Board of Governors of the Federal Reserve System

International Finance Discussion Papers

Number 610

April 1998

THE ROBUSTNESS OF IDENTIFIED VAR CONCLUSIONS ABOUT MONEY

Jon Faust

THE ROBUSTNESS OF IDENTIFIED VAR CONCLUSIONS ABOUT MONEY

Jon Faust*

Abstract: This paper presents a new way to assess robustness of claims from identified VAR work. All possible identifications are checked for the one that is worst for the claim, subject to the restriction that the VAR produce *reasonable* impulse responses to shocks. The statistic on which the claim is based need not be identified; thus, one can assess claims in large models using minimal restrictions. The technique reveals only weak support for the claim that monetary policy shocks contribute a small portion of the forecast error variance of post-War U.S. output in standard 6-variable and 13-variable models.

Keywords: identification, VAR, monetary policy.

* Faust is an economist in the Division of International Finance, Board of Governors of the Federal Reserve System (faustj@frb.gov). This paper was prepared for the fall 1997 Carnegie-Rochester Conference on Public Policy. I offer thanks for comments to David Bowman, Neil Ericsson, Jim Hamilton, Andy Levin, Adrian Pagan, Chris Sims, Jim Stock, and Harald Uhlig and to seminar participants at the Federal Reserve Bank of Minneapolis, the Federal Reserve Board, George Washington University, UCLA, UCSD and at the Carnegie-Rochester Conference. Special thanks to Eric Leeper and Tao Zha for discussions, data, and program code. Michael Sharkey provided excellent research assistance. The views in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

The robustness of identified VAR conclusions about money

Jon Faust
International Finance Division
Board of Governors of the Federal Reserve System
Washington, D.C. 20551
tel. (202) 452-2328 fax (202) 736-5638

October 1997
(revised May 1, 1998)

Abstract

This paper presents a new way to assess robustness of claims from identified VAR work. All possible identifications are checked for the one that is worst for the claim, subject to the restriction that the VAR produce *reasonable* impulse responses to shocks. The statistic on which the claim is based need not be identified; thus, one can assess claims in large models using minimal restrictions. The technique reveals only weak support for the claim that monetary policy shocks contribute a small portion of the forecast error variance of post-War U.S. output in standard 6-variable and 13-variable models.

Keywords: identification, VAR, monetary policy

This paper was prepared for the Fall 1997 Carnegie-Rochester Conference on Public Policy. I offer thanks for comments to David Bowman, Neil Ericsson, Jim Hamilton, Andy Levin, Adrian Pagan, Chris Sims, Jim Stock, and Harald Uhlig and to seminar participants at the Federal Reserve Bank of Minneapolis, the Federal Reserve Board, George Washington University, UCLA, UCSD and at the Carnegie-Rochester Conference. Special thanks to Eric Leeper and Tao Zha for discussions, data, and program code. Michael Sharkey provided excellent research assistance.

The robustness of identified VAR conclusions about money

Jon Faust*
Federal Reserve Board

Abstract: This paper presents a new way to assess robustness of claims from identified VAR work. All possible identifications are checked for the one that is worst for the claim, subject to the restriction that the VAR produce *reasonable* impulse responses to shocks. The statistic on which the claim is based need not be identified; thus, one can assess claims in large models using minimal restrictions. The technique reveals only weak support for the claim that monetary policy shocks contribute a small portion of the forecast error variance of post-War U.S. output in standard 6-variable and 13-variable models.

---

* I offer thanks for comments to David Bowman, Neil Ericsson, Jim Hamilton, Andy Levin, Adrian Pagan, Chris Sims, Jim Stock, and Harald Uhlig and to seminar participants at the Federal Reserve Bank of Minneapolis, the Federal Reserve Board, George Washington University, UCLA, UCSD and at the Carnegie-Rochester Conference. Special thanks to Eric Leeper and Tao Zha for discussions, data, and program code. Michael Sharkey provided excellent research assistance. The views in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or other members of its staff.

In “Macroeconomics and Reality,” Christopher Sims presented the first analysis of monetary policy in vector autoregressive (VAR) models and concluded with a warning about his six-variable, “small scale” example (1980, p.33): “A long road remains, however, between what has been displayed here and models in this style that compete seriously with existing large-scale models on their home ground—forecasting and policy projection.” Sims cited the need to increase the range of policy-relevant variables in the VAR and to improve methods for handling the large number of free parameters in the expanded models. After nearly 20 years on the road, the provocateur might well ask whether VARs can yet seriously compete with large-scale econometric models for analyzing monetary policy.

Important advances have been made. Most recently, Strongin (1995), Christiano, Eichenbaum, and Evans (CEE) (1996) and Bernanke and Mihov (1995) broadened the focus from monetary aggregates to bank reserve markets. Gordon and Leeper (1994) and Bernanke and Mihov (1995) emphasized the importance of taking account of different monetary policy regimes. Sims (1992) demonstrated the importance of including variables in the VAR such as commodity prices that the central bank might use in forecasting inflation.

Despite the advances, most published VARs are smaller, and few are larger, than Sims's originals. Sims and Zha (1996a,b) and Leeper, Sims and Zha (LSZ) (1996) recently made a major break in this respect, demonstrating how to use Bayesian methods to study identified VAR models much larger than those previously studied. Further, these authors greatly clarified the justification for identified VAR methods.

The recent progress has brought some modest claims of victory by VAR practitioners. Sims (1996) lists four conclusions:

1 Most variation in monetary policy instruments is accounted for by responses of policy to the state of the economy, not by random disturbances to policy behavior.

2 Responses of real variables to monetary policy shifts are estimated as modest or nill, depending on the specification.

3 Monetary policy has historically increased interest rates in response to non-policy shocks that increase inflationary pressure by more

than it would have under a policy of fixing the monetary stock.

4 A reasonable picture of the effects of monetary policy shifts emerges only under identifying assumptions of delay in the reaction of certain "sluggish" private sector variables to monetary policy shifts.

Bernanke (1996) and CEE (1997) give similar lists.

Before these advances, there was a common view that changes in sample period, information set, and time aggregation of the data lead to important changes in VAR conclusions (e.g., Todd, 1991; Pagan and Robertson, 1994). The assertions of success seem to have brought renewed vigor to the robustness discussions. Rudebusch (1997) raises several such issues, concluding that measures of monetary policy from VARs do not make sense. Bagliano and Favero (1997) find support for some of Rudebusch's conclusions, especially regarding instability of estimates spanning different operating procedures. They conclude, however, that taking account of such issues does not substantially change certain basic conclusions from the literature. Sims (1996) and CEE (1997) also argue that the basic conclusions are robust.

This paper analyzes a different dimension of robustness, focussing on the identifying assumptions. The approach is motivated by the possibility that Sims was right about small models in "Macroeconomics and Reality." It is straightforward to show that if the world is complicated even in simple ways then small VARs cannot get the right answer even asymptotically—this result is little more than a statement about omitted variable bias (e.g., Faust and Leeper, 1997). Thus, the identification of policy in small models is suspect because it rests on largely unmotivated zero restrictions on omitted variables. On the other hand, large models require more identifying restrictions than small models, inevitably leading to the use of less credible restrictions. Further, the very size of large models makes it difficult to implement the sort of informal checks on the identification that are an important part of small model work. Thus, if structural conclusions from small or large VAR models are to be persuasive, we need a way to assess the robustness of the results to alterations in questionable identifying restrictions.

The method in this paper takes a particular claim and checks all possible identifications of the VAR for the one that is worst case for the claim, subject to the restriction that the implied economic structure produce *reasonable* responses to policy shocks.

The procedure is most easily introduced by focusing on a particular claim, and I will use Sims's second claim, which is also the subject of the empirical application. I focus on this claim since it is arguably the most important of Sims's four conclusions. The question of whether erratic monetary policy has caused recessions has been a central point of contention in the study of business cycles for decades. It has been at the center of the VAR literature and of the real business cycle literature, which has argued that most variation in output is due to non-monetary factors. In the VAR literature, as elsewhere, parties differ on what result is to be expected. Strongin (1995), for example, considered the result that policy shocks have generated little output variance to be a puzzle.

This paper assesses the question from the Sims-Bernanke-CEE perspective:

> (*) For every reasonable identification of the VAR, the monetary policy shock accounts for a small share of the forecast error variance of output.[1]

For clarity, I want explicitly to concede a weaker claim: Given a sufficiently rich information set, there exists a reasonable identification of the VAR in which the policy shock accounts for a small part of the variance of output. Of course, to reach (*) from the weaker claim, one must rule out the existence of reasonable identifications in which the policy shock accounts for a large share of output variance. The general tenor of Sims and Bernanke's comments is consistent with the strong form of the claim. CEE (1997) make the strong claim explicitly and argue persuasively that supporting that claim is essential if we are to consider the matter decided.

It is not clear how strongly existing work supports (*). In particular, it is difficult to tell from published work how it is that reasonable identifications contradicting (*) have been ruled out. By checking all possible identifications for the one that is worst from the standpoint of the claim, the method of this paper offers a clearer means of support for such claims. If in the worst case the variance share is small, then the claim is supported. If the share is large, then either the identifying information—the characterization of a *reasonable* policy shock—must be sharpened or we must view the issue as unsettled.

The new technique reveals that (*) is not strongly supported by a work-horse six-variable model used in some variation by CEE (1996), Bernanke and Mihov (1995), Uhlig (1997), and LSZ. The claim receives more, though not unequivocal, support in the 13-variable model of LSZ. In both models, there is more support for (*) in a sample spanning 1965:01–1979:09 than in a longer sample of 1960:01–1996:03.

The technique of this paper can be applied quite broadly in VAR work, and has several nice features. All identifying restrictions are stated explicitly, in contrast with conventional use of *informal* restrictions. If the claim is not supported, the approach provides a counterexample and, in doing so, provides a concrete basis for further refinement of the issue. Surprisingly, perhaps, the method does not require that the parameter on which the claim is based be identified. Thus, one can assess whether very minimal commitments regarding the economy are sufficient to support the claim. This is of particular value when using large models in which we might not have sufficient economically credible assumptions to identify the economic quantity of interest. The cost of not identifying the parameter is that the procedure only provides bounds on the parameter of interest. Thus, one must interpret the results carefully: if large variance shares seem likely under the bound, it may simply be because the bound is not very tight.

This work is in the tradition of Cooley and LeRoy (1985) and Leeper and Gordon (1992), who also assess a broad range of identifications. It is most closely related in both motivation and technique to Uhlig (1997), which is discussed below.

The first section reviews why one should bother assessing the robustness of VAR identification. The following sections present a digression on side-stepping identification, and give a strategy for doing so. Next come the application and the conclusions.

