Empir Econ (2013) 45:831–844
DOI 10.1007/s00181-012-0632-y

Identification of monetary policy in SVAR models: a data-oriented perspective

Matteo Fragetta · Giovanni Melina

Received: 11 October 2011 / Accepted: 4 July 2012 / Published online: 15 August 2012

Abstract In the literature using short-run timing restrictions to identify monetary policy shocks in vector-auto-regressions (VAR) there is a debate on whether (i) contemporaneous real activity and prices or (ii) only data typically observed with high frequency should be assumed to be in the information set of the central bank when the interest rate decision is taken. This paper applies graphical modeling theory, a data-based tool, in a small-scale VAR of the US economy to shed light on this issue. Results corroborate the second type of assumption.

Keywords Monetary policy · SVAR · Graphical modeling · Identification

JEL Classification E43 · E52

1 Introduction

Vector-auto-regressions (VARs) are a widely used tool to provide stylized facts about responses of macroeconomic variables to structural shocks. These facts are useful *per se* and also serve as guidelines in evaluating or calibrating theoretical business cycle models. The literature employing VARs to identify and estimate the effects of monetary policy shocks using short-run timing restrictions typically distinguish among three sets of variables: (i) the information set, i.e., the set of variables known to

M. Fragetta
Department of Economics, University of Salerno, Via Ponte don Melillo, 84084 Fisciano, Italy
e-mail: mfragetta@unisa.it

G. Melina (✉)
School of Economics, Faculty of Business, Economics and Law, University of Surrey,
Ground Floor AD Building, Guildford, Surrey GU2 7XH, UK
e-mail: giovanni.melina@gmail.com

the monetary authorities when the policy decision is taken; (ii) the policy instrument; and (iii) the set of variables the value of which is known only after the policy is set. Such a distinction often suggests a block-recursive structure exploitable in identifying the VAR. Most of the existing empirical papers in the field can be classified into two broad groups, which differ in the content of the information set of the monetary authority.

The first group of papers, which can be thought of following a "*workhorse*" approach, include, among many others, Christiano and Eichenbaum (1992) and Christiano et al. (1996), as well as the influential paper by Christiano et al. (2005). These studies hold that the central bank has at its disposal sources of information about the economy well beyond the published data. In fact, policymakers have access to monthly or even daily estimates of a series of indicators on economic activity and prices sufficient to provide them with a clear and prompt indication of the state of the economy. Consistent with this argument, the assumption made is that, among other variables, the monetary authority is capable to observe the contemporaneous (within quarter) values of output and domestic prices (GDP deflator) at the time of the monetary policy decision.

The second group of papers can be thought of adopting an "*alternative*" approach. This approach is adopted, for instance, by Sims and Zha (1998), the extension proposed by Kim and Roubini (2000) with monthly data and international variables, and the macroecometric model of the UK proposed by Garratt et al. (2003). These papers argue that only high-frequency data should be assumed to be in the information set of the central bank. For example, Sims and Zha (1998) use quarterly data and find it more reasonable to assume that only contemporaneous money supply and commodity prices are known to the central bank when the interest rate is set, since such indices are released at monthly and daily frequencies, respectively. On the contrary, proper measures of variables such as the real GDP and the GDP deflator are assumed to be known to policymakers only with a lag.[1]

Both approaches make use of reasonable and convincing arguments; hence, in principle, there is no clear-cut reason why one should be preferred to the other. This makes the task of imposing a priori short-run identifying restrictions contentious and complex. In fact, especially in small-scale VARs, conditional also on the degree of correlation between reduced-form residuals, results depend (at least quantitatively) on the various possible timing restrictions imposed.

This paper applies Graphical Modeling (GM) theory to a small-scale VAR of the US economy to establish whether the data are informative on which of the two approaches is preferable. The methodology is well suited to establish short-run timing restrictions as it is able to characterize the relationship between contemporaneous variables in terms of linear predictability. It is, therefore, helpful in clarifying the issue from a statistical point of view. Reale and Wilson (2001) and Wilson and Reale (2008) show how the theory can be used in a VAR, while Oxley et al. (2009) and Fragetta and Melina (2011) are examples of how the method can be applied to macroeconomic analysis.

[1] For an extended survey of the literature, see Christiano et al. (1999).

Results are in line with the alternative approach. In other words, GM suggests that only high-frequency data are in the information set of the central bank when it sets the interest rate. For the sake of completeness, also impulse-response analysis is presented. This exercise unveils that the two approaches generate similar responses to an interest rate shock, featuring only minor quantitative differences although real output shows a faster and longer-lived response with the workhorse approach.

The remainder of the paper is structured as follows. Section 2 describes the econometric methodology. Section 3 presents the data. Section 4 illustrates the results. Finally, Section 5 concludes.

