Supplement to "Bayesian inference on structural impulse response functions"

Mikkel Plagborg-Møller
Department of Economics, Princeton University

This document contains Supplemental Material for the article "Bayesian inference on structural impulse response functions." It contains a simulation study illustrating the workings of the SVMA procedure; supporting results and discussion for the news shock application; technical details concerning SVMA likelihood evaluation, reweighting, the Hamiltonian Monte Carlo routine, and posterior consistency; and supplemental proofs.

Appendix C: Online appendix

C.1 *Simulation study*

To illustrate the workings of the SVMA approach, I conduct a small simulation study with two observed variables and two shocks. I show that prior information about the smoothness of the IRFs can substantially sharpen posterior inference. It is thus desirable to use an approach, like the SVMA approach, for which prior information about smoothness is directly controlled. I also illustrate the consequences of misspecifying the prior.

The illustration is based on the bivariate example from Section 2 with $n = 2$ and $q = 10$; cf. Figure 1. The number of parameters is $n^2(q + 1) = 2^2(10 + 1) = 44$, smaller than the dimensionality of realistic empirical applications but sufficient to elucidate the flexibility, transparency, and effectiveness of the SVAR approach.

C.1.1 *Parameters, prior, and simulation settings*   I consider a single parametrization, with a prior that is correctly centered but diffuse. The sample size is $T = 200$. The true IRF parameters Θ are the noninvertible ones plotted in Figure 1. The true shock standard deviations are $\sigma_1 = 1$ (monetary policy shock) and $\sigma_2 = 0.5$ (demand shock). I first show results for the prior specification in Figure 3 with $\rho_{ij} = 0.9$ for all (i, j). The prior is centered at the true values but it expresses significant prior uncertainty about the magnitudes of the individual impulse responses. The prior on $\sigma = (\sigma_1, \sigma_2)$ is median-unbiased for the true values but it is very diffuse, with prior standard deviation of $\log \sigma_j$ equal to $\tau_j^\sigma = 2$ for $j = 1, 2$.

I simulate a single sample of artificial data from the Gaussian SVMA model and then run the HMC algorithm using the Whittle likelihood (I do not reweight the draws as in Section C.3.2). I take 10,000 MCMC steps, storing every 10th step and discarding the first

Mikkel Plagborg-Møller: mikkelpm@princeton.edu

3,000 steps as burn-in.[1] The full computation takes less than 3 hours in Matlab 8.6 on a personal laptop with 2.3 GHz Intel CPU and 8 GB RAM. Below I provide graphical diagnostics on the convergence and mixing of the MCMC chain.

C.1.2 *Baseline results*   Figure C.1 shows that the posterior for the IRFs accurately estimates the true values and that the data serves to substantially reduce the prior uncertainty. The posterior means are generally close to the truth, although the means for two of the IRFs are slightly too low in this simulation. The 5–95 percentile posterior credible intervals are mostly much narrower than the prior 90% confidence bands, so this prior specification successfully allows the researcher to learn from the data about the magnitudes of the impulse responses. Figure C.2 shows the posterior draws for the shock standard deviations and compares them with the prior distribution. The posterior draws are tightly centered around the true values despite the very diffuse prior on σ. Overall, the inference method for this choice of prior works well, despite the noninvertibility of the true IRFs.

To illustrate the importance of prior information about the smoothness of the IRFs, I run the HMC algorithm with the same specification as above, except that I set $\rho_{ij} = 0.3$ for all (i, j) in the prior. Figure C.3 summarizes the posterior distribution of the IRFs corresponding to this alternative prior. Compared to Figure C.1, the posterior credible
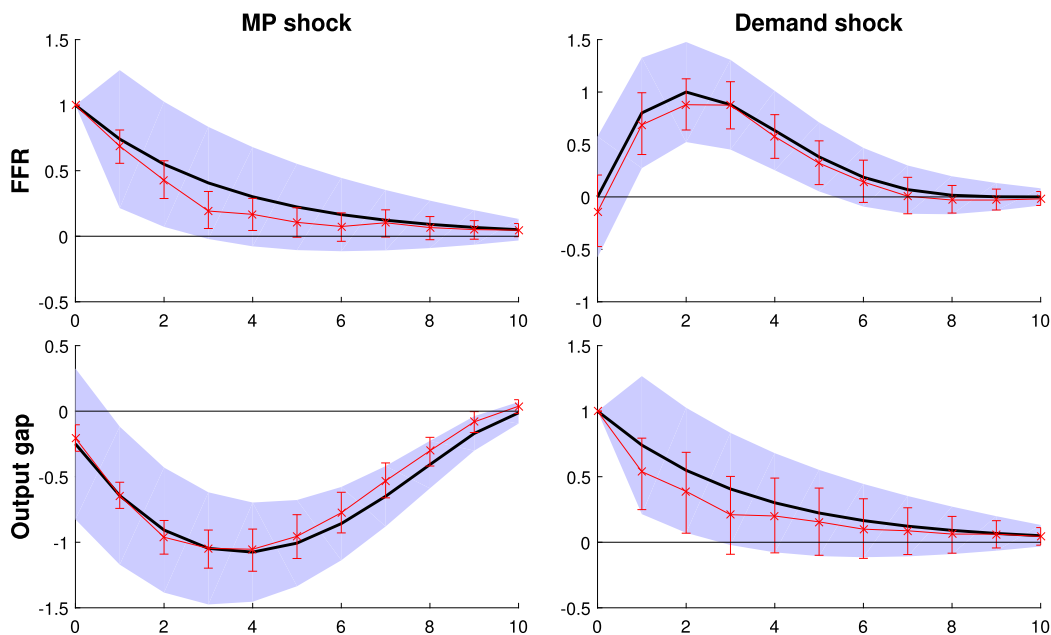

FIGURE C.1. Summary of posterior IRF (Θ) draws for the bivariate SVMA model with prior smoothness $\rho_{ij} = 0.9$. The plots show true values and prior means (thick lines), prior 90% confidence bands (shaded), posterior means (crosses), and posterior 5–95 percentile intervals (vertical bars).

---

[1]The results are virtually identical in simulations with 100,000 MCMC steps.
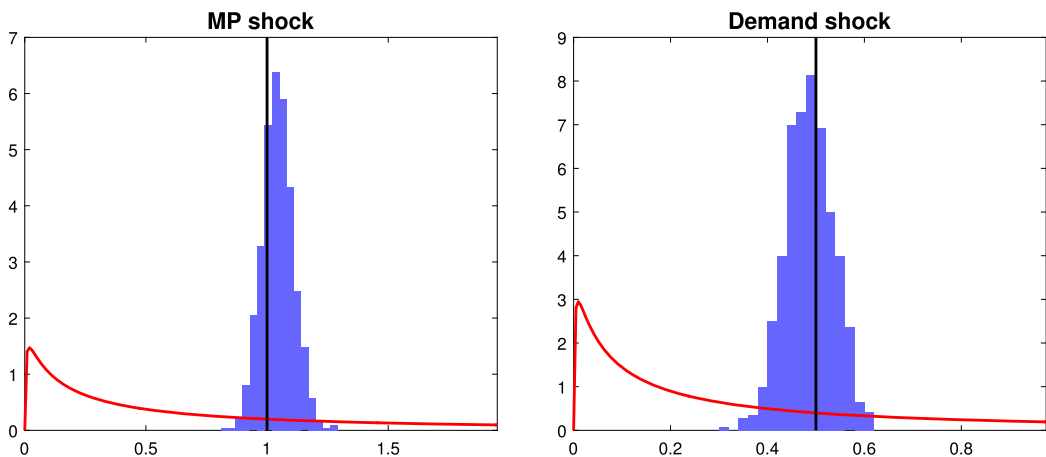
FIGURE C.2. Summary of posterior shock standard deviation (σ) draws for the bivariate SVMA model with prior smoothness $\rho_{ij} = 0.9$. The plots show the true value (thick vertical line), prior density (curve), and histogram of posterior draws, for each σ_j, $j = 1, 2$.
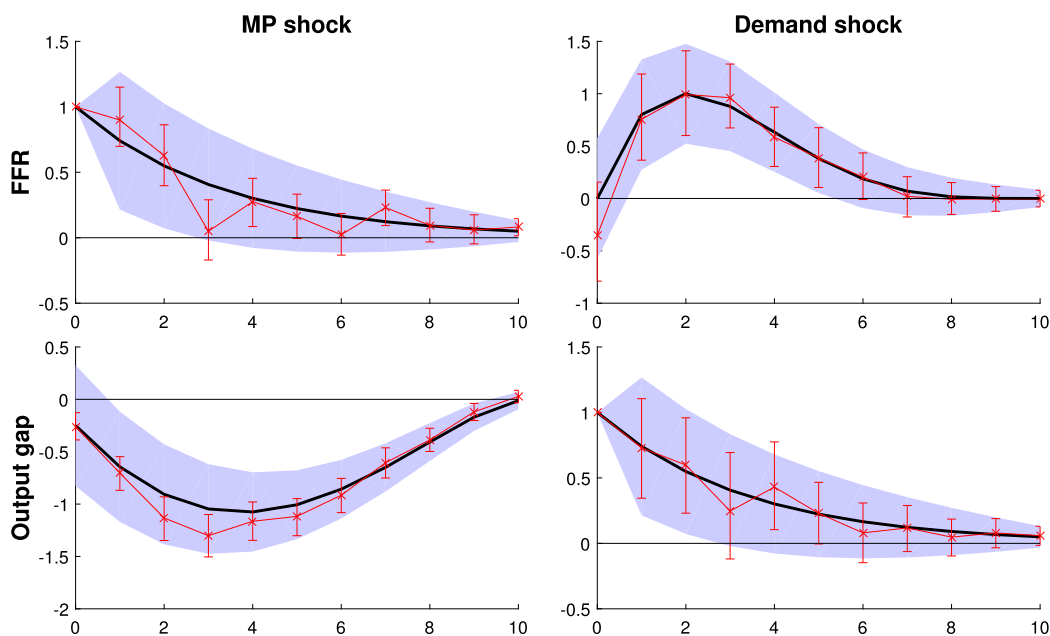

FIGURE C.3. Summary of posterior IRF (Θ) draws for the bivariate SVMA model with prior smoothness $\rho_{ij} = 0.3$. See the caption for Figure C.1.

intervals are much wider and the posterior means are less accurate estimates of the true IRFs.

The higher the degree of prior smoothness, the more do nearby impulse responses "learn from each other." Due to the prior correlation structure in equation (7), any feature of the data that is informative about the impulse response $\Theta_{ij,\ell}$ is also informa-

tive about $\Theta_{ij,\ell+k}$; more so for smaller values of $|k|$, and more so for larger values of the smoothness hyperparameter ρ_{ij}. Hence, a higher degree of prior smoothness reduces the effective number of free parameters in the model. If the true IRFs are not smooth but the prior imposes a lot of smoothness, posterior inference can be very inaccurate. It is therefore important to use a framework, like the SVMA approach, where prior smoothness is naturally parametrized and directly controlled. SVAR IRFs also impose smoothness a priori, but the degree of smoothness is implicitly controlled by the VAR lag length and restrictions on the VAR coefficients.

C.1.3 *Misspecified priors*   I now report results for modifications of the baseline simulation above, considering varying degrees of misspecified, that is, inaccurately centered priors. I consider four such experiments: one in which the shocks have less persistent effects than the prior indicates, one in which the true IRF of the output gap to a monetary policy shock is zero everywhere, one that combines the previous two experiments while increasing the prior variance, and finally one where the data generating parameters are observationally equivalent with the prior mode but economically very different. In the first three cases, the inaccurate prior is overruled by the data, delivering reasonably accurate posterior inference. This happens because the implied prior distribution of the ACF is inconsistent with the true ACF. Since the data is informative about the latter, the posterior distribution puts more weight than the prior on parameters that are consistent with the true ACF, as shown formally in Section 5.2. In the fourth experiment, the posterior inference is necessarily misleading, since the prior mode and the true parameters are associated with the same ACF.

I first consider an experiment in which the prior overstates the persistence of the shock effects, that is, the true IRFs die out quicker than indicated by the prior means $\mu_{ij,\ell}$ in Figure 3. The true IRFs are set to $\Theta_{ij,\ell} = c_{ij}e^{-0.25\ell}\mu_{ij,\ell}$ for all (i,j,ℓ), where $c_{ij} > 0$ is chosen so that $\max_\ell |\Theta_{ij,\ell}| = \max_\ell |\mu_{ij,\ell}|$ for each IRF. The true shock standard deviations, the prior ($\rho_{ij} = 0.9$), the sample size, and the HMC settings are exactly as in Section C.1.2. Figure C.4 compares these true IRFs to the prior distribution. The figure also summarizes the posterior distribution for the IRFs, for a single simulated data set. The posterior is not perfectly centered but is much closer to the truth than the prior is. Figure C.5 shows why this is the case: The prior distribution on (Θ, σ) implies a distribution for auto- and cross-correlations of observed variables that is at odds with the true ACF. Since the data is informative about the ACF, the posterior distribution for IRFs puts higher weight than the prior on IRFs that are consistent with the true auto- and cross-correlations.

The second experiment considers a prior that misspecifies the cross-correlations between the observed variables. I set the true IRFs equal to the prior means in Figure 3, except that the true IRF of the output gap to a monetary policy shock equals zero, that is, $\Theta_{21,\ell} = 0$ for $0 \leq \ell \leq q$. The true shock standard deviations, the prior ($\rho_{ij} = 0.9$), the sample size, and the HMC settings are as above. Figure C.6 shows that posterior inference is accurate despite the misspecified prior. Again, Figure C.7 demonstrates how the data corrects the prior distribution on auto- and cross-correlations, thus pulling the posterior on IRFs toward the true values (although here the true ACF is not estimated as accurately as in Figure C.5).
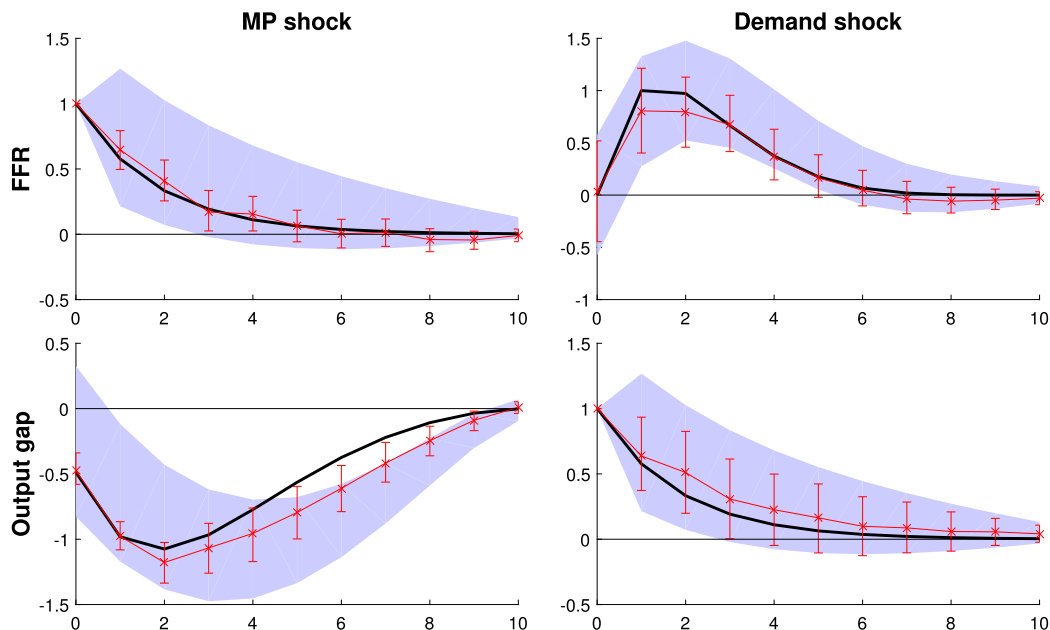
FIGURE C.4. Summary of posterior IRF (Θ) draws for the bivariate SVMA model with a prior that is too persistent relative to the true parameter values. The plots show true values (thick lines), prior 90% confidence bands (shaded), posterior means (crosses), and posterior 5–95 percentile intervals (vertical bars). The prior means (not shown) are the midpoints of the prior confidence bands, as in Figure 3.

The third experiment combines the previous two while increasing the prior variance. As true parameters, I use the IRFs from the first experiment (misspecified persistence) except that the IRF of the output gap to the monetary policy shock is set to zero at all horizons (misspecified cross-correlations, as in the second experiment). I use the baseline prior ($\rho_{ij} = 0.9$) from Section C.1.2, except that the prior standard deviations $\tau_{ij,\ell}$ are 50% larger to illustrate the consequences of having a more diffuse prior. The true shock standard deviations, the sample size, and the HMC settings are as in the baseline. Figure C.8 shows that the FFR IRFs are less accurately estimated in this experiment, yet the credible intervals are not misleading even through the prior is poorly centered and quite diffuse.