Identification in VARs

The identified VAR approach was born of Sims's criticism of the dominant approach to identification at the time. To their credit, participants in the VAR literature have remained close to these roots, paying careful attention to the difficult problem of identification in macroeconomics. In struggling for a credible identification scheme, the VAR literature has given us *informal identifying* restrictions, *partial* identification, *agnostic* identification, *tentative* identification, and *semi-structural* models. All of these approaches involve identifying certain coefficients in the conventional sense laid out by Koopmans (1953). Thus, the labels primarily reflect the self-critical stance taken to identification in this literature. As motivation for the robustness check that is the main purpose of the paper, this section reviews identification approaches in the VAR literature and some criticisms of the approaches.

The standard case

The traditional textbook case of identification begins with the model,

$$\Gamma Y_t = -BX_t + \varepsilon_t, \tag{1}$$

where Y_t $(n \times 1)$ is a vector of endogenous variables and X_t is a vector of exogenous variables and lagged endogenous variables.

The identification problem stems from the fact that if we premultiply the system by a full rank matrix, Q ,

$$Q\Gamma Y_t = -QBX_t + Q\varepsilon_t, \tag{2}$$

$$\tilde{\Gamma} Y_t = -\tilde{B}X_t + \tilde{\varepsilon}_t, \tag{3}$$

we get a system with the same reduced form as (1): $Y_t = -\Gamma^{-1}BX_t + \Gamma^{-1}\varepsilon_t$. In general, however, $\tilde{B}_{ij} \neq B_{ij}$. The data alone cannot help us choose between these values for B_{ij}, and while the data have the same distribution under the two structures, the dynamic effects of a shift in the intercept of the i^{th} equation, for example, will differ in the two cases.

A set of restrictions fully identifies the model if and only if it rules out all but one Q. This requires sufficient restrictions to pin down the n^2 elements of Q; n of the restrictions are normalizations that simply pick the units for the coefficients. In traditional simultaneous equations work, the model is identified exclusively using linear restrictions on the B and Γ coefficients.

Some important terms such as *identification*, *structure*, and *model* have been used in many ways. In this paper, any set of restrictions that picks out a unique structure for each reduced form *identifies* the model.[2] We can always write down arbitrary restrictions that achieve this end, and it is a relatively simple technical matter to resolve whether a given set of restrictions identify some parameter, say, Γ_{ij}. One can further ask whether the assumptions support a given economic interpretation of Γ_{ij} as, e.g., an *interest semi-elasticity of money demand.* Cooley and LeRoy (1985) echoed the Cowles Commission in arguing the answer to this question will generally be negative if the identifying restrictions are arbitrary but may be positive if the assumptions reflect beliefs about the causal mechanism operating in reality. Inferences requiring economically meaningful identifying restrictions are often called *structural inferences*, and (*) clearly involves such inference.

The standard VAR approach

Formal restrictions. VAR identification starts with a version of (1):

$$A_0 Y_t = -\sum_{j=1}^{k} A_j Y_{t-j} + \varepsilon_t \tag{4}$$

$$A(L) Y_t = \varepsilon_t \tag{5}$$

where $A(L) = \sum_{j=0}^{k} A_j L^j$, $L x_t = x_{t-1}$, and $E\varepsilon_t \varepsilon_t' = \Sigma$. The identification problem is just as before.

Those estimating identified VARs impose some linear restrictions on the As, typically in the form of zero restrictions on A_0. However, identified VAR work also always imposes the restriction that the shocks in (4) are orthogonal and imposes the normalization that the shocks have standard deviation one, $E\varepsilon_t \varepsilon_t' = I$. In some work, restrictions are placed on the long-run impulse response (e.g., Blanchard and Quah (1989)), that is, on elements of $C(1)$, where

$$C(L) = \sum_{j=0}^{\infty} C_j L^j = A(L)^{-1}. \tag{6}$$

The primary analytical difference between identification in VARs and more traditional approaches is the use of restrictions that are not linear restrictions on the slope parameters in (4). The orthogonality restrictions and long-run restrictions are examples nonlinear restrictions.

Informal restrictions. LSZ and Sims and Zha (1996a) seem to have been the first to attempt to explain and justify the use of informal restrictions in VARs. They argue that we have prior opinions about the dynamic response of the economy to a money supply shock. For example, short-term interest rates rise and the money stock falls in the short-run in response to a contractionary shock. Because these impulse response restrictions are difficult to impose, one identifies the model using the sort of restrictions discussed above. If the impulse responses do not look right, one then re-specifies the model in some way—either the formal identifying restrictions or the information set might be altered. Thus, the fact that standard VARs predicted that prices smoothly rise following a monetary contraction was declared a *price puzzle*, which was *solved* by adding an index of commodity prices to the model.

Once the informal restrictions have been used to settle on a specification and a set of formal restrictions, Bayesian coverage intervals are often computed. Generally (with the exception of Uhlig, (1997)), these intervals are computed imposing the formal restrictions and ignoring the informal ones. It is not the case that our belief in either the formal or informal restrictions is dogmatic (impervious to evidence). The *ad hoc* use of informal restrictions and the dogmatic application of the formal restrictions primarily reflects practical computational problems with any other course.[3]

Why we need to check the robustness of VAR identification

At policy institutions and elsewhere, structural inferences must be drawn, and they are certain to be drawn using approaches to identification that are fallible. Given this fact, one should at least hope to know the principle weaknesses of the approach used.

Some formal restrictions are only weakly credible. To attain identification, VAR analysts (and others) often impose restrictions that do not reflect strongly held convictions. The bulk of the VAR literature has stressed restrictions on contemporaneous interactions among variables.[4] Thus, it is common to impose that output and final goods prices do not react to money supply shocks within the smallest time period in the analysis, usually a month or quarter (CEE, 1996; Bernanke and Mihov, 1995; LSZ). Further, LSZ assume that policy does not respond to output shocks within the period. We can tell plausible arguments for many of these restrictions, but, as LSZ and CEE (1997) emphasize, we can easily imagine these restrictions not holding. Two brief examples serve to emphasize the point.

On March 14 and 15, 1980, credit controls were announced on the U.S. economy, bringing about the shortest recession in U.S. history. The unemployment rate, which had been unchanged for three months, jumped 0.6 percentage points in April—the second largest change in the post-1950s sample. On May 2, the Fed responded by cutting interest rates (see, e.g., Foldessy, 1980). Thus, over a period of 49 days, a policy was adopted, the real economy reacted, and the policy was altered. Little other than credit controls has been put forward to account for the sharp move in the unemployment rate, and news reports at the time make it clear that the Fed was responding to evidence about the real economy when it responded. Thus, it appears that the economy reacted to policy in less than a month, and the Fed countered in less than a month.[5] This episode clearly involved both large and unique changes in policy. The reactions to smaller policy changes are surely smaller, but are they slower?[6]

It is also common to treat federal funds rate innovations as due to policy decisions and not market forces during periods when Fed operating procedures focussed on the funds rate (e.g., Bernanke and Mihov, 1995; and CEE, 1996). This assumption generates contemporaneous restrictions on a standard VAR. Examination of the daily federal funds rate makes clear that the Fed has never attempted to control the funds rate tightly (Figure 1), and that it has allowed dramatic spikes in the rate at the end of settlement periods and when end-of-year "window dressing" demands for reserves arise (see, e.g., Goodfriend, 1983). Two single-day spikes of 200 basis points in a month can raise the monthly-average funds rate used in VARs by more than the standard deviation of the unpredictable portion of the funds rate. The year-end spike in 1986 led to an 87 basis point rise in the monthly December rate that was immediately reversed (top panel).[7] A more typical example came in July 1996 when the spikes at the beginning and end of the month led to about a 15 basis point monthly rise and fall (bottom panel). These spike-induced changes are largely unpredictable using the standard VAR information set and, hence, will be misclassified as policy-induced shocks to supply rather than to reserve demand shocks under the stated identifying assumptions.

Figure 1 about here

These two examples are little more than anecdotes and are intended only to emphasize that when contemporaneous restrictions are put forward tentatively it is because we have good reason to suspect them. Such suspicions naturally motivate testing the robustness of key conclusions to changes in less-than-fully-credible identifying restrictions.

Informal restrictions and the appearance of circularity. Uhlig (1997) has persuasively argued that the way informal restrictions are used may render the inference procedure circular. At the very least, the reader of VAR work will often find it difficult to tell if the procedure is circular. The problem arises because the informal restrictions are not only difficult to impose formally, their role is difficult to document thoroughly. Thus, when presented with results at the end of a paper, it is difficult for the reader to know which features were informally imposed as criteria for an acceptable model and which were freely estimated implications of the identifying restrictions.

LSZ note that the approach might appear to be "data mining," and they echo the arguments of self-confessed data miners Hendry (1995) and Leamer (1978) in responding. These authors all argue that they are merely being explicit about the sort of back-and-forth between data and model that is an essential part of all work with non-experimental data. This defense is unassailable but makes the problem no less vexing. This paper provides an additional tool for managing data mining-related problems in conducting and communicating VAR analysis.

Informal restrictions and confidence intervals. Neither the formal nor the informal restrictions used in VAR work are believed dogmatically. It is probably the case, however, that the most tenable of the informal restrictions are more strongly believed than the least tenable of the formal restrictions. In calculating confidence intervals, however, the formal restrictions are treated as dogmatic, while the informal ones are ignored. Thus, Bayesian coverage intervals for parameters are often computed by repeatedly drawing from the posterior for the parameters implied by some reference prior, the data, and the formal restrictions.

If some aspects of the informal prior are strongly held, this is problematic: any given draw from the posterior under the formal restrictions need not satisfy the informal restrictions; such draws should be assigned small posterior mass. In abstract, one knows little about the relation between the intervals arising from imposing the more dogmatic among the informal restrictions and those that do not. There are reasons to believe that the practical importance of this problem may be small.[8] In any case, the procedure below provides an imperfect check on this point by allowing a general loosening of restrictions and by allowing imposition of both formal and informal restrictions.

The curse of dimensionality. Most, if not all, of the arguments above have been appreciated in the literature, and they have motivated attempts to test the sensitivity of results to changes in the identification. If one limits consideration to fully recursive structures for the economy, there are only a finite number, and one can look at all of them and see if answers to key questions are sensitive to which is chosen. Work of this type is common, and as Cooley and LeRoy (1985) argued, the results tend to vary across recursive structures.

Moving beyond recursive systems, the set of possible identifications goes from finite to uncountable. In a bi-variate VAR under the assumption of orthogonal shocks, only one further identifying restriction is required, and one can still consider all possible identifications of the VAR. King and Watson (1992) show how to do this: one plots the outcome for the statistic of interest against a one-dimensional variable indexing the identification of the VAR.

As model size increases, however, the curse of dimensionality renders this process unwieldy. In a three-variable VAR, there are three free dimensions in the identification,[9] and it is already impossible to plot the parameter of interest against an index of the identification. Still, in models of three or four variables, one might be able informally to check all rotations visually by recombining the columns of a standard graph of the $n \times n$ impulse response function (e.g., Figure 4). LSZ carry out this process of robustness by ocular rotation. For models of six variables, this is extremely difficult, and in larger models, it may be impossible.

