Contagious Switching*

Michael T. Owyang[†], Jeremy Piger [‡], Daniel Soques[§]

PRELIMINARY AND INCOMPLETE – DO NOT CITE

December 10, 2018

Abstract

In this paper, we analyze recession propagation across countries. We extend the QualVAR framework of Dueker (2005) to allow for multiple qualitative state variables in a VAR. We consider two different versions of the model. One version assumes the discrete state of the economy (expansion or recession) is observed. In the other, the state of the economy is unobserved and inferred from movements in output growth. We apply the model to Canada, Mexico, and the U.S. to test if spillover effects were similar before and after NAFTA. We also estimate cross-country spillovers of U.S. monetary policy shocks.

JEL Codes: C32; E32

Keywords: NAFTA, business cycles

*Hannah G. Shell and Kate Vermann provided research assistance. Views expressed here are the authors' alone and do not reflect the opinions of the Federal Reserve Bank of St. Louis or the Federal Reserve System.

†Corresponding author. Research Division, Federal Reserve Bank of St. Louis. michael.t.owyang@stls.frb.org

‡Department of Economics, University of Oregon

§Department of Economics and Finance, University of North Carolina Wilmington

1 Introduction

[To be added.]

2 Empirical Setup

Consider the interaction between the business cycles of $n = 1, ..., N$ countries across $t = 1, ..., T$ periods. Let $S_{nt} = \{0, 1\}$ represent the discrete business cycle phase for country n at time t, where $S_{nt} = 0$ represents an expansionary phase and $S_{nt} = 1$ represents a recessionary phase. Collect the business cycle phases into a vector $S_t = [S_{1t}, ..., S_{Nt}]'$.

2.1 Observed Regimes

To model the interdependence of business cycle phases across countries, we must specify how S_{nt} affects S_{mt}, $m \neq n$. Assume, initially, that each S_{nt} is observed. Further, suppose that the business cycle phases propagate across countries with a lag. Let z_{nt} represent a continuous latent variable related to the binary observed variable S_{nt} through the deterministic relationship:

$$S_{nt} = \begin{cases} 1 \quad \text{if} \quad z_{nt} \geq 0 \\ 0 \quad \text{otherwise} \end{cases} .$$

Through the latent variable z_{nt}, we can study how other variables—both macro variables and the business cycle phases of other countries—affect the future business cycle phase of country n. Let $y_t = [y_{1t}, ..., y_{Jt}]'$ represent a $(J \times 1)$ vector of macro variables that could include country-specific policy variables or other economic indicators and let $z_t = [z_{1t}, ..., z_{Nt}]'$ collect the continuous latent business cycle indicators.

Define $Y_t = [z_t', y_t']'$, where the relationship between the contemporaneous Y_t and its lags follows a vector autoregression (VAR):

$$Y_t = B_0 + B(L) Y_{t-1} + u_t, \tag{1}$$

where $u_t = [u^z_{1t}, ..., u^z_{Nt}, u^y_{1t}, ..., u^y_{Jt}]'$ and $E_t [u_t u_t'] = \Sigma$. For exposition, it will be convenient to write (1) in a more detailed form:

$$\begin{bmatrix} z_t \\ y_t \end{bmatrix} = \begin{bmatrix} B^z_0 \\ B^y_0 \end{bmatrix} + \begin{bmatrix} B^{zz}(L) & B^{zy}(L) \\ B^{yz}(L) & B^{yy}(L) \end{bmatrix} \begin{bmatrix} z_{t-1} \\ y_{t-1} \end{bmatrix} + \begin{bmatrix} u^z_t \\ u^y_t \end{bmatrix},$$

where $B^{ij}(L)$ represents the lagged effect of j on i. Because z_t is latent, it is convenient to make scale assumptions by restricting their variances. In particular, we assume that the

$$\Sigma_{zz} = E_t [u^z_t u^{z\prime}_t]$$

has unit diagonal elements. In subsequent sections, we can impose additional restrictions on the decomposition of the VAR variance-covariance matrix that identifies the structural form of the VAR from its reduced form.

The current model has a form similar to a multi-binary-variable extension of Dueker's (2005) Qual-VAR. In that paper, a single binary variable indicates the state of the economy and can be affected—and, importantly, can affect—a vector of macroeconomic variables at lags. This version of our model with observed S_t collapses to the Qual-VAR when S_t is a scalar.[1] Because of the assumption that the reduced-form VAR errors are multivariate normal, the z_t equations in the VAR resemble a multivariate extension of the dynamic probit outlined in Eichengreen, Grossman, and Watson (1986). Chib and Greenberg (1998) develop methods of estimating the static multivariate probit, which is equivalent to the z_t equations in the VAR in our model with observed S_t with the additional assumption that $B^{zz}(L) = 0$. Our model is perhaps most similar to the multivariate dynamic probit of Candelon, Dumitrescu, Hurlin, and Palm (2013) in which we add a propagation mechanism for the covariates that allows the latents to affect macro variables at lags.

