Internet Appendix for "Monetary Policy and Asset Valuation"

FRANCESCO BIANCHI, MARTIN LETTAU, and SYDNEY C. LUDVIGSON *

This appendix provides additional content for the paper "Monetary Policy and Asset Valuation."Section I is a data appendix. Section II gives estimation details for cay^{MS}. Section III provides details of the Gibbs sampling algorithm. Section IV explains the most likely regime sequence. Section V discusses estimation of the present discounted values (PDVs) of future return premia. Section VI discusses variable selection for the Markov-switching vector autoregression (MS-VAR) used to compute PDVs, while Section VII covers estimation of the MS-VAR. Section VIII presents details of the Macro Block of the macro-finance model, Section IX covers the Asset Pricing block, and Section X provides details on the model solution. Section XI discusses details of the model estimation. Section XII discusses details of the model-based PDVs. Section XIII discusses ZLB robustness checks. Section XIV discusses the inflation target in the early- and late-dovish subperiods.

I. 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.

*Citation format: Bianchi, Francesco, Martin Lettau, and Sydney Ludvigson, Internet Appendix for "Monetary Policy and Asset Valuation," Journal of Finance [DOI STRING]. Please note: Wiley is not responsible for the content or functionality of any supporting information supplied by the authors. Any queries (other than missing material) should be directed to the authors of the article.

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 [wages and salaries/(wages and salaries + proprietors' income with IVA and CCADJ + rental income + personal dividends + 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 data set consists of $vwretx(t) = (P_t/P_{t-1}) - 1$, the return on a portfolio that does not 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

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

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

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

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: $d_{t+12} - d_t = d_{t+12} - d_{t+11} + d_{t+11} - d_{t+10} + \cdots + d_{t+1} - d_t$. 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 in this context. 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 $\ln\left(P_t^{Ax}/D_t^A\right)$. Note that this value for dividend growth is used only 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 Flow of Funds (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. 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: (i) the momentum return spread, that is, the difference between the excess return of the extreme winner (M10) portfolio and the excess return of the extreme loser (M1) portfolio; (ii) the value return spread (S1), that is, 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; (iii) the value return spread (S2), that is, the difference between the excess return of the size 2 high book-to-market (BM) portfolio and the excess return of the small size 2 low BM portfolio; (iv) the momentum BM spread: the difference between the logarithm of the BM ratio of the M10 and M1 portfolios; (v) the value BM spread (S1): the difference between the logarithm of the BM ratio of the small (size quintile 1) high BM portfolio and the logarithm of the BM ratio of the small (size 1) low BM portfolio; (vi) the value BM spread (S2): the difference between the logarithm of the BM ratio of the size quintile 2 high BM portfolio and the logarithm of the BM ratio of the size 2 low BM portfolio; (vii) the real federal funds rate (FFR) (FFR minus inflation); (viii) 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: (i) the market excess return, computed as the difference in the CRSP value-weighted stock market return (including dividend redistributions) and the three-month Treasury bill rate; (ii) $-cay^{MS}$; (iii) the small stock value spread (log-difference in the BM ratio of the S1 value and S1 growth portfolio); (iv) the SMB factor from Fama and French; (v) 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. Expected inflation is the mean of the one-year-ahead expected inflation based on the Michigan survey. All variables are annualized.

II. Computing cay^{MS}

Let $\boldsymbol{z}_t$ be a 3×1 vector of data on c_t, a_t, $\widetilde{y}_t$, and k leads and k lags of Δa_t and Δy_t and let $\boldsymbol{Z}_t = (\boldsymbol{z}_t, \boldsymbol{z}_{t-1}, ..., \boldsymbol{z}_1)$ 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}:

$$c_t = \alpha_{\xi_t} + \beta_a a_t + \beta_y y_t + \sum_{i=-k}^{k} b_{a,i} \Delta a_{t+i} + \sum_{i=-k}^{k} b_{y,i} \Delta y_{t+i} + \sigma^c \varepsilon_t^c, \tag{IA1}$$

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

$$mps_t = r_{\xi_t} + \epsilon_t^r, \tag{IA2}$$

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 Markov-switching 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 changes in both α_{ξ_t} and $\overline{r}_{\xi_t}$.

[1]The 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:

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

where ξ_t is a hidden variable that follows a Markov-switching process with transition matrix $\mathbf{H}$. Collect all model parameters into 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

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

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

Collect the conditional probabilities $\pi_{t|t}^i = p(\xi_t = i|Y^t;\boldsymbol{\theta})$ for $i = 1,..,m$ into an $m \times 1$ vector $\pi_{t|t} = p(\xi_t|Y^t;\boldsymbol{\theta})$. The filtered probabilities reflect the probability of a regime conditional on the data up to time t, $\pi_{t|t} = p(\xi_t|Y^t;\boldsymbol{\theta})$, for $t = 1,...,T$, 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:

$$\begin{aligned} \pi_{t|t} &= \frac{\pi_{t|t-1} \odot \eta_t}{\mathbf{1}'\left(\pi_{t|t-1} \odot \eta_t\right)} \\ \pi_{t+1|t} &= \mathbf{H}\pi_{t|t}, \end{aligned} \tag{IA3}$$

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

$$p(c_t, mps_t|\xi_t = j, x_{M,t}, x_{F,t}; \boldsymbol{\theta}) = \frac{1}{\sqrt{2\pi\sigma^{c,2}}}\frac{1}{\sqrt{2\pi\sigma^{r,2}}}\exp\left[-\frac{[c_t - (\alpha_j x_{M,t} + \beta x_{F,t})]^2}{2\sigma^{c,2}} - \frac{[mps_t - r_j x_{M,t}]^2}{2\sigma^{r,2}}\right], \tag{IA4}$$

the symbol $\odot$ denotes element-by-element multiplication, and $\mathbf{1}$ is a vector with all elements equal to one. To initialize the recursive calculation we need an assumption on the distribution of ξ_0. We assume that the two regimes have equal probabilities: $p(\xi_0 = 1) = 0.5 = p(\xi_0 = 2)$.

The smoothed probabilities reflect all of the information that can be extracted from the full 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. A recursive algorithm can then be implemented to derive the other probabilities:

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

where $(\div)$ denotes element-by-element division.

In using the DLS regression (IA1) to estimate cointegrating parameters, we lose six leads and six lags. For estimates of cay_t^{FC}, we take the estimated coefficients and we apply them to the full sample. To extend our estimates of cay_t^{MS} over the full sample, we can likewise apply the parameter estimates to the full sample but we need an estimate of the regime probabilities in the first six and last six 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 that no lagged values are available and the DLS regression omits all lags, but all leads are included. When at $t = -4$ we assume that 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 that all lags are available and all leads except one are available, when $t = T+2$ we assume that 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 $\mathbf{H}$.

III. 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 by $x_{M,t}$, while $x_{F,t}$ denotes the vector containing all of the variables whose coefficients are kept constant. We then obtain

$$c_t = \alpha_{\xi_t} x_{M,t} + \beta x_{F,t} + \sigma^c \varepsilon_t^c \tag{IA5}$$

$$mps_t = r_{\xi_t} x_{M,t} + \sigma^r \varepsilon_t^r, \tag{IA6}$$

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

Suppose the Gibbs sampling algorithm has reached the n^{th} iteration. We then have draws for $r_{\xi_t,n}, \sigma_n^r$, β_n, $\alpha_{\xi_t,n}$, σ_n^c, $\mathbf{H}_n$, and ξ_n^T, where $\xi_n^T = \{\xi_{1,n},\ \xi_{2,n},...,\xi_{T,n}\}$ denotes a draw for the whole regime sequence. The parameters for equations (IA5) and (IA6) 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:

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

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 $\beta \sim N\left(B_{\beta,0}, V_{\beta,0}\right)$), so that the conditional (on $\alpha_{\xi_t,n}$, σ_n^c, and ξ_n^T) posterior distribution is given by

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

with $V_{\beta,T} = \left(V_{\beta,0}^{-1} + \widetilde{X}'_F \widetilde{X}_F\right)^{-1}$ and $B_{\beta,T} = V_{\beta,T}\left[V_{\beta,0}^{-1} B_{\beta,0} + \widetilde{X}'_F \widetilde{C}\right]$, where $\widetilde{C} = (\widetilde{c}_1, ..., \widetilde{c}_T)'$, $\widetilde{X}_F = (x_{F,1}, ..., x_{F,T})'$, 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 we have $B_{\beta,0} = 0$, $V_{\beta,0}^{-1} = 0$, and $B_{\beta,T}$ and $V_{\beta,T}$ coincide with the maximum likelihood estimates (MLEs), 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:

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

   The above regression has standardized shocks and Markov-switching coefficients in the transformed data. Using ξ_n^T, we can group all of the observations that pertain to the same regime i. Given the prior $\alpha_i \sim N\left(B_{\alpha_i,0}, V_{\alpha_i,0}\right)$) 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:

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

   where $V_{\alpha_i,T} = \left(V_{\alpha_i,0}^{-1} + \widetilde{X}'_{M,i} \widetilde{X}_{M,i}\right)^{-1}$ and $B_{\alpha_i,T} = V_{\alpha_i,T}\left[V_{\alpha_i,0}^{-1} B_{\alpha_i,0} + \widetilde{X}'_{M,i} \widetilde{C}_i\right]$, with $\widetilde{C}_i$ and $\widetilde{X}_{M,i}$ collecting all 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 Markov-switching coefficients of the regression: $\alpha_i \sim N\left(B_{\alpha_i,0}, V_{\alpha_i,0}\right)$ 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 MLEs, 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:

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

   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 (σ_i^r and ξ_n^T) posterior distribution:

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

where $V_{r_i,T} = \left(V_{r_i,0}^{-1} + \widetilde{X}'_{M,i}\widetilde{X}_{M,i}\right)^{-1}$ and $B_{r_i,T} = V_{r_i,T}\left[V_{r_i,0}^{-1}B_{r_i,0} + \widetilde{X}'_{r,i}\widetilde{R}_i\right]$, with $\widetilde{R}_i$ and $\widetilde{X}_{r,i}$ collecting all 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 Markov-switching 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 MLEs, 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

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

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:

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

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/v_T. With flat priors we have $Q_0 = 0$ and $v_0 = 0$, and the mean of σ^c is therefore $(E^{c\prime}E^c)/T$, which coincides with the standard 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^r_{n+1} from the conditional (on $r_{\xi_t,n+1}$ and ξ^T_n) posterior inverse gamma distribution:

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

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 $\sigma^r \sim IG(Q^r_T, v^r_T)$ is Q^r_T / v^r_T. With flat priors we have $Q_0 = 0$ and $v_0 = 0$, and the mean of σ^r is therefore $(E^{r\prime}E^r)/T$, which coincides with the MLE of σ^r, conditional on the other parameters.

6. Sampling ξ^T_{n+1}: Given $r_{\xi_t,n+1}$, σ^r_{n+1}, β_{n+1}, $\alpha_{\xi_t,n+1}$, σ^c_{n+1}, and $\mathbf{H}_n$, we can treat equations (IA5) and (IA6) as a multivariate regression in which some parameters are restricted to zero. This allows us to obtain filtered probabilities for the regimes using the filter described in Hamilton (1994). Following Kim and Nelson (1999) we then use Multi-Move Gibbs sampling to draw a regime sequence ξ^T_{n+1}.

7. Sampling $\mathbf{H}_{n+1}$: Given the draws for the Markov-switching state variables ξ^T_{n+1}, 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.[2] 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^T_{n+1}, 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.

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

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 III.A below. We use the draws obtained with the Gibbs sampling algorithm to characterize parameter uncertainty. 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.

A. Convergence Checks

The 90% credible sets are obtained making 2,000,000 draws from the posterior using the Gibbs sampling algorithm. One in every 1,000 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 1% accuracy to be achieved with 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 Markov Chain Monte Carlo (MCMC) algorithm due to the fact that the likelihood can have multiple peaks. The tables below pertain to convergence of the Gibbs sampling algorithm.

IV. 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 $\xi_n^T = \{\widehat{\xi}_{1,n}, ..., \widehat{\xi}_{T,n}\}$ that is most likely to have occurred, given our estimated posterior mode parameter values for $\boldsymbol{\theta}$. This sequence is computed as follows. Let $P\left(\xi_t = i|\boldsymbol{Z}_{t-1};\boldsymbol{\theta}\right) \equiv \pi^i_{t|t-1}$. First, we run Hamilton's filter to get the vector of filtered probabilities $\pi_{t|t}$, $t = 1, 2, ..., T$. The Hamilton filter can be expressed iteratively as

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

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

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

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

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

We first take $\pi_{T|T}$ from the Hamilton filter and choose the regime that is associated with the largest probability, that is, if $\pi_{T|T} = (0.9, 0.1)$, where the first element corresponds to the probability of regime 1, we select $\widehat{\xi}_T = 1$, indicating that we are in regime 1 in period T. We then 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. Next we repeat the same procedure by choosing the regime for $T-1$ that has the largest probability at $T - 1$. For example, if $\pi_{T-1|T} = (0.2, 0.8)$, we select $\widehat{\xi}_{T-1} = 2$, indicating that we are in regime 2 in period $T - 1$, and we update to $\pi_{T-1|T} = (0, 1)$, which is used again on the right-hand side now

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

We proceed in this manner until we have a most likely regime sequence ξ^T for the entire sample $t = 1, 2, ..., T$. Two aspects of this procedure are worth noting. First, it fails if the updated probabilities are exactly $(0.5, 0.5)$. 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.

V. Book-to-Market Ratio and Present Discounted Values (PDVs)

We use the methods and assumptions of the previous section to obtain the present value decomposition of the BM 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 $\xi_t = 1, ..., m$, 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.

Section VII.A below on conditional expectations and volatility shows how to compute $\mathbb{E}_t (Z_{t+s}) = w q_{t+s|t}$, where

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

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

Now consider the formula from Vuolteenaho (1999):

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

Given that our goal is to assess whether assets with different risk profiles are differently affected by the breaks in the long-term interest rates, we focus on the *difference* between the

BM ratios. Specifically, given two portfolios x and y, we are interested in how the difference in their BM ratios, $\theta_{x,t} - \theta_{y,t}$, varies across the two regimes:

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

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

$$\underbrace{\theta_{x,t} - \theta_{y,t} + \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t \left(e^*_{x,t+1+j} - e^*_{y,t+1+j}\right)}_{\text{Spread in BM ratios corrected for earnings}} = \underbrace{\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t \left(r_{x,t+1+j} - r_{y,t+1+j}\right)}_{\text{PDV of the expected spread in excess returns}}. \qquad \text{(IA7)}$$

The spread in excess returns $r_{xy,t} \equiv r_{x,t} - r_{y,t}$. The right-hand side of (IA7) can then be computed as

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

We therefore have

$$\widetilde{\theta}_{xy,t} \equiv \underbrace{\widetilde{\theta}_{x,t} - \widetilde{\theta}_{y,t} + \sum_{j=0}^{\infty} \rho^j \mathbb{E}_t \left(e^*_{x,t+1+j} - e^*_{y,t+1+j}\right)}_{\text{Spread in BM ratios corrected for earnings}} = 1'_{r_{xy}} w \left(I - \rho\Omega\right)^{-1} \left[\Omega q_{t|t} + C \left(I - \rho\mathbf{H}\right)^{-1} \mathbf{H}\pi_{t|t}\right], \qquad \text{(IA8)}$$

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

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

$$\underbrace{premia_{z,t}}_{\text{Premia}} \equiv \underbrace{\sum_{j=0}^{\infty} \rho^j \mathbb{E}_t \left(r_{z,t+1+j}\right)}_{\text{PDV of excess returns}} = 1'_{r_z} w \left(I - \rho\Omega\right)^{-1} \left[\Omega q_{t|t} + C \left(I - \rho\mathbf{H}\right)^{-1} \mathbf{H}\pi_{t|t}\right], \qquad \text{(IA9)}$$

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 below. This implies that the vector $\pi_{t|t}$ assumes one of two values, $(1,0)'$ or $(0,1)'$.