2 Econometric methodology

This section presents the econometric strategy adopted in the analysis. Section 2.1 illustrates the basic tools of graphical modeling theory, while Section 2.2 shows how these tools can be applied in the identification of a SVAR.

2.1 Graphical modeling

GM is a statistical approach aiming at uncovering statistical causality from partial correlations observed in the data, which can be interpreted as linear predictability in the context of least-square estimation. Primal contributions to the methodology are due to Dempster (1972) and Darroch et al. (1980).

A *graph* is formally a pair $G = (V, E)$ where the elements of V are called *vertices* (or *nodes*) and the elements of E are called *edges* or *lines*. The most informative object of the procedure is the *Directed Acyclic Graph* (DAG), in which directed edges (arrows) link *initial nodes* (or *parents*) to *terminal nodes* (or *children*). Figure 1(C2) shows a typical and simple DAG, where nodes A, B, and C represent random variables and the directed edges connecting A and B, and B and C indicate the direction of a statistical causality. When *undirected edges* replace the arrows of a graph, a *Conditional Independence Graph* (CIG) is obtained. In a CIG, a link represents a significant partial correlation between any two random variables conditional on all the remaining variables of the model. Figure 1a shows an example of a CIG. The edge connecting nodes A and B represents a significant partial correlation between A and B conditional on C, while the edge connecting nodes B and C represents a significant partial correlation between B and C conditional on A. In Fig. 1a, the absence of an edge linking A and C implies that if A, B, and C are distributed as a multivariate Gaussian distribution, A and C are independent conditional on B, hence the name CIG.

DAGs and CIGs imply a different definition of joint probability. For example, if we consider a DAG such as the one in Fig. 1(C2), this has a joint distribution equal to

$$f_{A,B,C}(\cdot) = f_{C|B}(\cdot) f_{B|A}(\cdot) f_A(\cdot),$$

while if we take a CIG such as the one in Fig. 1a, we can assert that A and C are independent, conditional on B. Therefore, the implied joint distribution is

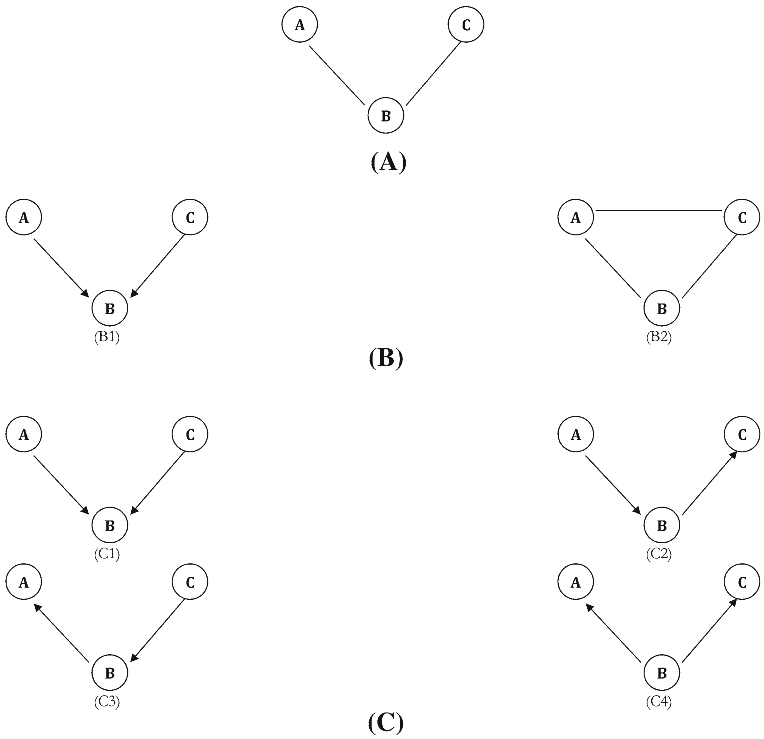

Fig. 1 Conditional independence graphs and directed acyclic graphs. **a** A CIG. **b** A DAG and its corresponding CIG. **c** Hypothetical DAGs deriving from CIG in **a**

$$f_{A,C|B}(\cdot) = f_{A|B}(\cdot) f_{C|B}(\cdot).$$

However, there is a correspondence between the two, represented by the so-called *moralization rule*, as first shown by Lauritzen and Spiegelhalter (1988), who introduced the verb "marrying" instead of "linking" two nodes and defined a graph where two parents of a common child are married (i.e., linked) to be *moral*. The moralization rule states that to derive a unique CIG from a given DAG, arrows should be transformed into undirected edges and unlinked parents of a common child should be linked with an edge. In other words, when two nodes jointly cause a third node and they do not cause each other, from a statistical point of view, there will be a significant correlation between the two. In the DAG shown in Fig. 1(B1), A and C are parents of B and do not cause each other. In order to obtain the corresponding unique CIG, arrows must be transformed into edges and a moral edge has to be added between parents A and C as in Fig. 1(B2). Putting it differently, when both A and C determine B, a significant partial correlation (due to moralization) should be observed between A and C.[2]

