

Assessing Structural VARs*

Lawrence J. Christiano,[†] Martin Eichenbaum,[‡] and Robert Vigfusson[§]

June 2006

Abstract

This paper analyzes the quality of VAR-based procedures for estimating the response of the economy to a shock. We focus on two key issues. First, do VAR-based confidence intervals accurately reflect the actual degree of sampling uncertainty associated with impulse response functions? Second, what is the size of bias relative to confidence intervals, and how do coverage rates of confidence intervals compare with their nominal size? We address these questions using data generated from a series of estimated dynamic, stochastic general equilibrium models. We organize most of our analysis around a particular question that has attracted a great deal of attention in the literature: How do hours worked respond to an identified shock? In all of our examples, as long as the variance in hours worked due to a given shock is above the remarkably low number of 1 percent, structural VARs perform well. This finding is true regardless of whether identification is based on short-run or long-run restrictions. Confidence intervals are wider in the case of long-run restrictions. Even so, long-run identified VARs can be useful for discriminating among competing economic models.

*The first two authors are grateful to the National Science Foundation for Financial Support. We thank Lars Hansen and our colleagues at the Federal Reserve Bank of Chicago and the Board of Governors for useful comments at various stages of this project. The views in this paper are solely those of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or its staff.

[†]Northwestern University, the Federal Reserve Bank of Chicago, and the NBER.

[‡]Northwestern University, the Federal Reserve Bank of Chicago, and the NBER.

[§]Federal Reserve Board of Governors.

1. Introduction

Sims's seminal paper *Macroeconomics and Reality* (1980) argued that procedures based on vector autoregression (VAR) would be useful to macroeconomists interested in constructing and evaluating economic models. Given a minimal set of identifying assumptions, structural VARs allow one to estimate the dynamic effects of economic shocks. The estimated impulse response functions provide a natural way to choose the parameters of a structural model and to assess the empirical plausibility of alternative models.¹

To be useful in practice, VAR-based procedures must have good sampling properties. In particular, they should accurately characterize the amount of information in the data about the effects of a shock to the economy. Also, they should accurately uncover the information that is there.

These considerations lead us to investigate two key issues. First, do VAR-based confidence intervals accurately reflect the actual degree of sampling uncertainty associated with impulse response functions? Second, what is the size of bias relative to confidence intervals, and how do coverage rates of confidence intervals compare with their nominal size?

We address these questions using data generated from a series of estimated dynamic, stochastic general equilibrium (DSGE) models. We consider real business cycle (RBC) models and the model in Altig, Christiano, Eichenbaum, and Linde (2005) (hereafter, ACEL) that embodies real and nominal frictions. We organize most of our analysis around a particular question that has attracted a great deal of attention in the literature: How do hours worked respond to an identified shock? In the case of the RBC model, we consider a neutral shock to technology. In the ACEL model, we consider two types of technology shocks as well as a monetary policy shock.

We focus our analysis on an unavoidable specification error that occurs when the data generating process is a DSGE model and the econometrician uses a VAR. In this case the true VAR is infinite ordered, but the econometrician must use a VAR with a finite number of lags.

We find that as long as the variance in hours worked due to a given shock is above the remarkably low number of 1 percent, VAR-based methods for recovering the response of hours to that shock have good sampling properties. Technology shocks account for a much larger fraction of the variance of hours worked in the ACEL model than in any of our estimated RBC models. Not surprisingly, inference about the effects of a technology shock on hours worked is much sharper when the ACEL model is the data generating mechanism.

Taken as a whole, our results support the view that structural VARs are a useful guide to constructing and evaluating DSGE models. Of course, as with any econometric procedure it is possible to find examples in which VAR-based procedures do not do well. Indeed, we present such an example based on an RBC model in which technology shocks account for less than 1 percent of the variance in hours worked. In this example, VAR-based methods work poorly in the sense that bias exceeds sampling uncertainty. Although instructive, the example is based on a model that fits the data poorly and so is unlikely to be of practical importance.

Having good sampling properties does not mean that structural VARs always deliver small

¹See for example Sims (1989), Eichenbaum and Evans (1995), Rotemberg and Woodford (1997), Gali (1999), Francis and Ramey (2004), Christiano, Eichenbaum, and Evans (2005), and Del Negro, Schorfheide, Smets, and Wouters (2005).

confidence intervals. Of course, it would be a Pyrrhic victory for structural VARs if the best one could say about them is that sampling uncertainty is always large and the econometrician will always know it. Fortunately, this is not the case. We describe examples in which structural VARs are useful for discriminating between competing economic models.

Researchers use two types of identifying restrictions in structural VARs. Blanchard and Quah (1989), Gali (1999), and others exploit the implications that many models have for the long-run effects of shocks.² Other authors exploit short-run restrictions.³ It is useful to distinguish between these two types of identifying restrictions to summarize our results.

We find that structural VARs perform remarkably well when identification is based on short-run restrictions. For all the specifications that we consider, the sampling properties of impulse response estimators are good and sampling uncertainty is small. This good performance obtains even when technology shocks account for as little as 0.5 percent of the variance in hours. Our results are comforting for the vast literature that has exploited short-run identification schemes to identify the dynamic effects of shocks to the economy. Of course, one can question the particular short-run identifying assumptions used in any given analysis. However, our results strongly support the view that if the relevant short-run assumptions are satisfied in the data generating process, then standard structural VAR procedures reliably uncover and identify the dynamic effects of shocks to the economy.

The main distinction between our short and long-run results is that the sampling uncertainty associated with estimated impulse response functions is substantially larger in the long-run case. In addition, we find some evidence of bias when the fraction of the variance in hours worked that is accounted for by technology shocks is very small. However, this bias is not large relative to sampling uncertainty as long as technology shocks account for at least 1 percent of the variance of hours worked. Still, the reason for this bias is interesting. We document that, when substantial bias exists, it stems from the fact that with long-run restrictions one requires an estimate of the sum of the VAR coefficients. The specification error involved in using a finite-lag VAR is the reason that in some of our examples, the sum of VAR coefficients is difficult to estimate accurately. This difficulty also explains why sampling uncertainty with long-run restrictions tends to be large.

The preceding observations led us to develop an alternative to the standard VAR-based estimator of impulse response functions. The only place the sum of the VAR coefficients appears in the standard strategy is in the computation of the zero-frequency spectral density of the data. Our alternative estimator avoids using the sum of the VAR coefficients by working with a nonparametric estimator of this spectral density. We find that in cases when the standard VAR procedure entails some bias, our adjustment virtually eliminates the bias.

Our results are related to a literature that questions the ability of long-run identified VARs to reliably estimate the dynamic response of macroeconomic variables to structural shocks.

²See, for example, Basu, Fernald, and Kimball (2004), Christiano, Eichenbaum, and Vigfusson (2003, 2004), Fisher (2006), Francis and Ramey (2004), King, Plosser, Stock and Watson (1991), Shapiro and Watson (1988) and Vigfusson (2004). Francis, Owyang, and Roush (2005) pursue a related strategy to identify a technology shock as the shock that maximizes the forecast error variance share of labor productivity at a long but finite horizon.

³This list is particularly long and includes at least Bernanke (1986), Bernanke and Blinder (1992), Bernanke and Mihov (1998), Blanchard and Perotti (2002), Blanchard and Watson (1986), Christiano and Eichenbaum (1992), Christiano, Eichenbaum and Evans (2005), Cushman and Zha (1997), Eichenbaum and Evans (1995), Hamilton (1997), Rotemberg and Woodford (1992), Sims (1986), and Sims and Zha (2006).

Perhaps the first critique of this sort was provided by Sims (1972). Although his paper was written before the advent of VARs, it articulates why estimates of the sum of regression coefficients may be distorted when there is specification error. Faust and Leeper (1997) and Pagan and Robertson (1998) make an important related critique of identification strategies based on long-run restrictions. More recently Erceg, Guerrieri, and Gust (2005) and Chari, Kehoe, and McGrattan (2005b) (henceforth CKM) also examine the reliability of VAR-based inference using long-run identifying restrictions.⁴ Our conclusions regarding the value of identified VARs differ sharply from those recently reached by CKM. One parameterization of the RBC model that we consider is identical to the one considered by CKM. This parameterization is included for pedagogical purposes only, as it is overwhelmingly rejected by the data.

The remainder of the paper is organized as follows. Section 2 presents the versions of the RBC models that we use in our analysis. Section 3 discusses our results for standard VAR-based estimators of impulse response functions. Section 4 analyzes the differences between short and long-run restrictions. Section 5 discusses the relation between our work and the recent critique of VARs offered by CKM. Section 6 summarizes the ACEL model and reports its implications for VARs. Section 7 contains concluding comments.

2. A Simple RBC Model

In this section, we display the RBC model that serves as one of the data generating processes in our analysis. In this model the only shock that affects labor productivity in the long-run is a shock to technology. This property lies at the core of the identification strategy used by King, et al (1991), Gali (1999) and other researchers to identify the effects of a shock to technology. We also consider a variant of the model which rationalizes short run restrictions as a strategy for identifying a technology shock. In this variant, agents choose hours worked before the technology shock is realized. We describe the conventional VAR-based strategies for estimating the dynamic effect on hours worked of a shock to technology. Finally, we discuss parameterizations of the RBC model that we use in our experiments.

2.1. The Model

The representative agent maximizes expected utility over per capita consumption, c_t , and per capita hours worked, l_t :

$$E_0 \sum_{t=0}^{\infty} (\beta(1+\gamma))^t \left[\log c_t + \psi \frac{(1-l_t)^{1-\sigma} - 1}{1-\sigma} \right],$$

subject to the budget constraint:

$$c_t + (1 + \tau_{x,t}) i_t \leq (1 - \tau_{l,t}) w_t l_t + r_t k_t + T_t,$$

where

$$i_t = (1 + \gamma) k_{t+1} - (1 - \delta) k_t.$$

⁴See also Fernandez-Villaverdez, Rubio-Ramirez, and Sargent (2005) who investigate the circumstances in which the economic shocks are recoverable from the VAR disturbances. They provide a simple matrix algebra check to assess recoverability. They identify models in which the conditions are satisfied and other models in which they are not.

Here, k_t denotes the per capita capital stock at the beginning of period t , w_t is the wage rate, r_t is the rental rate on capital, $\tau_{x,t}$ is an investment tax, $\tau_{l,t}$ is the tax rate on labor income, $\delta \in (0, 1)$ is the depreciation rate on capital, γ is the growth rate of the population, T_t represents lump-sum taxes and $\sigma > 0$ is a curvature parameter.

The representative competitive firm's production function is:

$$y_t = k_t^\alpha (Z_t l_t)^{1-\alpha},$$

where Z_t is the time t state of technology and $\alpha \in (0, 1)$. The stochastic processes for the shocks are:

$$\begin{aligned} \log z_t &= \mu_z + \sigma_z \varepsilon_t^z \\ \tau_{l,t+1} &= (1 - \rho_l) \tau_l + \rho_l \tau_{l,t} + \sigma_l \varepsilon_{t+1}^l \\ \tau_{x,t+1} &= (1 - \rho_x) \tau_x + \rho_x \tau_{x,t} + \sigma_x \varepsilon_{t+1}^x, \end{aligned} \tag{2.1}$$

where $z_t = Z_t/Z_{t-1}$. In addition, ε_t^z , ε_t^l , and ε_t^x are independently and identically distributed (i.i.d.) random variables with mean zero and unit standard deviation. The parameters, σ_z , σ_l , and σ_x are non-negative scalars. The constant, μ_z , is the mean growth rate of technology, τ_l is the mean labor tax rate, and τ_x is the mean tax on capital. We restrict the autoregressive coefficients, ρ_l and ρ_x , to be less than unity in absolute value.

Finally, the resource constraint is:

$$c_t + (1 + \gamma) k_{t+1} - (1 - \delta) k_t \leq y_t.$$

We consider two versions of the model, differentiated according to timing assumptions. In the *standard* or *nonrecursive version*, all time t decisions are taken after the realization of the time t shocks. This is the conventional assumption in the RBC literature. In the *recursive version* of the model the timing assumptions are as follows. First, $\tau_{l,t}$ is observed, and then labor decisions are made. Second, the other shocks are realized and agents make their investment and consumption decisions.

2.2. Relation of the RBC Model to VARs

We now discuss the relation between the RBC model and a VAR. Specifically, we establish conditions under which the reduced form of the RBC model is a VAR with disturbances that are linear combinations of the economic shocks. Our exposition is a simplified version of the discussion in Fernandez-Villaverde, Rubio-Ramirez, and Sargent (2005) (see especially their section III). We include this discussion because it frames many of the issues that we address. Our discussion applies to both the standard and the recursive versions of the model.

We begin by showing how to put the reduced form of the RBC model into a state-space, observer form. Throughout, we analyze the log-linear approximations to model solutions. Suppose the variables of interest in the RBC model are denoted by X_t . Let s_t denote the vector of exogenous economic shocks and let \hat{k}_t denote the percent deviation from steady state of the capital stock, after scaling by Z_t .⁵ The approximate solution for X_t is given by:

$$X_t = a_0 + a_1 \hat{k}_t + a_2 \hat{k}_{t-1} + b_0 s_t + b_1 s_{t-1}, \tag{2.2}$$

⁵Let $\tilde{k}_t = k_t/Z_{t-1}$. Then, $\hat{k}_t = (\tilde{k}_t - \tilde{k})/\tilde{k}$, where \tilde{k} denotes the value of \tilde{k}_t in nonstochastic steady state.

where

$$\hat{k}_{t+1} = A\hat{k}_t + Bs_t. \quad (2.3)$$

Also, s_t has the law of motion:

$$s_t = Ps_{t-1} + Q\varepsilon_t, \quad (2.4)$$

where ε_t is a vector of i.i.d. fundamental economic disturbances. The parameters of (2.2) and (2.3) are functions of the structural parameters of the model.

The ‘state’ of the system is composed of the variables on the right side of (2.2):

$$\xi_t = \begin{pmatrix} \hat{k}_t \\ \hat{k}_{t-1} \\ s_t \\ s_{t-1} \end{pmatrix}.$$

The law of motion of the state is:

$$\xi_t = F\xi_{t-1} + D\varepsilon_t, \quad (2.5)$$

where F and D are constructed from A , B , Q , P . The econometrician observes the vector of variables, Y_t . We assume Y_t is equal to X_t plus iid measurement error, v_t , which has diagonal variance-covariance, R . Then:

$$Y_t = H\xi_t + v_t. \quad (2.6)$$

Here, H is defined so that $X_t = H\xi_t$, that is, relation (2.2) is satisfied. In (2.6) we abstract from the constant term. Hamilton (1994, section 13.4) shows how the system formed by (2.5) and (2.6) can be used to construct the exact Gaussian density function for a series of observations, Y_1, \dots, Y_T . We use this approach when we estimate versions of the RBC model.

We now use (2.5) and (2.6) to establish conditions under which the reduced form representation for X_t implied by the RBC model is a VAR with disturbances that are linear combinations of the economic shocks. In this discussion, we set $v_t = 0$, so that $X_t = Y_t$. In addition, we assume that the number of elements in ε_t coincides with the number of elements in Y_t .

