

ARTICLE

Frequency-dependent regime-switching in VAR models

Youngjin Hwang

Department of Economics, Hanyang University ERICA, Ansan, South Korea
Email: youngjin@hanyang.ac.kr

Abstract

This study presents a simple frequency-dependent regime-switching vector autoregression (VAR) model, where each regime and its associated parameters in the VAR are characterized by their distinct spectral properties. Empirical applications to several key macroeconomic variables reveal clear frequency-dependent switching dynamics, with each regime exhibiting distinctive features regarding spectral properties, volatility, and impulse responses. We compare this model with a conventional regime-switching model (typically studied in the time domain) and highlight several key differences between the two approaches.

Keywords: regime-switching; frequency domain; VAR

JEL classifications: C32; E32

1. Introduction

1.1. Motivation and main findings

It is well established that macroeconomic time series exhibit various forms of time variation, which has led to a growing interest in capturing these features in vector autoregression (VAR) models. Various models have been developed for this purpose, including, among others, threshold and smooth-transition VARs (Balke and Fomby, 1997; Teräsvirta et al. 2010), Markov-switching VARs (Krolzig, 1997; Sims et al. 2008), and time-varying coefficient VARs (with stochastic volatilities) (Cogley and Sargent, 2005; Primiceri, 2005).

While these models have been successful in capturing the time-varying features of data for a variety of applications, most approaches in this line of research have been developed in the time domain, and there have been limited attempts to characterize such features in the frequency domain for VARs. Considering that most macroeconomic time series exhibit distinct dynamics across different frequencies, and distinguishing these features is crucial in various contexts,¹ the lack of research in this area is a significant gap that needs to be addressed.

Motivated by these observations, this study presents a simple frequency-dependent regime-switching VAR model for macroeconomic time series. The novel feature of the model is that each regime and its associated parameters in the VAR are characterized by their own distinct spectral properties. Specifically, by postulating the two regimes (in the frequency domain), low-frequency and business cycle (BC) frequency, we explore how the key spectral features of the data exhibit switching behavior between the two regimes over time.

I am very grateful to the two anonymous reviewers for their valuable and constructive comments, which have significantly improved this paper.

Empirical applications of this model to key macroeconomic series, including GDP, inflation rates, and interest rates, yield several notable results. First, the dynamics of these variables are well captured by frequency-dependent regime-switching behavior over time. Specifically, while the low-frequency regime tends to dominate overall, there are several distinct periods, particularly from the 1970s to the mid-1980s, when the dynamics of the variables are better characterized by the BC-frequency regime.

In addition, these changes in the frequency regimes appear to be associated with BCs, with recession periods often corresponding to the BC-frequency regime. However, the degree of this association varies across different recession episodes. We also find that volatilities generally increase with frequency, that is, the BC-frequency regime is associated with relatively larger shock variances, which is consistent with findings in the Great Moderation literature (e.g., Kim and Nelson, 1999; McConnell and Perez-Quiros, 2000).

The impulse response analysis also reveals distinct frequency-dependent dynamics across regimes, with the responses in each regime characterized by strong persistence and clear cyclical patterns, respectively. Beyond these differences, we also observe several notable features in the responses across the regimes. For example, following an interest rate shock, while both output and inflation exhibit strong and persistent increases in the low-frequency regime, their responses in the BC-frequency regime are either negative or relatively short-lived and insignificant.

To illustrate the key properties of the model, we compare it with several alternative specifications. By comparing it with other models having different volatility specifications or priors, we show that regime-switching heteroskedasticity is essential for making reasonable and sharp inferences about regimes.

Moreover, by comparing our (frequency-dependent) regime-switching model with a conventional (time domain) regime-switching VAR, which is usually used to identify recession regimes, we highlight some notable differences between the two types of models. First, while the overall features of the estimated regimes are somewhat similar between the two models, the conventional model identifies several recession episodes as a recession regime, which are classified as either a low-frequency regime or only a weakly identified BC-frequency regime in the frequency-dependent switching model. In addition, although the shapes of the spectral densities are largely similar for both the BC-frequency and recession regimes, the low-frequency regime contains relatively more low-frequency components and is therefore more persistent compared to the expansion regime. These results suggest that the regimes identified in each type of model have their distinct characteristics, and thus the expansion (recession) regime cannot be directly equated with or treated as the low- (BC-) frequency regime.

Finally, a simple comparison of forecasting performance between models shows that the frequency-dependent regime-switching model delivers reasonable predictive performance; it significantly outperforms the constant-parameter VAR, and its performance is roughly comparable to the conventional time-domain VAR.

1.2. Related literature

This study is closely related to several strands of empirical research.

First, it is widely recognized that most macroeconomic time series reflect both BC factors and low-frequency (or trend) forces, each with different characteristics. In the context of VARs, commonly used to analyze the responses of variables to structural shocks, recent studies have highlighted the importance of accounting for these frequency-domain properties. For example, some puzzling results regarding the effects of structural shocks can be explained by imposing restrictions on the frequency interval around zero for technology shocks (Francis and Ramey, 2005), or by isolating short-term and BC fluctuations from low-frequency effects in the context of news shocks (Barsky and Sims, 2011; Kurmann and Sims, 2021).

Meanwhile, another line of research explores the dynamic properties of DSGE models in the frequency domain. For instance, several recent studies have examined the properties of standard DSGE models (e.g., Smets and Wouters, 2007) over specific frequency ranges and found that key aspects of the models, such as parameter estimates, goodness of fit, and forecasting performance, vary significantly depending on the frequency bands over which the models are estimated (e.g., Tkachenko and Qu, 2012; Sala, 2015; Caraiiani, 2015). These studies suggest that, without careful consideration of the frequency-domain properties of the data, inference based on VARs (albeit relatively less stylized than DSGE models) can be potentially misleading.

Although the literature on estimating VARs at different frequencies is relatively limited, there are some notable exceptions. For example, Sargent and Surico (2011) examine the time-varying low-frequency relationships between money growth and inflation, while Kliem et al. (2016) study those between public deficits and inflation. Both studies find significant changes in these low-frequency relationships, especially around the 1970s, which is consistent with the findings of this study. While their focus is mainly on low-frequency behavior, with little attention to other frequencies, this study provides a more comprehensive analysis of how key macroeconomic variables exhibit time variations across the entire frequency spectrum, highlighting frequency-dependent regime shifts.

This study is also closely related to the literature on the prior elicitation for Bayesian VARs. Since VARs typically involve a large number of parameters, Bayesian VAR analysis has become increasingly popular, incorporating additional information into the estimation by adopting appropriate priors to sharpen the inference. Various priors have been proposed for this purpose, including the Minnesota prior (Litterman, 1986), (independent) normal-inverse Wishart prior (Kadiyala and Karlsson, 1997), the steady-state prior (Villani, 2009), and priors designed for large-scale VARs (Chan, 2020a).²

While these priors have proven useful in various contexts, most have been developed in the time domain, with limited efforts made to formulate priors for VARs in the frequency domain that capture the distinct features of data at different frequencies. Some studies have proposed useful approaches for eliciting priors over specific frequency ranges, such as the long-run (Giannone et al. 2019) or the BC frequencies (Andrle and Plašil, 2018; Planas et al. 2008; Jarociński and Lenza, 2018). However, few attempts have been made to develop a comprehensive set of priors covering the entire frequency spectrum, including cases involving frequency-dependent regime-switching, as explored in this study.

This study addresses this gap by proposing a set of priors tailored to the spectral properties of different frequency regimes and estimating the corresponding hyperparameters, thus contributing to the existing literature in this area. Our results show that several hyperparameters (in particular, those associated with co-integration) exhibit clear frequency-dependent characteristics, with each set well-suited to capture the distinct dynamics specific to its respective regime.

Finally, this study is closely related to wavelet-based analysis. While wavelet analysis provides a useful toolkit for exploring variations in time-frequency information in data and has yielded interesting results (e.g., Aguiar-Conraria et al 2012, 2018), it is technically complex, and interpreting or extracting economically meaningful insights can be challenging.³ This study, which allows the spectral properties of variables to change over time within a familiar parameterized model, complement wavelet-based approaches and can be easily extended or integrated with other widely used economic models.

The remainder of this paper is organized as follows: Section 2 outlines the econometric model, while Section 3 presents the data and the main empirical results. Section 4 examines several alternative model specifications and compares their results. Finally, Section 5 provides concluding remarks.

2. Econometric model

2.1. A regime-switching VAR

Consider a simple VAR model where the variables of interest, y_t , exhibit regime-switching behavior:

$$y_t = \beta_{S_t}^{(0)} + \beta_{S_t}^{(1)} y_{t-1} + \dots + \beta_{S_t}^{(p)} y_{t-p} + u_t, \quad u_t \sim N(0, \Sigma_{S_t}), \tag{1}$$

where y_t is an $n \times 1$ vector of data at time t ; $\beta_{S_t}^{(0)}$ is an $n \times 1$ vector of constants and $\beta_{S_t}^{(l)}$ is an $n \times n$ matrix associated with the lagged terms; and the error term, u_t is an n -dimensional zero mean white noise process with covariance matrix, Σ_{S_t} . Stacking the dependent variables into a $T \times n$ matrix, y , whose t 'th row is y_t' and the regressors into a $T \times (np + 1)$ matrix, X , where the t 'th row is $X_t = (1, y_{t-1}', \dots, y_{t-p}')$, respectively, we have

$$y = X\beta_{S_t} + u, \tag{2}$$

where u is a $T \times n$ matrix of innovations in which the t 'th row is u_t' and $\beta_{S_t} = [\beta_{S_t}^{(0)}, \beta_{S_t}^{(1)}, \dots, \beta_{S_t}^{(p)}]'$. The evolution of latent (R -state) Markov state variable S_t , is governed by the transition probability matrix, P . Specifically, the (r, r') element of P , denoted as $P_{(r,r')}$, represents the transition probability from state r at time $t - 1$ to state r' at time t , that is, $P_{(r,r')} = p(S_t = r' | S_{t-1} = r)$.

A novel feature of this model is that the key properties of each regime and the associated (regime-specific) parameters (i.e., β_{S_t} and Σ_{S_t}) are characterized by their own distinctive spectral properties. Specifically, by postulating a few regimes in the frequency domain such as low-frequency, BC-frequency, and/or high frequency, we assume that the dynamics of data at each time can be represented as one or a combination of these regimes.

Our main objective is to make an inference about the unobserved states as well as the model parameters. To do this, we rely on Bayesian inference, which provides a natural framework for incorporating the frequency-dependent behavior of the data. The following subsections describe our approach to prior specification and outline the estimation algorithm.

