On Weak Identification in Structural VAR(MA) Models

Abstract

We simulate synthetic data from known data generating processes (DGPs) that arise from economic theory, and compare the performance of VAR and VARMA models in fitting the dynamics of our DGPs. We show that in small samples typical of macroeconomic data, the moving average (MA) component of VARMA models is close to being non-identified. This in turn leads to an order reduction when identifying the lag structures of the VARMA models. As a result, VARMA models barely show any advantage over VARs in characterizing the known DGPs. This is a refinement of the conjecture that near non-stationarity, near non-invertibility or weak identification could be possible reasons for the failure of structural VARMA models in providing good estimates of theoretical impulse responses of particular DSGE models.

Keywords: VARMA, VAR, DSGE, impulse response analysis

JEL Classification: C15; C52; C32

1. Introduction

Theoretical dynamic stochastic general equilibrium (DSGE) models imply specific state space representations that characterize the dynamics of deviations of macroeconomic variables from their steady states. In many of these models, all or some of the state variables are unobserved. For example, many macro models have capital or debt as state variables but there are no reliable data on capital or debt. In these models, as long as there are as many observed variables as structural shocks, the unobserved state variables are substituted out, and as a result the reduced form of the observed variables will be a VARMA model with a non-trivial moving average component (King et al., 1988; Cooley and Dwyer, 1998; Fernández-Villaverde et al., 2007). However, in practice finite-lag VAR models are used as reduced-form approximations to locally linear solutions of the DSGE model (see e.g., Christiano et al., 2006; Bagliano and Favero, 1998; Erceg et al., 2005; Pagan and Pesaran, 2008). Chari et al. (2007) and Ravenna (2007) showed that a VAR is incapable of capturing the impulse response dynamics of the true VARMA representation of the DSGE model solution, because the VAR is only a truncated approximation of the true VARMA data generating process (DGP). We, as do others, show that the correct VARMA structure is not identified with realistic sample sizes, which in turn leads to unreliable inference of the structure of the true DGP.[1] Kascha and Mertens (2009) attribute it to the fact that the DGPs used in the previous literature (in particular Christiano et al., 2006; Chari et al., 2007) are:

> "[*n*]*early non-stationary, nearly non-invertible and the correct VARMA representation is close to being not identified.*"

We refine this claim and show that the identification problem also arises in a counterexample where the true VARMA DGP is strictly invertible and far away from being non-stationary. We provide some more general analytical understanding of this problem. We demonstrate that the difficulty in identifying the correct VARMA structure with small samples is caused by the fact that the roots of the AR and MA lag matrix polynomials are always very close to each other, and this is true regardless of whether the system is close to being non-stationary and non-invertible, or not. This near cancellation in the AR and MA dynamics leads to an order reduction when identifying the canonical VARMA structure. The same phenomenon has been noted by Cogley and Nason (1993) in the case of scalar ARMA processes. Here, we provide a multivariate generalization of this insight.

Our experimental design follows and extends the work of Kascha and Mertens (2009), Erceg et al. (2005), and Ravenna (2007). We consider as a benchmark, the stylized real business cycle (RBC) model by Hansen (1985) and derive its equilibrium-restricted VARMA representation (i.e., the true DGP taken by us to be known with certainty). Within our experiments, the researcher fits reduced-form VAR and VARMA models to simulated data

[1] We would like to emphasize that the objects of interest here are the responses to structural shocks, and not the deep parameters of their underlying DSGE model. We have nothing to say here about the issue of identifiability of structural parameters of the true DSGE. On this issue, we refer the reader to Fernández-Villaverde et al. (2016), and references therein.

from the DGP. As a further contribution, we conduct such a thought experiment across a wide class of DGPs, each with increasing layers of dynamic sophistication.

In this note we focus on the RBC model as the fixed DGP. At the end, we summarize similar conclusions from further extensions into alternative DSGE models. The results of these extensions are summarized in Table 1.