[1]The Prob-VAR outlined in Fornari and Lemke (2010) is a more restricted version of the Qual-VAR, where they assume $B^{zz}(L)$ and $B^{yz}(L)$ are both identically zero. Their model, then, is essentially a VAR in the macro variables y and a probit where lags of y determine a scalar z. Their application is to forecasting S using iterative multistep methods. The VAR forms forecasts for y, which in turn informs the forecast of S at longer horizons.

Two key features differentiate our model from a set of independent time-varying transition probability switching models. First, there is a lagged cross-regime effect that is embedded in the off-diagonal elements of $B^{zz}(L)$. The lagged effect represents the contagious switching, where a regime change in one country can spill over into the regimes of its neighbors. Further notice that the regime cross-series dependence is a function of the continuous latent rather than the binary latent. This means that z_{nt} may be thought of as representing the strength of the cycle Second, there is a contemporaneous cross-regime effect that is embedded in the tetrachoric correlation term in Σ. The tetrachoric effect can represent either simultaneity of shocks that cross country borders or within-quarter contagion effects. The model allows us to test for the presence of cyclical contagion, the speed at which it acts, and the degree to which countries affect each other. In addition, countercyclical or prophylactic policy can be included in the y_j's to determine whether, say, changes in fiscal or monetary policy can reduce the probability of recessions.

2.2 Unobserved Regimes

While we previously assumed that the S_t are observed, we can relax this assumption by including a vector of economic indicators whose means depends on the discrete regimes. Unobserved regimes can be relevant for a number of reasons. For example, one simply might not have the data available as all countries to not construct or announce business cycle turning points. On the other hand, some countries have more than one set of turning point dates, suggesting some uncertainty over the timing of the events. In the U.S., the NBER Business Cycle Dating Committee dates are widely accepted as the "official" business cycle turning points. However, these dates are not revised even in the presence of new or revised data. Moreover, other measures such as the OECD Recession Indicators may vary slightly from the NBER in the timing of the turning points. In some of these cases, it may be advantageous to estimate the regime changes directly from the data.

Suppose, then, that each of the N countries can be characterized by a single period$-t$ business cycle indicator, x_{nt}.[2] For the purposes of exposition, we can call x_{nt} the output

[2] A straightforward extension would allow the regimes to be identified from multiple series per country. Using multiple series can sharpen identification of the turning points, since the switch is assumed to affect

growth rate but it could be, in principle, any contemporaneous indicator of the cycle. Collect the period$-t$ output growth rates into a vector $x_t = [x_{1t}, ..., x_{Nt}]'$. We assume that output growth is a stochastic sampling from a mixture of normals, where μ_{n0} and μ_{n1} are the means of the two normal distributions and we impose $\mu_{n0} > 0 > \mu_{n1}$ for identification. Note that the mixtures can be potentially different for each country, as evidenced by the index n. The interpretation of our assumption is that each country's economy moves between two business cycle phases, a positive mean (expansion) and a negative mean (recession).

During each period, a country n's business cycle phase is represented by the latent variable S_{nt} that determines which of the two distributions y_{nt} is drawn from that period. The process can be summarized by

$$x_{nt} = \mu_{n0} + \Delta\mu_n S_{nt} + \phi_n (L) x_{n,t-1} + \varepsilon_{nt}, \tag{2}$$

where we can define $\Delta\mu_n = \mu_{n1} - \mu_{n0}$, $\Delta\mu_n < -\mu_{n0}$, as implied by our identifying assumption, and $\varepsilon_{nt} \sim N(0, \sigma_n^2)$. We impose that the output volatility is time invariant and that the output shocks are uncorrelated across countries, serially uncorrelated, and uncorrelated with the shocks to the variables in the VAR.[3]

3 Estimation and Data

3.1 The Sampler

We estimate the model using the Gibbs sampler, a Markov-Chain Monte Carlo algorithm that draws a block of the model parameters—including the underlying continuous states–conditional on the remaining parameters and the data. Let Ω_t represent the information (i.e., data) available at time t. We specify a standard set of priors for the model with observed regimes. The parameters in B are multivariate normal. We assume a standard inverse Wishart prior for Σ. In order to identify the state variables z_{nt}, we assume that

all of the series within the country.

[3]These assumptions are made for expositional clarity are straightforward to relax.

$var[u^z_{nt}] = 1$ for all $n = 1, ..., N$.[4] In the case with unobserved regimes, we also need to set priors for the intercepts, the AR terms, and the innovation variances in the x_t equation. We assume that the parameters in the x_t equation have a Normal-inverse Gamma prior.

