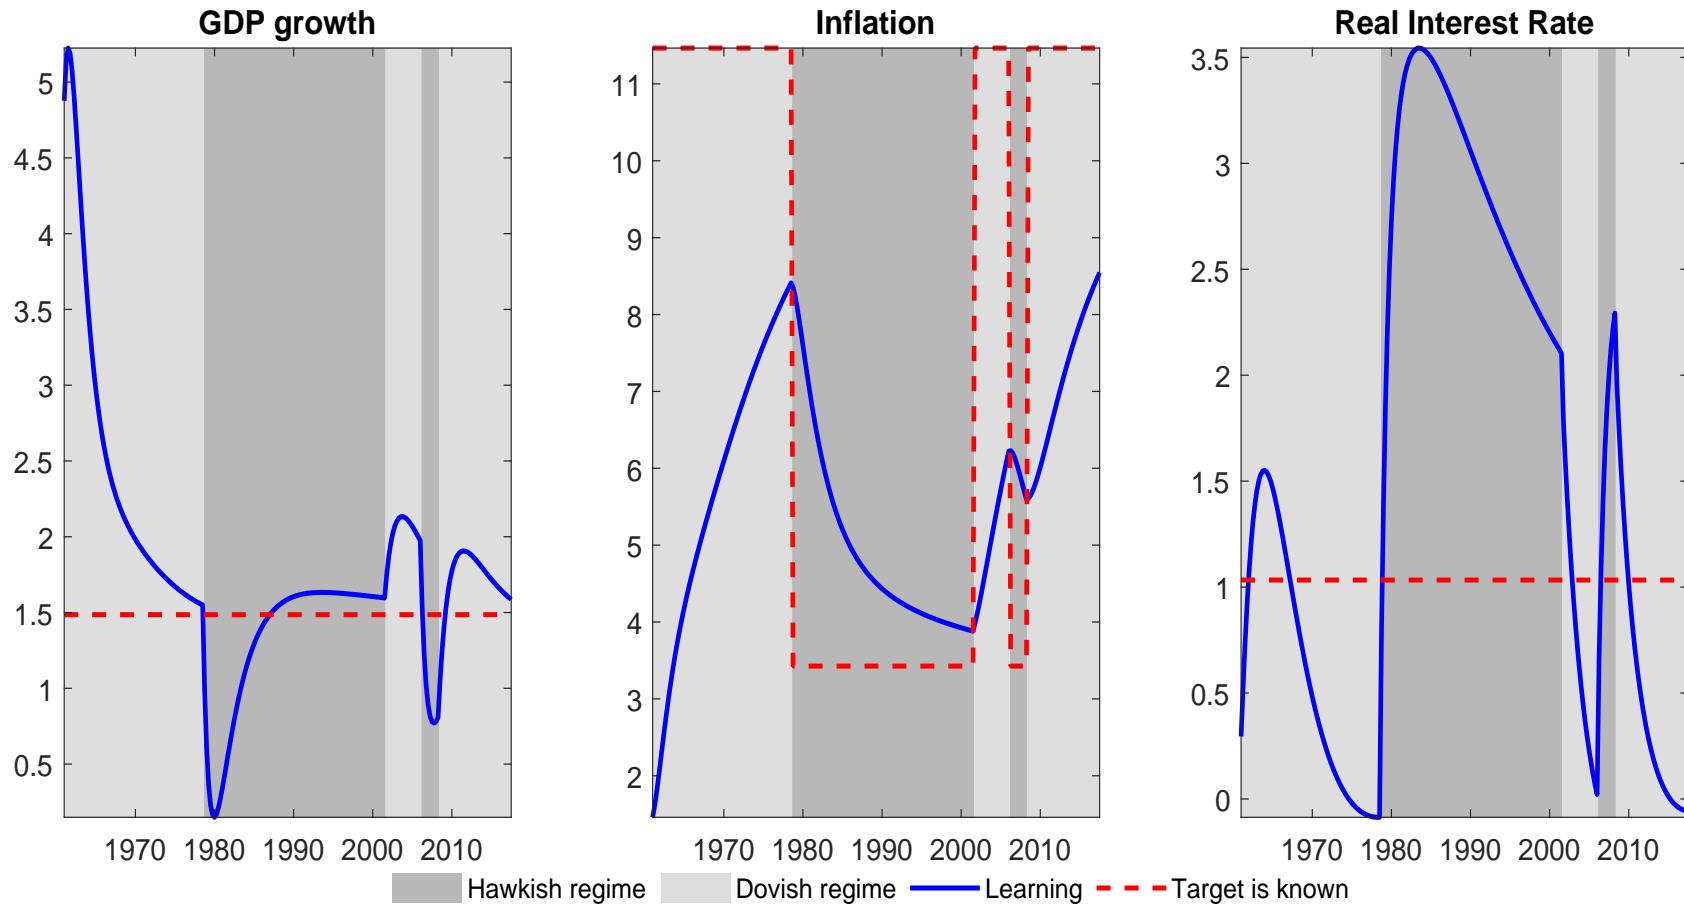


Figure A.4: Shadow rate estimates of the macro-finance model. Perfect information about the target. The blue solid line shows estimated fluctuations generated only by changes in the policy rule (inflation target and slope coefficients) when agents learn about trend inflation. The orange dashed line shows a counterfactual in which the policy rule shifts but agents observe the inflation target. Dovish regime has a high target π^T and low activism against deviations from the target π^T . Hawkish regime has a low π^T and high activism against deviations from π^T . The sample spans 1961:Q1 - 2017:Q3. Results are based on estimates obtained replacing the FFR with the Wu-Xia shadow rate when the ZLB is binding.