[2] While the reader is referred to Lauritzen and Spiegelhalter (1988) for a formal proof of the moralization rule, an example should provide an intuitive insight into the issue: if one wants to become a famous football player (P), he/she must be gifted with good skills (S), and/or must work hard (W). Therefore, S and W are

Fig. 2 Directed cyclic graph

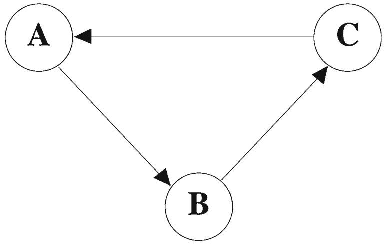

While there is a unique CIG deriving from a given DAG, the reverse is not true. What the econometrician can observe in the data is a CIG, where every edge can assume two possible directions. Therefore, for any given CIG, there are 2^n hypothetical DAGs, where n is the number of edges. Figure 1c shows all the hypothetical DAGs corresponding to the CIG in Fig. 1a. The DAG in Fig. 1(C1) is not compatible with the CIG because the moralization rule requires a moral edge between A and C, which is not captured by the CIG.[3]

Any DAG, by definition, has to satisfy the principle of acyclicality. Therefore, the graph depicted in Fig. 2 cannot be a DAG as it is clearly cyclic. The acyclicality in a DAG allows one to completely determine the distribution of a set of variables and implies a recursive ordering of the variables themselves, where each element in turn depends on none, one or more elements. For example, in the DAG in Fig. 1(C2), A depends on no other variables, B depends on A and C on B.

2.2 Identification of a SVAR with graphical modeling

GM theory can be applied to obtain identification of a structural VAR (SVAR), as shown by Reale and Wilson (2001) and Oxley et al. (2009) among others.

Any SVAR may be turned into a DAG where current and lagged variables are represented by nodes and causal dependence by arrows. After collecting the endogenous variables of interest in the k-dimensional vector X_t, the associated reduced form, or canonical, VAR can be written as

$$X_t = A(L)X_{t-1} + u_t, \tag{1}$$

Footnote 2 continued
the causes of P. Suppose that we know that one individual did not work hard. This per se does not provide any information on whether he/she had good skills. However, if the individual is a famous football player, the only thing we can conclude is that he/she had good skills. Therefore, observing P—which is the effect and not the cause of S and W—is crucial in establishing the partial correlation between S and W.

[3] In the process of obtaining plausible DAGs from an observed CIG, it may also be possible that some of the links captured by the CIG are due to moralization and hence must be eliminated in a corresponding DAG. Such *demoralization process*, in most cases, can be assessed by considering some quantitative rules. Let us suppose we observe a CIG such as the one in Fig. 1(B2). If the true corresponding DAG were the one in Fig. 1(B1), then the partial correlation between A and C, $\rho_{(A,C|B)}$, should be equal to $-\rho_{(A,B|C)} \times \rho_{(B,C|A)}$. In such a case, when tracing DAG Fig. 1(B1), the edge between A and C must be removed.

where $A(L)$ is a polynomial in the lag operator L and u_t is a k-dimensional vector of reduced-form disturbances with $E[u_t] = 0$ and $E[u_t u_t'] = \Sigma_u$.

As reduced-form disturbances are correlated, to identify structural shocks, the reduced-form model has to be transformed into a structural model. Pre-multiplying both sides of Eq. (1) by the $(k \times k)$ matrix A_0, yields the structural form

$$A_0 X_t = A_0 A(L) X_{t-1} + B e_t. \tag{2}$$

The relationship between the structural disturbances e_t and the reduced-form disturbances u_t is described by

$$A_0 u_t = B e_t, \tag{3}$$

where A_0 also describes the contemporaneous relations among the endogenous variables and B is a $(k \times k)$ matrix. In the structural model, disturbances are assumed to be uncorrelated with each other. In other words, the covariance matrix of the structural disturbances Σ_e is diagonal.

As it is, the model described by Eq. (2) is not identified because there may be possibly many matrices A and B that satisfy (2). Therefore, first matrix B can be restricted to be a $(k \times k)$ diagonal matrix. Second, to impose identifying restrictions on matrix A_0, graphical modeling theory can be applied to trace DAGs of the contemporaneous variables.

The acyclicality of DAGs implies a recursive ordering of the variables that makes A_0 a lower-triangular matrix. A_0 has generally zero elements also in its lower triangular part; hence, in general, the model is over identified. The GM methodology has the distinctive feature that the variable ordering and any further restrictions come from statistical properties of the data.