2.2. Prior

This subsection discusses the (regime-dependent) priors imposed on the model parameters. We adopt a normal-inverse Wishart prior for the VAR model parameters (β_{S_t} and Σ_{S_t}), along with the Minnesota prior, which provides shrinkage toward the specified prior means in each frequency regime. In the following, for the sake of notational simplicity, we drop the regime subscript, S_t , unless it causes confusion.

Eliciting priors for the VAR coefficients largely draws on the framework of the unobserved components model, widely used in the trend-cycle decomposition literature (Harvey, 1985; Watson, 1986; Clark, 1987). This approach is useful in that it allows us to impose priors that appropriately capture the key characteristics of different frequency regimes with a small set of hyperparameters.

In a typical unobserved component model, (log of) a macro time series is assumed to consist of three components, trend (y_t^τ), cyclical (y_t^c), and irregular components (y_t^ε). First, the trend is usually assumed to evolve as follows:

$$y_t^\tau = y_{t-1}^\tau + \mu_t, \tag{3}$$

$$\mu_t = \mu_{t-1} + \delta_t, \quad \delta_t \sim N(0, \sigma_\delta^2), \tag{4}$$

which implies that the trend follows an integrated random walk, or ARIMA(0, 2, 0):

$$\Delta^2 y_t^\tau = y_t^\tau - 2y_{t-1}^\tau + y_{t-2}^\tau = \delta_t. \tag{5}$$

Building on this, we impose an I(2) process for the VAR coefficients in the low-frequency regime ($S_t = 1$).⁴

Next, the prior for the dynamics of the BC-frequency regime ($S_t = 2$) is specified using trigonometric functions. Specifically, following Planas et al. (2008) and Jarociński and Lenza (2018), we assume:

$$(1 - 2\rho \cos(2\pi/\tau) L + \rho^2 L^2) y_t^c = u_t, \tag{6}$$

where L is the lag operator, ρ denotes the amplitude, and τ represents the periodicity.⁵

Finally, we set a simple white noise process for the fluctuations in the high-frequency regime ($S_t = 3$) (when additionally considered).⁶

To summarize, the prior means for the VAR coefficients are set as follows:

$$\underline{\beta}_{S_t} = \begin{cases} [0_{(1 \times n)}; 2I_n; -I_n; 0_{(n-2)p \times n}], & \text{for } S_t = 1; \\ [0_{(1 \times n)}; 2\rho \cos(2\pi/\tau) I_n; -\rho^2 I_n; 0_{(n-2)p \times n}], & \text{for } S_t = 2; \\ [0_{((np+1) \times n)}], & \text{for } S_t = 3. \end{cases} \tag{7}$$

We implement these priors by adding the following (regime-dependent) dummy observations, grouped into several blocks, to the actual data (Sims and Zha, 1998).

The first block is specified as:

$$\underbrace{y_1^D}_{(np \times n)} = [\underline{\beta}^{(1)} \odot \underline{S} \cdot 1^{\lambda_2}; \dots; \underline{\beta}^{(p)} \odot \underline{S} \cdot p^{\lambda_2}] / \lambda_1, \tag{8}$$

and

$$\underbrace{X_1^D}_{(np \times (np+1))} = [0_{(np \times 1)}, \text{diag}(1^{\lambda_2}, \dots, p^{\lambda_2}) \otimes \underline{S} / \lambda_1], \tag{9}$$

where $\underline{\beta}^{(l)}$ is an $n \times n$ matrix of (regime-dependent) prior means of the VAR coefficients at lag l , whose (i, j) element represents the coefficient for variable j in equation i at lag l ; \underline{S} is the diagonal scale matrix whose elements are set to the standard deviations of error terms from OLS estimates of the autoregression for each variable in the model; and \odot is the element-by-element multiplication operator.

Since this block implies that, for variable j in equation i at lag l ,

$$\beta_{(i,j)}^{(l)} \sim N \left(\underline{\beta}_{(i,j)}^{(l)}, \left(\frac{\lambda_1}{l^{\lambda_2}} \times \frac{\Sigma_{(i,i)}}{\underline{S}_{(i)}} \right)^2 \right), \tag{10}$$

it follows that λ_1 governs the overall tightness of the prior, while λ_2 controls the shrinkage of more distant lags. Specifically, a smaller λ_1 indicates a tighter prior while a larger value of λ_2 implies greater shrinkage for longer lags.

The second block is defined as:

$$\underbrace{y_2^D}_{(n \times n)} = \tilde{\beta} \bar{y} / \lambda_3 \quad \text{and} \quad \underbrace{X_2^D}_{(n \times (np+1))} = [0_{(n \times 1)}, \bar{y} / \lambda_3, \dots, \bar{y} / \lambda_3], \tag{11}$$

where $\tilde{\beta}$ is an $n \times n$ diagonal matrix whose elements are the sum of its own coefficients over lags, that is, $\tilde{\beta}_{(i,i)} = \sum_{l=1}^p \beta_{(i,i)}^{(l)}$; and \bar{y} is an $n \times n$ diagonal matrix whose i 'th element is equal to the sample mean of the variable y_i (as in Sims and Zha, 1998). This prior implies:

$$\sum_{l=1}^p \beta_{(i,i)}^{(l)} \sim N \left(\tilde{\beta}_{(i,i)}, \left(\frac{\lambda_3}{\bar{y}_{(i)}} \right)^2 \Sigma_{(i,i)} \right), \tag{12}$$

indicating that λ_3 controls the tightness of the sum of its own coefficients. This prior, a generalization of the random walk prior, allows for flexible forms of persistence (through autocorrelation), with a smaller λ_3 corresponding to a more informative prior.

The third block is given by:

$$\underbrace{y_3^D}_{(n \times n)} = \text{diag}(\bar{y}) / \lambda_4 \quad \text{and} \quad \underbrace{X_3^D}_{(n \times (np+1))} = [0_{(n \times 1)}, y_3^D, \dots, y_3^D], \tag{13}$$

and this prior pushes the variables toward their unconditional mean, ensuring that the series in the VAR share a common stochastic trend. Thus, this prior pertains to co-persistence or co-integration. The tightness of this prior is controlled by the hyperparameter λ_4 , where a smaller λ_4 indicates a tighter prior.

The last two blocks are:

$$\underbrace{y_4^D}_{(n \times n)} = 0_{(n \times n)} \quad \text{and} \quad \underbrace{X_4^D}_{(n \times (np+1))} = [1_{(n \times 1)} / \lambda_5, 0_{(n \times np)}], \tag{14}$$

and

$$\underbrace{y_5^D}_{(n \times n)} = \underline{S} \quad \text{and} \quad \underbrace{X_5^D}_{(n \times (np+1))} = 0_{(n \times (np+1))}, \tag{15}$$

and these two blocks relate to the priors on the constant terms and the covariances of the error terms, respectively. All these blocks are collected in $y^D = [y_1^D; y_2^D; y_3^D; y_4^D; y_5^D]$ and $X^D = [X_1^D; X_2^D; X_3^D; X_4^D; X_5^D]$, both with a length of T_D .

The dummy observation priors are characterized by a set of hyperparameters, $\lambda = (\lambda_1, \dots, \lambda_5)$. In setting their values, we follow the hierarchical approach of Giannone et al. (2015), by treating these hyperparameters as additional regime-dependent parameters to be estimated.⁷ In addition to its flexibility, this approach is useful since there seems no consensus on the appropriate values for frequency-specific priors (especially in models involving regime-switching, as in this study). In particular, given that the degree and pattern of shrinkage and persistence may differ across frequencies, it will be interesting to observe whether there are significant differences in the tightness of these priors across different frequency regimes.

Note that by assuming a normal-inverse Wishart distribution, our prior specification treats each equation symmetrically. To address this potentially restrictive property, several approaches involving asymmetric conjugate priors have been proposed (e.g., Carriero et al. 2019; Carriero et al. 2022; Chan, 2022). While these approaches offer greater flexibility and computational efficiency, it is unclear whether this feature would be beneficial in the context of our model. On the one hand, allowing for a flexible structure through an asymmetric prior may allow the data to better reflect equation- or variable-specific information. On the other hand, this flexibility may complicate the estimation of the joint switching behavior of the variables in the frequency domain. We revisit this issue later by examining a model in which the symmetric property of the prior is relaxed.

For the transition probabilities, we impose a Dirichlet distribution, as a conjugate prior, for each column of the transition matrix, P . Specifically, for column r , its density depends on $\alpha_r = (\alpha_{r1}, \dots, \alpha_{rR})$ and the prior mean of the transition probability (from regime r to r') is

$$P_{(r,r')} = \frac{\alpha_{rr'}}{\sum_{r'=1}^R \alpha_{rr'}}, \tag{16}$$

Collecting all the parameters in $\theta = (\beta, \Sigma, P, \lambda)$, the prior structure (for each regime) can be written as:

$$p(\theta) = p(\beta|\Sigma, \lambda) p(\Sigma|\lambda) p(\lambda) p(P). \tag{17}$$

2.3. Regime inference, posterior, and likelihood

This section briefly describes the algorithm for regime inference, posterior distribution of model parameters, and the computation of likelihood. In terms of notation, a variable with a superscript t denotes the entire history of the variable up to time t , for example, $y^t = (y_1, \dots, y_t)'$, and variables with a subscript S_t denote the actual observations (along with dummy observations) corresponding to the respective regime.

In essence, regime inference amounts to computing the probabilities of each regime at each time t , based on the actual observations, y^t ,

$$\xi_{t|t} = p(S_t|y^t, \theta), \tag{18}$$

where the t 'th element in $\xi_{t|t}$ represents the filtered probability of regime r , $\xi_{t|t}(r) = p(S_t = r|y^t, \theta)$. Starting with appropriate initial regime probabilities, $\xi_{0|0}$, this inference can be made recursively, through prediction and updating steps, as follows.

At the start of time t (conditional on the data up to time $t - 1$), based on the regime inference for the previous period, $\xi_{t-1|t-1}$, the predicted probabilities for time t , $\xi_{t|t-1}$ (of each regime) are given by:

$$\xi_{t|t-1} = P\xi_{t-1|t-1}. \tag{19}$$

Then, updating the information contained in y_t yields the filtered probabilities:

$$\xi_{t|t} = \frac{\xi_{t|t-1} \odot \eta_t}{\mathbf{1}' (\xi_{t|t-1} \odot \eta_t)}, \tag{20}$$

where the symbol \odot denotes element-by-element multiplication, $\mathbf{1}$ is a vector with all elements equal to 1, and η_t is a vector whose r 'th element contains the conditional likelihood for regime r :

$$\begin{aligned} \eta_t(r) &= p(y_t|S_t = r, y^{t-1}, \theta) \\ &= (2\pi)^{-n/2} |\Sigma_{S_t}|^{-1/2} \exp \left\{ -0.5 (y_t - \beta'_{S_t} X_t)' \Sigma_{S_t}^{-1} (y_t - \beta'_{S_t} X_t) \right\}. \end{aligned} \tag{21}$$

As a byproduct of this algorithm, the log likelihood function (of all the observed data) can be calculated as follows:

$$\text{Log}L(y^T|\theta) = \sum_{t=1}^T \log f(y_t|y^{t-1}, \theta), \tag{22}$$

where $f(y_t|y^{t-1}, \theta) = \mathbf{1}' (\xi_{t|t-1} \odot \eta_t)$.

