Executing Long-Run Restrictions*

Bill Dupor   Leonard Kiefer

September 9, 2008

Abstract

The structural vector-autoregression (SVAR) method uses restrictions from economic theory to identify shocks that have economic, and not simply statistical, meaning. Recent research criticizes this method by applying it to data generated from a stochastic macroeconomic model. The SVAR approach has been evaluated according to its accuracy in estimating the response of hours to a technology shock. This recent research finds that the estimated response of hours has: (i) large bias, (ii) high root mean squared error. This paper develops a new method that avoids using vector autoregressions and instead estimates a moving average representation directly. Our method is based on multi-step ahead forecasts. The method reduces, relative to the standard approach, the bias by 74% and the root mean square error by 7%. Finally, we compare our method performs to alternative, recently developed approaches.

- JEL Codes: C32, C53, E37

1 Introduction

There are few widely accepted theories in macroeconomics. The few that exist mainly concern the long run. Two examples are: money is neutral in the long run; technology (and little else) improves productivity in the long run.

Macroeconomists ought to seize on these rare opportunities in order to empirically identify their economic models. Existing research has show that it is possible to identify an economy's short-run behavior based upon long-run assumptions, also called long-run

*We thank Óscar Jorda for making available computer code on his website. We also thank Nathan Balke, Lawrence Christiano, Paul Evans, Óscar Jorda, James Nason, Masao Ogaki and Frank Schorfheide for helpful comments. We further thank seminar and conference participants at Ohio State University, the Federal Reserve Bank of Atlanta, the Midwest Economics Association Meeting and the Texas Monetary Conference for helpful comments. All errors and omissions in this paper are our own.

restrictions. Three early contributions are Blanchard and Quah (1989), King, Plosser, Stock and Watson (1991) and Shapiro and Watson (1988).

As an introduction to long-run restrictions, consider a simple example of an observer conducting a case study. Suppose an observer interested in the macroeconomy sees that the unemployment rate decreased by 0.25 percentage points in the past month. She knows that an exogenous event caused the decline and can narrow the event to either an increase in the money supply or an increase in the technology level. Moreover, she believes that technology innovations are permanent and that changes in the money supply have no real effects in the long run.

She believes that a technological improvement eventually results in, among other things, greater productivity. With this in mind, she could answer the question of what caused the unemployment rate decline in the previous month by waiting ten to fifteen years to see whether productivity does in fact increase. If it did increase, then she would infer that the original decline in unemployment had been to due to a technological improvement.

It is essential that the observer have a sufficiently long attention span in order to assess whether on not productivity eventually increases; otherwise, the approach would not work. In actual economies, this ”case study” approach is not feasible for a different reason: an economy is being buffeted by different exogenous events that occur sequentially over time. Thinking beyond case studies, Blanchard and Quah (1989) develops a statistical approach to implement this idea.

The method builds a linear combination of a set of reduced form shocks into a set of structural shocks, where some structural shocks have a lasting effects on observed variable(s), while some do not. Just as the observer conducting her ”case study” must have a long attention span, likewise, in the Blanchard-Quah approach, the statistical model must be able to capture covariances of variables at sufficiently long lags to have any hope of identifying the structural shocks.

In implementing Blanchard-Quah's long run restrictions technique (also called a long run structural VAR), many researchers have used vector autogressions to characterize the covariance between observed variables.[1,2]

In an important contribution, Faust and Leeper (1997) recognized econometric difficulties in conducting inference with the technique. They show that the long-run effect of a shock may be imprecisely estimated in a finite sample, possibly with large bias. More-

[1] These include Altig, et.al. (2003), Edge, Laubach and Williams (2003), Fisher (2006), Francis and Ramey (2004), Gali (1999), Gali, Lopez Salido and Valles (2003).

[2] Sometimes this VAR approach is formulated as an instrumental variables problem, even though the two are equivalent.

over, this imprecision and bias in turn applies to all other model parameter estimates, including the short run dynamics.

Chari, Kehoe and McGrattan (2006, hereafter CKM) and Ravenna (2007) study how long-run structural VARs perform in the context of a stochastic growth model (hereafter, the economic model). In doing so, they each provide concrete examples of the Faust and Leeper difficulty. Both CKM and Ravenna identify a technology shock by assuming non-technology shocks have no long run effects on labor productivity. The assumption is satisfied in both models. Each finds that when the employed VAR has a shorter order than the economic model being evaluated, the SVAR delivers biased estimates of impulse responses.

For example, in CKM's economic model, the true elasticity of hours with respect to a technology shock is 0.48 percentage points. However, the long-run SVAR applied to data from the economic model implies that this elasticity is 0.85 percentage points. Thus, an econometrician would on average overstate the impact response of hours to a technology shock by nearly double! Also, CKM show that the impact response of hours is very imprecisely estimated.

Based on this finding, one may conclude that long-run structural VARs are too unreliable to be used in empirical work. In light of this difficulty implementing long-run restrictions paired with the widespread support for the long-run restriction assumptions, we ask in this paper whether there are ways to improve the procedure.

We develop a new approach to executing long-run restrictions. Rather than estimating and then inverting a moving average, we estimate a moving average of the time series directly. We estimate the moving average using *local projections*, as developed by Jorda (2005). To find the s period moving average coefficient matrix, we regress Y_{t+s} on Y_t. and its lags. We construct the sequence of moving average terms by estimating the equation for $s = 1, 2, \ldots, h$.

The structural mapping, or rotation, is then done using the sum of the moving average coefficient matrices. Since the type of long-run restrictions we consider here are about the impulse response function this is more direct. We call the method 'direct rotation.' The direct rotation is done non-parametrically, allowing long horizon covariances to not be dictated by short horizon ones.

We compare our approach to the standard one via the accuracy in estimating the response of hours to a technology shock. We use the metric of bias and RMSE. Our direct rotation method reduces, relative to the standard approach, the bias by 74% and the root mean square error by 7%.

We also address why local projections outperform VARs as a way to execute long-run restrictions in the stochastic growth model. One reason is that VARs are specified with only a few lags and that the VARs are typically used for forecasting. Little forecasting power is lost from using only a few lags. Also, the number of parameters expands rapidly as the lag length increases. Estimating systems with long lags quickly becomes infeasible given the length of standard macroeconomic data sets. The issue with the standard method is evidenced by how the VAR length is usually chosen. Researchers often choose the VAR lag length according to either the Akaike's Information Criteria or Schwartz/Bayesian Information Criterion. These criteria select lag length based on one period ahead forecasting performance and not the specification's ability to capture medium and long horizon covariation. Yet, short-order VARs can severely restrict a stochastic processes' medium and long horizon dynamics.

There are two issues regarding long-run restrictions that our paper does not discuss. First, Christiano, Eicenbaum and Vigfusson (2006, hereafter CEV) have criticized CKM's and Ravenna's exercises because. CEV argue that those papers' calibrations of the economic model are unrealistic. For example, CKM's calibration implies that less than two percent of the variance of hours is explained by technology shocks. With a larger role for technology shocks, CEV show that the performance of the standard SVAR improves. In our calibration, we maintain the assumption of a small role for technology shocks in explaining hours. We do this to give our method a more challenging test, rather than because we necessarily believe that it fits U.S. data.

Second, the response of hours to a technology shock identified with a long-run SVAR, in empirical work, changes drastically depending on whether and how hours are detrended. One popular specification, which is implied by our economic model and which we use in our Monte Carlo analysis, is that hours are stationary. We do not consider other specifications.

The next section presents our data generating process: a two shock stochastic growth model. Section 3 reviews two useful time series representations: the structural and the Wold. Section 4 presents and then compares two methods for executing long-run restrictions: the existing approach based on vector autoregressions and ours based on multi-step ahead forecasts. Section 5 compares our method with several other approaches, and the final section concludes.

2 Data from a Dynamic Equilibrium Model

2.1 A Business Cycle Model

We create a data generating mechanism using a standard neoclassical growth model. There are two exogenous random variables: technology Z_t and a labor tax rate τ_t. Time is discrete: $t = 0, 1, 2, ...$ The corresponding laws of motion are:

$$Z_{t+1}/Z_t = 1 + \sigma_z \varepsilon_t^z$$

$$\tau_t = \rho_\tau \tau_{t-1} + (1 - \rho_\tau)\, \bar{\tau} + \sigma_l \varepsilon_t^l$$

where ε_t^i is iid $\mathcal{N}(0, 1)$ for $i = z, l$. The values Z_0 and τ_0 are given as initial conditions.