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

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

where $\overline{\pi}_i = 1_i$ and $\overline{q}_i \equiv [0,...,\overline{\mu}_i,...,0]$ is a column vector that contains the conditional steady state of the mean Z_t conditional on being in regime i, that is, $\mathbb{E}_i\left(Z_t\right) = \overline{\mu}_i = \left(I_n - A_i\right)^{-1} c_i$ and zero otherwise. Recall that the conditional steady state, $\overline{\mu}_i$, 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 is 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 $\overline{\overline{\theta}}_{xy}^{i}$ is computed by conditioning on the economy initially being at $Z_t = \overline{\mu}_i$ and in regime i, but taking into account the possibility of regime changes in the future. Therefore, we can also think about $\overline{\overline{\theta}}_{xy}^{i}$ as the expected value of $\widetilde{\theta}_{xy,t}$, 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

$$\overline{\overline{\theta}}_{xy}^{i} = \mathbb{E}\left(\widetilde{\theta}_{xy,t}|\xi_t = i, Z_t = \overline{\mu}_i\right). \tag{IA10}$$

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

$$\overline{premia}_{z}^{i} \equiv 1_{r_z}' w\left(I-\rho\Omega\right)^{-1}\left[\Omega\overline{q}_i + C\left(I-\rho\mathbf{H}\right)^{-1}\mathbf{H}\overline{\pi}_i\right]. \tag{IA11}$$

Formulas (IA8), (IA9), (IA10), and (IA11) are used in the paper to produce Figure 4 and Table IV. For each draw of the VAR parameters from the posterior distribution, we can compute the evolution of $\widetilde{\theta}_{xy,t}$ and individual portfolio $premia_{z,t}$, by using (IA8) and (IA9). We therefore obtain a posterior distribution for $\widetilde{\theta}_{xy,t}$ 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 $\overline{\overline{\widetilde{\theta}}}^i_{xy}$ and the difference $\overline{\overline{\widetilde{\theta}}}^1_{xy} - \overline{\overline{\widetilde{\theta}}}^2_{xy}$. We therefore obtain a posterior distribution for $\overline{\overline{\widetilde{\theta}}}^i_{xy}$ and for the difference $\overline{\overline{\widetilde{\theta}}}^1_{xy} - \overline{\overline{\widetilde{\theta}}}^2_{xy}$. The medians of the distribution of $\overline{\overline{\widetilde{\theta}}}^i_{xy}$ and $\overline{premia}^i_z$ for $i = 1, 2$ are reported in Figure 4 (red dashed line). Table IV reports the median and the 68% posterior credible sets both for the distribution of $\overline{\overline{\widetilde{\theta}}}^i_{xy}$ for $i = 1, 2$, and for the difference in these across regimes, $\overline{\overline{\widetilde{\theta}}}^1_{xy} - \overline{\overline{\widetilde{\theta}}}^2_{xy}$. Finally, the last row of Table IV reports the percentage of draws for which $\overline{\overline{\widetilde{\theta}}}^1_{xy} - \overline{\overline{\widetilde{\theta}}}^2_{xy} > 0$ and $\overline{premia}^1_z - \overline{premia}^2_z > 0$ as the probability that return premia are lower in the high-asset valuation/low-interest rate regime than they are in the low-asset valuation/high-interest rate regime.

VI. 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 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 we are conditioning to the regime sequence obtained when estimating cay^{MS}, the intercept for the corresponding equation will adjust 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 BM ratio of the small value portfolios and the

BM ratio of the small growth portfolios), Real FFR, term yield spread, four of the five Fama and French factors (SMB, HML, RMW, CMA), and *cay* (based on personal consumption expenditures, 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, and hence, this variable is automatically included in the MS-VAR.

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

2. MS-VAR for (i) 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; (ii) 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; and (iii) 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): (i) Momentum return spread; (ii) Value return spread (S1); (iii) Value return spread (S2); (iv) 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; (v) Value BM spread (S1): The difference between the logarithm of the BM ratio of the small (size quintile 1) high BM portfolio and the logarithm of the BM ratio of the small (size 1) low BM portfolio; and (vi) Value BM spread (S2): The difference between the logarithm of the BM ratio of the size quintile 2 high BM portfolio and the logarithm of the BM ratio of the size 2 low BM 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, and 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.

VII. 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 of 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

$$\begin{aligned} \underset{T\times n}{Y} &= \underset{(T\times k)(k\times n)}{XA_{\xi_t}} + \underset{T\times n}{\varepsilon},\ \xi_t = 1,2 \\ \varepsilon_t &\sim N\left(0, \Sigma_{\xi_t}\right), \end{aligned}$$

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

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

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

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

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

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

where T_i, Y_i, and 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 MLEs using observations in regime i:

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

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 $\mathbf{H}^A$ depends only 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.[3] 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 $\eta_{ij,r+1}$ is the number of transitions from regime i to regime j based on ξ_n^T, while a_{ii} and a_{ij} are 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 in general be different from the posterior distribution of $\mathbf{H}$ because the

[3]The Dirichlet distribution is a generalization of the beta distribution that allows one to potentially consider more than two regimes. See, for example, Sims and Zha (2006).

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).

A. 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 first-order MS-VAR

$$Z_t = c_{\xi_t} + A_{\xi_t} Z_{t-1} + V_{\xi_t} \varepsilon_t,\ \varepsilon_t \sim N(0, I) \tag{IA12}$$

and suppose that we are interested in $\mathbb{E}_0(Z_t) = \mathbb{E}(Z_t|\mathbb{I}_0)$, where $\mathbb{I}_0$ is 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 Section VI above. Let m be the number of Markov-switching states. Define the $mn \times 1$ column vector q_t as

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

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

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

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:

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

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

PROPOSITION 1: *Consider a Markov-switching model whose law of motion can be described by (IA12) and define* $q_t^i = \mathbb{E}_0(Z_t 1_{\xi_t=i})$ *for* $i = 1...m$*. 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 $\mu_{t+s|t} \equiv \mathbb{E}_t(Z_{t+s})$, that is, the expected value for the vector Z_{t+s} conditional on the information set available at time t. If we define

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

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

$$\mu_{t+s|t} = \mathbb{E}_t(Z_{t+s}) = wq_{t+s|t}, \tag{IA13}$$

where for $s \geq 1$, $q_{t+s|t}$ evolves according to

$$q_{t+s|t} = C\pi_{t+s|t} + \Omega q_{t+s-1|t} \tag{IA14}$$

$$\pi_{t+s|t} = \mathbf{H}\pi_{t+s-1|t} \tag{IA15}$$

with $\pi_{t+s|t} = \left[\pi_{t+s|t}^{1}, ..., \pi_{t+s|t}^{m}\right]'$, $\Omega = bdiag(A_1, ..., A_m)(\mathbf{H} \otimes I_n)$, and $\underset{mn\times m}{C} = bdiag(c_1, ..., c_m)$, where for example, 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 uses 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 also make use of the following result: $\varphi(X_1X_2X_3) = (X_3' \otimes X_1)\varphi(X_2)$.

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

$$Q_t = \left[Q_t^{1\prime}, ..., Q_t^{m\prime}\right]',$$

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

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

where the matrix $W = [I_{n^2}, ..., I_{n^2}]$ is obtained by 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 (IA12) and define* $Q_t^i = \varphi\left[\mathbb{E}_0\left(Z_t Z_t' 1_{\xi_t=i}\right)\right]$, $q_t^i = \mathbb{E}_0\left[Z_t 1_{\xi_t=i}\right]$, *and* $\pi_t^i = P_0\left(\xi_t = i\right)$, *for* $i = 1...m$. *Then* $Q_t^j = \left[\widehat{cc}_j + \widehat{VV}_j \varphi\left[I_k\right]\right]\pi_t^j + \sum_{i=1}^m \left[\widehat{AA}_j Q_{t-1}^i + \widehat{DAC}_j q_{t-1}^i\right] h_{ji}$, *where* $\widehat{cc}_j = (c_j \otimes c_j)$, $\widehat{VV}_j = (V_j \otimes V_j)$, $\widehat{AA}_j = (A_j \otimes A_j)$, *and* $\widehat{DAC}_j = (A_j \otimes c_j) + (c_j \otimes A_j)$.

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 $\mathbb{E}_t\left(Z_{t+s}Z_{t+s}'\right)$, that is, the second moment of the vector Z_{t+s} conditional on the information available at time t. If we define

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

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

$$\begin{aligned} Q_{t+s|t} &= \Xi Q_{t+s-1|t} + \widehat{DAC} q_{t+s-1|t} + \widehat{V}c\pi_{t+s|t} &\text{(IA16)}\\ q_{t+s|t} &= C\pi_{t+s|t} + \Omega q_{t+s-1|t},\ \pi_{t+s|t} = \mathbf{H}\pi_{t+s-1|t}, &\text{(IA17)}\end{aligned}$$

where

$$\begin{aligned}\Xi &= bdiag(\widehat{AA}_1, ..., \widehat{AA}_m)(\mathbf{H} \otimes I_{n^2}),\ \widehat{V}c = \left[\widehat{VV} + \widehat{cc}\right],\ \widehat{cc} = bdiag(\widehat{cc}_1, ..., \widehat{cc}_m),\\ \widehat{VV} &= bdiag(\widehat{VV}_1\varphi\left[I_k\right], ..., \widehat{VV}_m\varphi\left[I_k\right]),\ \widehat{DAC} = bdiag(\widehat{DAC}_1, ..., \widehat{DAC}_m)(\mathbf{H} \otimes I_n).\end{aligned}$$

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

$$\varphi\left[\mathbb{V}_t\left(Z_{t+s}\right)\right] = M_{t+s|t} - \varphi\left[\mu_{t+s|t}\mu_{t+s|t}'\right], \tag{IA18}$$

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

To report estimates of (IA13) and (IA18) we proceed as follows. Note that $\mu_{t+s|t} = \mathbb{E}_t(Z_{t+s}) = wq_{t+s|t}$ and $M_{t+s|t}$ depend only on $q_{t+s|t}$ and $Q_{t+s|t}$. Furthermore, we can express (IA14)-(IA15) and (IA16) to (IA17) in compact form as

$$\widetilde{Q}_{t+s|t} = \widetilde{\widetilde{\Xi}}^s \widetilde{Q}_{t|t} \text{ where } \widetilde{\widetilde{\Xi}} = \left[\begin{array}{c|cc} \Xi & \widehat{DAC} & \widehat{V}c\mathbf{H} \\ \hline & \Omega & C\mathbf{H} \\ & & \mathbf{H} \end{array}\right], \tag{IA19}$$

where $\widetilde{Q}_{t+s|t} = \left[Q'_{t+s|t}, q'_{t+s|t}, \pi'_{t+s|t}\right]'$. Armed with starting values $\widetilde{Q}_{t|t} = \left[Q'_{t|t}, q'_{t|t}, \pi'_{t|t}\right]'$, we can then compute (IA13) and (IA18) using (IA19). To obtain $\pi'_{t|t}$ recall that we assume $\mathbb{I}_t$ 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, $\pi^i_{t|t} \equiv P(\xi_t = i|\mathbb{I}_t)$ can assume only two values, 0 or 1. As a result $\pi'_{t|t}$ will be $(1,0)$ or $(0,1)$. Thus, given $Z_t \in \mathbb{I}_t$, $q'_{t|t} = \left[q^{1\prime}_{t|t}, q^{2\prime}_{t|t}\right]'$ with $q^i_{t|t} \equiv \mathbb{E}_t(Z_t|\xi_t = i)\,\pi^i_{t|t}$ will be $[Z'_t \cdot 1, Z'_t \cdot 0]'$ or $[Z'_t \cdot 0, Z'_t \cdot 1]'$. Analogously, $Q'_{t|t} = \left[Q^{1\prime}_{t|t}, Q^{2\prime}_{t|t}\right]'$ with $Q^i_{t|t} \equiv \varphi(\mathbb{E}_t(Z_t Z'_t|\xi_t = i))\,\pi^i_{t|t}$ will be $\left[\varphi(Z_t Z'_t \cdot 1)', \varphi(Z_t Z'_t \cdot 0)'\right]'$ or $\left[\varphi(Z_t Z'_t \cdot 0)', \varphi(Z_t Z'_t \cdot 1)'\right]'$.

B. Mean-Square Stability

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

$$Z_t = c_{\xi_t} + A_{1,\xi_t} Z_{t-1} + A_{2,\xi_t} Z_{t-2} + V_{\xi_t}\varepsilon_t,\ \varepsilon_t \sim N(0, I), \tag{IA20}$$

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} is an $n \times n$ matrix of coefficients for $l = 1,2$ and $V_{\xi_t} V'_{\xi_t}$ 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 $\overline{\mu}$ and an n^2-vector $\overline{M}$ such that:*

1) $\lim_{t\to\infty} \mathbb{E}_0[Z_t] = \overline{\mu}$

2) $\lim_{t\to\infty} \mathbb{E}_0 [Z_t Z_t'] = \overline{M}$

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 the innovation process ε_t is asymptotically covariance stationary, Costa, Fragoso, and Marques (2004) show that a multivariate Markov-switching model such as the one described by (IA20) 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 to establish mean-square stability of a process such as the one described by (IA20), it is enough to check mean-square 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 the following proposition.

PROPOSITION 3: *A Markov-switching process whose law of motion can be described by (IA20) 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 MS-VAR.

C. 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 $\xi_t = 1, ..., m$, 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

$$\mathbb{E}_i \left(Z_t \right) = \overline{\mu}_i = \left(I_n - A_i \right)^{-1} c_i, \tag{IA21}$$

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 $\mathbb{E}_i \left(Z_t \right)$ over a finite horizon. If there are no further switches, we can then expect the variables to fluctuate around $\mathbb{E}_i \left(Z_t \right)$. Therefore, the conditional steady states for the mean can also be thought of 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 $\mathbb{V}_i \left(Z_t \right)$ obtained imposing that a certain regime is in place forever:

$$\varphi \left(\mathbb{V}_i \left(Z_t \right) \right) = \left(I_{n^2} - A_i \otimes A_i \right)^{-1} \varphi \left(V_{\xi_t} V_{\xi_t}' \right), \tag{IA22}$$

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

VIII. Dynamic Macro-Finance Model: Macro Block

This section reports technical details about the macro-finance dynamic stochastic general equilibrium (MS-DSGE) model.

A. Constant-Gain Adaptive Learning

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

$$\pi_t = \alpha + \phi\pi_{t-1} + \eta_t. \tag{IA23}$$

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

$$\begin{aligned} R_t &= R_{t-1} + \gamma_t \left(x_{t-1}x'_{t-1} - R_{t-1} \right) \\ b_t &= b_{t-1} + \gamma_t R_t^{-1} x_{t-1} \left(\pi_t - b'_{t-1}x_{t-1} \right), \end{aligned} \tag{IA24}$$

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 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

$$R_t = 1 \text{ if } R_{t-1} = 1 \tag{IA25}$$

$$\alpha_t^m = \alpha_{t-1}^m + \gamma_t \left(\pi_t - \phi\pi_{t-1} - \alpha_{t-1}^m \right). \tag{IA26}$$

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 $R_t = R_{t-1} + \gamma_t \left(x_{t-1}x'_{t-1} - R_{t-1} \right) = R_t = R_{t-1} + \gamma_t \left(1 - R_{t-1} \right)$. Starting value for $R = R_0 => R_1 = R_0 + \gamma \left(1 - R_0 \right)$. Continuing to iterate, this converges to one no matter the value for R_0 as long as $0 < \gamma_t < 1$. Set $x_t = R_1 = 1$ in (IA24) to get (IA26).

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

$$\alpha_t^m = \alpha_{t-1}^m + \gamma \left(\pi_t - \phi \pi_{t-1} - \alpha_{t-1}^m \right). \tag{IA27}$$

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

Perceived trend inflation, $\overline{\pi}_t$, is defined as $\lim_{h\to\infty} \mathbb{E}_t^m \left(\pi_{t+h}\right)$. Observe that, since expectations obey the constant gain adaptive rule, $\overline{\pi}_t$ is not constant but varies with information at time t. This can be seen by taking expectations on both sides of equation (IA23) to find

$$\begin{aligned}
\overline{\pi}_t &= \lim_{h\to\infty} \mathbb{E}_t^m \left(\pi_{t+h}\right) \\
&= \lim_{h\to\infty} \mathbb{E}_t^m \left(\alpha_t^m + \phi \pi_{t+h-1}\right) \\
&= \lim_{h\to\infty} \mathbb{E}_t^m \left(\alpha_t^m + \phi \alpha_t^m + \phi^2 \alpha_t^m + \ldots \phi^{h-1} \alpha_t^m + \phi^h \pi_t\right) \\
&= \alpha_t^m / \left(1 - \phi\right),
\end{aligned}$$

where we plug in the value of α_t^m that agents perceive at t as the last step. Above we use 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, that is, 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 $\overline{\pi}_t$. Using the relation between $\overline{\pi}_t$ and α_t^m, we get

$$\begin{aligned}
\alpha_t^m &= \left(1 - \phi\right) \overline{\pi}_t => \\
\left(1 - \phi\right) \overline{\pi}_t &= \left(1 - \phi\right) \overline{\pi}_{t-1} + \gamma \left(\pi_t - \phi \pi_{t-1} - \left(1 - \phi\right) \overline{\pi}_{t-1}\right) => \\
\overline{\pi}_t &= \overline{\pi}_{t-1} + \gamma \left(1 - \phi\right)^{-1} \left(\pi_t - \phi \pi_{t-1} - \left(1 - \phi\right) \overline{\pi}_{t-1}\right),
\end{aligned} \tag{IA28}$$

where the second equation above follows from (IA27).

Finally, the unconditional mean of inflation as perceived by the agent is estimated recursively under the constant-gain adaptive rule and hence depends on the sample of data she uses at time t to estimate α. Denote this information by $\mathbb{I}_t$. Taking perceived unconditional means on both sides of (IA23), we find that the unconditional mean of inflation as perceived by the agent at time t is the same as perceived trend inflation:

$$\mathbb{E}^m\left(\pi_t|\mathbb{I}_t\right) = \alpha + \phi\mathbb{E}^m\left(\pi_t|\mathbb{I}_t\right) => \mathbb{E}^m_{T_t^m}\left(\pi_t\right) = \overline{\pi}_t = \alpha_t^m / \left(1-\phi\right).$$

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 shaped in part 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 $\overline{\pi}_t^{CG}$ to denote the beliefs implied by the constant-gain learning described above, we obtain modified updating rules for α_t^m and $\overline{\pi}_t$ that are weighted averages of two terms:

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

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

The first terms in square brackets, α_t^{mCG} and $\overline{\pi}_t^{CG}$, are the recursive updating rules implied by constant-gain learning as in (IA27) and (IA28). These terms are combined with two terms that involve the central bank's current inflation target $\pi_{\xi_t}^T$. Note that since $\alpha_t^m = \left(1-\phi\right)\overline{\pi}_t$ under the autoregressive model, the term $\left(1-\phi\right)\pi_{\xi_t}^T$ is simply the value of α_t^m that would arise if $\overline{\pi}_t = \pi_{\xi_t}^T$. 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 beliefs are a weighted average of what would arise under constant-gain learning and a term reflecting information about the current inflation target.