In addition to these filtered probabilities, we can compute the smoothed probabilities for each regime, $p(S_t = r|y^T, \theta)$, by incorporating the full sample information, y^T . The recursive algorithm for this is summarized as follows:

$$\xi_{t|T} = \xi_{t|t} \odot [P' \xi_{t+1|T} (\div) \xi_{t+1|t}], \tag{23}$$

where (\div) denotes element by element division.⁸

We simulate the posterior distribution of the parameters, using the algorithm of random walk Metropolis-Hastings within Gibbs. For the hyperparameters, we first draw λ^* from the normal proposal distribution, $N(\lambda', c_\lambda H_\lambda^{-1})$, for each regime, where λ' is the previous draw of λ , H_λ is the Hessian of the negative of the log posterior of the hyperparameters at the peak, and c_λ is a

scaling constant. Set $\lambda = \lambda^*$ with probability α_λ ; otherwise reject the draw and set $\lambda = \lambda'$, where the acceptance probability is

$$\alpha_\lambda = \min \left(1, \frac{p_\lambda (y_{S_t} | \lambda^*)}{p_\lambda (y_{S_t} | \lambda')} \right), \tag{24}$$

and the marginal likelihood, $p_\lambda(y)$, is as given by Del Negro and Schorfheide (2011). The scaling parameter c_λ is set so that acceptance probability is around 0.3 to 0.4.

Given the result of regime inference, we can select and group the sub-sample of actual data that pertains to a particular regime, y_{S_t} and X_{S_t} . Then, the conditional posterior distribution for the VAR parameters (which is normal-inverse Wishart due to the natural conjugate prior) is given by:

$$vec(\beta_{S_t}) | \Sigma_{S_t}, \lambda, S^T, y^T \sim N \left(vec(\bar{\beta}_{S_t}), \Sigma_{S_t} \otimes \left(X_{S_t}^* X_{S_t}^{*'} \right)^{-1} \right) \tag{25}$$

and

$$\Sigma_{S_t} | \lambda, S^T, y^T \sim iW \left(S_{S_t}^*, T_{S_t}^* \right), \tag{26}$$

where $y_{S_t}^*$ and $X_{S_t}^*$ are the regime-specific actual data, augmented with the dummy observations, that is, $y_{S_t}^* = [y_{S_t}; y_{S_t}^D]$ and $X_{S_t}^* = [X_{S_t}; X_{S_t}^D]$; $\bar{\beta}_{S_t} = (X_{S_t}^* X_{S_t}^{*'})^{-1} X_{S_t}^{*'} y_{S_t}^*$; $S_{S_t}^* = (y_{S_t}^* - X_{S_t}^* \beta_{S_t})(y_{S_t}^* - X_{S_t}^* \beta_{S_t})'$; and $T_{S_t}^*$ is the length of $y_{S_t}^*$. In each draw of β_{S_t} , we discard all unstable draws and keep only stable ones.

Finally, the transition probabilities (whose posterior is, conditional on S^T , independent of the data and other parameters in the model) are sampled from the following Dirichlet distribution:

$$p \left(P_r | S^T \right) \sim Dir \left(\alpha_{r1} + \varphi_{r1}, \dots, \alpha_{rR} + \varphi_{rR} \right), \tag{27}$$

where $r = 1, 2, \dots, R$ refers to the column number of the transition probability matrix. The parameter $\varphi_{rr'}$ refers to the number of times that regime r is followed by regime r' , which can be counted using the draw of S^T .

The full MCMC estimation algorithm is summarized as follows.

- Step 0. Initialize parameters and latent state variables, $\theta^{(0)}$ and $\xi_{0|0}$.
- Step 1. Sample S^T from $p(S^T | \beta, \Sigma, \lambda, y^T)$, as described in equations (20) and (23).
- Step 2. Draw P from $p(P | S^T)$, as described in equation (27).
- Step 3. Draw λ from the normal proposal distribution, $N(\lambda', c_\lambda H_\lambda^{-1})$, as described above.
- Step 4. Draw Σ from $p(\Sigma | \lambda, y^T, S^T)$, as described in equation (26).
- Step 5. Draw β from $p(\beta | \Sigma, \lambda, y^T, S^T)$, as described in equation (25).
- Step 6. Repeat Steps 1–5 for a large number of times.

In step 0, we initialize the VAR parameters at their OLS estimates, that is, $\beta^{(0)} = \hat{\beta}_{OLS}$ and $\Sigma^{(0)} = \hat{\Sigma}_{OLS}$. We set $P^{(0)}$ at its prior mean, $\xi_{0|0}$ at the corresponding ergodic (or unconditional) probabilities of the Markov chain, and $\lambda^{(0)}$ at their posterior mode.⁹ In drawing the unobserved states, we ensure that each regime has a minimum of ten observations.

In producing the results below, we first generate 50,000 draws and discard the first 10,000. From the remaining 40,000 draws, we retain every 20th draw, resulting in 2,000 draws on which our inference is based. According to the relative numerical efficiency (RNE) (Geweke, 1992) and Raftery and Lewis (1992) diagnostics, these draws exhibit good convergence of the Markov chain.

3. Empirical results

3.1. Data and model specification

In our empirical applications, the frequency-dependent regime-switching VAR model presented in Section 2 is applied to quarterly US data on log GDP (multiplied by 100), inflation rates (measured as percentage changes in the CPI), and the federal funds rate.¹⁰ The data are obtained from the Federal Reserve Bank of St. Louis' FRED database, covering the period from 1959:Q1 to 2023:Q2. To address the issue of outliers during COVID-19, which are associated with extremely high variance that may significantly affect parameter estimates, we include dummy variables in the error covariance matrix for the period from 2020:Q1 to 2020:Q3. This adjustment provides a slightly better fit than the model without the dummies, although the overall results remain largely unchanged.¹¹

A few decisions remain regarding the model specifications, such as the number of regimes and lag length. For the lag length, we consider values ranging from $p = 2$ to 5, covering most of the commonly used options for quarterly VAR models, with a minimum of 2 lags due to prior settings. Regarding the number of regimes, while the trend-cycle decomposition literature often considers three components and thus suggests three frequency regimes (i.e., low, BC, and high) as discussed in Section 2.2, it is not clear a priori whether the third or high-frequency regime is particularly relevant or economically meaningful in the context of our regime-switching VAR model. Therefore, we also explore models with two regimes, combining the high-frequency regime with the BC-frequency regime and setting the prior for the combined regime equal to that of the BC-frequency regime. This results in 8 possible specifications. Based on the log marginal likelihoods, we find that the model with 2 regimes and 4 lags provides the best fit, and this serves as our baseline model.¹²

The specific priors are set as follows. For the VAR coefficients in the BC-frequency regime, we follow Planas et al. (2008) and Jarociński and Lenza (2018) and set $\rho = 0.7$ and $\tau = 32$, implying a typical BC lasts 8 years. For the transition probabilities, we set $\alpha_{rr'} = 8$ if $r = r'$ and $\alpha_{rr'} = 2$ otherwise.¹³ This specification ensures that each regime is treated symmetrically and exhibits reasonable persistence, with the continuation probability of $P_{(r,r')} = 0.8$. Note that, except for the VAR coefficients, the priors imposed on all other parameters are either the same or symmetric across the regimes.

3.2. Labeling and regime identification

In Bayesian inference for regime-switching models, a practical and important issue is the labeling problem (Frühwirth-Schnatter, 2001). Since the likelihood in these models is invariant with respect to permutations of the regime labels, without careful and appropriate identifying restrictions, a model may not be globally identified, which can lead to biased posterior inferences.

To address this issue, a common solution is to impose an identification restriction on the model parameters, such as imposing an order on certain parameters. For models with frequency-domain features, such as ours, another useful approach is to examine the shape of the spectral density. To illustrate this point, Figure 1 plots the spectral densities for each regime implied by these priors, along with a standard (unit root type) Minnesota prior.¹⁴

As shown in the figure, while the spectral densities decrease with frequency (except for the high-frequency regime whose densities are flat), and each regime contains other frequency components beyond its target frequency ranges, each regime appears to have distinctive spectral characteristics in terms of shape. Specifically, the spectral densities are relatively more concentrated within their respective frequency ranges, with varying degrees of decline: the low-frequency regime exhibits the steepest decline, followed by the Minnesota prior, and then the BC-frequency regime. These features allow the construction of simple and reasonable identification criteria based on the properties of spectral densities.




Figure 1. Spectral densities implied by priors.
 Note: The figure presents the spectral densities implied by the regime-specific priors. The vertical lines indicate the frequency band associated with typical business cycles, that is, $[2\pi/32, 2\pi/6]$.

Specifically, once the normalized frequency domain is divided into two parts, $[0, \tilde{\omega}]$ and $[\tilde{\omega}, 0.5]$, each representing a relatively low- and high-frequency range, respectively, we can compute the difference in the average log spectral density for regime r between these two ranges, denoted as $D(\overline{\log(SD(r))}) = \overline{\log(SD(r))}_{[0, \tilde{\omega}]} - \overline{\log(SD(r))}_{[\tilde{\omega}, 0.5]}$. This difference can be roughly interpreted as the “slope” of the density and the identification condition can then be stated in terms of this slope as follows:

$$D(\overline{\log(SD(1))}) > D(\overline{\log(SD(2))}) \tag{28}$$

After some experimentation, we find that a cutoff frequency of $\tilde{\omega} = 0.083$ (corresponding to a 12-quarter cycle) works well and gives reasonable results.¹⁵

3.3. Estimated regimes and their key properties

We now present the main empirical results. Figure 2 shows the smoothed probabilities of the BC-frequency regime from the baseline model, along with the associated spectral densities (the 5%–95% intervals of the estimated log spectrum), while Table 1 contains the estimates of the model parameters. For comparison, the results for the constant parameter VAR (CVAR) are also included in Table 1.¹⁶

Figure 2 illustrates clear frequency-dependent regime-switching behavior over time. While the low-frequency regime generally dominates, the period from the 1970s to the mid-1980s is largely characterized by the BC-frequency movements. This pattern of frequency shifts is consistent with the results of other related studies. For example, Aguiar-Conraria et al. (2012) show that, based on wavelet analysis, there are notable changes in the frequency behavior of economic activity, inflation, and interest rates during this period. Similarly, Sargent and Surico (2011) show that the low-frequency relationship between inflation and money growth exhibited distinct behavior in the 1970s and early 1980s, while Kliem et al. (2016) find a similar pattern in the relationship between public deficits and inflation. Both studies suggest that these low-frequency changes are likely related to shifts in the conduct of monetary policy over time.