2. DGP 1: Hansen's RBC and VARMA Identification

We begin with the most used DGP assumption in the literature: The indivisible labor RBC model of Hansen (1985). There are two exogenous structural shocks in the model—i.e., a non-stationary technology shock Z_t, and a stationary labor supply shock D_t. The total output X_t is induced by a Cobb-Douglas production function. Denote C_t, K_t, and, N_t, respectively, as consumption, capital and labor. A technology shock has a permanent effect on the level of Z_t, and hence on C_t, K_t, and X_t. Therefore, we define the model in terms of the stationary variables $\{N_t,\, D_t,\, \hat{C}_t = C_t/Z_t,\, \hat{X}_t = X_t/Z_t,\, \hat{K}_t = K_t/Z_t,\, \hat{Z}_t = Z_t/Z_{t-1}\}_{t=1}^{\infty}$, and log-linearize around the steady state. For any variable S_t, we define its log-deviation from the steady state value $\bar{S}$, by the lower case letter $s_t = \ln(S_t/\bar{S})$. More details on this model are in our Online Appendix A. Following Blanchard and Quah (1989), the percentage deviations of hours worked n_t and output growth $\Delta \ln X_t = \hat{x}_t - \hat{x}_{t-1} + \hat{z}_t$ are taken as the observable variables, that is, $y_t := (n_t,\, \Delta \ln X_t)'$.

2.1. DGP and experimental design

The RBC model is parameterized following Erceg et al. (2005) and Ravenna (2007). This is reported in the first column of Table 1. This RBC instance implies an equilibrium-restricted VARMA(1,1) representation in terms of y_t:

$$y_t = \begin{pmatrix} 0.94 & 1.05 \\ 0 & 0.80 \end{pmatrix} y_{t-1} + u_t + \begin{pmatrix} -0.25 & -0.92 \\ -0.19 & -0.71 \end{pmatrix} u_{t-1}, \tag{1}$$

where u_t is the reduced form disturbance with zero mean and non-singular covariance matrix Σ_u. The VARMA process (1) is strictly stationary and invertible. The eigenvalues of the AR and MA coefficient matrices are given in Table 1.

Additionally, in the experiments, we let the researchers use a long-run identification restriction as in Blanchard and Quah (1989) on either their VAR or VARMA models, which would be consistent with the true DGP's long-run autocovariance structure. We then examine whether the reduced-form VAR or VARMA models are able to match the impulse dynamics of the true VARMA DGP (1).

We consider two cases in every experiment: From the DGP (1), we simulate large-sample-path ($T = 20,000$), and, small-sample-path ($T = 200$) scenarios. Each of such paths is replicated or sampled 1,000 times. The setting of $T = 200$ corresponds to 50 years of quarterly data, which is similar to the sample sizes examined in previous studies (see, e.g., Kascha and Mertens, 2009), and represents a typical sample size for macroeconomic data. This experimental design is repeated in all remaining DSGE models used as "true DGPs" in alternative experimental cases. The large-sample simulations in all our experiments will show that there

Table 1: Properties of the DSGE models, their implied VARMA DGPs and estimation results on simulated data

		(1) RBC	(2) HI[a]	(3) NEWS	(4) MS
Key parameter values of the models					
	α	0.35	0.36	0.36	1/3
	β	$1.03^{-0.25}$	0.99	0.99	0.99
	δ	2%	$\delta_0 = 2.5\%$	2%	2.5%
	$\bar{N}$	1/3	1/3	1/3	1/3
	μ_z	0.0037	0.0037	0.0037	1
	σ_z	0.0148	0.0148	0.009	0.007
	ρ_d	0.8	0.8	0.8	
	σ_d	0.009	0.009	0.005	
			$\bar{U} = 1$	$\rho_G = 0.99$, $\sigma_G = 0.018$	$\rho = 0.26$
			$\theta_c = 0.7$	$\rho_K = 0.97$, $\sigma_K = 0.025$	$\rho_m = 0.5857$
			$\theta_l = 0.5$	$\rho_L = 0.99$, $\sigma_L = 0.020$	$\rho_z = 0.6$
			$\delta_2 = 0.11$	$\tau_K = 0.39$, $\tau_L = 0.21$	$\sigma_m = 0.00397$
			$\varsigma = 4$	$r_{K,L} = 0.261$	
Properties of the DGPs:					
eig(AR)[b]		0.9459, 0.8	0.99, 0.79, 0.0002, 0	0.98, 0.90, 0.8, 0	0.9386, 0.6, 0.5857, 0
eig(MA)		0.9563, 0	0.91, 0.20	0.92, 0.92	0.9425, 0.0087
True VARMA orders		(1,1)	(2,1)	(2,1)	(2,1)
Correct SCM structure		(1,1)∼(1,0)	(2,1)∼(1,1)	(2,1)∼(1,1)	(2,1)∼(1,1)
VAR(MA) estimation results on simulated data from DGPs:					
$T = 20,000$	% of corr.id. SCM[c]	93.1%	85.9%	97.0%	93.5%
	Median VAR lag[d]	18	49	26	21
$T = 200$	% of corr.id. SCM	4.5%	1.5%	0.2%	1.5%
	Median VAR lag	1	4	1	2
	Most.id. SCM[e]	(1,0)∼(0,0)	(1,1)∼(1,0)	(1,0)∼(0,0)	(1,1)∼(1,0)
	% of most.id. SCM	84.1%	87.8%	79.2%	87.2%