B. Expected Inflation

Expected inflation from the point of view of the agents in the model is formed based on equation (IA23) and their beliefs about the constant α, that is, 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 assume that agents do not take into account the possibility that their beliefs might change in the future (i.e., they do not have anticipated utility).

Cumulative inflation over the next year is

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

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

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

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

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

where once again we use $\overline{\pi}_t = \alpha_t^m / (1 - \phi)$ and we use the matrix $e_{\pi\overline{\pi}}$ to extract both inflation π_t and the perceived long-term inflation $\overline{\pi}_t$ from the state vector S_t. The latter formulation is used for the solution of the model since it is $\overline{\pi}_t$ rather than α_t^m that appears in the state-space representation of the model. It follows that

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

where the vector e_π is used to extract inflation.

C. 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:

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

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

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

where we use the fact that $\overline{\pi}_t = \mathbb{E}[\pi_t] = \pi_{\xi_t}^T$. We therefore have $\mathbb{E}\left[y_{t-1}^*\right] = \mathbb{E}[y_{t-1}] = 0$. 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,

$\overline{\pi}_t$, align with the econometrician's beliefs and with the central bank's inflation target only in the long run. However, even with sticky expectations, if the central bank permanently changes the target, we continue to have $\lim_{h\to\infty}\mathbb{E}_t\left[\overline{\pi}_{t+h}\right]=\mathbb{E}\left[\overline{\pi}_t\right]=\pi^T_{\xi_t}=\mathbb{E}\left[\pi_t\right]$, where $\mathbb{E}_t\left[\cdot\right]$ denotes the expectations of the econometrician. Thus,

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

We therefore again have $\mathbb{E}\left[y^*_{t-1}\right]=\mathbb{E}\left[y_{t-1}\right]=0$.

D. Solution and Estimation of the Macro Block

We can rewrite the system of equations as

$$\widetilde{y}_t = \varrho\widetilde{y}_{t-1}-\sigma\left[i_t-\phi\pi_t-(1-\phi)\overline{\pi}_t-r\right]+f_t \tag{IA29}$$

$$\pi_t = \overline{\pi}_t+\frac{\kappa}{1-\beta\phi}\left[y_{t-1}-y^*_{t-1}\right] \tag{IA30}$$

$$\begin{aligned}
i_t-\left(r+\pi^T_{\xi_t}\right) &= (1-\rho_{i,\xi_t})\left[\psi_{\pi,\xi_t}\left(\pi_t-\pi^T_{\xi_t}\right)+\psi_{\Delta y,\xi_t}\left(\widetilde{y}_t-\widetilde{y}_{t-1}\right)\right]\\
&\quad+\rho_{i,\xi_t}\left[i_{t-1}-\left(r+\pi^T_{\xi_t}\right)\right]+\sigma_i\varepsilon_{i,t}
\end{aligned} \tag{IA31}$$

$$\widetilde{y}^*_t = \rho_{y^*}\widetilde{y}^*_{t-1}+\sigma_{y^*}\varepsilon_{y^*,t} \tag{IA32}$$

$$\begin{aligned}
\overline{\pi}_t &= \left[1-\gamma^T\right]\left[\overline{\pi}_{t-1}+\gamma(1-\phi)^{-1}\left(\pi_t-\phi\pi_{t-1}-(1-\phi)\overline{\pi}_{t-1}\right)\right]\\
&\quad+\gamma^T\pi^T_{\xi_t}
\end{aligned} \tag{IA33}$$

$$f_t = \rho_f f_{t-1}+\sigma_f\varepsilon_{f,t}. \tag{IA34}$$

Define the parameter vectors θ_{ξ_t} and $\theta^c_{\xi_t}$ as

$$\begin{aligned}
\theta_{\xi_t} &= \left[\varrho,\sigma,\beta,\kappa,\psi_{\pi,\xi_t},\psi_{\Delta y,\xi_t},\rho_{i,\xi_t},\rho_{y^*},\gamma^T,\gamma,\phi,\rho_f\right]'\\
\theta^c_{\xi_t} &= \left[\pi^T_{\xi_t},r\right]'.
\end{aligned}$$

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

$$\begin{aligned} S_t &= [\widetilde{y}_t, \widetilde{y}_t^*, \pi_t, i_t, \overline{\pi}_t, f_t]' \\ \varepsilon_t &= \left[\varepsilon_{i,t}, \varepsilon_{y_t^*}, \varepsilon_{f,t}\right]', \ \varepsilon_t \sim N(0, I) \end{aligned}$$

Let the matrix $Q = diag(\sigma_i, \sigma_{y^*}, \sigma_d)$ 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:

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

Note that the vector $\Gamma_c\left(\theta_{\xi_t}^c\right)$ 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:

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

where $C\left(\theta_{\xi_t}^c, \theta_{\xi_t}\right) = \Gamma_0^{-1}(\theta_{\xi_t}) \Gamma_c\left(\theta_{\xi_t}^c\right)$, $T(\theta_{\xi_t}) = \Gamma_0^{-1}(\theta_{\xi_t}) \Gamma_1(\theta_{\xi_t})$, 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 (IA36), below, varies over time. We thus have the following state-space representation:

$$X_t = D + Z\left[S_t', y_{t-1}\right]' + Uv_t \tag{IA35}$$

$$S_t = C\left(\theta_{\xi_t}^c, \theta_{\xi_t}\right) + T(\theta_{\xi_t})S_{t-1} + R(\theta_{\xi_t})Q\varepsilon_t \tag{IA36}$$

$$v_t \sim N(0, I), \tag{IA37}$$

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 (IA36) 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 three shocks for four observables.

Thus, the vector of data X_t is defined as

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

where in the last row we use the fact that expectations for an agent in the model are given by

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

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 thus

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

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

The matrix $\overline{\Sigma}$ corresponds to the inverse of the Hessian computed at the posterior mode $\overline{\theta}$. The parameter c is set to obtain an acceptance rate of around 30%. We use four chains of $540,000$ draws each (one in every 200 draws is saved). Convergence is checked by using the Brooks-Gelman-Rubin potential reduction scale factor using within and between variation 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 (IA36) 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 as follows.

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

1. Prediction:

$$S_{t|t-1} = C\left(\theta^c_{\xi_t}, \theta_{\xi_t}\right) + T(\theta_{\xi_t})S_{t-1|t-1} \tag{IA38}$$

$$P_{t|t-1} = T(\theta_{\xi_t})P_{t-1|t-1}T(\theta_{\xi_t})' + R(\theta_{\xi_t})Q^2R(\theta_{\xi_t})' \tag{IA39}$$

$$\eta_{t|t-1} = X_t - X_{t|t-1} = X_t - D - Z * S_{t|t-1} \tag{IA40}$$

$$f_{t|t-1} = ZP_{t|t-1}Z' + U^2. \tag{IA41}$$

2. Updating:

$$S_{t|t} = S_{t|t-1} + K_t\eta_{t|t-1} \tag{IA42}$$

$$P_{t|t} = P_{t|t-1} - K_tZP_{t|t-1}, \tag{IA43}$$

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

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

$$\ln L = -.5\sum_{t=1}^{T}\ln\left(2\pi f_{t|t-1}\right) - 0.5\sum_{t=1}^{T}\eta'_{t|t-1}f^{-1}_{t|t-1}\eta_{t|t-1}.$$

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

Equations:

$$\begin{aligned}
y_t &= \varrho y_{t-1} - \sigma\left[i_t - \phi\pi_t - (1-\phi)\,\overline{\pi}_t - r\right] + f_t \\
\pi_t &= \overline{\pi}_t + \frac{\kappa}{1-\beta\phi}\left[y_{t-1} - y^*_{t-1}\right] && \text{(IA44)} \\
i_t - \left(r + \pi^T_{\xi_t}\right) &= \left(1-\rho_{i,\xi_t}\right)\left[\psi_{\overline{\pi},\xi_t}\left(\overline{\pi}_t - \pi^T_{\xi_t}\right) + \psi_{\pi,\xi_t}\left(\pi_t - \pi^T_{\xi_t}\right) + \psi_{\Delta y,\xi_t}\left(y_t - y_{t-1}\right)\right] \\
&\quad + \rho_{i,\xi_t}\left[i_{t-1} - \left(r + \pi^T_{\xi_t}\right)\right] + \sigma_i\varepsilon_{i,t} && \text{(IA45)} \\
y^*_t &= \rho_{y^*}y^*_{t-1} + \sigma_{y^*}\varepsilon_{y^*,t} && \text{(IA46)} \\
\overline{\pi}_t &= \left[1-\gamma^T\right]\left[\overline{\pi}_{t-1} + \gamma\left(1-\phi\right)^{-1}\left(\pi_t - \phi\pi_{t-1} - (1-\phi)\,\overline{\pi}_{t-1}\right)\right] \\
&\quad + \gamma^T\pi^T_{\xi_t} + \sigma_{\overline{\pi}}\varepsilon_{\overline{\pi},t} && \text{(IA47)} \\
f_t &= \rho_f f_{t-1} + \sigma_f\varepsilon_{f,t}. && \text{(IA48)}
\end{aligned}$$

We get

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

Equations with state variables at t on the left-hand side, everything else on the right-hand side, and reordered to match the state variable vector are as follows:

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

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

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

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

First, Γ_0, which corresponds to the time t state variables on the left-hand side. Empty cells are zero.

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

Next, Γ_1, which corresponds to the time $t-1$ state variables on the right-hand. Empty cells are zero.

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

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

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

Finally, Γ_C collects all of the leftover constant terms on the right-hand side.

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

IX. 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 asset pricing (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 as to whether such a policy rule will persist for a long time or not.

Agents in the model are fully rational, conduct Bayesian learning, and know that they do not know. Therefore, when forming expectations, agents take into account the fact that their beliefs will evolve according to what they 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 differ only 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 data-generating process (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 over-extrapolate what this implies for future monetary policy and memories of previous regimes might fade away. Finally, consistent 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, coming to consider the regime change as a structural one only after spending enough time in the new regime.

A. Beliefs: Overview

The policy rule follows two regimes, $\xi_t = H$ for hawkish and $\xi_t = D$ for dovish. We assume that the asset pricing agent observes all variables of the economy in the current 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 $\widetilde{S}_t = [S_t, m_t, pd_t, \mathbb{E}_t^p(m_{t+1}), \mathbb{E}_t^p(pd_{t+1})]'$. Suppose first that agents can observe the monetary policy regime in place and that they form expectations based on the transition matrix $\mathbf{H}$ of the true DGP transitions across the two policy regimes. In this case, the model can be expressed as

$$\Gamma_{0,\xi_t}\widetilde{S}_t = \Gamma_{c,\xi_t} + \Gamma_{1,\xi_t}\widetilde{S}_{t-1} + \Psi_{\xi_t}\varepsilon_t + \Pi\eta_t, \tag{IA49}$$

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

$$\mathbf{H}=\left[\begin{array}{cc} p_{HH} & p_{HD} \\ p_{DH} & p_{DD} \end{array}\right],$$

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 AP 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, $\theta\left(\xi_t^p\right)$. 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: $b_1 = \{\xi_t^p \in \{1, 2\} : \theta\left(\xi_t\right) = \theta_{b_1}\}$, characterized by hawkish monetary policy ($\xi_t = 1$), while regimes $\xi_t^p = 3$ and $\xi_t^p = 4$ belong to a block 2: $b_2 = \{\xi_t^p \in \{3, 4\} : \theta\left(\xi_t\right) = \theta_{b_2}\}$, 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 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

$$\mathbf{H}^p = \left[\begin{array}{cc|cc} p_{11} & 0 & 0 & p_{14} \\ 0 & p_{22} & p_{23} & p_{24} \\ \hline 0 & p_{32} & p_{33} & 0 \\ p_{41} & p_{42} & 0 & p_{44} \end{array}\right], \tag{IA50}$$

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 time in a regime, they become convinced that this regime will persist indefinitely.[4]

Suppose that the economy is initially in a state in which 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 the $\xi_t^p = 2$ 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

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

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 long enough, 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. We therefore 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 probability 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, that is, beliefs about the persistence of the current monetary policy regime ξ_t. We then define an expanded set of regimes that keep track of both 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 proceed in two steps. First, we characterize the evolution of agents' beliefs *within* a block for given prior beliefs. This allows us to track the evolution of beliefs as agents observe more periods of the same policy rule regime. Second, we explain how agents' beliefs are pinned down once the economy moves *across* blocks. This allows us to characterize agents' beliefs when agents observe a change in the conduct of monetary policy. All results are based on Bayes' theorem. Finally, for each of these cases, we 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.

B. Evolution of Beliefs Within a Block

In what follows, we derive the law of motion of agents' beliefs conditional on being in a specific block, that is, on observing a certain policy rule. The formulas derived below 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 the moment the system enters the new block, that is, the moment agents observe a policy regime that is different from the one observed in the previous period. We characterize these initial beliefs in the next subsection.

As we note 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 $prob\left(\xi_{t-1}^p = 2|\tau_{t-1}^1\right) = 1 - prob\left(\xi_{t-1}^p = 1|\tau_{t-1}^1\right)$, the probability of being in regime 1 given that we have observed τ_t^1 consecutive realizations of block 1, $prob\left(\xi_t^p = 1|\tau_t^1\right)$, is given by

$$\begin{aligned}
prob\left(\xi_t^p = 1|\tau_t^1\right) &= \frac{prob\left(\xi_{t-1}^p = 1|\tau_{t-1}^1\right) p_{11}}{prob\left(\xi_{t-1}^p = 1|\tau_{t-1}^1\right) p_{11} + prob\left(\xi_{t-1}^p = 2|\tau_{t-1}^1\right) p_{22}} \\
&= \frac{prob\left(\xi_{t-1}^p = 1|\tau_{t-1}^1\right) p_{11}}{prob\left(\xi_{t-1}^p = 1|\tau_{t-1}^1\right)\left(p_{11} - p_{22}\right) + p_{22}}, \qquad \text{(IA51)}
\end{aligned}$$

where $\tau_t^1 = \tau_{t-1}^1 + 1$ and for $\tau_t^1 > 1$. Notice that for $\tau_t^1 = 1$, $prob\left(\xi_t^p = 1|\tau_t^1\right)$ denotes the initial beliefs that we discuss in Section IX.C. Equation (IA51) 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 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, $prob\left(\xi_t = 3|\tau_t^2\right)$, can be analogously derived:

$$\begin{aligned}
prob\left(\xi_t^p = 3|\tau_t^2\right) &= \frac{prob\left(\xi_{t-1}^p = 3|\tau_{t-1}^2\right) p_{33}}{prob\left(\xi_{t-1}^p = 3|\tau_{t-1}^2\right) p_{33} + prob\left(\xi_{t-1}^p = 4|\tau_{t-1}^2\right) p_{44}} \\
&= \frac{prob\left(\xi_{t-1}^p = 3|\tau_{t-1}^2\right) p_{33}}{prob\left(\xi_{t-1}^p = 3|\tau_{t-1}^2\right)\left(p_{33} - p_{44}\right) + p_{44}}. \qquad \text{(IA52)}
\end{aligned}$$

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

The recursive equations (IA51) and (IA52) characterize the dynamics of agents' beliefs in both blocks for a given set of prior beliefs. Bianchi and Melosi (2016) 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 by 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 $prob\left(\xi_t^p = 1|\tau^1\right)$ and $prob\left(\xi_t^p = 3|\tau^2\right)$ by $\widetilde{\lambda}_{b_1}$and $\widetilde{\lambda}_{b_2}$, respectively. The converging probabilities for the respecive long-lasting regimes are then $1 - \widetilde{\lambda}_{b_1}$and $1 - \widetilde{\lambda}_{b_2}$, 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.

