A Method for Taking Models to the Data

Peter N. Ireland*
Boston College and NBER

April 1999

Abstract

This paper develops a method for combining the power of a dynamic, stochastic, general equilibrium model with the flexibility of a vector autoregressive time-series model to obtain a hybrid that can be taken directly to the data. It estimates this hybrid model via maximum likelihood and uses the results to address a number of issues concerning the ability of a prototypical real business cycle model to explain movements in aggregate output and employment in the postwar US economy, the stability of the real business cycle model's structural parameters, and the performance of the hybrid model's out-of-sample forecasts.

JEL: C51, C52, E32, E37.

1. Introduction

Two distinct approaches to macroeconomic analysis emerged during the early 1980s and continue to yield insights today. First, work following Sims (1980) characterizes and attempts to explain the movements and co-movements of key aggregate variables using vector autoregressive (VAR) time-series models. Second, work following Kydland and Prescott (1982) characterizes and attempts to explain

*Please address correspondence to: Peter N. Ireland, Boston College, Department of Economics, Carney Hall, 140 Commonwealth Avenue, Chestnut Hill, MA 02467-3806. Tel: (617) 552-3687. Fax: (617) 552-2308. E-mail: peter.ireland@bc.edu. I would like to thank seminar participants at Boston University, Brandeis University, Brigham Young University, Clark University, the Federal Reserve Bank of Kansas City, Indiana University, and the University of Pennsylvania for very helpful comments and suggestions.

the movements and co-movements of many of the same variables using dynamic, stochastic, general equilibrium (DSGE) models.

These two distinct approaches to macroeconomics both have their distinct strengths and weaknesses. VAR models, for instance, are designed to be taken directly to the data: they are easy to estimate and, once estimated, can be used to perform statistical hypothesis tests as well as to generate out-of-sample forecasts. Moreover, since their specification requires little, if any, reference to detailed economic theory, VAR models remain flexible enough to address a wide range of questions regarding the nature and sources of business cycle fluctuations. Because they rely so loosely on economic theory, however, VAR models often fail to uncover parameters that are truly structural; thus, these models may exhibit instability during periods when monetary and fiscal policies change. Indeed, Stock and Watson (1996) find evidence of widespread instability in VAR models estimated with postwar US data.

DSGE models, by contrast, are firmly grounded in economic theory. These models draw tight links between the structural parameters describing private agents' tastes and technologies and the time-series behavior of endogenous variables such as aggregate output and employment; in principle, at least, these structural parameters should remain invariant to changes in policy regimes. Yet because they rely so heavily on economic theory, DSGE models are often regarded as being too stylized to be taken directly to the data, making traditional econometric methods for estimation, hypothesis testing, and forecasting inapplicable. Moreover, since they take such a strong stand on so many details concerning the structure of the economy, DSGE models often yield results that appear to be fragile, at least at first glance. When Kydland and Prescott (1982) report, for instance, that technology shocks can account for most of the observed output variation in the postwar US data, one is still left to wonder whether this result will survive modifications to their model, such as the introduction of other types of shocks.

This paper develops a method for combining the power of dynamic, stochastic, general equilibrium theory with the flexibility of vector autoregressive time-series models, in hopes of obtaining a hybrid that shares the desirable features of both approaches to macroeconomics. This method takes as its starting point a fully-specified DSGE model, but also admits that while this model may be powerful enough to account for and explain many key features of the US data, it remains too stylized to possibly capture all of the dynamics that can be found in those data. Hence, it augments the DSGE model so that its residuals—meaning the movements in the data that the theory cannot explain—are described by a VAR, making estimation, hypothesis testing, and forecasting feasible.

To illustrate how this method works, the remainder of the paper proceeds as follows. The next section outlines a prototypical DSGE model: Hansen's (1985) real business cycle model with indivisible labor. Section 3 then augments this model with VAR residuals to arrive at the hybrid specification described above. Section 4 estimates the hybrid model via maximum likelihood and uses the estimated model to address a number of key issues concerning the ability of the real business cycle model to explain movements in aggregate output, consumption, investment, and hours worked in the postwar US data, the stability of the real business cycle model's structural parameters, and the performance of the hybrid model's out-of-sample forecasts. Finally, section 5 concludes.

Before proceeding, however, mention should be made of some previous work that also estimates the structural parameters of DSGE models using maximum likelihood methods, including Christiano (1988), Altug (1989), Bencivenga (1992), McGrattan (1994), Kim (1995), Hall (1996), DeJong, Ingram, and Whiteman (1997*a*, 1997*b*), Ireland (1997), McGrattan, Rogerson, and Wright (1997), and Chow and Kwan (1998). In fact, a number of these earlier studies draw on a framework for combining models and data, originally developed by Sargent (1989), that also lies at the heart of the approach taken here. Thus, the connections between the current study and much of this previous work are noted below.