[a] HI stands for the RBC model with habit formation and investment adjustment cost presented in Appendix C.1, NEWS stands for the model with news shocks from Appendix C.2, and MS denotes for monetary search model discussed in Appendix C.3 and Appendix D.

[b] eig(AR) and eig(MA) denote roots of the AR and MA characteristic polynomials of the true (RBC and MS) and estimated (HI and NEWS) VARMA processes.The MA roots for NEWS model are calculated from the corresponding fundamental representation of the VARMA process.

[c] Represents the percentage of correct identification of the SCM structure when the sample size is $T = 20,000$.

[d] Represents the median of the lag lengths chosen by AIC for VAR using $T = 20,000$.

[e] Represents the mostly identified SCM structure when the sample size is $T = 200$. The percentage of identifying this SCM is shown in the row below.

is no identification problem—i.e, the true impulse responses of each given DSGE model are identified.

The specification of the lag order in VAR models allows for the estimation of all parameters. However, VARMA models are not always fully identified in terms of their specification of AR and MA orders. We need to specify the largest lag orders in each equation in order to estimate a VARMA model, because the lag orders in one equation can have implication for the identification of parameters in other equations. Consider the VARMA(1,1) example below:

$$\begin{pmatrix} y_{1,t} \\ y_{2,t} \end{pmatrix} = \begin{pmatrix} \phi_{11} & \phi_{12} \\ 0 & 0 \end{pmatrix} \begin{pmatrix} y_{1,t-1} \\ y_{2,t-1} \end{pmatrix} + \begin{pmatrix} u_{1,t} \\ u_{2,t} \end{pmatrix} + \begin{pmatrix} \theta_{11} & \theta_{12} \\ 0 & 0 \end{pmatrix} \begin{pmatrix} u_{1,t-1} \\ u_{2,t-1} \end{pmatrix}.$$

The fact that the longest lag in the second equation is zero implies that ϕ_{12} and θ_{12} are not separately identifiable.

Therefore, we need to determine the orders of each equation and impose additional zero and normalization restrictions before estimating the VARMA model. This can be done using the scalar components model (SCM) methodology developed by Athanasopoulos and Vahid (2008). This method specifies each row of the VARMA model as an SCM with certain orders (p,q), where p denotes the AR lag length and q denotes the MA lag length in that row. The highest SCM orders of the entire system are always the same as the overall VARMA orders.

2.2. Large-sample simulation

For each simulated large sample path, the structure of the VARMA models used by the researcher within each experiment is specified using the scalar components model (SCM) methodology. For example, most sample paths simulated from process (1) are identified as a VARMA(1,1) model with SCM(1,1) $\sim$ SCM(1,0). The canonical form for this SCM VARMA structure is

$$\begin{pmatrix} 1 & 0 \\ a_0 & 1 \end{pmatrix} y_t = \begin{pmatrix} \phi_{11} & \phi_{12} \\ \phi_{21} & \phi_{22} \end{pmatrix} y_{t-1} + u_t + \begin{pmatrix} \theta_{11} & \theta_{12} \\ 0 & 0 \end{pmatrix} u_{t-1}, \tag{2}$$

where the SCM(1,0) (second row) does not have an MA lag. There are also some zero and normalization restrictions imposed on the left-hand side transformation matrix, in order to ensure a unique identification of the unknown parameters.

In our Online Appendix A.1, we show the following for *large sample* simulations from the DGP: (i) the VAR model still fails to approximate the true DGP's dynamics; and (ii) the VARMA is identified and it approximates the true DGP well. Below we focus on the more realistic case of small samples from the DGP.