C. Evolution of Beliefs Across Blocks

In the previous subsection, we characterize the evolution of agents' beliefs conditional on being in a specific block, that is, on observing additional realizations of the same policy rule. The formulas derived above apply to *any set of initial beliefs*. In this subsection, we pin down agents' beliefs at the moment the economy moves *across* blocks, that is, 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 serve as starting points for the recursions (IA51) and (IA52) 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, 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

$$prob\left(\xi_t^p = 1 | \xi_{t-1}^p = 3, \tau_t^1 = 1\right) = \frac{p_{31}}{p_{31} + p_{32}} = 0 \tag{IA53}$$

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

$$prob\left(\xi_t^p = 1 | \xi_{t-1}^p = 4, \tau_t^1 = 1\right) = \frac{p_{41}}{p_{41} + p_{42}} = 1 \tag{IA54}$$

if the economy was under regime 4 (long-lasting dovish) 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

$$prob\left(\xi_t^p = 3 | \xi_{t-1}^p = 1, \tau_t^2 = 1\right) = \frac{p_{13}}{p_{13} + p_{14}} = 0 \tag{IA55}$$

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

$$prob\left(\xi_t^p = 3 | \xi_{t-1}^p = 2, \tau_t^2 = 1\right) = \frac{p_{23}}{p_{23} + p_{24}} = 1 \tag{IA56}$$

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

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

$$prob\left\{\xi_t^p = 1|\mathcal{I}_t\right\} = \frac{\left(1 - prob\left\{\xi_{t-1}^p = 3|\mathcal{I}_{t-1}\right\}\right) p_{41}}{prob\left\{\xi_{t-1}^p = 3|\mathcal{I}_{t-1}\right\} p_{32} + \left(1 - prob\left\{\xi_{t-1}^p = 3|\mathcal{I}_{t-1}\right\}\right)(p_{41} + p_{42})}, \tag{IA57}$$

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

$$prob\left\{\xi_t^p = 3|\mathcal{I}_t\right\} = \frac{\left(1 - prob\left\{\xi_{t-1}^p = 1|\mathcal{I}_{t-1}\right\}\right) p_{23}}{prob\left\{\xi_{t-1}^p = 1|\mathcal{I}_{t-1}\right\} p_{14} + \left(1 - prob\left\{\xi_{t-1}^p = 1|\mathcal{I}_{t-1}\right\}\right)(p_{23} + p_{24})}, \tag{IA58}$$

where $\mathcal{I}_t$ includes the history of policy regimes (blocks) up to time t. Because the above are recursive formulations, we find 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 agents' initial beliefs when the system entered the current block.

D. 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 of grid points approaches infinity, the approximation becomes arbitrarily accurate. Note that for each point in the grid, (IA57) and (IA58) tell us how beliefs will evolve if we observe a change in the conduct of monetary policy, while (IA51) and (IA52) 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 of both 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 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 $\mathcal{G}_{b_1} = \{\mathcal{G}_1, ..., \mathcal{G}_{g_1}\}$ and for beliefs $prob\{\xi_t^p = 3|\mathcal{I}_t\}$ as $\mathcal{G}_{b_2} = \{\mathcal{G}_{g_1+1}, ..., \mathcal{G}_{g_1+g_2}\}$, where $0 \leq \mathcal{G}_i \leq 1$, for 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 (IA51) and (IA52) and the initial conditions (IA57) and (IA58). Denote this expanded perceived transition matrix $\widehat{\mathbf{H}}^p$. The algorithm below illustrates how exactly to perform this task.

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

Step 1 For each of the two blocks, employ 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

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

where $prob\left\{\xi_{t-1}^p = 1|\mathcal{I}_{t-1}\right\} = \mathcal{G}_i$ and $j \leq g_1$ is set so as to min $\left|prob\left\{\xi_t^p = 1|\mathcal{I}_t\right\} - \mathcal{G}_j\right|$, where $prob\left\{\xi_t^p = 1|\mathcal{I}_t\right\}$ is computed using the recursive equation (IA51) by approximating $prob\left\{\xi_{t-1}^p = 1|\mathcal{I}_{t-1}\right\} = \mathcal{G}_i$. To ensure the convergence of beliefs, we correct j as follows: if $j = i$ and $\mathcal{G}_i \neq \widetilde{\lambda}_{b_1}$, then set $j = \min(j+1, g_1)$ if $\mathcal{G}_i < \widetilde{\lambda}_{b_1}$ or $j = \max(1, j-1)$ if $\mathcal{G}_i > \widetilde{\lambda}_{b_1}$.

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

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

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

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

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

Step 4 Write the transition equation as $\widetilde{H^p} = T \cdot \widehat{\mathbf{H}}^p \cdot T'$. If no column of $\widetilde{H^p}$ has all zero elements, set $\widehat{\mathbf{H}}^p = \widetilde{H^p}$ 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 $\widehat{\mathbf{H}}^p$ 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 $\widetilde{H^p}$. Steps 2 to 4 are not necessary but help 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., $\widetilde{\lambda}_{b_1}$ 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

$$\widehat{H}_{ij} = \Pr\left(\zeta_{t+1} = j | \zeta_t = i\right). \tag{IA59}$$

X. 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. Next, 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 $\widehat{\mathbf{H}}^p$, 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 $\widehat{\mathbf{H}}^p$. 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

$$\begin{aligned}
m_t &= \log(\delta) - \sigma_p e_c (S_t - S_{t-1}) + \vartheta_{p,t-1} && \text{(IA60)} \\
e_i S_t - \mathbb{E}_t^p [e_\pi S_{t+1}] &= -\mathbb{E}_t^p [m_{t+1}] \underbrace{-.5\mathbb{V}_t^p [m_{t+1} + e_i S_t - e_\pi S_{t+1}]}_{\text{Regime-dependent risk adjustment}} - lp && \text{(IA61)} \\
pd_t &= \kappa_0 + \mu + \underbrace{[.5\mathbb{V}_t^p [m_{t+1} + e_c (S_{t+1} - S_t) + \kappa_1 pd_{t+1}]]}_{\text{Regime-dependent risk adjustment}} && \text{(IA62)} \\
&\quad + \mathbb{E}_t^p [m_{t+1} + e_c (S_{t+1} - S_t) + \kappa_1 pd_{t+1}] \\
\eta_t^{pd} &= pd_t - \mathbb{E}_{t-1}^p (pd_t) && \text{(IA63)} \\
\eta_t^m &= m_t - \mathbb{E}_{t-1}^p (m_t) && \text{(IA64)} \\
S_t &= \mathbf{C}_{\xi_t} + \mathbf{T}_{\xi_t} S_{t-1} + \mathbf{R}_{\xi_{t+1}} \mathbf{Q} \varepsilon_t, && \text{(IA65)}
\end{aligned}$$

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 such that it 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 X.A below entitled "Solving the model with a risk adjustment" below provides details on lognormality in a setting with regime shifts.) Define the augmented state space as $\widetilde{S}_t = [S_t, m_t, pd_t, \mathbb{E}_t^p (m_{t+1}), \mathbb{E}_t^p (pd_{t+1})]$. For any variable x_t in the asset pricing block, conditional lognormality assumption implies

$$\mathbb{V}_t^p [x_{t+1}] \approx \mathbb{E}_t^p (\mathbb{V}_t [x_{t+1} | \xi_{t+1}]) = e_x \mathbb{E}_t^p \left[\mathbf{R}_{\xi_{t+1}} \mathbf{Q} \mathbf{Q}' \mathbf{R}'_{\xi_{t+1}}\right] e_x, \tag{IA66}$$

where e_x is a vector used to extract the desired linear combination of the variables in S_t. This assumption maintains conditional lognormality of the entire system and guarantees that 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 of both the behavior of the 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 (IA66), the one-step-ahead variance $\mathbb{V}_t^p\left[x_{t+1}\right]$ is a function only of the expanded regimes at time t, ζ_t. This leaves us with a new system to be solved, given by

$$\Gamma_0\left(\zeta_t\right)\widetilde{S}_t = \Gamma_c\left(\zeta_t\right) + \Gamma_1\left(\zeta_t\right)\widetilde{S}_{t-1} + \Psi\left(\zeta_t\right)\mathbf{Q}\varepsilon_t + \Pi\eta_t, \tag{IA67}$$

where the regime $\zeta_t \in \{1, ..., g_1 + g_2\}$ follows the transition matrix $\widehat{\mathbf{H}}^p$ and the terms $\Gamma_c\left(\zeta_t\right)$ now also contain the regime-specific risk adjustment terms $\mathbb{V}_t^p\left[x_{t+1}\right]$ that are part of the asset pricing block. Note that $\Gamma_c\left(\zeta_t\right)$ depends on $\mathbb{V}_t^p\left[x_{t+1}\right]$ as given in (IA66). For variables in the system (IA60) to (IA65) expressed in recursive form, like pd_t, the vector e_x is not known until we solve for $\widetilde{S}_t$. 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 once again takes the form of a MS-VAR:

$$\widetilde{S}_t = \widetilde{\mathbf{C}}\left(\zeta_t, \widehat{\mathbf{H}}^p\right) + \widetilde{\mathbf{T}}\left(\zeta_t, \widehat{\mathbf{H}}^p\right)S_{t-1} + \widetilde{\mathbf{R}}\left(\zeta_t, \widehat{\mathbf{H}}^p\right)\mathbf{Q}\varepsilon_t. \tag{IA68}$$

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

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

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

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

Let the model solution for the price-dividend ratio be denoted by $pd_t = pd\left(\widetilde{S}_t\right)$, 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:

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

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

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

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

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

where we use the output in (IA59) 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 $\widehat{\mathbf{H}}^p$.

A. Solving a Model with Risk Adjustment

This section provides more details on 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 lognormal 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 lognormality does not affect the rest of the log-linearized equations. When introducing regime changes, the process becomes conditionally lognormal, 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.

B. Conditional Lognormality

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

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

Then

$$\ln\left(\mathbb{E}_t[Z_{t+1}]\right) = \mathbb{E}_t[z_{t+1}] + 0.5\mathbb{V}_t[z_{t+1}] = c + az_t + 0.5\sigma^2. \tag{IA69}$$

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

$$z_{t+1} = c_{\xi_{t+1}} + a_{\xi_{t+1}} z_t + \sigma_{\xi_{t+1}}\varepsilon_{t+1}, \tag{IA70}$$

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 face log-linearization equations of the form

$$\mathbb{E}_t[e^{z_{t+1}}]. \tag{IA71}$$

We extend the approach in Bansal and Zhou (2002), who use conditional lognormality of the process in equation (IA70). Conditioning on the regime in the next period, lognormality 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 $e^{c_{\xi_{t+1}}+az_t+0.5\sigma^2_{\xi_{t+1}}} \approx 1+c_{\xi_{t+1}}+a_{\xi_{t+1}}z_t+0.5\sigma^2_{\xi_{t+1}}$. With this approximation, we have

$$\mathbb{E}_t[e^{z_{t+1}}] = \mathbb{E}_t[\mathbb{E}_t[e^{z_{t+1}}|\xi_{t+1}]] \approx \mathbb{E}_t[1+c_{\xi_{t+1}}+a_{\xi_{t+1}}z_t+0.5\sigma^2_{\xi_{t+1}}] = \tag{IA72}$$

$$= 1+\mathbb{E}_t[c_{\xi_{t+1}}+a_{\xi_{t+1}}z_t]+0.5\mathbb{E}_t[\sigma^2_{\xi_{t+1}}]. \tag{IA73}$$

We thus obtain

$$\ln(\mathbb{E}_t[Z_{t+1}]) \approx \mathbb{E}_t[c_{\xi_{t+1}}+a_{\xi_{t+1}}z_t]+0.5\mathbb{E}_t[\sigma^2_{\xi_{t+1}}], \tag{IA74}$$

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

Above we make 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 $\widetilde{z}_{t+1} \equiv z_{t+1}-z$ is the log-deviation of Z_{t+1} from its mean. We then have

$$\begin{aligned}e^z\mathbb{E}_t[e^{\widetilde{z}_{t+1}}] &= e^z\mathbb{E}_t[\mathbb{E}_t[e^{\widetilde{z}_{t+1}}|\xi_{t+1}]] && \text{(IA75)}\\ &= e^z\mathbb{E}_t[e^{\mathbb{E}_t[\widetilde{z}_{t+1}|\xi_{t+1}]+0.5\mathbb{V}_t[\widetilde{z}_{t+1}|\xi_{t+1}]}] && \text{(IA76)}\\ &\approx e^z\mathbb{E}_t[1+\mathbb{E}_t[\widetilde{z}_{t+1}|\xi_{t+1}]+0.5\mathbb{V}_t[\widetilde{z}_{t+1}|\xi_{t+1}]] && \text{(IA77)}\\ &= e^z[1+\mathbb{E}_t[\widetilde{z}_{t+1}]+0.5\mathbb{E}_t[\sigma^2_{\xi_{t+1}}]], && \text{(IA78)}\end{aligned}$$

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

$$\begin{aligned}
\log\left(\mathbb{E}_t[Z_{t+1}]\right) &= \log\left(e^{z}\mathbb{E}_t[e^{\widetilde{z}_{t+1}}]\right) && \text{(IA79)}\\
&\approx z + \mathbb{E}_t[\widetilde{z}_{t+1}] + 0.5\mathbb{E}_t[\sigma^2_{\xi_{t+1}}] && \text{(IA80)}\\
&= \mathbb{E}_t[z_{t+1}] + 0.5\mathbb{E}_t[\sigma^2_{\xi_{t+1}}] && \text{(IA81)}\\
&= \mathbb{E}_t[c_{\xi_{t+1}} + a_{\xi_{t+1}} z_t] + 0.5\mathbb{E}_t[\sigma^2_{\xi_{t+1}}]. && \text{(IA82)}
\end{aligned}$$

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

$$\begin{aligned}
P^D_t &= \mathbb{E}^p_t\left[M_{t+1}\left(P^D_{t+1} + D_{t+1}\right)\right]\\
\frac{P^D_t}{D_t} &= \mathbb{E}^p_t\left[M_{t+1}\frac{D_{t+1}}{D_t}\left(\frac{P^D_{t+1}}{D_{t+1}} + 1\right)\right].
\end{aligned}$$

Taking logs on both sides, we get

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

Applying the approximation implied by conditional lognormality, we have

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

where under the conditional lognormality approximation we have

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

XI. Estimating the Dynamic Macro-Finance Model

As explained in Section VIII.D above, 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, $\overline{er}$, 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 objective function:

$$\widehat{\theta}_p = \arg\min\left[\sum_{t=1}^{T}\left(cay_t^{MS} - cay_t^m\left(\theta_p, X^T, \xi^T\right)\right)^2 + 0.05\left(\left|\overline{er}\left(\theta_p, X^T, \xi^T\right) - 6\right|\right)\right],$$

where cay_t^m and the annualized average equity premium $\overline{er}$ 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. We thus ask the model to explain as much of the observed variation in $-cay_t^{MS}$ out of regime changes as possible.

XII. 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 $\widetilde{S}_t$ according to the linear transformation

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

In doing this, the econometrician can use the transition matrix reflecting the actual frequency of regime changes or the transition matrix used by the AP agent when forming expectations. In the first case, we obtain the actual path of the PDV of excess returns based on the DGP, 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

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

where we omit 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 AP 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 $\widehat{\mathbf{H}}^p$, 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 by $\widehat{\mathbf{H}}$.

As above, define the column vectors q_t and π_t,

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

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

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

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

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

$$\underbrace{premia_{z,t}}_{\text{Premia}}\equiv\underbrace{\sum_{j=0}^{\infty}\rho^j\mathbb{E}_t r_{t+1+j}}_{\text{PDV of excess returns}}=1_r'w(I-\rho\Omega)^{-1}\left[\Omega q_{t|t}+C\left(I-\rho\widehat{\mathbf{H}}\right)^{-1}\widehat{\mathbf{H}}\pi_{t|t}\right], \quad \text{(IA83)}$$

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.

XIII. 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 reestimate 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 FFR 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 results 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 IA.1 to IA.12.

In the Wu and Xia (2016) model, the short-term interest rate is the maximum of the shadow FFR 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 FFR 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 FFR to 25 to 50 basis points. when the lower bound is no longer binding, the shadow rate coincides with the actual FFR.

The results of Wu and Xia (2016) 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 nonlinearity 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 (2016) 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.

XIV. Inflation Target in the Early- and Late-Dovish Subperiods

This appendix shows that the early-dovish and late-dovish subperiods both rationalize a high value for $\pi^T_{\xi_t}$, but for different reasons.

To demonstrate this, we compute an alternative "third-regime" case of the model. The third-regime case has the same regime sequence as the baseline case, but we reestimate the model allowing the policy parameters to differ in all three regime subperiods. This third-regime case implies that perceived trend inflation $\overline{\pi}_t$ is 9.2% at the end of the post-millennial subsample, which is higher than in the baseline two-regime case. Why? After all, observed inflation is quite low at the end of our sample (almost zero) and had been trending down.

Perceived trend inflation is high at the end of our sample because the policy rule target inflation parameter $\pi^T_{\xi_t}$ is high, about 11.5%, in the dovish regime. The parameter $\overline{\pi}_t$ is mechanically linked to $\pi^T_{\xi_t}$ because, in addition to forming expectations with an adaptive rule, the macro agent receives a noisy signal about the inflation target $\pi^T_{\xi_t}$. Since these signals accumulate over time, $\overline{\pi}_t$ will increasingly approach $\pi^T_{\xi_t}$ the longer a regime subperiod lasts. Why is the inflation target high at the end of the sample? The policy rule parameter $\pi^T_{\xi_t}$ that drives $\overline{\pi}_t$ to be high at the end of the sample reflects the historical data we have used to pin down inflation, inflation expectations, and GDP growth. The post-millennial period is characterized by large 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. At the same time, data on on inflation expectations, which the model matches well, remain persistently high in relative terms throughout the post-millennial period. Expected inflation as measured by the SOC survey remains above 3% at the end of our sample in 2017:Q3 and was above 4% only a few years earlier. This set of facts is rationalized in the model by a central bank that changed objectives to implicitly implement extremely dovish policy in the post-millennial period in order to counteract two recessions, which in the stylized model of monetary policy shows up as a high value for target inflation, $\pi^T_{\xi_t}$. Had the implicit target not been this high, inflation expectations would have been lower than implied by survey evidence.