We begin by substituting (2.5) into (2.6) to obtain:

$$Y_t = HF\xi_{t-1} + C\varepsilon_t, \quad C \equiv HD.$$

Our assumption on the dimensions of Y_t and ε_t implies that the matrix C is square. In addition, we assume C is invertible. Then:

$$\varepsilon_t = C^{-1}Y_t - C^{-1}HF\xi_{t-1}. \quad (2.7)$$

Substituting (2.7) into (2.5), we obtain:

$$\xi_t = M\xi_{t-1} + DC^{-1}Y_t,$$

where

$$M = [I - DC^{-1}H] F. \quad (2.8)$$

As long as the eigenvalues of M are less than unity in absolute value,

$$\xi_t = DC^{-1}Y_t + MDC^{-1}Y_{t-1} + M^2DC^{-1}Y_{t-2} + \dots \quad (2.9)$$

Using (2.9) to substitute out for ξ_{t-1} in (2.7), we obtain:

$$\varepsilon_t = C^{-1}Y_t - C^{-1}HF [DC^{-1}Y_{t-1} + MDC^{-1}Y_{t-2} + M^2DC^{-1}Y_{t-3} + \dots],$$

or, after rearranging:

$$Y_t = B_1Y_{t-1} + B_2Y_{t-2} + \dots + u_t, \quad (2.10)$$

where

$$u_t = C\varepsilon_t \quad (2.11)$$

$$B_j = HFM^{j-1}DC^{-1}, \quad j = 1, 2, \dots \quad (2.12)$$

Expression (2.10) is an infinite-order VAR, because u_t is orthogonal to Y_{t-j} , $j \geq 1$.

Proposition 2.1. (Fernandez-Villaverde, Rubio-Ramirez, and Sargent) *If C is invertible and the eigenvalues of M are less than unity in absolute value, then the RBC model implies:*

- Y_t has the infinite-order VAR representation in (2.10)
- The linear one-step-ahead forecast error Y_t given past Y_t 's is u_t , which is related to the economic disturbances by (2.11)
- The variance-covariance of u_t is CC'
- The sum of the VAR lag matrices is given by:

$$B(1) \equiv \sum_{j=1}^{\infty} B_j = HF [I - M]^{-1} DC^{-1}.$$

We will use the last of these results below.

Relation (2.10) indicates why researchers interested in constructing DSGE models find it useful to analyze VARs. At the same time, this relationship clarifies some of the potential pitfalls in the use of VARs. First, in practice the econometrician must work with finite lags. Second, the assumption that C is square and invertible may not be satisfied. Whether C satisfies these conditions depends on how Y_t is defined. Third, significant measurement errors may exist. Fourth, the matrix, M , may not have eigenvalues inside the unit circle. In this case, the economic shocks are not recoverable from the VAR disturbances.⁶ Implicitly, the econometrician who works with VARs assumes that these pitfalls are not quantitatively important.

2.3. VARs in Practice and the RBC Model

We are interested in the use of VARs as a way to estimate the response of X_t to economic shocks, i.e., elements of ε_t . In practice, macroeconomists use a version of (2.10) with finite lags, say q . A researcher can estimate B_1, \dots, B_q and $V = Eu_tu_t'$. To obtain the impulse response functions, however, the researcher needs the B_i 's and the column of C corresponding to the shock in ε_t that

⁶For an early example, see Hansen and Sargent (1980, footnote 12). Sims and Zha (forthcoming) discuss the possibility that, although a given economic shock may not lie exactly in the space of current and past Y_t , it may nevertheless be 'close'. They discuss methods to detect this case.

is of interest. However, to compute the required column of C requires additional identifying assumptions. In practice, two types of assumptions are used. Short-run assumptions take the form of direct restrictions on the matrix C . Long-run assumptions place indirect restrictions on C that stem from restrictions on the long-run response of X_t to a shock in an element of ε_t . In this section we use our RBC model to discuss these two types of assumptions and how they are imposed on VARs in practice.

2.3.1. The Standard Version of the Model

The log-linearized equilibrium laws of motion for capital and hours in this model can be written as follows:

$$\log \hat{k}_{t+1} = \gamma_0 + \gamma_k \log \hat{k}_t + \gamma_z \log z_t + \gamma_l \tau_{l,t} + \gamma_x \tau_{x,t}, \quad (2.13)$$

and

$$\log l_t = a_0 + a_k \log \hat{k}_t + a_z \log z_t + a_l \tau_{l,t} + a_x \tau_{x,t}. \quad (2.14)$$

From (2.13) and (2.14), it is clear that all shocks have only a temporary effect on l_t and \hat{k}_t .⁷ The only shock that has a permanent effect on labor productivity, $a_t \equiv y_t/l_t$, is ε_t^z . The other shocks do not have a permanent effect on a_t . Formally, this *exclusion restriction* is:

$$\lim_{j \rightarrow \infty} [E_t a_{t+j} - E_{t-1} a_{t+j}] = f(\varepsilon_t^z \text{ only}). \quad (2.15)$$

In our linear approximation to the model solution f is a linear function. The model also implies the *sign restriction* that f is an increasing function. In (2.15), E_t is the expectation operator, conditional on the information set $\Omega_t = (\log \hat{k}_{t-s}, \log z_{t-s}, \tau_{l,t-s}, \tau_{x,t-s}; s \geq 0)$.

In practice, researchers impose the exclusion and sign restrictions on a VAR to compute ε_t^z and identify its dynamic effects on macroeconomic variables. Consider the $N \times 1$ vector, Y_t . The VAR for Y_t is given by:

$$\begin{aligned} Y_{t+1} &= B(L)Y_t + u_{t+1}, \quad Eu_t u_t' = V, \\ B(L) &\equiv B_1 + B_2 L + \dots + B_q L^{q-1}, \\ Y_t &= \begin{pmatrix} \Delta \log a_t \\ \log l_t \\ x_t \end{pmatrix}. \end{aligned} \quad (2.16)$$

Here, x_t is an additional vector of variables that may be included in the VAR. Motivated by the type of reasoning discussed in the previous subsection, researchers assume that the fundamental economic shocks are related to u_t as follows:

$$u_t = C\varepsilon_t, \quad E\varepsilon_t \varepsilon_t' = I, \quad CC' = V. \quad (2.17)$$

Without loss of generality, we assume that the first element in ε_t is ε_t^z . We can easily verify that:

$$\lim_{j \rightarrow \infty} [\tilde{E}_t a_{t+j} - \tilde{E}_{t-1} a_{t+j}] = \tau [I - B(1)]^{-1} C\varepsilon_t, \quad (2.18)$$

⁷Cooley and Dwyer (1998) argue that in the standard RBC model, if technology shocks have a unit root, then per capita hours worked will be difference stationary. This claim, which plays an important role in their analysis of VARs, is incorrect.

where τ is a row vector with all zeros, but with unity in the first location. Here:

$$B(1) \equiv B_1 + \dots + B_q.$$

Also, \tilde{E}_t is the expectation operator, conditional on $\tilde{\Omega}_t = \{Y_t, \dots, Y_{t-q+1}\}$. As mentioned above, to compute the dynamic effects of ε_t^z , we require B_1, \dots, B_q and C_1 , the first column of C .

The symmetric matrix, V , and the B_i 's can be computed using ordinary least squares regressions. However, the requirement that $CC' = V$ is not sufficient to determine a unique value of C_1 . Adding the exclusion and sign restrictions does uniquely determine C_1 . Relation (2.18) implies that these restrictions are:

$$\text{exclusion restriction: } [I - B(1)]^{-1} C = \begin{bmatrix} \text{number} & \underline{0} \\ \text{numbers} & \text{numbers} \end{bmatrix},$$

where $\underline{0}$ is a row vector and

$$\text{sign restriction: } (1, 1) \text{ element of } [I - B(1)]^{-1} C \text{ is positive.}$$

There are many matrices, C , that satisfy $CC' = V$ as well as the exclusion and sign restrictions. It is well-known that the first column, C_1 , of each of these matrices is the same. We prove this result here, because elements of the proof will be useful to analyze our simulation results. Let

$$D \equiv [I - B(1)]^{-1} C.$$

Let $S_Y(\omega)$ denote the spectral density of Y_t at frequency ω that is implied by the q^{th} -order VAR. Then:

$$DD' = [I - B(1)]^{-1} V [I - B(1)']^{-1} = S_Y(0). \quad (2.19)$$

The exclusion restriction requires that D have a particular pattern of zeros:

$$D = \begin{bmatrix} d_{11} & \mathbf{0} \\ 1 \times 1 & 1 \times (N-1) \\ D_{21} & D_{22} \\ (N-1) \times 1 & (N-1) \times (N-1) \end{bmatrix},$$

so that

$$DD' = \begin{bmatrix} d_{11}^2 & d_{11}D'_{21} \\ D_{21}d_{11} & D_{21}D'_{21} + D_{22}D'_{22} \end{bmatrix} = \begin{bmatrix} S_Y^{11}(0) & S_Y^{21}(0)' \\ S_Y^{21}(0) & S_Y^{22}(0) \end{bmatrix},$$

where

$$S_Y(\omega) \equiv \begin{bmatrix} S_Y^{11}(\omega) & S_Y^{21}(\omega)' \\ S_Y^{21}(\omega) & S_Y^{22}(\omega) \end{bmatrix}.$$

The exclusion restriction implies that

$$d_{11}^2 = S_Y^{11}(0), \quad D_{21} = S_Y^{21}(0) / d_{11}. \quad (2.20)$$

There are two solutions to (2.20). The sign restriction

$$d_{11} > 0 \quad (2.21)$$

selects one of the two solutions to (2.20). So, the first column of D , D_1 , is uniquely determined. By our definition of C , we have

$$C_1 = [I - B(1)] D_1. \quad (2.22)$$

We conclude that C_1 is uniquely determined.

2.3.2. The Recursive Version of the Model

In the recursive version of the model, the policy rule for labor involves $\log z_{t-1}$ and $\tau_{x,t-1}$ because these variables help forecast $\log z_t$ and $\tau_{x,t}$:

$$\log l_t = a_0 + a_k \log \hat{k}_t + \tilde{a}_l \tau_{l,t} + \tilde{a}'_z \log z_{t-1} + \tilde{a}'_x \tau_{x,t-1}.$$

Because labor is a state variable at the time the investment decision is made, the equilibrium law of motion for \hat{k}_{t+1} is:

$$\begin{aligned} \log \hat{k}_{t+1} = & \gamma_0 + \gamma_k \log \hat{k}_t + \tilde{\gamma}_z \log z_t + \tilde{\gamma}_l \tau_{l,t} + \tilde{\gamma}_x \tau_{x,t} \\ & + \tilde{\gamma}'_z \log z_{t-1} + \tilde{\gamma}'_x \tau_{x,t-1}. \end{aligned}$$

As in the standard model, the only shock that affects a_t in the long run is a shock to technology. So, the long-run identification strategy discussed in section 2.3.1 applies to the recursive version of the model. However, an alternative procedure for identifying ε_t^z applies to this version of the model. We refer to this alternative procedure as the ‘short-run’ identification strategy because it involves recovering ε_t^z using only the realized one-step-ahead forecast errors in labor productivity and hours, as well as the second moment properties of those forecast errors.

Let $u_{\Omega,t}^a$ and $u_{\Omega,t}^l$ denote the population one-step-ahead forecast errors in a_t and $\log l_t$, conditional on the information set, Ω_{t-1} . The recursive version of the model implies that

$$u_{\Omega,t}^a = \alpha_1 \varepsilon_t^z + \alpha_2 \varepsilon_t^l, \quad u_{\Omega,t}^l = \gamma \varepsilon_t^l,$$

where $\alpha_1 > 0$, α_2 , and γ are functions of the model parameters. The projection of $u_{\Omega,t}^a$ on $u_{\Omega,t}^l$ is given by

$$u_{\Omega,t}^a = \beta u_{\Omega,t}^l + \alpha_1 \varepsilon_t^z, \quad \text{where } \beta = \frac{\text{cov}(u_{\Omega,t}^a, u_{\Omega,t}^l)}{\text{var}(u_{\Omega,t}^l)}. \quad (2.23)$$

Because we normalize the standard deviation of ε_t^z to unity, α_1 is given by:

$$\alpha_1 = \sqrt{\text{var}(u_{\Omega,t}^a) - \beta^2 \text{var}(u_{\Omega,t}^l)}.$$

In practice, we implement the previous procedure using the one-step-ahead forecast errors generated from a VAR in which the variables in Y_t are ordered as follows:

$$Y_t = \begin{pmatrix} \log l_t \\ \Delta \log a_t \\ x_t \end{pmatrix}.$$

We write the vector of VAR one-step-ahead forecast errors, u_t , as:

$$u_t = \begin{pmatrix} u_t^l \\ u_t^a \\ u_t^x \end{pmatrix}.$$

We identify the technology shock with the second element in ε_t in (2.17). To compute the dynamic response of the variables in Y_t to the technology shock we need B_1, \dots, B_q in (2.16)

and the second column, C_2 , of the matrix C , in (2.17). We obtain C_2 in two steps. First, we identify the technology shock using:

$$\varepsilon_t^z = \frac{1}{\hat{\alpha}_1} \left(u_t^a - \hat{\beta} u_t^l \right),$$

where

$$\hat{\beta} = \frac{\text{cov}(u_t^a, u_t^l)}{\text{var}(u_t^l)}, \quad \hat{\alpha}_1 = \sqrt{\text{var}(u_t^a) - \hat{\beta}^2 \text{var}(u_t^l)}.$$

The required variances and covariances are obtained from the estimate of V in (2.16). Second, we regress u_t on ε_t^z to obtain:⁸

$$C_2 = \begin{pmatrix} \frac{\text{cov}(u_t^l, \varepsilon_t^z)}{\text{var}(\varepsilon_t^z)} \\ \frac{\text{cov}(u_t^a, \varepsilon_t^z)}{\text{var}(\varepsilon_t^z)} \\ \frac{\text{cov}(u_t^x, \varepsilon_t^z)}{\text{var}(\varepsilon_t^z)} \end{pmatrix} = \begin{pmatrix} 0 \\ \hat{\alpha}_1 \\ \frac{1}{\hat{\alpha}_1} \left(\text{cov}(u_t^x, u_t^a) - \hat{\beta} \text{cov}(u_t^x, u_t^l) \right) \end{pmatrix}.$$

2.4. Parameterization of the Model

We consider different specifications of the RBC model that are distinguished by the parameterization of the laws of motion of the exogenous shocks. In all specifications we assume, as in CKM, that:

$$\begin{aligned} \beta &= 0.98^{1/4}, \quad \theta = 0.33, \quad \delta = 1 - (1 - .06)^{1/4}, \quad \psi = 2.5, \quad \gamma = 1.01^{1/4} - 1 \\ \tau_x &= 0.3, \quad \tau_l = 0.242, \quad \mu_z = 1.016^{1/4} - 1, \quad \sigma = 1. \end{aligned} \quad (2.24)$$

2.4.1. Our MLE Parameterizations

We estimate two versions of our model. In the *two-shock maximum likelihood estimation (MLE) specification* we assume that $\sigma_x = 0$, so that there are two shocks, $\tau_{l,t}$ and $\log z_t$. We estimate the parameters ρ_l , σ_l , and σ_z , by maximizing the Gaussian likelihood function of the vector, $X_t = (\Delta \log y_t, \log l_t)'$, subject to (2.24).⁹ Our results are given by:

$$\begin{aligned} \log z_t &= \mu_z + 0.00953 \varepsilon_t^z, \\ \tau_{l,t} &= (1 - 0.986) \bar{\tau}_l + 0.986 \tau_{l,t-1} + 0.0056 \varepsilon_t^l. \end{aligned}$$

The *three-shock MLE specification* incorporates the investment tax shock, $\tau_{x,t}$, into the model. We estimate the three-shock MLE version of the model by maximizing the Gaussian likelihood function of the vector, $X_t = (\Delta \log y_t, \log l_t, \Delta \log i_t)'$, subject to the parameter values in (2.24). The results are:

$$\begin{aligned} \log z_t &= \mu_z + 0.00968 \varepsilon_t^z, \\ \tau_{l,t} &= (1 - 0.9994) \tau_l + 0.9994 \tau_{l,t-1} + 0.00631 \varepsilon_t^l, \\ \tau_{x,t} &= (1 - 0.9923) \tau_x + 0.9923 \tau_{x,t-1} + 0.00963 \varepsilon_t^x. \end{aligned}$$

⁸We implement the procedure for estimating C_2 by computing $CC' = V$, where C is the lower triangular Cholesky decomposition of V , and setting C_2 equal to the second column of C .