The last experiment illustrates that the ability of the data to correct a misspecified prior is limited, as must necessarily be the case due to partial identification (cf. Sections 2.4 and 5). In this experiment with a grossly misspecified prior, I choose the true IRFs and shock standard deviations to be observationally equivalent with—but economically very different from—the prior mode. This is accomplished by flipping a root of the SVMA polynomial and applying a rotation (cf. the terminology in Appendix A.2). The Matlab code that accompanies this paper provides the details of the chosen values of Θ and σ. The prior ($\rho_{ij} = 0.9$), the sample size, and the HMC settings are as in the baseline experiment in Section C.1.2. Figure C.9 plots the chosen true IRFs as well as the posterior distribution, for a simulated data set. Due to the observational equivalence of the
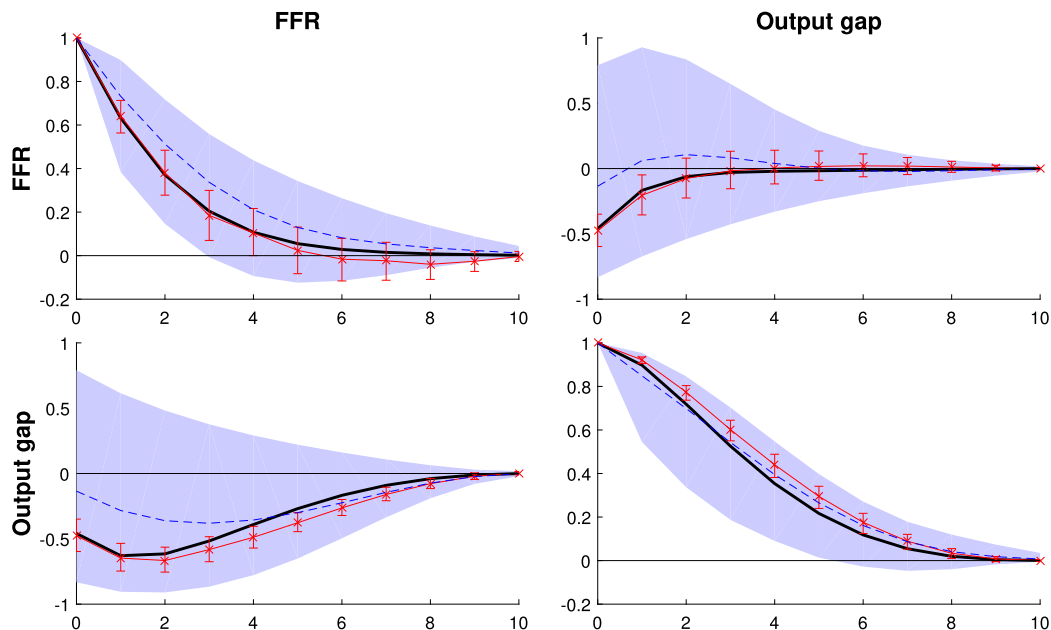
FIGURE C.5. Posterior auto- and cross-correlation draws for the bivariate SVMA model with a prior that misspecifies the persistence of the IRFs. The displays plot draws of $\text{Corr}(y_{i,t}, y_{j,t-k} | \Theta, \sigma)$, where i indexes rows, j indexes columns, and k runs along the horizontal axes. The top right display, say, concerns cross-correlations between the FFR and lags of the output gap. The plots show true values (thick lines), prior means (dashed lines) and 5–95 percentile confidence bands (shaded), and posterior means (crosses) and 5–95 percentile intervals (vertical bars).

true parameters to those used in Section C.1.2, the posterior distribution in Figure C.9 is very similar to Figure C.1. Because the ACF of the data here cannot distinguish between the true parameters and the prior mode, the data does not meaningfully correct the mistaken prior beliefs: As in all partially identified models, there exist very inaccurate priors that yield very inaccurate conclusions.

C.1.4 *MCMC diagnostics*   I report diagnostics for the baseline $\rho_{ij} = 0.9$ bivariate simulation in Section C.1.2, but diagnostics for other specifications in this section are similar. The average HMC acceptance rate is slightly higher than 0.60, which is the rate targeted by the NUTS algorithm when tuning the HMC step size. The score of the posterior was evaluated about 382,000 times. Figures C.10 and C.11 show the MCMC chains for the IRF and log shock standard deviation draws. Figures C.12 and C.13 show the autocorrelation functions of the draws.

C.2 *Application: Additional results and discussion*

This subsection presents additional details, results, and discussion related to the news shock application in Section 4. First, I detail the data construction. Second, I summarize
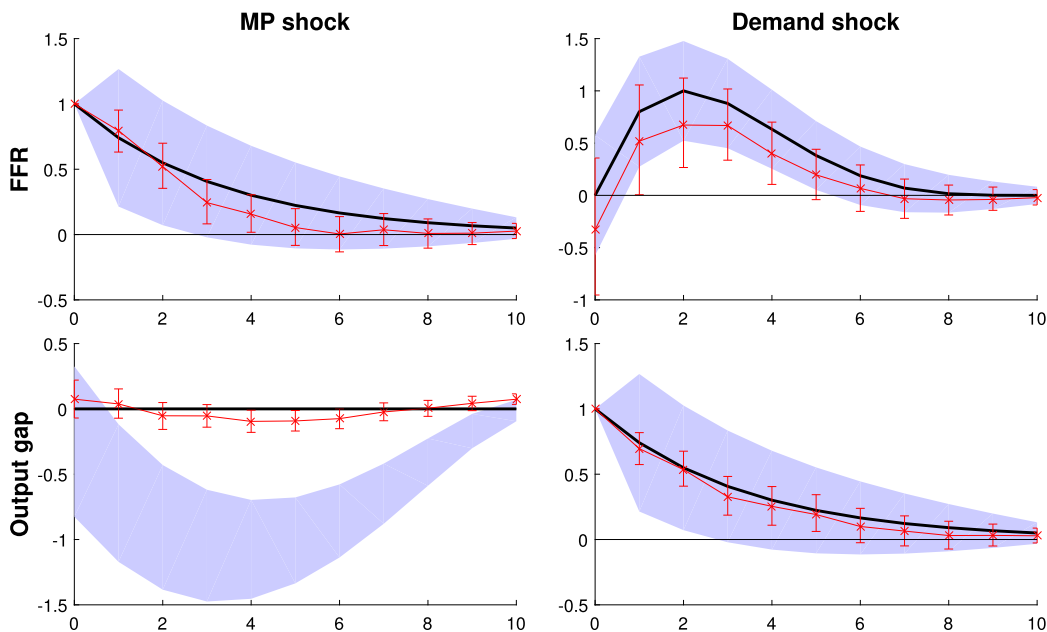
FIGURE C.6. Summary of posterior IRF (Θ) draws for the bivariate SVMA model with a prior that misspecifies the cross-correlations between variables. See the caption for Figure C.4.
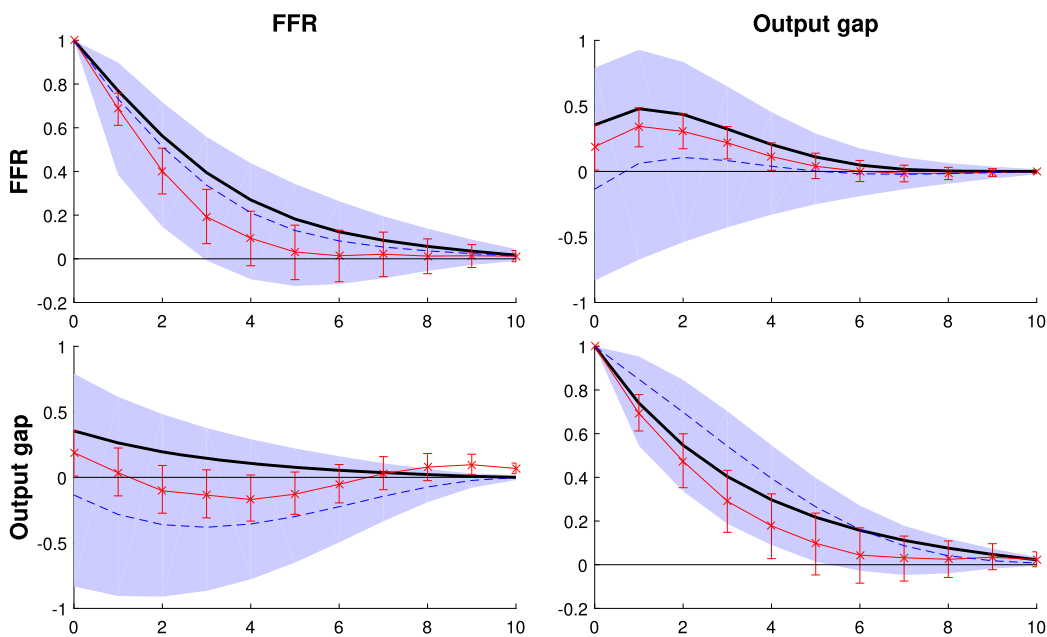

FIGURE C.7. Posterior autocorrelation draws for the bivariate SVMA model with a prior that misspecifies the cross-correlations between variables. See the caption for Figure C.5.
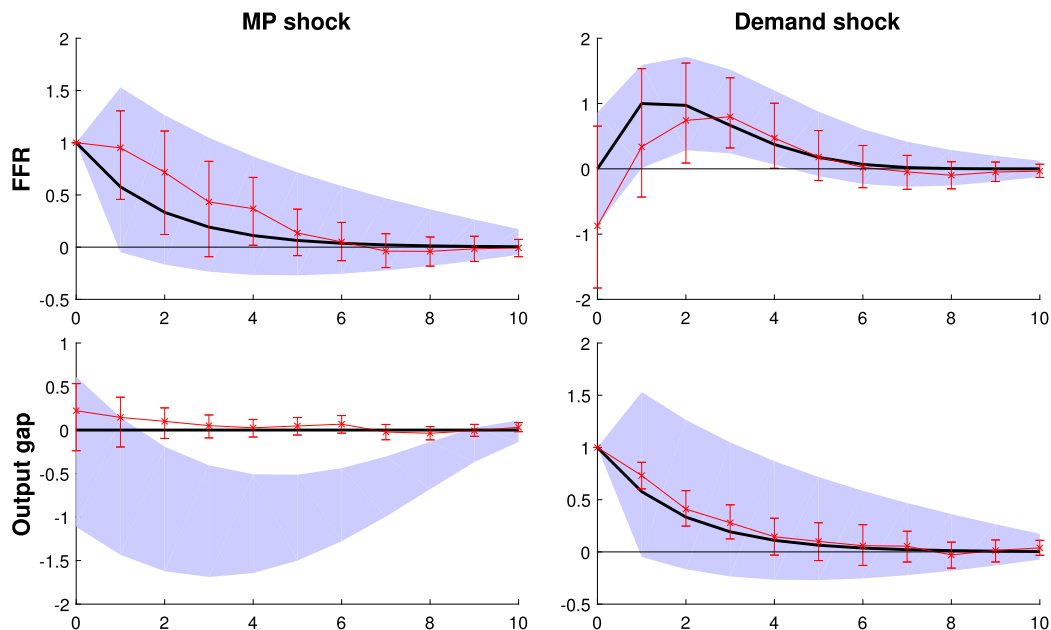
FIGURE C.8. Summary of posterior IRF (Θ) draws for the bivariate SVMA model with a prior that misspecifies both the persistence and cross-correlations of variables, while increasing the prior standard deviations by 50% relative to the baseline. See the caption for Figure C.4.
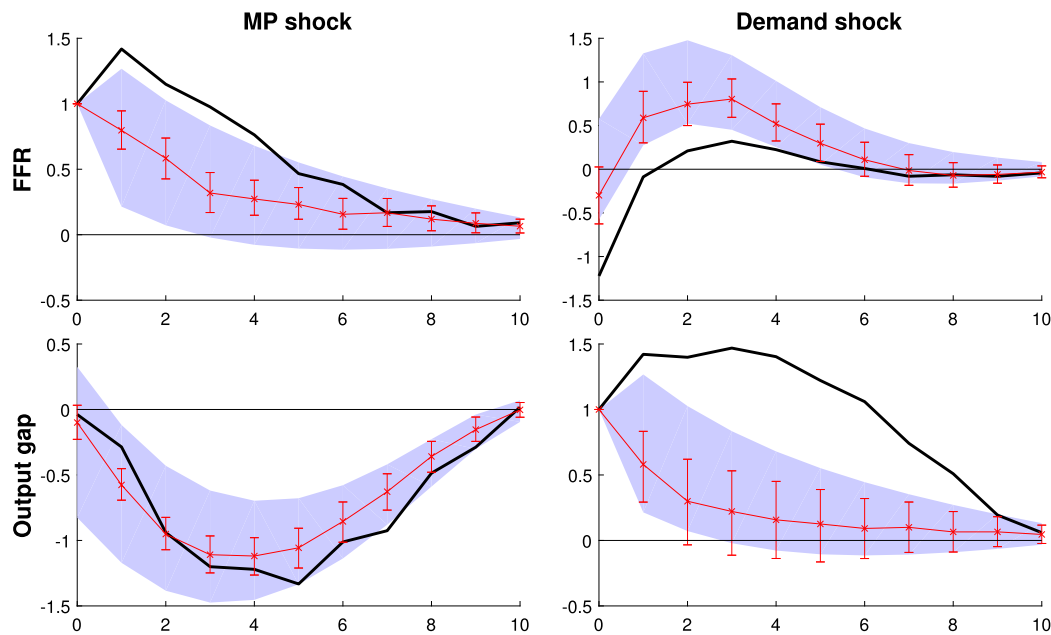

FIGURE C.9. Summary of posterior IRF (Θ) draws for the bivariate SVMA model with a grossly misspecified prior whose mode is nevertheless observationally equivalent with the true IRFs. See the caption for Figure C.4.

FIGURE C.10. MCMC chains for each IRF parameter (Θ) in the $\rho_{ij} = 0.9$ simulations in Section C.1.2. Each jagged line represents a different impulse response parameter (two of them are normalized at 1). The vertical dashed line marks the burn-in time, before which all draws are discarded. The horizontal axes are in units of MCMC steps, not stored draws (every 10th step is stored).

FIGURE C.11. MCMC chains for each log shock standard deviation parameter ($\log \sigma$) in the $\rho_{ij} = 0.9$ simulations in Section C.1.2. See the caption for Figure C.10.