First, as shown by Oxley et al. (2009), to construct the CIG among contemporaneous variables, one has to derive the sample partial correlation between each pair of contemporaneous variables, conditioned on the values of the remaining contemporaneous variables and the lagged values of all variables. This can be computed from the inverse $\hat{W}$ of the sample covariance matrix $\hat{V}$:

$$\hat{\rho}\left(x_{i,t}, x_{j,t} | \{x_{k,t}\}\right) = -\frac{\hat{W}_{ij}}{\sqrt{(\hat{W}_{ii}\hat{W}_{jj})}}, \tag{4}$$

where $\{x_{k,t}\}$ is the whole set of variables excluding the two considered. Whenever a sample partial correlation is statistically significant a link is retained. Swanson and Granger (1997) have applied a similar strategy to sort out causal flows among contemporaneous variables, i.e., applying a residual orthogonalization of the innovations from a canonical VAR. In particular, Swanson and Granger (1997) have also focused on testing the constraints implied by structural forms that have been used in practice. Their test is based on pairwise partial correlations, which are thus not directional and therefore do not give rise to a causal interpretation (or linear predictability interpretation in the case of least square estimation). This is why, once partial correlations are obtained, they suggest utilizing prior economic information to draw a causal order. As also remarked by Swanson and Granger themselves, the structural form of dependence between variables is equivalent to a DAG. With GM and its rules, starting from pairwise partial correlations, it is possible to construct a CIG which imply data-determined constraints on permissible DAGs. As a result, the approach offers a data-driven systematic procedure that leads to the selection of the best DAG, which has the interpretation of statistical causation (or linear predictability in the context of a SVAR).

All possible DAGs (satisfying the moralization rule) which represent alternative competitive models are compared via likelihood-based methods—such as the Akaike Information Criterion (AIC), the Hannan and Quinn Information Criterion (HIC) or the Schwarz Information Criterion (SIC)—and/or based on their out-of-sample forecasting performances and the best-performing one is chosen. In order to construct an empirically well-founded SVAR, one has to assure that the covariance matrix of the resulting residuals is diagonal. A first diagnostic check is thus inspecting the significance of such correlations. Further diagnostic checks are advisable. For instance, as this procedure typically entails the imposition of over-identifying restrictions, a χ^2 likelihood-ratio test should be conducted.[4]

3 Data

The empirical analysis presented in the remainder of the paper employs quarterly US data over the period 1965:1–2007:4. The starting year coincides with that used by Christiano et al. (1999) and Christiano et al. (2005) while the end date falls in a pre-crisis quarter.

The model is a four-variable VAR including (i) the log of real GDP, y_t; (ii) the effective federal funds rate (quarterly average), r_t; (iii) the log the GDP implicit price deflator, p_t; and (iv) the log of the quarterly average of a commodity price index (producer price index), cp_t. The variables are representative of the real activity, monetary policy, and price dynamics. Such a model specification represents a minimal setting similar to those adopted by Stock and Watson (2001)—for illustrative purposes—and by more recent contributions such as Primiceri (2005) and Koop et al. (2009). The addition of a commodity price proves helpful in ruling out the *price puzzle*.[5] Giordani (2004) argues that the commodity price index solves the price puzzle not because it is useful in forecasting inflation (as it is often argued in the literature), but because it is correlated with the output gap (typically omitted in VARs). In the context of this paper, the commodity price index represents a high-frequency variable the central bank looks at and, in accordance with Giordani (2004), this variable may act as an indicator of the

[4] In some cases, the distributional properties of the variables for different DAGs are likelihood equivalent although the residual series are different. In such cases, it is possible to construct DAG models by considering only the lagged variables that play a significant role in explaining contemporaneous variables determined by the significant partial correlation. This can help via comparison of information criteria determine the best DAG for contemporaneous variables.

[5] The term *price puzzle* is due to Sims (1992). Christiano et al. (1999) show that omitting a commodity price index from the VAR specification delivers a rise in the price level that lasts several years after a contractionary monetary policy shock.

state of the business cycle. The absence of monetary aggregates is due to a preference for parsimony coupled with the fading role of monetary aggregates in the conduct of monetary policy as empirically shown by Estrella and Mishkin (1997), among others, and theoretically explored by Woodford (2008).

A constant is included in the VAR and results are reported both for a VAR in levels, with and without a deterministic trend,[6] and for a VAR in which the logs of GDP, the GDP deflator, and the commodity price index have been first differenced. The sampling properties of GM are valid regardless of the presence of unit roots in the data, as shown by Wilson and Reale (2008). In fact, we show below that the three model specifications give rise to the same CIGs and DAGs.