⁹We use the standard Kalman filter strategy discussed in Hamilton (1994, section 13.4). We remove the sample mean from X_t prior to estimation and set the measurement error in the Kalman filter system to zero, i.e., $R = 0$ in (2.6).

The estimated values of ρ_x and ρ_l are close to unity. This finding is consistent with other research that also reports that shocks in estimated general equilibrium models exhibit high degrees of serial correlation.¹⁰

2.4.2. CKM Parameterizations

The *two-shock CKM specification* has two shocks, z_t and $\tau_{l,t}$. These shocks have the following time series representations:

$$\begin{aligned}\log z_t &= \mu_z + 0.0131\varepsilon_t^z, \\ \tau_{l,t} &= (1 - 0.952)\tau_l + 0.952\tau_{l,t-1} + 0.0136\varepsilon_t^l.\end{aligned}$$

The *three-shock CKM specification* adds an investment shock, $\tau_{x,t}$, to the model, and has the following law of motion:

$$\tau_{x,t} = (1 - 0.98)\tau_x + 0.98\tau_{x,t-1} + 0.0123\varepsilon_t^x. \quad (2.25)$$

As in our specifications, CKM obtain their parameter estimates using maximum likelihood methods. However, their estimates are very different from ours. For example, the variances of the shocks are larger in the two-shock CKM specification than in our MLE specification. Also, the ratio of σ_l^2 to σ_z^2 is nearly three times larger in the two-shock CKM specification than in our two-shock MLE specification. Section 5 below discusses the reasons for these differences.

2.5. The Importance of Technology Shocks for Hours Worked

Table 1 reports the contribution, V_h , of technology shocks to three different measures of the volatility in the log of hours worked: (i) the variance of the log hours, (ii) the variance of HP-filtered, log hours and (iii) the variance in the one-step-ahead forecast error in log hours.¹¹ With one exception, we compute the analogous statistics for log output. The exception is (i), for which we compute the contribution of technology shocks to the variance of the growth rate of output.

The key result in this table is that technology shocks account for a very small fraction of the volatility in hours worked. When V_h is measured according to (i), it is always below 4 percent. When V_h is measured using (ii) or (iii) it is always below 8 percent. For both (ii) and (iii), in the CKM specifications, V_h is below 2 percent.¹² Consistent with the RBC literature, the table also shows that technology accounts for a much larger movement in output.

¹⁰See, for example, Christiano (1988), Christiano, et al. (2004), and Smets and Wouters (2003).

¹¹We compute forecast error variances based on a four lag VAR. The variables in the VAR depend on whether the calculations correspond to the two or three shock model. In the case of the two-shock model, the VAR has two variables, output growth and log hours. In the case of the three-shock model, the VAR has three variables: output growth, log hours and the log of the investment to output ratio. Computing V_h requires estimating VARs in artificial data generated with all shocks, as well as in artificial data generated with only the technology shock. In the latter case, the one-step ahead forecast error from the VAR is well defined, even though the VAR coefficients themselves are not well defined due to multicollinearity problems.

¹²When we measure V_h according to (i), V_h drops from 3.73 in the two-shock MLE model to 0.18 in the three-shock MLE model. The analogous drop in V_h is an order of magnitude smaller when V_h is measured using (ii) or (iii). The reason for this difference is that ρ_l goes from 0.986 in the two-shock MLE model to 0.9994 in the three-shock MLE model. In the latter specification there is a near-unit root in $\tau_{l,t}$, which translates into a near-unit root in hours worked. As a result, the variance of hours worked becomes very large at the low frequencies. The near-unit root in τ_{lt} has less of an effect on hours worked at high and business cycle frequencies.

Figure 1 displays visually how unimportant technology shocks are for hours worked. The top panel displays two sets of 180 artificial observations on hours worked, simulated using the standard two-shock MLE specification. The volatile time series shows how log hours worked evolve in the presence of shocks to both z_t and $\tau_{l,t}$. The other time series shows how log hours worked evolve in response to just the technology shock, z_t . The bottom panel is the analog of the top figure when the data are generated using the standard two-shock CKM specification.

3. Results Based on RBC Data Generating Mechanisms

In this section we analyze the properties of conventional VAR-based strategies for identifying the effects of a technology shock on hours worked. We focus on the bias properties of the impulse response estimator, and on standard procedures for estimating sampling uncertainty.

We use the RBC model parameterizations discussed in the previous section as the data generating processes. For each parameterization, we simulate 1,000 data sets of 180 observations each. The shocks ε_t^z , ε_t^l , and possibly ε_t^x , are drawn from *i.i.d.* standard normal distributions. For each artificial data set, we estimate a four-lag VAR. The average, across the 1,000 data sets, of the estimated impulse response functions, allows us to assess bias.

For each data set we also estimate two different confidence intervals: a percentile-based confidence interval and a standard-deviation based confidence interval.¹³ We construct the intervals using the following bootstrap procedure. Using random draws from the fitted VAR disturbances, we use the estimated four lag VAR to generate 200 synthetic data sets, each with 180 observations. For each of these 200 synthetic data sets we estimate a new VAR and impulse response function. For each artificial data set the percentile-based confidence interval is defined as the top 2.5 percent and bottom 2.5 percent of the estimated coefficients in the dynamic response functions. The standard-deviation-based confidence interval is defined as the estimated impulse response plus or minus two standard deviations where the standard deviations are calculated across the 200 simulated estimated coefficients in the dynamic response functions.

We assess the accuracy of the confidence interval estimators in two ways. First, we compute the coverage rate for each type of confidence interval. This rate is the fraction of times, across the 1,000 data sets simulated from the economic model, that the confidence interval contains the relevant true coefficient. If the confidence intervals were perfectly accurate, the coverage rate would be 95 percent. Second, we provide an indication of the actual degree of sampling uncertainty in the VAR-based impulse response functions. In particular, we report centered 95 percent probability intervals for each lag in our impulse response function estimators.¹⁴ If the confidence intervals were perfectly accurate, they should on average coincide with the boundary of the 95 percent probability interval.

When we generate data from the two-shock MLE and CKM specifications, we set $Y_t =$

¹³Sims and Zha (1999) refer to what we call the percentile-based confidence interval as the ‘other-percentile bootstrap interval’. This procedure has been used in several studies, such as Blanchard and Quah (1989), Christiano, Eichenbaum, and Evans (1999), Francis and Ramey (2004), McGrattan (2006), and Runkle (1987). The standard-deviation based confidence interval has been used by other researchers, such as Christiano Eichenbaum, and Evans (2005), Gali (1999), and Gali and Rabanal (2004).

¹⁴For each lag starting at the impact period, we ordered the 1,000 estimated impulse responses from smallest to largest. The lower and upper boundaries correspond to the 25th and the 975th impulses in this ordering.

$(\Delta \log a_t, \log l_t)'$. When we generate data from the three-shock MLE and CKM specifications, we set $Y_t = (\Delta \log a_t, \log l_t, \log i_t/y_t)'$.

3.1. Short-Run Identification

Results for the two- and three- Shock MLE Specifications

Figure 2 reports results generated from four different parameterizations of the recursive version of the RBC model. In each panel, the solid line is the average estimated impulse response function for the 1,000 data sets simulated using the indicated economic model. For each model, the starred line is the true impulse response function of hours worked. In each panel, the gray area defines the centered 95 percent probability interval for the estimated impulse response functions. The stars with no line indicate the average percentile-based confidence intervals across the 1,000 data sets. The circles with no line indicate the average standard-deviation-based confidence intervals.

Figures 3 and 4 graph the coverage rates for the percentile-based and standard-deviation-based confidence intervals. For each case we graph how often, across the 1,000 data sets simulated from the economic model, the econometrician's confidence interval contains the relevant coefficient of the true impulse response function.

The 1,1 panel in Figure 2 exhibits the properties of the VAR-based estimator of the response of hours to a technology shock when the data are generated by the two-shock MLE specification. The 2,1 panel corresponds to the case when the data generating process is the three-shock MLE specification.

The panels have two striking features. First, there is essentially no evidence of bias in the estimated impulse response functions. In all cases, the solid lines are very close to the starred lines. Second, an econometrician would not be misled in inference by using standard procedures for constructing confidence intervals. The circles and stars are close to the boundaries of the gray area. The 1,1 panels in Figures 3 and 4 indicate that the coverage rates are roughly 90 percent. So, with high probability, VAR-based confidence intervals include the true value of the impulse response coefficients.

Results for the CKM Specification

The second column of Figure 2 reports the results when the data generating process is given by variants of the CKM specification. The 1,2 and 2,1 panels correspond to the two and three-shock CKM specification, respectively.

The second column of Figure 2 contains the same striking features as the first column. There is very little bias in the estimated impulse response functions. In addition, the average value of the econometrician's confidence interval coincides closely with the actual range of variation in the impulse response function (the gray area). Coverage rates, reported in the 1,2 panels of Figures 3 and 4, are roughly 90 percent. These rates are consistent with the view that VAR-based procedures lead to reliable inference.

A comparison of the gray areas across the first and second columns of Figure 2, clearly indicates that more sampling uncertainty occurs when the data are generated from the CKM specifications than when they are generated from the MLE specifications (the gray areas are wider). VAR-based confidence intervals detect this fact.

3.2. Long-run Identification

Results for the two- and three- Shock MLE Specifications

The first and second rows of column 1 in Figure 5 exhibit our results when the data are generated by the two- and three- shock MLE specifications. Once again there is virtually no bias in the estimated impulse response functions and inference is accurate. The coverage rates associated with the percentile-based confidence intervals are very close to 95 percent (see Figure 3). The coverage rates for the standard-deviation-based confidence intervals are somewhat lower, roughly 80 percent (see Figure 4). The difference in coverage rates can be seen in Figure 5, which shows that the stars are shifted down slightly relative to the circles. Still, the circles and stars are very good indicators of the boundaries of the gray area, although not quite as good as in the analog cases in Figure 2.

Comparing Figures 2 and 5, we see that Figure 5 reports more sampling uncertainty. That is, the gray areas are wider. Again, the crucial point is that the econometrician who computes standard confidence intervals would detect the increase in sampling uncertainty.

Results for the CKM Specification

The third and fourth rows of column 1 in Figure 5 report results for the two and three - shock CKM specifications. Consistent with results reported in CKM, there is substantial bias in the estimated dynamic response functions. For example, in the Two-shock CKM specification, the contemporaneous response of hours worked to a one-standard-deviation technology shock is 0.3 percent, while the mean estimated response is 0.97 percent. This bias stands in contrast to our other results.

Is this bias big or problematic? In our view, bias cannot be evaluated without taking into account sampling uncertainty. Bias matters only to the extent that the econometrician is led to an incorrect inference. For example, suppose sampling uncertainty is large and the econometrician knows it. Then the econometrician would conclude that the data contain little information and, therefore, would not be misled. In this case, we say that bias is not large. In contrast, suppose sampling uncertainty is large, but the econometrician thinks it is small. Here, we would say bias is large.

We now turn to the sampling uncertainty in the CKM specifications. Figure 5 shows that the econometrician's average confidence interval is large relative to the bias. Interestingly, the percentile confidence intervals (stars) are shifted down slightly relative to the standard-deviation-based confidence intervals (circles). On average, the estimated impulse response function is not in the center of the percentile confidence interval. This phenomenon often occurs in practice.¹⁵ Recall that we estimate a four lag VAR in each of our 1,000 synthetic data sets. For the purposes of the bootstrap, each of these VARs is treated as a true data generating process. The asymmetric percentile confidence intervals show that when data are generated by these VARs, VAR-based estimators of the impulse response function have a downward bias.

Figure 3 reveals that for the two- and three-shock CKM specifications, percentile-based coverage rates are reasonably close to 95 percent. Figure 4 shows that the standard deviation

¹⁵An extreme example, in which the point estimates roughly coincide with one of the boundaries of the percentile-based confidence interval, appears in Blanchard and Quah (1989).

based coverage rates are lower than the percentile-based coverage rates. However even these coverage rates are relatively high in that they exceed 70 percent.

In summary, the results for the MLE specification differ from those of the CKM specifications in two interesting ways. First, sampling uncertainty is much larger with the CKM specification. Second, the estimated responses are somewhat biased with the CKM specification. But the bias is small: It has no substantial effect on inference, at least as judged by coverage rates for the econometrician’s confidence intervals.

3.3. Confidence Intervals in the RBC Examples and a Situation in Which VAR-Based Procedures Go Awry

Here we show that the more important technology shocks are in the dynamics of hours worked, the easier it is for VARs to answer the question, ‘how do hours worked respond to a technology shock’. We demonstrate this by considering alternative values of the innovation variance in the labor tax, σ_l , and by considering alternative values of σ , the utility parameter that controls the Frisch elasticity of labor supply.

Consider Figure 6, which focuses on the long-run identification schemes. The first and second columns report results for the two-shock MLE and CKM specifications, respectively. For each specification we redo our experiments, reducing σ_l by a half and then by a quarter. Table 1 shows that the importance of technology shocks rises as the standard deviation of the labor tax shock falls. Figure 6 indicates that the magnitude of sampling uncertainty and the size of confidence intervals fall as the relative importance of labor tax shocks falls.¹⁶

Figure 7 presents the results of a different set of experiments based on perturbations of the two-shock CKM specification. The 1,1 and 2,1 panels show what happens when we vary the value of σ , the parameter that controls the Frisch labor supply elasticity. In the 1,1 panel we set $\sigma = 6$, which corresponds to a Frisch elasticity of 0.63. In the 2,1 panel, we set $\sigma = 0$, which corresponds to a Frisch elasticity of infinity. As the Frisch elasticity is increased, the fraction of the variance in hours worked due to technology shocks decreases (see Table 1). The magnitude of bias and the size of confidence intervals are larger for the higher Frisch elasticity case. In both cases the bias is still smaller than the sampling uncertainty.

We were determined to construct at least one example in which the VAR-based estimator of impulse response functions have bad properties, i.e., bias is larger than sampling uncertainty. We display such an example in the 3,1 panel of Figure 7. The data generating process is a version of the two-shock CKM model with an infinite Frisch elasticity and double the standard deviation of the labor tax rate. Table 1 indicates that with this specification, technology shocks account for a trivial fraction of the variance in hours worked. Of the three measures of V_h , two are 0.46 percent and the third is 0.66 percent. The 3,1 panel of Figure 7 shows that the VAR-based procedure now has very bad properties: the true value of the impulse response function lies outside the average value of both confidence intervals that we consider. This example shows that constructing scenarios in which VAR-based procedures go awry is certainly possible. However, this example seems unlikely to be of practical significance given the poor fit to the data of this version of the model.

¹⁶As σ_l falls, the total volatility of hours worked falls, as does the relative importance of labor tax shocks. In principle, both effects contribute to the decline in sampling uncertainty.

3.4. Are Long-Run Identification Schemes Informative?

Up to now, we have focused on the RBC model as the data generating process. For empirically reasonable specifications of the RBC model, confidence intervals associated with long-run identification schemes are large. One might be tempted to conclude that VAR-based long-run identification schemes are uninformative. Specifically, are the confidence intervals so large that we can never discriminate between competing economic models? Erceg, Guerrieri, and Gust (2005) show that the answer to this question is ‘no’. They consider an RBC model similar to the one discussed above and a version of the sticky wage-price model developed by Christiano, Eichenbaum, and Evans (2005) in which hours worked fall after a positive technology shock. They then conduct a series of experiments to assess the ability of a long-run identified structural VAR to discriminate between the two models on the basis of the response of hours worked to a technology shock.