the posteriors for long-run (cumulative) impulse responses. Third, I demonstrate the economic importance of noninvertibility. Fourth, I use the Kalman smoother to draw inference about the shocks. Fifth, I examine the sensitivity of posterior inference to the choice of prior. Sixth, I assess the model's fit. Seventh, I compare my empirical results
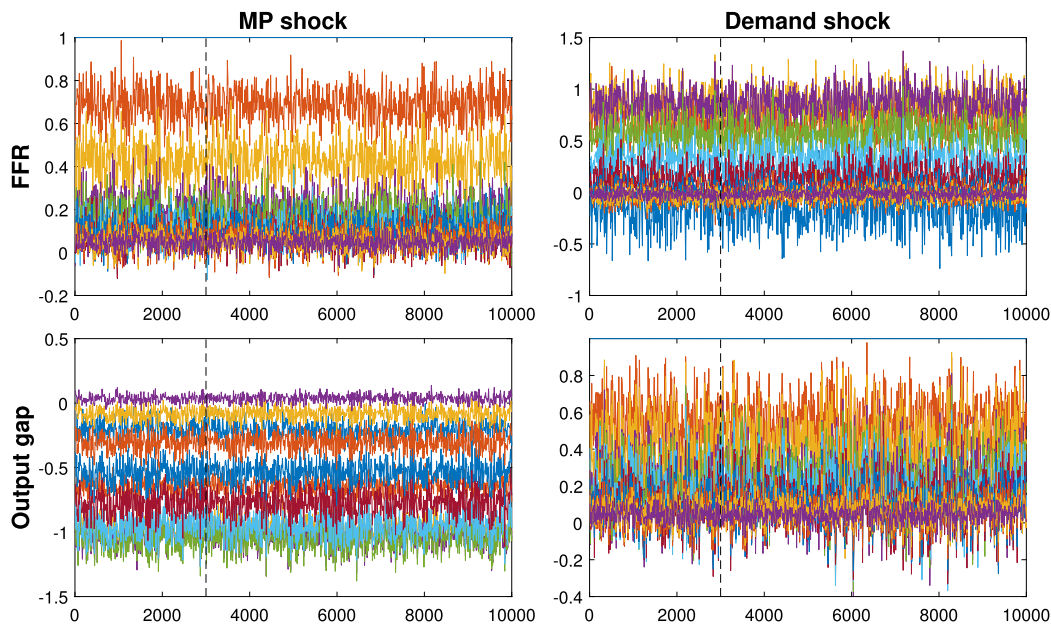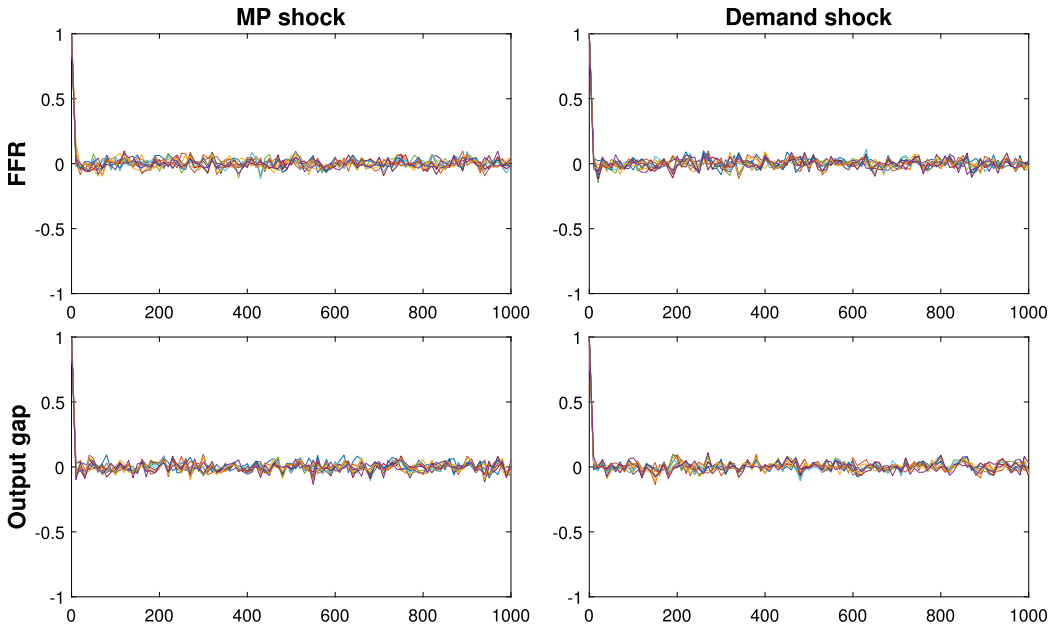
FIGURE C.12. Autocorrelation functions for HMC draws of each IRF parameter (Θ) in the $\rho_{ij} = 0.9$ simulations in Section C.1.2. Each jagged line represents a different impulse response parameter. Only draws after burn-in were used to computed these figures. The autocorrelation lag is shown on the horizontal axes in units of MCMC steps.
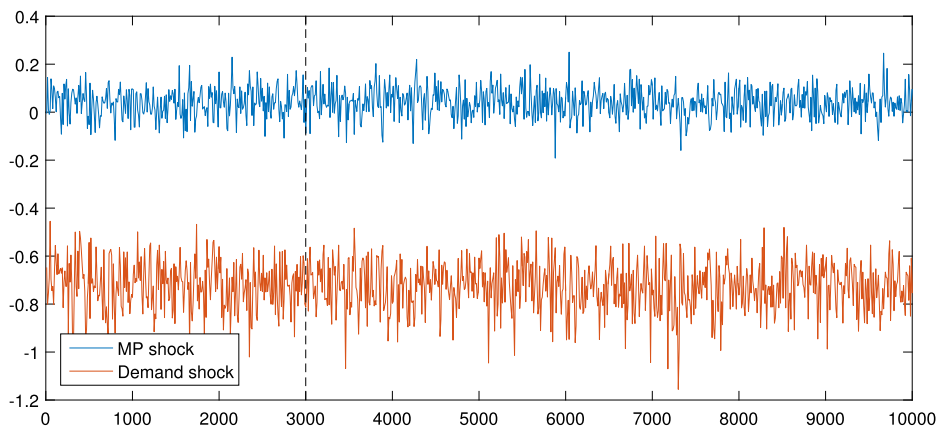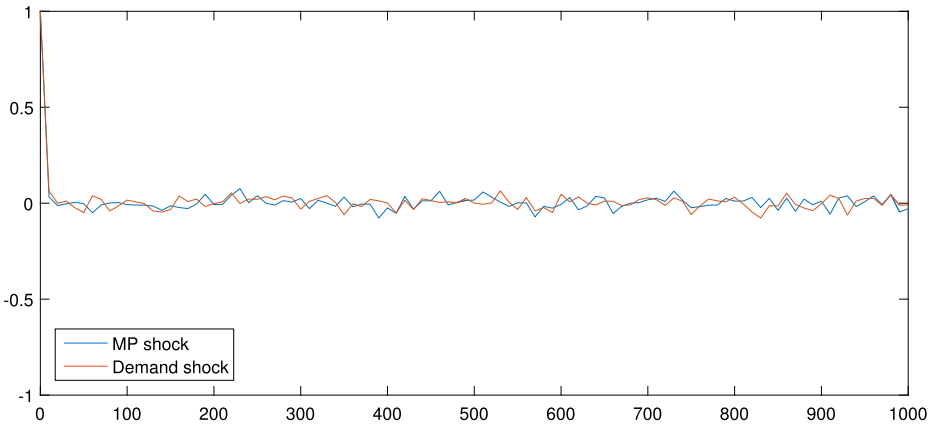

FIGURE C.13. Autocorrelation functions for HMC draws of each log shock standard deviation parameter ($\log \sigma$) in the $\rho_{ij} = 0.9$ simulations in Section C.1.2. See the caption for Figure C.12.

to the literature. Finally, I show that the SVMA procedure accurately estimates IRFs on simulated data.

C.2.1 *Data construction*   TFP growth equals 100 times the log growth rate of TFP and is taken from the data appendix to Fernald (2014).[2] The remaining data is from the St. Louis Federal Reserve's FRED database.[3] Real GDP growth is given by 100 times the log growth rate of seasonally adjusted GDP per capita in chained dollars, as measured by the Bureau of Economic Analysis (NIPA Table 7.1, line 10). My real interest rate series equals the nominal policy interest rate minus the contemporaneous inflation rate.[4] The nominal policy rate is the average effective federal funds rate, expressed as a quarterly rate. The inflation rate equals 100 times the log growth rate in the seasonally adjusted implicit price deflator for the nonfarm business sector, as reported by the Bureau of Labor Statistics.

I detrend the three data series to remove secular level changes that are arguably unrelated to the business cycle. Following Stock and Watson (2012), I estimate the trend in each series using a biweight kernel smoother with a bandwidth of 100 quarters; the trends are then subtracted from the raw series. The data is plotted in the Online Supplemental Material.

Figure C.14 plots the raw data and estimated trends.

C.2.2 *Long-run responses*   Figure C.15 plots the posterior distribution of long-run (i.e., cumulative) impulse responses $\sum_{\ell=0}^{q} \Theta_{ij,\ell}$ for each variable-shock combination (i, j).

C.2.3 *Economic significance of noninvertibility*   The noninvertibility of the estimated IRFs is economically significant. Figure C.16 summarizes the posterior distribution of those invertible IRFs that are closest to the actual (possibly noninvertible) IRFs. Specifically, for each posterior draw (Θ, σ) I compute the parameter vector $(\tilde{\Theta}, \tilde{\sigma})$ that minimizes the Frobenius distance $\|\Theta \operatorname{diag}(\sigma) - \tilde{\Theta} \operatorname{diag}(\tilde{\sigma})\|$ over parameters for which $\tilde{\Theta}$ is invertible and $(\tilde{\Theta}, \tilde{\sigma})$ generates the same ACF as (Θ, σ).[5] While the invertible IRFs for the productivity and monetary policy shocks are similar to the unrestricted IRFs, the invertible news shock IRFs look nothing like the actual estimated IRFs.[6] Thus, no SVAR identification scheme can deliver accurate inference about the effects of technological news shocks in this dataset.

---

[2] The TFP measure is based on a growth accounting method that adjusts for differing marginal products of capital across sectors as well as changes over time in labor quality and labor's share of income. Fernald (2014) also estimates utilization-adjusted TFP, but the adjustment is model-based and reliant on estimates from annual regressions on a separate dataset, so I prefer the simpler series. Data downloaded July 14, 2015.

[3] FRED series codes: A939RX0Q048SBEA (real GDP per capita), FEDFUNDS (effective federal funds rate), and IPDNBS (implicit price deflator, nonfarm business sector). Data was downloaded August 13, 2015.

[4] If agents form inflation expectations under the presumption that quarterly inflation follows a random walk, then my measure of the real interest rate equals the conventional ex ante real interest rate.

[5] According to Appendix A.2, $(\tilde{\Theta}, \tilde{\sigma})$ is obtained as follows. First, apply transformation (ii) in Proposition 2 several times to (Θ, σ) in order to flip all roots outside the unit circle. Denote the resulting invertible parameters by $(\check{\Theta}, \check{\sigma})$. Then $\tilde{\Theta} \operatorname{diag}(\tilde{\sigma}) = \check{\Theta} \operatorname{diag}(\check{\sigma}) Q$, where Q is the orthogonal matrix that minimizes $\|\Theta \operatorname{diag}(\sigma) - \check{\Theta} \operatorname{diag}(\check{\sigma}) Q\|$. This is an "orthogonal Procrustes problem," whose solution is well known.

[6] Figure C.16 cannot be interpreted as the posterior distribution corresponding to a prior which truncates the prior from Figure 5 to the invertible region. It is difficult to sample from this truncated posterior, as essentially none of the unrestricted posterior draws are invertible, so an accept–reject scheme is inapplicable.
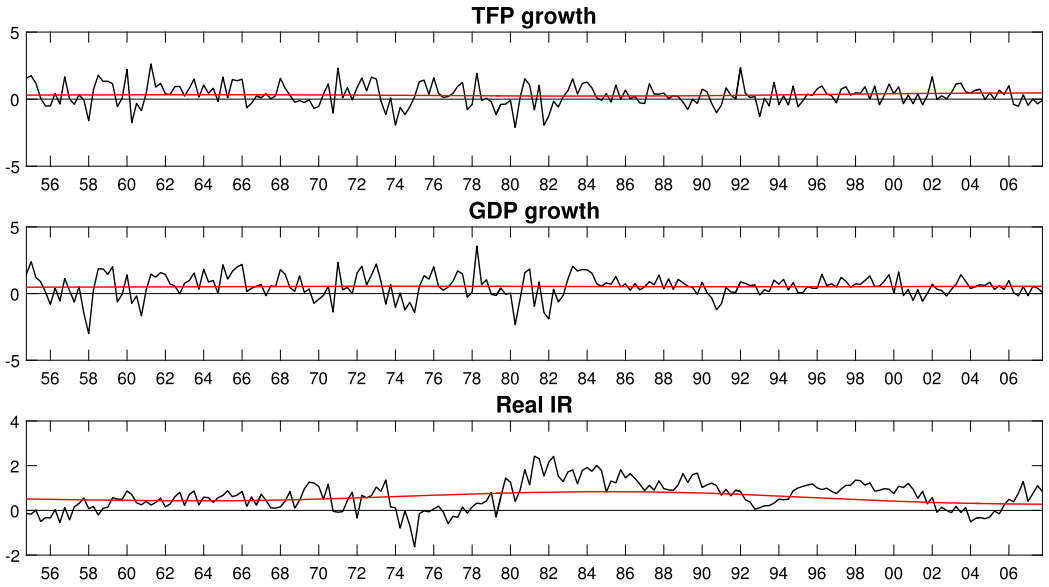
FIGURE C.14. Raw data on TFP growth, GDP growth, and the real interest rate (IR), along with estimated time-varying trends (smooth curves). The final data used in the empirical analysis are differences between the raw series and the trends.
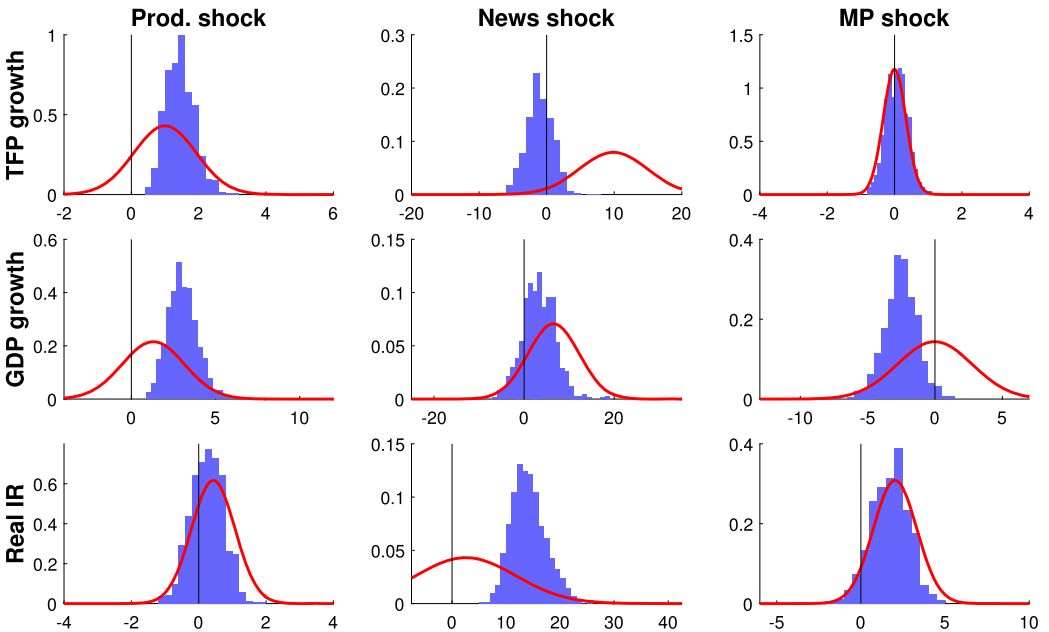
FIGURE C.15. Histograms of posterior draws of long-run impulse responses $\sum_{\ell=0}^{q} \Theta_{ij,\ell}$ for each (i, j), news shock application. Curves are prior densities. Histograms and curves each integrate to 1.
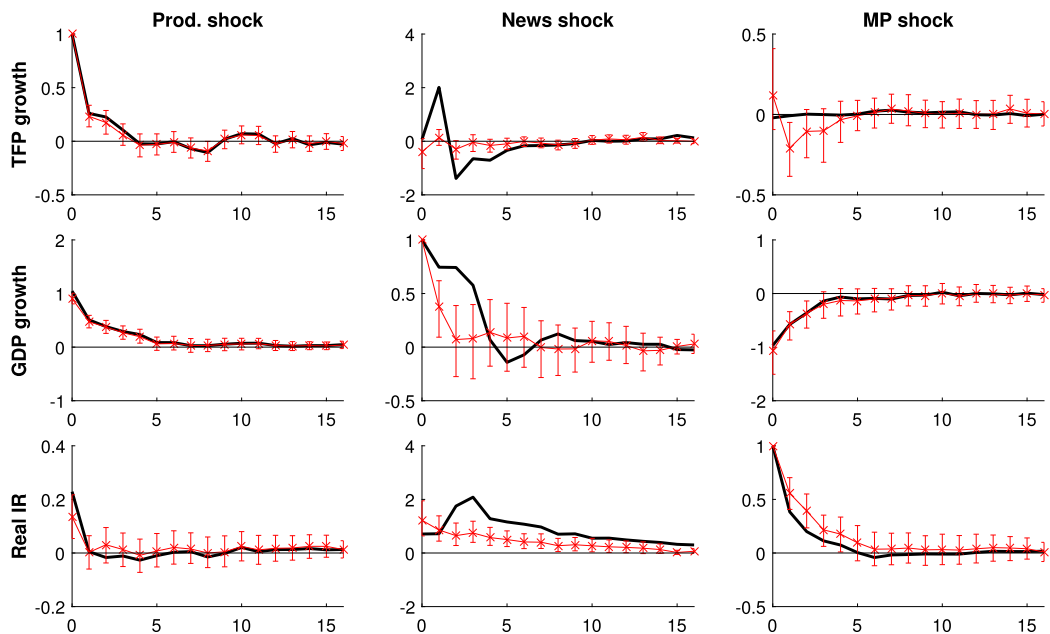
FIGURE C.16. Posterior distribution of the invertible IRFs that are closest to the actual IRFs, news shock application. The figure shows posterior means of actual IRFs from Figure 6 (thick lines), posterior means of the closest invertible IRFs (crosses), and posterior 5–95 percentile intervals for these invertible IRFs (vertical bars).

C.2.4 *Inference about shocks*   Figure C.17 shows the time series of posterior means for the structural shocks. For each posterior draw of the structural parameters (Θ, σ), I compute $E(\varepsilon_t | \Theta, \sigma, Y_T)$ using the smoothing recursions corresponding to the Gaussian state-space representation in Section C.3.1 (Durbin and Koopman (2012, p. 157)), and then I average over draws. If the structural shocks are in fact non-Gaussian, the smoother still delivers mean-square-error-optimal linear estimates of the shocks. If desired, draws from the full joint posterior distribution of the shocks can be obtained from a simulation smoother (Durbin and Koopman (2012, Chapter 4.9)). It is also straightforward to draw from the predictive distribution of future values of the data using standard methods for state-space models.

C.2.5 *Prior sensitivity*   Here, I consider the sensitivity of posterior inference with respect to local and global changes in the prior.