Thus, the final difficulty with the current approach to supporting claims like (*) is that in models of more than a few variables, the class of possible reasonable identifications is large and is difficult to search effectively.

A digression: side-stepping identification

We need a way to check that (*) is implied by every rotation of the VAR that is consistent with firmly held beliefs. We would like the method to be applicable in both small and large models. Our firm commitments may, however, be insufficient to identify the variance share in (*); thus, it would be best if we could test the claim even when the statistic upon which it is based is not identified. LSZ put forward the basic idea (the notation in the quote corresponds with (6)):

> [The assumption that the structural shocks are orthogonal] means that, in some circumstances, conclusions about model behavior are less dependent on identifying assumptions about A than in [traditional simultaneous equations models].... One might find that the rows of $C(L)$ that correspond to prices and interest rates (the first and second rows, say) mostly show prices and interest rates moving in the same direction,

> when they show any substantial movement.... One might expect that the response to a monetary policy shock should show the opposite sign pattern.... Then one could conclude that monetary policy disturbances cannot account for much of the observed variation in prices and interest rates, regardless of the specific identifying assumptions.

The most concise statement of the reasoning is that when using nonlinear restrictions, data may be informative about a parameter that is not identified. Because this notion may be unfamiliar, I provide a brief digression to clarify the issue.

Suppose that we have a reduced form parameterized by $\beta \in \mathcal{R}^k$ and an associated structural model parameterized by $\theta \in \mathcal{R}^k$. Each structure is associated with one reduced form so that there is a function $\beta = h(\theta)$ giving the reduced form parameter for each structure. We wish to estimate $g(\theta)$.

The identification problems stem from the fact that there may be more than one structural parameter associated with a single β. The standard definition states that the restriction $\theta \in \Theta_R$ *identifies* $g(\theta)$ if and only if $\theta_1, \theta_2 \in \Theta_R$ and $h(\theta_1) = h(\theta_2) = \beta$ implies $g(\theta_1) = g(\theta_2)$.[10] If two θs are consistent with the restriction and share the same reduced form, they must give the same value for the parameter of interest.

Now define that the restriction $\theta \in \Theta_R$ is *informative* about $g(\theta)$ if and only if $\theta_1, \theta_2 \in \Theta_R$ and $h(\theta_1) = h(\theta_2) = \beta$ implies

$$g(\theta_1), g(\theta_2) \in G_R(\beta) \tag{7}$$

where $G_R(\beta)$ is a proper subset of the unrestricted parameter space. If two structural parameters satisfy the restriction and have the same reduced form, then we know that $g(\theta)$ falls in a restricted subset of its parameter space. Identifying restrictions are the special case of informative restrictions when for each β, $G_R(\beta)$ has only one element.

This notion may be unfamiliar—I have never seen it stated—because the bulk of discussion of identification treats the case of identifying slope parameters of a linear model using linear restrictions on those parameters. In this case, restrictions are informative if and only if they are identifying (see the Appendix).

A simple example of the distinction between informative and identifying restrictions is the restriction that structural shocks are orthogonal. This assumption is informative about the share of the forecast error variance of output accounted for by the policy shock in a standard monetary VAR. To see this, note that without any restrictions, the variance share might fall anywhere between zero and one. Given a VAR under the assumption of orthogonal shocks, one can compute the maximum variance share attributable to *any* single shock. It is well-known (and see below) that this maximum share is given by the maximum eigen value of a matrix formed from the reduced form coefficients. Thus, the orthogonality assumption produces a bound on the parameter of interest without identifying it.

Assessing robustness to changes in identification

The goal is to see whether firmly held beliefs about the economy are sufficiently informative to support claims like (*). So long as we maintain the orthogonality assumption, we can proceed using the following machinery.

The easiest form of the VAR to work with is what I will call a generic orthonormal (GO) form, which is simply a transformation of the moving-average representation in which the variance covariance matrix is the identity matrix:

$$Y_t = C(L)\varepsilon_t,$$

where $E\varepsilon_t\varepsilon_t' = I$. Any recursive ordering gives a GO form. Under the assumption of orthogonal shocks, the impulse response of each variable to any shock in any identification of the VAR is given by the coefficients of the $(n \times 1)$ vector of lag polynomials:

$$C(L)\alpha,$$

for some α satisfying $\alpha'\alpha = 1$.[11] Every identification of the full set of impulse responses to all shocks is similarly of the form:

$$Y_t = [C(L)D][D'\varepsilon_t] = D(L)\nu_t, \tag{10}$$

where D is orthonormal: $D'D = I$.[12] The forecast error variance share of variable y at horizon h attributed to the shock defined by α is,

$$V_{yh}(\alpha) = \alpha' V_{yh} \alpha, \tag{11}$$

where V_{yh} is a positive definite matrix depending on the GO form parameters.[13]

Now consider how to impose restrictions on the impulse response to the shock defined by α. Suppose we have a VAR of interest rates, money, prices, and output (r, m, p, y). A sign restriction on the impulse response of variable, m, at lag i is of the form:

$$C_{\hat{m}i} \alpha \geq 0, \tag{12}$$

where $C_{\hat{m}i}$ is the row of C_i corresponding to m. Thus, to impose that a contractionary policy shock raises r and lowers m, p and y on impact, one writes,

$$C_R \alpha \geq 0, \tag{13}$$

where $\geq$ means each element of the vector satisfies the restriction and

$$C_R = \begin{bmatrix} C_{\hat{r}0} \\ -C_{\hat{m}0} \\ -C_{\hat{p}0} \\ -C_{\hat{y}0} \end{bmatrix}. \tag{14}$$

Restrictions on whether the impulse response function is rising or falling between particular lag horizons are also of this form.[14]

These results provide the basis of the approach. If important elements of our prior commitments about the economy can be cast as $C\alpha \geq 0$, then one can check all identifications of the VAR consistent with (13) to see if any are inconsistent with (*). Of course, one only need check the worst case, that is, the largest variance share. This suggests solving the following problem:

$$\bar{V}_{yh} = \max_{\alpha} \alpha' V_{yh} \alpha \tag{15}$$

subject to:

$$\begin{aligned} \alpha'\alpha &= 1 & (16) \\ C_R\alpha &\geq 0. & (17) \end{aligned}$$

Given the α^* solving the problem, the impulse response to the associated shock is $d^*(L) = C(L)\alpha^*$.

Without the second constraint, $\bar{V}$ would be the largest eigen value of V_{yh}. Without the first constraint, the problem has the form of quadratic programming. The full problem can be solved by computing a large but finite set of eigen value problems; thus, no general search algorithm is required (see the Appendix).

An algorithm for examining all relevant rotations of the VAR

Suppose initially that we are interested in point estimates only and ignore the more subtle questions raised by interval estimates. The following algorithm provides a way formally to assess claims like (*).

1 Impose some minimal set of restrictions regarding what a policy shock does.

2 Calculate $\bar{V}_{yh}$.

3 If $\bar{V}_{yh}$ is small, stop: the claim is confirmed. Otherwise,

4 Look at $d^*(L)$ to see if the shock with the large variance share looks reasonable. If it does, stop: the claim is contradicted. Otherwise,

5 Add a restriction ruling out whatever is unreasonable and return to 2.

This simple algorithm has several attractive features. It provides a formal way to check all possible identifications of the VAR, even in relatively large systems. When the claim is falsified, it is falsified constructively: a counter-example is provided. In producing potential counter-examples, the algorithm is likely to elicit the prior from believers in the claim. The method also provides a way to discover which restrictions are most informative about the variance share. Imposing some restrictions may not lower the bound, while others may lower it sharply.

The process is no substitute for identifying a full VAR. For example, it does not end with an estimate of the impulse response to a policy shock; the final $d^*(L)$

has no special claim to attention. Further, this discussion presumes that all the restrictions we would want to impose are of the form (13). While this need not be the case, a surprisingly large portion can be cast in this way. Even so, one might exhaust all of these in the search.

Confidence bounds

We might find that there is no rotation of the point estimate of the reduced form that gives a large variance share. There still might be reduced form parameters that are quite likely from the standpoint of the data and that do admit a large variance shares under the restrictions.

Following the Bayesian approach common in this literature, one way to take account of uncertainty regarding the reduced form is to posit a reference prior for the parameters of the reduced form, then evaluate the posterior distribution of $\bar{V}_{yh}$ under the chosen inequality restrictions. For reference priors such as the standard "RATS prior" (see, e.g., Uhlig (1997)) or the more complex Sims and Zha (SZ) prior (1996b), drawing from the posterior is straightforward, and evaluating $\bar{V}_{yh}$ can be completed as before. The result is a posterior for $\bar{V}_{yh}$.

Since $\bar{V}_{yh}$ bounds the parameter of interest, V_{yh}, for every reduced form, we have

$$\mathrm{pr}(V_{yh} > \gamma) < \mathrm{pr}(\bar{V}_{yh} > \gamma), \tag{18}$$

where V_{yh} and $\bar{V}_{yh}$ are a function of the reduced form and the probabilities are evaluated under a common posterior. Thus, the posterior for $\bar{V}$ gives a probability bound for the parameter of interest that is conservative (biased upwards) from the standpoint of evaluating (*).[15] As with any procedure that produces conservative estimates, one would like to know just how conservative they are. This remains an open question for now.

For a fixed reduced form, adding restrictions must either lower $\bar{V}$ or result in the restrictions being inconsistent with the model.[16] In calculating confidence bounds, this latter possibility means that the support of the prior may change with the set of restrictions. For this reason, say, the 66^{th} percentile of the posterior for $\bar{V}$ may rise or fall with added restrictions.

The Uhlig approach

Uhlig's (1997) approach is similar in motivation and implementation to the one here. His method also involves solving an optimization problem to compute an estimate of the variance share under sign restrictions on the impulse response to a policy shock. In Uhlig's case, the restrictions are that the response to a contractionary policy shock has the correct sign for each of the first 48 quarters on output (down), interest rates (up), money (down), and prices (down). These restrictions generally will not all be consistent with the reduced form, and Uhlig's method picks the shock that comes as close as possible to meeting the restrictions under a loss function that penalizes bigger impulse response coefficients of both signs. The penalty on coefficients of the wrong sign is 100 times larger.

While my approach chooses the identification that is worst from the standpoint of (*), Uhlig's picks the one that is best from the standpoint of the restrictions (under the specified loss). From the standpoint of identification, Uhlig's approach identifies the response to a policy shock using dogmatic restrictions (albeit of a novel form), just as is the case with standard approaches. Uhlig's approach shares with mine the absence of informal restrictions and will be attractive to anyone whose prior beliefs about policy are better captured by some loss function like Uhlig's than by traditional restrictions. Those skeptical about aspects of the loss function may still want to check the robustness of the result, and my approach provides one way to do so.