Using estimated versions of each of the economic models as a data generating process, they generate 10,000 synthetic data sets each with 180 observations. They then estimate a four-variable structural VAR on each synthetic data set and compute the dynamic response of hours worked to a technology shock using long-run identification. Erceg, Guerrieri, and Gust (2005) report that the probability of finding an initial decline in hours that persists for two quarters is much higher in the model with nominal rigidities than in the RBC model (93 percent versus 26 percent). So, if these are the only two models contemplated by the researcher, an empirical finding that hours worked decline after a positive innovation to technology will constitute compelling evidence in favor of the sticky wage-price model.

Erceg, Guerrieri, and Gust (2005) also report that the probability of finding an initial rise in hours that persists for two quarters is much higher in the RBC model than in the sticky wage-price model (71 percent versus 1 percent). So, an empirical finding that hours worked rises after a positive innovation to technology would constitute compelling evidence in favor of the RBC model versus the sticky wage-price alternative.

4. Contrasting Short- and Long- Run Restrictions

The previous section demonstrates that, in the examples we considered, when VARs are identified using short-run restrictions, the conventional estimator of impulse response functions is remarkably accurate. In contrast, for some parameterizations of the data generating process, the conventional estimator of impulse response functions based on long-run identifying restrictions can exhibit noticeable bias. In this section we argue that the key difference between the two identification strategies is that the long-run strategy requires an estimate of the sum of the VAR coefficients, $B(1)$. This object is notoriously difficult to estimate accurately (see Sims, 1972).

We consider a simple analytic expression related to one in Sims (1972). Our expression shows what an econometrician who fits a misspecified, fixed-lag, finite-order VAR would find in population. Let $\hat{B}_1, \dots, \hat{B}_q$ and \hat{V} denote the parameters of the q th-order VAR fit by the econometrician. Then:

$$\hat{V} = V + \min_{\hat{B}_1, \dots, \hat{B}_q} \frac{1}{2\pi} \int_{-\pi}^{\pi} \left[B(e^{-i\omega}) - \hat{B}(e^{-i\omega}) \right] S_Y(\omega) \left[B(e^{i\omega}) - \hat{B}(e^{i\omega}) \right]' d\omega, \quad (4.1)$$

where

$$\begin{aligned} B(L) &= B_1 + B_2L + B_3L^2 + \dots, \\ \hat{B}(L) &= \hat{B}_1 + \hat{B}_2L + \dots + \hat{B}_4L^3. \end{aligned}$$

Here, $B(e^{-i\omega})$ and $\hat{B}(e^{-i\omega})$ correspond to $B(L)$ and $\hat{B}(L)$ with L replaced by $e^{-i\omega}$.¹⁷ In (4.1), B and V are the parameters of the actual infinite-ordered VAR representation of the data (see (2.10)), and $S_Y(\omega)$ is the associated spectral density at frequency ω .¹⁸ According to (4.1), estimation of a VAR approximately involves choosing VAR lag matrices to minimize a quadratic form in the difference between the estimated and true lag matrices. The quadratic form assigns greatest weight to the frequencies for which the spectral density is the greatest. If the econometrician's VAR is correctly specified, then $\hat{B}(e^{-i\omega}) = B(e^{-i\omega})$ for all ω , and $\hat{V} = V$, so that the estimator is consistent. If there is specification error, then $\hat{B}(e^{-i\omega}) \neq B(e^{-i\omega})$ for some ω and $V > \hat{V}$.¹⁹ In our context, specification error exists because the true VAR implied by our data generating processes has $q = \infty$, but the econometrician uses a finite value of q .

To understand the implications of (4.1) for our analysis, it is useful to write in lag-operator form the estimated dynamic response of Y_t to a shock in the first element of ε_t

$$Y_t = [I + \theta_1L + \theta_2L^2 + \dots] \hat{C}_1\varepsilon_{1,t}, \quad (4.2)$$

where the θ_k 's are related to the estimated VAR coefficients as follows:

$$\theta_k = \frac{1}{2\pi} \int_{-\pi}^{\pi} [I - \hat{B}(e^{-i\omega})e^{-i\omega}]^{-1} e^{k\omega i} d\omega. \quad (4.3)$$

In the case of long-run identification, the vector \hat{C}_1 is computed using (2.22), and $\hat{B}(1)$ and \hat{V} replace $B(1)$ and V respectively. In the case of short-run identification, we compute \hat{C}_1 as the second column in the upper triangular Cholesky decomposition of \hat{V} .²⁰

¹⁷The minimization in (4.1) is actually over the trace of the indicated integral. One interpretation of (4.1) is that it provides the probability limit of our estimators – what they would converge to as the sample size increases to infinity. We do not adopt this interpretation, because in practice an econometrician would use a consistent lag-length selection method. The probability limit of our estimators corresponds to the true impulse response functions for all cases considered in this paper.

¹⁸The derivation of this formula is straightforward. Write (2.10) in lag operator form as follows:

$$Y_t = B(L)Y_{t-1} + u_t,$$

where $Eu_tu_t' = V$. Let the fitted disturbances associated with a particular parameterization, $\hat{B}(L)$, be denoted \hat{u}_t . Simple substitution implies:

$$\hat{u}_t = [B(L) - \hat{B}(L)]Y_{t-1} + u_t.$$

The two random variables on the right of the equality are orthogonal, so that the variance of \hat{u}_t is just the variance of the sum of the two:

$$\text{var}(\hat{u}_t) = \text{var}([B(L) - \hat{B}(L)]Y_{t-1}) + V.$$

Expression (4.1) in the text follows immediately.

¹⁹By $V > \hat{V}$, we mean that $V - \hat{V}$ is a positive definite matrix.

²⁰In the earlier discussion it was convenient to adopt the normalization that the technology shock is the second element of ε_t . Here, we adopt the same normalization as for the long-run identification – namely, that the technology shock is the first element of ε_t .

We use (4.1) to understand why estimation based on short-run and long-run identification can produce different results. According to (4.2), impulse response functions can be decomposed into two parts, the impact effect of the shocks, summarized by \hat{C}_1 , and the dynamic part summarized in the term in square brackets. We argue that when a bias arises with long-run restrictions, it is because of difficulties in estimating C_1 . These difficulties do not arise with short-run restrictions.

In the short-run identification case, \hat{C}_1 is a function of \hat{V} only. Across a variety of numerical examples, we find that \hat{V} is very close to V .²¹ This result is not surprising because (4.1) indicates that the entire objective of estimation is to minimize the distance between \hat{V} and V . In the long-run identification case, \hat{C}_1 depends not only on \hat{V} but also on $\hat{B}(1)$. A problem is that the criterion does not assign much weight to setting $\hat{B}(1) = B(1)$ unless $S_Y(\omega)$ happens to be relatively large in a neighborhood of $\omega = 0$. But, a large value of $S_Y(0)$ is not something one can rely on.²² When $S_Y(0)$ is relatively small, attempts to match $\hat{B}(e^{-i\omega})$ with $B(e^{-i\omega})$ at other frequencies can induce large errors in $\hat{B}(1)$.

The previous argument about the difficulty of estimating C_1 in the long-run identification case does not apply to the θ'_k s. According to (4.3) θ_k is a function of $\hat{B}(e^{-i\omega})$ over the whole range of ω 's, not just one specific frequency.

We now present a numerical example, which illustrates Proposition 1 as well as some of the observations we have made in discussing (4.1). Our numerical example focuses on population results. Therefore, it provides only an indication of what happens in small samples.

To understand what happens in small samples, we consider four additional numerical examples. First, we show that when the econometrician uses the true value of $B(1)$, the bias and much of the sampling uncertainty associated with the Two-shock CKM specification disappears. Second, we demonstrate that bias problems essentially disappear when we use an alternative to the standard zero-frequency spectral density estimator used in the VAR literature. Third, we show that the problems are attenuated when the preference shock is more persistent. Fourth, we consider the recursive version of the two-shock CKM specification in which the effect of technology shocks can be estimated using either short- or long-run restrictions.

A Numerical Example

Table 2 reports various properties of the two-shock CKM specification. The first six B_j 's in the infinite-order VAR, computed using (2.12), are reported in Panel A. These B_j 's eventually converge to zero, however they do so slowly. The speed of convergence is governed by the size of the maximal eigenvalue of the matrix M in (2.8), which is 0.957. Panel B displays the \hat{B}_j 's that solve (4.1) with $q = 4$. Informally, the \hat{B}_j 's look similar to the B_j 's for $j = 1, 2, 3, 4$. In line with this observation, the sum of the true B_j 's, $B_1 + \dots + B_4$ is similar in magnitude to the sum of the estimated \hat{B}_j 's, $\hat{B}(1)$ (see Panel C). But the econometrician using long-run restrictions needs a good estimate of $B(1)$. This matrix is very different from $B_1 + \dots + B_4$. Although the

²¹This result explains why lag-length selection methods, such as the Akaike criterion, almost never suggest values of q greater than 4 in artificial data sets of length 180, regardless of which of our data generating methods we used. These lag length selection methods focus on \hat{V} .

²²Equation (4.1) shows that $\hat{B}(1)$ corresponds to only a single point in the integral. So other things equal, the estimation criterion assigns *no* weight at all to getting $\hat{B}(1)$ right. The reason $B(1)$ is identified in our setting is that the $B(\omega)$ functions we consider are continuous at $\omega = 0$.

remaining B_j 's for $j > 4$ are individually small, their sum is not. For example, the 1,1 element of $B(1)$ is 0.28, or six times larger than the 1,1 element of $B_1 + \dots + B_4$.

The distortion in $\hat{B}(1)$ manifests itself in a distortion in the estimated zero-frequency spectral density (see Panel D). As a result, there is distortion in the estimated impact vector, \hat{C}_1 (Panel F).²³ To illustrate the significance of the latter distortion for estimated impulse response functions, we display in Figure 8 the part of (4.2) that corresponds to the response of hours worked to a technology shock. In addition, we display the true response. There is a substantial distortion, which is approximately the same magnitude as the one reported for small samples in Figure 5. The third line in Figure 8 corresponds to (4.2) when \hat{C}_1 is replaced by its true value, C_1 . Most of the distortion in the estimated impulse response function is eliminated by this replacement. Finally, the distortion in \hat{C}_1 is due to distortion in $\hat{B}(1)$, as \hat{V} is virtually identical to V (panel E).

This example is consistent with our overall conclusion that the individual B_j 's and V are well estimated by the econometrician using a four-lag VAR. The distortions that arise in practice primarily reflect difficulties in estimating $B(1)$. Our short-run identification results in Figure 2 are consistent with this claim, because distortions are minimal with short-run identification.

Using the True Value of $B(1)$ in a Small Sample

A natural way to isolate the role of distortions in $\hat{B}(1)$ is to replace $\hat{B}(1)$ by its true value when estimating the effects of a technology shock. We perform this replacement for the two-shock CKM specification, and report the results in Figure 9. For convenience, the 1,1 panel of Figure 9 repeats our results for the two-shock CKM specification from the 3,1 panel in Figure 5. The 1,2 panel of Figure 9 shows the sampling properties of our estimator when the true value of $B(1)$ is used in repeated samples. When we use the true value of $B(1)$ the bias completely disappears. In addition, coverage rates are much closer to 95 percent and the boundaries of the average confidence intervals are very close to the boundaries of the gray area.

Using an Alternative Zero-Frequency Spectral Density Estimator

In practice, the econometrician does not know $B(1)$. However, we can replace the VAR-based zero-frequency spectral density in (2.19) with an alternative estimator of $S_Y(0)$. Here, we consider the effects of using a standard Bartlett estimator:²⁴

$$S_Y(0) = \sum_{k=-(T-1)}^{T-1} g(k)\hat{C}(k), \quad g(k) = \begin{cases} 1 - \frac{|k|}{r} & |k| \leq r \\ 0 & |k| > r \end{cases}, \quad (4.4)$$

where, after removing the sample mean from Y_t :

$$\hat{C}(k) = \frac{1}{T} \sum_{t=k+1}^T Y_t Y'_{t-k}.$$

²³A similar argument is presented in Ravenna (2005).

²⁴Christiano, Eichenbaum and Vigfusson (2006) also consider the estimator proposed by Andrews and Monahan (1992).

We use essentially all possible covariances in the data by choosing a large value of r , $r = 150$.²⁵ In some respects, our modified estimator is equivalent to running a VAR with longer lags.

We now assess the effect of our modified long-run estimator. The first two rows in Figure 5 present results for cases in which the data generating mechanism corresponds to our two- and three-shock MLE specifications. Both the standard estimator (the left column) and our modified estimator (the right column) exhibit little bias. In the case of the standard estimator, the econometrician's estimator of standard errors understates somewhat the degree of sampling uncertainty associated with the impulse response functions. The modified estimator reduces this discrepancy. Specifically, the circles and stars in the right column of Figure 5 coincide closely with the boundary of the gray area. Coverage rates are reported in the 2,1 panels of Figures 3 and 4. In Figure 3, coverage rates now exceed 95 percent. The coverage rates in Figure 4 are much improved relative to the standard case. Indeed, these rates are now close to 95 percent. Significantly, the degree of sampling uncertainty associated with the modified estimator is not greater than that associated with the standard estimator. In fact, in some cases, sampling uncertainty declines slightly.

The last two rows of column 1 in Figure 5 display the results when the data generating process is a version of the CKM specification. As shown in the second column, the bias is essentially eliminated by using the modified estimator. Once again the circles and stars roughly coincide with the boundary of the gray area. Coverage rates for the percentile-based confidence intervals reported in Figure 3 again have a tendency to exceed 95 percent (2,2 panel). As shown in the 2,2 panel of Figure 4, coverage rates associated with the standard deviation based estimator are very close to 95 percent. There is a substantial improvement over the coverage rates associated with the standard spectral density estimator.

Figure 5 indicates that when the standard estimator works well, the modified estimator also works well. When the standard estimator results in biases, the modified estimator removes them. These findings are consistent with the notion that the biases for the two CKM specifications reflect difficulties in estimating the spectral density at frequency zero. Given our finding that \hat{V} is an accurate estimator of V , we conclude that the difficulties in estimating the zero-frequency spectral density in fact reflect problems with $B(1)$.

The second column of Figure 7 shows how our modified VAR-based estimator works when the data are generated by the various perturbations on the Two-shock CKM specification. In every case, bias is substantially reduced.

Shifting Power to the Low Frequencies

Formula (4.1), suggests that, other things being equal, the more power there is near frequency zero, the less bias there is in $\hat{B}(1)$ and the better behaved is the estimated impulse response function to a technology shock. To pursue this observation we change the parameterization of the non-technology shock in the two-shock CKM specification. We reallocate power toward frequency zero, holding the variance of the shock constant by increasing ρ_l to 0.998 and suitably lowering σ_l in (2.1). The results are reported in the 2,1 panel of Figure 9. The bias associated with the two-shock CKM specification almost completely disappears. This result is

²⁵The rule of always setting the bandwidth, r , equal to sample size does not yield a consistent estimator of the spectral density at frequency zero. We assume that as sample size is increased beyond $T = 180$, the bandwidth is increased sufficiently slowly to achieve consistency.

consistent with the notion that the bias problems with the two-shock CKM specification stem from difficulties in estimating $B(1)$.

The previous result calls into question conjectures in the literature (see Erceg, Guerrieri, and Gust, 2005). According to these conjectures, if there is more persistence in a non-technology shock, then the VAR will produce biased results because it will confuse the technology and non-technology shocks. Our result shows that this intuition is incomplete, because it fails to take into account all of the factors mentioned in our discussion of (4.1). To show the effect of persistence, we consider a range of values of ρ_l to show that the impact of ρ_l on bias is in fact not monotone.