To gauge the local robustness of posterior inference with respect to the choice of prior, I compute the sensitivity measure "PS" of Müller (2012). This measure captures the first-order approximate effect on the posterior means of changing the prior mean hyperparameters. Let θ denote the vector containing all impulse responses and log shock standard deviations of the SVMA model, and let e_k denote the k-th unit vector. Because my prior for θ is a member of an exponential family, the Müller (2012) PS measure for
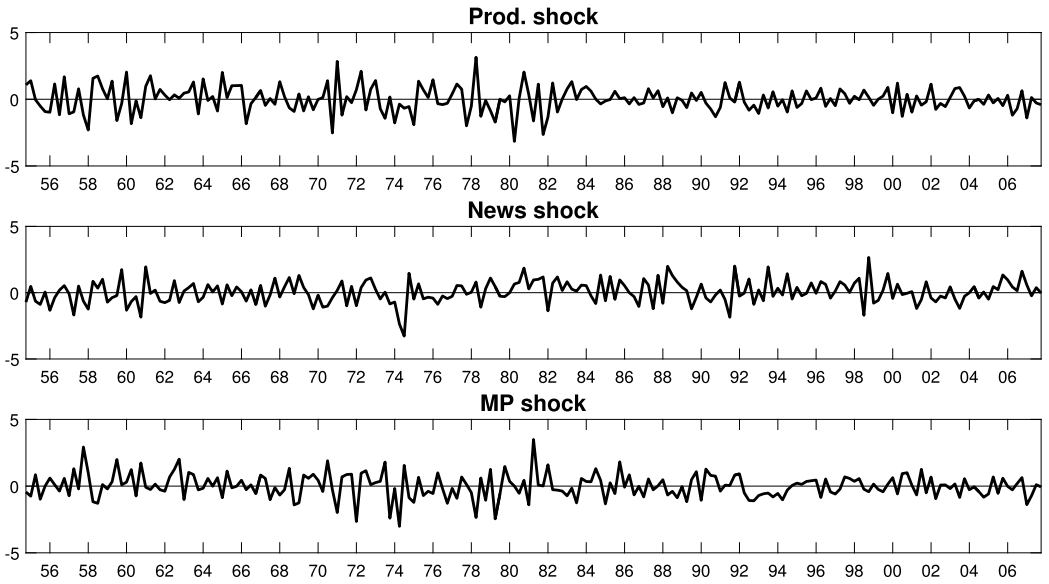
FIGURE C.17. Posterior means of *standardized* structural shocks ($\varepsilon_{jt}/\sigma_j$) at each point in time, news shock application.

parameter θ_k equals

$$PS_k = \max_{\nu:\ \sqrt{\nu'\,\mathrm{Var}(\theta)^{-1}\nu}=1} \frac{\partial E(\theta_k|Y_T)}{\partial E(\theta)'}\nu = \sqrt{e_k'\,\mathrm{Var}(\theta|Y_T)\,\mathrm{Var}(\theta)^{-1}\,\mathrm{Var}(\theta|Y_T)e_k}. \qquad \text{(C.1)}$$

This is the largest (local) change that can be induced in the posterior mean of θ_k from changing the prior means of the components of θ by the multivariate equivalent of 1 prior standard deviation.[7] PS_k depends only on the prior and posterior variance matrices $\mathrm{Var}(\theta)$ and $\mathrm{Var}(\theta|Y_T)$, which are easily obtained from the HMC output.

Figure C.18 plots the posterior means of the impulse responses along with $\pm PS_k$ intervals (where the index k corresponds to the (i, j, ℓ) combination for each impulse response). The wider the band around an impulse response, the more sensitive is the posterior mean of that impulse response to (local) changes in the prior. In economic terms, most of the posterior means are seen to be insensitive to changes in the prior means of magnitudes smaller than 1 prior standard deviation. The most prior-sensitive posterior inferences, economically speaking, concern the IRF of GDP growth to a news shock, but large changes in the prior means are necessary to alter the qualitative features of the posterior mean IRF.

As an indication of *global* prior sensitivity, I additionally present posterior results for an alternative prior which differs from the baseline by doubling the prior standard deviation of the IRFs of the real interest rate to technology and monetary policy shocks. I thank a referee for pointing out that the priors on these two IRFs may be too tight for

---

[7]In particular, $PS_k \geq \max_b |\partial E(\theta_k|Y_T)/\partial E(\theta_b)|\sqrt{\mathrm{Var}(\theta_b)}$. Whereas PS_k is a local measure, the effects of large changes in the prior can be evaluated using reweighting (Lopes and Tobias (2011, Section 2.4)).
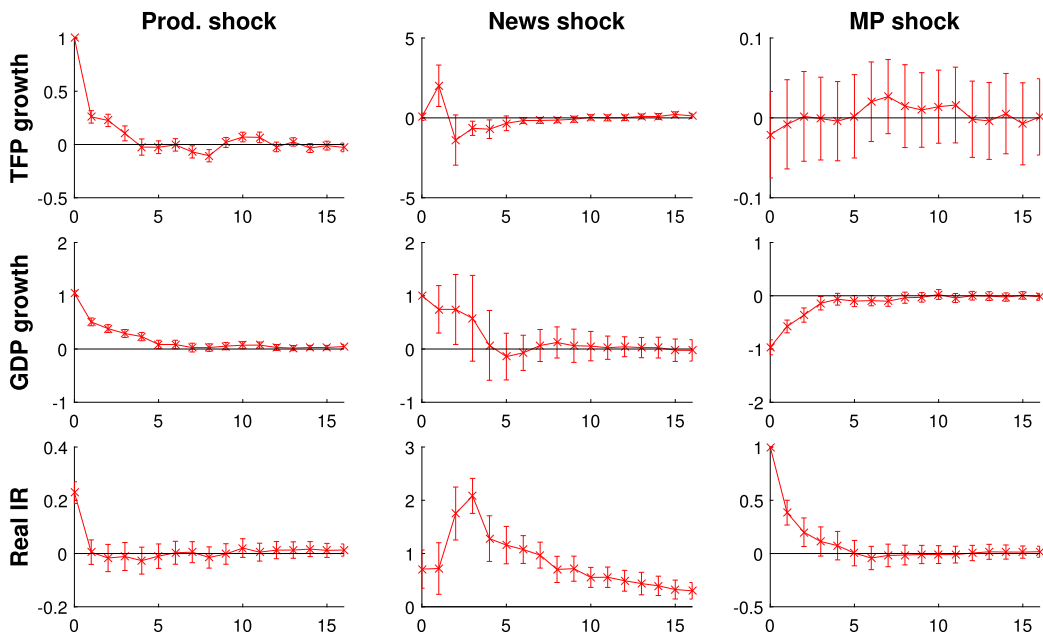
FIGURE C.18. PS_k measure of the sensitivity of the posterior IRF means with respect to changes in the prior means of all parameters (cf. (C.1)) in the news shock application. The symmetric vertical bars have length $2PS_k$ and are centered around the corresponding posterior means (crosses).

some readers. Figure C.19 shows that the posterior inference is mostly unchanged under this less informative prior, although the posterior credible intervals are somewhat wider, particularly for the IRFs of the real interest rate to news and monetary policy shocks. The mixing of the MCMC chain is slower with this prior, as the posterior appears to be slightly multimodal; hence, I here report results based on 30,000 MCMC steps after burn in (storing every 10th draw), which seems sufficient to explore the full posterior distribution.

C.2.6 *Posterior predictive analysis*   I conduct a posterior predictive analysis to identify ways to improve the fit of the Gaussian SVMA model (Geweke (2010, Chapter 2.4.2)). For each posterior parameter draw produced by HMC, I simulate an artificial data set of sample size $T = 213$ from a Gaussian SVMA model with the given parameters. On each artificial dataset I compute four checking functions. First and second, the skewness and excess kurtosis of each series. Third, the long-run autocorrelation of each series, defined as the Newey-West long-run variance estimator (20 lags) divided by the sample variance. Fourth, I run a reduced-form VAR regression of the three-dimensional data vector y_t on its 8 first lags and a constant; then I compute the first autocorrelation of the squared VAR residuals for each of the three series. The third measure captures persistence, while the fourth measure captures volatility clustering in forecast errors.

Figure C.20 shows the distribution of checking function values across simulated datasets, as well as the corresponding checking function values for the actual data. The
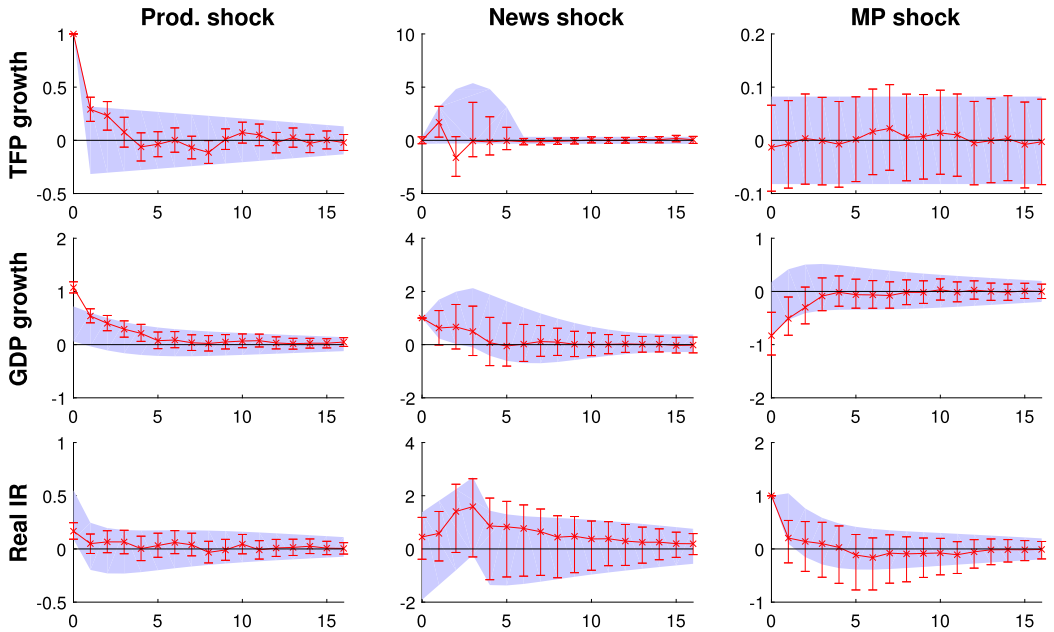
FIGURE C.19. Summary of posterior IRF (Θ) draws, news shock application, alternative prior. The plots show prior 90% confidence bands (shaded), posterior means (crosses), and posterior 5–95 percentile intervals (vertical bars).
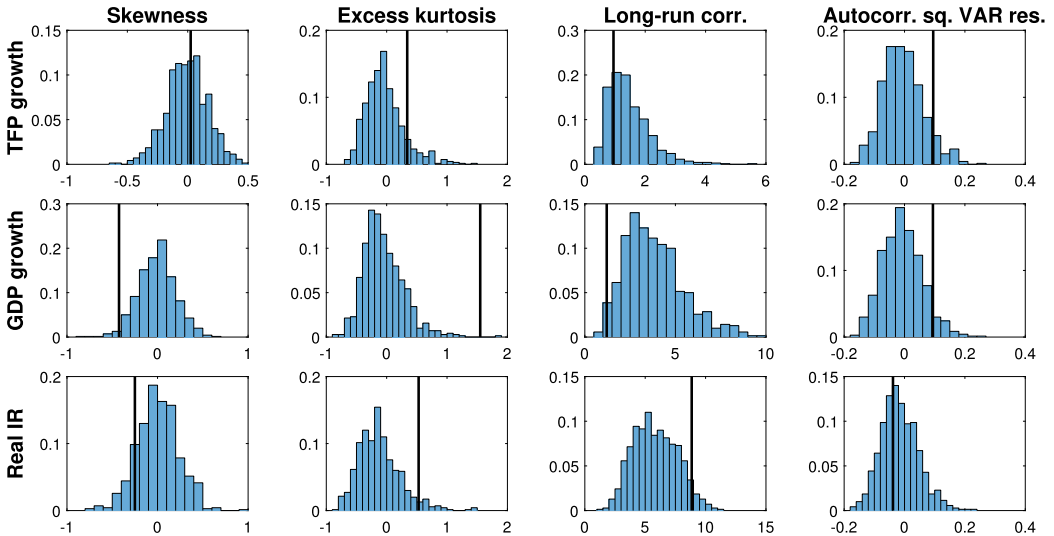
FIGURE C.20. Posterior predictive checks, news shock application. Observed variables along rows, checking functions along columns. Histograms show the distribution of checking function values on simulated data sets based on the posterior parameter draws; thick vertical lines mark checking function values on actual data. Checking functions from left to right: skewness; excess kurtosis; Newey–West long-run variance estimate (20 lags) divided by sample variance; first autocorrelation of squared residuals from a VAR regression of y_t on a constant and 8 lags.

Gaussian SVMA model does not capture the skewness and kurtosis of GDP growth; essentially, the model does not generate recessions that are sufficiently severe relative to the size of booms. The model somewhat undershoots the persistence and kurtosis of the real interest rate. The fourth column suggests that forecast errors for TFP and GDP growth exhibit volatility clustering in the data, which is not captured by the Gaussian SVMA model.

The results point to three fruitful model extensions. First, introducing stochastic volatility in the SVMA model would allow for better fit along the dimensions of kurtosis and forecast error volatility clustering. Second, nonlinearities or skewed shocks could capture the negative skewness of GDP growth. Finally, increasing the MA lag length q would allow the model to better capture the persistence of the real interest rate, although this is not a major concern, as I am primarily interested in shorter-run impulse responses.

C.2.7 *Comparison with the literature*   My conclusion that technological news shocks are not important for explaining business cycles is consistent with the literature, but my method is the first to allow for noninvertibility without additional assumptions. Forni, Gambetti, and Sala (2014) estimate small effects of technological news shocks in a factor-augmented SVAR. Their empirical strategy may overcome the noninvertibility issue if technological news are well captured by the first few principal components of their large macroeconomic panel data set. They confirm that low-dimensional systems (without factors) are noninvertible. Papers that estimate fully-specified DSGE models with news shocks also tend to find a limited role for technological news, cf. the review by Beaudry and Portier (2014, Section 4.2.2). Unlike these papers, I do not dogmatically impose restrictions implied by a particular structural model.

Several SVAR papers on news shocks have used stock market data in an attempt to overcome the invertibility problem; cf. Beaudry and Portier (2014, Section 3). Such SVAR specifications may be valid if the stock market is a good proxy for the news shock, i.e., if the market responds immediately and forcefully upon arrival of technological news. On the other hand, if market movements are highly contaminated by other types of shocks, incorporating stock market data may lead to biased SVAR estimates.

C.2.8 *Consistency check with simulated data*   I show that the SVMA approach, with the same prior and HMC settings as in Section 4, can recover the true IRFs when applied to data generated by the log-linearized Sims (2012) DSGE model. I simulate data for the three observed variables from an SVMA model with i.i.d. Gaussian shocks. The true IRFs are those implied by the log-linearized Sims (2012) model (baseline calibration) out to horizon $q = 16$, yielding a noninvertible representation. The true shock standard deviations are set to $\sigma = (0.5, 0.5, 0.5)'$. Note that the prior for the IRF of TFP growth to the news shock is not centered at the true IRF, as explained in Section 4. The sample size is the same as for the actual data ($T = 213$).

Figures C.21 and C.22 summarize the posterior draws produced by the HMC algorithm when applied to the simulated data set. The posterior means accurately locate the true parameter values. The equal-tailed 90% posterior credible intervals are tightly concentrated around the truth in most cases. In particular, inference about the shock standard deviation parameters is precise despite the very diffuse prior.
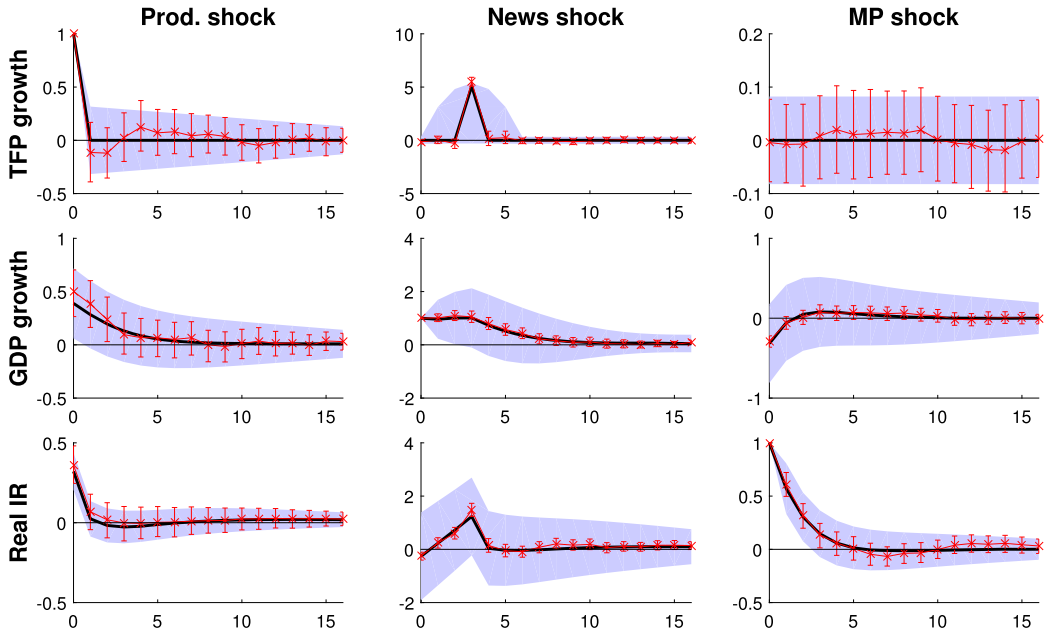
FIGURE C.21. Summary of posterior IRF (Θ) draws, simulated news shock data. See the caption for Figure C.1.
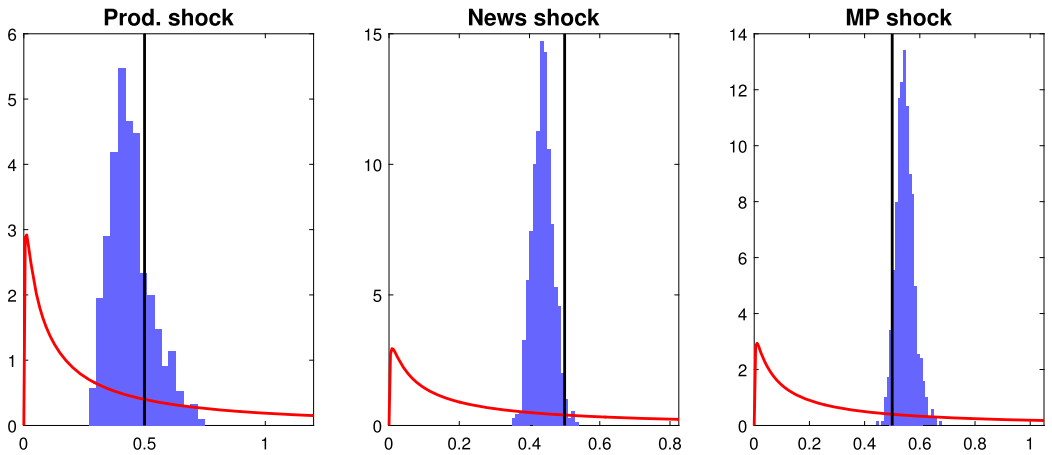

FIGURE C.22. Summary of posterior shock standard deviation (σ) draws, simulated news shock data. See the caption for Figure C.2.