A *representative household* chooses hours worked L_t, consumption C_t and next period's capital K_t to maximize utility

$$\sum_{t=0}^{\infty} \beta^t E_0 \left[\log(C_t) + \phi \log(1 - L_t)\right]$$

subject to: $C_t + K_t = (1 - \tau_t)\, w_t L_t + r_t K_{t-1} + (1 - \delta)\, K_{t-1} + T_t$.

The initial capital stock K_{-1} is given as an initial conditions. The wage and capital rental rates are w_t and r_t. Proceeds from labor taxes are rebated lump sum to the agent through T_t. [3] The household's Euler equations are:

$$\frac{(1 - \tau_t)\, w_t}{C_t} = \frac{\phi}{1 - L_t} \tag{1}$$

$$1 = \beta E_t \left[\frac{C_t}{C_{t+1}} (r_{t+1} + 1 - \delta)\right] \tag{2}$$

together with a standard transversality condition.

A *representative firm* hires labor and capital for use in production. In each period, the firm maximizes its profits, output net of labor and capital costs. The production function is: $Y_t = (K_{t-1})^\theta (Z_t L_t)^{1-\theta}$. The implied first-order conditions for profit maximization are:

$$w_t = Z_t \left(\frac{K_{t-1}}{Z_t L_t}\right)^\theta \tag{3}$$

$$r_t = Z_t \left(\frac{K_{t-1}}{Z_t L_t}\right)^{\theta - 1} \tag{4}$$

[3] Standard non-negativity constraints on l_t, c_t, k_t apply and k_0, τ_0, Z_0 are given as an initial conditions.

The *government* has the above exogenous paths for τ_t. Taking these as given, the government chooses T_t to satisfy the following budget constraint in each period:

$$T_t = \tau_t w_t L_t$$

The definition of a competitive equilibrium is standard. Then, *economy-wide resource constraint*, which holds in equilibrium, is $C_t + K_t = Y_t + (1-\delta) K_{t-1}$.

2.2 Calibrations and Impulse Response Functions

As described in the introduction, critics of SVAR identification find the estimators have large bias and high RMSE in samples of length typically used in macroeconomics. It is worth repeating that this criticism hinges on the quantitative assumption that technology shocks have a very small role in explaining business cycles. For example, if technology shocks explain 30% of the variance in hours, the bias and RMSE of long-run structural VAR methods shrink substantially. On the other hand, McGrattan demonstrates poor properties of long-run SVAR identification by studying a model where technology shocks explain only 4% of the variance in hours.[4] To provide the greatest challenge for our new approach, we choose a calibration according to the latter.

Our calibration appears in table 1 and is nearly identical to CKM's.[5] The production function and preference parameters are standard. The exogenous stochastic processes imply that technology shocks explain only 3.5% of the variance in hours. Instead, shocks to the labor tax rate account for most of the variance in hours.

Figure 1 plots the model's impulse responses. Panels (a) and (b) plot the responses of hours to each shock. Each shock is of one standard deviation. A tax rate cut increases hours because it increases a household's private return to supplying labor. A technological improvement increases hours because a household increases savings to accumulate capital, which has become more productive. The impact response of hours is 0.8% and 0.4% to the tax rate and technology shocks, respectively.

Panels (c) and (d) of figure 1 plot the response of labor productivity to each shock. Labor productivity falls in response to the tax rate cut in the short run. This occurs because hours increases (as described above) combined with diminishing returns to labor. After twenty quarters, labor productivity rises above the steady-state level. This occurs because previous capital accumulation works to increase labor productivity. Although la-

[4]See Christiano, Eichenbaum and Vigfussion (2006).

[5]We differ from CKM in that did not include population or technology growth in our model. We do this for simplicity.

Table 1: Benchmark calibration of the economic model

Variable	Meaning	
σ_τ	tax rate shock innovation std. dev.	0.0136
σ_z	tech. shock innovation std. dev	0.0131
$\bar{\tau}$	steady-state labor tax rate	0.243
ρ_τ	persistence of labor tax rate shock	0.952
Contribution of tech. shocks to variance of hours		3.4%
β	discount factor	0.995
θ	depreciation rate	0.015
δ	capital's share	0.333
ϕ	labor pref parameter	2.50

bor productivity eventually does return to the steady-state level which happens (beyond the number of periods plotted), even ten years after labor productivity remains high.

In response to a technological improvement, labor productivity jumps upward, as seen in figure 1(d). Moreover, as every household accumulates capital, labor productivity rises further. Because the technology shock is permanent, labor productivity is permanently higher following the shock. The tax rate shock does not change the long-run level of labor productivity. This special feature of the technology shock provides the identification used in many SVAR analyses, including the ones presented in the following sections.

3 Two Useful Representations

3.1 Wold Representation

Consider a zero mean purely linearly indeterministic covariance stationary 2 by 1 vector stochastic process X_t. This is given by: $X_t = \begin{bmatrix} p_t & h_t \end{bmatrix}'$. Here $p_t = \Delta \log (Y_t / L_t)$ and $h_t = \log (L_t)$, where each is demeaned. This time series has a Wold representation:

$$X_t = \Gamma (L) u_t \tag{5}$$

where $\Gamma (L) = \Gamma_0 + \Gamma_1 L + \Gamma_2 L^2 + ...$ and L is a 2 by 2 lag operator. Let $E (u_t u_t') = \Omega$.

Because (5) is a Wold representation, u_t is uncorrelated at non-zero leads and lags, $\Gamma_0 = I$ and $\sum_{k=0}^{\infty} \Gamma_k^2 < \infty$. Also, the $(i,j)th$ element of Γ_k is denoted $\Gamma_{i,j}^{(k)}$.[6] The errors u_t represent the forecast error from a linear function of past X :

[6]The $\Gamma (L), \Gamma_k, \Gamma_{i,j}^{(k)}$ convention is used throughout the paper for lag polynomials.

Figure 1
True impulse responses to each shock.

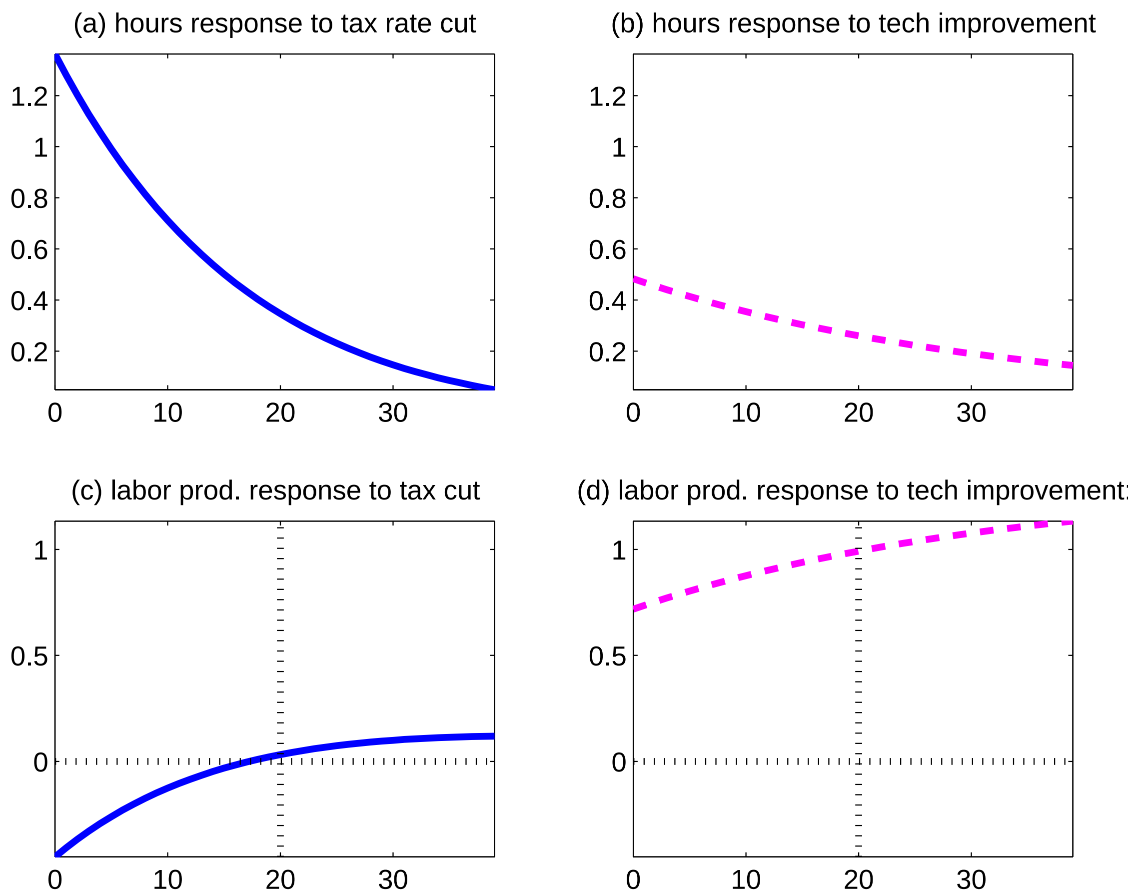

Notes: say more here.

$$u_t = X_t - \widehat{E}\left(X_t | X_{t-1}, X_{t-2}, ...\right). \tag{6}$$

The forecast errors u_t are devoid of any economic meaning.