We divide the exposition of the sampler into two parts. In the first part, we outline the sampler for the case where S_t is observed. In this case, there are three blocks for estimation: (1) the coefficient matrices for the VAR, $B = \{B_0, ..., B_P\}$; (2) the VAR variance-covariance matrix, Σ; and (3) the latent states, $\{z_t\}_{t=1}^T$. The first two blocks are conjugate normal and inverse Wishart, respectively. The third block is executed by drawing the continuous latent state variable recursively from smoothed Kalman posterior distributions.[5]

Aside from sampling the additional parameters in the x_t equation, the case of unobserved regimes adds a wrinkle that warrants more explanation. Because the sign of z_{nt} is determined by the value of S_{nt} and the past z_t determine the transition probabilities for S_{nt}, these two values must be sample simultaneously. Thus, the sampler for the unobserved state case has five blocks: (1) the coefficient matrices for the VAR, $B = \{B_0, ..., B_P\}$; (2) the VAR variance-covariance matrix, Σ; (3) the coefficients for the measurement equation, $\Psi = \{\mu_0', \mu_1', \phi'\}$; (4) the measurement innovation variances, $\{\sigma_n^2\}_{n=1}^N$; and (5) the latent states, $\{z_t\}_{t=1}^T$. Blocks (3) and (4) are conjugate. We outline the filter used to obtain draws of block (5) below; other draws are detailed in the Technical Appendix.

3.1.1 Drawing $\{z_t\}_{t=1}^T, \{S_t\}_{t=1}^T$ conditional on B, Σ, $\{\sigma_n^2\}_{n=1}^N$, Ψ

Unfortunately, we cannot draw the sequences of the two latent variables in separate blocks. The value of S_{nt} is directly related to the sign of z_{nt}. One might posit a draw in which the full sequence $\{S_\tau\}_{\tau=1}^T$ is drawn, conditional on the past iteration of $\{z_\tau\}_{\tau=1}^{t-1}$; then, a draw of the full sequence of $\{z_\tau\}_{\tau=1}^{t-1}$, conditional on the new draw of $\{S_\tau\}_{\tau=1}^T$,

[4] Alternatively, we could draw the restricted matrix as outlined by Chan and Jeliazkov (2009).

[5] This differs from Dueker's original sampler. In this sampler, Dueker used exact conditional distributions for the interior $T - 2P$ periods. The first P periods were drawn using Metropolis-Hastings. The last P periods are drawn by iterating forward the mean of the exact conditional distribution for the $T - P - 1$ period.

where each S_{nt} determines the direction of the truncation of z_{nt}. However, any draw that changes S_{nt} across Gibbs iteration invalidates the last draw of z_{nt}, as the truncation would be improper. Drawing the full sequence $\{z_\tau\}_{\tau=1}^{t-1}$ first also would be invalid. While we can obtain a Kalman posterior for z_{nt}, the exact conditional distribution will be truncated. Simply drawing z_{nt} from the Kalman posterior and then assigning S_{nt} based on the sign of z_{nt} would ignore information in the x's that inform S_{nt}.

We adopt an alternative approach that takes advantage of both the Kalman filter and Hamilton's Markov switching filter to draw candidates for a Metropolis-in-Gibbs step. Because we need to use the draws of lagged z_t to form the transition probabilities for the Hamilton filter, we cannot draw the candidates using smoothed probabilities. Instead, for each t, we draw a candidate S_t^*, conditional on lags of z_t, using the forward component of the Hamilton filter. We then draw a candidate z_t^* from the posterior obtained by the forward component of the Kalman filter.

Specifically, start with set of initialization probabilities, $\Pr[S_{n0}]$, which could be the steady state regime probability, and initialize the vector of latents, z_0. The goal is to obtain (jointly) a candidate pair of vectors (z_t^*, S_t^*) for each t. We can form the joint proposal density as

$$p(z_t^*, S_t^*|\Omega_t) = p\left(z_t^*|\Omega_t, S_{nt}^*, \{z_\tau\}_{\tau=1}^{t-1}\right) \prod_{n=1}^{N} p\left(S_{nt}^*|\Omega_t, \{z_\tau\}_{\tau=1}^{t-1}\right).$$

We draw the candidate S_{nt}^* from

$$\Pr[S_{nt}^* = 1|\Omega_t] = \frac{\sum_{S_{n,t-1}} \ell\left(S_{nt}^* = 1, S_{n,t-1}|\Omega_t, \{z_\tau\}_{\tau=1}^{t-1}\right) \Pr\left[S_{nt}^* = 1|S_{n,t-1}, \{z_\tau\}_{\tau=1}^{t-1}\right] \Pr\left[S_{n,t-1}|\Omega_{t-1}, \{z_\tau\}\right.}{\sum_{S_{nt}} \sum_{S_{n,t-1}} \ell\left(S_{nt}^*, S_{n,t-1}|\Omega_t, \{z_\tau\}_{\tau=1}^{t-1}\right) \Pr\left[S_{n,t}^*|S_{n,t-1}, \{z_\tau\}_{\tau=1}^{t-1}\right] \Pr\left[S_{n,t-1}|\Omega_{t-1}, \{z_\tau\}\right.}$$

where $\ell(.,.|.)$ is the likelihood and

$$\Pr\left[S_{nt}^* = 1|S_{n,t-1}, \{z_\tau\}_{\tau=1}^{t-1}\right] = \Pr[z_{nt} > 0|z_{t-1}]$$

are the transition probabilities, which depend on the lagged continuous latent variable for all n.