C.3 *Exact likelihood and reweighting*

I now describe the Kalman filter for SVMA models and the reweighting procedure for translating Whittle posterior parameter draws into draws from the exact Gaussian posterior.

C.3.1 *Kalman filter*   The state space representation of the SVMA model is

$$y_{i,t} = \Psi_i \alpha_t, \quad i = 1, \ldots, n, t = 1, \ldots, T,$$

$$\alpha_t = \begin{pmatrix} 0 & 0 \\ I_{nq} & 0 \end{pmatrix} \alpha_{t-1} + \begin{pmatrix} \tilde{\varepsilon}_t \\ 0 \end{pmatrix}, \quad \tilde{\varepsilon}_t \overset{\text{i.i.d.}}{\sim} N(0, I_n), t = 2, 3, \ldots, T,$$

$$\alpha_1 \sim N(0, I_{n(q+1)}),$$

where Ψ_i is the $n(q+1)$-dimensional i-th row vector of $\Psi = (\Psi_0, \Psi_1, \ldots, \Psi_q)$ (with $\Psi_\ell = \Theta_\ell \operatorname{diag}(\sigma)$), $\tilde{\varepsilon}_t$ is the n-dimensional standardized structural shock vector (each element has variance 1), and $\alpha_t = (\tilde{\varepsilon}_t', \tilde{\varepsilon}_{t-1}', \ldots, \tilde{\varepsilon}_{t-q}')'$ is the $n(q+1)$-dimensional state vector.

I use the "univariate treatment of multivariate series" Kalman filter in Durbin and Koopman (2012, Ch. 6.4), since that algorithm avoids inverting large matrices. For my purposes, the algorithm is as follows.

1. Initialize the state forecast mean $a_{1,1} = 0$ and state forecast variance $Z_{1,1} = I_{n(q+1)}$. Set $t = 1$.

2. For each $i = 1, \ldots, n$:

   (a) Compute the forecast error $v_{i,t} = y_{i,t} - \Psi_i a_{i,t}$, forecast variance $\lambda_{i,t} = \Psi_i Z_{i,t} \Psi_i'$, and Kalman gain $g_{i,t} = (1/\lambda_{i,t}) Z_{i,t} \Psi_i'$.

   (b) Compute the log likelihood contribution: $L_{i,t} = -\frac{1}{2}(\log \lambda_{i,t} + v_{i,t}^2/\lambda_{i,t})$.

   (c) Update the state forecast mean: $a_{i+1,t} = a_{i,t} + g_{i,t} v_{i,t}$.

   (d) Update the state forecast variance: $Z_{i+1,t} = Z_{i,t} - \lambda_{i,t} g_{i,t} g_{i,t}'$.

3. Let $\tilde{a}_{n+1,t}$ denote the first nq elements of $a_{n+1,t}$, and let $\tilde{Z}_{n+1,t}$ denote the upper left $nq \times nq$ block of $Z_{n+1,t}$. Set

$$a_{1,t+1} = \begin{pmatrix} 0 \\ \tilde{a}_{n+1,t} \end{pmatrix}, \qquad Z_{1,t+1} = \begin{pmatrix} I_n & 0 \\ 0 & \tilde{Z}_{n+1,t} \end{pmatrix}.$$

4. If $t = T$, stop. Otherwise, increment t by 1 and go to step 2.

The log likelihood $\log p_{Y|\Psi}(Y_T|\Psi)$ is given by $\sum_{t=1}^{T} \sum_{i=1}^{n} L_{i,t}$, up to a constant.

C.3.2 *Reweighting*   An optional reweighting step may be used to translate draws obtained from the Whittle-based HMC algorithm into draws from the exact Gaussian posterior density $p_{\Theta,\sigma|Y}(\Theta, \sigma|Y_T)$. The Whittle HMC algorithm yields draws $(\Theta^{(1)}, \sigma^{(1)}), \ldots, (\Theta^{(N)}, \sigma^{(N)})$ (after discarding a burn-in sample) from the Whittle posterior density $p_{\Theta,\sigma|Y}^W(\Theta, \sigma|Y_T)$. If desired, apply the following reweighting procedure to the Whittle draws:

1. For each Whittle draw $k = 1, 2, \ldots, N$, compute the relative likelihood weight

$$w_k = \frac{p_{Y|\Psi}(Y_T|\Psi(\Theta^{(k)}, \sigma^{(k)}))}{p_{Y|\Psi}^W(Y_T|\Psi(\Theta^{(k)}, \sigma^{(k)}))} \propto \frac{p_{\Theta,\sigma|Y}(\Theta^{(k)}, \sigma^{(k)}|Y_T)}{p_{\Theta,\sigma|Y}^W(\Theta^{(k)}, \sigma^{(k)}|Y_T)}.$$

2. Compute normalized weights $\tilde{w}_k = w_k / \sum_{b=1}^{N} w_b$, $k = 1, \dots, N$.

3. Draw N samples $(\tilde{\Theta}^{(1)}, \tilde{\sigma}^{(1)}), \dots, (\tilde{\Theta}^{(N)}, \tilde{\sigma}^{(N)})$ from the multinomial distribution with mass points $(\Theta^{(1)}, \sigma^{(1)}), \dots, (\Theta^{(N)}, \sigma^{(N)})$ and corresponding probabilities $\tilde{w}_1, \dots, \tilde{w}_N$.

Then $(\tilde{\Theta}^{(1)}, \tilde{\sigma}^{(1)}), \dots, (\tilde{\Theta}^{(N)}, \tilde{\sigma}^{(N)})$ constitute N draws from the exact posterior distribution. This reweighting procedure is a Sampling-Importance-Resampling procedure (Rubin (1988)) that uses the Whittle posterior as a proposal distribution. The reweighting step is fast, as it only needs to compute the exact likelihood—not the score—for N different parameter values, where N is typically orders of magnitude smaller than the required number of likelihood/score evaluations during the HMC algorithm.

C.4 *Hamiltonian Monte Carlo implementation*

I here describe my implementation of the posterior simulation algorithm. First, I outline my method for obtaining an initial value. Then I discuss the modifications I make to the Hoffman and Gelman (2014) algorithm. The calculations below require evaluation of the log prior density, its gradient, the log likelihood, and the score. Evaluation of the multivariate Gaussian log prior and its gradient is straightforward; this is also the case for many other choices of priors. Evaluation of the Whittle likelihood and its score is described in Appendix A.3.

C.4.1 *Initial value*   The HMC algorithm produces draws from a Markov Chain whose long-run distribution is the Whittle posterior of the SVMA parameters, regardless of the initial value used for the chain. However, using an initial value near the mode of the posterior distribution can significantly speed up the convergence to the long-run distribution. I approximate the posterior mode using the following computationally cheap procedure:

1. Compute the empirical ACF of the data.

2. Run q steps of the Innovations Algorithm to obtain an invertible SVMA representation that approximately fits the empirical ACF (Brockwell and Davis (1991, Prop. 11.4.2)).[8] Denote these invertible parameters by $(\hat{\Theta}, \hat{\sigma})$.

3. Let $\mathcal{C}$ denote the (finite) set of complex roots of the SVMA polynomial corresponding to $(\hat{\Theta}, \hat{\sigma})$, cf. Proposition 2.

4. For each root γ_j in $\mathcal{C}$ (each complex conjugate pair of roots is treated as one root):

(a) Let $(\check{\Theta}^{(j)}, \check{\sigma}^{(j)})$ denote the result of flipping root γ_j, i.e., of applying transformation (ii) in Proposition 2 to $(\hat{\Theta}, \hat{\sigma})$ with this root.

(b) Determine the orthogonal matrix $Q^{(j)}$ such that $\check{\Theta}^{(j)} \operatorname{diag}(\check{\sigma})^{(j)} Q^{(j)}$ is closest to the prior mean $E(\Theta \operatorname{diag}(\sigma))$ in Frobenius norm, cf. Footnote 5.

---

[8]In principle, the Innovations Algorithm could be run for more than q steps, but this tends to lead to numerical instability in my trials. The output of the first q steps is sufficiently accurate in my experience.

(c) Obtain parameters $(\tilde{\Theta}^{(j)}, \tilde{\sigma}^{(j)})$ such that $\check{\Theta}^{(j)} \operatorname{diag}(\check{\sigma}^{(j)}) Q^{(j)} = \tilde{\Theta}^{(j)} \operatorname{diag}(\tilde{\sigma}^{(j)})$, i.e., apply transformation (a) in Proposition 2. Calculate the corresponding value of the prior density $\pi(\tilde{\Theta}^{(j)}, \tilde{\sigma}^{(j)})$.

5. Let $\tilde{j} = \operatorname{argmax}_j \pi(\tilde{\Theta}^{(j)}, \tilde{\sigma}^{(j)})$.

6. If $\pi(\tilde{\Theta}^{(\tilde{j})}, \tilde{\sigma}^{(\tilde{j})}) \le \pi(\hat{\Theta}, \hat{\sigma})$, go to Step 7. Otherwise, set $(\hat{\Theta}, \hat{\sigma}) = (\tilde{\Theta}^{(\tilde{j})}, \tilde{\sigma}^{(\tilde{j})})$, remove γ_j (and its complex conjugate) from $\mathcal{C}$, and go back to Step 4.

7. Let the initial value for the HMC algorithm be the parameter vector of the form $((1-x)\check{\Theta} + xE(\Theta), (1-x)\check{\sigma} + xE(\sigma))$ that maximizes the posterior density, where x ranges over the grid $\{0, 0.01, \dots, 0.99, 1\}$, and $(E(\Theta), E(\sigma))$ is the prior mean of (Θ, σ).

Step 2 computes a set of invertible parameters that yields a high value of the likelihood. Steps 3–6 find a set of possibly noninvertible parameters that yields a high value of the prior density while being observationally equivalent with the parameters from Step 2 (I use a "greedy" search algorithm since it is computationally prohibitive to consider all combinations of root flips). Because Steps 2–6 lexicographically prioritize maximizing the likelihood over maximizing the prior, Step 7 allows the parameters to shrink toward the prior means.

C.4.2 *Modifications to NUTS algorithm*   I use the HMC variant NUTS from Hoffman and Gelman (2014), which automatically tunes the step size and trajectory length of HMC. See their paper for details on the NUTS algorithm. I downloaded the code from Hoffman's website.[9] I make two modifications to the basic NUTS algorithm, neither of which are essential, although they do tend to improve the mixing speed of the Markov chain in my trials: step size jittering and diagonal mass matrix adaptation. These modifications are also used in the NUTS-based statistics software Stan (Stan Development Team (2015)).

Each step I draw a new HMC step size from a uniform distribution over some interval (Neal (2011, Section 5.4.2.2)). The jittering is started after the stepsize has been tuned as described in Hoffman and Gelman (2014, Section 3.2). For the applications in this paper, the step size is chosen uniformly at random from the interval $[0.5\hat{\epsilon}, 1.5\hat{\epsilon}]$, where $\hat{\epsilon}$ is the tuned step size.

I allow for a diagonal "mass matrix," where the entries along the diagonal are estimates of the posterior standard deviations of the SVMA parameters (Neal (2011, Section 5.4.2.4)). I first run the NUTS algorithm for a number of steps with an identity mass matrix. Then I calculate the sample standard deviations of the parameter draws over a window of subsequent steps, after which I update the mass matrix accordingly.[10] I update the mass matrix twice more using windows of increasing length. Finally, I freeze the mass matrix for the remainder of the NUTS algorithm. In this paper, the mass matrix is estimated over steps 300–400, steps 401–600, and steps 601–1,000, and it is fixed after step 1,000.

---

[9]http://matthewdhoffman.com.

[10]The sample standard deviations are partially shrunk toward 1 before updating the mass matrix.

C.5 *Asymptotic analysis*

Here, I provide supplemental discussion and results concerning posterior consistency of the ACF or Wold parameters, as well as a more abstract consistency lemma.

C.5.1 *Posterior consistency: Autocovariance function*   I now show that the posterior consistency assumption for the reduced-form parameter Γ in Lemma 1 is satisfied in almost all stationary time series models for which Γ can be chosen to be the ACF, such as the SVMA model. The result below supposes that the posterior measure for the ACF is computed using the Whittle likelihood, but posterior consistency obtains under general model misspecification.[11]

To state the posterior consistency result, I first define the posterior measure. Let $\mathbb{T}_{n,q}$ be the space of ACFs for n-dimensional full-rank nondeterministic q-dependent processes, and let $p^W_{Y|\Gamma}(Y_T|\Gamma)$ denote the Whittle approximation to the likelihood of a stationary Gaussian process with ACF Γ (these objects are defined in detail in Section C.5.2). Let $\Pi_\Gamma(\cdot)$ be a prior measure on $\mathbb{T}_{n,q}$. The associated Whittle posterior measure for $\{\Gamma_0(k)\}_{0 \le k \le q}$ is

$$P^W_{\Gamma|Y}(\mathcal{A}|Y_T) = \frac{\int_{\mathcal{A}} p^W_{Y|\Gamma}(Y_T|\Gamma)\Pi_\Gamma(d\Gamma)}{\int_{\mathbb{T}_{n,q}} p^W_{Y|\Gamma}(Y_T|\Gamma)\Pi_\Gamma(d\Gamma)}, \tag{C.2}$$

for any measurable subset $\mathcal{A}$ of $\mathbb{T}_{n,q}$.

LEMMA C.1. *Let Assumption 3 hold. Assume that $\{\Gamma_0(k)\}_{0 \le k \le q}$ is in the support of $\Pi_\Gamma(\cdot)$. Then the Whittle posterior for $\{\Gamma_0(k)\}_{0 \le k \le q}$ is consistent, i.e., for any neighborhood $\mathcal{U}$ of $\{\Gamma_0(k)\}_{0 \le k \le q}$ in $\mathbb{T}_{n,q}$, we have $P^W_{\Gamma|Y}(\mathcal{U}|Y_T) \overset{p}{\to} 1$ as $T \to \infty$, under the true probability measure of the data.*

The SVMA model with Gaussian shocks is an example of a model with a stationary Gaussian and q-dependent likelihood. Hence, when applied to the SVMA model, Lemma C.1 states that if the prior measure on the SVMA parameters induces a prior measure on Γ that has the true ACF $\{\Gamma_0(k)\}_{0 \le k \le q}$ in its support, then the model-implied Whittle posterior for Γ pins down the true ACF in large samples. Lemma C.1 places no restrictions on the prior $\Pi_\Gamma(\cdot)$ on the ACF, except that the true ACF Γ_0 lies in its support.

Posterior consistency for the true ACF (up to lag q) holds even for time series that are not Gaussian or q-dependent. This is true even though the measure $P^W_{\Gamma|Y}(\mathcal{A}|Y_T)$ is computed using the Whittle likelihood and, therefore, exploits the working assumption that

---

[11]In the case of i.i.d. data, posterior consistency in misspecified models has been investigated in detail; see Kleijn and van der Vaart (2012), Ramamoorthi, Sriram, and Martin (2015), and references therein. Shalizi (2009) proved posterior consistency under misspecification with dependent data, placing high-level assumptions on the prior and likelihood. Müller (2013) discussed decision theoretic properties of Bayes estimators when the model is misspecified. Tamaki (2008) derived an asymptotic approximation to the Whittle posterior under correct specification.

the data is Gaussian and q-dependent. The only restrictions placed on the true distribution of the data are the stationarity and weak dependence conditions in Assumption 3. Lemma C.1 is silent about posterior inference on autocovariances at lags higher than q.