This shows that, in the baseline two-regime case, the early-dovish and late-dovish subperiods both rationalize a high value for $\pi^T_{\xi_t}$, but for different reasons. In the early subperiod, observed inflation and inflation expectations were commensurately high, which the model rationalizes with a high value for $\pi^T_{\xi_t}$. In the late subperiod actual inflation is quite low, but a persistent gap opens up in between expected inflation and actual inflation, which the model again rationalizes with a high value for $\pi^T_{\xi_t}$. To verify that this gap between expected inflation and actual inflation is indeed the source of the high value for $\pi^T_{\xi_t}$ in the post-millennial dovish subperiod, we reestimated the third-regime model using only GDP growth, inflation, and the FFR, while excluding inflation expectations. In this case the estimated value for $\pi^T_{\xi_t}$ and $\overline{\pi}_t$ at the end of our sample is much lower, with the latter only 1.2%.

REFERENCES

Bansal, Ravi, and Hao Zhou, 2002, Term structure of interest rates with regime shifts, *The Journal of Finance* 57, 1997–2043.

Bianchi, Francesco, 2016, Methods for measuring expectations and uncertainty in Markov-switching models, *Journal of Econometrics* 190, 79–99.

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

Black, Fischer, 1995, Interest rates as options, *Journal of Finance* 50, 1371–1376.

Costa, Osvaldo L.V., Marcelo D. Fragoso, and Ricardo P. Marques, 2004, *Discrete-Time Markov Jump Linear Systems* (Springer, New York).

Evans, George W., and Seppo Honkapohja, 2001, *Learning and Expectations In Macroeconomics* (Princeton University Press, Princeton, NJ).

Geweke, John F., 1992, Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments, *Bayesian Statistics* 4, 169–193.

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

Kim, Chang-Jin, and Charles R. Nelson, 1999, *State-Space Models with Regime Switching* (The MIT Press, Cambridge, Massachusetts).

Raftery, Adrian E., and Steven M. Lewis, 1992, How many iterations in the Gibbs sampler?, *Bayesian Statistics* 4, 763–773.

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

Stock, James H., and Mark W. Watson, 1993, A simple estimator of cointegrating vectors in higher order integrated systems, *Econometrica* 61, 783–820.

Wu, Jing C., and Fan D. Xia, 2016, Measuring the macroeconomic impact of monetary policy at the zero lower bound, *Journal of Money, Credit and Banking* 48, 253–291.

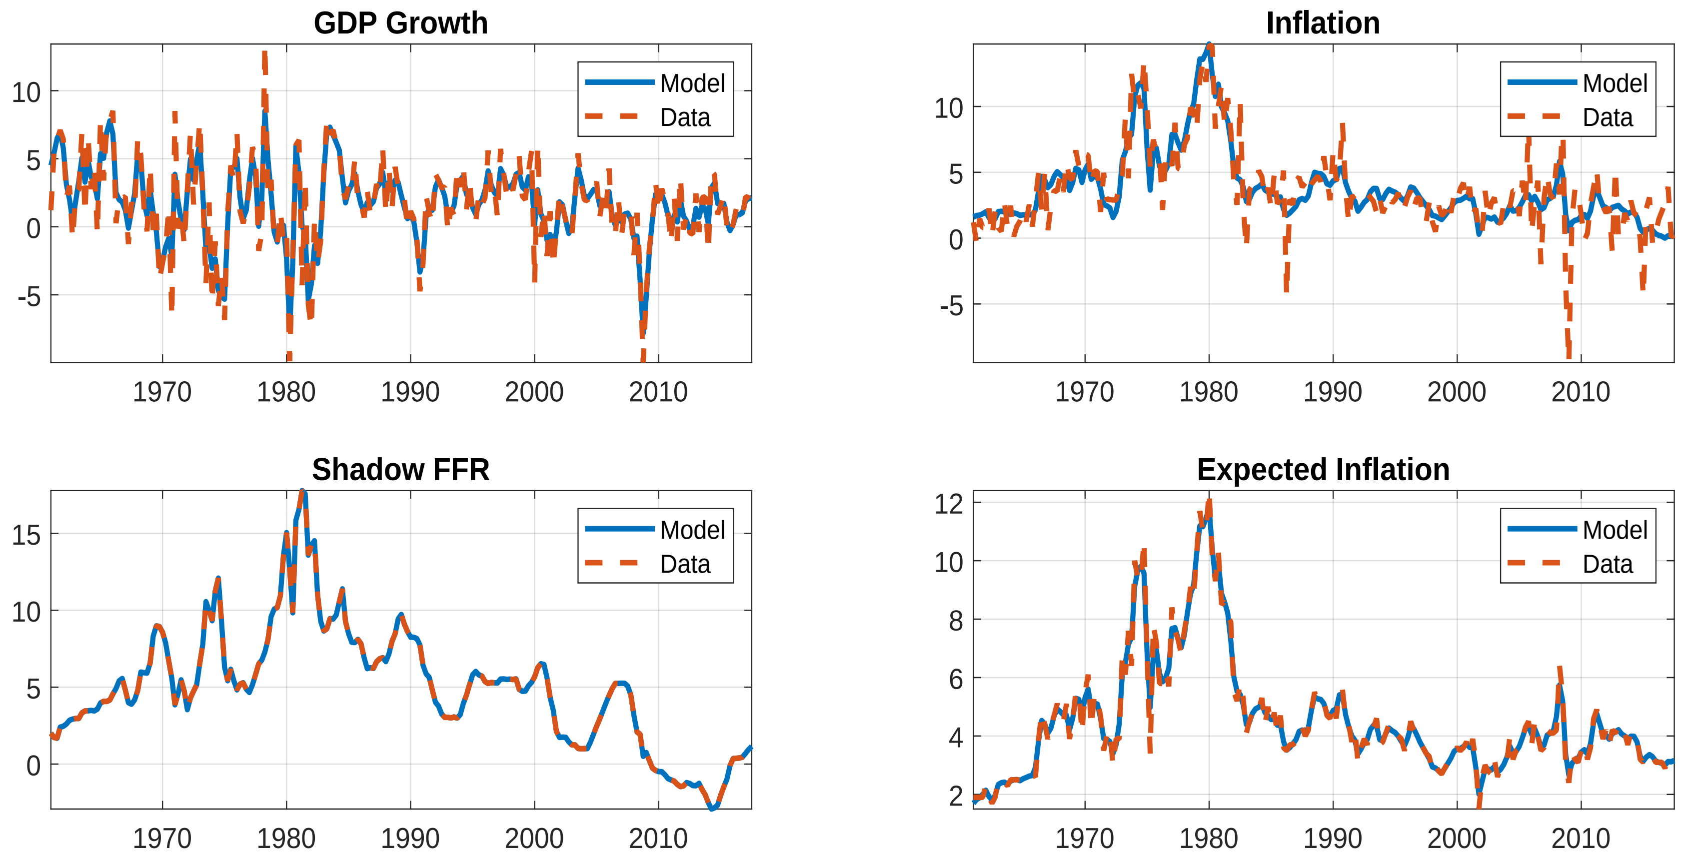

Figure IA.1. Shadow rate estimates of the macro-finance model. The figure reports the model-implied series and the corresponding observed series. Expected inflation comes from the Michigan Survey of Consumers. The difference is due to observation errors. The sample spans 1961:Q1 to 2017:Q3. Results are based on estimates obtained replacing the FFR with the Wu-Xia (2016) shadow rate when the ZLB is binding.

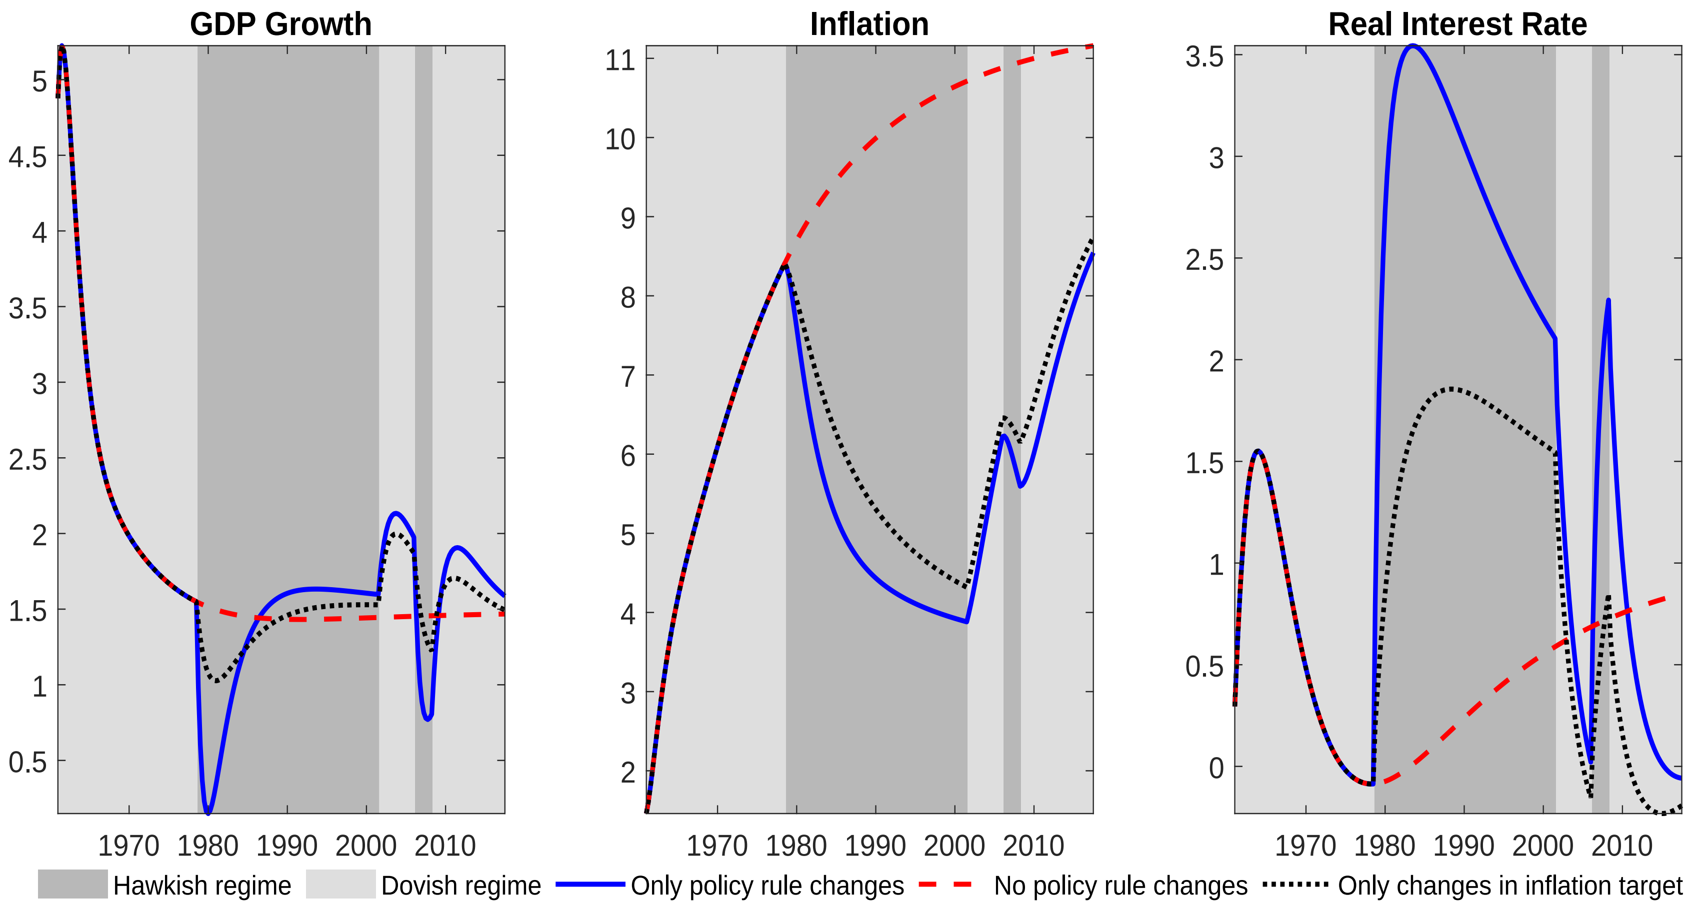

Figure IA.2. Shadow rate estimates of the macro-finance model. The blue line corresponds to the fluctuations generated by changes in both the target and the slope coefficients. The orange line assumes that monetary policy starts under the dovish regime and no regime change occurs. Finally, the black 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 (2016) shadow rate when the ZLB is binding.

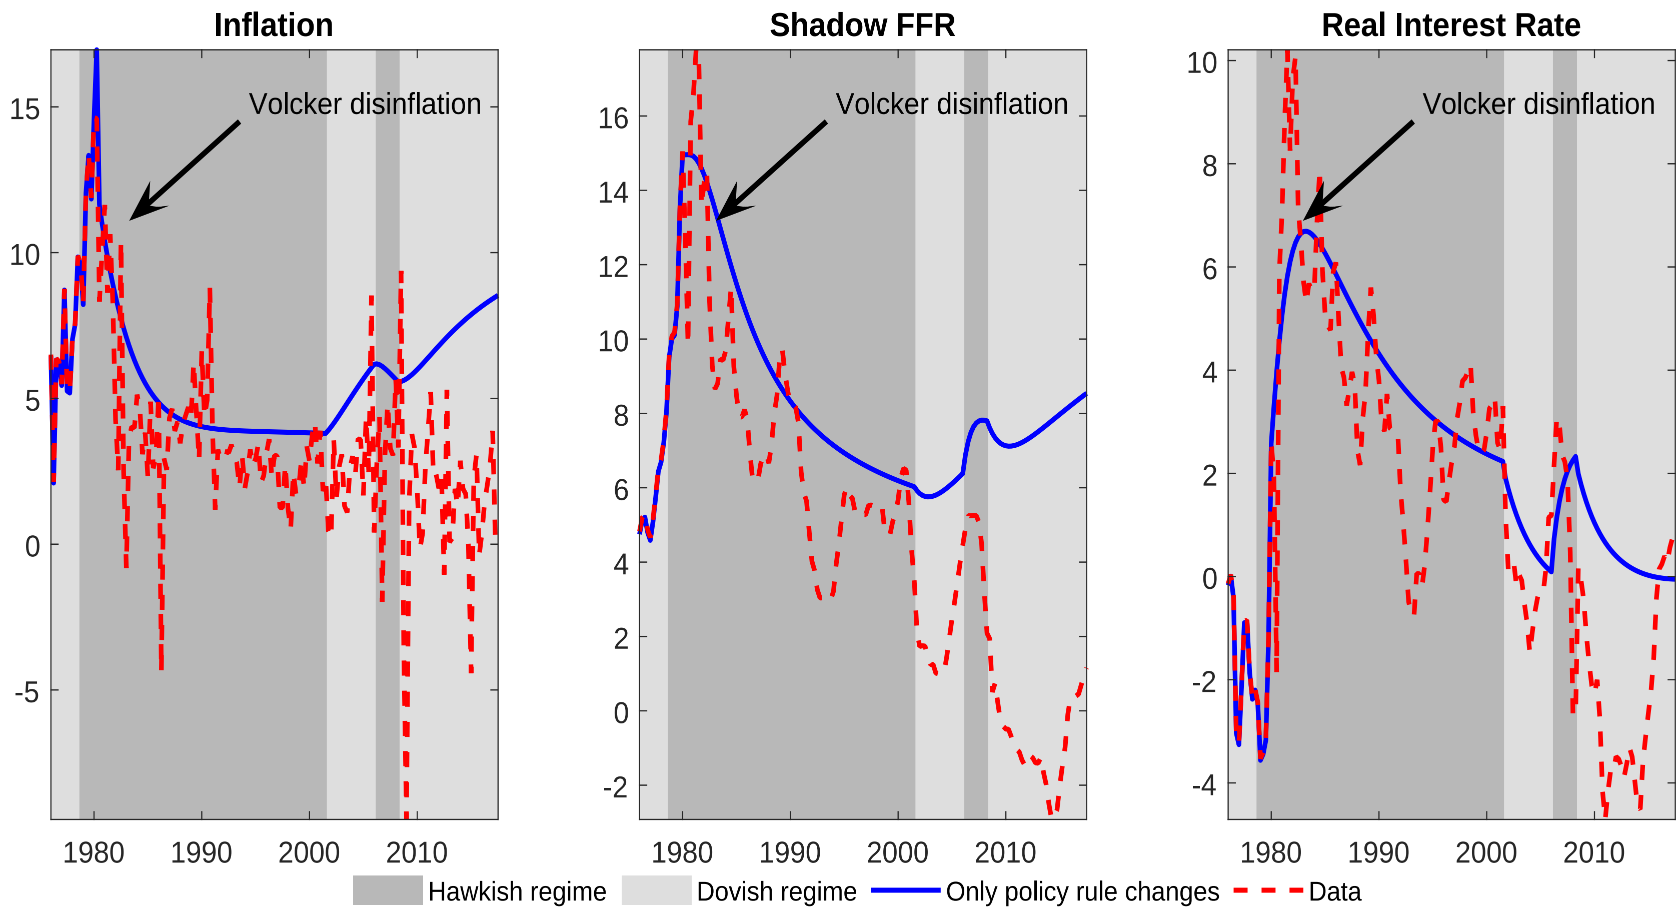

Figure IA.3. Shadow rate estimates of the macro-finance model. The Volcker disinflation. We start the economy as it was in 1980:Q1 and remove all Gaussian shocks that occured after that period but keep the estimated regime sequence. The dashed line corresponds to the data. The real interest rate is computed 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 (2016) shadow rate when the ZLB is binding.