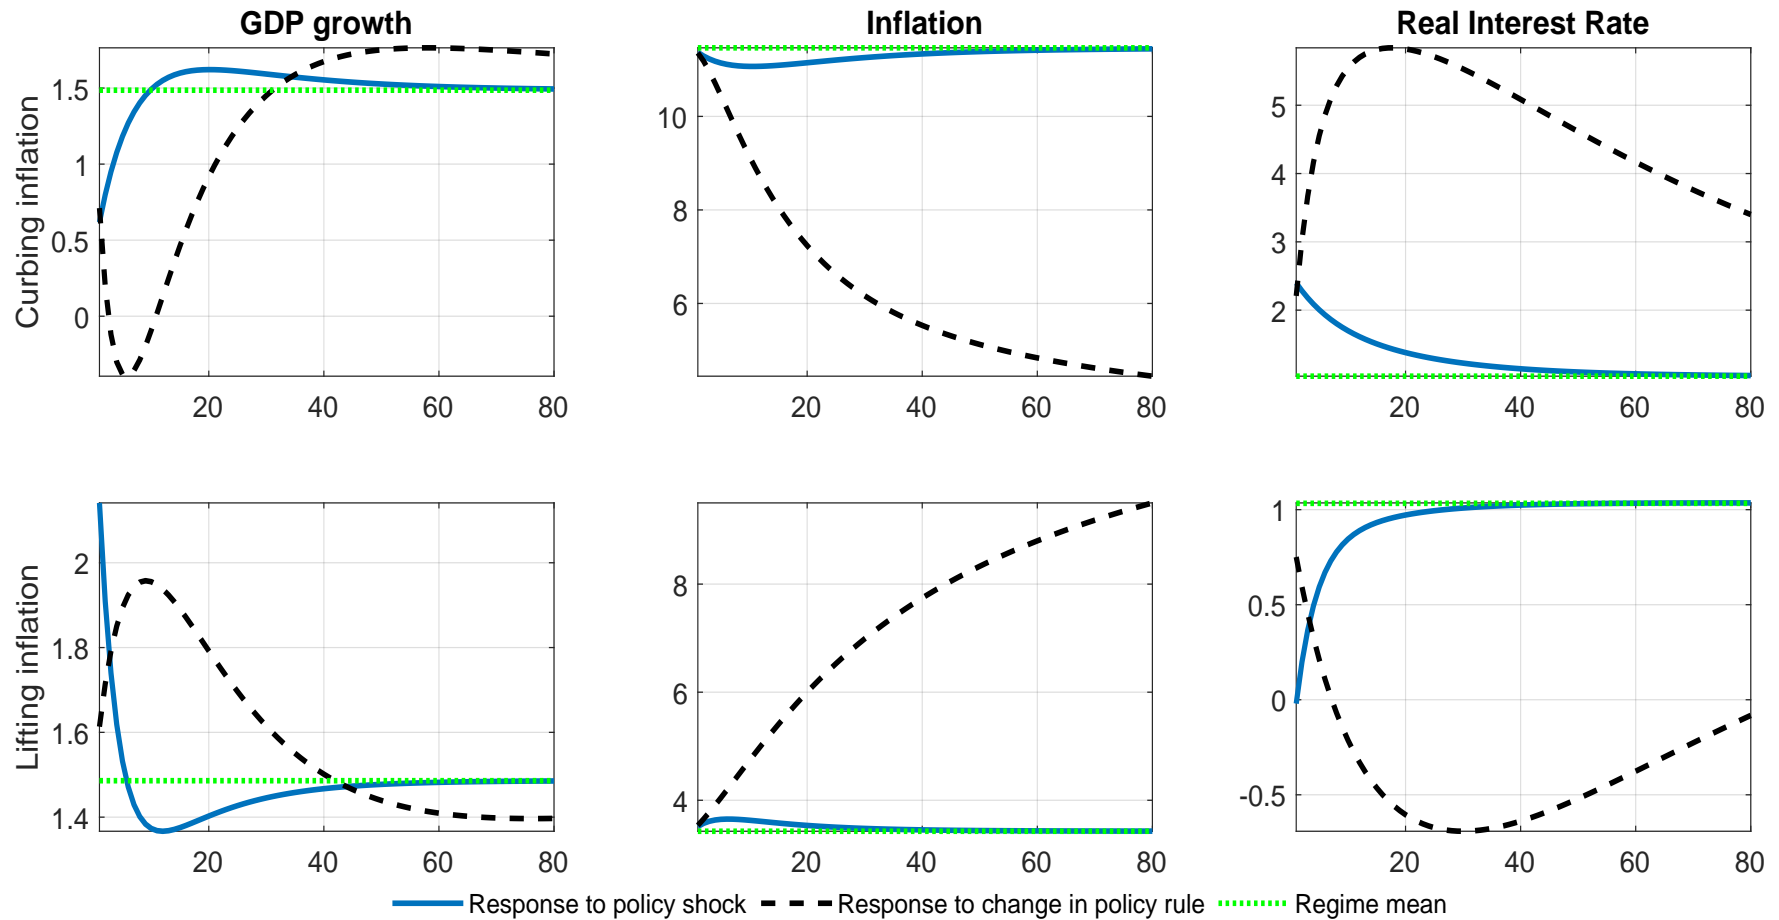


Figure A.5: Shadow rate estimates of the macro-finance model. Top row: Curbing inflation. The economy is initially in the Dovish regime. The blue solid line presents the evolution of the macro variables in response to a two standard deviation contractionary monetary policy shock and no regime change. The black dashed line presents the evolution of the macro variables in response to a regime change from the Dovish regime to the Hawkish regime. **Lower row: Lifting inflation.** The economy is initially in the Hawkish regime. The blue solid line presents the evolution of the macro variables 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 in response to a regime change from the Hawkish regime to the Dovish regime. Results are based on estimates obtained replacing the FFR with the Wu-Xia shadow rate when the ZLB is binding.




Figure A.6: Shadow rate estimates of the macro-finance model. Excess returns and policy rule changes. The figure reports the time series of the present discounted value of expected excess returns for different portfolios (dashed line, right axis) together with fluctuations of the real interest rate due to changes in the monetary policy rule (solid line, left axis). Results are based on estimates obtained replacing the FFR with the Wu-Xia shadow rate when the ZLB is binding.