C.5.2 *Whittle likelihood for a q-dependent process*   I now define the Whittle ACF likelihood for q-dependent processes used in Section C.5.1. Consider the spectral density for a q-dependent process parametrized in terms of its ACF:

$$f(\omega; \Gamma) = \frac{1}{2\pi}\left(\Gamma(0) + \sum_{k=1}^{q}\{e^{-\iota k\omega}\Gamma(k) + e^{\iota k\omega}\Gamma(k)'\}\right), \quad \Gamma \in \mathbb{T}_{n,q},$$

$$\mathbb{T}_{n,q} = \Big\{\{\Gamma(k)\}_{0\leq k\leq q} : \Gamma(\cdot) \in \mathbb{R}^{n\times n}, \Gamma(0) = \Gamma(0)',$$

$$\inf_{\omega\in[0,\pi)}\det\left(\Gamma(0) + \sum_{k=1}^{q}\{e^{-\iota\omega k}\Gamma(k) + e^{\iota\omega k}\Gamma(k)'\}\right) > 0\Big\}.$$

Let $\hat{\Gamma}(k) = T^{-1}\sum_{t=1}^{T-k} y_{t+k}y_t'$, $k = 0, 1, \ldots, T-1$, be the k-th sample autocovariance, and set $\hat{\Gamma}(k) = \hat{\Gamma}(-k)'$ for $k = -1, -2, \ldots, 1-T$. Define the periodogram

$$\hat{I}(\omega) = \frac{1}{2\pi T}\left(\sum_{t=1}^{T}e^{-\iota t\omega}y_t\right)\left(\sum_{t=1}^{T}e^{\iota t\omega}y_t'\right) = \frac{1}{2\pi}\sum_{k=-(T-1)}^{T-1}e^{-\iota k\omega}\hat{\Gamma}(k), \quad \omega \in [-\pi, \pi].$$

The Whittle ACF log likelihood is (up to a constant) given by

$$\log p^{W}_{Y|\Gamma}(Y_T|\Gamma) = -\frac{T}{4\pi}\int_{-\pi}^{\pi}\log\det(f(\omega; \Gamma))\, d\omega - \frac{T}{4\pi}\int_{-\pi}^{\pi}\text{tr}\{f(\omega; \Gamma)^{-1}\hat{I}(\omega)\}\, d\omega.$$

As in Appendix A.3, it is common to use a discretized Whittle log likelihood that replaces integrals with corresponding discretized sums. The proof of Lemma C.1 shows that posterior consistency also holds when the discretized Whittle likelihood is used.

C.5.3 *Posterior consistency: Wold parameters*   The proof of Lemma C.1 relies on the following posterior consistency result for the Wold IRFs and prediction covariance matrix. The result in this subsection concerns (invertible) reduced-form IRFs, not (possibly noninvertible) structural IRFs.

Fix a finite $q \in \mathbb{N}$, and let $\beta_0(L) = I_n + \sum_{\ell=1}^{q}\beta_{0,\ell}L^{\ell}$ and Σ_0 denote the MA lag polynomial and prediction covariance matrix, respectively, in the Wold decomposition (Hannan (1970, Thm. 2″, p. 158)) of a q-dependent stationary n-dimensional process with ACF (out to lag q) given by $\{\Gamma_0(k)\}_{0\leq k\leq q} \in \mathbb{T}_{n,q}$. That is, $\beta_0 = (\beta_{0,1}, \ldots, \beta_{0,q}) \in \mathbb{B}_{n,q}$ and $\Sigma_0 \in \mathbb{S}_n$ are the unique parameters such that $\Gamma_0(k) = \sum_{\ell=0}^{q-k}\beta_{0,\ell+k}\Sigma_0\beta_{0,\ell}'$ for $0 \leq k \leq q$, defining $\beta_{0,0} = I_n$. Here $\mathbb{S}_n$ denotes the space of symmetric positive definite $n \times n$ matrices, while

$$\mathbb{B}_{n,q} = \{\beta = (\beta_1, \ldots, \beta_q) \in \mathbb{R}^{n\times nq} : \det(\Phi(z; \beta)) \neq 0\, \forall|z| \leq 1\},$$

$$\Phi(z;\beta) = I_n + \sum_{\ell=1}^{q} \beta_\ell z^\ell, \quad z \in \mathbb{C}.$$

Define the MA spectral density parametrized in terms of (β, Σ):

$$\tilde{f}(\omega;\beta,\Sigma) = \frac{1}{2\pi}\Phi(e^{-\iota\omega};\beta)\Sigma\Phi(e^{\iota\omega};\beta)', \quad \omega \in [-\pi, \pi], (\beta, \Sigma) \in \mathbb{B}_{n,q} \times \mathbb{S}_n.$$

Using the notation in Section C.5.2 for the periodogram $\hat{I}(\omega)$, the Whittle MA log likelihood is (up to a constant) given by

$$\log p_{Y|\beta,\Sigma}^{W}(Y_T|\beta,\Sigma) = -\frac{T}{4\pi}\int_{-\pi}^{\pi}\log\det(\tilde{f}(\omega;\beta,\Sigma))\,d\omega$$

$$-\frac{T}{4\pi}\int_{-\pi}^{\pi}\text{tr}\{\tilde{f}(\omega;\beta,\Sigma)^{-1}\hat{I}(\omega)\}\,d\omega.$$

The result below also holds for the discretized Whittle likelihood (cf. Section C.5.2).

I now state the posterior consistency result for (β_0, Σ_0). Let $\Pi_{\beta,\Sigma}(\cdot)$ be a prior measure for (β_0, Σ_0) on $\mathbb{B}_{n,q} \times \mathbb{S}_n$. Define the Whittle posterior measure

$$P_{\beta,\Sigma|Y}^{W}(\mathcal{A}|Y_T) = \frac{\int_{\mathcal{A}} p_{Y|\beta,\Sigma}^{W}(Y_T|\beta,\Sigma)\Pi_{\beta,\Sigma}(d(\beta,\Sigma))}{\int_{\mathbb{B}_{n,q}\times\mathbb{S}_n} p_{Y|\beta,\Sigma}^{W}(Y_T|\beta,\Sigma)\Pi_{\beta,\Sigma}(d(\beta,\Sigma))}$$

for any measurable set $\mathcal{A} \subset \mathbb{B}_{n,q} \times \mathbb{S}_n$. Note that the lemma below does not require the true data distribution to be Gaussian or q-dependent.

LEMMA C.2. *Let Assumption* 3 *hold. Assume that the pseudo-true parameters* $(\beta_0, \Sigma_0) \in \mathbb{B}_{n,q} \times \mathbb{S}_n$ *are in the support of the prior* $\Pi_{\beta,\Sigma}(\cdot)$. *Then, for any neighborhood* $\tilde{\mathcal{U}}$ *of* (β_0, Σ_0) *in* $\mathbb{B}_{n,q} \times \mathbb{S}_n$, *we have* $P_{\beta,\Sigma|Y}^{W}(\tilde{\mathcal{U}}|Y_T) \xrightarrow{P} 1$ *under the true probability measure of the data.*

C.5.4 *Posterior consistency: General lemma*   Following Ghosh and Ramamoorthi (2003, Theorem 1.3.4), I give general sufficient conditions for assumption (ii) of Lemma 1. These are used in the proof of Lemma C.2.

Let $\Pi_\Gamma(\cdot)$ denote the marginal prior measure for parameter Γ, with Euclidean parameter space Ξ_Γ. Let $p_{Y|\Gamma}(Y_T|\Gamma)$ denote the (possibly misspecified) likelihood function. The posterior measure is given by

$$P_{\Gamma|Y}(\mathcal{A}|Y_T) = \frac{\int_{\mathcal{A}} p_{Y|\Gamma}(Y_T|\Gamma)\Pi_\Gamma(d\Gamma)}{\int_{\Xi_\Gamma} p_{Y|\Gamma}(Y_T|\Gamma)\Pi_\Gamma(d\Gamma)}$$

for measurable sets $\mathcal{A} \subset \Xi_\Gamma$.

LEMMA C.3. *Define the normalized log likelihood ratio $\hat{\phi}(\Gamma) = T^{-1} \log \frac{p_{Y|\Gamma}(Y_T|\Gamma)}{p_{Y|\Gamma}(Y_T|\Gamma_0)}$ for all $\Gamma \in \Xi_\Gamma$. Assume there exist a function $\phi : \Xi_\Gamma \to \mathbb{R}$, a compact neighborhood $\mathcal{K}$ of Γ_0 in Ξ_Γ, and a scalar $\zeta < 0$ such that the following conditions hold.*

(i) $\sup_{\Gamma \in \mathcal{K}} |\hat{\phi}(\Gamma) - \phi(\Gamma)| \overset{p}{\to} 0$.

(ii) $\phi(\Gamma)$ *is continuous at all $\Gamma \in \mathcal{K}$.*

(iii) $\phi(\Gamma) < 0$ *for all $\Gamma \neq \Gamma_0$.*

(iv) $\sup_{\Gamma \in \mathcal{K}^c} \hat{\phi}(\Gamma) < \zeta$ *w.p.a.* 1.

(v) Γ_0 *is in the support of $\Pi_\Gamma(\cdot)$.*

Then for any neighborhood $\mathcal{U}$ of Γ_0 in Ξ_Γ, $P_{\Gamma|Y}(\mathcal{U}|Y_T) \overset{p}{\to} 1$.

C.6 *Supplemental proofs*

C.6.1 *Proof of Proposition 2*   As in Lippi and Reichlin (1994, p. 311), define the rational matrix function

$$R(\gamma, z) = \begin{pmatrix} \dfrac{z - \gamma}{1 - \bar{\gamma}z} & 0 \\ 0 & I_{n-1} \end{pmatrix}, \quad \gamma, z \in \mathbb{C}.$$

Transformation (ii) corresponds to the transformation $\check{\Psi}(z) = \Psi(z)QR(\gamma_k, z)^{-1}$ if γ_k is real. If γ_k is not real, the transformation corresponds to $\check{\Psi}(z) = \tilde{\Psi}(z)\check{Q}$, where $\tilde{\Psi}(z) = \Psi(z)QR(\gamma_k, z)^{-1}\tilde{Q}R(\overline{\gamma_k}, z)^{-1}$ and $\check{Q} = \tilde{\Psi}(0)^{-1}J$ is a unitary matrix. I proceed in three steps.

Step 1   Consider the first claim of the proposition. Let $f(\omega; \Gamma) = (2\pi)^{-1} \sum_{k=-q}^{q} \Gamma(k) \times e^{-\imath k \omega}$, $\omega \in [-\pi, \pi]$, denote the spectral density matrix function associated with the ACF $\Gamma(\cdot)$. Since $\Psi(z) = \Theta(z)\operatorname{diag}(\sigma)$ with $(\Theta, \sigma) \in \mathcal{S}(\Gamma)$, we must have $\Psi(e^{-\imath\omega})\Psi(e^{-\imath\omega})^* = 2\pi f(\omega; \Gamma)$ for all ω by the usual formula for the spectral density of a vector MA process (Brockwell and Davis (1991, Example 11.8.1)). Because $R(\gamma, e^{-\imath\omega})R(\gamma, e^{-\imath\omega})^* = I_n$ for any (γ, ω), it is easy to verify that $\check{\Psi}(z)$—constructed by applying transformation (i) or transformation (ii) to $\Psi(z)$—also satisfies $\check{\Psi}(e^{-\imath\omega})\check{\Psi}(e^{-\imath\omega})^* = 2\pi f(\omega; \Gamma)$. Hence, $\check{\Psi}(z) = \sum_{\ell=0}^{q} \check{\Psi}_\ell z^\ell$ is a matrix MA polynomial satisfying $\sum_{\ell=0}^{q-k} \check{\Psi}_{\ell+k}\check{\Psi}_\ell^* = \Gamma(k)$ for all $k = 0, 1, \ldots, q$. In Step 2 below I show that $\check{\Psi}(z)$ is a matrix polynomial with real coefficients. By construction of $\check{\Theta}(z) = \sum_{\ell=0}^{q} \check{\Theta}_\ell z^\ell$ and $\check{\sigma}$, we then have $\sum_{\ell=0}^{q-k} \check{\Theta}_{\ell+k}\operatorname{diag}(\check{\sigma})^2\check{\Theta}_\ell' = \Gamma(k)$ for all $k = 0, 1, \ldots, q$, so $(\check{\Theta}, \check{\sigma}) \in \mathcal{S}(\Gamma)$, as claimed.

Step 2   I now show that transformation (ii) yields a real matrix polynomial $\check{\Psi}(z)$. This fact was asserted by Lippi and Reichlin (1994, pp. 317–318). I am grateful to Professor Marco Lippi for providing me with the proof arguments for Step 2; all errors are my own.

$\check{\Psi}(z)$ is clearly real if the flipped root γ_k is real (since η and Q can be chosen to be real in this case), so consider the case where we flip a pair of complex conjugate roots γ_k and $\overline{\gamma_k}$. Recall that in this case, $\check{\Psi}(z) = \tilde{\Psi}(z)\check{Q}$, where $\tilde{\Psi}(z) = \Psi(z)QR(\gamma_k, z)^{-1}\tilde{Q}R(\overline{\gamma_k}, z)^{-1}$

and $\check{Q}$ is unitary. It follows from the same arguments as in Step 1 that the complex-valued matrix polynomial $\tilde{\Psi}(z) = \sum_{\ell=0}^{q} \tilde{\Psi}_{\ell} z^{\ell}$ satisfies $\sum_{\ell=0}^{q-k} \tilde{\Psi}_{\ell+k} \tilde{\Psi}_{\ell}^* = \Gamma(k)$ for all $k = 0, 1, \ldots, q$.

Let $\overline{\tilde{\Psi}}(z) = \sum_{\ell=0}^{q} \overline{\tilde{\Psi}_{\ell}} z^{\ell}$ denote the matrix polynomial obtained by conjugating the coefficients of the polynomial $\tilde{\Psi}(z)$. By construction, the roots of $\det(\overline{\tilde{\Psi}}(z))$ are real or appear as complex conjugate pairs, so $\det(\overline{\tilde{\Psi}}(z))$ has the same roots as $\det(\tilde{\Psi}(z))$. Furthermore, for $k = 0, 1, \ldots, q$,

$$\sum_{\ell=0}^{q-k} \overline{\tilde{\Psi}}_{\ell+k} \overline{\tilde{\Psi}}_{\ell}^{\,*} = \overline{\Gamma(k)} = \Gamma(k) = \sum_{\ell=0}^{q-k} \tilde{\Psi}_{\ell+k} \tilde{\Psi}_{\ell}^*.$$

By Theorem 3(b) of Lippi and Reichlin (1994), there exists a unitary $n \times n$ matrix $\tilde{\tilde{Q}}$ such that $\overline{\tilde{\Psi}}(z) = \tilde{\Psi}(z)\tilde{\tilde{Q}}$ for $z \in \mathbb{R}$. The matrix polynomial $\tilde{\Psi}(z)\tilde{\Psi}(0)^{-1}$ then has real coefficients:[12] For all $z \in \mathbb{R}$,

$$\tilde{\Psi}(z)\tilde{\Psi}(0)^{-1} = \big(\tilde{\Psi}(z)\tilde{\tilde{Q}}\big)\big(\tilde{\Psi}(0)\tilde{\tilde{Q}}\big)^{-1} = \overline{\tilde{\Psi}}(z)\overline{\tilde{\Psi}}(0)^{-1} = \overline{\tilde{\Psi}(z)\tilde{\Psi}(0)^{-1}}.$$

Consequently, with the real matrix J defined as in the proposition, $\check{\Psi}(z) = \tilde{\Psi}(z)\tilde{\Psi}(0)^{-1}J$ is a matrix polynomial with real coefficients. Note that, since $\tilde{\tilde{Q}}$ is unitary, the matrix

$$\tilde{\Psi}(0)\tilde{\Psi}(0)^* = \big(\overline{\tilde{\Psi}}(0)\tilde{\tilde{Q}}\big)\big(\overline{\tilde{\Psi}}(0)\tilde{\tilde{Q}}\big)^* = \overline{\tilde{\Psi}(0)\tilde{\Psi}(0)^*}$$

is real, symmetric, and positive definite, so J is well-defined.