All series are extracted from the ALFRED database of the Federal Reserve Bank of St. Louis. The commodity price index was adjusted for seasonality by the Census X12 method, while the other variables were seasonally adjusted by the source.

4 Results

DAGs are obtained by fitting the data to Eq. (1). The lag order is selected via the AIC.[7] Table 1 reports the estimated partial correlation matrices of the series and their significance at 0.10, 0.05, and 0.01 levels. Partial correlation matrices are constructed by computing the sample correlations between each pair of contemporaneous variables, conditioned on the values of the remaining contemporaneous variables and the lagged values of all variables.

Both the matrix coming from the model in first differences and those coming from the model in levels (with and without trend) translate into the same CIG depicted in Fig. 3. The three edges in the CIG cannot be moral, as moral edges link parents of a common child. The $2^3 = 8$ possible DAGs implied by the CIG are reported in Fig. 4. The moralization rule implies that DAGs (A), (E), (G), and (H) can be discarded as they are not compatible with the observed CIG. In fact, in (A) and (E), r_t and p_t are parents of common child cp_t, which would imply a moral edge between r_t and p_t that does not appear in the observed CIG, respectively. In (G) and (H), y_t and cp_t are parents of common child r_t, which would imply a moral edge between y_t and cp_t that again does not appear in the observed CIG, respectively. The four remaining models are compared via the information criteria mentioned in Section 2. Table 2 shows that the three information criteria for all model specifications are minimized by the model implied by DAG (C), which in turn implies that, within the same quarter, the Federal funds rate is not affected by shocks to the general price level and the real output, while it is affected by shocks to the commodity price.

[6] We prefer to report results for both cases, as in the literature both options are explored. For instance, while Bernanke (1986) includes a deterministic trend in the level specification, Christiano et al. (2005) carry out the estimation including only the levels of the variables.

[7] The AIC typically selects a larger number of lags with respect to SIC and HIC, which we prefer based on the view that the consequences of overestimation of the order are less serious than underestimation (Kilian 2001).

Table 1 Estimated partial correlations of the variables

	r_t	y_t	p_t	cp_t
(a) Model in first differences				
r_t	1.000			
y_t	0.183*	1.000		
p_t	0.121	−0.096	1.000	
cp_t	0.202***	−0.067	0.387***	1.000
(b) Model in levels				
r_t	1.000			
y_t	0.185**	1.000		
p_t	0.062	−0.121	1.000	
cp_t	0.211***	−0.016	0.435***	1.000
(c) Model in levels with deterministic trend				
r_t	1.000			
y_t	0.219**	1.000		
p_t	0.011	−0.088	1.000	
cp_t	0.220***	−0.026	0.439***	1.000

*, ** and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively. The corresponding threshold values for the baseline model are 0.1270, 0.1504 and 0.1963, respectively

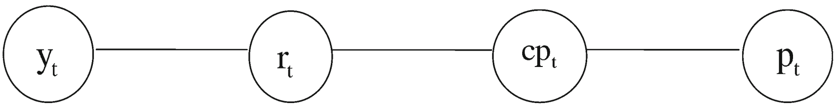

Fig. 3 Sample CIG

Table 3 indicates that the performance of model (C) is highest also as far as out-of-sample predictability is concerned. Retaining approximately the first third of observations as the training period, first, one-step ahead forecasts were recursively computed for the period 1980:1–2007:4, i.e., conditional only on the information up to the date of the forecast and with subsequent re-estimation every time a new observation was included in the sample. Second, following Clarida et al. (2006), the cross-sectional mean of square forecast errors (MSFE) of each variable of the SVAR was computed for each model.

Table 3(a) reports the ratios between the average MSFE (A-MSFE) of each model relative to that of model (C). The forecast accuracy of model (C) is the highest in every specification given that the ratios are all larger than unity. To take the possible uncertainty around parameter estimates into account, the models are compared also by means of the Diebold–Mariano (DM) test (Diebold and Mariano 1995). Table 3(b) reports the DM test statistics computed on the differences between the MSFE of each competing model and that of model (C). In accordance with Table 3(a), the test statistics are systematically positive. The null hypothesis of zero difference is rejected in most cases, at least at a 0.10 significance level. In particular, for the models in first