3.2 Structural Representation

To give economic meaning to the innovations to X_t, we introduce the 'structural representation.' For concreteness, we apply the structural representation implied by the economic model. Here, $\varepsilon_t = \left[\begin{array}{cc} \varepsilon_{z,t} & \varepsilon_{l,t} \end{array} \right]'$.

We rotate the u_t vector according to:

$$P\varepsilon_t = u_t$$

The structural representation is then

$$X_t = \Psi(L)\,\varepsilon_t \tag{7}$$

The structural representation satisfies $\Psi(L) = \Gamma(L)\,P$, where $PP' = \Omega$ by construction.

Mapping from the Wold to structural representation requires an identification restriction. Our economic model provides one such restriction: the cumulative effect over the infinite horizon of a non-technology shock on labor productivity equals zero. Examining panels (c) and (d) of figure 1, labor productivity converges to the steady-state following tax rate shock, but increases permanently in response to a technological improvement.

One powerful feature of this restriction is that it holds in a more general class of models than the one above. For example, in a standard sticky price or cash in advance model, a permanent increase in the nominal money stock also has no long-run effect on labor productivity.

The identifying restriction implies:

$$\sum_{k=0}^{\infty} \psi_{1,2}^{(k)} = 0 \tag{8}$$

Equation (8) is referred to as a *long-run restriction*. Next, we assume the structural innovations are orthogonal, i.e. $E(\varepsilon_t \varepsilon_t') = I$. As a sign restriction, we follow existing research and assume that a positive technology shock leads to a non-negative long-run response of labor productivity.

4 Two Methods to Execute Long Run Restrictions

4.1 VAR Inversion

The first method, which we call **VAR Inversion**, is the most common way of implementing long-run restrictions.[7] This first method is criticized by CKM (2005), McGrattan (2006) and Ravenna (2007). The method consists of two steps.

Step one uses the VAR representation for X_t:

[7] It is used, for example, in Gust et.al. (2006) and CKM and is also a built in procedure in the statistical program Eviews.

$$X_t = \sum_{i=1}^{\infty} \bar{B}_i X_{t-i} + u_t. \tag{9}$$

VAR Inversion, the standard method for executing long-run restrictions, uses (9) instead of the Wold representation or general ARMA(p,q) representation to estimate the process. This assumes that the MA process is in fact invertible. The actual VAR for a covariance stationary process is, in general, infinite order. Typically researchers will estimate a short order VAR, with lags based on some selection criteria such as Akaike or Schwartz. As such, we estimate a truncated version of (9). We denote this as

$$X_t = B(L)\, X_{t-1} + u_t \tag{10}$$

where $B_k = 0$ for $k \geq p$ and therefore $B(1) = I - \Gamma(1)^{-1}$.[8] The validity of (10) rests on the assumption that $E(X_{t-k}u_t') = 0$, or that p lags captures all of the dynamic correlation between elements of X.

As is standard, one can estimate $B(L)$ with equation-by-equation least squares. Assuming that the moving average operator is invertible and the researcher increases the lag length as the sample size grows, the sample covariance function of the regression residuals provides a consistent estimate of Ω.

Step two inverts (10) to get the moving average representation. With the moving average representation, the long run restriction is then implemented.

$$B(z) = \left[I - \Gamma(z)^{-1} z\right]^{-1}$$

Then, the mapping P is estimated as:

$$\hat{P}^* = \left[I - \hat{B}(1)\right] chol\left(\left[I - \hat{B}(1)\right]^{-1} \hat{\Omega} \left(\left[I - \hat{B}(1)\right]^{-1}\right)'\right) \tag{11}$$

---

[8] $X_t = B(L)X_{t-1} + u_t$

$(I - B(L)L)\, X_t = u_t$

$X_t = (I - B(L)L)^{-1}\, u_t, \forall L$

$X_t = \Gamma(L)u_t = (I - B(L)L)^{-1}\, u_t$
Letting $L = I$ generates the equation in the text.

Figure 2
Impulse response of hours to technology improvement: VAR inversion.

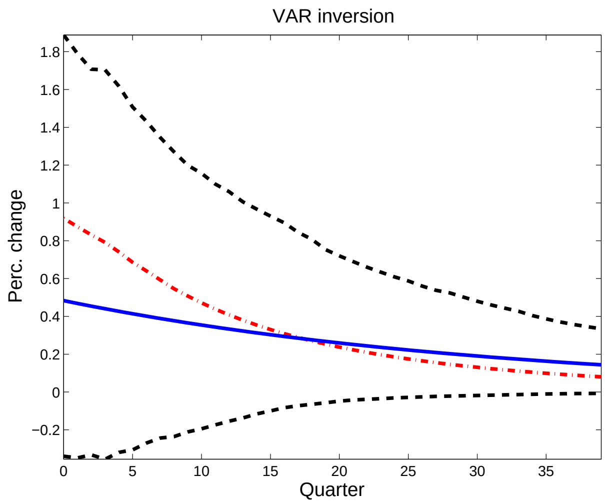

Notes: True impulse response (solid), mean impulse response implied by VAR inversion (dash-dot), 90% range of estimates implied by VAR inversion (dashed). Generated from 1000 simulations, each using data of sample length 200.

Here, chol() is the Choleski decomposition.[9],[10]

By construction, $\hat{P}^*$ implies that the restrictions on the structural representation are satisfied. To see this, verify the equality of the moments generated by the two representations: (10), after substituting in (11) and (12), is equivalent to (7).

Finally, the estimate of the structural representation is:

$$\hat{\Psi}^*(L) = \left[I - \hat{B}(L)\right]^{-1} \hat{P}^* \tag{12}$$

4.2 Bias and RMSE under VAR Inversion

In this subsection, we estimate the response of hours to a technology shock using VAR inversion and data from the economic model. This will reiterate the finding of large bias and RMSE that has been demonstrated by previous authors. Following this, we will present the new method and quantitatively compare it to VAR inversion.

For VAR inversion, we assume the number of lags in $B(L)$ equals 4. This is both

[9]To avoid confusion, we use * notation, when necessary, to distinguish estimates based on VAR Inversion from those based on our first new approach, Direct Rotation.

[10]Letting $\left[I - \hat{B}(1)\right]^{-1} \hat{\Omega} \left(\left[I - \hat{B}(1)\right]^{-1}\right)' = A$, we know A is positive definite and real valued, in which case there is a unique lower triangular matrix L such that $A = LL'$. $L = $chol$(A)$.

typical for existing studies.[11] Recall that the actual lag polynomial $\bar{B}(L)$ may be of infinite order and is also the source of the misspecification.

As a first step, we draw one thousand random samples of $\left(\varepsilon_t^z, \varepsilon_t^l\right)$, each of length two hundred. We feed these shocks through the economic model to generate one thousand time series of productivity growth and hours. Third, we estimate the hours response to a technology shock via VAR inversion for each sample.

In figure 2, the dash-dotted line plots the impulse response using VAR inversion averaged across all simulations and the solid line plots the true response. VAR inversion implies a large upward bias in the first two years following the shock. Table 2 gives the numerical values for the hours responses in the initial period. While the true response is 0.48%, the average response under VAR inversion is 0.92%.