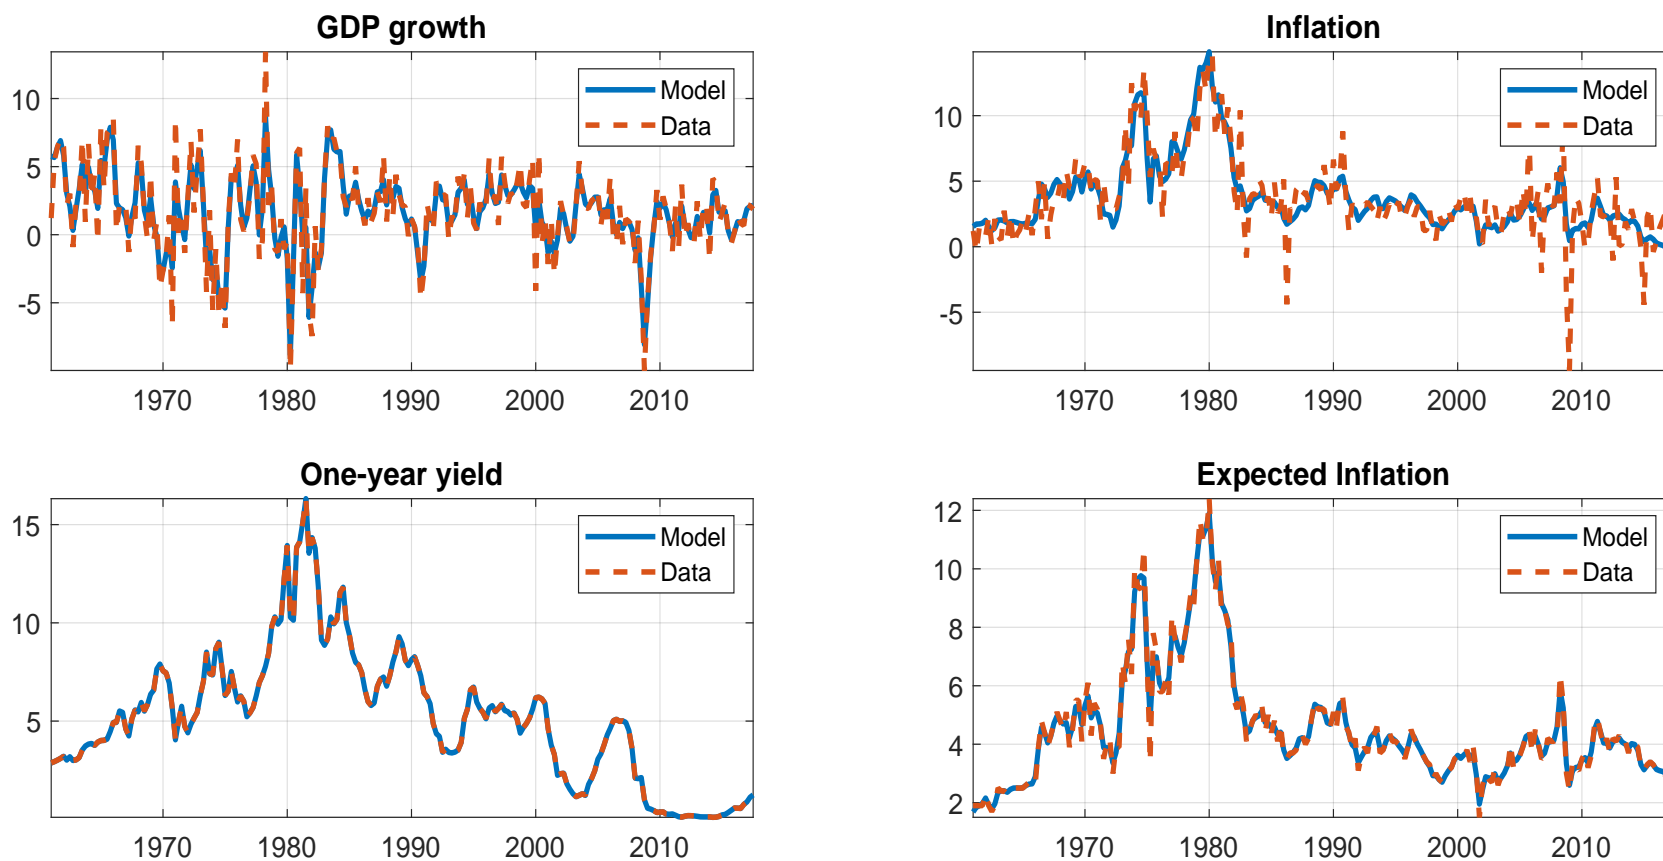


Figure A.7: One-year yield estimates of the macro-finance model. The figure reports the model implied series and the corresponding observed series. Expected inflation comes from the Michigan Survey of Consumers. The difference is due to observation errors. The sample spans 1961:Q1 - 2017:Q3. Results are based on estimates obtained using the one-year treasury yield instead of the FFR.