The 2,2 panel of Figure 9 displays the econometrician's estimator of the contemporaneous impact on hours worked of a technology shock against ρ_l . The dashed line indicates the true contemporaneous effect of a technology shock on hours worked in the two-shock CKM specification. The dot-dashed line in the figure corresponds to the solution of (4.1), with $q = 4$, using the standard VAR-based estimator.²⁶ The star in the figure indicates the value of ρ_l in the two-shock CKM specification. In the neighborhood of this value of ρ_l , the distortion in the estimator falls sharply as ρ_l increases. Indeed, for $\rho_l = 0.9999$, essentially no distortion occurs. For values of ρ_l in the region, $(-0.5, 0.5)$, the distortion increases with increases in ρ_l .

The 2,2 panel of Figure 9 also allows us to assess the value of our proposed modification to the standard estimator. The line with diamonds displays the modified estimator of the contemporaneous impact on hours worked of a technology shock. When the standard estimator works well, that is, for large values of ρ_l the modified and standard estimators produce similar results. However, when the standard estimator works poorly, e.g. for values of ρ_l near 0.5, our modified estimator cuts the bias in half.

A potential shortcoming of the previous experiments is that persistent changes in $\tau_{l,t}$ do not necessarily induce very persistent changes in labor productivity. To assess the robustness of our results, we also considered what happens when there are persistent changes in $\tau_{x,t}$. These do have a persistent impact on labor productivity. In the two-shock CKM model, we set $\tau_{l,t}$ to a constant and allowed $\tau_{x,t}$ to be stochastic. We considered values of ρ_x in the range, $[-0.5, 1]$, holding the variance of $\tau_{x,t}$ constant. We obtain results similar to those reported in the 2,2 panel of Figure 9.

Short- and Long-Run Restrictions in a Recursive Model

We conclude this section by considering the recursive version of the two-shock CKM specification. This specification rationalizes estimating the impact on hours worked of a shock to technology using either the short- or the long-run identification strategy. We generate 1,000 data sets, each of length 180. On each synthetic data set, we estimate a four lag, bivariate VAR. Given this estimated VAR, we can estimate the effect of a technology shock using the short- and long-run identification strategy. Figure 10 reports our results. For the long-run identification strategy, there is substantial bias. In sharp contrast, there is no bias for the short-run identification strategy. Because both procedures use the same estimated VAR parameters, the bias in the long-run identification strategy is entirely attributable due to the use of $\hat{B}(1)$.

²⁶Because (4.1) is a quadratic function, we solve the optimization problem by solving the linear first-order conditions. These are the Yule-Walker equations, which rely on population second moments of the data. We obtain the population second moments by complex integration of the reduced form of the model used to generate the data, as suggested by Christiano (2002).

5. Relation to Chari-Kehoe-McGrattan

In the preceding sections we argue that structural VAR-based procedures have good statistical properties. Our conclusions about the usefulness of structural VARs stand in sharp contrast to the conclusions of CKM. These authors argue that, for plausibly parameterized RBC models, structural VARs lead to misleading results. They conclude that structural VARs are not useful for constructing and evaluating structural economic models. In this section we present the reasons we disagree with CKM.

CKM's Exotic Data Generating Processes

CKM's critique of VARs is based on simulations using particular DSGE models estimated by maximum likelihood methods. Here, we argue that their key results are driven by assumptions about measurement error. CKM's measurement error assumptions are overwhelmingly rejected in favor of alternatives under which their key results are overturned.

CKM adopt a state-observer setup to estimate their model. Define:

$$Y_t = (\Delta \log a_t, \log l_t, \Delta \log i_t, \Delta \log G_t)',$$

where G_t denotes government spending plus net exports. CKM suppose that

$$Y_t = X_t + v_t, \quad E v_t v_t' = R, \quad (5.1)$$

where R is diagonal, v_t is a 4×1 vector of i.i.d. measurement errors and X_t is a 4×1 vector containing the model's implications for the variables in Y_t . The two-shock CKM specification has only the shocks, $\tau_{l,t}$ and z_t . CKM model government spending plus net exports as:

$$G_t = g_t \times Z_t,$$

where g_t is in principle an exogenous stochastic process. However, when CKM estimate the parameters of the technology and preferences processes, $\tau_{l,t}$ and z_t , they set the variance of the government spending shock to zero, so that g_t is a constant. As a result, CKM assume that

$$\Delta \log G_t = \log z_t + \text{measurement error}.$$

CKM fix the elements on the diagonal of R exogenously to a "small number", leading to the remarkable implication that government purchases plus net exports.

To demonstrate the sensitivity of CKM's results to their specification of the magnitude of R , we consider the different assumptions that CKM make in different drafts of their paper. In the draft of May 2005, CKM set the diagonal elements of R to 0.0001. In the draft of July 2005, CKM set the i^{th} diagonal element of R equal to 0.01 times the variance of the i^{th} element of Y_t .

The 1,1 and 2,1 panels in Figure 11 report results corresponding to CKM's two-shock specifications in the July and May drafts, respectively.²⁷ These panels display the log likelihood

²⁷To ensure comparability of results we use CKM's computer code and data, available on Ellen McGrattan's webpage. The algorithm used by CKM to form the estimation criterion is essentially the same as the one we used to estimate our models. The only difference is that CKM use an approximation to the Gaussian function by working with the steady state Kalman gain. We form the exact Gaussian density function, in which the Kalman gain varies over dates, as described in Hamilton (1994). We believe this difference is inconsequential.

value (see *LLF*) of these two models and their implications for VAR-based impulse response functions (the 1,1 panel is the same as the 3,1 panel in Figure 5). Surprisingly, the log-likelihood of the July specification is orders of magnitude worse than that of the May specification.

The 3,1 panel in Figure 11 displays our results when the diagonal elements of R are included among the parameters being estimated.²⁸ We refer to the resulting specification as the “CKM free measurement error specification”. First, both the May and the July specifications are rejected relative to the free measurement error specification. The likelihood ratio statistic for testing the May and July specifications are 428 and 6,266, respectively. Under the null hypothesis that the May or July specification is true, these statistics are realizations of a chi-square distribution with 4 degrees of freedom. The evidence against CKM’s May or July specifications of measurement error is overwhelming.

Second, when the data generating process is the CKM free measurement error specification, the VAR-based impulse response function is virtually unbiased (see the 3,1 panel in Figure 11). We conclude that the bias in the two-shock CKM specification is a direct consequence of CKM’s choice of the measurement error variance.

As noted above, CKM’s measurement error assumption has the implication that $\Delta \log G_t$ is roughly equals to $\log z_t$. To investigate the role played by this peculiar implication, we delete $\Delta \log G_t$ from Y_t and reestimate the system. We present the results in the right column of Figure 11. In each panel of that column, we re-estimate the system in the same way as the corresponding panel in the left column, except that $\Delta \log G_t$ is excluded from Y_t . Comparing the 2,1 and 2,2 panels, we see that, with the May measurement error specification, the bias disappears after relaxing CKM’s $\Delta \log G_t = \log z_t$ assumption. Under the July specification of measurement error, the bias result remains even after relaxing CKM’s assumption (compare the 1,1 and 1,2 graphs of Figure 11). As noted above, the May specification of CKM’s model has a likelihood that is orders of magnitude higher than the July specification. So, in the version of the CKM model selected by the likelihood criterion (i.e., the May version), the $\Delta \log G_t = \log z_t$ assumption plays a central role in driving the CKM’s bias result.

In sum, CKM’s examples which imply that VARs with long-run identification display substantial bias, are not empirically interesting from a likelihood point of view. The bias in their examples is due to the way CKM choose the measurement error variance. When their measurement error specification is tested, it is overwhelmingly rejected in favor of an alternative in which the CKM bias result disappears.

Stochastic Process Uncertainty

CKM argue that there is considerable uncertainty in the business cycle literature about the values of parameters governing stochastic processes such as preferences and technology. They argue that this uncertainty translates into a wide class of examples in which the bias in structural VARs leads to severely misleading inference. The right panel in Figure 12 summarizes their argument. The horizontal axis covers the range of values of $(\sigma_l/\sigma_z)^2$ considered by CKM. For each value of $(\sigma_l/\sigma_z)^2$ we estimate, by maximum likelihood, four parameters of the two-shock

²⁸When generating the artificial data underlying the calculations in the 3,1 panel of Figure 11, we set the measurement error to zero. (The same assumption was made for all the results reported here.) However, simulations that include the estimated measurement error produce results that are essentially the same.

model: μ_z , τ_l , σ_l and ρ_l .²⁹ We use the estimated model as a data generating process. The left vertical axis displays the small sample mean of the corresponding VAR-based estimator of the contemporaneous response of hours worked to a one-standard deviation technology shock.

Based on a review the RBC literature, CKM report that they have a roughly uniform prior over the different values of $(\sigma_l/\sigma_z)^2$ considered in Figure 12. The figure indicates that for many of these values, the bias is large (compare the small sample mean, the solid line, with the true response, the starred line). For example, there is a noticeable bias in the 2-shock CKM specification, where $(\sigma_l/\sigma_z)^2 = 1.1$.

We emphasize three points. First, as we stress repeatedly, bias cannot be viewed in isolation from sampling uncertainty. The two dashed lines in the figure indicate the 95 percent probability interval. These intervals are enormous relative to the bias. Second, not all values of $(\sigma_l/\sigma_z)^2$ are equally likely, and for the ones with greatest likelihood there is little bias. On the horizontal axis of the left panel of Figure 12, we display the same range of values of $(\sigma_l/\sigma_z)^2$ as in the right panel. On the vertical axis we report the log-likelihood value of the associated model. The peak of this likelihood occurs close to the estimated value in the two-shock MLE specification. Note how the log-likelihood value drops sharply as we consider values of $(\sigma_l/\sigma_z)^2$ away from the unconstrained maximum likelihood estimate. The vertical bars in the figure indicate the 95 percent confidence interval for $(\sigma_l/\sigma_z)^2$.³⁰ Figure 12 reveals that the confidence interval is very narrow relative to the range of values considered by CKM, and that within the interval, the bias is quite small.

Third, the right axis in the right panel of Figure 12 plots V_h , the percent of the variance in log hours due to technology, as a function of $(\sigma_l/\sigma_z)^2$. The values of $(\sigma_l/\sigma_z)^2$ for which there is a noticeable bias correspond to model economies where V_h is less than 2 percent. Here, identifying the effects of a technology shock on hours worked is tantamount to looking for a needle in a haystack.

The Metric for Assessing the Performance of Structural VARs

CKM emphasize comparisons between the true dynamic response function in the data generating process and the response function that an econometrician would estimate using a four-lag VAR with an infinite amount of data. In our own analysis in section 4, we find population calculations with four lag VARs useful for some purposes. However, we do not view the probability limit of a four lag VAR as an interesting metric for measuring the usefulness of structural VARs. In practice econometricians do not have an infinite amount of data. Even if they did, they would certainly not use a fixed lag length. Econometricians determine lag length endogenously and, in a large sample, lag length would grow. If lag lengths grow at the appropriate rate with sample size, VAR-based estimators of impulse response functions are consistent. The interesting issue (to us) is how VAR-based procedures perform in samples of the size that practitioners have at their disposal. This is why we focus on small sample properties like bias and sampling uncertainty.

Over-Differencing

²⁹We use CKM's computer code and data to ensure comparability of results.

³⁰The bounds of this interval are the upper and lower values of $(\sigma_l/\sigma_z)^2$ where twice the difference of the log-likelihood from its maximal value equals the critical value associated with the relevant likelihood ratio test.

The potential power of the CKM argument lies in showing that VAR-based procedures are misleading, even under circumstances when everyone would agree that VARs should work well, namely when the econometrician commits no avoidable specification error. The econometrician does, however, commit one unavoidable specification error. The true VAR is infinite ordered, but the econometrician assumes the VAR has a finite number of lags. CKM argue that this seemingly innocuous assumption is fatal for VAR analysis. We have argued that this conclusion is unwarranted.

CKM present other examples in which the econometrician commits an avoidable specification error. Specifically, they study the consequences of over differencing hours worked. That is, the econometrician first differences hours worked when hours worked are stationary.³¹ This error gives rise to bias in VAR-based impulse response functions that is large relative to sampling uncertainty. CKM argue that this bias is another reason not to use VARs.

However, the observation that avoidable specification error is possible in VAR analysis is not a problem for VARs per se. The possibility of specification error is a potential pitfall for any type of empirical work. In any case, CKM's analysis of the consequences of over differencing is not new. For example, Christiano, Eichenbaum and Vigfusson (2003, hereafter, CEV) study a situation in which the true data generating process satisfies two properties: Hours worked are stationary and they rise after a positive technology shock. CEV then consider an econometrician who does VAR-based long-run identification when Y_t in (2.16) contains the growth rate of hours rather than the log level of hours. CEV show that the econometrician would falsely conclude that hours worked fall after a positive technology shock. CEV do not conclude from this exercise that structural VARs are not useful. Rather, they develop a statistical procedure to help decide whether hours worked should be first differenced or not.

CKM Ignore Short-Run Identification Schemes

We argue that VAR-based short-run identification schemes lead to remarkably accurate and precise inference. This result is of interest because the preponderance of the empirical literature on structural VARs explores the implications of short-run identification schemes. CKM are silent on this literature. McGrattan (2006) dismisses short-run identification schemes as “hokey.” One possible interpretation of this adjective is that McGrattan can easily imagine models in which the identification scheme is incorrect. The problem with this interpretation is that all models are a collection of strong identifying assumptions, all of which can be characterized as “hokey”. A second interpretation is that in McGrattan (2006)'s view, the type of zero restrictions typically used in short run identification are not compatible with dynamic equilibrium theory. This view is simply incorrect (see Sims and Zha (2006)). A third possible interpretation is that no one finds short-run identifying assumptions interesting. However, the results of short-run identification schemes have had an enormous effect on the construction of dynamic, general equilibrium models. See Woodford (2003) for a summary in the context of monetary models.

Sensitivity of Some VAR Results to Data Choices

³¹For technical reasons, CKM actually consider ‘quasi differencing’ hours worked using a differencing parameter close to unity. In small samples this type of quasi differencing is virtually indistinguishable from first differencing.

CKM argue that VARs are very sensitive to the choice of data. Specifically, they review the papers by Francis and Ramey (2004), CEV, and Gali and Rabanal (2004), which use long-run VAR methods to estimate the response of hours worked to a positive technology shock. CKM note that these studies use different measures of per capita hours worked and output in the VAR analysis. The bottom panel of Figure 13 displays the different measures of per capita hours worked that these studies use. Note how the low frequency properties of these series differ. The corresponding estimated impulse response functions and confidence intervals are reported in the top panel. CKM view it as a defect in VAR methodology that the different measures of hours worked lead to different estimated impulse response functions. We disagree. Empirical results *should* be sensitive to substantial changes in the data. A constructive response to the sensitivity in Figure 13 is to carefully analyze the different measures of hours worked, see which is more appropriate, and perhaps construct a better measure. It is not constructive to dismiss an econometric technique that signals the need for better measurement.

CKM note that the principle differences in the hours data occur in the early part of the sample. According to CKM, when they drop these early observations they obtain different impulse response functions. However, as Figure 13 shows, these impulse response functions are not significantly different from each other.