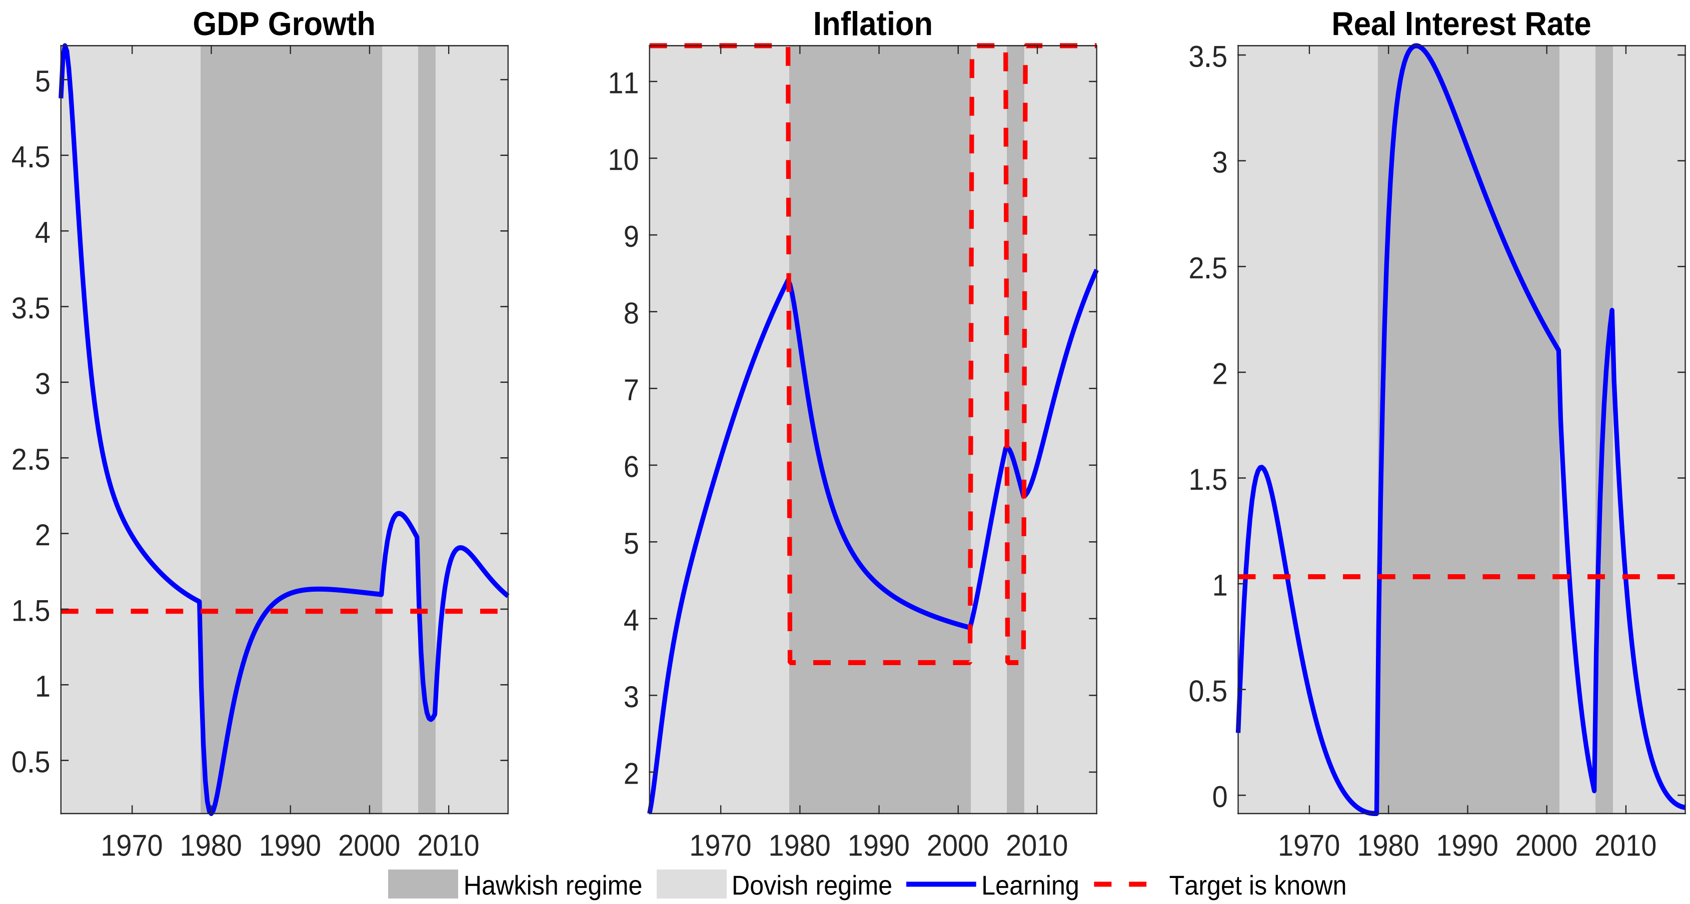

Figure IA.4. Shadow rate estimates of the macro-finance model. Perfect information about the target. The blue solid line shows estimated fluctuations generated only by changes in the policy rule (inflation target and slope coefficients) when agents learn about trend inflation. The orange dashed line shows a counterfactual in which the policy rule shifts but agents observe the inflation target. The dovish regime has a high target π^T and low activism against deviations from the target π^T. The hawkish regime has a low π^T and high activism against deviations from π^T. The sample spans 1961:Q1 to 2017:Q3. Results are based on estimates obtained replacing the FFR with the Wu-Xia (2016) shadow rate when the ZLB is binding.

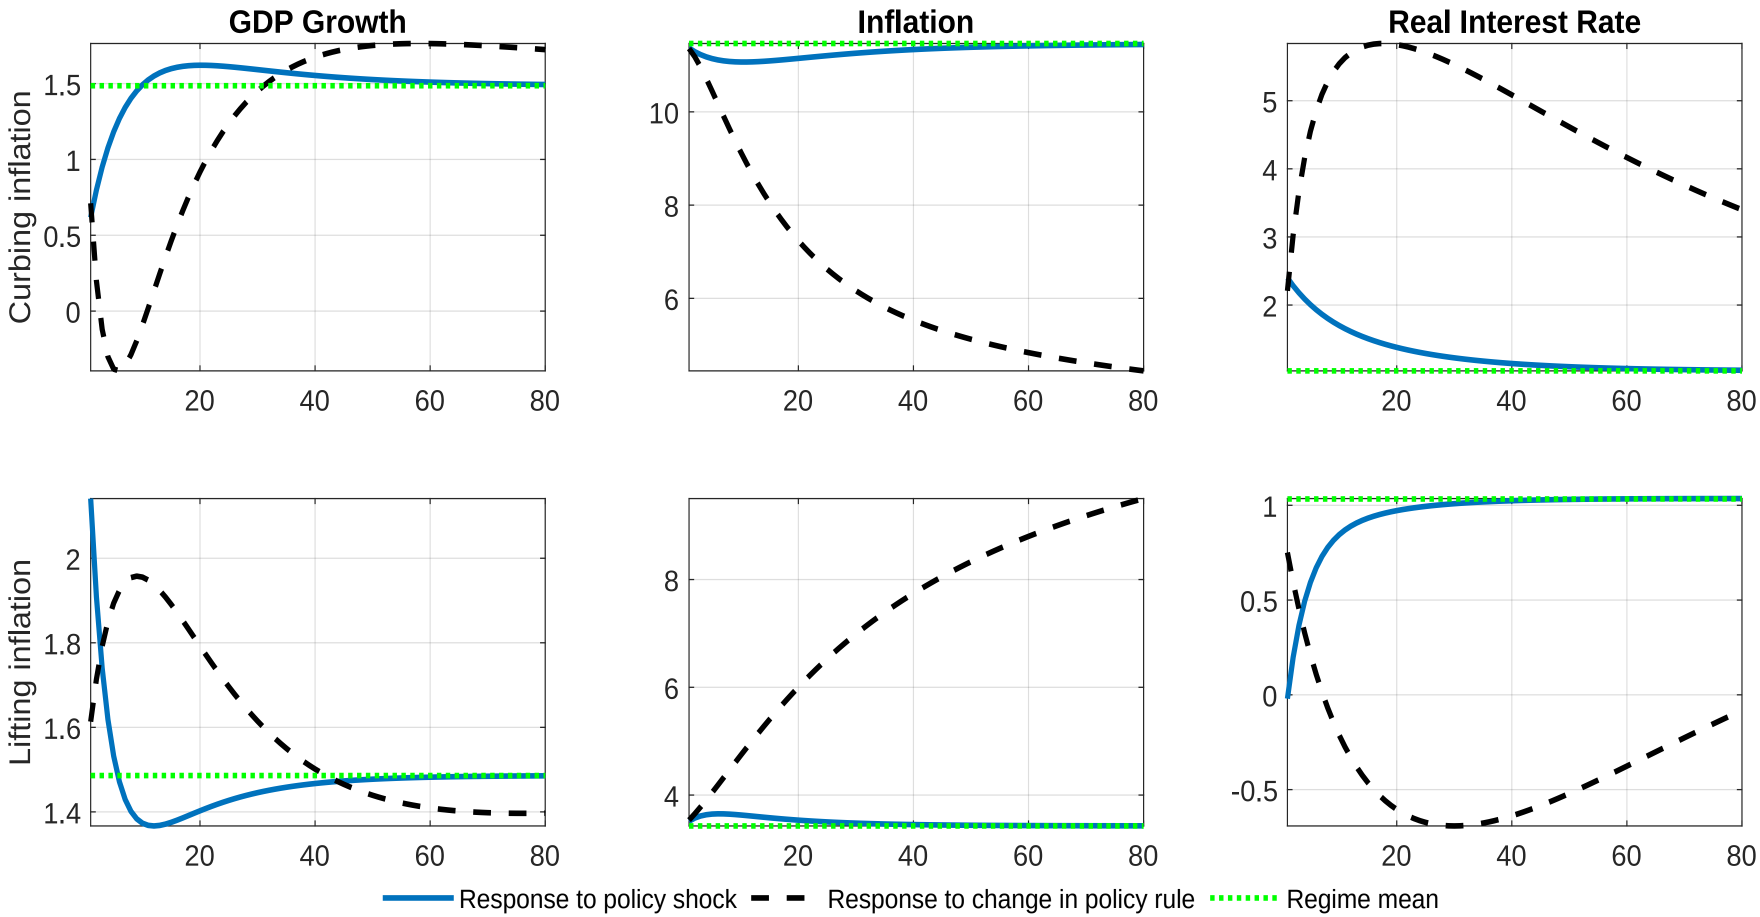

Figure IA.5. Shadow rate estimates of the macro-finance model. Top row: Curbing inflation. The economy is initially in the dovish regime. The blue solid line presents the evolution of the macro variables in response to a two-standard-deviation contractionary monetary policy shock and no regime change. 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 (2016) shadow rate when the ZLB is binding.

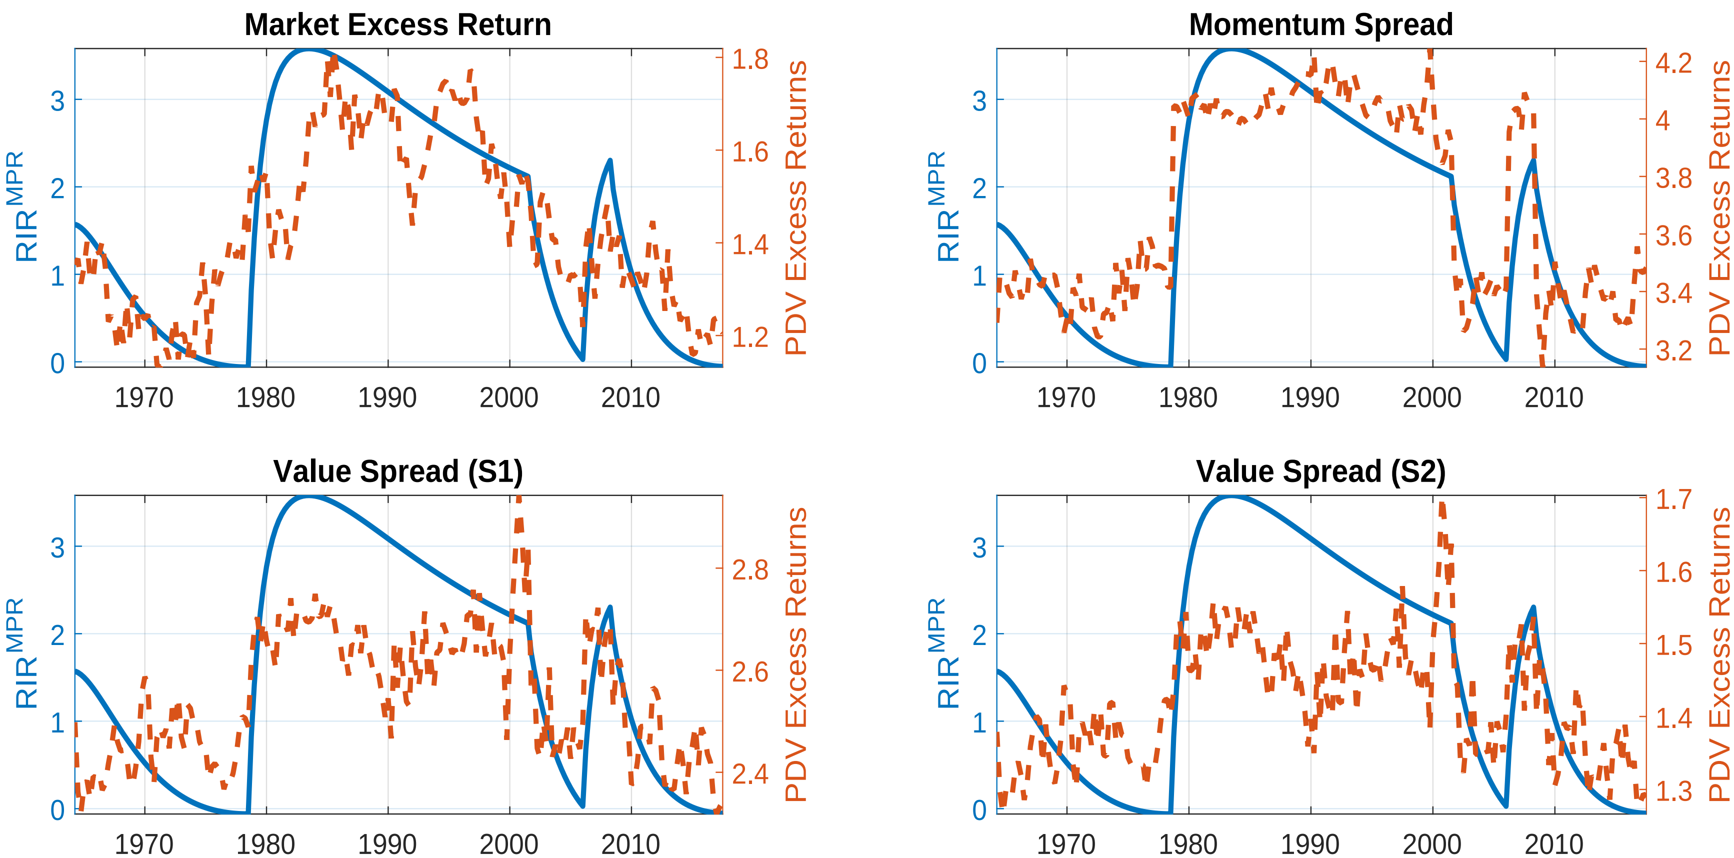

Figure IA.6. Shadow rate estimates of the macro-finance model. Excess returns and policy rule changes. The figure reports the time series of the present discounted value of expected excess returns for different portfolios (dashed line, right axis) together with fluctuations of the real interest rate 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 (2016) shadow rate when the ZLB is binding.