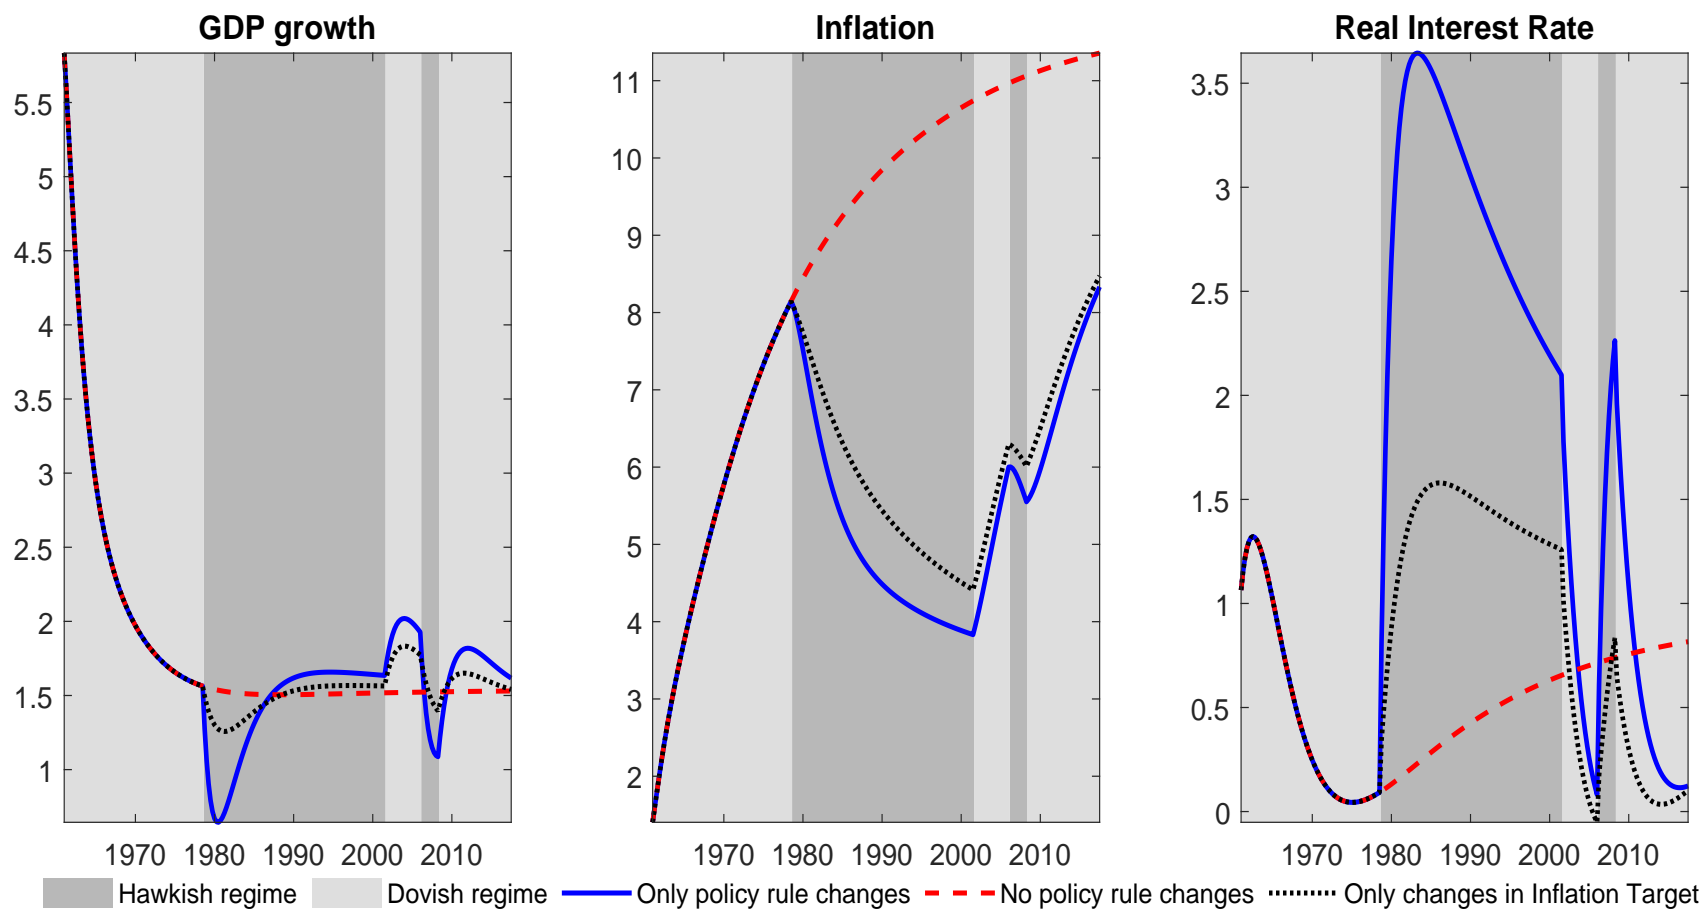


Figure A.8: One-year yield estimates of the macro-finance model. The blue line corresponds to the fluctuations generated by changes in both the target and the slope coefficients. The orange line assumes that monetary policy starts under the Dovish regime and no regime change occurs. Finally, the black dotted line assumes that changes in the target occurred, but that the slope coefficients in the Taylor rule coefficients always remain as in the Dovish-high target regime. Results are based on estimates obtained using the one-year treasury yield instead of the FFR.




Figure A.9: One-year yield estimates in 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 FFR and expected inflation. Expected inflation is obtained based on the model solution. Results are based on estimates obtained using the one-year treasury yield instead of the FFR.