While several recession episodes roughly align with the BC-frequency regime, the degree of their association varies across episodes, suggesting that their synchronization is not particularly

Table 1. Estimates for model parameters

	low frequency	BC frequency	CVAR
A. Error Covariances			
$Var(u_y)$	0.392 (0.308, 0.462)	0.968 (0.784, 1.346)	0.573 (0.522, 0.633)
$Var(u_p)$	0.145 (0.119, 0.209)	0.342 (0.224, 0.461)	0.194 (0.177, 0.214)
$Var(u_i)$	0.127 (0.073, 0.148)	0.870 (0.630, 1.208)	0.375 (0.322, 0.445)
$Cov(u_y, u_p)$	-0.027 (-0.049, -0.004)	0.160 (0.061, 0.279)	0.032 (0.011, 0.051)
$Cov(u_y, u_i)$	0.039 (0.016, 0.062)	0.240 (0.131, 0.401)	0.128 (0.092, 0.169)
$Cov(u_p, u_i)$	0.011 (-0.002, 0.028)	0.132 (0.059, 0.256)	0.055 (0.035, 0.082)
B. Hyperparameters			
λ_1	0.230 (0.126, 0.664)	0.189 (0.121, 0.450)	
λ_2	5.472 (1.655, 8.330)	2.035 (0.742, 4.141)	
λ_3	0.184 (0.118, 0.449)	0.208 (0.123, 0.598)	
λ_4	0.327 (0.115, 0.526)	4.005 (0.533, 7.724)	
C. Transition probability			
p_{11}		0.955 (0.934, 0.972)	
p_{12}		0.045 (0.028, 0.066)	
p_{21}		0.127 (0.070, 0.207)	
p_{22}		0.873 (0.793, 0.930)	

Note: The table presents the median estimates for the model parameters (in each frequency regime) along with the 16%–84% credible intervals. u_y , u_p , and u_i represent the (reduced-form) shocks in the GDP, inflation, and interest rates equations, respectively.




Figure 2. Estimated regimes and spectral densities.

Note: The left panel of the figure shows the (mean) smoothed probabilities of the BC-frequency regimes. The shaded areas represent NBER recession dates. The right panel shows the 5%–95% intervals of the estimated log spectra for each frequency regime. The vertical lines indicate the frequency band associated with typical business cycles, that is, $[2\pi/32, 2\pi/6]$.

strong. For example, the 1990–1991 and 2001 recessions are either barely or weakly identified as the BC-frequency regime by the model. Meanwhile, the 1969–1970 recession and the 2007–2009 Global Financial Crisis are identified as the BC-frequency regime, though with some lag.

The spectral densities indicate that, while a large portion of the spectra for both regimes overlaps in the low frequency regions, there is a clear distinction in their slopes. Specifically, the low-frequency regime exhibits a steeper slope and contains relatively more low-frequency components.

The sharp shift in frequency regimes (from BC-frequency to low-frequency) around the mid-1980s is reminiscent of the Great Moderation. As Panel A of Table 1 shows, the BC-frequency regime is indeed associated with higher volatility, consistent with conventional findings. The covariances also show clear differences, with shocks being more highly correlated in the

BC-frequency regime. Notably, while the covariance between output and inflation shocks is positive in both the BC-frequency regime and the CVAR, it becomes significantly negative in the low-frequency regime.

Panel B of Table 1 shows that, while the other hyperparameters are roughly similar across the two frequency regimes, there are clear differences in the two specific hyperparameters. Specifically, λ^2 , which controls for shrinkage at longer lags, is much larger in the low-frequency regime, while λ^4 , which is associated with co-persistence, is significantly smaller in the low-frequency regime. These results are intuitive and suggest that: (i) the low-frequency regime is associated with greater persistence, possibly in the form of an integrated process, where tighter shrinkage at longer lags may improve the representation of the dynamics in this regime; and (ii) co-persistence or co-integration properties between variables are more suitable for the low-frequency behavior of the series.

The transition probabilities (Panel C in Table 1) show strong persistence for both regimes, with the low-frequency regime having a higher median continuation probability of 0.955, compared to 0.873 for the BC-frequency regime.

To further illustrate the regime-dependent dynamics implied by the model, Figure 3 presents the impulse responses across the regimes, along with those from the CVAR. The identification is based on a simple Cholesky decomposition, with the ordering of output, inflation, and interest rate (the last being the most exogenous). While it is somewhat unclear to discern the economic implications for some of the responses under this identification scheme, the figure overall reveals sharply contrasting results across the regimes. Specifically, the responses in each regime are characterized by strong persistence and clear cyclical patterns, respectively, with the responses from the CVAR falling between those of the two regimes. Beyond these differences, there are several notable features across the regimes, which we briefly discuss below.

First, in the case of an interest rate shock (which can be viewed as a rudimentary form of a monetary policy shock), both output and inflation increase persistently and significantly in the low-frequency regime. In contrast, in the BC-frequency regime, output decreases significantly for up to six years, while the responses of inflation, aside from an initial short-term increase, remain largely insignificant over a 2- to 10-year horizon. While the literature often attributes the price puzzle, that is, the strong positive response of prices following an interest rate shock, to missing information or omitted variables in VARs, this study suggests that another contributing factor may be the failure to properly extract relevant frequency-domain information about the shock and the associated variable dynamics.

Next, following an inflation shock, while output declines persistently and significantly in the BC-frequency regime (much more so than in the CVAR), its responses in the low-frequency regime are much smaller and become insignificant after the first several quarters. Finally, for an output shock, while the responses of output in the low-frequency regime are strongly persistent (more so than those in the CVAR), the responses in the BC-frequency regime are relatively transitory and become insignificant after a few years. Interestingly, in the low-frequency regime, the short-run responses of inflation to an output shock are negative, although they quickly turn positive after a few quarters, in contrast to the BC-frequency regime and the CVAR, where its short-run responses are significantly positive. This difference in the initial responses reflects the results of the regime-dependent error covariance matrix, as reported in Table 1.

4. Alternative model specifications

In this section, we explore several alternative model specifications, including models with different volatility specifications and priors, as well as conventional (time-domain) regime-switching models. This analysis serves two main purposes. First, by exploring alternative models with different variance specifications and priors, and comparing them to the baseline model, we aim to better understand the key features of our model and the factors driving the results.




Figure 3. Impulse responses across frequency regimes.

Note: The figure shows the 16%–84% bands of the impulse responses across frequency regimes (blue solid line), along with those of the constant parameter VAR (red dashed line).

Second, because regime-switching models, typically studied in the time domain, are widely used to identify recession periods, and the estimated regimes from the baseline model show some (albeit rather weak) association with recession periods, it is important to compare this frequency-dependent switching model with conventional time-domain regime-switching models. This comparison will emphasize the similarities and differences between the two approaches and provide further insights into the nature of the frequency-dependent regime-switching model.

4.1. List of alternative models

This subsection briefly outlines the main features of the alternative models, with detailed descriptions provided in the Appendix. First, to assess the role of regime-dependent heteroskedasticity in the disturbances, we consider regime-switching VAR models with either homoskedastic shocks or

stochastic volatility. In the homoskedastic model (denoted as RS-VAR-homo), the error covariances are constrained to be the same in both regimes (i.e., $\Sigma_{S_t} = \Sigma$ for all t), while all other features, including the prior means of the VAR coefficients, remain the same as in the baseline model (hereafter denoted as RS-VAR-FD).

For the model with stochastic volatility (denoted as RS-VAR-SV), we follow the approach outlined in the literature (e.g., Primiceri, 2005). In this model, the (co)variances of the innovations are assumed to (gradually) evolve over time rather than switching between regimes while all other features remain the same as in the baseline model. Specifically, the model relies on the standard decomposition of the (time-varying) error covariance matrix, $\Sigma_t = A_t^{-1} H_t A_t^{-1'}$, where A_t is a lower triangular matrix and H_t is a diagonal matrix. We assume that the evolution of the diagonal elements of H_t and the elements of A_t follows a (geometric) random walk.

To address the issue of symmetry across the equations associated with the natural conjugate prior, discussed in Section 2.2, we also consider a model where this restrictive structure is relaxed. In specifying this model (denoted as RS-VAR-AP), we largely follow the approach of Carriero et al. (2019) and Carriero et al. (2022) and allow each equation to be treated and estimated individually. Thus, in this model, each equation as well as each regime is governed by a set of separate hyperparameters.

We also estimate a conventional regime-switching VAR model with standard time-domain features (denoted as RS-VAR-TD) to identify expansion and recession regimes. While retaining the same essential structure as the RS-VAR-FD, this model differs in several key respects. First, for the VAR coefficients, following the literature (e.g., Sims and Zha, 2006; Sims et al. 2008), we impose a unit root-type Minnesota prior on the VAR coefficients in both regimes. With this symmetric prior, the dynamic properties of each regime are primarily determined by the data. Second, regime identification in this model relies on typical time-domain characteristics of recessions, such as lower output growth and higher volatility. Thus, the associated identification condition (in terms of model parameters) is that, in the recession regime, the intercept in the GDP equation is lower and the determinant of the error covariance matrix is larger.

Diebold et al. (1994) and Filardo (1994) argue that constant transition probabilities in regime-switching models may be too restrictive for many empirical contexts. In the spirit of these studies, we also examine models with time-varying transition probabilities which are specified as functions of lagged output growth, in both the baseline and conventional (time-domain) models (each denoted as RS-VAR-FD-TVTP and RS-VAR-TD-TVTP, respectively). Given that both the transition probabilities and the regime inferences in these models are based on past output dynamics, this approach may provide additional insights into the nature of regime-switching dynamics.

The alternative models are summarized as follows:

- RS-VAR-homo: Frequency-dependent regime-switching VAR model with homoskedasticity
- RS-VAR-SV: Frequency-dependent regime-switching VAR model with stochastic volatility
- RS-VAR-AP: Frequency-dependent regime-switching VAR model with an asymmetric prior
- RS-VAR-TD: Conventional time-domain regime-switching VAR models with constant transition probabilities
- RS-VAR-TD-TVTP: Conventional time-domain regime-switching VAR models, both with constant and time-varying transition probabilities
- RS-VAR-FD-TVTP: Baseline frequency-dependent regime-switching VAR model with time-varying transition probabilities




Figure 4. Results for alternative models.

Note: Each panel shows the result for the three alternative models. Panel (a) The model with homoskedasticity; panel (b) The model with stochastic volatility; panel (c) the model with asymmetric prior. See notes to Figure 2.