2.3. Small-sample simulation

Now consider the more realistic case of limited sample sizes. Unfortunately, the SCM identification procedure for VARMA models always fails to detect the MA component with sample size $T = 200$: It chooses SCM(1,0) $\sim$ SCM(0,0) 84.1% of the time as shown in Table 1, which is equivalent to a VAR(1). The three information criteria, AIC, HQ, and BIC, only choose lag one for most of the estimated VAR models.

Figure 1 depicts the mean and 97.5/2.5 percentile of the estimated impulse responses with 200 observations. It shows that VAR models based on small samples tend to overestimate the initial impact of the technology shock, and underestimate the initial impact of the labor supply shock. This phenomenon is widely found in all cases, even with a larger sample, or a much higher AR lag length. Kilian (2011) suggests that finite VAR approximation to the VARMA process may be poor for realistic sample sizes for any feasible choice of lag length, particularly when the VARMA representation has a large MA root. This is the case here as shown in the first column of Table 1 that one of the MA roots (0.9563) is very close to the unit circle. We also estimate VARMA models with several different SCM structures to validate this. The resulting impulse responses do not show visible differences from those in Figure 1, which is consistent with the conclusion of Kascha and Mertens (2009).[2]

Figure 1: The impulse responses generated from the estimated VARs (small sample)

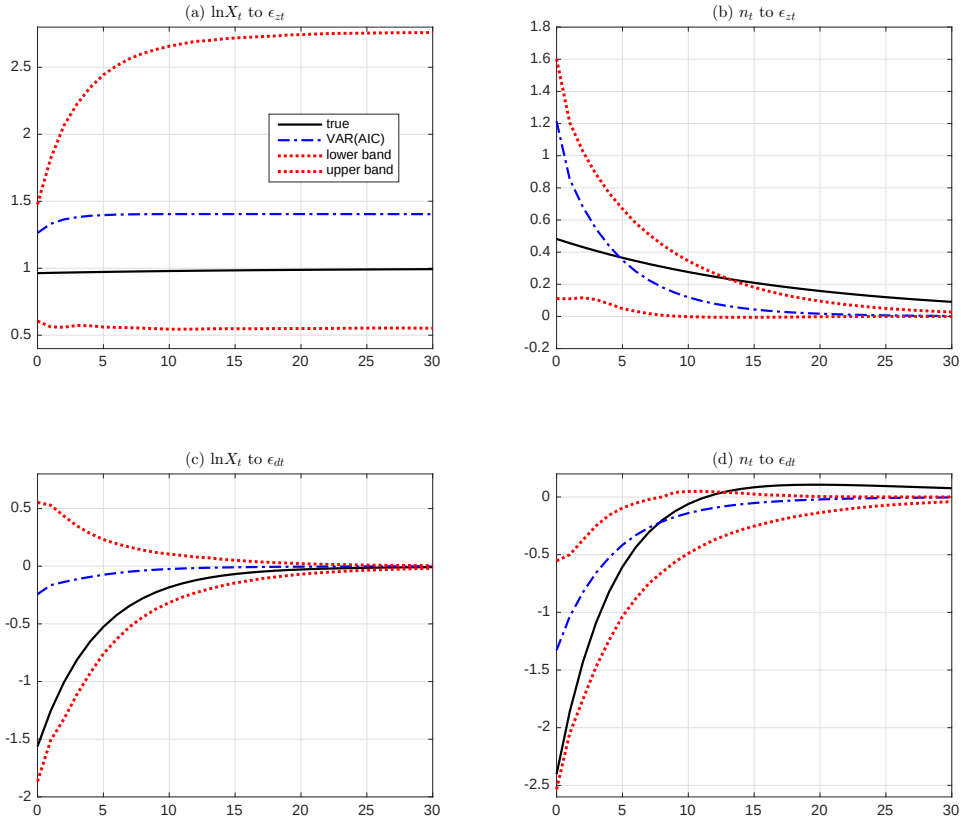

3. Examining the AR and MA roots

We would like to understand with small sample sizes, why there is non-identification of the AR and MA components of the true DGP in (2). Here we decompose analytically and

[2]Plots of the impulse responses generated from the estimated VARMA models are omitted here as they are almost exactly the same as those presented in Figure 1.

show why that is the case by simulating over many parametric instances of the DGP. To do so, we can rotate each parametric instance of the DGP VARMA system in such a way that we can focus on the two-dimensional space containing an equivalent scalar AR and a scalar MA parameter. Note that this rotation still preserves the original SCM(1,1) $\sim$ SCM(1,0) structure. This makes our analysis algebraically and graphically more instructive.