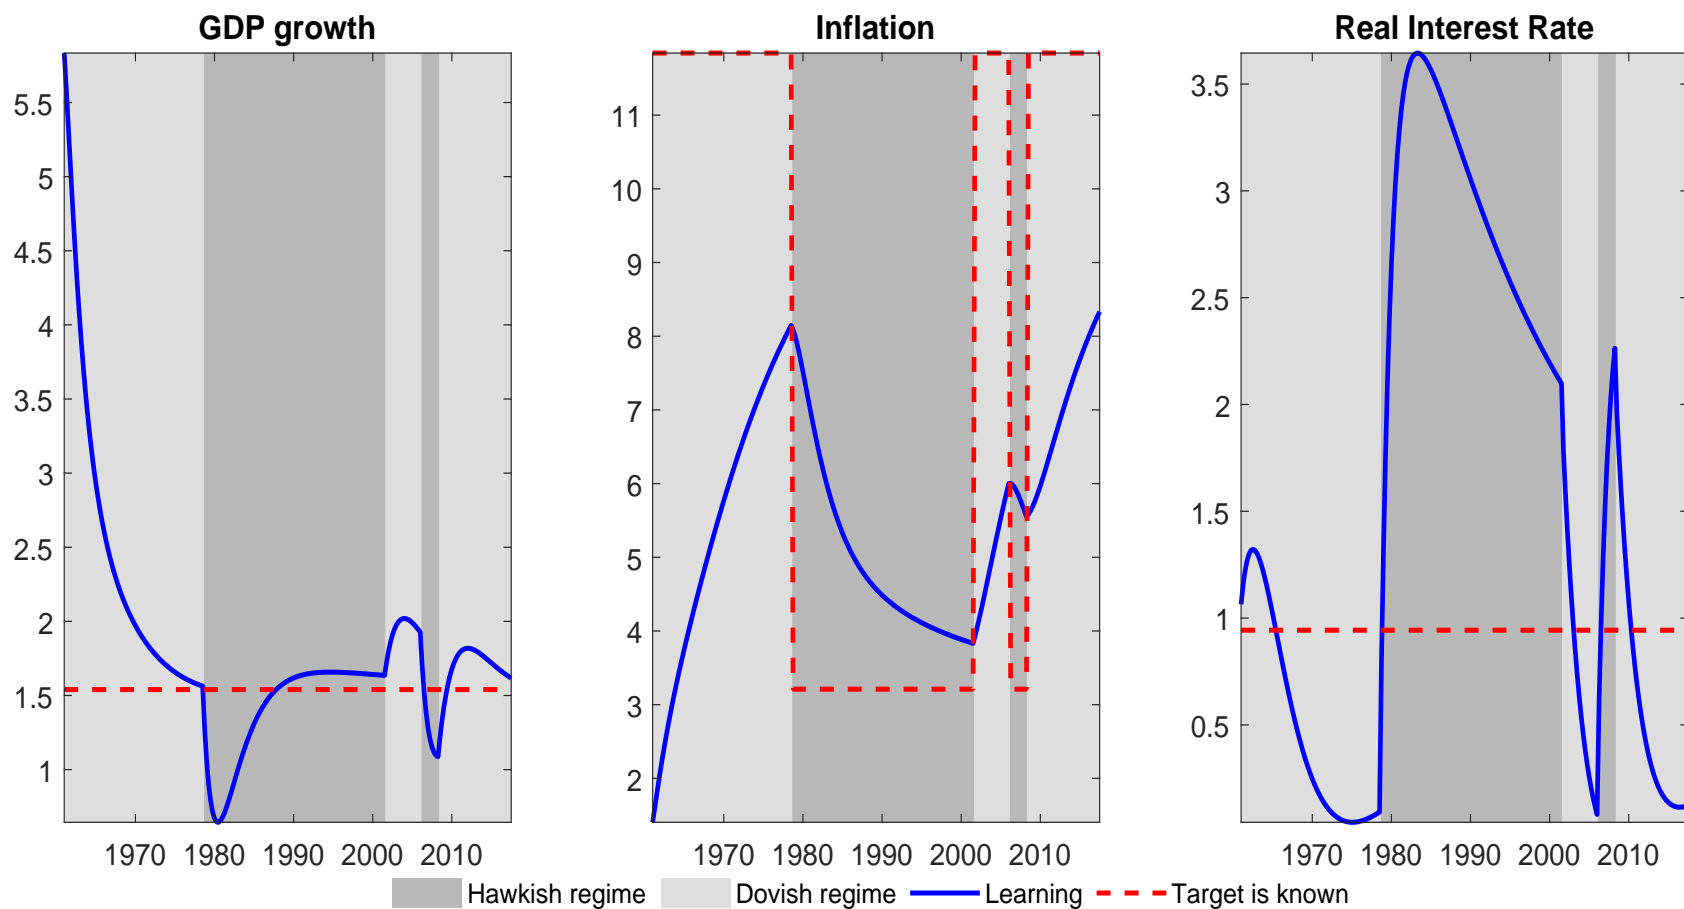


Figure A.10: One-year yield estimates of the macro-finance model. Perfect information about the target. The blue solid line shows estimated fluctuations generated only by changes in the policy rule (inflation target and slope coefficients) when agents learn about trend inflation. The orange dashed line shows a counterfactual in which the policy rule shifts but agents observe the inflation target. Dovish regime has a high target π^T and low activism against deviations from the target π^T . Hawkish regime has a low π^T and high activism against deviations from π^T . The sample spans 1961:Q1 - 2017:Q3. Results are based on estimates obtained using the one-year treasury yield instead of the FFR.




Figure A.11: One-year yield estimates of the macro-finance model. Top row: Curbing inflation. The economy is initially in the Dovish regime. The blue solid line presents the evolution of the macro variables 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 in response to a regime change from the Dovish regime to the Hawkish regime. **Lower row: Lifting inflation.** The economy is initially in the Hawkish regime. The blue solid line presents the evolution of the macro variables 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 in response to a regime change from the Hawkish regime to the Dovish regime. Results are based on estimates obtained using the one-year treasury yield instead of the FFR.




Figure A.12: One-year yield estimates of the macro-finance model. Excess returns and policy rule changes. The figure reports the time series of the present discounted value of expected excess returns for different portfolios (dashed line, right axis) together with fluctuations of the real interest rate due to changes in the monetary policy rule (solid line, left axis). Results are based on estimates obtained using the one-year treasury yield instead of the FFR.