Example: a common 6-variable model

As a first application, I consider the six-variable model of originated by CEE (1996) and Bernanke and Mihov (1995) and used by both LSZ and Uhlig. The model contains output (Y), prices (CPI), commodity prices (PC), nonborrowed reserves (NBR), the federal funds rate (RF), total reserves (TR) and a constant.[17] For comparison with later results, I use the version of the model estimated by LSZ using monthly data from 1960:01 to 1996:03 using six lags (so that initial lags are from 1959). The model is identified by the recursive ordering given in the variable list above, and the fourth shock is the policy shock.

While it is not clear what horizon forecast error variance is of greatest interest, I focus on the 108 month horizon. If we wish to know how much policy shocks contribute to business cycle variation in output, a horizon of several years at minimum is required. The results do not seem to be very sensitive to horizon in this range. The point estimate of $V_{y,108}$ under my replication of LSZ's estimates is 18 percent.[18]

Following the algorithm above, I first calculated a bound on the 108-month forecast error variance share of $\bar{V}_{y,108} = 91$ percent under the restrictions that the contemporaneous effects of a contractionary shock all had the right signs: all variables down except RF which rises. The computational algorithm evaluates many identifications and I saved and inspected the responses to any shock the algorithm found with a share greater than 30 percent.[19] These shocks generally showed output moving very sharply on impact, much more sharply than the other variables, and the interest rate effect sometimes changed signs very quickly. Thus, I imposed that the interest rate effect was positive in the third month after the shock and that the contemporaneous growth rate effect on output was no more than one-quarter of the effect on the interest rate. Neither of these assumptions is uncontroversial, of course.

The bound was lowered to 63 percent, and the algorithm found two identifications worth discussing. I call these identifications A and B; they give variance shares of 36 percent and 61 percent, respectively. The character of the response of the system to these shocks is very similar to that of LSZ's point estimates (Figures 2 and 3). The primary differences other than the larger output effects are that output moves a bit contemporaneously in the counter-examples, and in counter-example B, total reserves do not change contemporaneously.

Figure 2 about here

Figure 3 about here

Are these cases reasonable? One might suppose that total reserves should fall contemporaneously in response to a contractionary money shock, at least during the contemporaneous reserve accounting period. It is unclear whether the change might be near zero, however. Strongin (1995) argues for treating total reserves as fixed in the short-run.[20] The small contemporaneous output effect seems difficult to rule out *a priori*.

Of course, part of the informal check usually applied is an assessment of how policy variables like NBR and RF react to other shocks. Since the point estimates are from a recursive identification, it is easy to substitute a counter-example policy shock as the fourth shock (as it is in the recursive identification) and choose the other shocks to conform as closely as possible to the recursive model.[21]

The three models have impulse responses that look very much alike except, perhaps, in the response to the 5^{th} shock (Figures 4 and 5). The fifth shock in the recursive ordering raises output, lowers prices, raises commodity prices, and lowers reserves and interest rates. It is not clear *a priori* that this makes better sense than the 5^{th} shock in case A. As for case B, the fifth shock looks like a permanent jump in the level of reserves that raises prices, but does little else. This is similar to the third shock in LSZ's 13 variable model, and also seems difficult to rule out *a priori*.

Figure 4 about here

Figure 5 about here

The exercise illustrates that fairly subtle changes in the identification can have substantial effects on the apparent validity of claims like (*). While results for the recursive identification establish that there exist identifications of reduced form point estimates giving reasonable results and in which the variance share is small, such results offer little support for (*). More generally, it is clear that the method used here gives us something concrete to talk about. We know what restrictions generate the case B result; we know that there are no other identifications with much worse results for (*) under these restrictions. Either we rule out the examples like cases A and B *a priori*, or our prior commitments about the economy in conjunction with the estimated reduced form do not support (*).

Variations on a 6-variable and a 13-variable model

Description of the 13-variable model

The variables in the 13-variable model are the 3-month Treasury bill rate ($R3$), TR, M1 ($M1$), CPI, Y, the unemployment rate (U), non-residential investment (INR), residential investment (IR), consumption (C), the 10-year Treasury constant maturity rate ($R10$), the Standard and Poor's 500 index (S), PC, and a trade-weighted index of the value of the dollar (DOL). Once again, the model has six lags and a constant. LSZ use a non-recursive identification, but the principle policy shock produced is very close to what would come from the first shock in a recursive ordering as the variables are listed here.

Three versions of each model

I consider three variations on both the 6-variable and 13-variable models. In particular, I assess (*) in the 1960:01-1996:03 sample using both the SZ prior and the RATS *flat* prior. As noted above, a number of authors have claimed results are sensitive to the sample period. Thus, I also assess (*) in the longest of the sample periods studied by Bernanke and Mihov (1995): 1965:01-1979:09. These short-sample estimates use only the SZ prior. In the shorter sample especially, I would worry about the number of free parameters under the RATS prior; the variance reduction aspects in the SZ prior seem most important in this case.

The restrictions

For both models, I always impose that the policy shock has the *right* sign on impact. For the 6-variable model, these signs were given above. For the thirteen variable model, the restrictions are that interest rates and the dollar rise; the other variables fall. These are consistent with what LSZ find to be reasonable; LSZ and CEE (1997) argue that policy shocks should have such effects. Further, I always assume that the impulse response of the short-term interest rate is positive in the third month after the shock.

I investigate combinations of 4 other restrictions:

1 The short-term interest rate effect is positive in the 9^{th} month.

2 The impulse response of the CPI is negative in the 60^{th} month,

3 The impulse response of Y is no larger in the 108^{th} month than in the 60^{th} month,

4 a) The contemporaneous Y effect in the 6-variable model is less than 1/4 of the contemporaneous RF effect.

b) The contemporaneous Y effect in the 13-variable model is less than 1/2 the effect on S.

Various authors have found that the interest rate response in some models disappears in just a few months, which is inconsistent with their views of the persistence of the response of policy to policy shocks. Restriction 1 will help assess whether imposing this view is informative about (*). Restriction 2 rules out policy shocks in which prices smoothly rise, the classic price puzzle result. In both of the models, the contemporaneous restrictions alone sometimes allow shocks that appear permanently to alter the growth rate of output. While one can make an argument for this result, 3 rules out (some of) such shocks.[22] Finally, 4 rules out shocks for which output moves contemporaneously more than certain financial variables that might be thought to be quicker moving.

While none of these restrictions is uncontroversial,[23] they each seem to be the sort of thing that might have been imposed informally in conventional VAR work. Further, not all reasonable restrictions have been imposed.[24]

The 6-variable model

In the 6-variable model, the previous section displayed an identification in which the policy shock accounts for 61 percent of the variance of output. This shock satisfies all the mandatory and optional restrictions. It is still of some interest to see what the posterior bound for the variance share is under various combinations of restrictions and for the two versions of the model not yet discussed—the shorter sample and the RATS prior (Table 1).

The results for the full sample under the SZ prior illustrate an important point. The percentiles for $\bar{V}$ are much smaller when the first restriction (interest rate effect positive for 9 months) holds than in any other case. Dropping this one restriction

Table 1: Posterior for $\bar{V}_{y,108}$ in the 6-variable model

prior:	SZ prior		SZ prior		RATS prior	
sample:	60:01–96:03		65:01–79:09		60:01–96:03	
restriction	50^{th}	66^{th}	50^{th}	66^{th}	50^{th}	66^{th}
1111	42	52	12	22	41	51
0111	71	79	36	53	46	57
1011	42	52	23	31	41	51
0011	71	79	42	56	46	57
1101	48	57	27	38	57	64
0101	77	83	66	75	69	75
1001	48	57	36	46	57	64
0001	77	83	67	76	69	75
1110	42	52	13	23	45	54
0110	74	82	40	59	54	65
1010	42	52	23	32	45	54
0010	74	82	45	62	54	65
1100	48	57	28	40	59	66
0100	80	86	74	84	76	81
1000	48	57	37	47	59	66
0000	80	86	75	84	76	81

Notes: The priors are described in the text. 50^{th} and 66^{th} stand for the obvious percentile of the posterior distribution of the bound on the forecast error variance share in output (in percent, at horizon 108 months) attributed to any shock meeting the specified restrictions. In all cases, the mandatory restrictions apply. The column labeled *restriction* is a four digit number indicating which of the four optional restrictions are imposed. A one in the j^{th} digit (numbering from left to right) indicates that the j^{th} restriction is imposed; a zero means it was not imposed. The three sets of 50^{th} and 66^{th} columns are from three independent Monte Carlo experiments evaluating the posterior; each pair of columns is from the same experiment; each experiment involved 2000 draws from the posterior. For the "1111" restrictions 27, 156, and 13 draws, respectively, were inconsistent with the restrictions.

from the full set raises the 66^{th} percentile by over 30 percentage points. This highlights the fact that the approach allows us to discover just which restrictions are most informative in the sample at hand, and, hence, we know which restrictions should be scrutinized. If one strongly believes that policy shocks have a persistent effect on interest rates, ones posterior bound for the variance share will be substantially lower. Since I am skeptical regarding the persistence of the liquidity effect on interest rates, my own interpretation of these results is that the full sample estimates for the 6-variable model do not provide much support for (*).

For the full-sample estimates under the two priors, the results show that the 66^{th} percentile is always above 50 percent; the 50^{th} percentile is generally above 50 percent and always above 40 percent.[25] The sample period and choice of prior seem to affect the results. The full sample, RATS prior estimates are very similar to those for the SZ prior when the first optional restriction is imposed, but they do not rise as sharply when this restriction is removed. The results for the shorter sample are generally much more favorable to (*).

Some economists may have a strong belief that policy shocks have no contemporaneous effect on prices or output. Thus, I imposed those exact restrictions and followed the standard algorithm until I found the policy shock in Figure 6, for which the variance share is 24 percent. This shock looks reasonable by the standards applied throughout, and while I cannot guarantee that this is the largest variance share attributable to a reasonable shock with the two contemporaneous zero effects, it appears that imposing such zeros does lower the bound in the point estimates considerably.

Figure 6 about here

The search revealed another interesting feature of the data. Under the two zero restrictions, one could push the bound up to well over 30 percent with responses of essentially the same shape as those in Figure 6. The primary difference is that the shape of the output effects *and* the commodity price effects are magnified. Indeed, for all the policy shocks displayed, if the share is higher than in the recursive model, the commodity price effect is exaggerated. Of course, the effect in the recursive identification is already quite large, and it would be difficult to rule out a slightly larger effect *a priori*. Despite this fact, further investigation of the role of commodity prices in this model is warranted: the variable solves the price puzzle, but reacts far stronger and with more persistence than any other variable.

The 13-variable model

One might suppose that if minimal restrictions do not tightly bound the variance share in a small model, then they are unlikely to do so in a large model.[26] The intuition is that in a large model the algorithm has many more shocks to combine in creating a reasonable policy shock that has a large variance share. The alternative intuition, of course, is that with more variables controlled for, the variance share due to policy may be more sharply estimated. This later view receives some support. Under all four optional restrictions, the maximum variance shares at the posterior mode for the reduced form were 39, 7, and 21 percent for the three versions, respectively. Thus, for the shorter sample and under the RATS prior for the full sample, the reduced-form point estimates are not consistent with very large values for $\bar{V}$.