4.2. Comparison with alternative specifications in variances and priors

The main results for these alternative models are summarized in Figures 4 and 5, which show the smoothed regime probabilities and the associated spectral densities, and in Tables 2 and 3, which show the estimates of selected model parameters.

In the case of the RS-VAR-homo (Figure 4(a)), the number of periods identified as belonging to the BC-frequency regime is significantly smaller compared to the RS-VAR-FD; many periods previously classified as part of the BC-frequency regime in the RS-VAR-FD are now identified as belonging to the low-frequency regime. This is because, unlike in the RS-VAR-FD (where the likelihood is relatively flatter and more dispersed in the BC-frequency regime, while more concentrated in the low-frequency regime), imposing homoskedasticity increases the likelihood of a typical shock or data point favoring the low-frequency regime. In other words, data that might have been identified as belonging to the BC-frequency regime in the RS-VAR-FD are now more likely to be classified as belonging to the low-frequency regime.

In the case of the RS-VAR-SV (Figure 4(b)), although some weak patterns of frequency switching are observed during the 1970s to 1980s, the distinction between regimes is less clear. This result

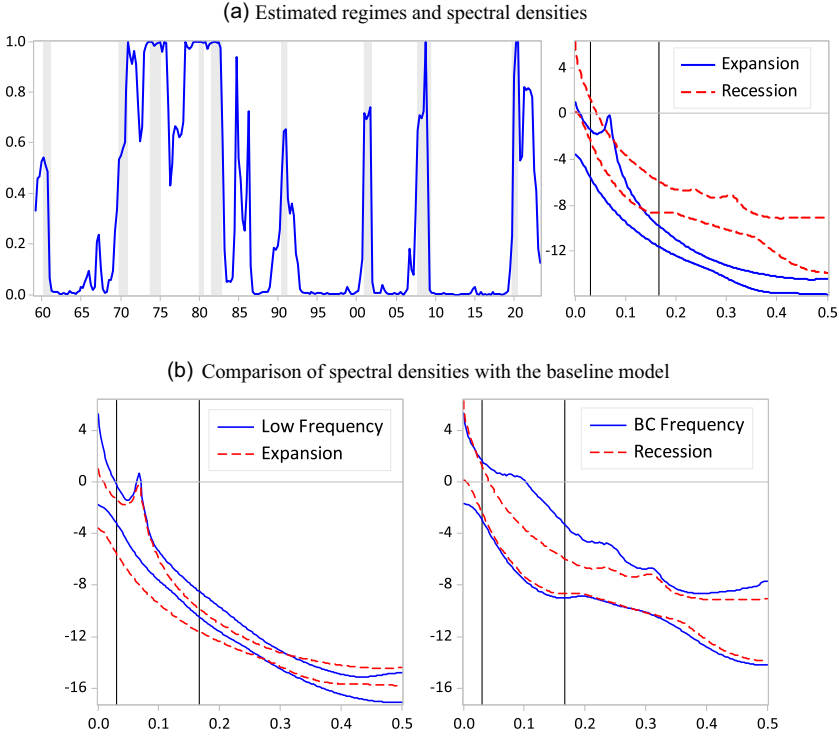


Figure 5. Results for the conventional (Time-domain) regime-switching model (RS-VAR-TD).
Notes: See notes to Figure 2. The figure plots the 5%–95% intervals of the estimated log spectra for the low-frequency and the expansion regime (left panel) and the BC-frequency and the recession regime (right panel).

seems to be due to two factors. First, the regime-specific errors in this model are likely to have non-normal and asymmetric distributions. As noted by Hwu and Kim (2024), when the distribution of the error term deviates from normality, especially in cases of asymmetry, the accuracy of regime probability inference deteriorates significantly. In addition, the lack of tightly synchronized heteroskedasticity across the series (as shown in Appendix Figure A3) may further complicate clear regime inference.¹⁷

Figure 4(c) shows the results for the RS-VAR-AP. Similar to the RS-VAR-homo and the RS-VAR-SV, this model does not provide clear evidence of distinct frequency switching. As with the RS-VAR-SV, the lack of a clear distinction between regimes is likely because the dynamics of each series are not tightly synchronized in terms of their frequency-domain behavior. To illustrate this, we apply the baseline model to each series, with the estimated regimes shown in Appendix Figure A4. Each series exhibits distinct frequency-dependent switching patterns, with common movements observed only during limited periods.

As shown in the figure, in all these models, the spectral densities of the two regimes are not clearly separated and tend to overlap across the entire frequency range, which is consistent with the poor or unclear regime inference.

4.3. Comparison with conventional regime-switching models

We now compare the results of the RS-VAR-TD (shown in Figure 5 and Table 3) with those of the RS-VAR-FD. While the overall results appear similar in terms of the parameter estimates and the estimated latent regimes, a closer inspection reveals several notable and important differences.

Table 2. Estimates of parameters for alternative models

(a) RS-VAR-homo				
	A. Hyperparameters		B. Transition probability	
	Low frequency	BC frequency		
λ_1	0.229 (0.122, 0.879)	0.185 (0.119, 0.475)	p_{11}	0.956 (0.937, 0.971)
λ_2	7.088 (2.639, 8.976)	4.561 (1.433, 7.989)	p_{12}	0.044 (0.029, 0.063)
λ_3	0.220 (0.119, 0.682)	0.226 (0.123, 0.709)	p_{21}	0.512 (0.335, 0.647)
λ_4	1.779 (0.196, 7.047)	5.000 (1.686, 8.221)	p_{22}	0.488 (0.353, 0.665)
C. Error Covariances				
$Var(u_y)$		0.586 (0.530, 0.650)	$Cov(u_y, u_p)$	-0.001 (-0.015, 0.012)
$Var(u_p)$		0.066 (0.059, 0.085)	$Cov(u_y, u_i)$	0.100 (0.070, 0.130)
$Var(u_i)$		0.377 (0.331, 0.430)	$Cov(u_p, u_i)$	0.011 (0.002, 0.022)
(b) RS-VAR-SV				
	A. Hyperparameters		B. Transition probability	
	Low frequency	BC frequency		
λ_1	0.165 (0.121, 0.530)	0.147 (0.114, 0.280)	p_{11}	0.917 (0.778, 0.956)
λ_2	4.899 (1.351, 8.392)	2.422 (0.760, 7.228)	p_{12}	0.083 (0.044, 0.222)
λ_3	0.181 (0.125, 0.741)	0.167 (0.113, 0.424)	p_{21}	0.284 (0.160, 0.420)
λ_4	0.311 (0.114, 0.820)	1.875 (0.396, 5.980)	p_{22}	0.716 (0.580, 0.840)

Note: The table presents the median estimates for parameters in each frequency regime, along with the 16%–84% credible intervals. See the notes to Table 1.

Table 3. Estimates of parameters for the RS-VAR-TD

	Expansion	Recession
A. Error Covariances		
$Var(u_y)$	0.332 (0.261, 0.417)	0.944 (0.797, 1.116)
$Var(u_p)$	0.117 (0.104, 0.131)	0.329 (0.271, 0.440)
$Var(u_i)$	0.126 (0.101, 0.144)	0.853 (0.632, 1.217)
$Cov(u_y, u_p)$	-0.033 (-0.049, -0.017)	0.151 (0.085, 0.257)
$Cov(u_y, u_i)$	0.051 (0.035, 0.068)	0.265 (0.164, 0.425)
$Cov(u_p, u_i)$	0.007 (-0.002, 0.017)	0.140 (0.075, 0.255)
B. Hyperparameters		
λ_1	0.232 (0.130, 0.764)	0.193 (0.124, 0.517)
λ_2	6.651 (2.462, 8.880)	3.902 (1.066, 7.902)
λ_3	0.209 (0.120, 0.560)	0.227 (0.127, 0.730)
λ_4	0.774 (0.209, 6.354)	2.982 (0.633, 7.820)
C. Transition probability		
p_{11}	0.943 (0.916, 0.964)	
p_{12}	0.057 (0.036, 0.084)	
p_{21}	0.124 (0.081, 0.183)	
p_{22}	0.876 (0.817, 0.919)	

Note: The table presents the median estimates for parameters in each regime (expansion and recession), along with the 16%–84% credible intervals. See the notes to Table 1.

First, the identified recession regimes (in the RS-VAR-TD, Panel (a) of Figure 5) are more closely aligned with the NBER reference cycles. For example, while the RS-VAR-FD classifies the periods of the 1990–1991 and 2001 recessions as either in the low-frequency regime or in between the two frequency regimes, respectively, the RS-VAR-TD identifies these periods as the recession regime with much higher probabilities. A similar observation can be made for the early periods of the 1969–1970 and 2007–2009 recessions.

The differences in regime inference between the two models can be attributed to their different approaches. In the RS-VAR-TD, regime identification is essentially based on the first two moments of the data (i.e., mean and volatility), whereas in the RS-VAR-FD, the transition dynamics around recessions play a more important role in regime inference. For example, if the transition into and out of a recession is relatively smooth and entails more low-frequency features, the recession period is more likely to be classified as a low-frequency regime.¹⁸

In this regard, comparing the spectral properties of the two models reveals further differences. Panel (b) of Figure 5 compares the spectral densities of both models in each regime (i.e., low-frequency/expansion and BC-frequency/recession, respectively). While the shapes of the spectral densities for the second regime (BC-frequency/recession) are largely similar across the two models (with slightly wider intervals in the RS-VAR-FD), there are notable differences in the spectral densities for the first regime (low-frequency/expansion). Specifically, the spectral densities in the expansion regime are relatively flatter compared to those in the low-frequency regime. This suggests that, although both regimes contain a substantial amount of low-frequency components, the low-frequency regime contains relatively more low-frequency components, implying relatively more persistence, whereas the expansion regime features relatively more high-frequency components, indicating relatively less persistence.^{19,20}

In addition, as shown in Figure 5, although the spectral densities in the RS-VAR-TD do not show clear differences in slope between the two regimes, a substantial portion of the spectra in the recession regime appear to lie above those in the expansion regime. This suggests that, in a rough sense, the spectral densities in the recession regime can be seen as a parallel upward shift from those in the expansion regime. In other words, the differences in the spectral characteristics between the two regimes in the RS-VAR-TD are influenced more by the error covariances than by the VAR coefficients.

Given the several differences between the two types of regime-switching models, an interesting question arises: What would happen if we applied the recession identification criteria (in the time-domain model) to the RS-VAR-FD? Conversely, what would happen if we applied the frequency-domain identification criteria (based on spectral density shapes) to the RS-VAR-TD? Exploring these questions could provide valuable insights into the role of prior specifications and regime identification schemes, potentially revealing further differences between the two types of models.