Step 3   Finally, I prove the second claim of the proposition. Suppose we have a fixed element $(\check{\Theta}, \check{\sigma})$ of the identified set that we want to end up with after transforming the initial element (Θ, σ) appropriately. Define $\check{\Psi}(z) = \check{\Theta}(z)\,\mathrm{diag}(\check{\sigma})$. Since $(\Theta, \sigma), (\check{\Theta}, \check{\sigma}) \in \mathcal{S}(\Gamma)$, the two sets of SVMA parameters correspond to the same spectral density, i.e., $\Psi(e^{-\iota\omega})\Psi(e^{-\iota\omega})^* = \check{\Psi}(e^{-\iota\omega})\check{\Psi}(e^{-\iota\omega})^*$ for all $\omega \in [-\pi, \pi]$. As in the proof of Theorem 2 in Lippi and Reichlin (1994), we can apply transformation (ii) finitely many (say, b) times to $\Psi(z)$, flipping all the roots that are inside the unit circle, thus ending up with a polynomial

$$B(z) = \Psi(z)Q_1 R(\gamma_{k_1}, z)^{-1} \cdots Q_b R(\gamma_{k_b}, z)^{-1} Q_{b+1}$$

for which all roots of $\det(B(z))$ lie on or outside the unit circle. Likewise, denote the (finitely many) roots of $\det(\check{\Psi}(z))$ by $\check{\gamma}_k$, $k = 1, 2, \ldots$, and apply to $\check{\Psi}(z)$ a finite sequence of transformation (ii) to arrive at a polynomial

$$\check{B}(z) = \check{\Psi}(z)\check{Q}_1 R(\check{\gamma}_{\check{k}_1}, z)^{-1} \cdots \check{Q}_{\check{b}} R(\check{\gamma}_{\check{k}_{\check{b}}}, z)^{-1} \check{Q}_{\check{b}+1}$$

for which all roots of $\det(\check{B}(z))$ lie on or outside the unit circle. Since $\det(B(z))$ and $\det(\check{B}(z))$ have all roots on or outside the unit circle, and we have $B(e^{-\iota\omega})B(e^{-\iota\omega})^* =$

---

[12] $\tilde{\Psi}(0)$ is nonsingular because $\det(\Psi(0)) \neq 0$.

$\check{B}(e^{-\iota\omega})\check{B}(e^{-\iota\omega})^* = 2\pi f(\omega;\Gamma)$ for all ω, there must exist an orthogonal matrix Q such that $\check{B}(z) = B(z)Q$ (Lippi and Reichlin (1994, p. 313), Hannan (1970, p. 69)). Thus,

$$\check{\Psi}(z) = \Psi(z)Q_1 R(\gamma_{k_1}, z)^{-1} \cdots Q_b R(\gamma_{k_b}, z)^{-1} Q_{b+1} Q \check{Q}_{b+1}^* R(\check{\gamma}_{\check{k}_{\check{b}}}, z)\check{Q}_{\check{b}}^* \cdots R(\check{\gamma}_{\check{k}_1}, z)\check{Q}_1^*,$$

and

$$\det\!\big(\check{\Psi}(z)\big) = \det\!\big(\Psi(z)\big)\frac{(z - \check{\gamma}_{\check{k}_1})\cdots(z - \check{\gamma}_{\check{k}_{\check{b}}})(1 - \overline{\gamma_{k_1}}z)\cdots(1 - \overline{\gamma_{k_b}}z)}{(z - \gamma_{k_1})\cdots(z - \gamma_{k_b})(1 - \overline{\check{\gamma}_{\check{k}_1}}z)\cdots(1 - \overline{\check{\gamma}_{\check{k}_{\check{b}}}}z)},$$

so any root of $\det(\check{\Psi}(z))$ must either equal γ_k or it must equal $1/\overline{\gamma_k}$, where γ_k is some root of $\det(\Psi(z))$. It follows that we can apply a finite sequence of transformation (ii) (i.e., an appropriate sequence of root flips) to $\Psi(z)$ to obtain a real matrix polynomial $\tilde{\tilde{\Psi}}(z)$ satisfying $\det(\tilde{\tilde{\Psi}}(z)) = \det(\check{\Psi}(z))$ for all $z \in \mathbb{C}$. Theorem 3(b) in Lippi and Reichlin (1994) then implies that $\check{\Psi}(z)$ can be obtained from $\tilde{\tilde{\Psi}}(z)$ through transformation (i) (i.e., an orthogonal rotation, which clearly must be real). Finally, obtain $(\check{\Theta}, \check{\sigma})$ from $\check{\Psi}(z)$ by transformation (a).

C.6.2 *Proof of Lemma 2*   Suppressing the arguments (Ψ), let $L_k = \log\det(f_k) + \tilde{y}_k^* \times f_k^{-1}\tilde{y}_k$. Then

$$\frac{\partial L_k}{\partial(f_k')} = f_k^{-1} - f_k^{-1}\tilde{y}_k\tilde{y}_k^* f_k^{-1} = C_k.$$

Writing $f_k' = \overline{\tilde{\Psi}_k}\tilde{\Psi}_k'$, we have

$$\frac{\partial\,\mathrm{vec}(f_k')}{\partial\,\mathrm{vec}(\Psi_\ell)'} = (\tilde{\Psi}_k \otimes I_n)e^{\iota\omega_k\ell} + (I_n \otimes \overline{\tilde{\Psi}_k})K_n' e^{-\iota\omega_k\ell},$$

where K_n is the $n^2 \times n^2$ commutation matrix such that $\mathrm{vec}(B') = K_n\,\mathrm{vec}(B)$ for any $n \times n$ matrix B (Magnus and Neudecker (2007, Ch. 3.7)). Using $\mathrm{vec}(ABC) = (C' \otimes A)\,\mathrm{vec}(B)$,

$$\frac{\partial L_k}{\partial\,\mathrm{vec}(\Psi_\ell)'} = \frac{\partial L_k}{\partial\,\mathrm{vec}(f_k')}\frac{\partial\,\mathrm{vec}(f_k')}{\partial\,\mathrm{vec}(\Psi_\ell)'} = \mathrm{vec}\big(C_k\tilde{\Psi}_k e^{\iota\omega_k\ell} + \overline{C_k^*\tilde{\Psi}_k e^{\iota\omega_k\ell}}\big)'.$$

Since $C_k^* = C_k$, we get $\partial L_k/\partial\Psi_\ell = 2\,\mathrm{Re}(C_k\tilde{\Psi}_k e^{\iota\omega_k\ell})$, so

$$\frac{\partial\log p_{Y|\Psi}^W(Y_T|\Psi)}{\partial\Psi_\ell} = -\frac{1}{2}\sum_{k=0}^{T-1}\frac{\partial L_k}{\partial\Psi_\ell}$$

$$= -\sum_{k=0}^{T-1}\mathrm{Re}\!\left(C_k\sum_{\tilde{\ell}=1}^{q+1} e^{-\iota\omega_k(\tilde{\ell}-1)}\Psi_{\tilde{\ell}-1} e^{\iota\omega_k\ell}\right)$$

$$= -\sum_{\tilde{\ell}=0}^{q}\mathrm{Re}\!\left(\sum_{k=0}^{T-1} C_k e^{-\iota\omega_k(\tilde{\ell}-\ell)}\right)\Psi_{\tilde{\ell}}.$$

Finally, $\sum_{k=0}^{T-1} C_k e^{-\iota\omega_k(\tilde{\ell}-\ell)} = \sum_{k=0}^{T-1} C_k e^{-\iota\omega_{\tilde{\ell}-\ell}k} = \tilde{C}_{\tilde{\ell}-\ell}$ for $\tilde{\ell} \geq \ell$, and $\sum_{k=0}^{T-1} C_k e^{-\iota\omega_k(\tilde{\ell}-\ell)} = \sum_{k=0}^{T-1} C_k e^{-\iota\omega_k(T+\tilde{\ell}-\ell)} = \sum_{k=0}^{T-1} C_k e^{-\iota\omega_{T+\tilde{\ell}-\ell}k} = \tilde{C}_{\tilde{\ell}-\ell}$ for $\tilde{\ell} < \ell$.

C.6.3 *Proof of Lemma* C.1

The proof exploits the one-to-one mapping between the ACF Γ_0 and the Wold parameters (β_0, Σ_0) defined in Section C.5.3, which allows me to use Lemma C.2 to infer posterior consistency for Γ_0 under the Whittle likelihood.

Let $M : \mathbb{T}_{n,q} \to \mathbb{B}_{n,q} \times \mathbb{S}_n$ denote the function that maps a q-dependent ACF $\Gamma(\cdot)$ into its Wold representation $(\beta(\Gamma), \Sigma(\Gamma))$ (Hannan (1970, Theorem 2″, p. 158)). By construction, the map $M(\cdot)$ is continuous (and measurable). The inverse map $M^{-1}(\cdot)$ is given by $\Gamma(k) = \sum_{\ell=0}^{q-k} \beta_{\ell+k} \Sigma \beta_\ell'$ (with $\beta_0 = I_n$) and so also continuous. The prior $\Pi_\Gamma(\cdot)$ for the ACF Γ induces a particular prior measure for the Wold parameters (β, Σ) on $\mathbb{B}_{n,q} \times \mathbb{S}_n$ given by $\Pi_{\beta,\Sigma}(\mathcal{A}) = \Pi_\Gamma(M^{-1}(\mathcal{A}))$ for any measurable set $\mathcal{A}$. Let $P^W_{\beta,\Sigma|Y}(\cdot|Y_T)$ be the posterior measure for (β, Σ) computed using the induced prior $\Pi_{\beta,\Sigma}(\cdot)$ and the Whittle MA likelihood $p^W_{Y|\beta,\Sigma}(Y_T|\beta, \Sigma)$, cf. Section C.5.3.

I first show that the induced posterior for (β_0, Σ_0) is consistent. Let $\tilde{\mathcal{U}}$ be any neighborhood of $(\beta_0, \Sigma_0) = M(\{\Gamma_0(k)\}_{0\leq k\leq q})$ in $\mathbb{B}_{n,q} \times \mathbb{S}_n$. Since $M(\cdot)$ is continuous, $M^{-1}(\tilde{\mathcal{U}})$ is a neighborhood of $\{\Gamma_0(k)\}_{0\leq k\leq q}$ in $\mathbb{T}_{n,q}$. Hence, since $\{\Gamma_0(k)\}_{0\leq k\leq q}$ is in the support of $\Pi_\Gamma(\cdot)$, (β_0, Σ_0) is in the support of $\Pi_{\beta,\Sigma}(\cdot)$: $\Pi_{\beta,\Sigma}(\tilde{\mathcal{U}}) = \Pi_\Gamma(M^{-1}(\tilde{\mathcal{U}})) > 0$. Due to Assumption 3 and the fact that (β_0, Σ_0) is in the support of $\Pi_{\beta,\Sigma}(\cdot)$, Lemma C.2 implies that $P_{\beta,\Sigma|Y}(\tilde{\mathcal{U}}|Y_T) \xrightarrow{p} 1$ for any neighborhood $\tilde{\mathcal{U}}$ of (β_0, Σ_0) in $\mathbb{B}_{n,q} \times \mathbb{S}_n$.

I now prove posterior consistency for Γ_0. Since $\tilde{f}(\omega; M(\Gamma)) = f(\omega; \Gamma)$ for all $\omega \in [-\pi, \pi]$ and $\Gamma \in \mathbb{T}_{n,q}$, we have $p^W_{Y|\beta,\Sigma}(Y_T|M(\Gamma)) = p^W_{Y|\Gamma}(Y_T|\Gamma)$ for all $\Gamma \in \mathbb{T}_{n,q}$. Consequently, $P^W_{\Gamma|Y}(\mathcal{A}|Y_T) = P^W_{\beta,\Sigma|Y}(M(\mathcal{A})|Y_T)$ for all measurable sets $\mathcal{A}$. Let $\mathcal{U}$ be an arbitrary neighborhood of $\{\Gamma_0(k)\}_{0\leq k\leq q}$ in $\mathbb{T}_{n,q}$. Since $M^{-1}(\cdot)$ is continuous at (β_0, Σ_0), the set $\tilde{\mathcal{U}} = M(\mathcal{U})$ is a neighborhood of (β_0, Σ_0) in $\mathbb{B}_{n,q} \times \mathbb{S}_n$. It follows that

$$P^W_{\Gamma|Y}(\mathcal{U}|Y_T) = P^W_{\beta,\Sigma|Y}(\tilde{\mathcal{U}}|Y_T) \xrightarrow{p} 1.$$

The proof of Lemma C.2 shows that the discretized Whittle posterior is also consistent.

C.6.4 *Proof of Lemma* C.2

The proof closely follows the steps in Dunsmuir and Hannan (1976, Section 3) for proving consistency of the Whittle maximum likelihood estimator in a reduced-form identified VARMA model. Note that the only properties of the data generating process used in Dunsmuir and Hannan (1976, Section 3) are covariance stationarity and ergodicity for second moments, as in Assumption 3. Dunsmuir and Hannan also need $T^{-1} y_t y'_{t+T-k} \xrightarrow{p} 0$ for fixed t and k, which follows from Markov's inequality under covariance stationarity. Where Dunsmuir and Hannan (1976) appeal to almost sure convergence, I substitute convergence in probability.

Define the normalized log likelihood ratio

$$\hat{\phi}(\beta, \Sigma) = T^{-1} \log \frac{p^W_{\beta,\Sigma}(Y_T|\beta, \Sigma)}{p^W_{\beta,\Sigma}(Y_T|\beta_0, \Sigma_0)}.$$

By the Kolmogorov–Szegö formula, for any $(\beta, \Sigma) \in \mathbb{B}_{n,q} \times \mathbb{S}_n$,

$$\frac{1}{2\pi} \int_{-\pi}^{\pi} \log \det(\tilde{f}(\omega; \beta, \Sigma)) \, d\omega = \log \det(\Sigma) - n \log(2\pi). \tag{C.3}$$

Hence,

$$\hat{\phi}(\beta, \Sigma) = \frac{1}{2} \log \det(\Sigma_0 \Sigma^{-1}) + \frac{1}{4\pi} \int_{-\pi}^{\pi} \operatorname{tr}\{[\tilde{f}(\omega; \beta_0, \Sigma_0)^{-1} - \tilde{f}(\omega; \beta, \Sigma)^{-1}]\hat{I}(\omega)\} d\omega. \tag{C.4}$$

Define also the function

$$\phi(\beta, \Sigma) = \frac{1}{2} \log \det(\Sigma_0 \Sigma^{-1}) + \frac{1}{4\pi} \int_{-\pi}^{\pi} \operatorname{tr}\{I_n - \tilde{f}(\omega; \beta, \Sigma)^{-1} \tilde{f}(\omega; \beta_0, \Sigma_0)\} d\omega.$$

$\phi(\beta, \Sigma)$ is continuous. By the argument in Dunsmuir and Hannan (1976, p. 342) (see also Brockwell and Davis (1991, Prop. 10.8.1), for the univariate case), we have $\phi(\beta, \Sigma) \leq \phi(\beta_0, \Sigma_0) = 0$ for all $(\beta, \Sigma) \in \mathbb{B}_{n,q} \times \mathbb{S}_n$, with equality if and only if $(\beta, \Sigma) = (\beta_0, \Sigma_0)$.

The remainder of the proof verifies the conditions of Lemma C.3 in several steps.

Step 1   I first show that there exists a compact neighborhood $\mathcal{K}$ of (β_0, Σ_0) in $\mathbb{B}_{n,q} \times \mathbb{S}_n$ on which $\phi(\beta, \Sigma)$ is continuous and such that

$$\sup_{(\beta, \Sigma) \in \mathcal{K}} |\hat{\phi}(\beta, \Sigma) - \phi(\beta, \Sigma)| = o_p(1). \tag{C.5}$$

By definition of the Wold decomposition of a time series with a non-singular spectral density, all the roots of $z \mapsto \det(\Phi(\beta_0; z))$ lie strictly outside the unit circle. $\tilde{f}(\omega; \beta, \Sigma)^{-1} = \Phi(\beta; e^{\iota\omega})^{-1\prime} \Sigma^{-1} \Phi(\beta; e^{-\iota\omega})^{-1}$ is therefore uniformly continuous in (ω, β, Σ) for all $\omega \in [-\pi, \pi]$ and (β, Σ) in a small compact neighborhood of (β_0, Σ_0). Denoting this neighborhood by $\mathcal{K}$, the discussion around Lemma 1 in Dunsmuir and Hannan (1976, p. 350) implies (C.5).

Step 2   For any $(\beta, \Sigma) \in \mathbb{B}_{n,q} \times \mathbb{S}_n$ and $z \in \mathbb{C}$, define the adjoint of $\Phi(\beta; z)$ as

$$\Phi_{\mathrm{adj}}(\beta; z) = \Phi(\beta; z)^{-1} \det(\Phi(\beta; z)),$$

so $\tilde{f}(\omega; \beta, \Sigma) = |\det(\Phi(\beta; e^{-\iota\omega}))|^2 \Phi_{\mathrm{adj}}(\beta; e^{-\iota\omega})^{-1} \Sigma \Phi_{\mathrm{adj}}(\beta; e^{-\iota\omega})^{-1*}$. The elements of $\Phi_{\mathrm{adj}}(\beta; z)$ are polynomials in z, each polynomial of order at most $\kappa \leq q(n-1)$ (Dunsmuir and Hannan (1976, p. 354)). Write the matrix polynomial as $\Phi_{\mathrm{adj}}(\beta; z) = I_n + \sum_{\ell=1}^{\kappa} \beta_{\mathrm{adj},\ell} z^\ell$, and define $\tilde{\Phi}_{\mathrm{adj}}(\beta) = (\sum_{\ell=1}^{\kappa} \|\beta_{\mathrm{adj},\ell}\|^2)^{1/2}$.