As for the probability bounds (Table 2), for the RATS prior, four combinations of restrictions push the 66^{th} percentile to 33 percent or below. In the short sample, SZ-prior estimates, two combinations are sufficient to do the same.

For the SZ prior in the full sample, the 66^{th} percentiles are all above 50 percent.[27] It is of interest to know why the two priors give such different answers for the full sample; I currently have little to offer on this count. Those who accept the arguments of Gordon and Leeper (1994) and Bernanke and Mihov (1995) would tend to place less emphasis on these estimates that span clear changes in operating procedures. While I have sympathy with this view, Sims (1996) argues in favor of full sample estimates.

Results summary

Overall, the results regarding (*) are mixed: both the choice of reference prior and sample period seem to matter. The short sample generally provides more sup-

Table 2: Posterior for $\bar{V}_{y,108}$ in the 13-variable model

prior:	SZ prior		SZ prior		RATS prior	
sample:	60:01–96:03		65:01–79:09		60:01–96:03	
restriction	50^{th}	66^{th}	50^{th}	66^{th}	50^{th}	66^{th}
1111	49	58	19	30	26	31
0111	54	63	31	44	28	33
1011	50	58	32	42	31	37
0011	55	63	43	52	34	40
1101	56	63	44	54	30	35
0101	63	69	57	66	33	37
1001	56	63	53	61	36	41
0001	63	70	62	70	39	44
1110	49	58	19	30	26	32
0110	55	63	31	44	28	33
1010	50	58	32	42	32	38
0010	55	63	43	52	34	40
1100	56	63	44	54	31	35
0100	63	70	58	67	33	38
1000	56	63	53	61	36	41
0000	63	70	63	70	39	44

Notes: See Table 1. The simulations for the three versions each had 1000 draws from the posterior. For the "1111" restrictions, 1,5, and 2 draws, respectively, were inconsistent with the restrictions.

port for (*). The larger model also provides somewhat stronger support. Further, imposing that the liquidity effect on the short-term interest rate lasts at least through the ninth month after the shock or that the contemporaneous price and output effects are exactly zero seems to lower the bound on the variance share markedly.

It is important to remember that what is being displayed is a posterior for $\bar{V}$, which bounds the variance share of interest. Thus, viewed as probability bounds on the share of interest, these numbers are conservative from the perspective of evaluating (*). Of course, if one were interested in evaluating the smallest possible variance share of output attributable to the policy shock, one could alter the procedure to calculate a lower bound on the minimum share. Some experimentation suggests that this would likely result in a bound near zero. This would be further evidence in favor of the weak form of (*): current VAR models are consistent with the variance share in output of the money shock being trivial.

It is not the case that the few restrictions considered here characterize all we believe about policy shocks. Some nonsystematic experimentation has convinced me that imposing more restrictions would probably lower the bounds. I am currently pursuing some ideas in this regard.

Discussion

We are unlikely ever to have enough uncontroversial restrictions to clearly identify important macroeconomic phenomena such as the role of monetary policy. Thus, given results for a particular identification, we would like to know how strongly the results depend on suspect restrictions. Ultimately, we would like to be able to have confidence in general claims like (*): for every reasonable identification, the variance share of output due to unpredictable shifts in policy is small. Current formal and informal methods for verifying such claims break down quickly as model size increases.

This paper develops and applies a scheme that can work in both large and small models. One can check the validity of the claim under minimal "firm" commitments about the response of the economy to shock. Further, one can determine which restrictions if any are most informative, and such restrictions can be singled out for special scrutiny. Further, the scheme yields a constructive procedure for eliciting aspects of prior beliefs that may support claims like (*). Specifically, the algorithm generates counter-examples to the claim, and one may invite a believer to suggest which prior belief about policy shocks rules out such examples. My limited experience with this algorithm suggests that it is quite useful. I very quickly convinced myself that the 6-variable model offered very little support for (*). Others may disagree, and the algorithm invites them to specify (and perhaps justify) restrictions that rule out the examples I find persuasive.

In the 13-variable model, the approach seems to offer somewhat more support for the claim. This support is not unequivocal and important issues for further consideration include understanding differences in the results under the two reference priors, and more generally, assessing what prior for the variance share is implied by the reference priors.

The results illustrate that certain minimal commitments about the economy may more tightly bound structural estimates in large models than in small. They call into question the view that structural inference in large models is hopeless because we simply do not have enough economically meaningful identifying information to fully identify the economic structure. I hope that they bolster the impetus to study larger models of the sort introduced by LSZ.

Appendix

Data

The data are from LSZ and were kindly provided by Tao Zha. All data are in units of natural logarithms multiplied by 100, except interest rates and unemployment rates, which are stated in percentage points. The following are the definitions of the variables as provided in LSZ.

C Personal consumption expenditures, seasonally adjusted, billions of chain 1992 dollars.

CPI Consumer price index for all urban consumers, total, seasonally adjusted.

DOL Trade-weighted value of the U.S. dollar, Atlanta Fed index, 1980=100.

INR Real private non-residential fixed investment, seasonally adjusted, billions of chain 1992 dollars; monthly series interpolated using Chow-Lin procedure with monthly data on real value of new construction of privately owned nonresidential industrial structures, total equipment component of industrial production, industrial machinery and equipment component of industrial production, intermediate products and business supplies component of industrial production, manufacturers' shipments to capital goods industries, and manufacturers' shipments of construction materials, supplies, and intermediate products.

IR Real residential fixed investment, seasonally adjusted, billions of chain 1992 dollars; interpolated using Chow-Lin procedure with monthly data on housing starts, construction supplies component of industrial production, manufacturers' shipments of construction materials, supplies, and intermediate products, and real value of new construction of privately owned residential buildings.

$M1$ M1 money stock, seasonally adjusted, billions of dollars.

NBR Non-borrowed reserves plus extended credit, seasonally adjusted billions of dollars.

PC Crude materials component of the producers' price index, seasonally adjusted.

RF Federal funds rate, effective rate, percent per annum.

$R3$ 3-month Treasury bill rate, secondary market, percent per annum.

$R10$ 10-year Treasury bond yield, constant maturity, percent per annum.

S Standard and Poor's 500 composite stock price index, 1943=100.

TR Total reserves, adjusted for breaks due to changes in reserve requirements, seasonally adjusted, billions of dollars.

U Civilian unemployment rate, seasonally adjusted, percent.

Y Real gross domestic product, seasonally adjusted, billions of chain 1992 dollars; interpolated from national income and product accounts quarterly series using Chow-Lin procedure with monthly data on total industrial production, civilian employment 16 years or older, retail sales deflated by consumer prices, real personal consumption expenditures, and the National Association of Purchasing Managers Composite Index.

Informative implies identification in the linear case

Take the structure,

$$\Gamma_0 Y_t = -B_0 X_t + \varepsilon_t. \tag{19}$$

Define $\Pi_0 = [B_0 : \Gamma_0]$ and $\pi_0 = \mathrm{vec}(\Pi_0')$, where vec(.) means the vector formed by stacking the columns of the argument matrix. Suppose the parameter of interest is the k^{th} element of π_0, π_{0k}, and that all the identifying restrictions are linear restrictions on π:

$$\Phi \pi_0 = \phi. \tag{20}$$

The standard result (e.g., Rothenberg, 1985) is that π_{0k} is identified if and only if the rank of $\Xi = \Phi(I_n \otimes \Pi_0')$ is the same as the rank of,

$$\Xi^* = \begin{bmatrix} \Phi \\ i_k' \end{bmatrix} (I_n \otimes \Pi_0'), \tag{21}$$

where i_k is a conformable vector with 1 in the k^{th} element and zero otherwise, I_n is an $n \times n$ identity matrix, and $\otimes$ denotes Kronecker product.

Now the claim in the text can be verified. If restrictions are identifying they are clearly informative. What must be shown is that, given the structure (19), if the restrictions (20) are not identifying then for any τ there exists a representation of the model with parameter $\Pi_1 = Q\Pi_0$, consistent with the restrictions and with $\pi_{1k} = \tau$.

Add the restriction that $\pi_k = \tau$ to (20):

$$\begin{bmatrix} \Phi \\ i_k \end{bmatrix} \pi_1 = \begin{bmatrix} \phi \\ \tau \end{bmatrix}. \tag{22}$$

The vector $\pi_1 = \text{vec}((Q\Pi_0)') = (I_n \otimes \Pi_0')\text{vec}(Q')$. Thus, there exists a Q giving a π_1 satisfying (22) if

$$\begin{bmatrix} \Phi \\ i_k \end{bmatrix} (I_n \otimes \Pi_0')\text{vec}(Q') = \Xi^* \text{vec}(Q') = \begin{bmatrix} \phi \\ \tau \end{bmatrix} \tag{23}$$

can be solved for Q. Since the rank of Ξ^* is greater than that of Ξ (and the matrices have the same column dimension) we know that the number of rows of Ξ^* is less than or equal to its rank, and the equation has a solution.

Maximizing the variance share

The maximization problem is stated in (15)-(17). I assume that any matrix of n rows taken from C_R is of full rank; thus, between zero and $n-1$ of the constraints $C_R\alpha \geq 0$ must hold with equality at the solution. Define C as the matrix made up of the rows of C_R representing constraints that hold with equality at the solution. By Kuhn-Tucker theory, we know that the solution satisfies the first order conditions,

$$\begin{aligned} V\alpha^* - \lambda^*\alpha^* - C'\mu^* &= 0 && (24) \\ C\alpha^* &= 0 && (25) \\ {\alpha^*}'\alpha^* &= 1 && (26) \end{aligned}$$

for positive Lagrange multipliers λ^* and μ^*. (I drop the yh subscript on V for simplicity.)

The following shows that the $\bar{V}$ and α^* that solve the problem are the maximum eigen value and associated eigen vector of $(I-P)V$, where $P = C'(CC')^{-}C$. The superscript "$-$" indicates the generalized inverse; if C has no rows, P is a matrix of zeros.

Pre-multiplying the initial first order condition by $I-P$ gives,

$$(I-P)V\alpha^* - \lambda^*\alpha^* = 0, \tag{27}$$

which is satisfied by a unit-length α^* only when α^* is an eigen vector of $(I-P)V$. What remains is to show that α^* is the eigen vector associated with the largest eigen value.

As Rao (1964) notes, the eigen values of $(I-P)V$ correspond to the eigen values of $V^{1/2}(I-C'(CC')^{-}C)V^{1/2'}$ where $V^{1/2'}V^{1/2} = V$. This is useful computationally, since symmetric eigen value problems are well-understood. Further, $V^{1/2}(I-C'(CC')^{-}C)V^{1/2'}$ is positive semi-definite, since $(I-P)$ is idempotent and $V^{1/2}$ is full rank. Thus, $(I-P)V$ is positive semi-definite and its eigen vectors can be chosen to be mutually orthogonal.