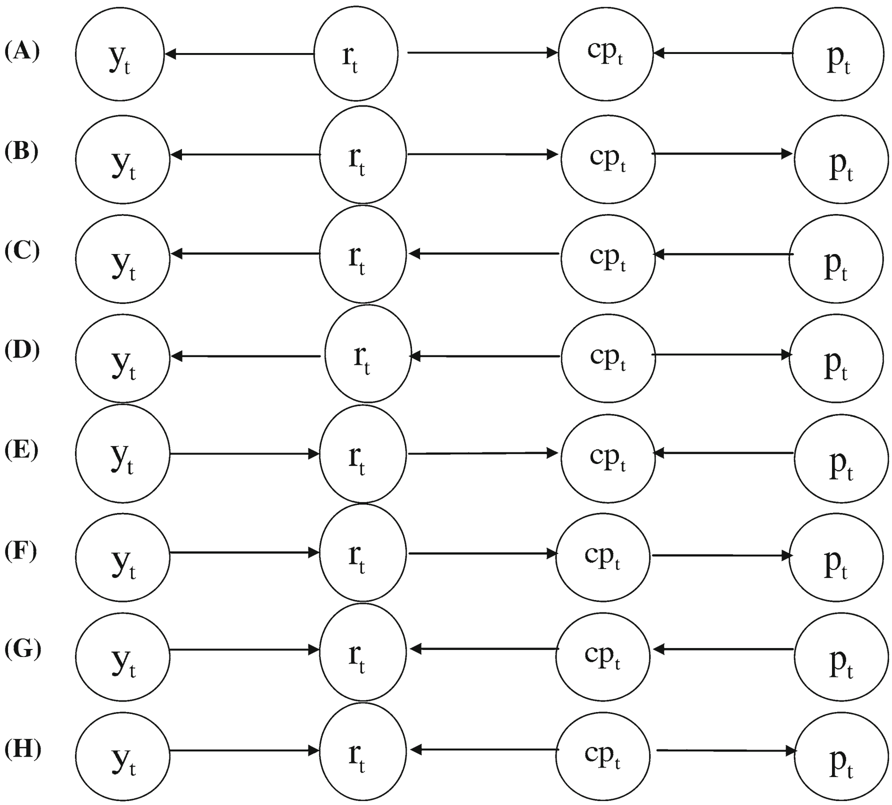

Fig. 4 All possible DAGs deriving from the estimated CIG

Table 2 Information criteria associated to feasible DAGs

Model	AIC	HIC	SIC
(a) Model in first differences			
B	−418.56	−398.24	−368.48
C	−453.05	−433.17	−403.42
D	−358.26	−337.94	−308.19
F	−405.32	−385.00	−355.24
(b) Model in levels			
B	−466.06	−445.74	−415.98
C	−521.79	−501.46	−471.71
D	−484.43	−464.11	−434.35
F	−471.50	−451.17	−421.42
(c) Model in levels with deterministic trend			
B	−469.87	−444.47	−407.27
C	−525.34	−499.94	−462.74
D	−488.20	−462.79	−425.60
F	−463.66	−438.26	−401.06

AIC Akaike Information Criterion, *HIC* Hannan-Quinn Information Criterion, *SIC* Schwarz Information Criterion

Table 3 Out-of-sample predictability associated to feasible DAGs relative to model (C) over 1980: 1–2007:4

	FD	LEV	LEV-TR
(a) Ratios of A-MSFEs			
B/C	1.24	1.14	1.13
D/C	1.18	1.02	1.02
F/C	1.26	1.04	1.06
(b) Diebold–Mariano test statistics			
B–C	2.84**	1.70*	1.30
D–C	2.22**	3.44**	4.13**
F–C	3.05**	1.03	1.15

FD models in first differences, *LEV* models in levels, *LEV-TR* models in levels with deterministic trend, *A-MSFE* average mean square forecast error

*, ** Significance of the Diebold–Mariano test statistics at 0.05 and 0.10, respectively

Table 4 Correlations between residuals of the DAGs fitted to the VAR estimated innovations

	ϵ_t^r	ϵ_t^y	ϵ_t^p	ϵ_t^{cp}
(a) Model in first differences				
ϵ_t^r	1.000			
ϵ_t^y	0.026	1.000		
ϵ_t^p	0.092	−0.112	1.000	
ϵ_t^{cp}	−0.043	−0.048	0.000	1.000
(b) Model in levels				
ϵ_t^r	1.000			
ϵ_t^y	0.022	1.000		
ϵ_t^p	0.036	−0.144	1.000	
ϵ_t^{cp}	−0.018	−0.020	0.000	1.000
(c) Model in levels with deterministic trend				
ϵ_t^r	1.000			
ϵ_t^y	0.020	1.000		
ϵ_t^p	0.035	−0.146	1.000	
ϵ_t^{cp}	−0.018	−0.010	0.000	1.000

The two-standard-error band for a sample size of 204 is $\pm$ 0.1538

differences, the null is always rejected at a 0.05 level. As shown by Inoue and Kilian (2006), a biunivocal correspondence between model rankings based on (in-sample) information criteria and (out-of-sample) forecast errors does not always hold. In the specific case of this paper, however, it is reassuring to observe that model comparisons made with out-of-sample methods clearly go into the direction of corroborating the results obtained via in-sample criteria.