The conditional distributions for the z_{nt}'s can be obtained by the forward component of the Kalman filter. Based on the state equation, (1), the Kalman filter obtains the forecast density for the vector z_t, conditional on its lags. Then, the filter updates the forecast density using information from the current realization of y_t to obtain

$$p\left(z_t|\Omega_t\right) \sim N\left(\widehat{z}_{t|t}, P^z_{t|t}\right),$$

where $\widehat{z}_{t|t}$ is the mean of the conditional distribution and $P^z_{t|t}$ is the covariance matrix.[6] Then, conditional on S^*_t, we can draw the candidate from the truncated normal, where the truncation direction depends on S^*_t:

$$p\left(\varsigma_t|\Omega_t, S_t\right) \sim TN\left(\varsigma_{t|t}, P_{t|t}, S_t\right).$$

Finally, we validate the candidate (z^*_t, S^*_t)—drawn jointly for all n—using standard MH acceptance probabilities. The candidate (z^*_t, S^*_t) is accepted with probability α, where

$$\alpha = \min\left[1, \frac{\pi\left(S^*, z^*\right) f(x, y|S^*, z^*) q\left(S^{[i-1]}|z^{[i-1]}\right)}{\pi\left(S^{[i-1]}, z^{[i-1]}\right) f(x, y|S^{[i-1]}, z^{[i-1]}) q\left(S^*, z^*\right)}\right],$$

where $\pi\left(.,.\right)$ is the prior, $f(x, y|.,.)$ is the likelihood, and $q\left(,|,\right)$ are the move probabilities. Because we have an independence chain, the ratio of the move probabilities collapses to 1. Using the fact that y does not depend on S and the identity $P(S|z) = 1$, the posterior likelihood is

$$\pi\left(S, z\right) f(x, y|S, z) = f(y, z|S) \prod_{n=1}^{N} f(x_n|S_n),$$

where

$$f(y, z|S) = \frac{1}{\left|2\pi\Sigma\right|^{1/2}} \prod_{t=1}^{T} \exp\left\{-\frac{1}{2}\left(Y_t - B\left(L\right)Y_{t-1}\right)' \Sigma^{-1}\left(Y_t - B\left(L\right)Y_{t-1}\right)\right\}$$

[6] Each of these quantities will be a subvector and submatrix, respectively, of the output of the Kalman filter.

and

$$f(x_n|S_n) = \frac{1}{\sqrt{2\pi}\sigma_n^2}\prod_{t=1}^{T} \exp\left\{-\frac{1}{2\sigma_n^2}\left(x_{nt} - \mu_{n0} - \Delta\mu_n S_{nt} - \phi_n\left(L\right) x_{n,t-1}\right)'\left(x_{nt} - \mu_{n0} - \Delta\mu_n S_{nt} - \phi_n\left(L\right) x\right.\right.$$

3.2 Data

We apply the model to the NAFTA member countries (Canada, Mexico, and the U.S.). We perform two experiments. First, we consider the model with observed recessions. This model requires two sets of data: (1) the recession indicator S_{nt}, and (2) the macroeconomic variables in y_t. For the recession indicators S_{nt}, we use the Reference Turning Points from the OECD Composite of Leading Indicators. In our application, we use the U.S. effective federal funds rate or the Wu-Xia shadow short rate. The effective federal funds rate comes from FRED and the shadow short rate is available from the Atlanta Fed.

Next, we consider the model with unobserved regimes. This model also requires two sets of data: (1) x_t, the variable that informs S_t, and (2) the macroeconomic variables in y_t. For x_t, we include real GDP growth for each of the three countries. The real GDP growth data comes from the OECD Quarterly National Accounts. We use the same macroeconomic variables for the model with unobserved regimes.

The data for all three countries is available for the period 1961:Q2-2018:Q1.

4 Results

4.1 S_t is Observed

We first consider the model with observed recessions. Figure 1 shows the posterior median for the latent continuous recession variables z_{nt}.

Figure 2 shows the IRFs to a shock to the Canadian recession variable, z_{CA}. The second and third panels show the response of the recession variables in Mexico and the U.S., respectively. The negative shock to the Canadian economy significantly increases the recession variables in the other two countries as the recession quickly spreads. The

U.S. recession variable increases by more than double the increase in Mexico's, implying the U.S. is relatively more sensitive to Canadian economic fluctuations. The second row shows the response of the U.S. policy rate and inflation. U.S. monetary policy responds by cutting the federal funds rate and inflation pressures soften as the U.S. economy deteriorates.

Figure 3 displays IRFs for a negative economic shock in Mexico. The Canadian and U.S. economies weaken due to the propagation of recessionary pressures. Note that Canada is relatively less affected than the U.S. Similar to the negative shock to Canada, both the federal funds rate and U.S. inflation fall in response to a negative shock to the Mexican economy.

Figure 4 presents the IRFS of a negative U.S. economic shock. The economies of Canada and Mexico move with the U.S. with a higher recession variable. However, we find a significant increase in the federal funds rate in response to this deterioration. One potential explanation for this monetary tightening is the resulting higher inflation seen in the last panel.