2. A Prototypical DSGE Model

In Hansen's (1985) real business cycle model with indivisible labor, a representative consumer has preferences defined over consumption C_t and hours worked H_t during each period $t = 0, 1, 2, ...$, as described by the expected utility function

$$E \sum_{t=0}^{\infty} \beta^t [\ln(C_t) - \gamma H_t], \tag{1}$$

where the discount factor satisfies $1 > \beta > 0$ and where $\gamma > 0$. The linearity of utility in hours worked can be motivated, following Hansen (1985) and Rogerson (1988), by assuming that the economy consists of many individual consumers, each of whom either works full time or remains unemployed.

The representative consumer produces output Y_t with capital K_t and labor H_t according to the constant-returns-to-scale technology described by

$$Y_t = A_t K_t^{\theta} (\eta^t H_t)^{1-\theta}, \tag{2}$$

where $\eta > 1$ denotes the gross rate of labor-augmenting technological progress and where $1 > \theta > 0$. The technology shock A_t follows the first-order autoregressive process

$$\ln(A_{t+1}) = (1-\rho)\ln(A) + \rho\ln(A_t) + \varepsilon_{t+1}, \tag{3}$$

where $A > 0$ and $1 > \rho > -1$. The serially uncorrelated innovation ε_{t+1} is normally distributed with mean zero and standard deviation σ.

During each period $t = 0, 1, 2, ...$, the representative consumer divides output Y_t between consumption C_t and investment I_t, subject to the resource constraint

$$Y_t = C_t + I_t. \tag{4}$$

By investing I_t units of output during period t, the consumer increases the capital stock K_{t+1} available during period $t+1$ according to

$$K_{t+1} = (1-\delta)K_t + I_t, \tag{5}$$

where the depreciation rate satisfies $1 > \delta > 0$.

Equilibrium allocations for this economy can be characterized by solving the representative consumer's problem: choose sequences $\{Y_t, C_t, I_t, H_t, K_{t+1}\}_{t=0}^{\infty}$ to maximize the utility function (1) subject to the constraints (2)-(5) for all $t = 0, 1, 2, ...$ This problem lacks a closed-form solution, but approximate solutions may be constructed numerically as follows.

Define $y_t = Y_t/\eta^t$, $c_t = C_t/\eta^t$, $i_t = I_t/\eta^t$, $h_t = H_t$, $k_t = K_t/\eta^t$, and $a_t = A_t$. The first-order conditions for the consumer's problem imply that in the absence of technology shocks, when $\varepsilon_t = 0$ for all $t = 0, 1, 2, ...$, the economy converges to a steady state in which each of these detrended variables is constant, with $y_t = y$, $c_t = c$, $i_t = i$, $h_t = h$, $k_t = k$, and $a_t = a$ for all $t = 0, 1, 2, ...$ The first-order conditions can be log-linearized about this steady state, and the methods of Blanchard and Kahn (1980) can be applied to this log-linear system to obtain a solution of the form

$$\mathbf{s}_{t+1} = \mathbf{A}\mathbf{s}_t + \mathbf{B}\varepsilon_{t+1} \tag{6}$$

and

$$\mathbf{f}_t = \mathbf{C}\mathbf{s}_t \tag{7}$$

for all $t = 0, 1, 2, ...$, where the vectors $\mathbf{s}_t$ and $\mathbf{f}_t$ keep track of percentage deviations of each detrended variable from its steady-state level, with

$$\mathbf{s}_t = \left[\begin{array}{cc} \ln(k_t/k) & \ln(a_t/a) \end{array}\right]'$$

and

$$\mathbf{f}_t = \left[\begin{array}{cccc} \ln(y_t/y) & \ln(c_t/c) & \ln(i_t/i) & \ln(h_t/h) \end{array} \right]'.$$

In (6) and (7), the elements of the matrices **A**, **B**, and **C** depend on the real business cycle model's eight structural parameters β, γ, θ, η, δ, A, ρ, and σ.

3. The Hybrid Model

In principle, one could use data on aggregate output, consumption, investment, and hours worked, along with the solution described by (6) and (7), to estimate each of the real business cycle model's structural parameters. Many researchers, however, including Kydland and Prescott (1982), argue that models of this type are simply too stylized to explain many features of the data, making traditional econometric methods inapplicable.