6. A Model with Nominal Rigidities

In this section we use the model in ACEL to assess the accuracy of structural VARs for estimating the dynamic response of hours worked to shocks. This model allows for nominal rigidities in prices and wages and has three shocks: a monetary policy shock, a neutral technology shock, and a capital-embodied technology shock. Both technology shocks affect labor productivity in the long run. However, the only shock in the model that affects the price of investment in the long run is the capital-embodied technology shock. We use the ACEL model to evaluate the ability of a VAR to uncover the response of hours worked to both types of technology shock and to the monetary policy shock. Our strategy for identifying the two technology shocks is similar to the one proposed by Fisher (2006). The model rationalizes a version of the short-run, recursive identification strategy used by Christiano, Eichenbaum and Evans (1999) to identify monetary shocks. This strategy corresponds closely to the recursive procedure studied in section 2.3.2.

6.1. The Model

The details of the ACEL model, as well as the parameter estimates, are reported in Appendix A. Here, we limit our discussion to what is necessary to clarify the nature of the shocks in the ACEL model. Final goods, Y_t , are produced using a standard Dixit-Stiglitz aggregator of intermediate goods, $y_t(i)$, $i \in (0, 1)$. To produce a unit of consumption goods, C_t , one unit of final goods is required. To produce one unit of investment goods, I_t , Υ_t^{-1} units of final goods are required. In equilibrium, Υ_t^{-1} is the price, in units of consumption goods, of an investment good. Let $\mu_{\Upsilon,t}$ denote the growth rate of Υ_t , let μ_{Υ} denote the nonstochastic steady state value of $\mu_{\Upsilon,t}$, and let $\hat{\mu}_{\Upsilon,t}$ denote the percent deviation of $\mu_{\Upsilon,t}$ from its steady state value:

$$\mu_{\Upsilon,t} = \frac{\Upsilon_t}{\Upsilon_{t-1}}, \quad \hat{\mu}_{\Upsilon,t} = \frac{\mu_{\Upsilon,t} - \mu_{\Upsilon}}{\mu_{\Upsilon}}. \quad (6.1)$$

The stochastic process for the growth rate of Υ_t is:

$$\hat{\mu}_{\Upsilon,t} = \rho_{\mu_{\Upsilon}} \hat{\mu}_{\Upsilon,t-1} + \sigma_{\mu_{\Upsilon}} \varepsilon_{\mu_{\Upsilon},t}, \quad \sigma_{\mu_{\Upsilon}} > 0. \quad (6.2)$$

We refer to the i.i.d. unit variance random variable, $\varepsilon_{\mu_{\Upsilon},t}$, as the capital-embodied technology shock. ACEL assume that the intermediate good, $y_t(i)$, for $i \in (0, 1)$ is produced using a Cobb-Douglas production function of capital and hours worked. This production function is perturbed by a multiplicative, aggregate technology shock denoted by Z_t . Let z_t denote the growth rate of Z_t , let z denote the nonstochastic steady state value of z_t , and let \hat{z}_t denote the percentage deviation of z_t from its steady state value:

$$z_t = \frac{Z_t}{Z_{t-1}}, \quad \hat{z}_t = \frac{z_t - z}{z}. \quad (6.3)$$

The stochastic process for the growth rate of Z_t is:

$$\hat{z}_t = \rho_z \hat{z}_{t-1} + \sigma_z \varepsilon_t^z, \quad \sigma_z > 0, \quad (6.4)$$

where the i.i.d. unit variance random variable, ε_t^z , is the neutral shock to technology.

We now turn to the monetary policy shock. Let x_t denote M_t/M_{t-1} , where M_t denotes the monetary base. Let \hat{x}_t denote the percentage deviation of x_t from its steady state, i.e., $(\hat{x}_t - x)/x$. We suppose that \hat{x}_t is the sum of three components. One, \hat{x}_{Mt} , represents the component of \hat{x}_t reflecting an exogenous shock to monetary policy. The other two, \hat{x}_{zt} and $\hat{x}_{\Upsilon t}$, represent the endogenous response of \hat{x}_t to the neutral and capital-embodied technology shocks, respectively. Thus monetary policy is given by:

$$\hat{x}_t = \hat{x}_{zt} + \hat{x}_{\Upsilon t} + \hat{x}_{Mt}. \quad (6.5)$$

ACEL assume that

$$\begin{aligned} \hat{x}_{M,t} &= \rho_{xM} \hat{x}_{M,t-1} + \sigma_M \varepsilon_{M,t}, \quad \sigma_M > 0 \\ \hat{x}_{z,t} &= \rho_{xz} \hat{x}_{z,t-1} + c_z \varepsilon_t^z + c_z^p \varepsilon_{t-1}^z \\ \hat{x}_{\Upsilon,t} &= \rho_{x\Upsilon} \hat{x}_{\Upsilon,t-1} + c_{\Upsilon} \varepsilon_{\mu_{\Upsilon},t} + c_{\Upsilon}^p \varepsilon_{\mu_{\Upsilon},t}. \end{aligned} \quad (6.6)$$

Here, $\varepsilon_{M,t}$ represents the shock to monetary policy and is an i.i.d. unit variance random variable.

Table 3 summarizes the importance of different shocks for the variance of hours worked and output. Neutral and capital-embodied technology shocks account for roughly equal percentages of the variance of hours worked (40 percent each), while monetary policy shocks account for the remainder. Working with HP-filtered data reduces the importance of neutral technology shocks to about 18 percent. Monetary policy shocks become much more important for the variance of hours worked. A qualitatively similar picture emerges when we consider output.

It is worth emphasizing that neutral technology shocks are much more important in hours worked in the ACEL model than in the RBC model. This fact plays an important role in determining the precision of VAR-based inference using long-run restrictions in the ACEL model.

6.2. Results

We use the ACEL model to simulate 1,000 data sets each with 180 observations. We report results from two different VARs. In the first VAR, we simultaneously estimate the dynamic effect on hours worked of a neutral technology shock and a capital-embodied technology shock. The variables in this VAR are:

$$Y_t = \begin{pmatrix} \Delta \ln p_{It} \\ \Delta \ln a_t \\ \ln l_t \end{pmatrix},$$

where p_{It} denotes the price of capital in consumption units. The variable, $\ln(p_{It})$, corresponds to $\ln(\Upsilon_t^{-1})$ in the model. As in Fisher (2006), we identify the dynamic effects on Y_t of the two technology shocks, using a generalization of the strategy in section 2.3.1.³² The details are provided in Appendix B.

The 1,1 panel of Figure 14 displays our results using the standard VAR procedure to estimate the dynamic response of hours worked to a neutral technology shock. Several results are worth emphasizing. First, the estimator is essentially unbiased. Second, the econometrician's estimator of sampling uncertainty is also reasonably unbiased. The circles and stars, which indicate the mean value of the econometrician's standard-deviation-based and percentile-based confidence intervals, roughly coincide with the boundaries of the gray area. However, there is a slight tendency, in both cases, to understate the degree of sampling uncertainty. Third, confidence intervals are small, relative to those in the RBC examples. Both sets of confidence intervals exclude zero at all lags shown. This result provides another example, in addition to the one provided by Erceg, Guerrieri, and Gust (2005), in which long-run identifying restrictions are useful for discriminating between models. An econometrician who estimates that hours drop after a positive technology shock would reject our parameterization of the ACEL model. Similarly, an econometrician with a model implying that hours fall after a positive technology shock would most likely reject that model if the actual data were generated by our parameterization of the ACEL model.

The 2,1 panel in Figure 14 shows results for the response to a capital-embodied technology shock as estimated using the standard VAR estimator. The sampling uncertainty is somewhat higher for this estimator than for the neutral technology shock. In addition, there is a slight amount of bias. The econometrician understates somewhat the degree of sampling uncertainty.

We now consider the response of hours worked to a monetary policy shock. We estimate this response using a VAR with the following variables:

$$Y_t = \begin{pmatrix} \Delta \log a_t \\ \log l_t \\ R_t \end{pmatrix}.$$

As discussed in Christiano, Eichenbaum, and Evans (1999), the monetary policy shock is identified by choosing C to be the lower triangular decomposition of the variance covariance matrix, V , of the VAR disturbances. That is, we choose a lower triangular matrix, C with positive diagonal terms, such that $CC' = V$. Let $u_t = C\varepsilon_t$. We then interpret the last element of ε_t as the monetary policy shock. According to the results in the 1,2 panel of Figure 14, the VAR-based

³²Our strategy differs somewhat from the one pursued in Fisher (2006), who applies a version of the instrumental variables strategy proposed by Shapiro and Watson (1988).

estimator of the response of hours worked displays relatively little bias and is highly precise. In addition, the econometrician's estimator of sampling uncertainty is virtually unbiased. Suppose the impulse response in hours worked to a monetary policy shock were computed using VAR-based methods with data generated from this model. We conjecture that a model in which money is neutral, or in which a monetary expansion drives hours worked down, would be easy to reject.

7. Concluding Remarks

In this paper we study the ability of structural VARs to uncover the response of hours worked to a technology shock. We consider two classes of data generating processes. The first class consists of a series of real business cycle models that we estimate using maximum likelihood methods. The second class consists of the monetary model in ACEL. We find that with short-run restrictions, structural VARs perform remarkably well in all our examples. With long-run restrictions, structural VARs work well as long as technology shocks explain at least a very small portion of the variation in hours worked.

In a number of examples that we consider, VAR-based impulse response functions using long-run restrictions exhibit some bias. Even though these examples do not emerge from empirically plausible data generating processes, we find them of interest. They allow us to diagnose what can go wrong with long-run identification schemes. Our diagnosis leads us to propose a modification to the standard VAR-based procedure for estimating impulse response functions using long-run identification. This procedure works well in our examples.

Finally, we find that confidence intervals with long-run identification schemes are substantially larger than those with short-run identification schemes. In all empirically plausible cases, the VARs deliver confidence intervals that accurately reflect the true degree of sampling uncertainty. We view this characteristic as a great virtue of VAR-based methods. When the data contain little information, the VAR will indicate the lack of information. To reduce large confidence intervals the analyst must either impose additional identifying restrictions (i.e., use more theory) or obtain better data.

References

- Altig, David, Lawrence Christiano, Martin Eichenbaum, and Jesper Linde (2005). “Firm-Specific Capital, Nominal Rigidities, and the Business Cycle,” NBER Working Paper Series 11034. Cambridge, Mass.: National Bureau of Economic Research, January.
- Andrews, Donald W. K., and J. Christopher Monahan (1992). “An Improved Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimator,” *Econometrica*, vol. 60 (July), pp. 953–66.
- Basu, Susanto, John Fernald, and Miles Kimball (2004). “Are Technology Improvements Contractionary?” NBER Working Paper Series 10592. Cambridge, Mass.: National Bureau of Economic Research, June.
- Bernanke, Ben S. (1986). “Alternative Explanations of the Money-Income Correlation,” *Carnegie Rochester Conference Series on Public Policy*, vol. 25 (Autumn), pp. 49–99.
- Bernanke, Ben S., and Alan S. Blinder (1992). “The Federal Funds Rate and the Channels of Monetary Transmission,” *American Economic Review*, vol. 82 (September), pp. 901–21.
- Bernanke, Ben S., and Ilian Mihov (1998). “Measuring Monetary Policy,” *Quarterly Journal of Economics*, vol. 113 (August), pp. 869–902.
- Blanchard, Olivier, and Roberto Perotti (2002). “An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output,” *Quarterly Journal of Economics*, vol. 117 (November), pp. 1329–68.
- Blanchard, Olivier, and Danny Quah (1989). “The Dynamic Effects of Aggregate Demand and Supply Disturbances,” *American Economic Review*, vol. 79 (September), pp. 655–73.
- Blanchard, Olivier, and Mark Watson (1986). “Are Business Cycles All Alike?” in Robert Gordon, ed., *Continuity and Change in the American Business Cycle*. Chicago: University of Chicago Press, pp.123–56.
- Chari, V. V., Patrick J. Kehoe, and Ellen McGrattan (2005a). “A Critique of Structural VARs Using Real Business Cycle Theory,” Working Paper Series 631. Minneapolis: Federal Reserve Bank of Minneapolis, May.
- (2005b). “A Critique of Structural VARs Using Real Business Cycle Theory,” Working Paper Series 631. Minneapolis: Federal Reserve Bank of Minneapolis, July.
- Christiano, Lawrence J. (1988). “Why Does Inventory Investment Fluctuate So Much?” *Journal of Monetary Economics*, vol. 21 (March–May), pp. 247–80.
- (2002). “Solving Dynamic Equilibrium Models by a Method of Undetermined Coefficients,” *Computational Economics*, vol. 20 (October), pp. 21–55.
- Christiano, Lawrence J., and Martin Eichenbaum (1992). “Identification and the Liquidity Effect of a Monetary Policy Shock,” in Alex Cukierman, Zvi Hercowitz, and Leonardo Leiderman, eds., *Political Economy, Growth, and Business Cycles*. Cambridge, Mass.: MIT Press, pp. 335–70.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles Evans (1999). “Monetary Policy Shocks: What Have We Learned and to What End?,” in John B. Taylor and Michael D. Woodford, eds., *Handbook of Macroeconomics*. Volume 1A, Amsterdam; Elsevier Science, pp. 65–148
- Christiano, Lawrence J., Martin Eichenbaum, and Charles Evans (2005). “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, vol. 113 (February), pp. 1–45.
- Christiano, Lawrence J., Martin Eichenbaum, and Robert Vigfusson (2003). “What Happens after a Technology Shock?” NBER Working Paper Series 9819. Cambridge, Mass.: National Bureau of Economic Research, July.

- (2004). “The Response of Hours to a Technology Shock: Evidence Based on Direct Measures of Technology,” *Journal of the European Economic Association*, vol. 2 (April), pp. 381–95.
- (forthcoming). “Alternative Procedures for Estimating Long-Run Identified Vector Autoregressions,” *Journal of the European Economic Association*.
- Cooley, Thomas F., and Mark Dwyer (1998). “Business Cycle Analysis without Much Theory: A Look at Structural VARs,” *Journal of Econometrics*, vol. 83 (March–April), pp. 57–88.
- Cushman, David O., and Tao Zha (1997). “Identifying Monetary Policy in a Small Open Economy under Flexible Exchange Rates,” *Journal of Monetary Economics*, vol. 39 (August), pp. 433–48.
- Del Negro, Marco, Frank Schorfheide, Frank Smets, and Raf Wouters (2005). “On the Fit and Forecasting Performance of New Keynesian Models,” unpublished paper, University of Pennsylvania.
- Eichenbaum, Martin, and Charles Evans (1995). “Some Empirical Evidence on the Effects of Shocks to Monetary Policy on Exchange Rates,” *Quarterly Journal of Economics*, vol. 110 (November), pp. 975–1009.
- Erceg, Christopher J., Luca Guerrieri, and Christopher Gust (2005). “Can Long-Run Restrictions Identify Technology Shocks?” *Journal of the European Economic Association*, vol. 3 (December), pp. 1237–78.
- Faust, Jon, and Eric Leeper (1997). “When Do Long-Run Identifying Restrictions Give Reliable Results?” *Journal of Business and Economic Statistics*, vol. 15 (July), pp. 345–53.
- Fernandez-Villaverde, Jesus, Juan F. Rubio-Ramirez, and Thomas J. Sargent (2005). “A, B, C’s (and D’s) for Understanding VARs,” unpublished paper, New York University.
- Fisher, Jonas (2006). “The Dynamic Effects of Neutral and Investment-Specific Technology Shocks,” *Journal of Political Economy*, vol. 114, no. 3, June.
- Francis, Neville, Michael T. Owyang, and Jennifer E. Roush (2005). “A Flexible Finite-Horizon Identification of Technology Shocks,” Working Paper Series 2005-024A. St. Louis: Federal Reserve Bank of St. Louis, April.
- Francis, Neville, and Valerie A. Ramey (2005). “Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited,” *Journal of Monetary Economics*, vol. 52 (November), pp. 1379–99.
- Gali, Jordi (1999). “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?” *American Economic Review*, vol. 89 (March), pp. 249–71.
- Gali, Jordi, and Pau Rabanal (2005). “Technology Shocks and Aggregate Fluctuations: How Well Does the Real Business Cycle Model Fit Postwar U.S. Data?” in Mark Gertler and Kenneth Rogoff, eds., *NBER Macroeconomics Annual 2004*. Cambridge, Mass.: MIT Press.
- Hamilton, James D. (1994). *Time Series Analysis*. Princeton: Princeton University Press.
- (1997). “Measuring the Liquidity Effect,” *American Economic Review*, vol. 87 (March), pp. 80–97.
- Hansen, Lars, and Thomas Sargent (1980). “Formulating and Estimating Dynamic Linear Rational Expectations Models,” *Journal of Economic Dynamics and Control*, vol. 2 (February), pp. 7–46.
- King, Robert G., Charles I. Plosser, James H. Stock, and Mark W. Watson (1991). “Stochastic Trends and Economic Fluctuations,” *American Economic Review*, vol. 81 (September), pp. 819–40.
- McGrattan, Ellen (2006). “A Critique of Structural VARs Using Business Cycle Theory,” presentation at the annual meeting of the American Economic Association, Boston, January 6–8, <http://minneapolisfed.org/research/economists/mcgrattan/CEV/assa.htm>.
- Pagan, Adrian R., and John C. Robertson (1998). “Structural Models of the Liquidity Effect,” *Review of Economics and Statistics*, vol. 80 (May), pp. 202–17.