Lastly, figure 5 shows the IRFs to an increase in the federal funds rate. The U.S. recession variable significant increases as the policy rate rises. However, we do not find that the policy shock has significant effects on either the Canadian or Mexican economy. The last panel also shows inflation increasing in response to a higher federal funds rate. This result is commonly referred to as the "price puzzle."[7]

4.2 S_t is Unobserved

[To be added.]

5 Conclusions

[To be added.]

[7] See Rusnak et al. (2013) and Estrella (2015).

References

[1] Bertrand Candelon, Elena Ivona Dumitrescu, Christophe Hurlin, and Franz C. Palm. “Multivariate Dynamic Probit Models: An Application to Financial Crises Mutation,” SSRN.

[2] Candelon, Bertrand, Elena-Ivona Dumitrescu, Christophe Hurlin, and Franz C. Palm. ”Multivariate dynamic probit models: an application to financial crises mutation.” In *VAR Models in Macroeconomics–New Developments and Applications: Essays in Honor of Christopher A. Sims*, pp. 395-427. Emerald Group Publishing Limited, 2013.

[3] Joshua Chi-Chun Chan & Ivan Jeliazkov “MCMC Estimation of Restricted Covariance Matrices,” *Journal of Computational and Graphical Statistics,* 18(2), 2009, pp. 457-480.

[4] Siddhartha Chib and Edward Greenberg. “Analysis of Multivariate Probit Models,” *Biometrika* 85(2) pp. 347-361.

[5] Louis N. Christofides, Thanasis Stengos, and Robert Swidinsky “On the Calculation of Marginal Effects in the Bivariate Probit Model,” *Economics Letters*, 54(3), July 1997, pp. 203–208.

[6] Dueker Dynamic Forecasts of Qualitative Variables: A Qual VAR Model of U.S. Recessions Journal of Business & Economic Statistics Vol. 23, No. 1 (Jan., 2005), pp. 96-104

[7] Eichengreen, Barry, Mark W. Watson, and Richard S. Grossman. ”Bank rate policy under the interwar gold standard: A dynamic probit model.” *The Economic Journal* 95.379 (1985): 725-745.

[8] Estrella, Arturo. ”The price puzzle and VAR identification.” *Macroeconomic Dynamics* 19, no. 8 (2015): 1880-1887.

[9] Fornari, Fabio and Lemke, Wolfgang, (2010), Predicting recession probabilities with financial variables over multiple horizons, No 1255, Working Paper Series, European Central Bank.

[10] William H. Greene “Marginal Effects in the Bivariate Probit Model,” June 1996. NYU Working Paper No. EC-96-11.

[11] Rusnák, Marek, Tomas Havranek, and Roman Horvaáth. ”How to solve the price puzzle? A meta-analysis.” *Journal of Money, Credit and Banking* 45, no. 1 (2013): 37-70.