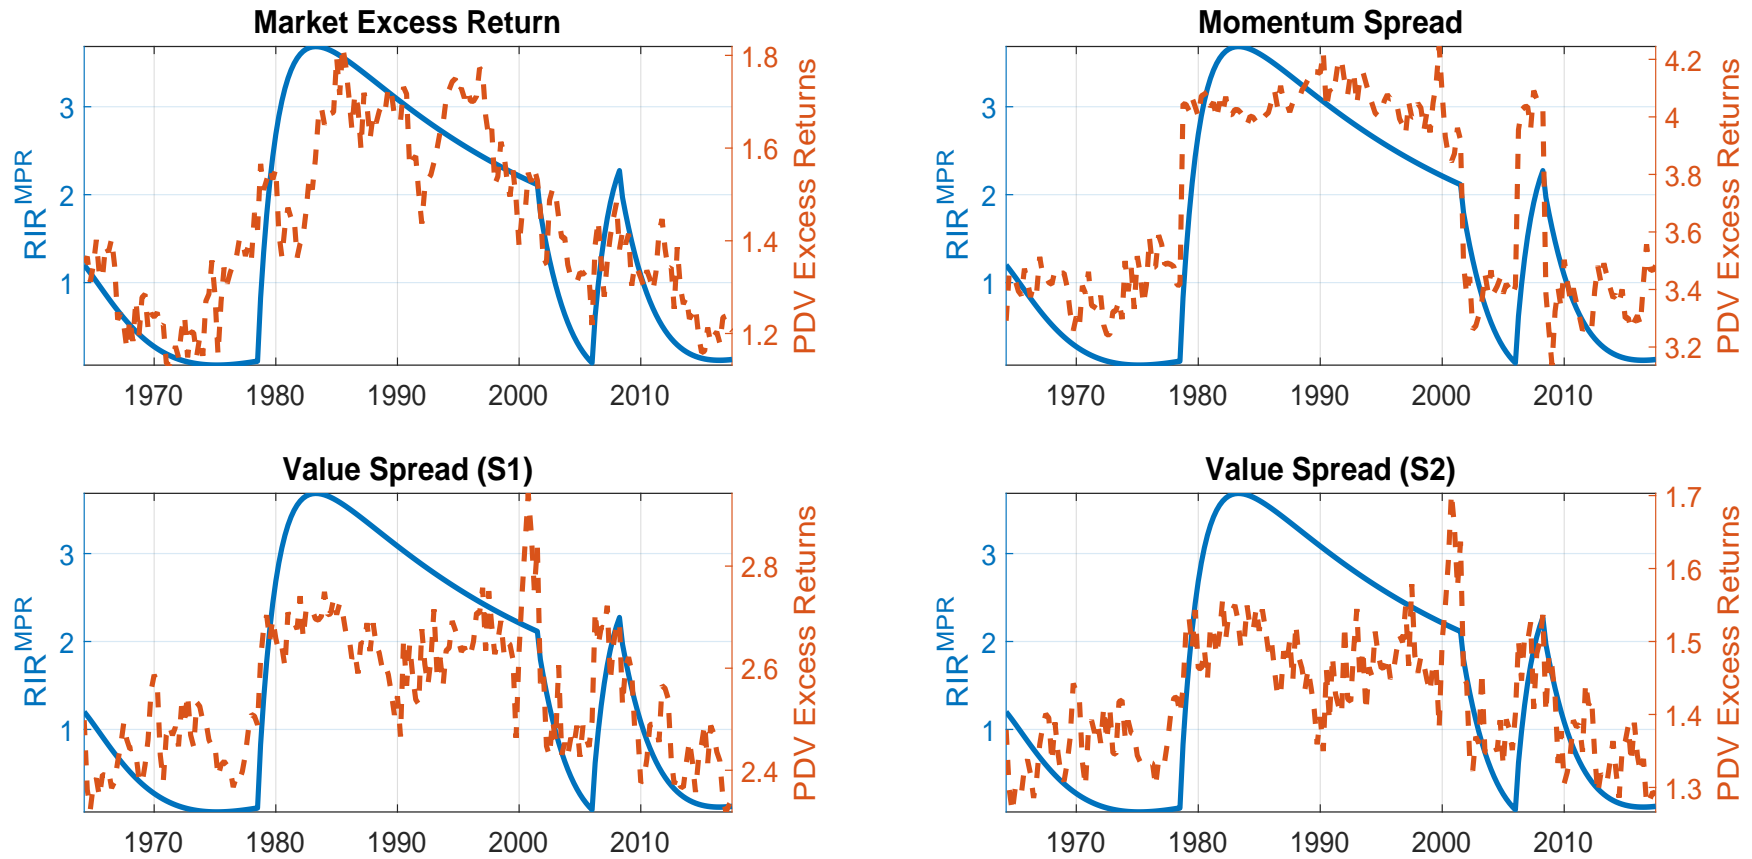


Figure A.13: Distribution of observation errors. The figure reports mean, median, and 90% error bands for the distribution of observation errors over time.

Estimating the Dynamic Macro-Finance Model

As explained in subsection 5, the macro block is put into state space form and estimated using standard Bayesian methods. The solution of the macro block at the estimated mode parameter values are then taken as inputs into the asset pricing block. To pin down the parameters of the asset pricing block, we take the estimates for the macro block as given and search for the parameters that minimize the distance between the data valuation ratio $-cay_t^{MS}$ and its model implied counterpart, cay_t^m . We also require the model to deliver an average annualized equity premium, $\bar{e}r$, of around 6%. Thus, we introduce a penalty for deviations of the average annualized equity premium from the 6% target. The distance between the two valuation ratios is defined as the sum of squared differences between the two ratios. Thus, we search for the set of parameters $\theta_p = \{k, \sigma_p, \beta_p, lp, p_{11}, p_{33}, p_{23}/(p_{23} + p_{24}), p_{41}/(p_{41} + p_{42})\}$ that minimizes the following object function:

$$\hat{\theta}_p = \arg \min \left[\sum_{t=1}^T (cay_t^{MS} - cay_t^m(\theta_p, X^T, \xi^T))^2 + .05 (|\bar{e}r(\theta_p, X^T, \xi^T) - 6|) \right]$$

where cay_t^m and the annualized average equity premium $\bar{e}r$ depend on the parameters of the model, the data used in the macro block estimation X^T , and the regime sequence in our sample ξ^T . The path for the model implied $-cay_t^m$ is computed based on the estimated regime sequence and the estimated initial conditions. Thus, we ask the model to explain as much as possible of the observed variation in $-cay_t^{MS}$ out of regime changes.