- Ravenna, Federico, 2005, 'Vector Autoregressions and Reduced Form Representations of Dynamic Stochastic General Equilibrium Models,' unpublished manuscript.
- Rotemberg, Julio J., and Michael Woodford (1992). "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity," *Journal of Political Economy*, vol. 100 (December), pp. 1153–1207.
- (1997). "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," in Ben S. Bernanke and Julio Rotemberg, eds., *NBER Macroeconomics Annual 1997*. Cambridge, Mass.: MIT Press.
- Runkle, David E. (1987). "Vector Autoregressions and Reality," *Journal of Business and Economic Statistics*, vol. 5 (October), pp. 437–442.
- Shapiro, Matthew, and Mark Watson (1988). "Sources of Business Cycle Fluctuations," in Stanley Fischer, ed., *NBER Macroeconomics Annual 1988*. Cambridge, Mass.: MIT Press.
- Sims, Christopher (1972). "The Role of Approximate Prior Restrictions in Distributed Lag Estimation," *Journal of the American Statistical Association*, vol. 67 (March), pp. 169–75.
- (1980). "Macroeconomics and Reality," *Econometrica*, vol. 48 (January), pp. 1–48.
- (1986). "Are Forecasting Models Usable for Policy Analysis?" Federal Reserve Bank of Minneapolis, *Quarterly Review*, vol. 10 (Winter), pp. 2–16.
- (1989). "Models and Their Uses," *American Journal of Agricultural Economics*, vol. 71 (May), pp. 489–94.
- Sims, Christopher, and Tao Zha (1999). "Error Bands for Impulse Responses," *Econometrica*, vol. 67 (September), pp. 1113–55.
- (forthcoming). "Does Monetary Policy Generate Recessions?" *Macroeconomic Dynamics*.
- Smets, Frank, and Raf Wouters (2003). "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area," *Journal of the European Economic Association*, vol. 1 (September), pp. 1123–75.
- Vigfusson, Robert J. (2004). "The Delayed Response to a Technology Shock: A Flexible Price Explanation," International Finance Discussion Paper Series 2004-810. Washington: Board of Governors of the Federal Reserve System, July.
- Woodford, Michael M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press.

A. Appendix A: A Model with Nominal Wage and Price Rigidities

This appendix describes the ACEL model used in section 6. The model economy is composed of households, firms, and a monetary authority.

There is a continuum of households, indexed by $j \in (0, 1)$. The j^{th} household is a monopoly supplier of a differentiated labor service, and sets its wage subject to Calvo-style wage frictions. In general, households earn different wage rates and work different amounts. A straightforward extension of arguments in Erceg, Henderson, and Levin (2000) and in Woodford (1996) establishes that in the presence of state contingent securities, households are homogeneous with respect to consumption and asset holdings.³³ Our notation reflects this result. The preferences of the j^{th} household are given by:

$$E_t^j \sum_{l=0}^{\infty} \beta^{l-t} \left[\log(C_{t+l} - bC_{t+l-1}) - \psi_L \frac{h_{j,t+l}^2}{2} \right],$$

where $\psi_L \geq 0$ and E_t^j is the time t expectation operator, conditional on household j 's time t information set. The variable, C_t , denotes time t consumption and h_{jt} denotes time t hours worked. The household's asset evolution equation is given by:

$$\begin{aligned} M_{t+1} = & R_t [M_t - Q_t + (x_t - 1)M_t^a] + A_{j,t} + Q_t + W_{j,t}h_{j,t} \\ & + D_t - (1 + \eta(V_t))P_t C_t. \end{aligned}$$

Here, M_t and Q_t denote, respectively, the household's stock of money, and cash balances at the beginning of period t . The variable $W_{j,t}$ represents the nominal wage rate at time t . In addition D_t and $A_{j,t}$ denote firm profits and the net cash inflow from participating in state-contingent security markets at time t . The variable, x_t , represents the gross growth rate of the economy-wide per capita stock of money, M_t^a . The quantity $(x_t - 1)M_t^a$ is a lump-sum payment to households by the monetary authority. The household deposits $M_t - Q_t + (x_t - 1)M_t^a$ with a financial intermediary. The variable, R_t , denotes the gross interest rate. The variable, V_t , denotes the time t velocity of the household's cash balances:

$$V_t = \frac{P_t C_t}{Q_t}, \tag{A.1}$$

where $\eta(V_t)$ is increasing and convex.³⁴ For the quantitative analysis of our model, we must specify the level and the first two derivatives of the transactions function, $\eta(V)$, evaluated

³³Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin (2000). "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics*, vol. 46 (October), pp. 281–313.

Woodford, Michael M. (1996). "Control of the Public Debt: A Requirement for Price Stability?" NBER Working Paper Series 5684. Cambridge, Mass.: National Bureau of Economic Research, July.

³⁴Similar specifications have been used by authors such as Sims (1994) and Schmitt-Grohé and Uribe (2004). (Schmitt-Grohé, Stefanie, and Martin Uribe (2004). "Optimal Fiscal and Monetary Policy under Sticky Prices," *Journal of Economic Theory*, vol. 114 (February), pp. 198–230. Sims, Christopher, (1994), "A Simple Model for Study of the Determination of the Price Level and the Interaction of Monetary and Fiscal Policy," *Economic Theory*, vol. 4 (3), 381–99.)

in steady state. We denote these by η , η' , and η'' , respectively. Let ϵ denote the interest semi-elasticity of money demand in steady state:

$$\epsilon \equiv -\frac{100 \times d \log\left(\frac{Q_t}{P_t}\right)}{400 \times dR_t}.$$

Let V and η denote the values of velocity and $\eta(V_t)$ in steady state. ACEL parameterize the second-order Taylor series expansion of $\eta(\cdot)$ about steady state. The values of η , η' , and η'' , are determined by ACEL's estimates of ϵ , V , and η .

The j^{th} household is a monopoly supplier of a differentiated labor service, h_{jt} . It sells this service to a representative, competitive firm that transforms it into an aggregate labor input, L_t , using the technology:

$$H_t = \left[\int_0^1 h_{j,t}^{\frac{1}{\lambda_w}} dj \right]^{\lambda_w}, \quad 1 \leq \lambda_w < \infty.$$

Let W_t denote the aggregate wage rate, i.e., the nominal price of H_t . The household takes H_t and W_t as given.

In each period, a household faces a constant probability, $1 - \xi_w$, of being able to re-optimize its nominal wage. The ability to re-optimize is independent across households and time. If a household cannot re-optimize its wage at time t , it sets $W_{j,t}$ according to:

$$W_{j,t} = \pi_{t-1} \mu_{z^*} W_{j,t-1},$$

where $\pi_{t-1} \equiv P_{t-1}/P_{t-2}$. The presence of μ_{z^*} implies that there are no distortions from wage dispersion along the steady state growth path.

At time t a final consumption good, Y_t , is produced by a perfectly competitive, representative final good firm. This firm produces the final good by combining a continuum of intermediate goods, indexed by $i \in [0, 1]$, using the technology

$$Y_t = \left[\int_0^1 y_t(i)^{\frac{1}{\lambda_f}} di \right]^{\lambda_f}, \quad (\text{A.2})$$

where $1 \leq \lambda_f < \infty$ and $y_t(i)$ denotes the time t input of intermediate good i . The firm takes its output price, P_t , and its input prices, $P_t(i)$, as given and beyond its control.

Intermediate good i is produced by a monopolist using the following technology:

$$y_t(i) = \begin{cases} K_t(i)^\alpha (Z_t h_t(i))^{1-\alpha} - \phi z_t^* & \text{if } K_t(i)^\alpha (Z_t h_t(i))^{1-\alpha} \geq \phi z_t^* \\ 0 & \text{otherwise} \end{cases} \quad (\text{A.3})$$

where $0 < \alpha < 1$. Here, $h_t(i)$ and $K_t(i)$ denote time t labor and capital services used to produce the i^{th} intermediate good. The variable Z_t represents a time t shock to the technology for producing intermediate output. The growth rate of Z_t , Z_t/Z_{t-1} , is denoted by μ_{z_t} . The non-negative scalar, ϕ , parameterizes fixed costs of production. To express the model in terms of a stochastic steady state, we find it useful to define the variable z_t^* as:

$$z_t^* = \Upsilon_t^{\frac{\alpha}{1-\alpha}} Z_t, \quad (\text{A.4})$$

where Υ_t represents a time t shock to capital-embodied technology. The stochastic process generating Z_t is defined by (6.3) and (6.4). The stochastic process generating Υ_t is defined by (6.1) and (6.2).

Intermediate good firms hire labor in perfectly competitive factor markets at the wage rate, W_t . Profits are distributed to households at the end of each time period. We assume that the firm must borrow the wage bill in advance at the gross interest rate, R_t .

In each period, the i^{th} intermediate goods firm faces a constant probability, $1 - \xi_p$, of being able to re-optimize its nominal price. The ability to re-optimize prices is independent across firms and time. If firm i cannot re-optimize, it sets $P_t(i)$ according to:

$$P_t(i) = \pi_{t-1} P_{t-1}(i). \quad (\text{A.5})$$

Let $\bar{K}_t(i)$ denote the physical stock of capital available to the i^{th} firm at the beginning of period t . The services of capital, $K_t(i)$ are related to stock of physical capital, by:

$$\bar{K}_t(i) = u_t(i) \bar{K}_t(i).$$

Here $u_t(i)$ is firm i 's capital utilization rate. The cost, in investment goods, of setting the utilization rate to $u_t(i)$ is $a(u_t(i)) \bar{K}_t(i)$, where $a(\cdot)$ is increasing and convex. We assume that $u_t(i) = 1$ in steady state and $a(1) = 0$. These two conditions determine the level and slope of $a(\cdot)$ in steady state. To implement our log-linear solution method, we must also specify a value for the curvature of a in steady state, $\sigma_a = a''(1)/a'(1) \geq 0$.

There is no technology for transferring capital between firms. The only way a firm can change its stock of physical capital is by varying the rate of investment, $I_t(i)$, over time. The technology for accumulating physical capital by intermediate good firm i is given by:

$$F(I_t(i), I_{t-1}(i)) = \left(1 - S \left(\frac{I_t(i)}{I_{t-1}(i)}\right)\right) I_t(i),$$

where

$$\bar{K}_{t+1}(i) = (1 - \delta) \bar{K}_t(i) + F(I_t(i), I_{t-1}(i)).$$

The adjustment cost function, S , satisfies $S = S' = 0$, and $S'' > 0$ in steady state. Given the log-linearization procedure used to solve the model, we need not specify any other features of the function S .

The present discounted value of the i^{th} intermediate good's net cash flow is given by:

$$E_t \sum_{j=0}^{\infty} \beta^j v_{t+j} \left\{ P_{t+j}(i) y_{t+j}(i) - R_{t+j} W_{t+j} h_t(i) - P_{t+j} \Upsilon_{t+j}^{-1} [I_{t+j}(i) + a(u_{t+j}(i)) \bar{K}_{t+j}(i)] \right\}, \quad (\text{A.6})$$

where R_t denotes the gross nominal rate of interest.

The monetary policy rule is defined by (6.5) and (6.6). Financial intermediaries receive $M_t - Q_t + (x_t - 1) M_t$ from the household. Our notation reflects the equilibrium condition,

$M_t^a = M_t$. Financial intermediaries lend all of their money to intermediate good firms, which use the funds to pay labor wages. Loan market clearing requires that:

$$W_t H_t = x_t M_t - Q_t. \quad (\text{A.7})$$

The aggregate resource constraint is:

$$(1 + \eta(V_t))C_t + \Upsilon_t^{-1} [I_t + a(u_t)\bar{K}_t] \leq Y_t. \quad (\text{A.8})$$

Tables B1 and B2 report the parameter values of the model. We refer the reader to ACEL for a description of how the model is solved and for the methodology used to estimate the model parameters. The data and programs, as well as an extensive technical appendix, may be found at the following website:

www.faculty.econ.northwestern.edu/faculty/christiano/research/ACEL/accelweb.htm.