Figure 1: Posterior Medians for the Continuous Recession Variables

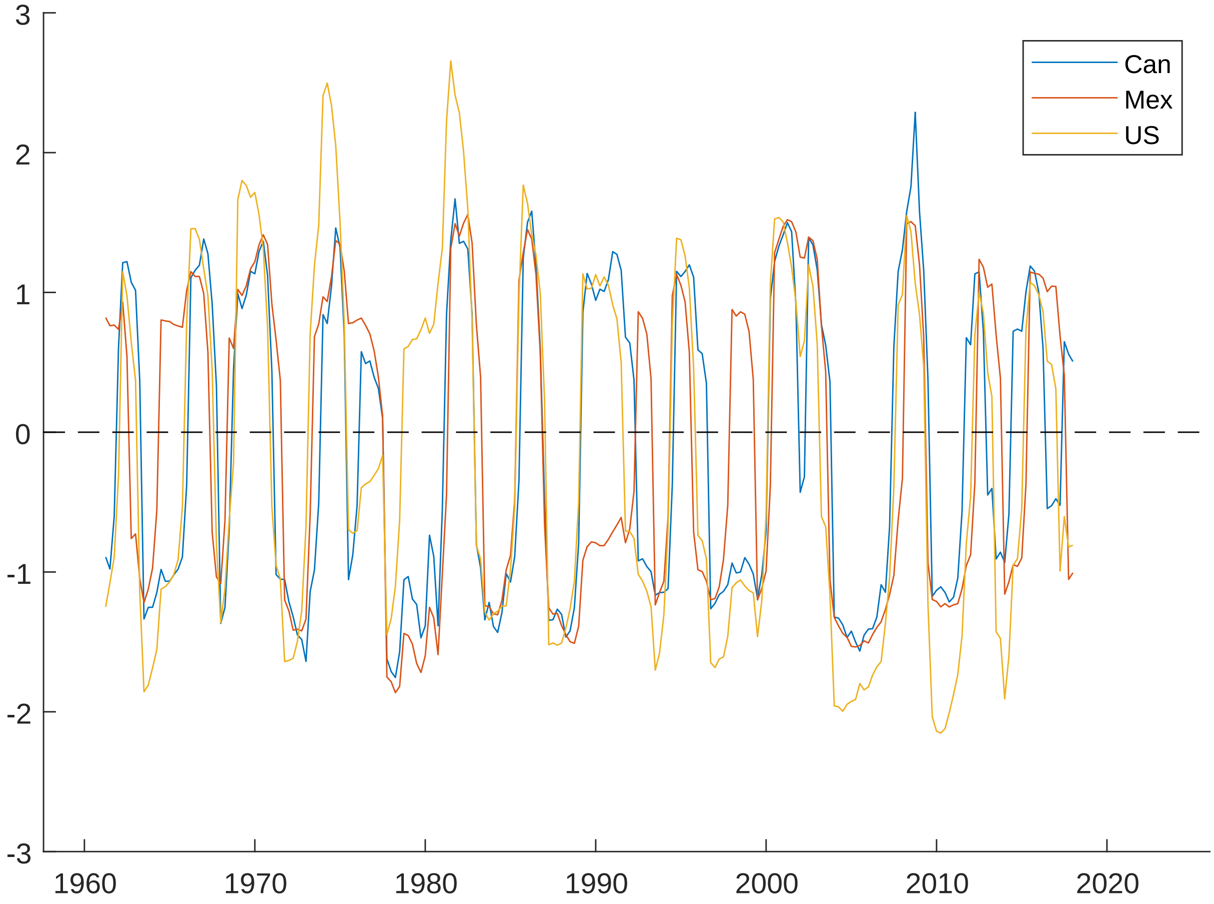

This figure shows the posterior medians for the latent continuous recession variable z_{nt}.

Figure 2: Impulse Response to an Increase in Canada's Recession Variable

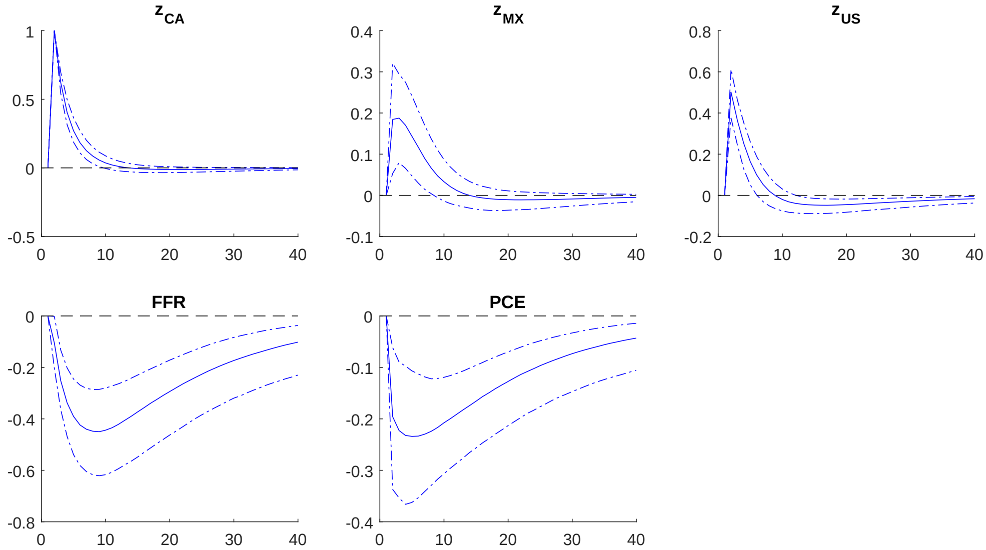

This figure shows the impulse response functions to a one standard deviation increase in the continuous recession variable for Canada, z_{CANt}. The solid lines are posterior medians and the dotted lines are the 67% HPD.

Figure 3: Impulse Response to an Increase in Mexico's Recession Variable

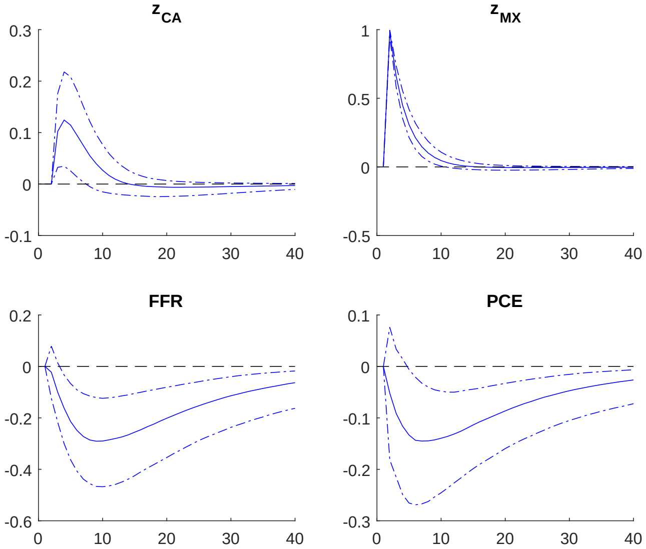

This figure shows the impulse response functions to a one standard deviation increase in the continuous recession variable for Mexico, z_{MEXt}. The solid lines are posterior medians and the dotted lines are the 67% HPD.