Specifically, consider the instance of the RBC DGP, which is the same as (2), except where $\phi_{22} = 0$:

$$\begin{pmatrix} 1 & 0 \\ a_0 & 1 \end{pmatrix} y_t = \begin{pmatrix} \phi_{11} & \phi_{12} \\ \phi_{21} & 0 \end{pmatrix} y_{t-1} + u_t + \begin{pmatrix} \theta_{11} & \theta_{12} \\ 0 & 0 \end{pmatrix} u_{t-1}. \tag{2'}$$

Recall that the coefficients a_0, ϕ_{ij} and θ_{ij}, $i, j = 1, 2$, are functions of the deep parameters of the RBC model. Note that the MA component only appears in the first row of the system of equations.

The true VARMA form is not identified whenever there is a rotation that yields a shorter lag order. Consider a rotation resulting from pre-multiplying equation (2') by the matrix $\begin{pmatrix} 1 & \kappa \\ 0 & 1 \end{pmatrix}$, where κ is some scalar to be determined. This preserves the original SCM(1,0) form for the second-row equation of (2'), whereas the first row is rotated as:

$$\Big[(1+\kappa a_0) - (\phi_{11} + \kappa\phi_{21})\, L\Big] n_t + (\kappa - \phi_{12}\, L)\Delta \ln X_t = u_{1t} + \kappa\, u_{2t} + \theta_{11} u_{1,t-1} + \theta_{12} u_{2,t-1}. \tag{3}$$

Equation (3) in principle will have a SCM(1,1) structure. However, order reduction can occur if all of the AR and MA roots in (3) turn out to be equal to the same value.

From equation (3), the two AR(1) coefficients for n_t and $\Delta \ln X_t$ are $AR_n^{(1)} = \frac{\phi_{11}+\kappa\phi_{21}}{1+\kappa a_0}$, and $AR_x^{(1)} = \frac{\phi_{12}}{\kappa}$, respectively. The terms on the left-hand side of equation (3) has an MA(1) structure. Hence, we can redefine it as $(1 - \gamma\, L)e_t$, where e_t is a univariate error term, and $|\gamma| < 1$ is the MA(1) coefficient that guarantees the invertibility of this process. Any value of κ corresponds to a point in the three-dimensional space $(AR_n^{(1)}(\kappa),\, AR_x^{(1)}(\kappa),\, \gamma(\kappa))$. In order to reduce the problem to two dimensions with only one AR coefficient and one MA coefficient, we find the value of κ that makes the two AR(1) coefficients the same, and then use this value of κ to calculate the MA(1) coefficient γ.[3] This is done for each experimental or parametric instance of the RBC DGP.

In the exact instance where $AR_n^{(1)} = AR_x^{(1)} = \gamma$, the first row equation (3) degenerates to a static equation with no lagged variables involved. As a result, the MA component is not detectable in the system (2'), whose overall orders will be dominated by the second row equation as a VARMA(1,0) process. More generally, even when each estimated instance of the triplets are not equal to each other but they are close, it still poses a challenge to the identification of the correct VARMA structure, particularly in small samples. Cogley and Nason (1993) came across the same situation in the setting of a univariate ARMA process.

[3] In general this will yield two different values of κ. We choose the one that generates the smaller distance between the AR(1) and MA(1) coefficients.

We provide a multivariate generalization of this insight.

Next, we put this insight into action. We vary the true DGP by simulating the values of the deep parameters in the RBC model randomly.[4] The value of κ is obtained by equating the two AR(1) coefficients $\frac{\phi_{11}+\kappa\phi_{21}}{1+\kappa a_0} = \frac{\phi_{12}}{\kappa}$, which in turn allows us to record their transforms as equivalent scalar AR(1) and MA(1) coefficients.