Indeed, one dimension along which the real business cycle model is quite stylized lies in the assumption that just one shock—the aggregate technology shock—drives all business cycle fluctuations. As emphasized by Ingram, Kocherlakota, and Savin (1994), this one-shock assumption makes the real business cycle model singular: the model predicts that certain linear combinations of the endogenous variables will be deterministic. If, in the data, these exact linear relationships do not hold, any attempt to estimate (6) and (7) via maximum likelihood will fail.

To facilitate estimation, therefore, consider augmenting each equation in (7) with a serially correlated residual, or error term, so that the empirical model consists of (6),

$$\mathbf{f}_t = \mathbf{C}\mathbf{s}_t + \mathbf{u}_t, \tag{8}$$

and

$$\mathbf{u}_{t+1} = \mathbf{D}\mathbf{u}_t + \boldsymbol{\xi}_{t+1} \tag{9}$$

for all $t = 0, 1, 2, ...$, where the vector $\boldsymbol{\xi}_{t+1}$ of zero-mean, serially uncorrelated innovations is normally distributed with covariance matrix $E\boldsymbol{\xi}_{t+1}\boldsymbol{\xi}_{t+1}' = \mathbf{V}$ and is uncorrelated with the innovation ε_{t+1} to technology.

This approach—adding error terms to the observation equation (7)—is also used by Altug (1989), McGrattan (1994), Hall (1996), and McGrattan, Rogerson, and Wright (1997) to estimate what would otherwise be singular real business cycle models. Each of these earlier studies follows Sargent (1989) by interpreting $\mathbf{u}_t$ as a vector of measurement errors in each variable and by assuming that the matrices **D** and **V** are diagonal, so that the measurement errors are uncorrelated across variables. Here, however, no such restrictions are imposed: the residuals in $\mathbf{u}_t$ are allowed to follow a general, first-order vector autoregression. Thus, the residuals may still capture measurement errors, but they can also be interpreted more liberally as capturing all of the movements and co-movements in the data that the real business cycle model, because of its elegance and simplicity, cannot explain. In this way, the hybrid model consisting of (6), (8), and (9) combines the power of the DSGE model with the flexibility of a VAR.

4. Results from the Hybrid Model

4.1. Estimation

The empirical model consisting of (6), (8), and (9) is in state-space form; it can be estimated via maximum likelihood, using methods described by Hamilton (1994, Ch.13), once analogs to the model's Y_t, C_t, I_t, and H_t are found in the US data. Thus, in the data, consumption C_t is defined as real personal consumption expenditures in chained 1992 dollars, investment I_t is defined as real gross private domestic investment, also in chained 1992 dollars, and output Y_t is defined by the sum $C_t + I_t$. Hours worked H_t is defined as hours of wage and salary workers on private, nonfarm payrolls. Each series is converted to per-capita terms by dividing by the civilian, noninstitutional population, age 16 and over.

All data, except for population, are seasonally adjusted. Since the real business cycle model implies that output, consumption, and investment grow at the common rate η in steady state, the data are automatically detrended as part of the estimation procedure; they are not filtered in any other way. Data for consumption, investment, output, and population are taken from the Federal Reserve Bank of St. Louis' FRED database; data for hours worked come from the Bureau of Labor Statistics' Establishment Survey. The series are quarterly and run from 1960:1 through 1997:3.

The resource constraint (4) holds by construction in the data. Thus, only the series for Y_t, C_t, and H_t are used in estimating the model; the series for I_t is redundant. For the purposes of estimation, therefore, $\mathbf{f}_t$, $\mathbf{u}_t$, and $\boldsymbol{\xi}_{t+1}$ reduce to 3×1 vectors, with

$$\mathbf{f}_t = \left[\begin{array}{ccc} \ln(y_t/y) & \ln(c_t/c) & \ln(h_t/h) \end{array} \right]',$$

$$\mathbf{u}_t = \left[\begin{array}{ccc} u_{yt} & u_{ct} & u_{ht} \end{array} \right]',$$

and

$$\boldsymbol{\xi}_{t+1} = \left[\begin{array}{ccc} \xi_{yt+1} & \xi_{ct+1} & \xi_{ht+1} \end{array} \right]'$$

for all $t = 0, 1, 2, ...,$ and the matrices $\mathbf{D}$ and $\mathbf{V}$ can be written as

$$\mathbf{D} = \begin{bmatrix} d_{yy} & d_{yc} & d_{yh} \\ d_{cy} & d_{cc} & d_{ch} \\ d_{hy} & d_{hc} & d_{hh} \end{bmatrix}$$

and

$$\mathbf{V} = \begin{bmatrix} v_y^2 & v_{yc} & v_{yh} \\ v_{yc} & v_c^2 & v_{ch} \\ v_{yh} & v_{ch} & v_h^2 \end{bmatrix}.$$

In estimating the hybrid model, the real business cycle model's structural parameters are constrained to satisfy the theoretical restrictions listed in section 2, above. In addition, the eigenvalues of the matrix $\mathbf{D}$ are constrained to lie inside the unit circle, so that the residuals in $\mathbf{u}_t$ are stationary. Finally, the covariance matrix $\mathbf{V}$ is constrained to be positive definite.