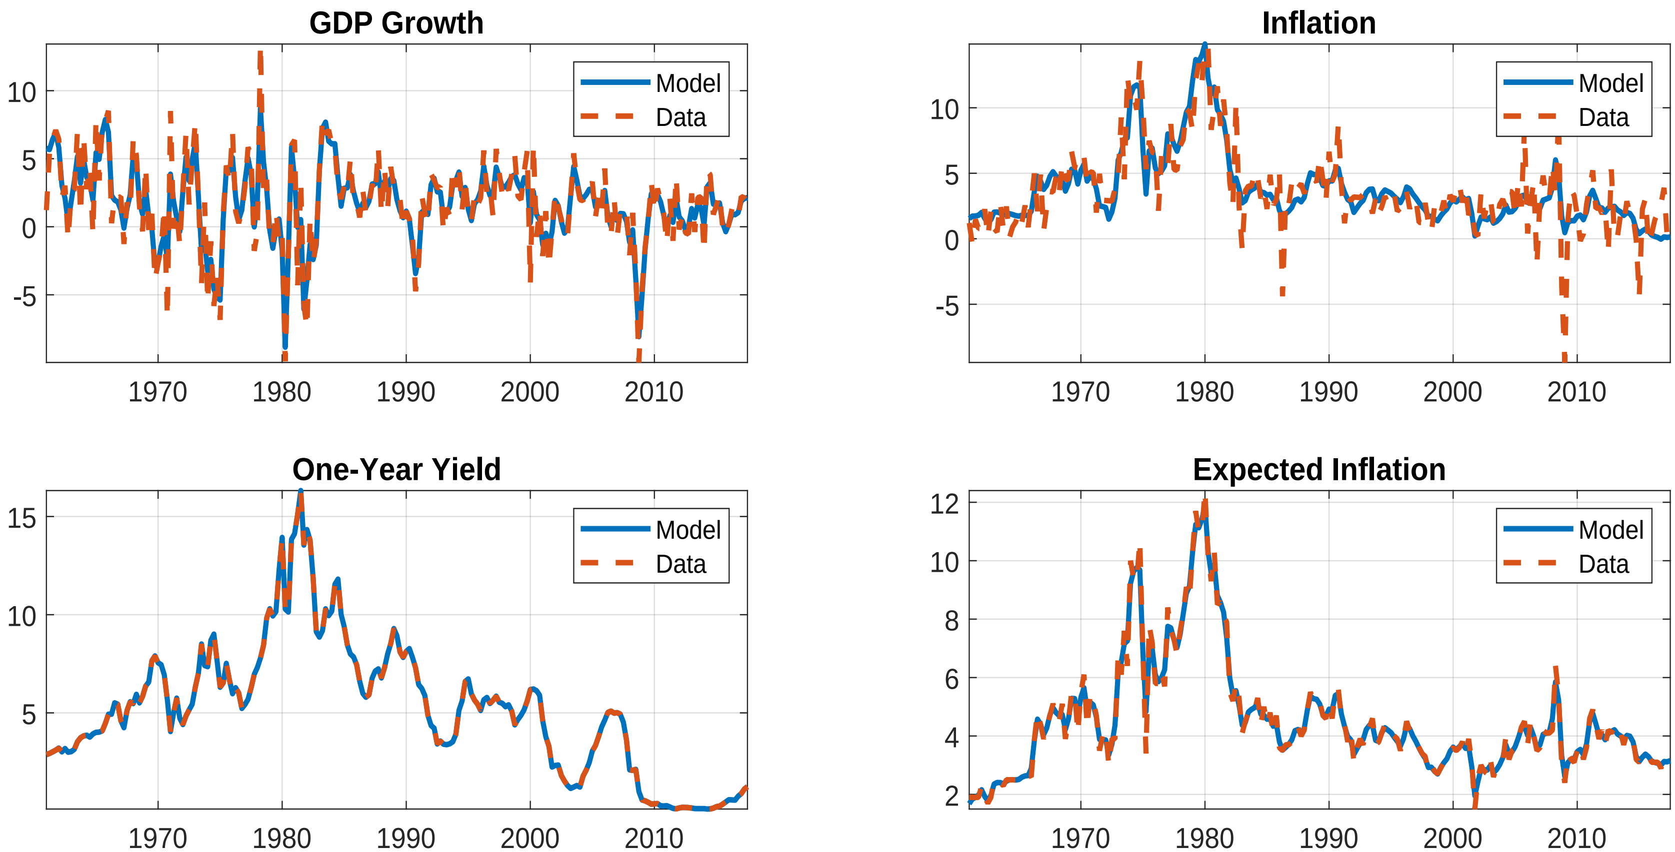

Figure IA.7. One-year yield estimates of the macro-finance model. The figure reports the model-implied series and the corresponding observed series. Expected inflation comes from the Michigan Survey of Consumers. The difference is due to observation errors. The sample spans 1961:Q1 to 2017:Q3. Results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

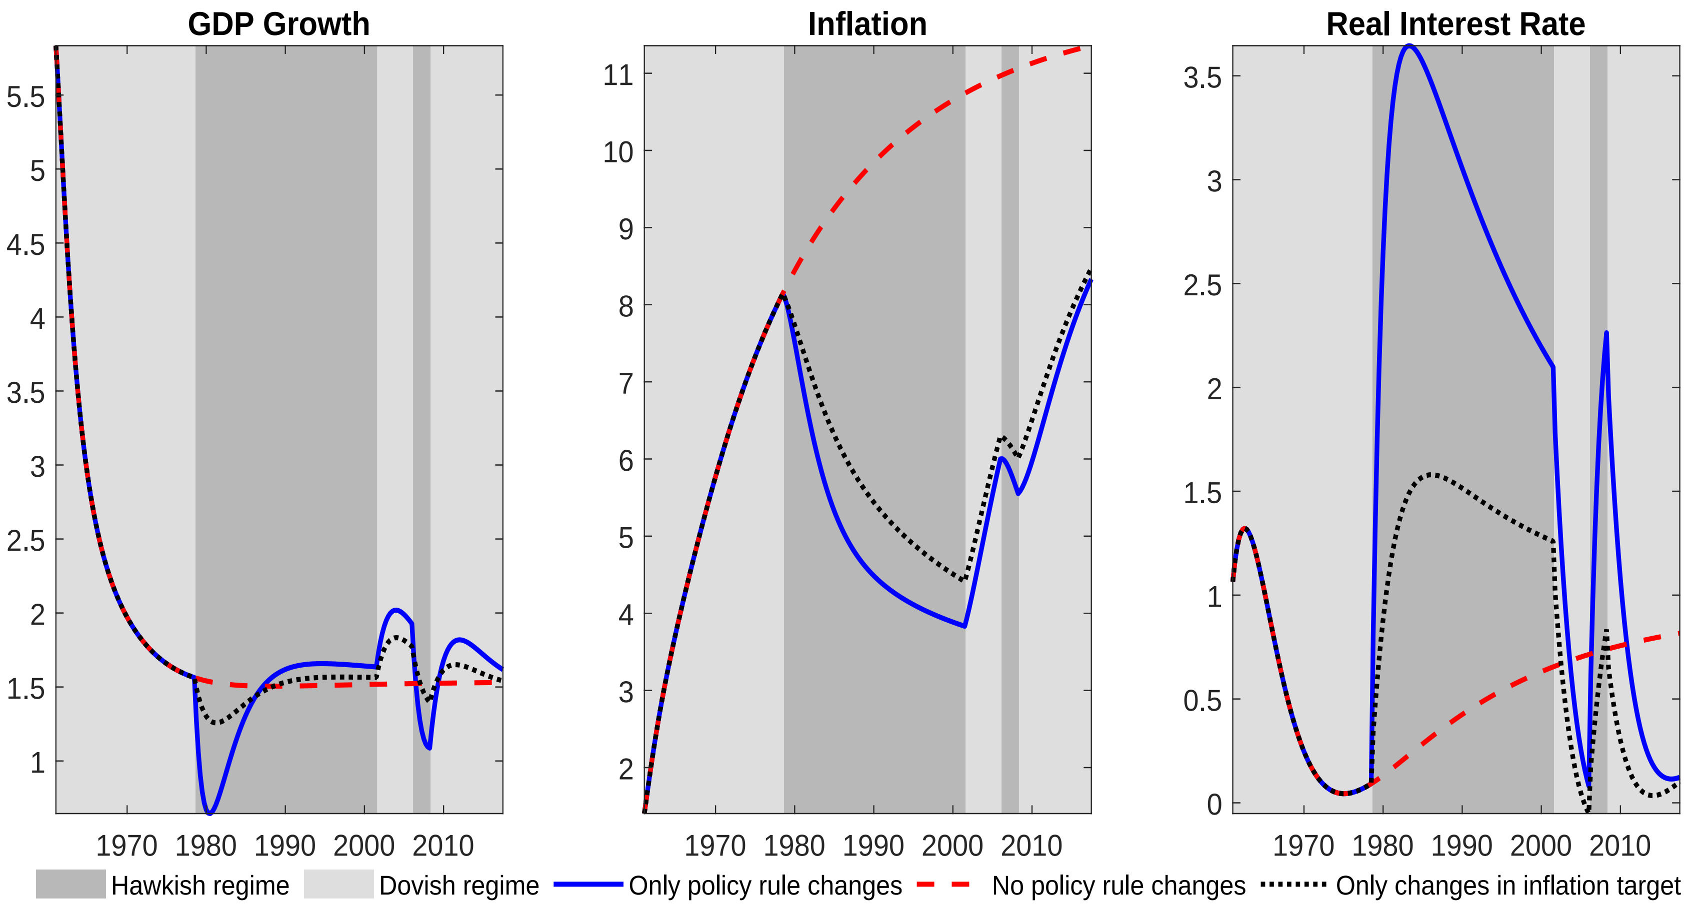

Figure IA.8. One-year yield estimates of the macro-finance model. The blue line corresponds to the fluctuations generated by changes in both the target and the slope coefficients. The orange line assumes that monetary policy starts under the dovish regime and no regime change occurs. Finally, the black 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.

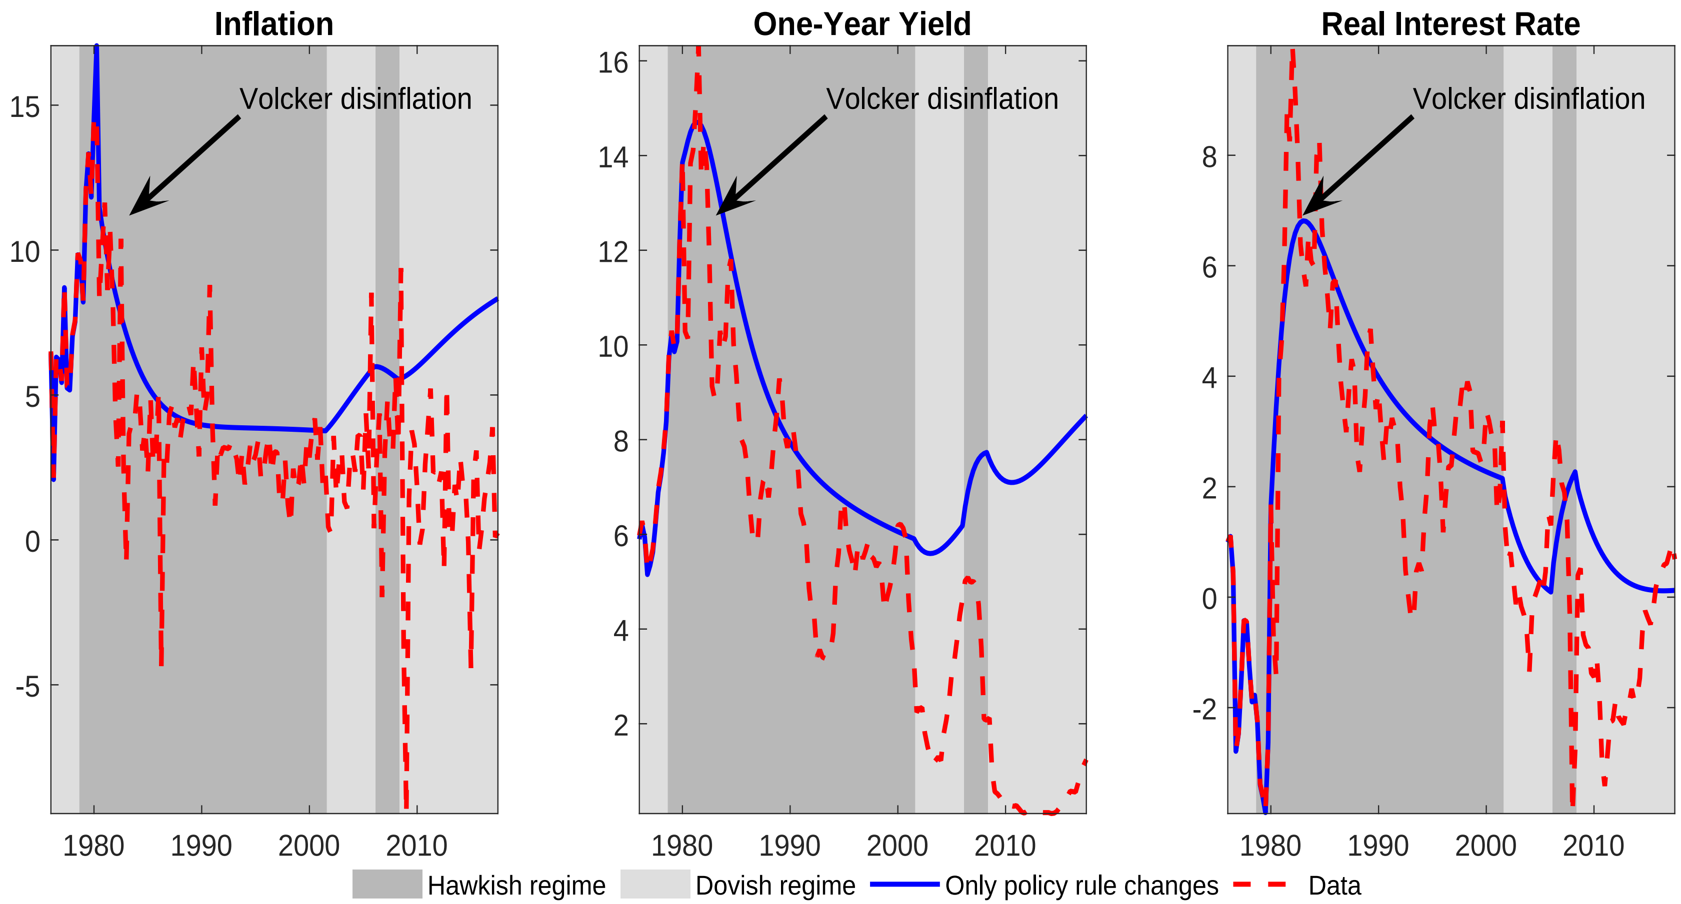

Figure IA.9. One-year yield estimates in the Volcker disinflation. We start the economy as it was in 1980:Q1 and remove all Gaussian shocks that occured after that period but keep the estimated regime sequence. The dashed line corresponds to the data. The real interest rate is computed as the difference 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.

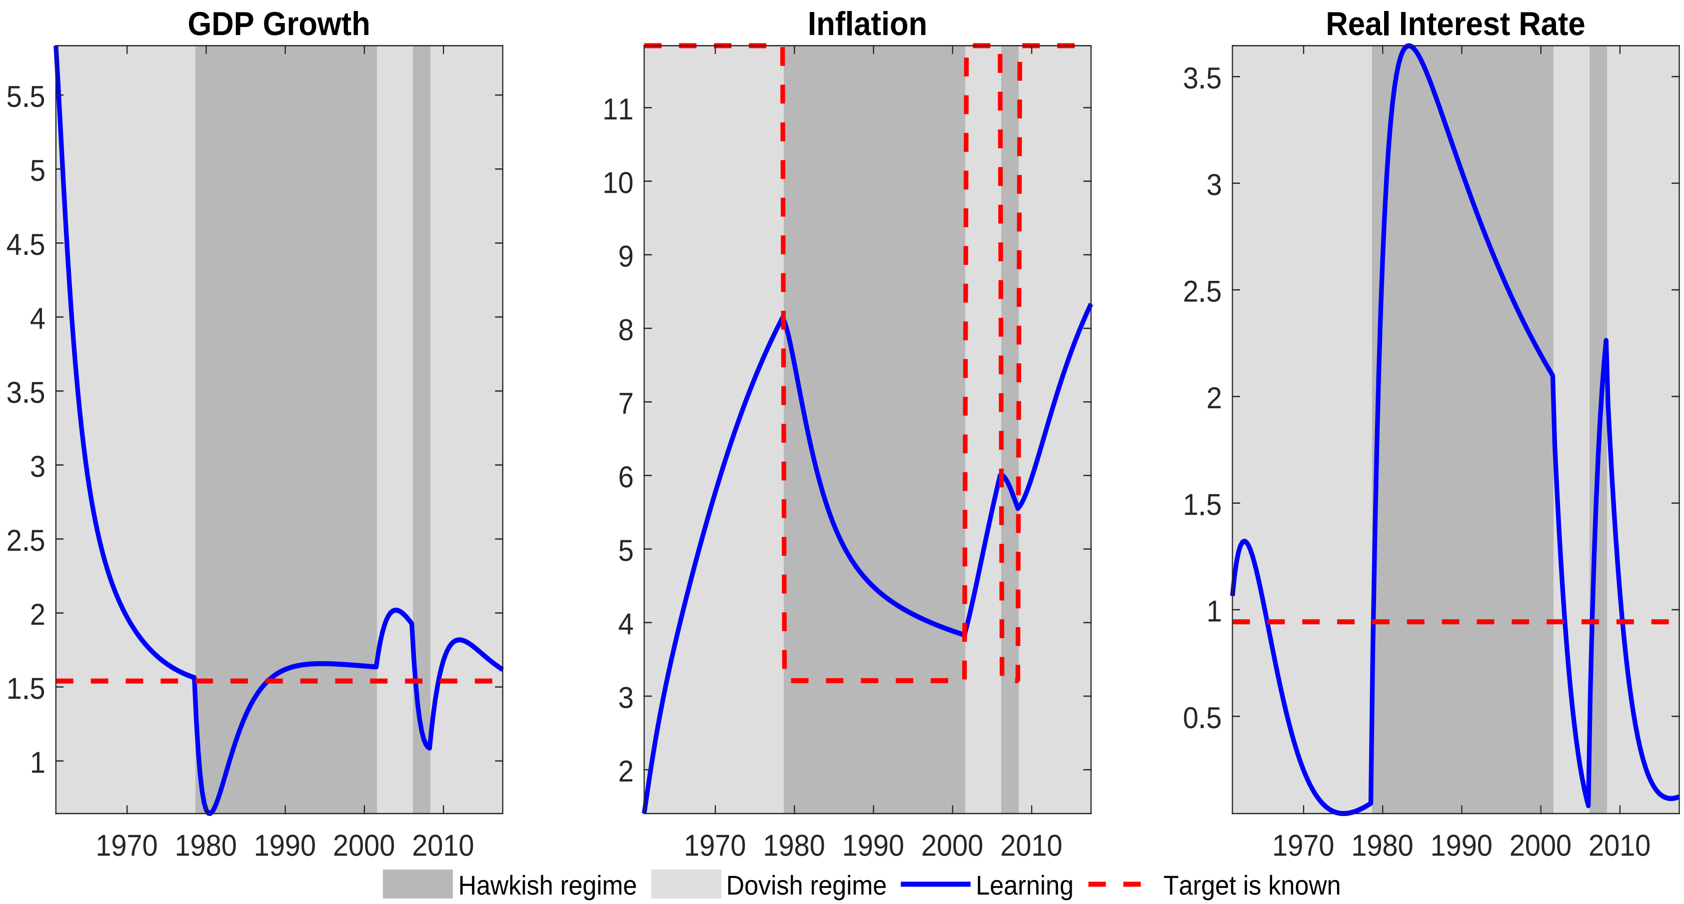

Figure IA.10. One-year yield estimates of the macro-finance model. Perfect information about the target. The blue solid line shows estimated fluctuations generated only by changes in the policy rule (inflation target and slope coefficients) when agents learn about trend inflation. The orange dashed line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