One way to express the magnitude of the bias is the difference between the two responses divided by the true response (all on impact). This value is 68%. CKM (2006) use this metric and conclude that such a large bias renders long run restrictions useless for macroeconomics.[12]

Why does VAR inversion have a large bias? We must distinguish between the theoretical ('true') responses and those generated by VAR inversion. First, following a tax rate cut, the true response of labor productivity rises above the steady-state at the medium horizon, as seen in figure 1(c). The rise in labor productivity, in the true model, happens only gradually due to capital accumulation. On the other hand, a VAR with only four lags cannot capture this gradual increase. Thus, only a small fraction of the volatility of hours using VAR inversion (relative to the true amount) will be due to the tax shock.

Because the structural VAR must account for all of the volatility of observed variables, more volatility of hours (using VAR inversion) must be due to another shock, besides the tax rate. A larger hours response to the technology shock accomplishes this.

In addition to large bias, VAR inversion also implies an imprecise estimate of the hours response to the technology shock. The upper and lower dashed lines in figure 2 represent the boundaries between with 90% of the 1000 simulated impulse response functions lay. In terms of the impact response, this range varies from -0.4% to nearly 1.85%. Based on a figure similar to this one and a similar calibration, CKM conclude that this range of values is too large to be useful in distinguishing between different economic models.

Reviewing the two criticisms of VAR inversion: (i) there is a large bias, (ii) there is a great deal of imprecision. These criticisms are devastating, according to CKM, because the structural VAR procedure has been applied by other researchers to develop the find-

[11]See, for example, Altig, et. al. (2004), Edge, Laubach and Williams (2003) and Gali (199).

[12]See also Ravenna's discussion of the bias generated by VAR inversion.

Table 2: Initial period responses of hours to positive technology shock: by method

Procedure	Mean Initial Period Response	Reduction in Bias (%)	Reduction in RMSE (%)
True	0.48		
VAR Inversion	0.92		
Direct Rotation	0.60	73.97	7.34
Spectral	0.57	79.07	27.64
Aux. Rotation	0.85	16.84	-2.32
Hybrid	-0.09	-31.85	24.62

ing that hours falls in response to a technological improvement. Gali (1999) and Francis and Ramey (2005) make this very strong claim. Both papers state that their finding is consistent with the predictions of sticky price models, but inconsistent with those of real business cycle models.

According to criticisms (i) and (ii) above, the structural VAR is unreliable and therefore the conclusions drawn from the structural VAR deserve little currency.

We have one deep concern about the exercise in CKM and Ravenna. To illustrate our concern simply, suppose that we move from our calibration where technology shocks account for 3.4% of hours volatility to the case where technology shocks account for 0% of hours variance.

Specifically, suppose that hours does not respond to a technology shock at all. How would criticisms (i) and (ii) apply in this case? First, according to the metric used in (i), any amount of bias in population (even if the VAR inversion implies a response of 0.00001 percentage points) would imply an *infinite* percentage error. Second, according to the metric used in (ii), any imprecision in small sample would that hours may either rise or fall after a technology improvement.

Our simple example is the razor's edge between hours rising (corresponding to the RBC case) and hours falling (corresponding to the sticky price case). Thus it is not surprising that in our benchmark calibration, which is quantitatively very similar to our simple example, is unsuccessful in distinguishing between the two cases. Along the same line of argument, CEV show that by changing the calibration so that technology shocks explain 20% or more of the volatility of hours, the bias and imprecision of VAR inversion are reduced dramatically.

Figure 3
Hours responses to a technology shock

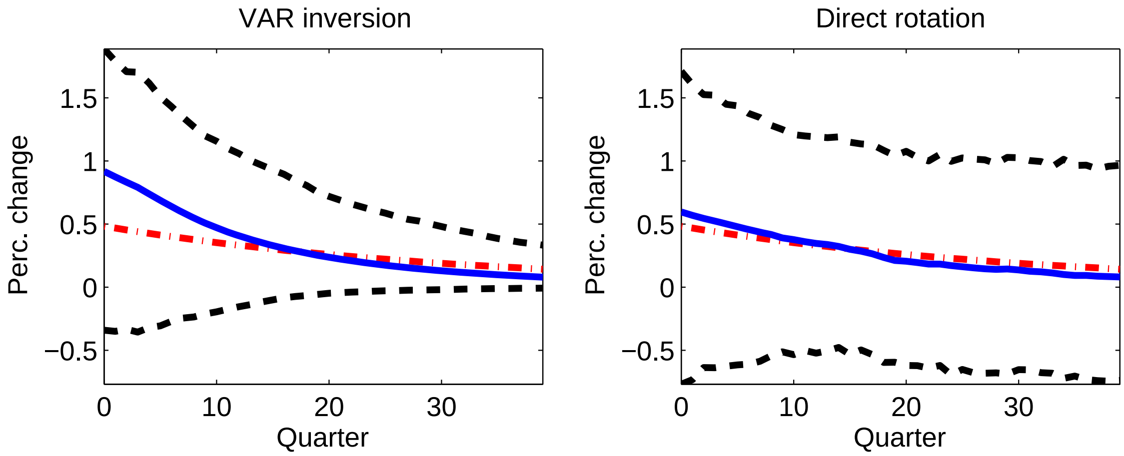

Notes: True impulse response (solid), mean impulse response implied by method (dash-dot), 90% range of estimates implied by method (dashed). Generated from 1000 simulations, each using data of sample length 200.

4.3 Direct Rotation

Our new method, which we dub **direct rotation**, estimates the Wold representation and applies the identification restriction to find the structural representation. To estimate the moving average coefficients, we use Jorda's (2005) Local Projections method, which constructs a "collection of projections local to each forecast horizon." The moving average coefficient matrix at lag s is estimated by a regression of the endogenous variables shifted $s+1$ periods forward on themselves:

$$Y_{t+s} = G_s Y_t + C_s\left(L\right) Y_{t-1} + v_{t+s}^{s} \tag{13}$$

for $s \geq 1$. The lag operator $C_s\left(L\right)$ has finite order and is estimated separately for each s.

Then, an estimate of the reduced representation moving average coefficient matrix is:

$$\hat{\Gamma}_s = \hat{G}_s$$

for $s \geq 1$. As previously assumed, $\Gamma_0 = I$.

The choice between executing a long run restriction via a VAR, as in the previous section, versus via a sequence of multi-step ahead regressions has an analogue in the forecasting literature. Suppose one wishes to forecast three period ahead inflation based upon current inflation. One could estimate an AR(1) and then iterate the AR(1) three periods ahead. Marcellino, Stock and Watson (pg. 1, 2008) state that this ”involves a trade-off between bias and estimation variance: the iterated method produces more efficient parameter estimates than the direct method, but it is prone to bias if the one-step ahead model is misspecified.” In our context, the former corresponds to VAR inversion and the latter corresponds to direct rotation.

Returning to our neoclassical growth model, the slow moving nature of labor productivity and the capital stock imply that a short-order VAR is insufficient in capturing long horizon dynamics. As such, the resulting structural impulse responses may be have significant bias. If bias is the primary concern, one might expect that direct rotation performs better than VAR inversion.

To apply the procedure, we must select: (i) the maximum lag length of $C_s(L)$, which we denote a, and (ii) the largest lag h of the operator $\Gamma(L)$ that is estimated. The latter amounts to assuming that $\sum_{s=h+1}^{\infty} \Gamma_{12}^{(s)} = 0$. Requirement (ii) has a natural economic interpretation. The original long run restriction, described in the introduction and expressed mathematically by (8), stated that the effect of a non-technology shock on labor productivity is zero after an arbitrarily long time (and hence the infinite summation).

Requirement (ii) means that after a finite time h, labor productivity growth and hours return and remain at their pre-shock values h periods following either shock. Interestingly, Faust and Leeper (1997) suggest a related idea. They state that one way to avoid econometric difficulties with the Blanchard-Quah method is to replace infinite horizon long-run restrictions with finite horizon ones. Faust and Leeper do not pursue this idea further.