Suppose, contrary to the desired result, that α^* is not the eigen vector associated with the largest eigen value. We can write α^* as:

$$\alpha^* = \omega\alpha_1 + (1-\omega^2)^{1/2}\tilde{\alpha} \tag{28}$$

where $\tilde{\alpha}$ and α_1 satisfy (27), are mutually orthonormal, and α_1 is associated with the largest eigen value. Parameterize α^* as,

$$\alpha^*(\delta) = (1+\delta)\omega\alpha_1 + (1-(1+\delta)^2\omega^2)^{1/2}\tilde{\alpha} \tag{29}$$

which is α^* for $\delta = 0$ and satisfies the unit length and equality restrictions for small δ. Since the eigen vectors are orthogonal, the value of the criterion function can be written:

$$\alpha^*(\delta)'V\alpha^*(\delta) = (1+\delta)^2\omega^2\alpha_1'V\alpha_1 + (1-(1+\delta)^2\omega^2)\tilde{\alpha}'V\tilde{\alpha} \tag{30}$$

The derivative of this expression with respect to δ is $2(1+\delta)\omega^2(\alpha_1' V \alpha_1 - \tilde{\alpha}' V \tilde{\alpha})$. This is positive for small δ since $\alpha_1' V \alpha_1$ maximizes the quadratic form under the unit length and C restrictions. Finally, since all the restrictions that hold with equality at $\alpha^*(0)$ also hold with equality for small δ, and since $\alpha^*(0)$ satisfies the full set of inequality restrictions, it follows by continuity that there must be a small $\delta > 0$ such that all restrictions are satisfied at $\alpha(\delta)$. This proves the desired result.

Given this result, we can calculate the solution to the full problem in the following way: Compute the maximum eigen value, λ and associated vector α of $(I-P)V$ for *all possible* P matrices—that is, all possible sets of constraints that might hold with equality. The solution is the largest λ such that the associated eigen vector satisfies $C_R \alpha \geq 0$. If none of these potential solutions satisfies the full set of inequality constraints, then the constraints are inconsistent.

If one follows the algorithm described above, one must solve the eigen value problem for $\sum_{i=0}^{M} \frac{R!}{i!(R-i)!}$ cases, where M is the minimum of $n-1$ and R, the number of rows in C_R. For 20 restrictions in a 13 variable model, the number is near 1 million. While this calculation is feasible, a quicker approach would be useful for calculating confidence intervals, which involves doing the maximization for each draw from the posterior. I am investigating various approaches to speeding this calculation.

NOTES

1 Actually, we might like to verify the that (*) holds *for every sufficiently rich information set* used in the VAR. Except for an *ad hoc* examination of some alternative information sets, I do not take up this stronger claim in this paper.

2 In context, I treat the following phrases as synonymous: an "identification of", a "structure consistent with" and a "model consistent with" a given reduced form.

3 For example, LSZ experiment with non-dogmatic application of formal restrictions, but do not report results due to computational problems.

4 The primary alternative has been long-run restrictions. Faust and Leeper [1997] discuss the difficulties with this approach.

5 For a discussion, see the *Economic Report of the President*, 1981.

6 Of course, if part of the reason for delay is due to signal extraction problems, then the reactions to large shocks (which are easily discerned) would be faster. Under the Gaussian likelihoods used in most empirical analysis, however, large innovations are the most informative, and the identification is most suspect at the time of large innovations.

7 In the 6-variable model below, this generates a 68 basis point innovation. Sims [1996] reports a funds rate regression with a standard error of about 14 basis points.

8 Chris Sims argued this in editorial comments on this paper.

9 Three restrictions are required after orthogonality of the errors.

10 Rothenberg [1977] gives a treatment of identification in a general setting.

11 The only restriction is that the shock have variance one, requiring $\alpha'\alpha = 1$.

12 If the shocks of the transformed model are to be orthogonal with unit variance, it must be that $D'ID = I$.

13

$$V_{yh}(\alpha) = \frac{\sum_{i=0}^{h}(C_{\hat{y}i}\alpha)^2}{\sigma_{yh}^2} \tag{31}$$

$$= \alpha' \left[\frac{\sum_{i=0}^{h} C'_{\hat{y}i}C_{\hat{y}i}}{\sigma_{yh}^2}\right]\alpha \tag{32}$$

where σ_{yh}^2 is the full forecast error variance and $C_{\hat{y}i}$ is the row of C_i corresponding to y.

14 The restriction that the effect on output grows from period zero to one is imposed by adding a row to C_R composed of $-C_{\hat{y}1} + C_{\hat{y}0}$.

15 From a Bayesian perspective, we can be more explicit about the way in which the statments are conservative. Specifically, having calculated the posterior probability, $\text{pr}(\bar{V}_{yh} > \gamma) < p$, one can conclude that

$$\text{pr}(V_{yh} > \gamma) \leq p \tag{33}$$

(given the data) under any prior that (i) shares the reference prior as the marginal prior for the reduced form parameters and that (ii) is such that any inequality restrictions imposed in calculating $\bar{V}$ hold with probability one.

16 If there are sufficiently many restrictions, then some draws from the posterior will be inconsistent with the restrictions. Since I am treating the restrictions as dogmatic, these draws are abandoned. Thus, the reference prior is interpreted as proportional to the prior density conditional on the reduced form and restrictions being consistent. The ratio of the number of draws consistent with the restrictions and those inconsistent is also of interest. This quantity is the posterior odds ratio in favor of the restrictions when viewing the reference prior as an unconditional prior. It may make sense to check this value, for if the posterior odds are too low, one might wish to consider whether the

reference prior is appropriate. Alternatively, the data may not be supportive of the shape restrictions.

17 For details on the data, see the Appendix. The data are from LSZ and were kindly provided by Tao Zha.

18 As in LSZ, the point estimate is the posterior mode under the reference prior documented in Sims and Zha [1996b]. This prior is a modification of the "Minnesota prior" that favors cointegration and allows imposition of identifying restrictions.

19 This is a generalization of step 4 above. It is useful to look at any extreme cases one comes across during the process, rather than just the worst case.

20 If businesses find it more profitable to exercise fixed-rate loan commitments when rates rise, then there is a short-run upward pressure on total reserves that the Fed might partially accomodate.

21 There are several ways to do this, but a simple way comes from viewing the identified recursive form under the ordering given above to be the GO form. The recursive identification can be viewed as being completed by choosing $D = I$ in (10). The α for, say, case A gives a fourth column for a new D. I pick the remaining columns from 1 through 6 in turn so that the sum of the squared off diagonal elements is as small as possible and so that the columns are mutually orthogonal. This procedure forces the earlier numbered shocks to look most like the recursive ordering, and later numbered ones have less flexibility to do so.

22 The restriction is actually that the 108-month response minus the 60-month response is positive, which rules out shocks for which the effect is negative at both points and larger in absolute terms at 108 months.

23 There is no reason to suppose that the price effect will be negative for 60

months unless the restrictive policy is sustained. The effect should, of course, be positive if restrictive policy shocks lead to future expansionary policy.

24 For example, the restrictions do not require that the effect of a monetary contraction on output is negative after the initial month. In some VARs, paradoxically large positive effects are found. This did not seem to be the case in the work reported.

25 For the full sample with the SZ prior, the 66^{th} percentile under all the restrictions is lower than 63 percent, the figure for $\bar{V}$ obtained from the posterior mode for the reduced form. Since the variance share is a nonlinear (and non-monotonic) transform of the reduced form parameters, there is no clear relation between the posterior mode for $\bar{V}$ and the value of $\bar{V}$ at the posterior mode for the reduced form parameters. Uhlig [1997] finds similar phenomena.

26 Notice that $\bar{V}$ is a forecast error variance *share* and need not rise or fall as model size increases. Indeed for several comparable cells under minimal restrictions in the two tables, the 66^{th} percentile rises in going to the larger model.

27 If one ignores the first optional restriction, which I view as most suspect, the full sample, SZ prior results are more favorable to (*) for the 13-variable model than are those for the smaller model.

References

Bagliano, F. and Favero, C.

(1997) Measuring Monetary Policy with VAR Models: an Evaluation. Manuscript prepared for the 20^{th} annual *International Seminar on Macroeconomics*, Gerzensee, Switzerland.

Bernanke, B.

(1996) Comment on "What Does Monetary Policy Do." *Brookings Papers on Economic Activity*, 2:69–73.

Bernanke, B., and Gertler, M.

(1995) Inside the Black Box: the Credit Channel of Monetary Policy Transmission. *Journal of Economic Perspectives*, 9:27–48.

Bernanke, B.S. and Mihov, I.

(1995) Measuring Monetary Policy, NBER Working Paper No. 5145.

Blanchard, O., and Quah, D.

(1989) The Dynamic Effects of Aggregate Demand and Supply Disturbances. *American Economic Review*, 79:655–73.

Christiano, L.J., Eichenbaum, M., and Evans, C.

(1996) The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds. *Review of Economics and Statistics*, 78:16–34.

________, ________, and ________

(1997) Monetary Policy Shocks: What Have We Learned and to What End? Manuscript, Federal Reserve Bank of Chicago.

Cooley, T. and LeRoy, S.

(1985) Atheoretical Macroeconomics: A Critique. *Journal of Monetary Economics*, 16:283–308.

Economic Report of the President

(1981) U.S. Government Printing Office.

Faust, J., and Leeper, E.

(1997) When Do Long-Run Identifying Restrictions Give Reliable Results? *Journal of Business Economics and Statistics*, 15:345–354.

Foldessy, E.

(1980) Fed Relaxes Money-Market Credit Curbs as Pressures Mount for End to Restraints, *Wall Street Journal*, May 5, 3.

Goodfriend, M.

(1983) Discount Window Borrowing, Monetary Policy, and the Post-October 1979 Operating Procedures. *Journal of Monetary Economics*, 12:345–356.

Gordon, D. and Leeper, E.

(1994) The Dynamic Impacts of Monetary Policy: an Exercise in Tentative Identification. *Journal of Political Economy*, 102 1228–47.

Hendry, D.

(1995) *Dynamic Econometrics.* New York:Oxford University Press.

King, R. and Watson, M.

(1992) Testing Long-Run Neutrality. Manuscript, Northwestern University.

Koopmans, T.

(1953) Identification Problems in Economic Model Construction. *Studies in Econometric Method*, eds. W.C. Hood and T.C. Koopmans. New York: Wiley.

Leamer, E.

(1978) *Specification Searches: Ad hoc Inferences with Nonexperimental Data.* New York:Wiley.

Leeper, E. and Gordon, D.

(1992) In Search of the Liquidity Effect. *Journal of Monetary Economics*, 29:341–69.

________, Sims, C. and Zha, T.

(1996) What Does Monetary Policy Do? *Brookings Papers on Economic Activity*, 2:1–78.

Pagan, A., and Robertson, J.