In sum, GM selects only data available at high frequencies for the information set of the central bank, providing support for the alternative approach.

A diagnostic check on the cross-correlation matrix of the resulting residuals reported in Table 4 unveils that all cross-correlations lie within two standard errors from zero. In addition, DAG (C) implies three over-identifying restrictions, which are not rejected at any conventional significance level.

For the sake of completeness, Fig. 5 reports the impulse responses to a positive Federal funds rate shock obtained by adopting both the workhorse and the alternative

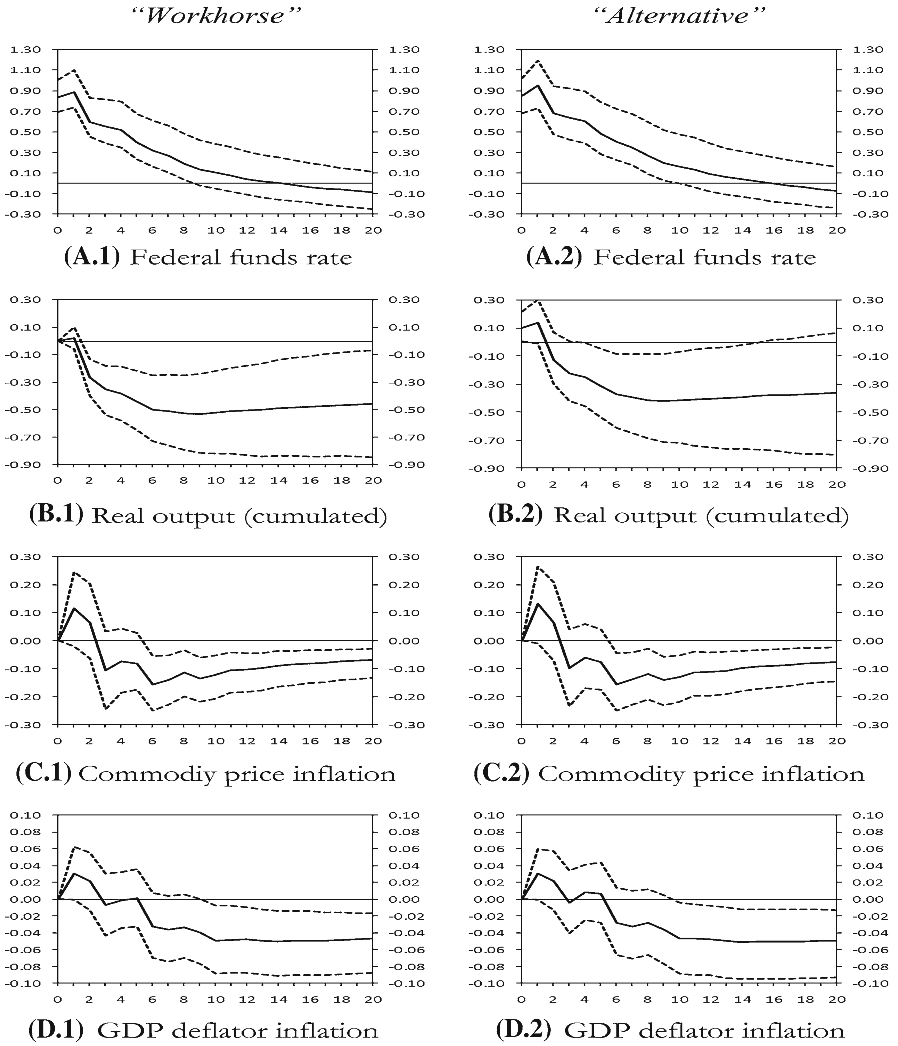

Fig. 5 Impulse responses to a Federal funds rate shock: “Workhorse” versus “Alternative” (GM-consistent) identification. *Dashed lines* represent 90 % confidence intervals computed according to Hall (1992)’s algorithm with 2000 bootstrap replications. Responses are shown for a 20-quarter horizon

identification approach, the latter being consistent with GM. The two approaches generate impulse responses with small quantitative differences although real output shows a faster and longer-lived response with the workhorse approach compared to the alternative approach.

5 Conclusion

The empirical approaches aiming at identifying monetary policy shocks can be classified into two groups: the “*workhorse*” approach, which assumes that the central bank has sufficient information to accurately infer what contemporaneous real output and GDP deflators are when it takes the monetary policy decision; and the "*alternative*" approach, which assumes that only variables observed with high frequency, such as commodity prices, are in the information set of the central bank at the time of policy setting. This paper makes use of GM theory to identify a small-scale VAR of the US economy and finds that the application of such a data-based tool give rise to identifying restrictions consistent with the alternative approach. When impulse-response analysis is concerned, however, the workhorse approach and the model identified by imposing restrictions suggested by GM—coinciding with the alternative approach—generate responses to a Federal funds rate shock featuring only small quantitative differences although real output shows a faster and longer-lived response with the workhorse approach.