For any $c_1, c_2, c_3 > 0$, define the set

$$\tilde{\mathcal{K}}(c_1, c_2, c_3) = \{(\beta, \Sigma) \in \mathbb{B}_{n,q} \times \mathbb{S}_n : \lambda_{\min}(\Sigma) \geq c_1, \|\Sigma\| \leq c_2, \tilde{\Phi}_{\mathrm{adj}}(\beta) \leq c_3\},$$

where $\lambda_{\min}(\Sigma)$ is the smallest eigenvalue of Σ.

Then the proof of Theorem 4(i) in Dunsmuir and Hannan (1976, pp. 354–355) (see also the beginning of the proof of their Theorem 3, pp. 352–353) shows that there exist constants $c_1, c_2, c_3 > 0$ and $\zeta_1 < 0$ such that

$$\sup_{(\beta, \Sigma) \notin \tilde{K}(c_1, c_2, c_3)} \hat{\phi}(\beta, \Sigma) \leq \zeta_1 + o_p(1).$$

Step 3   Now define, for $\delta \geq 0$,

$$\tilde{f}_\delta(\omega; \beta, \Sigma) = (|\det(\Phi(\beta; e^{-\iota\omega}))|^2 + \delta)\Phi_{\mathrm{adj}}(\beta; e^{-\iota\omega})^{-1}\Sigma\Phi_{\mathrm{adj}}(\beta; e^{-\iota\omega})^{-1*},$$

$$\hat{\phi}_\delta(\beta, \Sigma) = \frac{1}{2}\log\det(\Sigma_0\Sigma^{-1}) + \frac{1}{4\pi}\int_{-\pi}^{\pi}\mathrm{tr}\{[\tilde{f}(\omega; \beta_0, \Sigma_0)^{-1} - \tilde{f}_\delta(\omega; \beta, \Sigma)^{-1}]\hat{I}(\omega)\}\, d\omega,$$

and

$$\phi_\delta(\beta, \Sigma) = \frac{1}{2}\log\det(\Sigma_0\Sigma^{-1}) + \frac{1}{4\pi}\int_{-\pi}^{\pi}\mathrm{tr}\{I_n - \tilde{f}_\delta(\omega; \beta, \Sigma)^{-1}\tilde{f}(\omega; \beta_0, \Sigma_0)\}\, d\omega.$$

Because $\hat{I}(\omega)$ is positive semidefinite for each $\omega \in [-\pi, \pi]$, we have $\hat{\phi}(\beta, \Sigma) \leq \hat{\phi}_\delta(\beta, \Sigma)$ for all $(\beta, \Sigma) \in \mathbb{B}_{n,q} \times \mathbb{S}_n$ and $\delta > 0$.

The discussion surrounding Lemma 1 in Dunsmuir and Hannan (1976, p. 350) gives

$$\sup_{(\beta, \Sigma) \in \tilde{\mathcal{K}}(c_1, c_2, c_3)} |\hat{\phi}_\delta(\beta, \Sigma) - \phi_\delta(\beta, \Sigma)| = o_p(1),$$

for the constants $c_1, c_2, c_3 > 0$ found in Step 2. Moreover, for the set $\mathcal{K}$ found in Step 1, the arguments on pp. 355–356 in Dunsmuir and Hannan (1976) imply

$$\inf_{\delta > 0}\sup_{(\beta, \Sigma) \in \tilde{\mathcal{K}}(c_1, c_2, c_3) \cap \mathcal{K}^c} \phi_\delta(\beta, \Sigma) < \phi(\beta_0, \Sigma_0) = 0.$$

Hence, there exist constants $\delta > 0$ and $\zeta_2 < 0$ such that

$$\sup_{(\beta, \Sigma) \in \tilde{\mathcal{K}}(c_1, c_2, c_3) \cap \mathcal{K}^c} \hat{\phi}(\beta, \Sigma) \leq \sup_{(\beta, \Sigma) \in \tilde{\mathcal{K}}(c_1, c_2, c_3) \cap \mathcal{K}^c} \hat{\phi}_\delta(\beta, \Sigma) \leq \zeta_2 + o_p(1).$$

Step 4   Steps 2–3 imply that

$$\sup_{(\beta, \Sigma) \notin \mathcal{K}} \hat{\phi}(\beta, \Sigma) \leq \max\left\{\sup_{(\beta, \Sigma) \in \tilde{\mathcal{K}}(c_1, c_2, c_3) \cap \mathcal{K}^c} \hat{\phi}(\beta, \Sigma), \sup_{(\beta, \Sigma) \notin \tilde{\mathcal{K}}(c_1, c_2, c_3)} \hat{\phi}(\beta, \Sigma)\right\} \leq \zeta + o_p(1),$$

where $\zeta = \max\{\zeta_1, \zeta_2\} < 0$.

Step 5   Steps 1 and 4 imply that the sufficient conditions in Lemma C.3 hold. I conclude that $P^W_{\beta, \Sigma|Y}(\tilde{\mathcal{U}}|Y_T) \overset{P}{\to} 1$ for any neighborhood $\tilde{\mathcal{U}}$ of (β_0, Σ_0) in $\mathbb{B}_{n,q} \times \mathbb{S}_n$.

REMARK.   Finally, I prove an assertion in Section C.5.3: Lemma C.2 holds for the *discretized* Whittle likelihood that replaces integrals $(2\pi)^{-1}\int_{-\pi}^{\pi} g(\omega)\, d\omega$ in the definition of $\log p^W_{Y|\beta, \Sigma}(Y_T|\beta, \Sigma)$ with sums $T^{-1}\sum_{k=0}^{T-1} g(\omega_k)$, $\omega_k = 2\pi k/T$.

The proof of Theorem 4(ii) of Dunsmuir and Hannan (1976, p. 356) shows that steps 1–3 above carry through if the integral in expression (C.4) is replaced with a discretized sum. The only other effect of discretizing the integrals in the Whittle likelihood is that

the discretized version of the Kolmogorov–Szegö formula (C.3) does not hold exactly. Instead,

$$T^{-1} \sum_{j=0}^{T-1} \log \det(\tilde{f}(\omega_j; \beta, \Sigma)) = \log \det(\Sigma) - n \log(2\pi) + T^{-1} \sum_{j=0}^{T-1} \log \big| \det(\Phi(\beta; e^{-\iota\omega_j})) \big|^2.$$

The posterior consistency result for the discretized Whittle posterior follows from steps 1–4 above if I show

$$\sum_{j=0}^{T-1} \log \big| \det(\Phi(\beta; e^{-\iota\omega_j})) \big|^2 \leq 2nq \log 2 \tag{C.6}$$

for all $(\beta, \Sigma) \in \mathbb{B}_{n,q} \times \mathbb{S}_n$ and, furthermore,

$$\sum_{j=0}^{T-1} \log \big| \det(\Phi(\beta; e^{-\iota\omega_j})) \big|^2 = o_p(1) \tag{C.7}$$

uniformly in a small neighborhood of (β_0, Σ_0) in $\mathbb{B}_{n,q} \times \mathbb{S}_n$.

For any $\beta \in \mathbb{B}_{n,q}$ and $z \in \mathbb{C}$, write $\det(\Phi(z; \beta)) = \det(I_n + \sum_{\ell=1}^q \beta_\ell z^\ell) = \prod_{b=1}^{nq}(1 - a_b(\beta)z)$ for some complex scalars $\{a_b(\beta)\}_{1 \leq b \leq nq}$ that depend on β and satisfy $|a_b(\beta)| < 1$ (Brockwell and Davis (1991, p. 191)). From the Taylor series $\log(1 - z) = -\sum_{s=1}^\infty z^s/s$ (valid for $z \in \mathbb{C}$ inside the unit circle) we get, for all $\beta \in \mathbb{B}_{n,q}$,

$$\sum_{k=0}^{T-1} \log \det(\Phi(e^{-\iota\omega_k}; \beta)) = -\sum_{k=0}^{T-1} \sum_{b=1}^{nq} \sum_{s=1}^\infty \frac{(a_b(\beta)e^{-\iota\omega_k})^s}{s} = -\sum_{b=1}^{nq} \sum_{s=1}^\infty \frac{(a_b(\beta))^s}{s} \sum_{k=0}^{T-1} e^{-\iota\omega_k s}.$$

Since $\sum_{k=0}^{T-1} e^{-\iota\omega_k s}$ equals T when s is an integer multiple of T, and equals 0 otherwise,

$$\sum_{k=0}^{T-1} \log \det(\Phi(e^{-\iota\omega_k}; \beta)) = -\sum_{b=1}^{nq} \sum_{s=1}^\infty \frac{(a_b(\beta))^{sT}}{s} = \sum_{b=1}^{nq} \log(1 - (a_b(\beta))^T).$$

Hence,

$$\sum_{k=0}^{T-1} \log \big| \det(\Phi(e^{-\iota\omega_k}; \beta)) \big|^2 = \sum_{k=0}^{T-1} \log \det(\Phi(e^{-\iota\omega_k}; \beta)) + \sum_{k=0}^{T-1} \log \det(\Phi(e^{-\iota\omega_k}; \beta)^*)$$

$$= \sum_{b=1}^{nq} \log \big| 1 - (a_b(\beta))^T \big|^2$$

$$\leq nq \log 4,$$

where the inequality uses $|1 - (a_b(\beta))^T| < 2$. Claim (C.6) follows. For β in a small neighborhood of β_0, $\max_b |a_b(\beta)|$ is uniformly bounded away from 1. This implies claim (C.7).

C.6.5 *Proof of Lemma C.3* I follow the proof of Theorem 1.3.4 in Ghosh and Ramamoorthi (2003). Fix the neighborhood $\mathcal{U}$, assumed open without loss of generality. Define $\kappa_2 = \sup_{\Gamma \in \mathcal{K} \cap \mathcal{U}^c} \phi(\Gamma)$. Notice that $\hat{\phi}(\Gamma_0) = 0$ and assumption (i) together imply $\phi(\Gamma_0) = 0$. By assumptions (ii)–(iii) and compactness of $\mathcal{K} \cap \mathcal{U}^c$, we therefore have $\kappa_2 < 0$. Moreover, we can find a small neighborhood $\mathcal{V}$ of Γ_0 in Ξ_Γ such that $\kappa_1 = \inf_{\Gamma \in \mathcal{V}} \phi(\Gamma)$ satisfies $\max\{\kappa_2, \zeta\} < \kappa_1 < 0$. We may shrink $\mathcal{V}$ to ensure that it also satisfies $\mathcal{V} \subset \mathcal{U} \cap \mathcal{K}$. Choose $\delta > 0$ such that $\kappa_1 - \delta > \max\{\kappa_2 + \delta, \zeta\}$. Write

$$
P_{\Gamma|Y}(\mathcal{U}|Y_T) = \left(1 + \frac{\int_{\mathcal{U}^c} e^{T\hat{\phi}(\Gamma)} \Pi_\Gamma(d\Gamma)}{\int_{\mathcal{U}} e^{T\hat{\phi}(\Gamma)} \Pi_\Gamma(d\Gamma)}\right)^{-1}
$$

$$
\geq \left(1 + \frac{\int_{\mathcal{K}^c} e^{T\hat{\phi}(\Gamma)} \Pi_\Gamma(d\Gamma) + \int_{\mathcal{U}^c \cap \mathcal{K}} e^{T\hat{\phi}(\Gamma)} \Pi_\Gamma(d\Gamma)}{\int_{\mathcal{V}} e^{T\hat{\phi}(\Gamma)} \Pi_\Gamma(d\Gamma)}\right)^{-1}.
$$

Assumptions (i) and (iv) imply that the following three inequalities hold w.p.a. 1:

$$
\inf_{\Gamma \in \mathcal{V}} \hat{\phi}(\Gamma) > \kappa_1 - \delta, \qquad \sup_{\Gamma \in \mathcal{U}^c \cap \overline{\mathcal{K}}} \hat{\phi}(\Gamma) < \kappa_2 + \delta, \qquad \sup_{\Gamma \in \mathcal{K}^c} \hat{\phi}(\Gamma) < \zeta.
$$

We then have

$$
P_{\Gamma|Y}(\mathcal{U}|Y_T) \geq \left(1 + \frac{\int_{\mathcal{K}^c} e^{\zeta T} \Pi_\Gamma(d\Gamma) + \int_{\mathcal{U}^c \cap \mathcal{K}} e^{(\kappa_2 + \delta)T} \Pi_\Gamma(d\Gamma)}{\int_{\mathcal{V}} e^{(\kappa_1 - \delta)T} \Pi_\Gamma(d\Gamma)}\right)^{-1}
$$

$$
\geq \left(1 + \frac{e^{\zeta T} + e^{(\kappa_2 + \delta)T}}{\Pi_\Gamma(\mathcal{V}) e^{(\kappa_1 - \delta)T}}\right)^{-1}
$$

w.p.a. 1. Since $\Pi_\Gamma(\mathcal{V}) > 0$ by assumption (v), and $\kappa_1 - \delta > \max\{\kappa_2 + \delta, \zeta\}$, I conclude that $P_{\Gamma|Y}(\mathcal{U}|Y_T) \xrightarrow{p} 1$ as $T \to \infty$.

References

Beaudry, P. and F. Portier (2014), "News-driven business cycles: Insights and challenges." *Journal of Economic Literature*, 52 (4), 993–1074. [17]

Brockwell, P. J. and R. A. Davis (1991), *Time Series: Theory and Methods*, second edition. Springer Series in Statistics. Springer. [20, 25, 29, 31]

Dunsmuir, W. and E. J. Hannan (1976), "Vector linear time series models." *Advances in Applied Probability*, 8 (2), 339–364. [28, 29, 30]

Durbin, J. and S. Koopman (2012), *Time Series Analysis by State Space Methods*, second edition. Oxford Statistical Science Series. Oxford University Press. [13, 19]

Fernald, J. (2014), "A quarterly, utilization-adjusted series on total factor productivity." Federal Reserve Bank of San Francisco Working Paper No. 2012-19. [11]

Forni, M., L. Gambetti, and L. Sala (2014), "No news in business cycles." *Economic Journal*, 124 (581), 1168–1191. [17]

Geweke, J. (2010), *Complete and Incomplete Econometric Models*. The Econometric and Tinbergen Institutes Lectures. Princeton University Press. [15]

Ghosh, J. K. and R. V. Ramamoorthi (2003), *Bayesian Nonparametrics*. Springer. [24, 32]

Hannan, E. (1970). *Multiple Time Series*. Wiley Series in Probability and Statistics. John Wiley & Sons. [23, 27, 28]

Hoffman, M. D. and A. Gelman (2014), "The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo." *Journal of Machine Learning Research*, 15, 1593–1623. [20, 21]

Kleijn, B. and A. van der Vaart (2012), "The Bernstein–Von–Mises theorem under misspecification." *Electronic Journal of Statistics*, 6, 354–381. [22]

Lippi, M. and L. Reichlin (1994), "VAR analysis, nonfundamental representations, Blaschke matrices." *Journal of Econometrics*, 63 (1), 307–325. [25, 26, 27]

Lopes, H. F. and J. L. Tobias (2011), "Confronting prior convictions: On issues of prior sensitivity and likelihood robustness in Bayesian analysis." *Annual Review of Economics*, 3 (1), 107–131. [14]

Magnus, J. and H. Neudecker (2007), *Matrix Differential Calculus With Applications in Statistics and Econometrics*, third edition. Wiley Series in Probability and Statistics. John Wiley & Sons. [27]

Müller, U. K. (2012), "Measuring prior sensitivity and prior informativeness in large Bayesian models." *Journal of Monetary Economics*, 59 (6), 581–597. [13]

Müller, U. K. (2013), "Risk of Bayesian inference in misspecified models, and the sandwich covariance matrix." *Econometrica*, 81 (5), 1805–1849. [22]

Neal, R. M. (2011), "MCMC using Hamiltonian dynamics, Chapter 5." In *Handbook of Markov Chain Monte Carlo* (S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng, eds.), Handbooks of Modern Statistical Methods, 113–162, CRC Press. [21]

Ramamoorthi, R. V., K. Sriram, and R. Martin (2015), "On posterior concentration in misspecified models." *Bayesian Analysis*, 10 (4), 759–789. [22]

Rubin, D. B. (1988), "Using the SIR algorithm to simulate posterior distributions." In *Bayesian Statistics 3* (J. M. Bernardo, M. H. DeGroot, D. V. Lindley, and A. F. M. Smith, eds.), 395–402, Oxford University Press. [20]

Shalizi, C. R. (2009), "Dynamics of Bayesian updating with dependent data and misspecified models." *Electronic Journal of Statistics*, 3, 1039–1074. [22]

Sims, E. R. (2012), "News, non-invertibility, and structural VARs." In *DSGE Models in Macroeconomics: Estimation, Evaluation, and New Developments* (N. Balke, F. Canova, F. Milani, and M. A. Wynne, eds.), Advances in Econometrics, Vol. 28, 81–135, Emerald Group Publishing. [17]

Stan Development Team (2015), "Stan modeling language users guide and reference manual, Version 2.8.0." [21]

Stock, J. H. and M. W. Watson (2012), "Disentangling the channels of the 2007–09 recession." *Brookings Papers on Economic Activity (Spring)*, 81–135. [11]

Tamaki, K. (2008), "The Bernstein–von Mises theorem for stationary processes." *Journal of the Japan Statistical Society*, 38 (2), 311–323. [22]