Figure 4: Impulse Response to an Increase in the U.S. Recession Variable

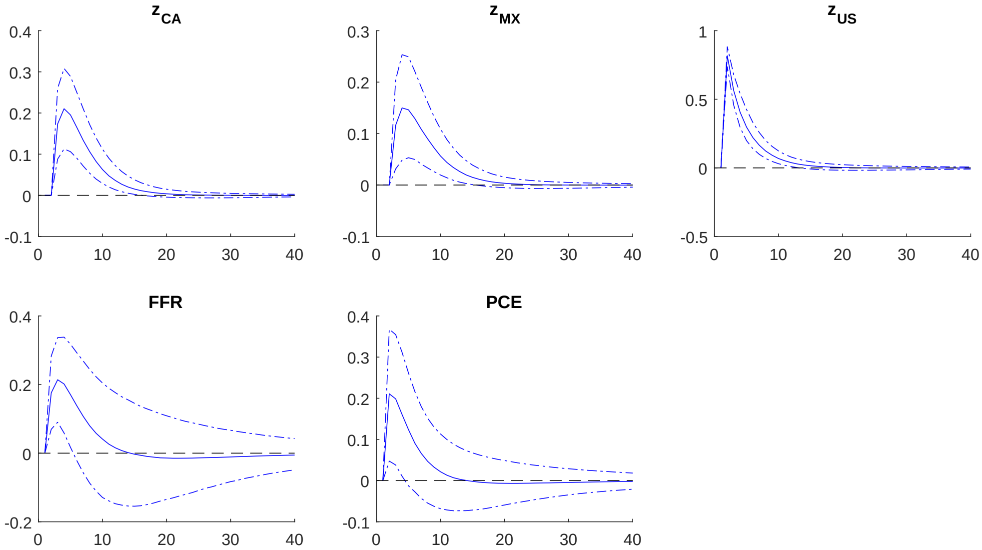

This figure shows the impulse response functions to a one standard deviation increase in the U.S. continuous recession variable, z_{USt}. The solid lines are posterior medians and the dotted lines are the 67% HPD.

Figure 5: Impulse Response to an Increase in Federal Funds Rate

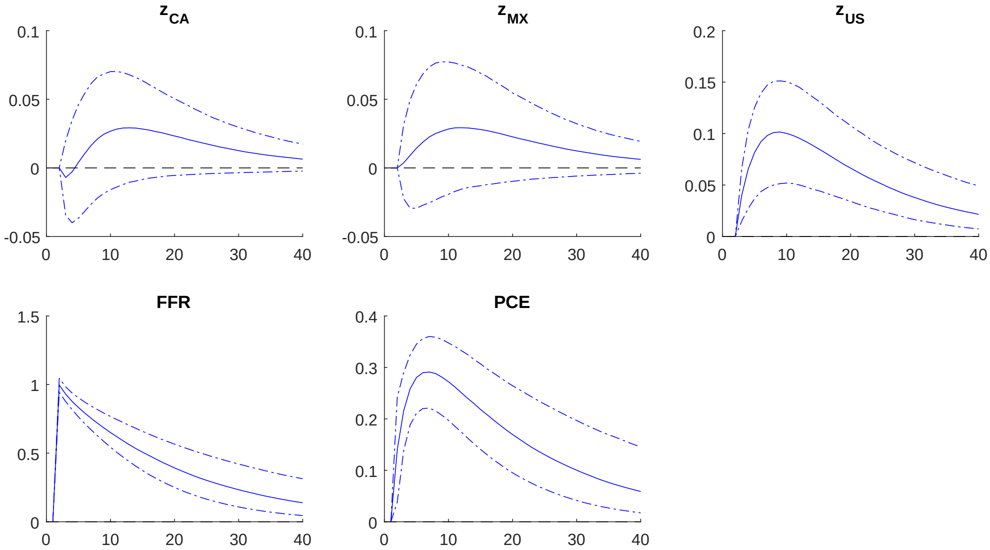

This figure shows the impulse response functions to a 100 basis point increase in the U.S. federal funds rate. The solid lines are posterior medians and the dotted lines are the 67% HPD.

A Technical Appendix

The following appendix describes in detail the draws for the parameters. We first outline the state space representation of the VAR. We then describe the two draws that are invariant to whether we observe the regime. These draws condition only on the continuous latent state, z_t. We then describe the draw for the continuous latent variable z_t when S_t is observed. Finally, we describe the draws for the parameters in the measurement equation that relates the discrete regime to the growth variable, x_t.