Our method explicitly assumes that after h periods the net effect of any non-technology shock on productivity is zero, and that any further movements in productivity from period $h+1$ forward will net out to be zero. Essentially we assert that all interesting dynamics in productivity resulting from non-technology shocks end after h periods. This method imposes an operational meaning of the long run. It offers two strengths and one weakness. First by imposing a finite interpretation of the long run, one which furthermore must be less than the sample size, we gain in the number of observations of the long run. Second, we could test directly whether or not shocks do indeed die out after h periods. The weakness is that many models only imply that the long-run restriction holds in the limit and thus we risk misspecification by imposing a finite horizon.

To construct an estimate for the covariance matrix Ω, define the residuals $\hat{v}_{t+s}^{s}$ as

$$\hat{v}_{t+s}^{s} = Y_{t+s} - \hat{G}_s Y_t - \sum_{i=0}^{a} \hat{c}_{s,i} Y_{t-1}$$

These residuals v_{t+s}^{s} are an average of forecast errors (from time t to $t+s$) weighted as follows:

$$v_{t+s}^{s} = \sum_{j=0}^{s-1} \Gamma_j u_{t+s-j} \tag{14}$$

As in Jorda (2007), we can construct an estimate of Ω by computing forecast errors for $s = 1$. Noting that $\Gamma_0 = I$, we have

$$\widehat{\Omega} = \frac{1}{T-k-h} \sum_{j=1}^{T-k-h} \widehat{u}_t \widehat{u}_t' \tag{15}$$

where $\widehat{u}_t = Y_{t+1} - \widehat{G}_2 Y_t - \widehat{C}_2 (L) Y_{t-1}$.

By estimating $\Gamma(L)$ directly, we avoid specifying the VAR(k) in the first method. To recover the structural innovations, we must compute the long run impact of the reduced form shocks on Y_t. This estimate is:

$$\hat{\Gamma}(1) = \sum_{s=0}^{h} \hat{\Gamma}_s \tag{16}$$

Then, the estimates of the rotation matrix and structural impulse response function are:

$$\hat{P} = \hat{\Gamma}(1)^{-1} \operatorname{chol}\left(\hat{\Gamma}(1)\, \hat{\Omega} \hat{\Gamma}(1)'\right) \tag{17}$$

and

$$\hat{\Psi}(L) = \hat{\Gamma}(L)\, \hat{P}$$

Figure 4 describes the difference between VAR inversion and direct rotation. The long-run identification restriction is identical in the two cases. The difference arises in which statistical model is used to execute the long-run restriction. Direct rotation estimates the reduced form moving average directly (via local projections) while VAR inversion estimates, and then inverts, a VAR to get the reduced form moving average.

In order for any execution of a long-run restriction to perform well, one must approximate the sum of the moving average coefficients accurately. One case see this by examining equation (16). Estimating moving average coefficients at medium and long horizons

Figure 4
Two methods to execute a long-run restriction: VAR inversion and direct rotation.

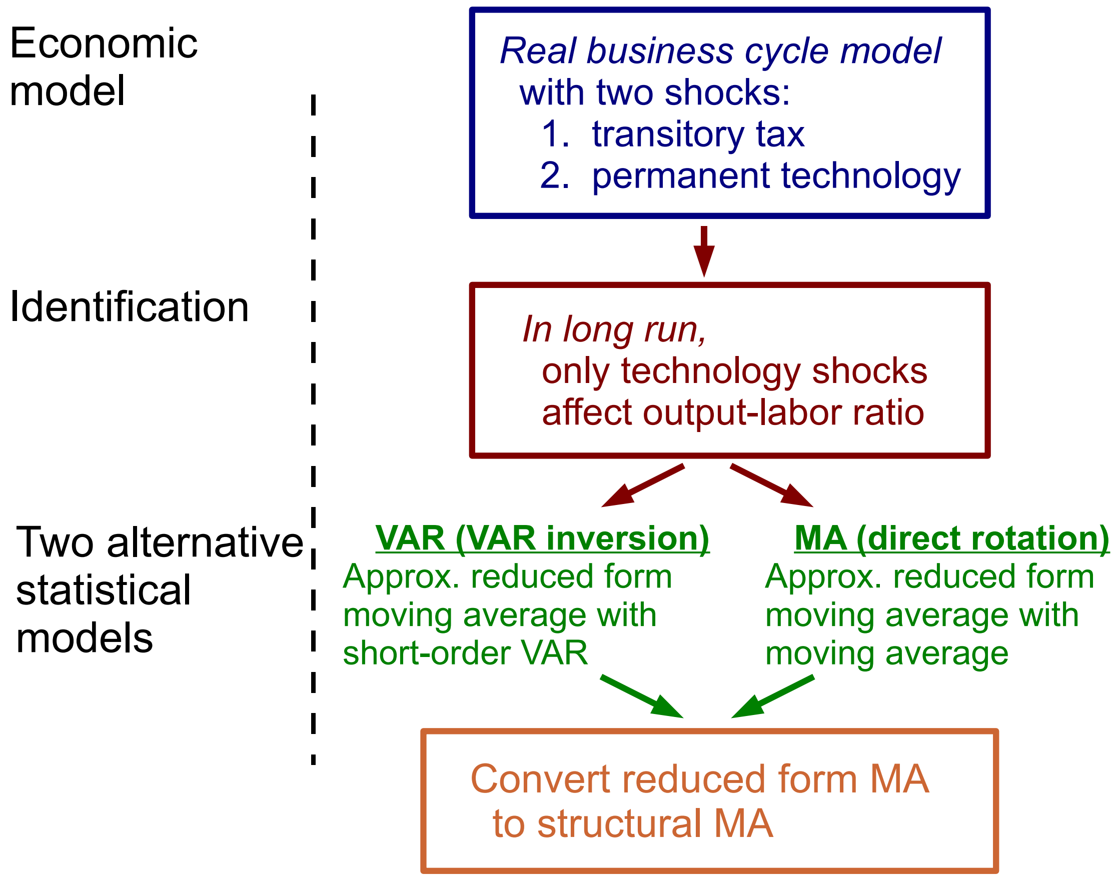

is as important as estimating those at short horizons. Using direct rotation, fitting all of these coefficients well is built into the method via local projections. On the other hand, a VAR seeks only to minimize the one period ahead forecast error variance. If a VAR has too few lags, then it may approximate the intermediate and long horizon covariances poorly.

VAR inversion works relatively well in terms of bias and efficiency if the true stochastic process can be described by a VAR with a number of lags equal to that of the approximating VAR. Then, using the less restrictive approach of direct rotation implies efficiency loss. For examples with slow moving unobserved variables, such as the capital stock in the neoclassic growth model, the short-order VAR is a poor approximation.

4.4 Direct Rotation vs. VAR Inversion: Bias and RMSE

This section compares the results for direct rotation with those from VAR inversion. For direct rotation, we set the lag length a for the forecasting term $C_s(L)$ in (13) equal to 4. The truncation point h for computing the long run responses in (16) is set equal to 40 in our benchmark specification.

Figure 3 plots the impulse responses (averaged across 1000 samples) for direct rotation and VAR inversion as well as for the true model. First, there is a dramatic reduction in the size of the bias moving from VAR inversion to direct rotation. For direct rotation, the average of the estimated impulse response lies on top of the true response. Examining table 2, direct rotation reduces the bias in the impact response of hours by 74%.

While the bias is reduced substantially, there is less improvement in the RMSE. Direct rotation reduces the RMSE by 7.3%. Another way to see this is by comparing the 90% intervals in figure 3. These are very similar across direct rotation and VAR inversion.

Figure 5 contains two histograms for the period zero response of hours. One histogram is constructed using VAR inversion and the other is constructed using direct rotation. Both histograms have very wide ranges. The result should not be surprising. Since technology shocks account for only 3.7% of the variance of hours, estimation strategies based only on identification of such an insignificant shock are likely to be imprecise in small samples. Moreover, it takes many years for labor productivity to return to its steady-state following the non-technology shocks. As such, it would take a great deal of data to precisely distinguish between the technology and non-technology shocks.

As a sensitivity analysis, tables 3 and 4 compare bias and RMSE for different values of the parameters used in direct rotation (the horizon h and number of lags a) and VAR inversion (the number of lags k). The sensitivity is instructive for understanding how the two methods differ.

First, using direct rotation, the bias is reduced relative to VAR inversion except when the truncation horizon is very short. That is, in table 3 bias increases only when $h = 10$. This is intuitive. Recall that VAR inversion exhibits large bias because it inaccurately captures covariance between variables at long leads and lags. Similarly, when direct rotation has a low value of h, it restricts covariances between t and $t + j$ to equal zero for $j > h$.