Preliminary attempts to apply maximum likelihood to (6), (8), and (9) led to an unreasonably low estimate of $\beta = 0.7679$ for the discount factor and an unreasonably high estimate of $\delta = 0.2165$ for the depreciation rate in this quarterly model; here, as in Altug (1989), more sensible results obtain when values of β and δ are fixed prior to estimation. Thus, table 1 reports maximum likelihood estimates of the six remaining structural parameters γ, θ, η, A, ρ, and σ, along with the 15 distinct elements of the matrices $\mathbf{D}$ and $\mathbf{V}$, with β fixed at 0.99 and δ fixed at 0.025, the values originally suggested by Hansen (1985). The standard errors, also reported in table 1, correspond to the square roots of the diagonal elements of the inverted matrix of second derivatives of the maximized log-likelihood function.

The estimates of the real business cycle model's parameters are sensible and precise. The estimate $\gamma = 0.0045$ matches steady-state hours worked in the model with average hours worked in the data; the estimate $A = 6.0952$ does the same for detrended output. The estimate $\theta = 0.2342$ implies that capital's share in production is just slightly less than 25 percent. The estimate $\eta = 1.0039$ makes the annualized, steady-state growth rate of real, per-capita output in the model equal to 1.57 percent. Finally, the estimate of $\sigma = 0.0050$ is of the same order of magnitude used throughout the literature, while the estimate $\rho = 0.9983$ implies that technology shocks are extremely persistent.

The other estimates in table 1 reveal, however, that there are important features of the data that the real business cycle model cannot explain. The estimates imply, for instance, that the matrix $\mathbf{D}$ has one real eigenvalue of modulus 0.9711 and two complex eigenvalues of modulus 0.8764; evidently, the residuals in $\mathbf{u}_t$ are nearly as persistent as the technology shock. Furthermore, the innovations in $\boldsymbol{\xi}_{t+1}$ have standard deviations of 0.0060, 0.0057, and 0.0006; two of these three figures exceed the estimated standard deviation of the innovation ε_{t+1} to technology.

4.2. Explanatory Power of the Real Business Cycle Model

What fraction of the observed output variation in the postwar US economy is explained by the real business cycle model? This question, first considered by Kydland and Prescott (1982), can also be addressed here by using (6), (8), and (9), together with the maximum likelihood estimates presented in table 1, to decompose the k-step-ahead forecast error variances in output, consumption, investment, and hours worked into two orthogonal components: one attributable to the real business cycle model's technology shock and the other attributable to the three residuals in $\mathbf{u}_t$. Table 2 displays the results of these forecast error variance decompositions.

In table 2, the last line of panel A, with $k = \infty$, indicates that the technology shock accounts for nearly 85 percent of the unconditional variance of detrended output. Kydland and Prescott also find that the real business cycle model accounts for most of the observed variation in output, but here, this result obtains despite the fact that the hybrid model also allows shocks to the elements of $\mathbf{u}_t$ to help explain the behavior of output. Presumably, the residuals in $\mathbf{u}_t$ pick up the combined effects of shocks not modelled in the real business cycle framework: shocks to monetary policy, fiscal policy, and so forth. Here, therefore, Kydland and Prescott's finding is shown to be robust to the inclusion of these other shocks.

The robustness of Kydland and Prescott's finding can also be assessed, in the context of the hybrid model estimated here, by attaching standard errors to each of the statistics reported in table 2. Thus, standard errors also appear in the table, where they are computed by expressing each statistic as a function g of the vector $\mathbf{\Theta}$ of estimated parameters and by calculating $[\partial g(\mathbf{\Theta})/\partial\mathbf{\Theta}]'\mathbf{H}[\partial g(\mathbf{\Theta})/\partial\mathbf{\Theta}]$, where $\mathbf{H}$ is the covariance matrix of the estimated parameters in $\mathbf{\Theta}$ and the derivatives $\partial g(\mathbf{\Theta})/\partial\mathbf{\Theta}$ are evaluated numerically, as suggested by Runkle (1987).