Figure 6 shows the results of these exercises. When the (time-domain) recession identification is applied to the RS-VAR-FD, the identified regimes are broadly similar to those in the original RS-VAR-FD (Figure 6(a)). This is not surprising, as the BC-frequency regime is also associated with higher volatilities, with the minor differences due to the additional consideration of the intercept term in the GDP equation. In contrast, applying the frequency-domain identification to the RS-VAR-TD results in poor regime inference, making the results difficult to interpret (Figure 6(b)). This result is largely due to the fact that, as discussed earlier, the two regimes in the RS-VAR-TD model are not clearly distinguished in terms of the slope of the spectral densities.²¹

These results suggest that, despite some similarities between the two types of models in terms of the volatilities and persistence of each regime, one cannot directly associate expansion/recession periods with low-/BC-frequency regimes, as each of their spectral properties appears to be different between the two models. Moreover, because of these differences, natural frequency-domain identification cannot simply be applied to identify the recession regime as in conventional time-domain models; careful and model-specific identification is required in each case.




Figure 6. Results from alternative identification.

Note: In each panel, the figure shows the probability of the BC-frequency (panel (a)) or recession regime (panel (b)).

Finally, when the transition probabilities are allowed to vary over time, the estimated regimes remain broadly similar to those obtained with constant transition probabilities (Appendix Figure A7). However, the distinction between regimes becomes clearer and more consistent with the intended regime properties of each model, which is also reflected in a slightly improved model fit, as indicated by the log marginal likelihoods.

The spectral densities for these models show that, in the second regime (BC-frequency/recession), there are notable peaks around 0.1 for the BC-frequency and 0.3 for the recession regime, corresponding to 10-quarter and 4-quarter cycles, respectively. This suggests that time-varying transition probabilities introduce more uncertainty around these frequency ranges, where the associated frequencies are also relatively lower in the frequency-dependent switching model (RS-VAR-FD-TVTP).

4.4. Forecasting performance comparison

Finally, this subsection presents a simple comparison of forecasting performance. Although the main objective of the study is not to examine the predictive ability of the model, evaluating the predictive performance of econometric models has become standard practice, and the forecast evaluation can be thought of as a useful tool for model comparison.

Therefore, we construct a small pool of models for forecasting comparison, including CVAR, RS-VAR-FD, RS-VAR-TD, RS-VAR-TVTP, and RS-VAR-TD-TVTP.²² For all models, we use a recursive estimation window to generate pseudo out-of-sample forecasts for $h = 1, 2, 3$, and

Table 4. Forecast comparison

	RMSE			log predictive likelihood
	GDP	inflation	interest rates	
<i>h = 1</i>				
CVAR	<i>1.530</i>	<i>0.580</i>	<i>0.708</i>	-3.959
RS-VAR-FD	0.996	0.974*	0.945**	0.176**
RS-VAR-TD	0.980*	0.953**	0.917**	0.146**
RS-VAR-FD-TVTP	0.983†	1.042	0.954**	0.334†**
RS-VAR-TD-TVTP	0.989*	0.962**	0.930**	0.142*
<i>h = 2</i>				
CVAR	<i>1.762</i>	<i>0.970</i>	<i>0.991</i>	-4.604
RS-VAR-FD	0.998	0.956*	0.870**	0.503**
RS-VAR-TD	0.964**	0.943*	0.872**	0.465**
RS-VAR-FD-TVTP	1.041	1.094	1.002	-0.416
RS-VAR-TD-TVTP	0.980	0.940*	0.883**	0.421**
<i>h = 3</i>				
CVAR	<i>2.278</i>	<i>1.755</i>	<i>1.253</i>	-6.661
RS-VAR-FD	1.089	0.943*	0.981*	-1.670
RS-VAR-TD	0.974†**	0.914**	0.958†**	-0.453
RS-VAR-FD-TVTP	1.137	0.928*	0.990*	-0.043†
RS-VAR-TD-TVTP	0.981†*	0.945*	0.956†**	-0.257†
<i>h = 4</i>				
CVAR	<i>2.890</i>	<i>4.012</i>	<i>1.481</i>	-9.198
RS-VAR-FD	1.076	0.921**	0.988	-3.314
RS-VAR-TD	0.976†**	0.912**	0.967*	-1.812
RS-VAR-FD-TVTP	1.117	0.774††**	0.990	0.363††
RS-VAR-TD-TVTP	0.975†**	0.952*	0.953†**	-1.266

Notes: At each forecast horizon, the row corresponding to CVAR contains the RMSEs and the average log predictive likelihood (in italics) while other entries are values relative to them. Asterisks (*) and daggers (†) indicate the significance of the Giacomini and White (2006) predictive ability test against the CVAR at the * 0.10 and ** 0.05 levels, and against the RS-VAR-FD at the † 0.10 and †† 0.05 levels, respectively.

4-step-ahead horizons. Starting with the estimation sample up to 1983:Q4, the first available forecasts are for 1983:Q4+h. We then proceed by expanding the sample by a quarter and generating forecasts through the end of the sample. In line with conventional practices and the interests of policymakers, h quarter-ahead forecasts for GDP and inflation are the cumulative quarterly rates of change (i.e., the percentage change from period t through t + h). All the forecasts from each model are transformed to be consistent with this cumulative data measure. The exercises involve multi-step forecasting, and we use iterated forecasting to obtain the result, which generally outperforms direct forecasting (Marcellino et al. 2006). In evaluating the forecasting performance of the models, we use the root mean squared error (RMSE) and the predictive likelihood as evaluation metrics. In the latter, we measure the log-predictive score using a Gaussian approximation of the predictive density.

The results of this forecast comparison are summarized in Table 4. We normalize all the results in the table relative to the CVAR. Specifically, the numbers represent the ratio of each model’s RMSE to that of the CVAR and the average of the log predictive likelihoods of a given model minus that of the CVAR, respectively. In addition, to provide a rough indication of the statistical

significance of the differences in predictive performance, we report the results of the predictive ability test (Giacomini and White, 2006) for each model against both the CVAR and the RS-VAR-FD.

As shown in the table, forecasting with the RS-VAR-FD leads to substantial improvement over the CVAR in most cases. The improvement is more pronounced for inflation and interest rates, as well as for the density forecast over short horizons, and is often accompanied by strong statistical significance over the CVAR. Compared to the RS-VAR-TD, the overall performance of the RS-VAR-FD is somewhat worse, especially for point forecasts. However, the predictive ability test results show that their performances are not significantly different in most cases and that the RS-VAR-FD remains strong and competitive over short horizons ($h = 1$ and $h = 2$). Despite their better in-sample fit, the models with time-varying transition probabilities perform slightly worse than their constant probability counterparts in several cases, indicating the possibility of overfitting.

5. Concluding remarks

This study presents a simple VAR model to analyze the frequency-dependent regime-switching dynamics of macroeconomic time series. Empirical applications to some key macroeconomic variables demonstrate clear regime-switching behavior between the low- and BC-frequency regimes. We also show that each frequency regime exhibits distinctive characteristics in terms of spectral properties, persistence, volatility, and impulse responses.

We conclude this section by suggesting some avenues for future research. First, while several recent studies, including this one, have highlighted a significant shift in the frequency-domain dynamics of key macroeconomic variables from the 1970s to the mid-1980s, the nature and source of this change remain unclear. In particular, although this shift appears to be associated with the Great Moderation, the reduction in volatility is only one aspect of the broader changes in the frequency domain. Therefore, a more in-depth investigation of the nature and underlying mechanisms driving this shift is warranted.

Second, while this study assumes discrete changes in the frequency domain, the time variation might be better characterized by smoother or gradual changes. Examples of such approaches include VARs with threshold and smooth-transition and (gradually changing) time-varying coefficients (along with stochastic volatilities). A rigorous comparison between the model presented in this study and these alternative models would be a valuable area for future research and could provide deeper insights into the nature of time-varying changes in macroeconomic time series.

Notes

1 For example, in structural VARs, neglecting these features leads to poor identification of shocks or the compounding of the different effects of shocks at different frequencies (Francis and Ramey, 2009).

2 For a comprehensive survey of these priors, see Koop and Korobilis (2010), Karlsson (2013), and Chan (2020b).

3 For an overview of wavelet-based analysis and its applications in economics, see Aguiar-Conraria and Soares (2014).

4 In empirical trend-cycle decomposition, the evolution of the trend component in (3) is sometimes augmented with an additional disturbance term, ψ_t , that is, $y_t^r = y_{t-1}^r + \mu_t + \psi_t$, so that (5) becomes $\Delta^2 y_t^r = \delta_t + \psi_t - \psi_{t-1}$. In this case, the trend component becomes relatively less smooth. Harvey and Jaeger (1993) discuss the empirical relevance of including this additional disturbance term and show that a model without this term provides a better fit for U.S. output. For comparison, we also consider later the case of an I(1) prior instead of an I(2) prior for the low-frequency VAR coefficients.

5 As noted by Jarociński and Lenza (2018), specifying the priors in the frequency domain using the cosine function, rather than in terms of the VAR coefficients, is more intuitive and imposes the dynamics implied by the priors in a more natural way.

6 While we assume three regimes here, we discuss the choice of the appropriate number of regimes and consider other possibilities in Section 3.1.