Second, according to table 4, a very large truncation value ($h = 90$) leads to a large increase in RMSE. This is intuitive. Given that the sample length equals 200, the amount of data available to estimate the 90*th* autocorrelation is limited. As such, these long-run autocovariances are estimated imprecisely. All autocovariances enter into the rotation matrix that maps reduced form shocks into structural shocks. As such, even the short-run structural impulse response functions are also estimated more imprecisely.

Table 3: Bias comparison between direct rotation and VAR inversion: varying specifications

		Number of lags (a or k)			
		2	6	10	20
	10	-156.90	-153.29	-148.72	-137.58
Truncation value	30	70.54	70.18	70.15	66.65
	50	42.05	46.33	51.04	68.11
	70	27.69	34.05	40.99	72.85
	90	32.62	43.58	56.16	95.90
VAR inversion		-5.91	7.30	21.44	50.03

Notes: Percentage reduction in bias relative to VAR inversion with 4 lags. A negative number means that bias increases relative to VAR inversion with 4 lags. Generated from 1000 simulations, each using data of sample length 200.

Table 4: RMSE comparison between direct rotation and VAR inversion: varying specifications

		Number of lags (a or k)			
		2	6	10	20
	10	-47.70	-46.22	-44.05	-38.70
Truncation value	30	11.33	11.78	10.66	6.80
	50	-4.63	-5.25	-4.28	-3.28
	70	-17.38	-15.69	-15.38	-13.15
	90	-25.70	-26.36	-25.03	-16.41
VAR inversion		1.96	-1.44	-3.94	-17.58

Notes: Percentage reduction in RMSE relative to VAR inversion with 4 lags. A negative number means that RMSE increases relative to VAR inversion with 4 lags. Generated from 1000 simulations, each using data of sample length 200.

Figure 5
Distribution of period zero hours responses to technology shock: by method.

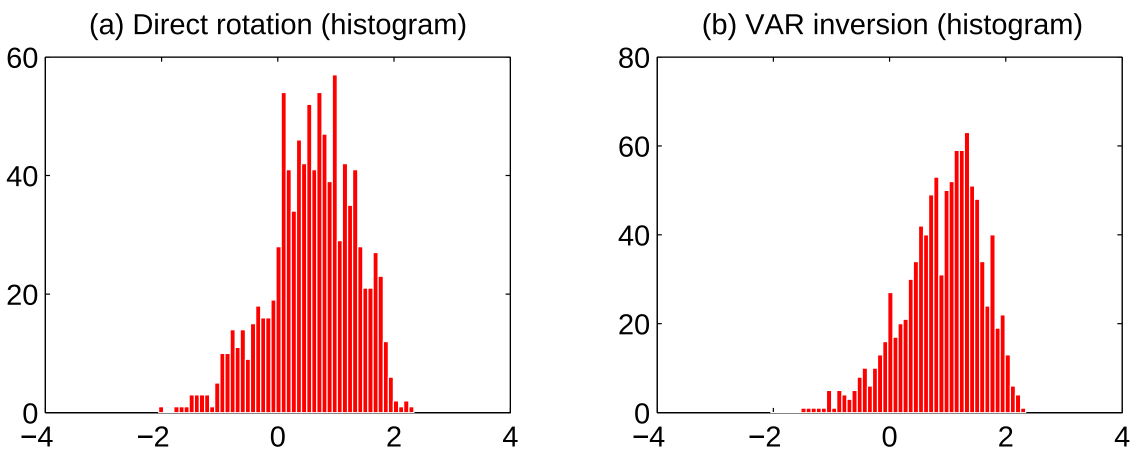

Notes: Generated from 1000 simulations, each using data of sample length 200.

Third, holding fixed the truncation value, increasing the number of lags in direct rotation has little effect on RMSE. For example, in table 4 and $h = 30$, using 6, 10 and 20 lags leads to a RMSE reduction equal to 12%, 11% and 6.8%. This is surprising since one may expect that increasing the number of parameters required in estimation would make the structural estimates less precise.

Since adding lags of the observed variables to the multi-period forecasts has little effect on RMSE using direct rotation, one might suspect that adding more lags to the VAR in the VAR inversion method might have a similar result. The final row of 4 address this question. As the number of lags increases using VAR inversion, the RMSE increases substantially. Using 6, 10 and 20 lags leads to RMSE increases of 1%, 4% and 18%. We do not have an intuition for why more lags hurts VAR inversion RMSE but has little effect for direct rotation.

5 Alternative Approaches

5.1 Auxiliary Rotation

Auxiliary rotation, introduced by Jorda (2007), also estimates the moving average coefficients directly.[13] The moving average coefficients are computed in exactly the same fash-

[13]This is our name for Jorda's procedure. We explain our rationale for the name later in this section.

ion as the direct rotation method. For this reason, the question 'Are short-run dynamics estimated using a VAR?' in table 5 is answered 'No' for both auxiliary and direct rotation.

Table 5: A Taxonomoy of Methods to Execute Long-Run Restrictions

Method	Is rotation estimated using a VAR[b]?	Are SRD[a] estimated using a VAR?	Originators
VAR inversion	Yes	Yes	Blanchard and Quah (1989)
Direct rotation	No	No	this paper
Auxiliary rotation	Yes	No	Jorda (2007)
Spectral	No	Yes	Christiano, Eichenbaum and Vigfusson (2006)
Direct rotation-VAR inversion hybrid	Partially[c]	No	this paper

Notes: [a] VAR Inversion is the standard method, which is employed by Chari, Kehoe & McGrattan (2005). [b] SRD stands for short run dynamics. [c] Hybrid method estimates reduced form moving average at short horizons using local projections and VAR inversion at long horizons.

Direct rotation uses the moving average coefficients to estimate the rotation matrix. In contrast, the auxiliary rotation does not use these coefficients. Instead, the rotation matrix P is computed via an auxiliary regression. The auxiliary rotation, in effect, uses a short order VAR to find the rotation. Specifically, $\bar{B}(1)$ is estimated from the representation (9) and then inverted to estimate $\Gamma(1)$.

We can rewrite (9) as

$$Y_t = \sum_{j=1}^{\infty} \theta_j \Delta Y_{t-j} + \Pi Y_{t-1} + u_t \tag{18}$$

where $\theta_j = -\sum_{i=j}^{\infty} B_i$ and $\Pi = B(1)$.

In practice, we truncate equation (18) at some finite horizon and estimate the coefficients by least squares. Then we can substitute $\left(I - \widehat{\Pi}\right)^{-1}$ for $\hat{\Gamma}(1)$ in (17) to estimate P.

For this reason, the question 'Is rotation estimated using a VAR?' is answered 'Yes' for auxiliary rotation and 'No' for direct rotation. The difficulties due to short-order VARs that arise in VAR inversion are, then, likely to be present using the auxiliary rotation method. This is born out by our simulations. Table 2 contains the bias and RMSE using the auxiliary rotation. There is a bias reduction relative to VAR inversion of only 17% and a negligible change in the RMSE. On the other hand, recall that under direct rotation, the bias was reduced by 74%.

5.2 Spectral Methods

Existing research has recognized a relationship between the rotation matrix P based on long-run restrictions and the spectral density function at frequency zero. These papers are first CEV (2006) and followed by Stock and Watson (2008). Based on this relationship, they present a new way to execute long-run restrictions.