(1994) Resolving the Liquidity Effect. Manuscript, University of Rochester.

Rao, C.

(1964) The Use and Interpretation of Principal Component Analysis in Applied Research. *Sankhya*, series A, 26:329–58.

Rothenberg, T.

(1977) Identification in Parametric Models. *Econometrica*, 39:577–91.

____

(1985) Notes on Identification. Manuscript, Berkeley.

Rudebusch, G.

(1997) Do Measures of Monetary Policy in a VAR Make Sense? Manuscript, Federal Reserve Bank of San Francisco.

Sims, C.

(1980) Macroeconomics and Reality. *Econometrica*, 48:1–48.

____

(1992) Interpreting the Macroeconomic Time Series Facts: the Effects of Monetary Policy. *European Economic Review*, 36:975–1011.

____

(1996) Comment on Glenn Rudebusch's "Do Measures of Monetary Policy in a VAR Make Sense?" Manuscript, Yale University.

________and Zha, T.

(1995) Error Bands for Impulse Responses. Manuscript, Yale University.

________and ________

(1996a) Does Monetary Policy Generate Recessions? Manuscript, Yale University.

________and ________

(1996b) Bayesian Methods for Dynamic Multivariate Models. Manuscript, Yale University.

Strongin, S.

(1995) The Identification of Monetary Policy Disturbances: Explaining the Liquidity Puzzle. *Journal of Monetary Economics*, 35:463–97.

Todd, R.M.

(1991) Vector Autoregression Evidence on Monetarism: Another Look at the Robustness Debate. *Quarterly Review of the Federal Reserve Bank of Minneapolis*, 14:19–37.

Uhlig, H.

(1997) What are the Effects of Monetary Policy? Results from an Agnostic Identification Procedure. Manuscript, Tilburg University.

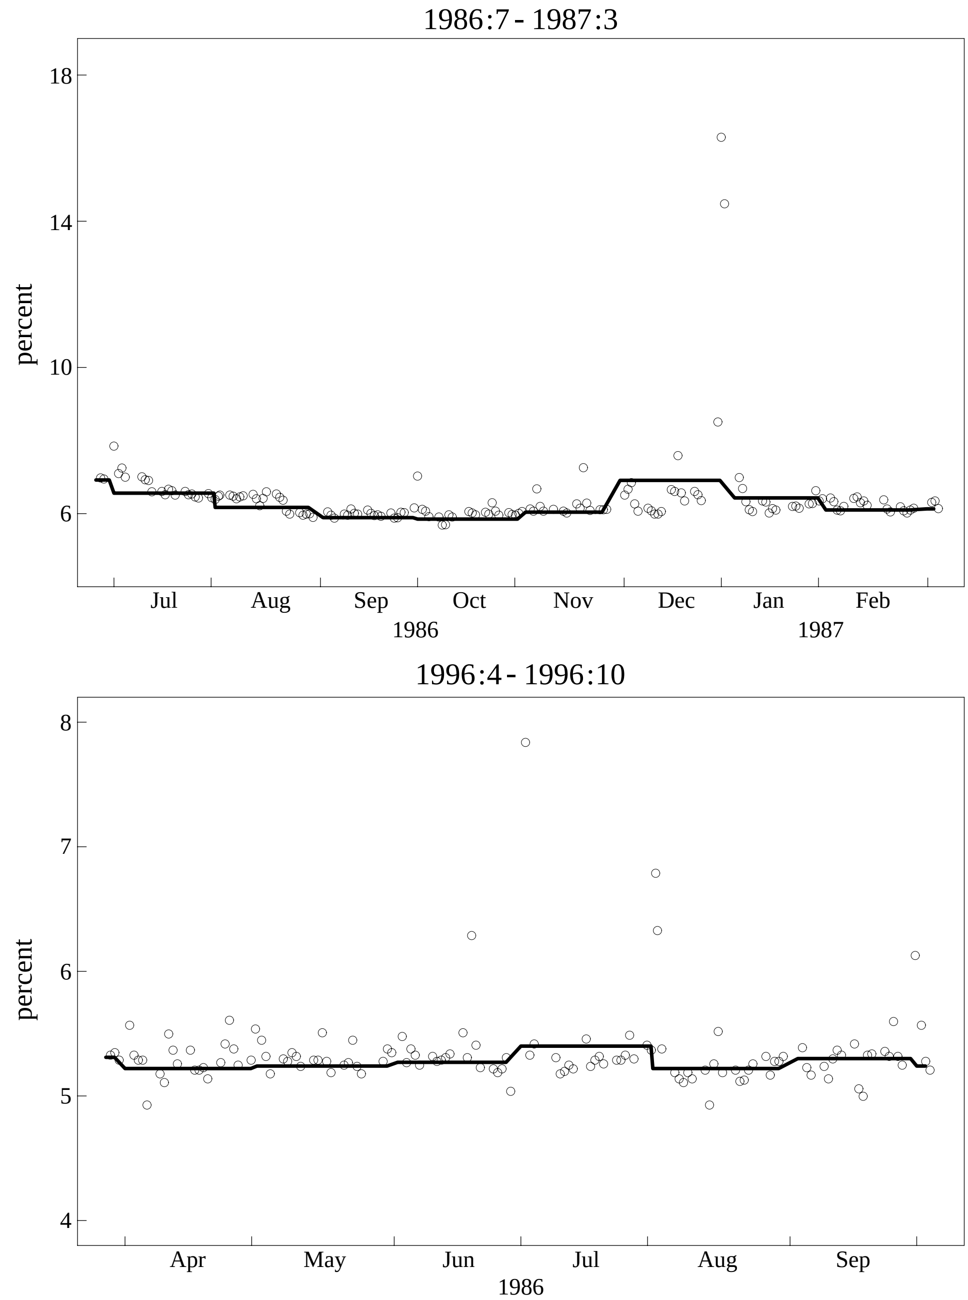

Figure 1: Daily and monthly-average federal funds rate for two time periods.

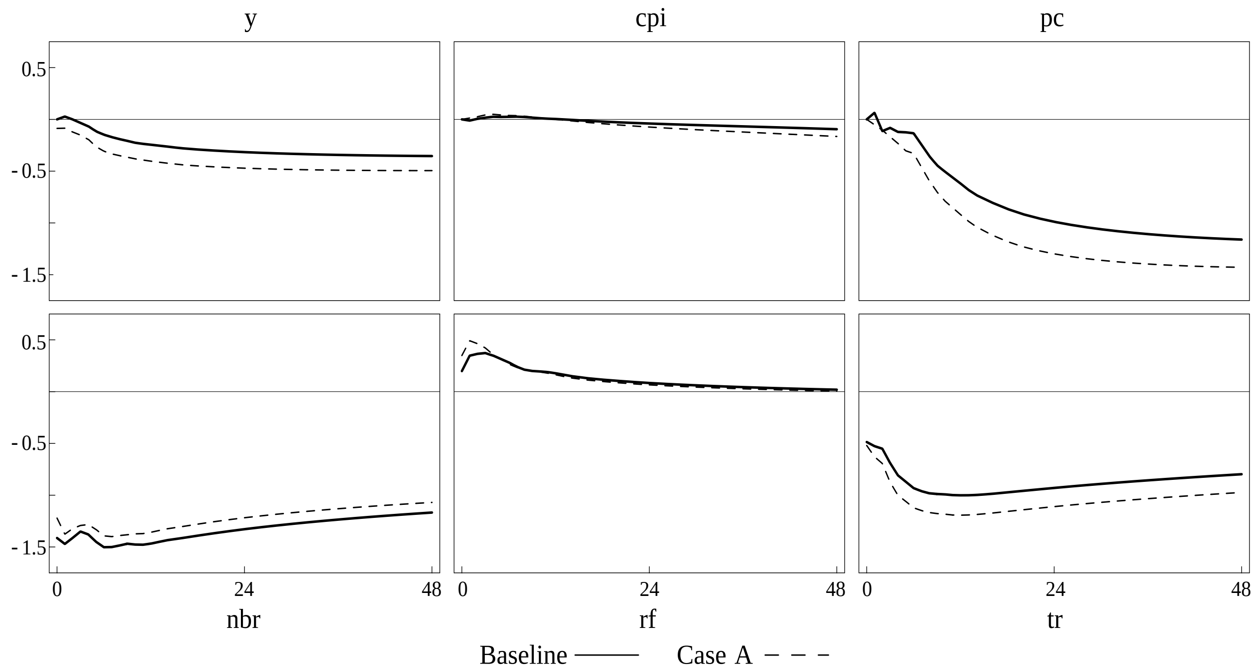

Figure 2: Response to policy shock for baseline and case A identifications of 6-variable model. Each panel reports the response of the labelled variable to a shock graphed against the response horizon in months on the horizontal axis. Except for the interest rate, the units of the vertical axes are approximate percent -- 100 times the difference from the "no shock" path in logarithms. For the interest rate, the figure shows the difference from the "no shock" path in annual percentage rates. The scales on each panel are the same.

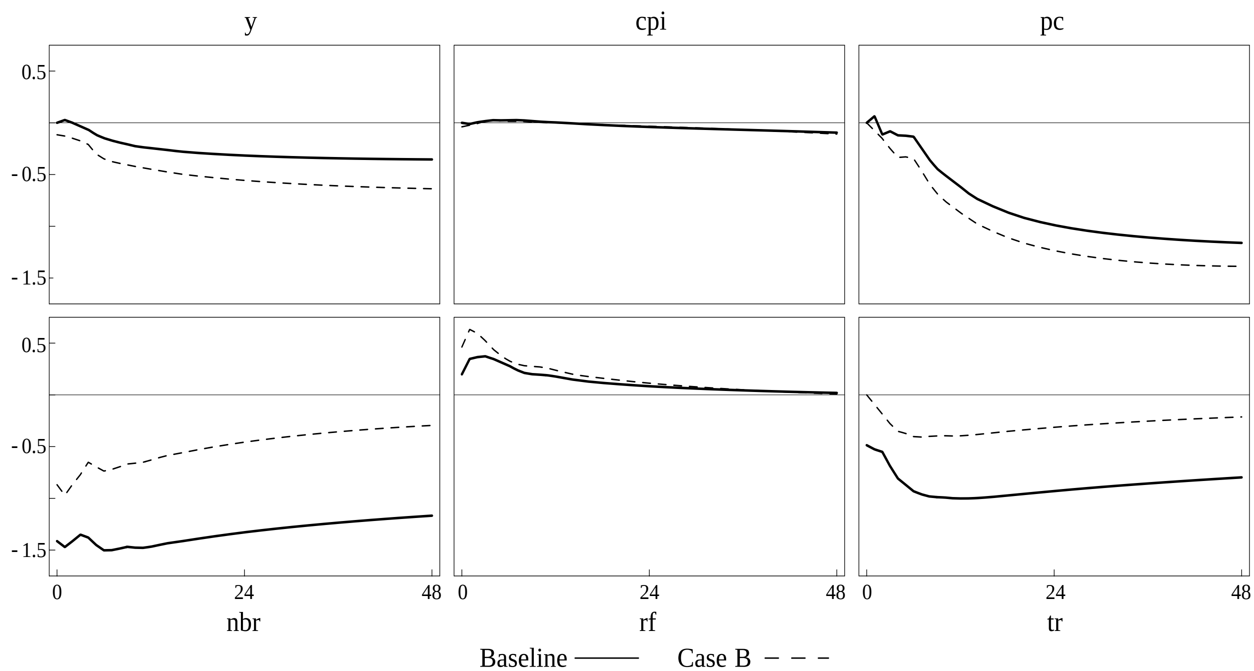

Figure 3: Response to policy shock for baseline and case B identifications of 6-variable model. See notes to Figure 2.