- 7 The only exception is λ_5 , which is fixed to 100, indicating a very loose prior.
- 8 For a more detailed discussion of the filtering and smoothing algorithm, refer to Frürwirth-Schnatter (2006) and references therein.
- 9 We have also tried other values such as prior means for β and Σ , and equal probabilities for $P^{(0)}$ and $\xi_{0|0}$, but found little difference in the results.
- 10 Considering that the GDP deflator is also a commonly used measure of the overall price level and may have different dynamic properties compared to the CPI, we also estimate models using inflation measured by the GDP deflator. The results, as shown in the Appendix (Table A2 and Figures A1–A2), are broadly similar to those obtained using the CPI.
- 11 In Section 4, we also examine a model in which the changes in error variances over time are specified in terms of stochastic volatility (without COVID-19 dummies).
- 12 The log marginal likelihoods for these alternative specifications are reported in Appendix Table A1.
- 13 For the three-regime case ($R = 3$), we set $\alpha_{rr'} = 16$ if $r = r'$ and $\alpha_{rr'} = 2$ otherwise.
- 14 The spectral densities for all models are computed by evaluating the determinant of the spectral density matrix (implied by each prior or model) at 200 equally spaced points in the normalized frequency domain, $[0, 0.5]$.
- 15 We also find that other cutoffs, such as $\tilde{\omega} = 0.05$ (corresponding to a 20-quarter cycle), give very similar results. As will be seen later, because the frequency spectra occasionally exhibit non-decreasing behavior (such as local peaks around the medium frequency ranges), it seems advisable to use a conservative cutoff value to account for these variations.
- 16 The CVAR model is a standard VAR without regime-switching features; further details are provided in the Appendix.
- 17 As shown in Table 2, the contrasting patterns in hyperparameters across the regimes observed in the RS-VAR-FD are still present in these alternative models.
- 18 In fact, the two recessions of 1990–1991 and 2001 are generally considered to be relatively mild and short and, thus, standard regime-switching and recession forecasting models often fail to detect or predict them. See Dueker (2002), Stock and Watson (2003), and Kauppi and Saikkonen (2008) for further discussions.
- 19 These differences between the two types of models are also reflected in the impulse responses (shown in Appendix Figure A5), which are largely consistent with their respective spectral properties. While the responses in the second regime (BC-frequency/recession) are similar in shape across both models, with slightly wider credible intervals in the RS-VAR-FD, there are several differences in persistence in the first regime (low-frequency/expansion), depending on the response variables. Specifically, output responses tend to be significantly stronger and more persistent in the low-frequency regime, while the interest rate exhibits slightly more persistent behavior in the expansion regime.
- 20 As part of a robustness check and to assess the effect of different persistence in the prior for the first regime (i.e., I(2) for the low frequency regime and I(1) for the expansion regime), we also estimate the baseline model with the low-frequency prior means replaced by a unit root-type Minnesota prior instead of I(2) (denoted as RS-VAR-FD-I(1)). The results (shown in Appendix Figure A6) are very similar to those of the baseline model, suggesting that the dynamics around several of the recession episodes discussed above are not consistent with the BC-frequency characteristics.
- 21 Therefore, the results do not change significantly when the cutoff frequency is adjusted to other values. One way to obtain results similar to the RS-VAR-TD using frequency-domain criteria is to impose an identification condition over relatively high-frequency ranges. For example, if the regime identification is imposed so that the average log spectral density in the recession regime is higher than in the expansion regime over the frequency band such as $[0.167, 0.5]$ (i.e., cycles with periodicity of less than 6 quarters), the resulting regimes are essentially the same as those of the RS-VAR-TD. However, as noted above, this identification is essentially based on differences in volatility rather than on the spectral shape of each regime.
- 22 While the previous sections also examine other specifications, including models with stochastic volatility and asymmetric priors, these models exhibit poor performance overall—not only in terms of in-sample fit, as shown above, but also in out-of-sample forecasting (often performing worse than the CVAR). Therefore, to save space, we focus here on models with forecasting performance that is reasonable and comparable to the baseline model.

References

- Aguiar-Conraria, L., M. M. F. Martins and M. J. Soares (2012) The yield curve and the macro-economy across time and frequencies. *Journal of Economic Dynamics and Control* 36(12), 1950–1970.
- Aguiar-Conraria, L., M. M. F. Martins and M. J. Soares (2018) Estimating the Taylor rule in the time-frequency domain. *Journal of Macroeconomics* 57, 122–137.
- Aguiar-Conraria, L. and M. J. Soares (2014) The continuous wavelet transform: moving beyond uni- and bivariate analysis. *Journal of Economic Surveys* 28(2), 344–375.
- Albert, J. and Chib, S. (1993) Bayesian Analysis of Binary and Polychotomous Response Data. *Journal of the American Statistical Association* 88(422), 669–679.
- Andrle, M. and M. Plašil (2018) Econometrics with system priors. *Economics Letters* 172(C), 134–137.
- Blake, N. S. and T. B. Fomby (1997) Threshold cointegration. *International Economic Review* 38(3), 627–645.
- Barsky, R. B. and E. R. Sims (2011) News shocks and business cycles. *Journal of Monetary Economics* 58(3), 273–289.

- Caraiani, P. (2015) Estimating DSGE models across time and frequency. *Journal of Macroeconomics* 44, 33–49.
- Carriero, A., J. Chan, T. Clark and M. Marcellino (2022) Corrigendum to, large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors, [J. Econometrics 212 (1)(2019) 137-154]. *Journal of Econometrics* 227 (2), 506–512.
- Carriero, A., T. E. Clark and M. Marcellino (2019) Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors. *Journal of Econometrics* 212(1), 137–154.
- Chan, J. C. C. (2020a) Large Bayesian VARs: a flexible kronecker error covariance structure. *Journal of Business & Economic Statistics* 38(1), 68–79.
- Chan, J. (2020b) Large Bayesian vector autoregressions, In: Chan, J. (ed.), *Macroeconomic Forecasting in the Era of Big Data*, pp. 95–125. Springer.
- Chan, J. C. C. (2022) Asymmetric conjugate priors for large Bayesian VARs. *Quantitative Economics* 13(3), 1145–1169.
- Chan, J. and Eisenstat, E. (2018) Bayesian model comparison for time-varying parameter VARs with stochastic volatility. *Journal of Applied Econometrics* 33(4), 509–532.
- Clark, P. K. (1987) The cyclical component of U. S. economic activity. *Quarterly Journal of Economics* 102(4), 797–814.
- Cogley, T. and T. Sargent (2005) Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics* 8(2), 262–302.
- Del Negro, M. and F. Schorfheide (2011). Bayesian macroeconometrics. In: Del Negro, M. and F. Schorfheide. (eds.), *The Oxford Handbook of Bayesian Econometrics*, pp. 293–389. Oxford: Oxford University Press.
- Diebold, F., J. Lee and G. Weinbach (1994) Regime switching with time-varying transition probabilities. In: Diebold, F., J. Lee and G. Weinbach. (ed.), *Nonstationary Time Series Analysis and Cointegration*, pp. 283–302, Oxford University Press.
- Dueker, M. J. (2002) Regime-dependent recession forecasts and the 2001 recession. *Federal Reserve Bank of Saint Louis Review* 84(6), 29–36.
- Filardo, A. J. (1994) Business-cycle phases and their transitional dynamics. *Journal of Business & Economic Statistics* 12(3), 299–308.
- Filardo, A. and Gordon, S. (1998) Business cycle durations. *Journal of Econometrics* 85(1), 99–123.
- Francis, N. and V. A. Ramey (2005) Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited. *Journal Monetary Economics* 52(8), 1379–1399.
- Frühwirth-Schnatter, S. (2001) Markov chain monte carlo estimation of classical and dynamic switching and mixture models. *Journal of the American Statistical Association* 96(453), 194–209.
- Frühwirth-Schnatter, S. (2006) Finite mixture and Markov switching models. New York, NY: Springer.
- Geweke, J. (1992) Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In: Geweke, J. (eds.), *Bayesian Statistics*, pp.169–193. Oxford: Clarendon Press.
- Giacomini, R. and H. White (2006) Tests of conditional predictive ability. *Econometrica* 74(6), 1545–1578.
- Giannone, D., M. Lenza and G. Primiceri (2015) Prior selection for vector autoregressions. *Review of Economics and Statistics* 97(2), 436–451.
- Giannone, D., M. Lenza and G. E. Primiceri (2019) Priors for the long run. *Journal of the American Statistical Association* 114(526), 565–580.
- Harvey, A. C. (1985) Trends and cycles in macroeconomic time series. *Journal of Business and Economic Statistics* 3(3), 216–227.
- Harvey, A. C. and A. Jaeger (1993) Detrending, stylized facts and the business cycle. *Journal of Applied Econometrics* 8(3), 231–247.
- Hwu, S.-T. and C.-J. Kim (2024) Markov-switching models with unknown error distributions: identification and inference within the bayesian framework. *Studies in Nonlinear Dynamics & Econometrics* 28(2), 177–199.
- Jarociński, M. and M. Lenza (2018) An inflation-predicting measure of the output gap in the euro area. *Journal of Money, Credit and Banking* 50(6), 1189–1224.
- Kadiyala, K. and S. Karlsson (1997) Numerical methods for estimation and inference in bayesian VAR-models. *Journal of Applied Econometrics* 12(2), 99–132.
- Karlsson, S. (2013) Forecasting with Bayesian vector autoregressions. In: Karlsson, S. (eds.), *Handbook of Economic Forecasting*, pp. 791–897. Amsterdam: Elsevier.
- Kauppi, H. and P. Saikkonen (2008) Predicting US recessions with dynamic binary response models. *Review of Economic Statistics* 90(4), 777–791.
- Kim, C.-J. and C. R. Nelson (1999) Has the U.S. economy become more stable? A bayesian approach based on a Markov-switching model of the business cycle. *The Review of Economics and Statistics* 81(4), 608–616.
- Kliem, M., A. Kriwoluzky and S. Sarferaz (2016) On the low-frequency relationship between public deficits and inflation. *Journal of Applied Econometrics* 31(3), 566–583.
- Koop, G. and D. Korobilis (2010) Bayesian multivariate time series methods for empirical macroeconomics. *Foundations and Trends in Econometrics* 3(4), 267–358.
- Krolzig, H. (1997) *Markov-Switching Vector Autoregressions: Modelling, Statistical Inference, and Application to Business Cycle Analysis*. Berlin: Springer-Verlag.

- Kurmann, A. and E. Sims (2021) Revisions in utilization-adjusted TFP and robust identification of news shocks. *Review of Economic Statistics* 103(2), 216–235.
- Litterman, R. B. (1986) Forecasting with bayesian vector autoregressions: five years of experience. *Journal of Business & Economic Statistics* 4(1), 25–38.
- Marcellino, M., J. H. Stock and M. W. Watson (2006) A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series. *Journal of Econometrics* 135(1-2), 499–526.
- McConnell, M. M. and G. Perez-Quiros (2000) Output fluctuations in the United States: what has changed since the early 1980's? *American Economic Review* 90(5), 1464–1476.
- Planas, C., A. Rossi and G. Fiorentini (2008) Bayesian analysis of the output gap. *Journal of Business & Economic Statistics* 26(1), 18–32.
- Primiceri, G. E. (2005) Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies* 72(3), 821–852.
- Raftery, A. and S. Lewis (1992) How many iterations in the Gibbs sampler?. In: Raftery, A. and S. Lewis. (eds.), *Bayesian Statistics*, pp. 763–773, Oxford University Press.
- Sala, L. (2015) DSGE models in the frequency domains. *Journal of Applied Econometrics* 30(2), 219–240.
- Sargent, T. J. and P. Surico (2011) Two illustrations of the quantity theory of money: breakdowns and revivals. *American Economic Review* 101(1), 109–128.
- Sims, C. A., D. F. Waggoner and T. Zha (2008) Methods for inference in large multiple-equation Markov-switching models. *Journal of Econometrics* 146(2), 255–274.
- Sims, C. A. and T. Zha (1998) Bayesian methods for dynamic multivariate models. *International Economic Review* 39(4), 949–968.
- Sims, C. A. and T. Zha (2006) Were there regime switches in U.S. monetary policy? *American Economic Review* 96(1), 54–81.
- Smets, F. and R. Wouters (2007) Shocks and frictions in US business cycles: a bayesian DSGE approach. *American Economic Review* 97(3), 586–606.
- Stock, J. and M. Watson (2003) How did leading indicator forecasts perform during the 2001 recession? *Federal Reserve Bank of Richmond Economic Quarterly* 89(3), 71–90, Summer.
- Teräsvirta, T., D. Tjøstheim and C. Granger (2010) *Modelling Nonlinear Economic Time Series*. Oxford: Oxford University Press.
- Tkachenko, D. and Z. Qu (2012) Frequency domain analysis of medium scale DSGE models with application to Smets and Wouters, 2007. In: *DSGE Models in Macroeconomics: Estimation, Evaluation, and New Developments*, pp. 319–385, Emerald Group Publishing Limited.
- Villani, M. (2009) Steady-state priors for vector autoregressions. *Journal of Applied Econometrics* 24(4), 630–650.
- Watson, M. W. (1986) Univariate detrending methods with stochastic trends. *Journal of Monetary Economics* 18(1), 49–75.