From equations (11) and (12), the VAR inversion estimator can be expressed as:

$$X_t = \left[I - \hat{B}\left(L\right)\right]^{-1} \left[I - \hat{B}\left(1\right)\right] chol\left(\left[I - \hat{B}\left(1\right)\right]^{-1} \hat{\Omega}\left(\left[I - \hat{B}\left(1\right)\right]^{-1}\right)'\right) \varepsilon_t \tag{19}$$

CEV (2006) recognized that the term inside the chol operator is an estimate of the spectral density function evaluated at zero. To see this, note that the autocovariance function of X_t is:

$$G_x\left(z\right) = \left[I - \bar{B}\left(z\right)\right]^{-1} \Omega\left(\left[I - \bar{B}\left(z\right)\right]^{-1}\right)'$$

Next, the spectral density function is:

$$\begin{aligned} S_x\left(\omega\right) &= \left(2\pi\right)^{-1} G_x\left(e^{-i\omega}\right) \\ &= \left(2\pi\right)^{-1} \left[I - \bar{B}\left(e^{-i\omega}\right)\right]^{-1} \Omega\left(\left[I - \bar{B}\left(e^{-i\omega}\right)\right]^{-1}\right)' \end{aligned}$$

Evaluating the spectral density function at zero, we have:

$$S_x\left(0\right) = \left(2\pi\right)^{-1} \left[I - \bar{B}\left(1\right)\right]^{-1} \Omega\left(\left[I - \bar{B}\left(1\right)\right]^{-1}\right)' \tag{20}$$

The term inside the chol operator in (19) is an estimate of $S_x\left(0\right)$, modulus the $\left(2\pi\right)^{-1}$ term. As such, any consistent estimator for the frequency zero spectral density could be replace the corresponding term in (19) to create a new estimator of the structural impulse response function. While CEV estimate the long-run dynamics using the spectrum, they estimate the short-run dynamics using a VAR.

We estimate the spectrum using the Newey-West approach (1987) with a truncation parameter equal to 20. This is larger than their rule of thumb approach to the parameter which suggests 5.[14] The small sample results appear in table 2. There is a reduction in both bias and RMSE relative to VAR inversion.

Note that the spectral method performs even better than direct rotation. The spectral method results in a 28% reduction in RMSE relative to VAR inversion, whereas direct rotation implies only a 7% reduction. It results in a 79% reduction in bias, which is very similar to the reduction under direct rotation.

One minor drawback of the spectral method is that it requires the econometrician specify a bandwidth type and also a window size. This selection is not typically informed by economic theory.[15] Moreover, CEV document in a Monte Carlo study of data from a simulated business cycle model that the performance of the spectral method is sensitive to changes in the bandwidth specification. Nonetheless, spectral methods are a useful tool, and perhaps even more useful than our method, based on the results of this section.

The relative performances of the methods are consistent with our hypothesis: structural impulse responses are more reliable when long-run covariances are matched directly (either using a spectral method or local projections) rather than when they are fit based upon one period ahead forecasting models (VAR inversion and auxiliary rotation). We note that our hypothesis should be checked against alternative DSGE models besides the one considered in this paper.

5.3 Direct Rotation-VAR Inversion Hybrid

Next, we consider a hybrid of direct rotation and VAR inversion. With this approach, we estimate the first h reduced form MA coefficients with the direct rotation methodology, but instead of truncating the summation at h, we use the implied MA coefficients from the VAR inversion at horizons beyond h. Recall that the reduced form vector process has the following representation:

$$X_t = \Gamma\left(L\right) u_t. \tag{21}$$

To estimate the structural model, we must compute the sum of the reduced form MA coefficients $\Gamma(1)$. Using direct rotation, we assume that the first h MA coefficients estimated by local projections capture all of the dynamics of labor productivity that are due

[14]CEV (2006) use an untruncated Bartlett kernel. Stock (2006), in a comment on CEV, explains that an untruncated kernel is non-standard and asymptotically inconsistent.

[15]On the other hand, the key parameter selected in direct rotation, h, can be informed by economic theory. This would be the length of time that the researcher believes that non-technology shocks influence labor productivity.

to non technology shocks. Using the hybrid approach, we do not assume that these dynamics die out after h periods. Instead we estimate the sum of the reduced form MA coefficients from two components. First we estimate the first h coefficients from the direct rotation method. Second we estimate the sum of the MA coefficients by using a short order VAR. Finally, we replace the first h MA coefficients implied by the VAR with the coefficients estimated from direct rotation.[16]

First, we express the VAR in companion form:

$$Y_t = \Phi Y_{t-1} + \begin{bmatrix} I_r \\ 0_{k*(r-1),r} \end{bmatrix} u_t, \tag{22}$$

where $Y_t = \left[X_t', X_{t-1}', ...X_{t-k}'\right]'$.

Next, let $\Gamma_{VAR}^{H}(1)$ represent the sum of the first h reduced form coefficients implied by the VAR. This expression can be easily computed from the companion form VAR and is given by:

$$\Gamma_{VAR}^{H}(1) = \left[(I-\Phi)^{-1}\left(I-\Phi^{H}\right)\right]_{1:r,1:r} \tag{23}$$

This expression is the submatrix defined by the $1st$ through rth row and column of the matrix inside of the brackets. Then we construct an estimate of the sum of MA coefficients by the following sum:

$$\widehat{\Gamma}_{HYBRID}^{H}(1) = \left[I - \hat{B}(1)\right]^{-1} - \widehat{\Gamma}_{VAR}^{H}(1) + \widehat{\Gamma}_{MADR}^{H}(1), \tag{24}$$

where the hats denote that these terms are constructed from the estimated coefficients. The first two terms of the summation are constructed from the short order VAR while the last term is constructed from the direct rotation method.

Table 2 presents the bias and RMSE findings for the hybrid method. The three required parameters are set as: $h = 40$, $k = 4$ and $a = 4$. Relative to VAR inversion, the hybrid method increases bias by 32% and reduces RMSE by 25%.

As a summary, we provide a taxonomy of the above methods for executing long-run restrictions in table 5. Quantitatively, the key difference across methods is whether a VAR is used to estimate the rotations (i.e. long run variances). Table 2 compares bias and RMSE pairs for each estimator. In the growth model under consideration, the VAR poorly captures the long-run dynamics and therefore leads to bias and high RMSE. These

[16]We thank Frank Schorfheide for his suggestion of combining the standard method with our approach. The suggestion was motivated by his own work (Schorfheide 2008) that compares multi-step ahead forecasting with forecasts based on iterating forward a one-period ahead forecast. Direct rotation is based on the former, while VAR inversion is based on the latter.

are auxiliary rotation and VAR inversion.

Several methods use estimators with the explicit objective of fitting long-run covariances. Generally, these achieve smaller bias and lower RMSE. These are direct rotation and the spectral approach. The hybrid approach combines VAR inversion and direct rotation. In small sample, it performs similarly to direct rotation in terms of RMSE, however, it has a significantly larger bias.

5.4 Further Approaches to Executing Long-Run Restrictions

We are aware of three other papers that present alternative ways to execute long-run restrictions.

Finite Horizon VAR Identification

Francis, Owyang and Roush (2007), hereafter FOR, consider a related approach to long run identification. In their approach, a short order VAR is estimated, but identification is achieved by requiring that the effects of technology shocks on economic variables other than labor productivity die out after a finite time period. Instead of estimating the long run impact, FOR impose identification by maximizing the forecast error variance for productivity attributable to technology shocks at some finite horizon h.

FOR begin with a structural representation:

$$X_t = \Gamma\left(L\right) P\varepsilon_t. \tag{25}$$

The forecast error for X at horizon h is then:

$$\left(X_{t+h} - E_t\left(X_{t+h}\right)\right) = \sum_{i=1}^{h} \Gamma_{h-i} P\varepsilon_{t+i}, \tag{26}$$

and the forecast error variance is given by:

$$E\left[\left(X_{t+h} - E_t\left(X_{t+h}\right)\right)\left(X_{t+h} - E_t\left(X_{t+h}\right)\right)'\right] = E\left[\left(\sum_{i=1}^{h} \Gamma_{h-i} P\varepsilon_{t+i}\right)\left(\sum_{i=1}^{h} \Gamma_{h-i} P\varepsilon_{t+i}\right)'\right], \tag{27}$$

or,

$$E\left[\left(X_{t+h} - E_t\left(X_{t+h}\right)\right)\left(X_{t+h} - E_t\left(X_{t+h}\right)\right)'\right] = \sum_{i=1}^{h} \Gamma_{h-i} PP'\Gamma'_{h-i}. \tag{28}$$

In order to identify the effects of structural technology shocks we need a mapping from the reduced form shocks to the structural shocks, provided by the rotation matrix P.

If we only wish to know the effects of technology shocks the first column of the rotation matrix is sufficient. FOR identify the impact of technology shocks by choosing a column of P that maximizes the contribution of technology shocks to the h horizon forecast error variance (28).

FOR differs from our approach in two important respects. First, they still use a short order VAR to estimate the covariance properties of the underlying process. Second, they do not require that labor productivity has a unit root or that its unit root can be fully characterized by the technology process. Instead they use a short order VAR to capture the process' covariance properties and require only that technology shocks account for the maximum amount of forecast error variance for labor productivity at some finite horizon h.