Acknowledgments We are grateful to Heather Anderson (editor), Sergio Destefanis, Emanuel Gasteiger, Marco Reale, Lucio Sarno, and Granville Wilson, an anonymous associate editor and an anonymous referee for helpful comments. Moreover, we thank Marco Reale and Granville Wilson for kindly sending some of their Matlab codes, and the developers of JMulTi for providing their software as an open source. Giovanni Melina acknowledges financial support from ESRC project RES-062-23-2451. The usual disclaimer applies.

References

Bernanke BS (1986) Alternative explanations of the money-income correlation. Carnegie-Rochester Conf Ser Public Policy 25(1):49–99

Christiano LJ, Eichenbaum M (1992) Liquidity effects, the monetary transmission mechanism, and monetary policy. Econ Perspect (Nov) 16(6):2–14

Christiano LJ, Eichenbaum M, Evans C (1996) The effects of monetary policy shocks: evidence from the flow of funds. Rev Econ Stat 78(1):16–34

Christiano LJ, Eichenbaum M, Evans C (1999) Monetary policy shocks: what have we learned and to what end?. In: Woodford M, Taylor JB (eds) Handbook of macroeconomics, vol 1, chap 2. Elsevier, Amsterdam, pp 65–148

Christiano LJ, Eichenbaum M, Evans C (2005) Nominal rigidities and the dynamic effects of a shock to monetary policy. J Polit Econ 113(1):1–45

Clarida RH, Sarno L, Taylor MP, Valente G (2006) The role of asymmetries and regime shifts in the term structure of interest rates. J Bus 79(3):1193–1224

Darroch JN, Lauritzen SL, Speed TP (1980) Markov fields and log–linear interaction models for contingency tables. Ann Stat 8(3):522–539

Dempster AP (1972) Covariance selection. Biometrics 28(1):157–175

Diebold FX, Mariano RS (1995) Comparing predictive accuracy. J Bus Econ Stat 13(3):253–263

Estrella A, Mishkin FS (1997) Is there a role for monetary aggregates in the conduct of monetary policy?. J Monet Econ 40(2):279–304

Fragetta M, Melina G (2011) The effects of fiscal policy shocks in SVAR models: a graphical modelling approach. Scott J Polit Econ 58(4):539–567

Garratt A, Lee K, Pesaran MH, Shin Y (2003) A long run structural macroeconometric model of the UK. Econ J 113(487):412–455

Giordani P (2004) An alternative explanation of the price puzzle. J Monet Econ 51:1271–1296

Hall P (1992) The Bootstrap and Edgeworth expansion. Springer, New York

Inoue A, Kilian L (2006) On the selection of forecasting models. J Econom 130(2):273–306

Kilian L (2001) Impulse response analysis in vector autoregressions with unknown lag order. J Forecast 20(3):161–179

Kim S, Roubini N (2000) Exchange rate anomalies in the industrial countries: a solution with a structural VAR approach. J Monet Econ 45(3):561–586

Koop G, Leon-Gonzalez R, Strachan RW (2009) On the evolution of the monetary policy transmission mechanism. J Econ Dyn Control 33(4):997–1017

Lauritzen SL, Spiegelhalter DJ (1988) Local computations with probabilities on graphical structures and their application to expert systems. J R Stat Soc Ser B (Methodological) 50(2):157–224

Oxley L, Reale M, Wilson G (2009) Constructing structural VAR models with conditional independence graphs. Math Comput Simul 79(9):2910–2916

Primiceri GE (2005) Time varying structural vector autoregressions and monetary policy. Rev Econ Stud 72(3):821–852

Reale M, Wilson G (2001) Identification of vector AR models with recursive structural errors using conditional independence graphs. Stat Methods Appl 10:49–65

Sims CA (1992) Interpreting the macroeconomic time series facts: the effects of monetary policy. Eur Econ Rev 36(5):975–1000

Sims CA, Zha TA (1998) Does monetary policy generate recessions? Federal Reserve Bank of Atlanta working paper 98-12

Stock JH, Watson MW (2001) Vector autoregressions. J Econ Perspect 15(4):101–115

Swanson NR, Granger CWJ (1997) Impulse response functions based on a causal approach to residual orthogonalization in vector autoregressions. J Am Stat Assoc 92(437):357–367

Wilson GT, Reale M (2008) The sampling properties of conditional independence graphs for I(1) structural VAR models. J Time Ser Anal 29(5):802–810

Woodford M (2008) How important is money in the conduct of monetary policy?. J Money Credit Banking 40(8):1561–1598