B. Appendix B: Long-Run Identification of Two Technology Shocks

This appendix generalizes the strategy for long-run identification of one shock to two shocks, using the strategy of Fisher (2006). As before, the VAR is:

$$\begin{aligned} Y_{t+1} &= B(L)Y_t + u_t, \quad E u_t u_t' = V, \\ B(L) &\equiv B_1 + B_2 L + \dots + B_q L^{q-1}, \end{aligned}$$

We suppose that the fundamental shocks are related to the VAR disturbances as follows:

$$u_t = C\varepsilon_t, \quad E\varepsilon_t \varepsilon_t' = I, \quad CC' = V,$$

where the first two element in ε_t are $\varepsilon_{\mu_{\Upsilon,t}}$ and ε_t^z , respectively. The exclusion restrictions are:

$$\begin{aligned} \lim_{j \rightarrow \infty} \left[\tilde{E}_t a_{t+j} - \tilde{E}_{t-1} a_{t+j} \right] &= f_z(\varepsilon_{\mu_{\Upsilon,t}}, \varepsilon_t^z, \text{only}) \\ \lim_{j \rightarrow \infty} \left[\tilde{E}_t \log p_{I,t+j} - \tilde{E}_{t-1} \log p_{I,t+j} \right] &= f_{\Upsilon}(\varepsilon_{\mu_{\Upsilon,t}}, \text{only}). \end{aligned}$$

That is, only technology shocks have a long-run effect on the log-level of labor productivity, whereas only capital-embodied shocks have a long-run effect on the log-level of the price of investment goods. According to the sign restrictions, the slope of f_z with respect to its second argument and the slope of f_{Υ} are non-negative. Applying a suitably modified version of the logic in Section 2.3.1, we conclude that, according to the exclusion restrictions, the indicated pattern of zeros must appear in the following 3 by 3 matrix:

$$[I - B(1)]^{-1} C = \begin{bmatrix} a & 0 & 0 \\ b & c & 0 \\ \text{number} & \text{number} & \text{number} \end{bmatrix}$$

The sign restrictions are $a, c > 0$. To compute the dynamic response of Y_t to the two technology shocks, we require the first two columns of C . To obtain these, we proceed as follows. Let $D \equiv [I - B(1)]^{-1} C$, so that:

$$DD' = [I - B(1)]^{-1} V [I - B(1)']^{-1} = S_Y(0), \quad (\text{B.1})$$

where, as before, $S_Y(0)$ is the spectral density of Y_t at frequency-zero, as implied by the estimated VAR. The exclusion restrictions require that D have the following structure:

$$D = \begin{bmatrix} d_{11} & 0 & 0 \\ d_{21} & d_{22} & 0 \\ d_{31} & d_{32} & d_{33} \end{bmatrix}.$$

Here, the zero restrictions reflect our exclusion restrictions, and the sign restrictions require $d_{11}, d_{22} \geq 0$. Then,

$$DD' = \begin{bmatrix} d_{11}^2 & d_{11}d_{21} & d_{11}d_{31} \\ d_{21}d_{11} & d_{21}^2 + d_{22}^2 & d_{21}d_{31} + d_{22}d_{32} \\ d_{31}d_{11} & d_{31}d_{21} + d_{32}d_{22} & d_{31}^2 + d_{32}^2 + d_{33}^2 \end{bmatrix} = \begin{bmatrix} S_Y^{11}(0) & S_Y^{21}(0) & S_Y^{31}(0) \\ S_Y^{21}(0) & S_Y^{22}(0) & S_Y^{32}(0) \\ S_Y^{31}(0) & S_Y^{32}(0) & S_Y^{33}(0) \end{bmatrix}$$

and

$$d_{11} = \sqrt{S_Y^{11}(0)}, \quad d_{21} = S_Y^{21}(0) / d_{11}, \quad d_{31} = S_Y^{31}(0) / d_{11}$$

$$d_{22} = \sqrt{\frac{S_Y^{11}(0) S_Y^{22}(0) - (S_Y^{21}(0))^2}{S_Y^{11}(0)}}, \quad d_{32} = \frac{S_Y^{32}(0) - S_Y^{21}(0) S_Y^{31}(0) / d_{11}^2}{d_{22}}.$$



The sign restrictions imply that the square roots should be positive. The fact that $S_Y(0)$ is positive definite ensures that the square roots are real numbers. Finally, the first two columns of C are calculated as follows:

$$\begin{bmatrix} C_1 \\ C_2 \end{bmatrix} = [I - B(1)] \begin{bmatrix} D_1 \\ D_2 \end{bmatrix},$$

where C_i is the i^{th} column of C and D_i is the i^{th} column of D , $i = 1, 2$.





To construct our modified VAR procedure, simply replace $S_Y(0)$ in (B.1) by (4.4).

Figure 1: A Simulated Time Series for Hours



— Both Shocks
— Only Technology Shocks

Figure 2: Short-run Identification Results





—* True Response
— Estimated Response



Sampling Distribution
 Average CI Standard Deviation Based
★ Average CI Percentile Based

Figure 3: Coverage Rates, Percentile-Based Confidence Intervals

Short-Run Identification



Long-Run Identification





—●— 2 Shock Standard
 —■— 3 Shock Standard



—○— 2 Shock Bartlett
 —□— 3 Shock Bartlett

Figure 4: Coverage Rates, Standard Deviation–Based Confidence Intervals

Short-Run Identification



Long-Run Identification



● 2 Shock Standard
 ■ 3 Shock Standard

○ 2 Shock Bartlett
 □ 3 Shock Bartlett

Figure 5: Long-run Identification Results











Figure 6: Analyzing Precision in Inference

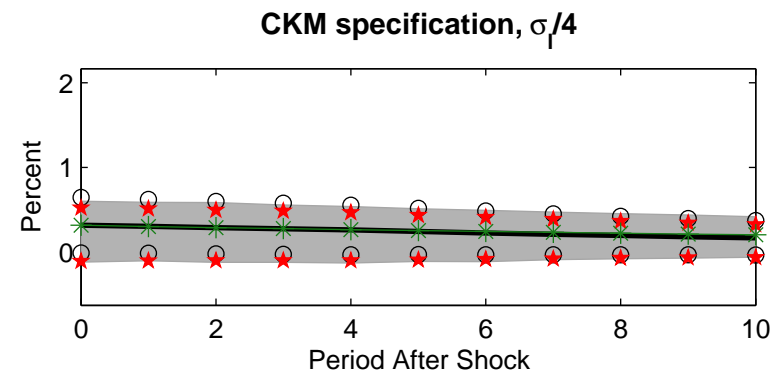







Figure 7: Varying the Labor Elasticity in the Two-shock CKM Specification





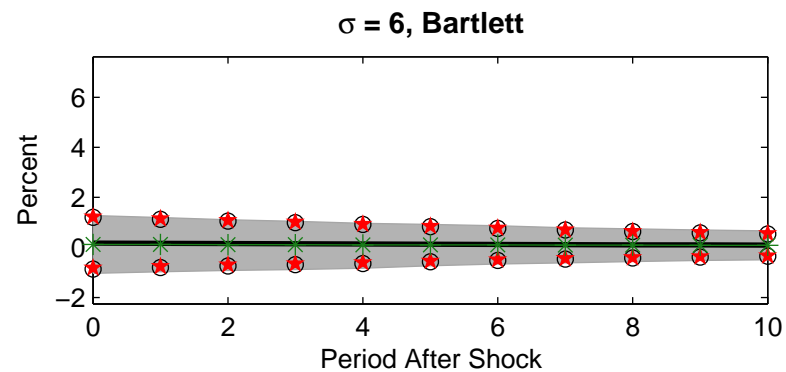
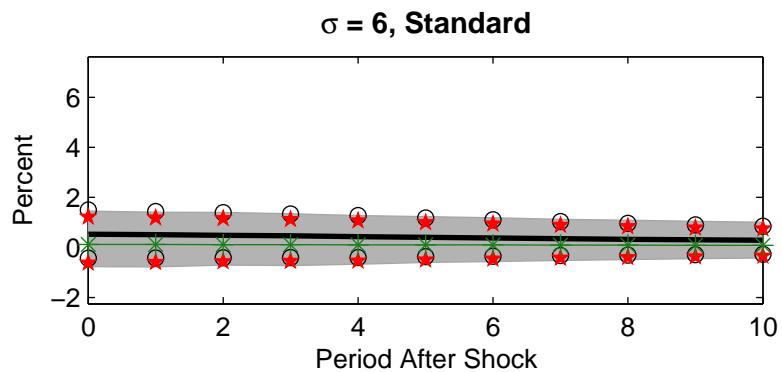





Figure 8: Impact of C_1 on Distortions




- Standard Long-Run Identification
- ◆ Response Using Estimated B(L) and True C_1
- * True Response

Figure 9: Analysis of Long-run Identification Results



Increased Persistence in Preference Shock ($\rho_1 = .998$ $\sigma_1 = 0.0028$)



Contemporaneous Impact of Technology on Hours






Figure 10: Comparing Long- and Short-Run Identification

Recursive Two-Shock CKM Specification



—*— True Response
— Estimated Response

■ Sampling Distribution
○ Average CI Standard Deviation Based
★ Average CI Percentile Based

Figure 11: The Treatment of CKM Measurement Error



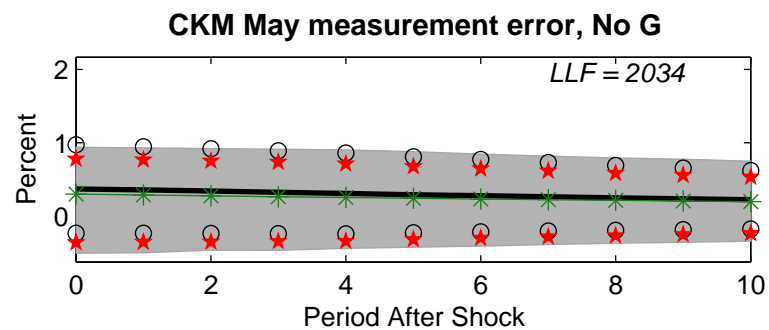


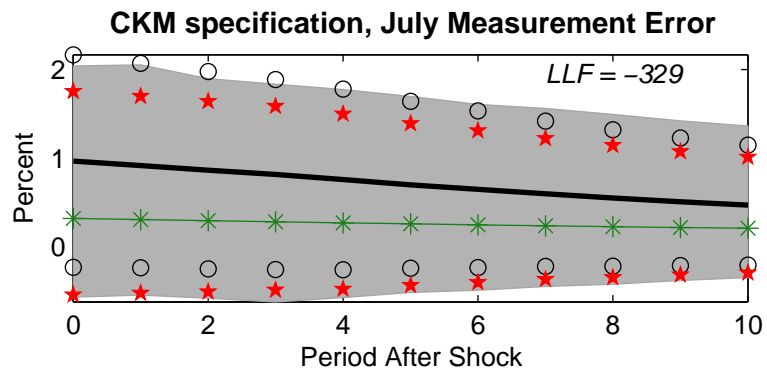


Figure 12: Stochastic Process Uncertainty







Figure 13: Data Sensitivity and Inference in VARs


Estimated Hours Response Starting in 1948



Estimated Hours Response Starting in 1959



Hours Per Capita



--- FR — CEV —○— GR

Figure 14: Impulse Response Results when the ACEL Model is the DGP






Table 1: Contribution of Technology Shocks to Volatility

Model specification		Measure of Variation					
		Unfiltered		HP-filtered		One-step-ahead forecast error	
		$\ln l_t$	$\Delta \ln y_t$	$\ln l_t$	$\ln y_t$	$\ln l_t$	$\Delta \ln y_t$
MLE							
Base	Nonrecursive	3.73	67.16	7.30	67.14	7.23	67.24
	Recursive	3.53	58.47	6.93	64.83	0.00	57.08
$\sigma_l/2$	Nonrecursive	13.40	89.13	23.97	89.17	23.77	89.16
	Recursive	12.73	84.93	22.95	88.01	0.00	84.17
$\sigma_l/4$	Nonrecursive	38.12	97.06	55.85	97.10	55.49	97.08
	Recursive	36.67	95.75	54.33	96.68	0.00	95.51
$\sigma = 6$	Nonrecursive	3.26	90.67	6.64	90.70	6.59	90.61
	Recursive	3.07	89.13	6.28	90.10	0.00	88.93
$\sigma = 0$	Nonrecursive	4.11	53.99	7.80	53.97	7.73	54.14
	Recursive	3.90	41.75	7.43	50.90	0.00	38.84
Three	Nonrecursive	0.18	45.67	3.15	45.69	3.10	45.72
	Recursive	0.18	36.96	3.05	43.61	0.00	39.51
CKM							
Base	Nonrecursive	2.76	33.50	1.91	33.53	1.91	33.86
	Recursive	2.61	25.77	1.81	31.41	0.00	24.93
$\sigma_l/2$	Nonrecursive	10.20	66.86	7.24	66.94	7.23	67.16
	Recursive	9.68	58.15	6.88	64.63	0.00	57.00
$\sigma_l/4$	Nonrecursive	31.20	89.00	23.81	89.08	23.76	89.08
	Recursive	29.96	84.76	22.79	87.91	0.00	84.07
$\sigma = 6$	Nonrecursive	0.78	41.41	0.52	41.33	0.52	41.68
	Recursive	0.73	37.44	0.49	40.11	0.00	37.42
$\sigma = 0$	Nonrecursive	2.57	20.37	1.82	20.45	1.82	20.70
	Recursive	2.44	13.53	1.73	18.59	0.00	12.33
$\sigma = 0$ and $2\sigma_l$	Nonrecursive	0.66	6.01	0.46	6.03	0.46	6.12
	Recursive	0.62	3.76	0.44	5.41	0.00	3.40
Three	Nonrecursive	2.23	30.73	1.71	31.11	1.72	31.79
	Recursive	2.31	23.62	1.66	29.67	0.00	25.62

Note: (a) V_h corresponds to the columns denoted by $\ln(l_t)$.

(b) In each case, the results report the ratio of two variances: the numerator is the variance for the system with only technology shocks and the denominator is the variance for the system with both technology shock and labor tax shocks. All statistics are averages of the ratios, based on 300 simulations of 5000 observations for each model.

(c) ‘Base’ means the two-shock specification, whether MLE or CKM, as indicated. Three’ means the three-shock specification.

(d) For a description of the procedure used to calculate the forecast error variance, see footnote 13.

(e) ‘MLE’ and ‘CKM’ refer, respectively, to our and CKM’s estimated models.

Table 2: Properties of Two-Shock CKM Specification

Panel A: First Six Lag Matrices in Infinite-Order VAR Representation

$$B_1 = \begin{bmatrix} 0.013 & 0.041 \\ 0.0065 & 0.94 \end{bmatrix}, B_2 = \begin{bmatrix} 0.012 & -0.00 \\ 0.0062 & -0.00 \end{bmatrix}, B_3 = \begin{bmatrix} 0.012 & -0.00 \\ 0.0059 & -0.00 \end{bmatrix},$$

$$B_4 = \begin{bmatrix} 0.011 & -0.00 \\ 0.0056 & -0.00 \end{bmatrix}, B_5 = \begin{bmatrix} 0.011 & -0.00 \\ 0.0054 & -0.00 \end{bmatrix}, B_6 = \begin{bmatrix} 0.010 & -0.00 \\ 0.0051 & -0.00 \end{bmatrix}$$

Panel B: Population Estimate of Four-lag VAR

$$\hat{B}_1 = \begin{bmatrix} 0.017 & 0.043 \\ 0.0087 & 0.94 \end{bmatrix}, \hat{B}_2 = \begin{bmatrix} 0.017 & -0.00 \\ 0.0085 & -0.00 \end{bmatrix}, \hat{B}_3 = \begin{bmatrix} 0.012 & -0.00 \\ 0.0059 & -0.00 \end{bmatrix},$$

$$\hat{B}_4 = \begin{bmatrix} 0.0048 & -0.0088 \\ 0.0025 & -0.0045 \end{bmatrix}$$

Panel C: Actual and Estimated Sum of VAR Coefficients

$$\hat{B}(1) = \begin{bmatrix} 0.055 & 0.032 \\ 0.14 & 0.94 \end{bmatrix}, B(1) = \begin{bmatrix} 0.28 & 0.022 \\ 0.14 & 0.93 \end{bmatrix}, \sum_{j=1}^4 B_j = \begin{bmatrix} 0.047 & 0.039 \\ 0.024 & 0.94 \end{bmatrix}$$

Panel D: Actual and Estimated Zero-Frequency Spectral Density

$$S_Y(0) = \begin{bmatrix} 0.00017 & 0.00097 \\ 0.00097 & 0.12 \end{bmatrix}, \hat{S}_Y(0) = \begin{bmatrix} 0.00012 & 0.0022 \\ 0.0022 & 0.13 \end{bmatrix}.$$

Panel E: Actual and Estimated One-Step-Ahead Forecast Error Variance

$$V = \hat{V} = \begin{bmatrix} 0.00012 & -0.00015 \\ -0.00015 & -0.00053 \end{bmatrix}$$

Panel F: Actual and Estimated Impact Vector

$$C_1 = \begin{pmatrix} 0.00773 \\ 0.00317 \end{pmatrix}, \hat{C}_1 = \begin{pmatrix} 0.00406 \\ 0.01208 \end{pmatrix}$$

Table 3: Percent Contribution of Shocks in the ACEL model to the Variation in Hours and in Output

Statistic	Types of shock		
	Monetary Policy	Neutral Technology	Capital-Embodied
variance of logged hours	22.2	40.0	38.5
variance of HP filtered logged hours	37.8	17.7	44.5
variance of Δy	29.9	46.7	23.6
variance of HP filtered logged output	31.9	32.3	36.1

Note: Results are average values based on 500 simulations of 3100 observations each.

ACEL: Altig Christiano, Eichenbaum and Linde (2005).