Since FOR does not require that technology and labor productivity follows a unit root and does not require that non-technology shocks have no effect on labor productivity at horizon h, FOR achieve identification by minimizing the forecast error variance attributable to technology shocks, represented by the following sum:

$$E_t\left[\left(X_t - X_{t+h|t}\right)\left(X_t - X_{t+h|t}\right)'\right]_{ij} = \frac{e_i'\left[\sum_{i=1}^{h}\widehat{\Gamma}_{h-i}\widehat{P}e_je_j'\widehat{P}'\widehat{\Gamma}_{h-i}'\right]e_i}{e_i'\left[\widehat{\Gamma}_{h-i}'\widehat{\Omega}\widehat{\Gamma}_{h-i}'\right]e_i}. \tag{29}$$

Finally, to achieve identification, FOR must find an expression for P. To do so they require that the technology shock provides the maximum h-step ahead forecast error variance share for productivity[17]:

$$\max_{\alpha}\frac{e_i'\left[\sum_{i=1}^{h}\widehat{\Gamma}_{h-i}chol\left(\widehat{\Omega}\right)\alpha\alpha' chol\left(\widehat{\Omega}\right)'\widehat{\Gamma}_{h-i}'\right]e_i}{e_i'\left[\widehat{\Gamma}_{h-i}'\widehat{\Omega}\widehat{\Gamma}_{h-i}'\right]e_i}, \tag{30}$$

$$\text{s.t. } \alpha\alpha' = 1 \tag{31}$$

Given the value of α that maximizes the expression (30) FOR construct impulse responses for the ith variable to the jthshock as: $e_i\widehat{\Gamma}(L)\, chol\left(\widehat{\Omega}\right)\alpha$, where $chol\left(\widehat{\Omega}\right)\alpha$ represents the jth column of the rotation matrix P.

Full Spectral Technique

Mertens (2007) estimates structural impulse responses by first estimating a spectral

[17]Note that $\widehat{P}e_je_j'\widehat{P}' = chol\left(\widehat{\Omega}\right)\alpha\alpha' chol\left(\widehat{\Omega}\right)'$, where α is the vector associated with the maximum forecast error varaince share for productivity. See Faust (1998).

density function from the data and then constructing the implied moving average representation. It is related to CEV, who instead estimate the long-run dynamics of the system using the spectrum and the short-run dynamics using a VAR. Given the similarity of the Merten's approach to CEVs, we do not include it in our above comparison analysis.

State Space and VARMA Approaches

Kascha and Mertens (2008) use the state-space and VARMA approaches, in turn, to execute long-run restrictions. They note that there exists no model misspecification using either approach. They find, nonetheless, that the resulting estimates are heavily biased and very imprecise. Given the generally negative conclusion of these approaches, we did not use them in our comparison analysis.

6 Conclusion

Economists often represent macroeconomic time series as being determined by a sequence of random shocks. One way to give economic, rather than simply statistical, meaning to the shocks is with an identification assumption. One popular identification approach is to assume that a subset of shocks have no long-run effect on one or more of the macroeconomic variables.

Researchers have presented examples where techniques that use long-run restrictions can generate estimates with large bias and high root mean square error. We have shown how to modify the procedure such that it reduces the bias dramatically without a large impact on the root mean square error. This method, which we dub *direct rotation*, is based on estimating a sequence of multi-step ahead linear regressions. Since it relies only on linear regressions, it is easy to implement. Implementing the method is intuitive since the econometric specification decision amounts to the econometrician taking a concrete stand on the question: 'how long is the long run?'

As with other existing approaches, the direct rotation technique does lead to imprecise estimates in the stochastic growth model studied here. We believe that the imprecision is not a problem with the method perse. Rather, the imprecision arises because, for our particular calibration of the stochastic growth model, the shock that we seek to identify accounts for a very small amount of the variance of the relevant endogenous variables.

7 Bibliography

Altig, D. L. Christiano, M. Eichenbaum and J. Linde (2004), "Firm-Specific Capital, Nominal Rigidities and the Business Cycle," Northwestern University, manuscript.

Blanchard, O. and D. Quah (1989), "The Dynamic Effects of Aggregate Demand and Supply Disturbances," American Economic Review 79, 655-680.

Blanchard, O. and M. Watson (1986), "Are All Business Cycles Alike?" in *The American Business Cycle: Continuity and Change*, 123-156.

Chari, V.V., P. Kehoe and E. McGrattan (2005), "A Critique of Structural VARs Using Real Business Cycle Theory," Federal Reserve Bank of Minneapolis, manuscript.

Christiano, L., M. Eichenbaum and R. Vigfusson (2006), "Assessing Structural VARs," NBER Macroeconomics Annual.

Christiano, L., M. Eichenbaum and R. Vigfusson (2006), "Alternative Procedures for Estimating Vector Autoregressions Identified with Long-Run Restrictions,"Journal of the European Economics Association 4, 475-483.

Edge, R., T. Laubach and J. Williams (2003), "The Responses of Wages and Prices to Technology Shocks,"Federal Reserve Board, manuscript.

Erceg, C., L. Guerrierri and C. Gust (2003), "Can Long-Run Restrictions Identify Technology Shocks?"Journal of the European Economic Association 3, 1237-1278.

Faust, Jon. "The Robustness of Identified VAR Conclusions about Money." Carnegie-Rochester Conference Series on Public Policy, December 1998, 49(0), pp. 207-44.

Faust, J. and E. Leeper (1997), "When Do Long-Run Identifying Restrictions Give Reliable Results?" Journal of Business and Economic Statistics 15, 345-353.

Fisher, J. (2006), "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks,"Journal of Political Economy 114, 413-52.

Francis, N., M. Owyang and J. Roush (2007), "A Flexible Finite-Horizon Identification of Technology Shocks," Federal Reserve Bank of St. Louis, manuscript.

Francis, N. and V. Ramey (2005), "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited", Journal of Monetary Economics, 1379-1399.

Francis, N. and V. Ramey (2004), “A New Measure of Hours per Capita with Implications for the Hours Technology Debate,” Lehigh University, manuscript.

Gali, J. (1999), “Technology, Employment and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?” American Economic Review 89, 249-271.

Gali, J. D., Lopez Salido and J. Valles (2003), “Technology Shocks and Monetary Policy: Assessing the Fed’s Performance,” Journal of Monetary Economics 50, 723-743.

Hansen, G. and E. Prescott, (1995), “Economic Growth and Business Cycles,” from Frontiers in Business Cycle Research, Princeton Press.

Jorda, Ò. (2005), “Estimation and Inference of Impulse Responses by Local Projection,” American Economic Review 95, 161-182.

Jorda, Ò. (2007), “Simultaneous Confidence Regions for Impulse Responses,”Review of Economics and Statistics, forthcoming.

Kascha, C. and K. Mertens (2008), “Business Cycle Analysis and VARMA Models,” Journal of Economic Dynamics and Control, forthcoming.

Kehoe, P. (1993), “Comment on ’Assessing Structural VARs’,”NBER Macroeconomics Annual 2006, 73-96.

King, R., C. Plosser, J. Stock and M. Watson (1991), “Stochastic Trends and Economic Fluctuations,”American Economic Review 81, 819-840.

Lippi, M. and L. Reichlin (1993), “The Dynamic Effects of Aggregate Demand and Supply Disturbances: Comment,”American Economic Review 83, 644-652.

Marcellino, M., J. Stock and M. Watson (2006), “A Comparison of Direct and Iterated Multistep AR Methods for Forecasting Macroeconomic Time Series,”Journal of Econometrics 135, 499-526.

McGrattan, E. (2006), “Measurement with Minimal Theory,”FRB Minneapolis Working Paper.

Mertens, E. (2007), “Are Spectral Estimators Useful for Implementing Long-Run Restrictions in SVARs,”Study Center Gerzensee.

Ravenna, F. (2007), “Vector Autoregressions and Reduced form representations of DSGE Models,” Journal of Monetary Economics, forthcoming.

Schorfheide, F. (2008), "VAR forecasting under misspecification, "Journal of Econometrics (forthcoming).

Stock, J. and M. Watson (2008), "Notes from Summer Mini-course on Advances in Macroeconometrics."