Appendix A. Details of Alternative Models

This section describes the alternative models presented in Section 4 in more detail. First, CVAR and RS-VAR-homo are nested versions of the baseline model where the regime dependence of (some) parameters is shut-off. The CVAR is a standard VAR model with constant parameters and no regime-switching. That is, $\beta_{S_t} = \beta$ and $\Sigma_{S_t} = \Sigma$ for all t . In this model, the prior means of the VAR coefficients are set as in the Minnesota prior, i.e., the means of the first lag of each variable are set to one, while the means of all other coefficients are set to zero, with other specifications remaining the same as in the baseline model (RS-VAR-FD).

In the model with homoskedasticity (RS-VAR-homo), we assume that the error covariances are identical across regimes, i.e., $\Sigma_{S_t} = \Sigma$ for all t . All other features, including the prior means of the VAR coefficients and the estimation procedure, remain the same as in the RS-VAR-FD.

For the model with stochastic volatility (RS-VAR-SV), we follow the conventional framework outlined in the literature (e.g., Primiceri, 2005). In this model, the (co)variances of the innovations are assumed to (gradually) evolve over time rather than switching between regimes. Specifically, it relies on the standard decomposition of the (time-varying) error covariance matrix:

$$\Sigma_t = A_t^{-1} H_t A_t^{-1'}, \quad (\text{A.1})$$

where A_t is a lower triangular matrix with (i, j) element $\alpha_{ij,t}$ and H_t is a diagonal matrix with i th element $h_{i,t}$, and we assume the evolution of $h_{i,t}$ follows geometric random walks:

$$\ln h_{i,t} = \ln h_{i,t-1} + w_{i,t}, \quad i = 1, \dots, n, \tag{A.2}$$

and the elements of the matrix A_t , collected in $\alpha_t = (\alpha_{21,t}, \dots, \alpha_{n(n-1),t})$, evolve as driftless random walks:

$$\alpha_t = \alpha_{t-1} + \zeta_t, \tag{A.3}$$

where $w_t = (w_{1,t}, \dots, w_{n,t})' \sim N(0, V_w)$ and $\zeta_t \sim N(0, V_\zeta)$. For the priors on the new hyperparameters, V_w and V_ζ , we use the same values as in Primiceri (2005). All other features, including the prior means of the VAR coefficients, remain the same as in the RS-VAR-FD. While in this specification we assume that all elements in Σ_t are time-varying, we also explore alternative scenarios where the elements in A_t are either constant or regime-dependent. We find that the overall results are largely unchanged under these alternative specifications.

In specifying the model with asymmetric prior (RS-VAR-AP), we largely follow Carriero et al., (2019). First, using the factorization of a regime-dependent (reduced-form) error covariance matrix, $\Sigma_{S_t} = \tilde{A}_{S_t} \tilde{A}'_{S_t}$, where \tilde{A}_{S_t} is a triangular matrix, we can recover the structural innovations, $\varepsilon_t = \tilde{A}_{S_t}^{-1} u_t$, with $\varepsilon_t \sim N(0, I_n)$. The equations in the (reduced-form) VAR model then can be rewritten in terms of structural shocks:

$$y_t = \beta'_{S_t} X_t + \tilde{A}_{S_t} \varepsilon_t, \tag{A.4}$$

with the i th equation being

$$y_{it} - \sum_{j=1}^{i-1} \tilde{\alpha}_{ij,S_t} \varepsilon_{jt} = \beta'_{i,S_t} X_t + \tilde{\alpha}_{ii,S_t} \varepsilon_{it}, \tag{A.5}$$

where $\tilde{\alpha}_{ij,S_t}$ is the (i, j) element of \tilde{A}_{S_t} , and β_{i,S_t} is the part of the VAR coefficients associated with the equation i . This model then can be easily estimated by sampling the posterior of the VARs coefficients equation by equation.

In estimating this model, we essentially follow the algorithm in Carriero et al., (2019) and Carriero et al., (2022). Unlike the baseline model, where the hyperparameters are only regime-specific, this model features distinct sets of hyperparameters for each equation and each regime. Therefore, the step of drawing hyperparameters is divided into two blocks, where the first block draws λ_1, λ_2 , and λ_3 (i.e., those not related to co-persistence) for individual equations, and given these draws, the second block draws λ_4 (related to co-persistence) for the entire VAR model.

For the RS-VAR-TD, we impose the same unit root-type Minnesota priors on the VAR coefficients in both regimes, while the other prior specifications and estimation procedure remain the same as in the RS-VAR-FD. Letting $\beta_{1,S_t}^{(0)}$ denote the intercept term in the GDP equation, the regime identification condition can be stated as follows: $\beta_{1,S_t=2}^{(0)} < \beta_{1,S_t=1}^{(0)}$ and $|\Sigma_{S_t=2}| > |\Sigma_{S_t=1}|$.

In both the RS-VAR-FD-TVTP and RS-VAR-TD-TVTP models, we follow Filardo and Gordon (1998) and specify the transition probabilities using a probit model, $\Pr(S_t = 1) = \Pr(S_t^* < 0)$, where S_t^* is a continuous unobserved variable with

$$S_t^* = \gamma_0 + \sum_{l=1}^p \gamma_l z_{t-l} + \phi S_{t-1} + v_t, \quad v_t \sim N(0, 1) \tag{A.6}$$

and z_t represents the GDP growth rate. The lag order is set to 2 based on the log marginal likelihoods. The sampling algorithm for this model includes additional steps to draw S_t^* and the parameters $(\gamma_0, \gamma_1, \dots, \gamma_p, \phi)$. We sample the latent variable from truncated normal distributions (Albert and Chib, 1993) and draw the coefficients in the probability equation using the standard approach, as in a simple linear regression with a known error variance.

Appendix B. Additional Tables and Figures

Table A1. Log Marginal Likelihoods

A. The frequency-dependent switching VARs (with varying R and p)			
$(R = 2, p = 2)$	$(R = 2, p = 3)$	$(R = 2, p = 4)$	$(R = 2, p = 5)$
-1291.11	-1297.39	-1285.24	-1317.07
$(R = 3, p = 2)$	$(R = 3, p = 3)$	$(R = 3, p = 4)$	$(R = 3, p = 5)$
-1354.87	-1334.19	-1369.60	-1406.67
B. Other models			
RS-VAR-homo	RS-VAR-SV	RS-VAR-AP	CVAR
-1398.50	-1383.29	-1378.36	-1375.24
RS-VAR-TD	RS-VAR-FD-I(1)	RS-VAR-FD-TVTP	RS-VAR-TD-TVTP
-1298.42	-1289.22	-1264.69	-1289.39

Note: The log marginal likelihood for each model is computed using the algorithm proposed by Chan and Eisenstat (2018), which is based on the integrated likelihood. A large part of the computation relies on the MATLAB codes that accompany their paper.

Table A2. Estimates for Model Parameters: Model with GDP Deflator Inflation

	A. Error Covariances		
	low frequency	BC frequency	CVAR
$Var(u_y)$	0.396 (0.304, 0.474)	0.993 (0.816, 1.248)	0.601 (0.548, 0.663)
$Var(u_p)$	0.050 (0.043, 0.060)	0.106 (0.088, 0.132)	0.065 (0.059, 0.071)
$Var(u_i)$	0.118 (0.078, 0.142)	0.842 (0.672, 1.101)	0.392 (0.337, 0.447)
$Cov(u_y, u_p)$	-0.010 (-0.021, 0.003)	0.010 (-0.031, 0.057)	-0.001 (-0.013, 0.012)
$Cov(u_y, u_i)$	0.054 (0.029, 0.080)	0.173 (0.073, 0.293)	0.135 (0.092, 0.178)
$Cov(u_p, u_i)$	0.009 (0.004, 0.016)	0.012 (-0.024, 0.044)	0.013 (0.003, 0.024)
	B. Hyperparameters		C. Transition prob.
	low frequency	BC frequency	
λ_1	0.176 (0.117, 0.471)	0.151 (0.112, 0.342)	p_{11} 0.954 (0.926, 0.974)
λ_2	4.808 (1.540, 8.352)	1.969 (0.659, 3.782)	p_{12} 0.046 (0.026, 0.074)
λ_3	0.267 (0.133, 0.721)	0.180 (0.118, 0.468)	p_{21} 0.118 (0.064, 0.203)
λ_4	0.270 (0.096, 0.464)	0.651 (0.227, 2.472)	p_{22} 0.882 (0.797, 0.936)

Note: See notes to Table 1.




Figure A1. Estimated regimes and spectral densities (with GDP deflator inflation).

Note: See notes to Figure 2.




Figure A2. Impulse responses from the model with GDP deflator inflation (Full result).

Note: See notes to Figure 3.




Figure A3. Estimates of stochastic volatilities.

Notes: The figure plots the median estimates of the time-varying stochastic volatilities, $\sqrt{h_{i,t}}$, in each series in RS-VAR-SV, along with 16%–84% credible intervals.




Figure A4. Estimated regimes for individual series.
 Note: See notes to Figure 2.




Figure A5. Impulse responses of the conventional regime-switching model.

Note: In each panel, blue solid lines represent the 16%–84% credible interval for the impulse responses for the low-frequency and the BC-frequency regime, respectively (in the RS-VAR-FD); red dashed lines represent those for the expansion and recession regime, respectively (in the RS-VAR-TD).




Figure A6. Estimated regimes and spectral densities for the RS-VAR-FD-I(1).

Note: See the notes to Figure 2.




Figure A7. Models with time-varying transition probabilities.

Note: Each panel in the figure shows the estimated probabilities of the second regime (BC-frequency and recession) with the time-varying transition probabilities. On the right of each panel are 5%–95% intervals for the spectral densities for each regime.