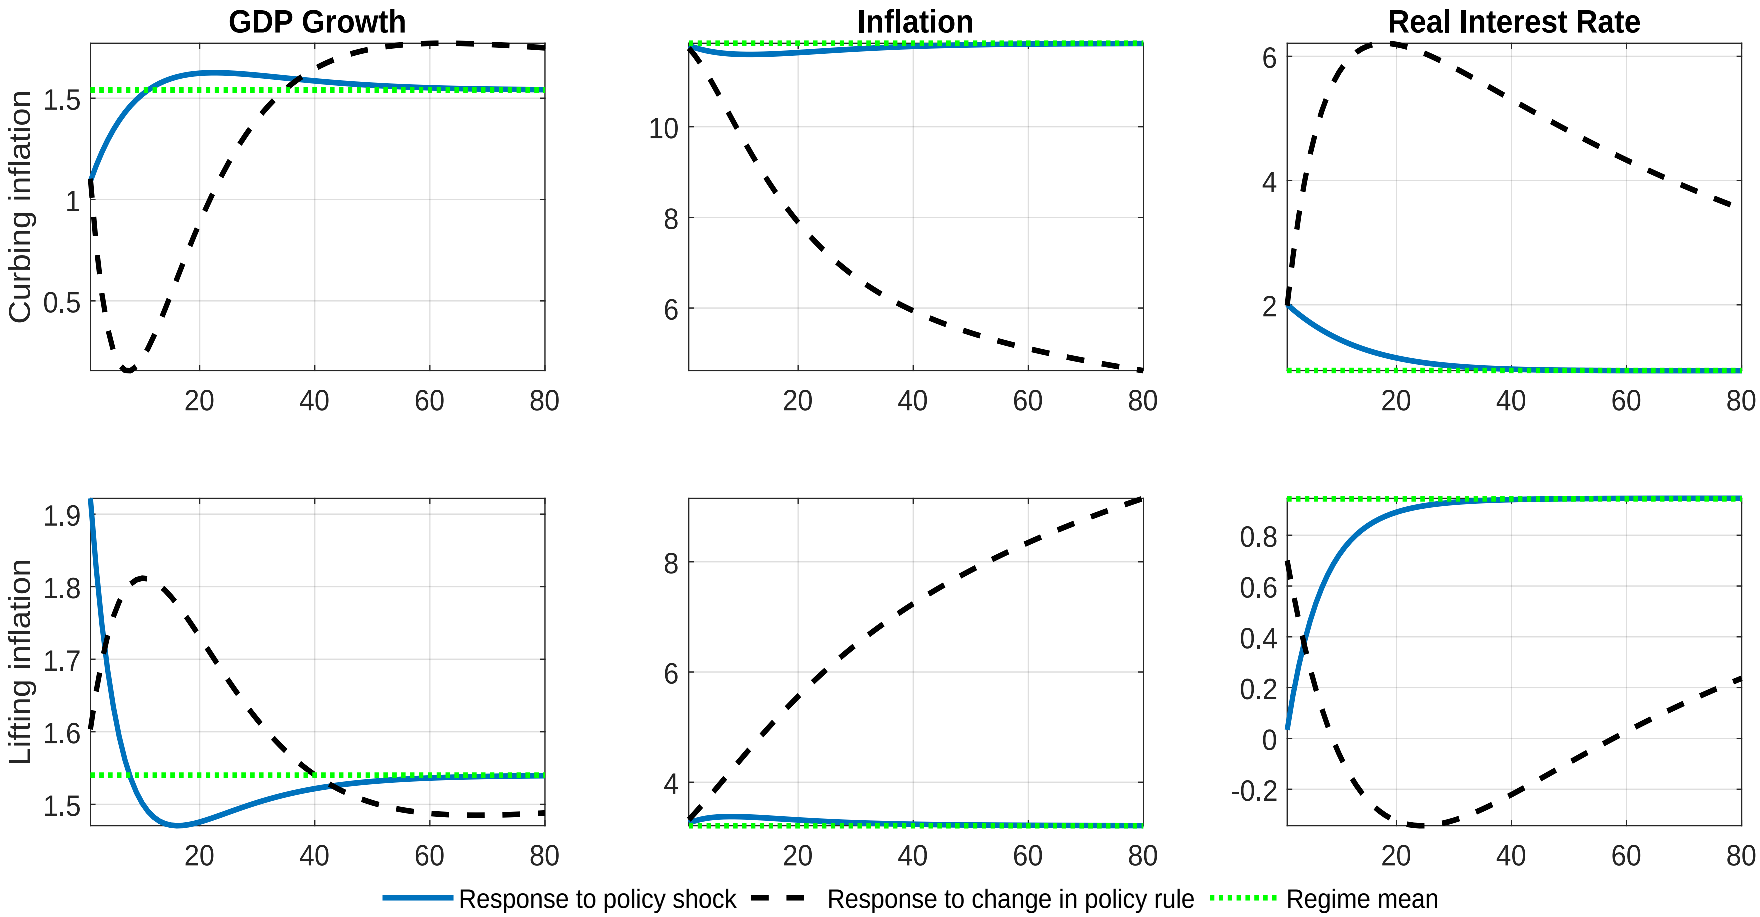

Figure IA.11. One-year yield estimates of the macro-finance model. Top row: Curbing inflation. The economy is initially in the dovish regime. The blue solid line presents the evolution of the macro variables in response to a two-standard-deviation contractionary monetary policy shock and no regime change. 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.

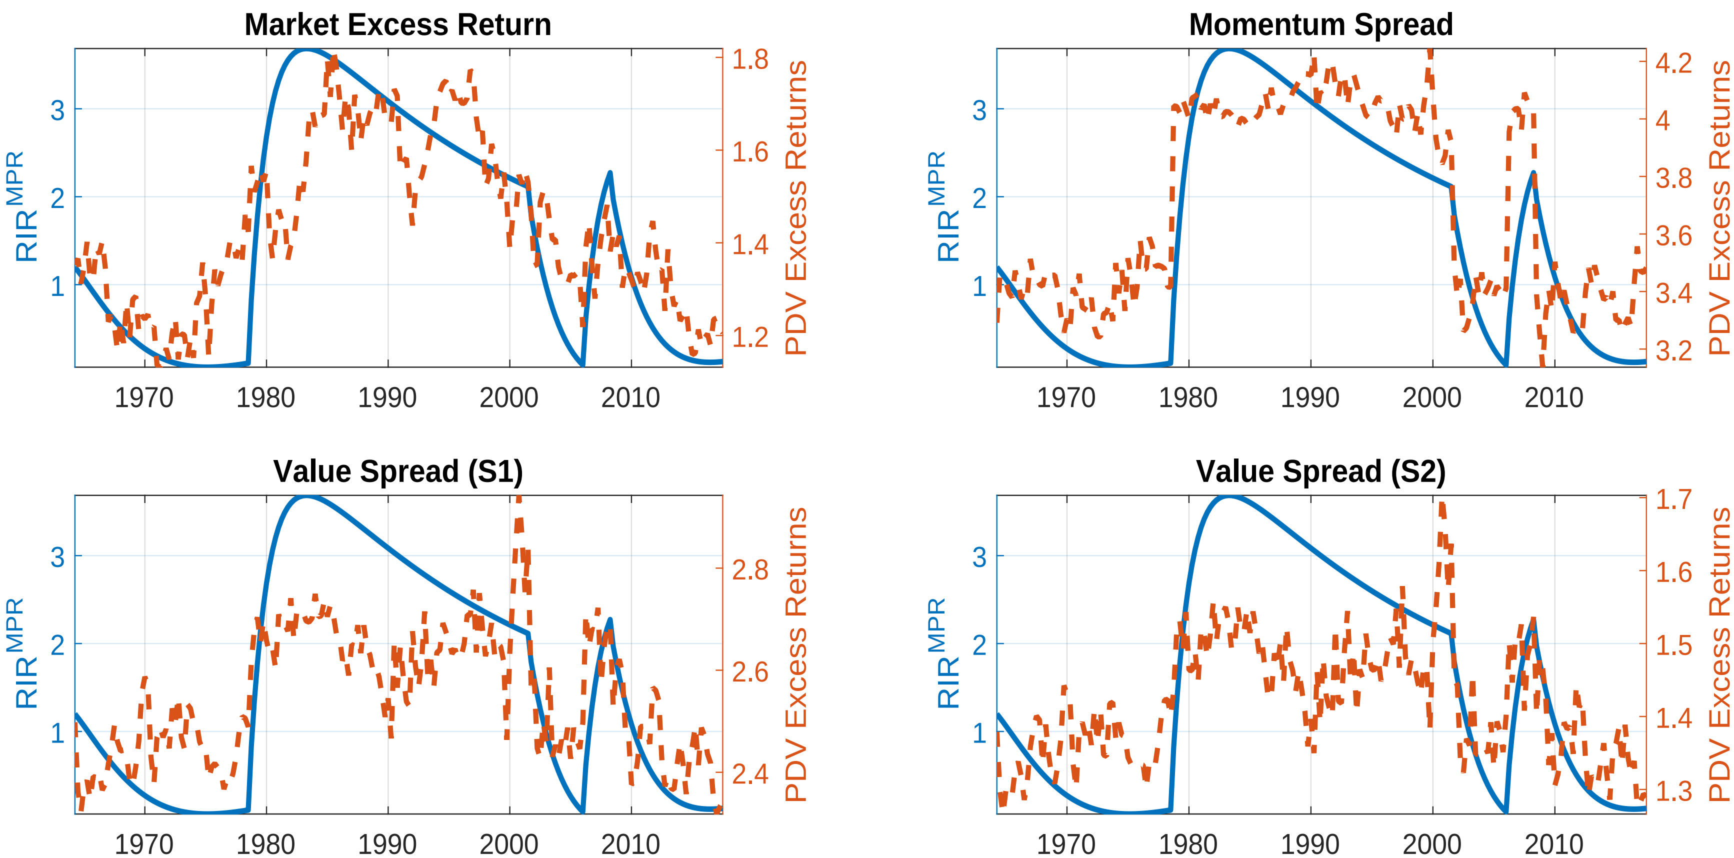

Figure IA.12. One-year yield estimates of the macro-finance model. Excess returns and policy rule changes. The figure reports the time series of the present discounted value of expected excess returns for different portfolios (dashed line, right axis) together with fluctuations of the real interest rate 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.

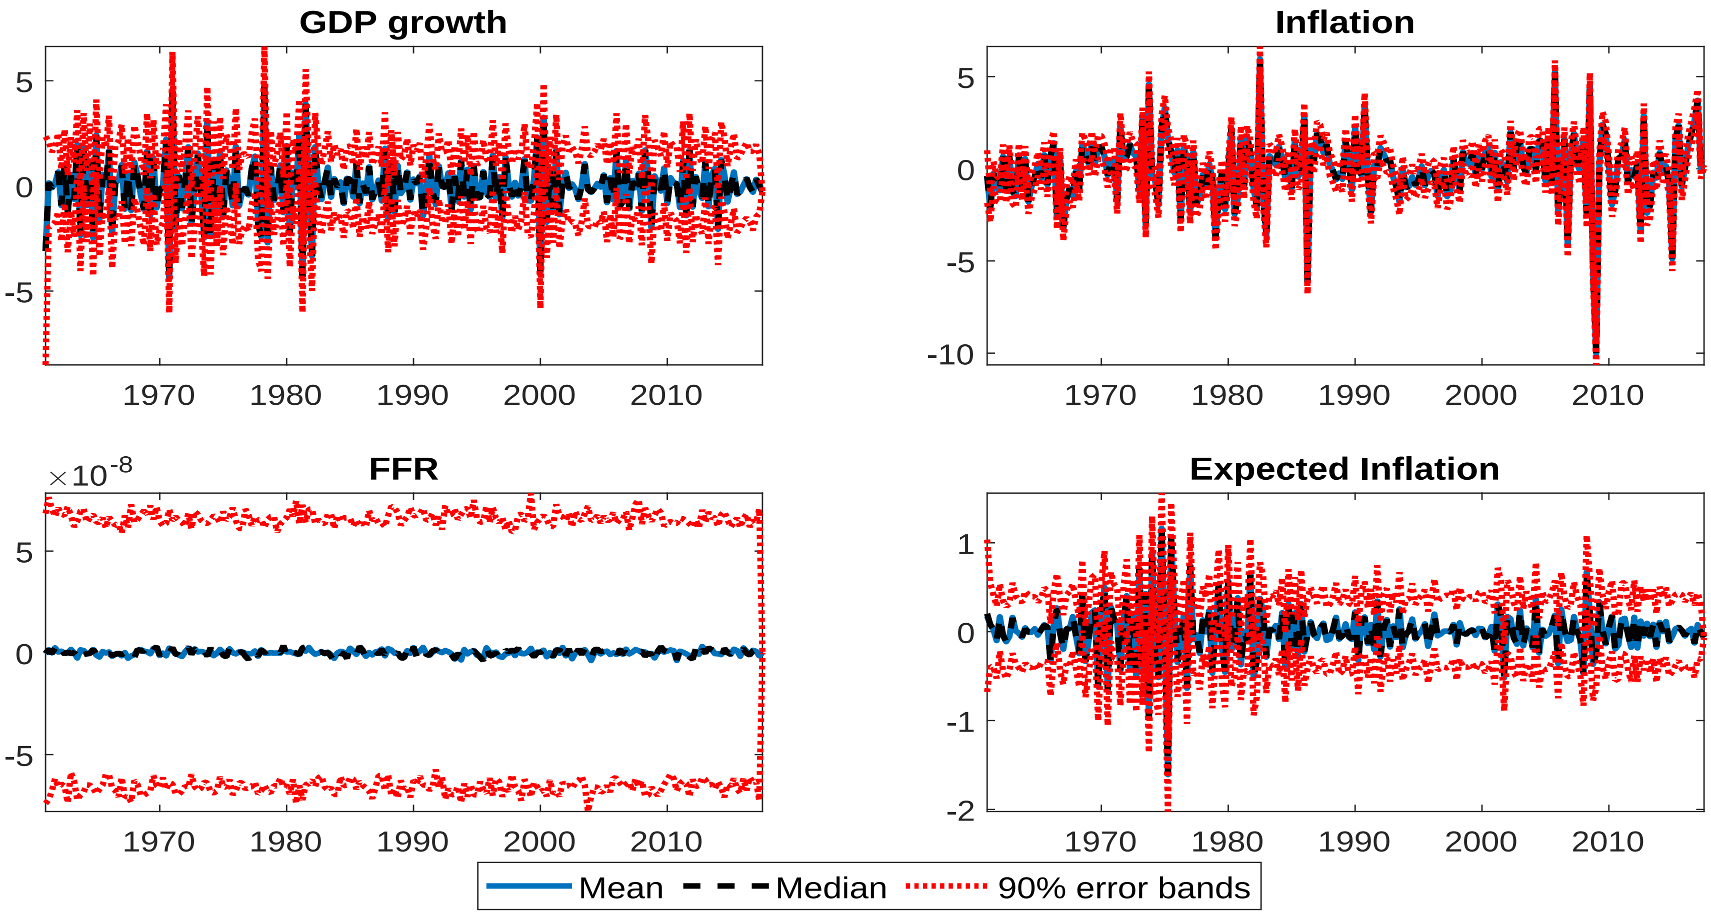

Figure IA.13. Distributuion of observation errors. The figure reports mean, median, and 90% error bands for the distribution of observation errors over time.