Figure 2: Instances of the RBC-DGPs projected as AR(1)-MA(1) coefficient pairs

(a) Entire feasible space

(b) Not "near-unit root" feasible subspace

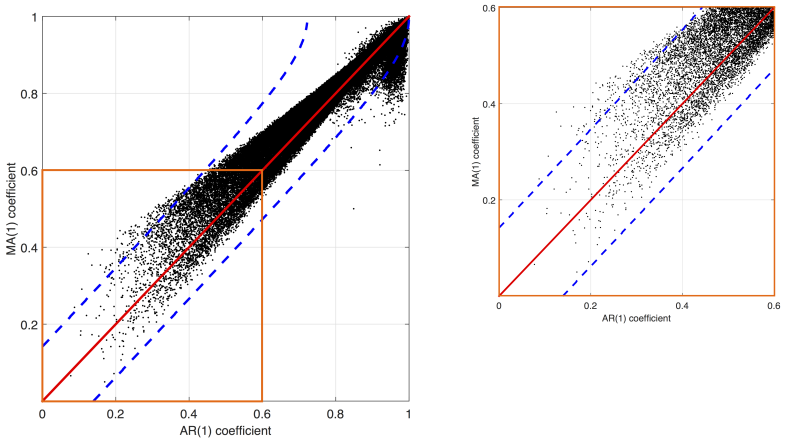

Figure 2 plots corresponding AR(1) and MA(1) coefficient pairs, together with the 45-degree line. The latter is the set of all instances in which the DGP is exactly unidentifiable. The two dashed bands in Figure 2 outlines the region in the parameter space where the true first order autocorrelation coefficient of this ARMA(1,1) process is smaller than $1.96/\sqrt{200}$. That is, with a sample size $T = 200$, given combinations of AR(1) and MA(1) coefficients within these bands, the true first order autocorrelation will not be recognized as statistically significantly different from zero. Consequently, the first row equation (3) is likely to be identified as a SCM(0,0) structure, and hence the second row equation of SCM(1,0) dictates the overall orders of the system (2') as VARMA(1,0).

We can see that from the experimental DGPs, the AR(1) and MA(1) coefficients are always close to each other for most simulated values of the structural parameters. From 1,000,000 simulations, the difference between the AR(1) and MA(1) coefficients is greater than 0.1 in absolute value for only 2.5% of the time. The pair of coefficients almost always falls inside the region where the first order serial correlation is statistically insignificant. This conclusively shows that the inherent structure of the RBC model itself always gives rise to a data generating mechanism, in which this type of near cancellation is very likely to occur. This still hold in our more general setting: Changing the value of structural parameters in the RBC model barely affects the fact that the MA root stays very close to one of the AR roots.

We ensured the invertibility of the VARMA processes in the simulation by setting $|\gamma| < 1$

[4] See the details in Online Appendix B.

in the rotated representation. Further, we can also check that the identification problem is not because of the roots of the RBC model's VARMA DGP being close to unit roots. To illustrate this point, we used $(0.6, 0.6)$ as an arbitrary cutoff point to demarcate roots that are sufficiently far away from being near unit root. From the experiment, out of 1 million simulated DGP instances, about 7% have both roots fall in the set $(0, 0.6]^2$, whereas about 33.5% have both roots fall within $(0, 0.8]^2$. Thus, even conditioning on DGP cases where the true processes are nowhere near unit root, the near cancellation problem is still quite prevalent: from Figures 2a and 2b we can see that a majority of these DGP instances are still inside the non-identification region, irrespective of whether the DGPs have roots that are near unity or not.

In our Online Appendix C, we show that the same problem arises with different DGPs which are solutions to different economic theories that imply different VARMA structures, different variables and different cross-equation restrictions. The results are summarized in Table 1, respectively, in column (2) for a RBC model with habit formation, in column (3) for a RBC model with news shocks, and in column (4) for a search-theoretic monetary model.

4. Conclusion

In structural VAR analysis, in order to estimate the response of macroeconomic variables to policy, demand, and/or supply shocks, a practitioner would start from a reduced form VAR model fitted to macroeconomic time series data. She then imposes reasonable restrictions that technology shocks have permanent effects and purely transitory demand shocks do not. These long-run restrictions are embedded in nearly all theoretical models. With such minimal, and reasonable, restrictions, one would hope to get credible estimates of structural impulse responses.