5.1 Constructing the PDV of Expected Returns from the Model

Suppose that we want to build the PDV of a vector of variables Y_t based on the model solution, where Y_t depends on \tilde{S}_t according to the following linear transformation:

$$Y_t = D + Z\tilde{S}_t.$$

In doing this, the econometrician can use the transition matrix reflecting the actual frequency of regime changes or the transition matrix used by the asset pricing agent when forming expectations. In the first case, we obtain the actual path of PDV of excess returns based on the data generating process, in the second case we obtain the PDV perceived by agents in the economy. For the main results in the paper, we compute the PDV that an econometrician would find if the dynamic macro model proposed in the paper generated the data. In this case we have:

$$\begin{aligned} \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t Y_{t+1+j} &= \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (D + Z\tilde{S}_{t+1+j}) \\ &= (1 - \rho)^{-1} D + Z \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t (\tilde{S}_{t+1+j}) \end{aligned}$$

where we have omitted the superfix p on the expectation operator because the probability assigned by the econometrician to moving across regimes is not in general the same as that

implied by the transition matrix used by the asset pricing agent. The transition matrix \mathbf{H} of the econometrician coincides with what was estimated in the first part of the paper and differs from \mathbf{H}^p and $\hat{\mathbf{H}}^p$, the transition matrices that enter the solution of the asset pricing block. To use \mathbf{H} in the expanded regime space, we expand it to cover the same number of regimes and reflect the probability of moving across them as implied by \mathbf{H} . We denote this expanded transition matrix consistent with the original transition matrix $\hat{\mathbf{H}}$.

As above, Define the column vectors q_t and π_t :

$$q_t = [q_t^1, \dots, q_t^m]' , q_t^i = \mathbb{E}_0 \left(\tilde{S}_t 1_{\zeta_t=i} \right) , \pi_t = [\pi_t^1, \dots, \pi_t^m]' ,$$

where $\pi_t^i = P_0(\zeta = i)$ and $1_{\zeta_t=i}$ is an indicator variable that is equal to 1 when regime i is in place and zero otherwise. The law of motion for $\tilde{q}_t = [q_t', \pi_t']'$ is then given by

$$\underbrace{\begin{bmatrix} q_t \\ \pi_t \end{bmatrix}}_{\tilde{q}_t} = \underbrace{\begin{bmatrix} \Omega & C\hat{\mathbf{H}} \\ & \hat{\mathbf{H}} \end{bmatrix}}_{\hat{\Omega}} \begin{bmatrix} q_{t-1} \\ \pi_{t-1} \end{bmatrix}$$

where $\pi_t = [\pi_{1,t}, \dots, \pi_{m,t}]'$, $\Omega = bdiag(A_1, \dots, A_m)\hat{\mathbf{H}}$, and $C = bdiag(c_1, \dots, c_m)$.

Similar formulas are used to compute risk premia for the individual portfolios. The premium for a portfolio z coincides with the present discounted value of its excess returns:

$$\underbrace{premia_{z,t}}_{\text{Premia}} \equiv \underbrace{\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t r_{t+1+j}}_{\text{PDV of excess returns}} = 1'_r w (I - \rho\Omega)^{-1} \left[\Omega q_{t|t} + C (I - \rho\hat{\mathbf{H}})^{-1} \hat{\mathbf{H}} \pi_{t|t} \right] , \quad (\text{A85})$$

where $1'_r$ is a vector used to extract the PDV of excess returns from a vector containing the PDV of all variables included in the VAR.

ZLB Robustness Checks for the Dynamic Macro Model

In this appendix, we conduct two robustness checks to verify that our results are not distorted by the time spent at the zero lower bound (ZLB) in the aftermath of the financial crisis. First, we re-estimate our MS-DSGE model using the Wu-Xia (Wu and Xia (2016)) shadow rate. Second, we use the one-year Treasury yield instead of the federal funds rate in our estimation. The shadow rate is downloaded from Professor Wu's website, while the one-year Treasury yield is downloaded from FRED. The figures presented in this appendix show that the main result of the paper are not affected by using these alternative measures for the interest rate. The figures pertaining to these estimates are found in Figures A.1-A.12.

In the Wu and Xia model, the short-term interest rate is the maximum of the shadow federal funds rate and the lower bound on interest rates. Wu and Xia set this lower bound to 25 basis points because that was the rate paid on both required and excess reserve balances during the

December 16, 2008, to December 15, 2015, period when the Federal Open Market Committee (FOMC) set the target range for the federal funds rate at 0 to 25 basis points. On December 16, 2015, the FOMC increased the rate paid on reserve balances to 50 basis points and the target range for the federal funds rate to 25 to 50 basis points. Once the lower bound is not binding anymore, the shadow rate coincides with the actual FFR.

The results of Wu and Xia are based on a multivariate version of the shadow rate term structure model (SRTSM) introduced by Black (1995). In the SRTSM, the observed short term rate is the maximum between a lower bound and the shadow rate. The shadow rate, in turn, is an affine function of a vector of state variables that follow a VAR process. Absent the lower bound, the model would be fully linear. Thus, the lower bound introduces a non-linearity in the mapping from the factors to the observed short term interest rate. The key idea behind the model and the work of Wu and Xia is that by observing the behavior of forward rates at different maturities, the researcher can back out a measure of the shadow short term interest rate. In other words, forward rates reflect the overall monetary policy stance and can be used to recover the implicit behavior of the shadow interest rate.