A.1 The State Space Representation

Define $Y_t = [z_t', y_t']'$ and $\varsigma_t = \left[Y_t', Y_{t-1}', \cdots, Y_{t-P+1}'\right]'$ as the state in the state-space representation of the model with measurement equation:

$$X_t = H\varsigma_t,$$

and transition equation:

$$\varsigma_t = M + F\varsigma_{t-1} + e_t,$$

where

$$e_t \sim N(\mathbf{0}_{(N+J)P\times 1}, Q).$$

The parameters of the state space are defined as follows:

$$Q = \begin{bmatrix} \Sigma & \mathbf{0}_{(N+J)\times(N+J)(P-1)} \\ \mathbf{0}_{(N+J)(P-1)\times(N+J)(P)} & \end{bmatrix},$$

$$H = \begin{bmatrix} \mathbf{0}_{N\times N} & \boldsymbol{I}_J & \mathbf{0}_{(N+J)\times(N+J)} & \cdots & \mathbf{0}_{(N+J)\times(N+J)} \end{bmatrix},$$

$$M = \begin{bmatrix} B_0 \\ \mathbf{0}_{(N+J)(P-1)\times 1} \end{bmatrix},$$

and

$$F = \begin{bmatrix} B_1 & B_2 & \cdots & B_{P-1} & B_P \\ \boldsymbol{I}_{(N+J)} & \mathbf{0}_{(N+J)\times(N+J)} & \cdots & \mathbf{0}_{(N+J)\times(N+J)} & \mathbf{0}_{(N+J)\times(N+J)} \\ \mathbf{0}_{(N+J)\times(N+J)} & \boldsymbol{I}_{(N+J)} & \cdots & \mathbf{0}_{(N+J)\times(N+J)} & \mathbf{0}_{(N+J)\times(N+J)} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0}_{(N+J)\times(N+J)} & \mathbf{0}_{(N+J)\times(N+J)} & & \boldsymbol{I}_{(N+J)} & \mathbf{0}_{(N+J)\times(N+J)} \end{bmatrix}.$$

Notice that the measurement equation is deterministic.

A.2 Drawing B conditional on Σ, $\{z_\tau\}_{\tau=1}^{t-1}$

Conditional on Σ, the VAR parameters B are conjugate Normal. Define $\varsigma_t = \left[Y_t', Y_{t-1}', \cdots, Y_{t-P+1}'\right]'$. Then, the VAR can be written as:

$$Y_t = Bx_t + u_t,$$

where $x_t = \left[1, Y_{t-1}', ..., Y_{t-P}'\right]'$. Stacking the observations, we get:

$$Y = XB' + U$$

where $Y = [Y_{P+1}, Y_{P+2}, ..., Y_T]'$ and $X = [X_{P+1}, X_{P+2}, ..., X_T]'$. Let $\hat{B}' = (X'^{-1}X'Y$ be the OLS estimates for B'.

Define $b = vec(B')$ and $\hat{b} = vec(\hat{B}')$. Assuming a normal prior $b \sim N(\bar{b}_0, \bar{B}_0)$, the posterior distribution for b is:

$$b \sim N(b_1, B_1),$$

where

$$b_1 = B_1^{-1}(\bar{B}_0^{-1}\bar{b}_0 + \Sigma^{-1} \otimes X'X\hat{b}),$$

$$B_1 = (\bar{B}_0^{-1} + \Sigma^{-1} \otimes X'^{-1}.$$

We redraw B if the usual stationarity condition is violated.

A.3 Drawing Σ conditional on B, $\{z_t\}_{t=1}^T$

The draw of Ω is straightforward given the coefficient parameters B. Let $\hat{u}_t = Y_t - Bx_t$ where $x_t = \left[1, Y'_{t-1}, ..., Y'_{t-P}\right]'$. Stack the residuals in the matrix $U = [u'_{P+1}, ..., u'_T]'$. Given the prior distribution $\Omega^{-1} \sim W(\nu_0, \omega_0)$, the posterior distribution for Ω is:

$$\Sigma^{-1} \sim W(\nu, \omega),$$

where $\nu = \nu_0 + T/2$ and $\omega = (\omega_0 + U'U)/2$.

For $n = 1, ..., N$, we divide each row n and column n by the square root of the $[n, n]$ element of Ω. This ensures each latent variable z_{nt} is identified by the normalization $\mathrm{Var}(z_{nt}) = 1$ for all $n = 1, ..., N$.

A.4 Drawing z_t conditional on B, Σ, and observed $\{S_\tau\}_{\tau=1}^T$

We implement the Kalman filter with smoothing to draw the vector z_t given the state vector $S_t = [S_{1t}, ..., S_{Nt}]'$. If the sign of the draw for z_t does not match the state implied by S_t, we redraw until the condition is met. Since Q is singular, we use the modification outlined by Kim and Nelson (1999) that simplifies the backwards smoother to only the relevant conditioning factors.

A.5 Drawing $\Psi = \{\mu'_0, \mu'_1, \phi'\}$ conditional on $\{S_\tau\}_{\tau=1}^T$ and $\left\{\sigma_n^2\right\}_{n=1}^N$

[To be added.]

A.6 Drawing $\left\{\sigma_n^2\right\}_{n=1}^N$ conditional on $\{S_t\}_{\tau=1}^T$ and $\Psi = \{\mu'_0, \mu'_1, \phi'\}$

[To be added.]