In this paper, we have shown that in situations where the theoretical model implies a VARMA reduced form for the observed variables, even if we were endowed with the correct long-run insights, we would still end up with incorrect conclusions if we start with a finite order VAR to model the reduced form; and this problem persists even if we had a long sample of data. With the luxury of a long sample of data, starting with a VARMA reduced form produces accurate estimates of the structural impulse responses. Unfortunately, however, we show that with the typical size of macroeconomic data sets, the correct order of the underlying VARMA reduced form cannot be identified, and conclusions made on the basis of the structural VARMA are no better than those resulting from a structural VAR. This confirms the results of Kascha and Mertens (2009), and Ravenna (2007). The additional insight that we provide here is that this is not because the theoretical models imply highly persistent variables or nearly non-invertible moving average components. Rather, this is because the theoretical DSGE models imply reduced forms with autoregressive and moving average roots that are too close to each other, which makes the VARMA model not identifiable in small samples. Our experiments with a variety of theoretical data generating processes show that this problem is quite endemic.

Acknowledgements

We thank the Editor, an anonymous referee, Heather Anderson, Mardi Dungey, Christian Kascha, Eric Leeper, Don Poskitt, and Shaun Vahey for their valuable comments. This research was supported by the Australian Research Council Grant DP0984399. An Online Appendix containing more details and results can be accessed from https://github.com/phantomachine/_varmita.

Reference

Athanasopoulos, G. and Vahid, F. (2008), "A Complete VARMA Modelling Methodology Based on Scalar Components", *J. of Time Ser. Anal.* **29**(3), 533–554.

Bagliano, F. C. and Favero, C. A. (1998), "Measuring Monetary Policy with VAR Models: An Evaluation", *Eur. Econ. Rev.* **42**(6), 1069–1112.

Blanchard, O. J. and Quah, D. (1989), "The Dynamic Effects of Aggregate Demand and Supply Disturbances", *Am. Eco. Rev.* **79**(4), 655–73.

Chari, V. V., Kehoe, P. J. and McGrattan, E. R. (2007), "Are Structural VARs with Long-run Restrictions Useful in Developing Business Cycle Theory?", Technical report, Federal Reserve Bank of Minneapolis.

Christiano, L. J., Eichenbaum, M. and Vigfusson, R. (2006), "Assessing Structural VARs", NBER Working Papers 12353, National Bureau of Economic Research, Inc.

Cogley, T. and Nason, J. M. (1993), "Impulse Dynamics and Propagation Mechanisms in a Real Business Cycle Model", *Econ. Lett.* **43**(1), 77–81.

Cooley, T. and Dwyer, M. (1998), "Business Cycle Analysis without Much Theory. A Look at Structural VARs", *J. of Econom.* **83**, 57–88.

Erceg, C. J., Guerrieri, L. and Gust, C. (2005), "Can Long-Run Restrictions Identify Technology Shocks?", *J. of the Eur. Econ. Assoc.* **3**(6), 1237–1278.

Fernández-Villaverde, J., Rubio-Ramírez, J. F., Sargent, T. J. and Watson, M. W. (2007), "ABCs (and Ds) for understanding VARs", *Am. Econ. Rev.* **97**(3), 1021–1026.

Fernández-Villaverde, J., Rubio-Ramírez, J. F., Schorfheide, F. (2016), "Solution and Estimation Methods for DSGE Models", *forthcoming in* Handbook of Macroeconomics, Volume 2A, John B. Taylor and Harald Uhlig (eds), Elsevier.

Hansen, G. D. (1985), "Indivisible Labor and the Business Cycle", *J. of Monetary Econ.* **16**(3), 309–327.

Kascha, C. and Mertens, K. (2009), "Business Cycle Analysis and VARMA Models", *J. of Econ. Dyn. and Control* **33**(2), 267–282.

Kilian, L. (2011), "Structural Vector Autoregressions", CEPR Discussion Papers 8515, C.E.P.R. Discussion Papers.

King, R. G., Plosser, C. I. and Rebelo, S. T. (1988), "Production, Growth and Business Cycles: II. New Directions", *J. of Monetary Econ.* **21**(2-3), 309–341.

Pagan, A. and Pesaran, M. H. (2008), "Econometric Analysis of Structural Systems with Permanent and Transitory Shocks", *J. of Econ. Dyn. and Control* **32**(10), 3376–3395.

Ravenna, F. (2007), "Vector Autoregressions and Reduced Form Representations of DSGE Models", *J. of Monetary Econ.* **54**(7), 2048–2064.