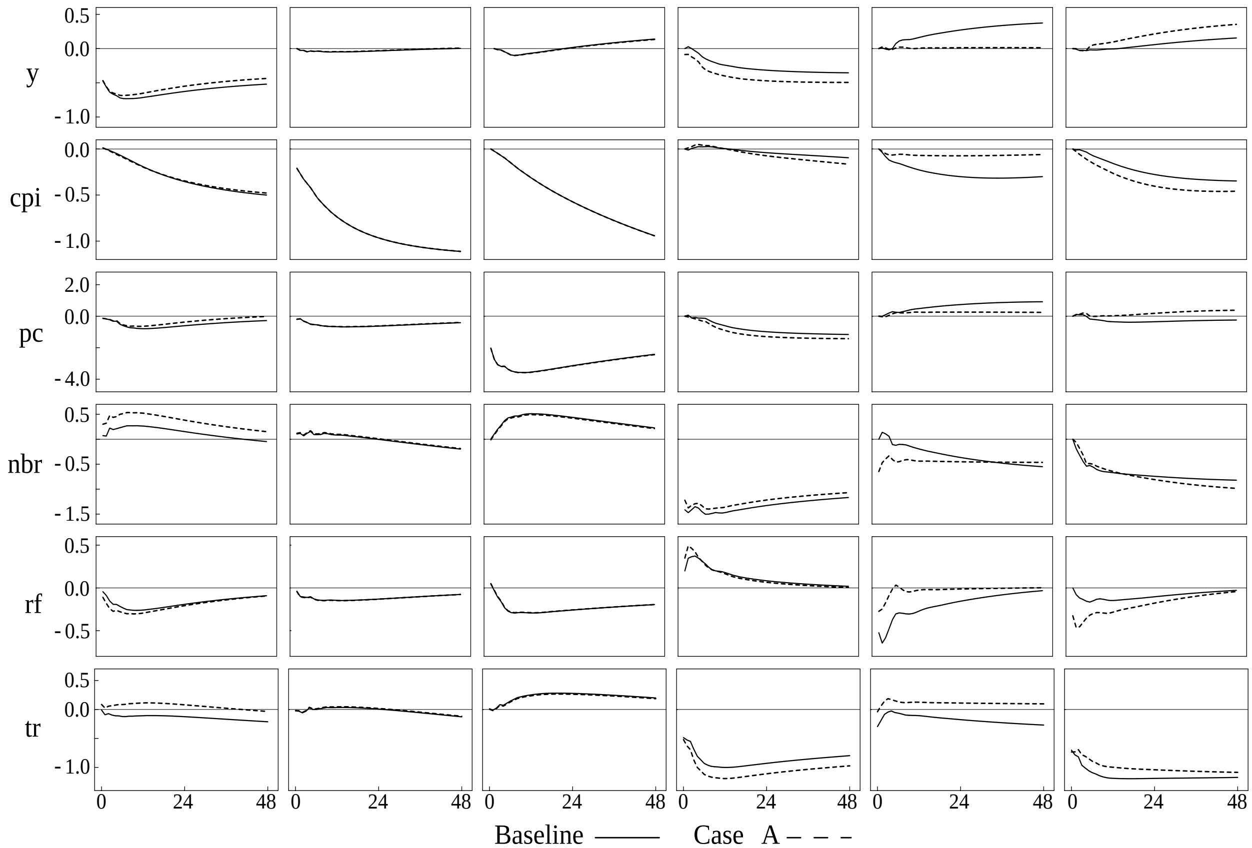

Figure 4: Response of all variables to all shocks for baseline and case A identifications of 6-variable model. Each column reports the response to a particular shock; each panel reports the response of the variable labelled on the row to the shock. The units of all axes are as in Figure 2. The vertical scale for each panel on a given row is the same; all horizontal scales are the same. The fourth column is the response to the policy shock.

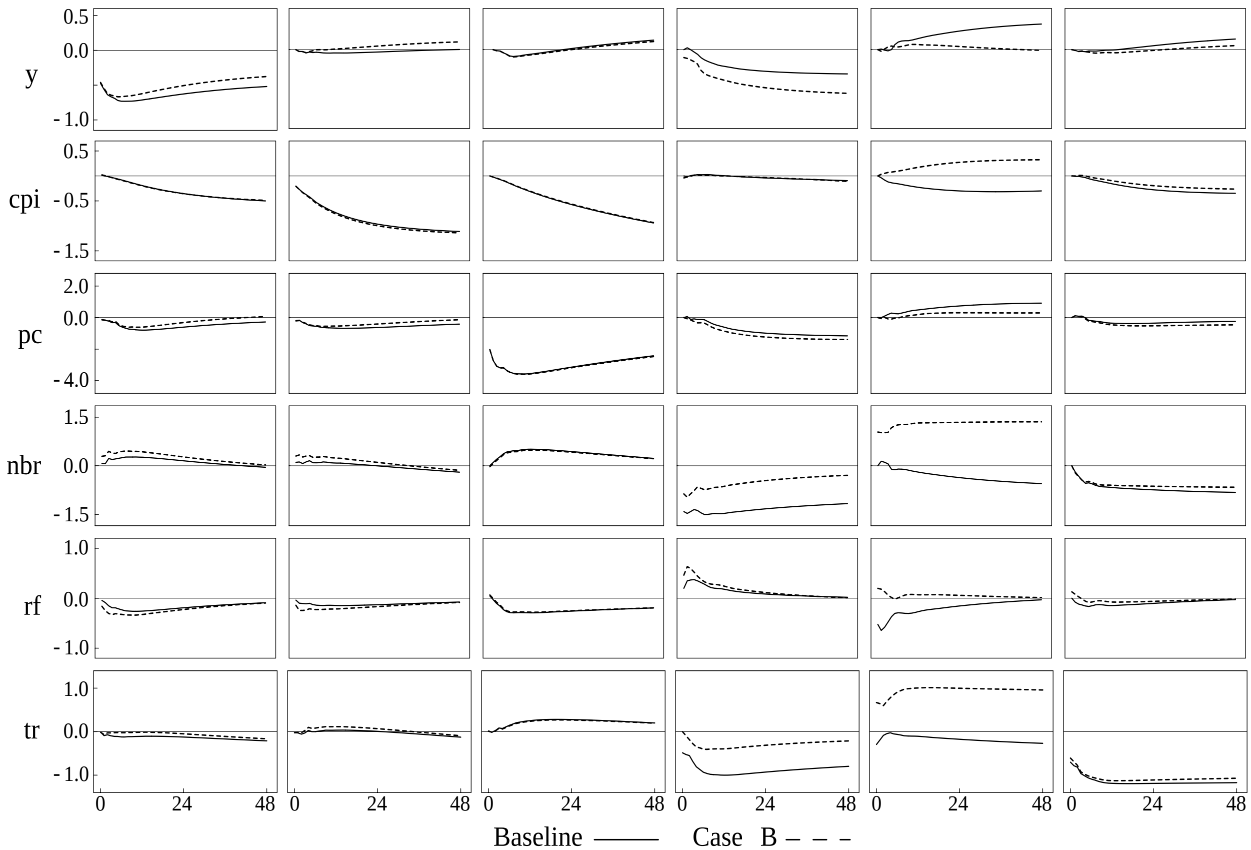

Figure 5: Response of all variables to all shocks for baseline and case B identifications of 6-variable model. See notes to Figure 4.

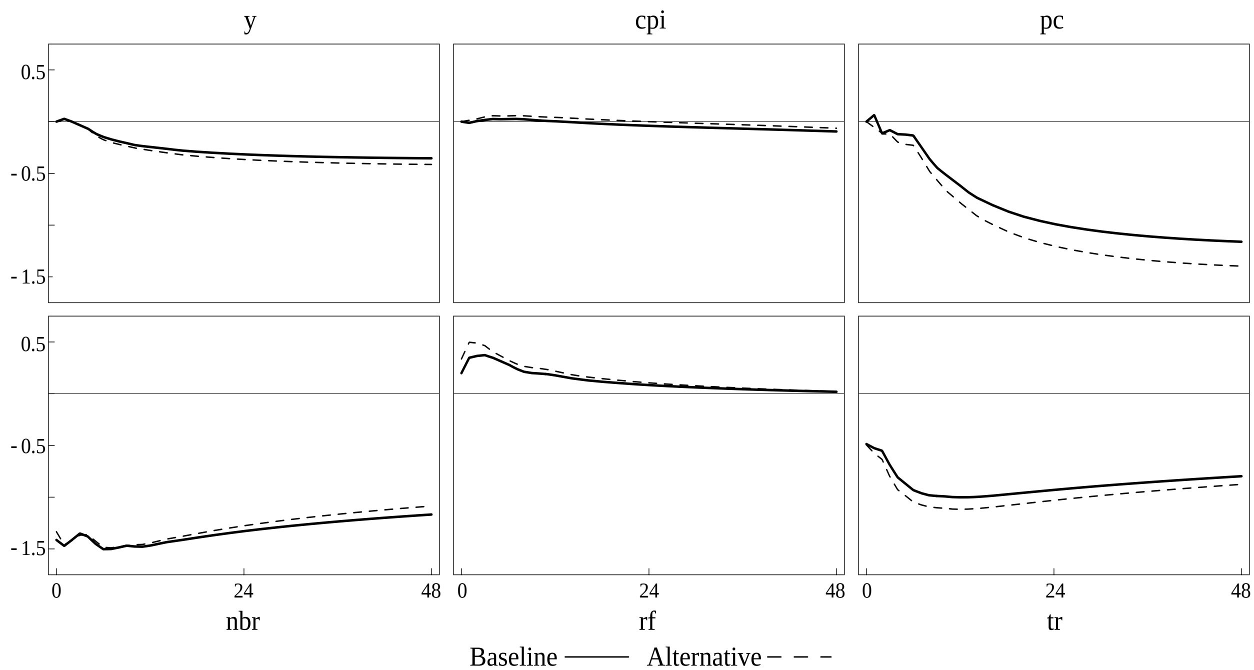

Figure 6: Response to policy shock for baseline and alternative identification of 6-variable model. In the alternative, neither price nor output move contemporaneously in response to the shock. See notes to Figure 2.