The standard errors shown in table 2 indicate that the statistical uncertainty surrounding the real business cycle model's ability to explain a substantial fraction of the observed output variation in the US data is large, though perhaps not as large as first suggested by Eichenbaum (1991), who estimates the model's parameters using a generalized method of moments procedure instead of the more efficient maximum likelihood technique used here. Even if the true fraction of output variation explained by the real business cycle model is two standard errors less than the point estimate of 85 percent, that fraction remains greater than 45 percent.

Other results displayed in table 2 show that the technology shock accounts for more than 90 percent of the unconditional variance of detrended consumption and more than 50 percent of the unconditional variance of detrended investment, but almost none of unconditional variance of hours worked. Thus, as noted by Cooley and Prescott (1995) among others, the real business cycle model does a much better job in explaining the behavior of output and its components than it does in explaining the behavior of hours worked.

As noted above, the technology shock accounts for almost 85 percent of the unconditional variance in aggregate output, and as shown in table 2, it also accounts for more than 60 percent of the one-quarter-ahead forecast error variance in output. On the other hand, the technology shock accounts for less than half of the k-step-ahead forecast error variances for values of k ranging from 4 to 40, implying that the real business cycle model has difficulty explaining output fluctuations over horizons between one and ten years. This result is, of course, consistent with previous findings reported by Watson (1993) and Rotemberg and Woodford (1996); Watson, in particular, finds that while the real business cycle model explains very high and very low frequency movements in output, it is less successful at explaining those movements that take place at business cycle frequencies.

Finally, table 2 reports a surprising result. Although, as noted above, technology shocks account for almost none of the unconditional variance of hours worked, they explain almost all of the one-quarter-ahead forecast error variance in the hours series. This result is encouraging, since it suggests that the real business cycle model may still have some success at forecasting quarter-to-quarter movements in aggregate hours worked, even if it fares less well at explaining movements over longer horizons.

4.3. Tests for Parameter Stability

One great strength of the real business cycle model is that it is structural: it links the behavior of aggregate output and employment to parameters describing private agents' tastes and technologies—parameters that should remain invariant to changes in monetary and fiscal policy regimes. Here, the hybrid model consisting of (6), (8), and (9) can be used to perform statistical tests of the hypothesis that the structural parameters have, in fact, remained stable over time.

To test for parameter stability, the hybrid model is estimated over two disjoint subsamples: the first running from 1960:1 through 1979:4 and the second running from 1980:1 through 1997:3. The 1980 breakpoint serves to divide the full sample into subsamples of roughly equal length and, more important, corresponds to a date around which major changes in US monetary and fiscal policies are widely thought to have occurred.

Table 3 reports estimates of the hybrid model's parameters for the two subsamples, along with their standard errors. Focusing on the six estimated parameters from the real business cycle model, only small differences in the estimates of γ, θ, and A appear across subsamples. The estimate of $\eta = 1.0046$ for the pre-1980 subsample exceeds the estimate of $\eta = 1.0033$ for the post-1980 subsample, reflecting the productivity slowdown. Meanwhile, for both subsamples, the estimates of ρ and σ lie below their full-sample counterparts. This result—that aggregate shocks appear smaller and less persistent when a break in the trend rate of growth is allowed for—can also be found in work by Perron (1989) and Rappoport and Reichlin (1989).

Andrews and Fair (1988) describe procedures for testing for parameter stability across the two subsamples. Let the vector $\mathbf{\Theta}_q^1$ consist of any q parameters estimated from the first subsample, let $\mathbf{\Theta}_q^2$ consist of the same q parameters estimated from the second subsample, and let $\mathbf{H}_q^1$ and $\mathbf{H}_q^2$ denote the covariance matrices of the estimates in $\mathbf{\Theta}_q^1$ and $\mathbf{\Theta}_q^2$. Andrews and Fair show that the Wald statistic

$$W = (\mathbf{\Theta}_q^1 - \mathbf{\Theta}_q^2)'(\mathbf{H}_q^1 + \mathbf{H}_q^2)^{-1}(\mathbf{\Theta}_q^1 - \mathbf{\Theta}_q^2)$$

is asymptotically distributed as a chi-square random variable with q degrees of freedom under the null hypothesis of parameter stability: $\mathbf{\Theta}_q^1 = \mathbf{\Theta}_q^2$.

Table 4 reports Wald statistics for the stability of all 21 estimated parameters, the stability of the six structural parameters γ, θ, η, A, ρ, and σ identified by the real business cycle model, and the stability of the 15 parameters in the matrices $\mathbf{D}$ and $\mathbf{V}$. The tests reject the null of stability for all 21 parameters as well as for the 15 distinct elements of $\mathbf{D}$ and $\mathbf{V}$. On the other hand, the test fails to reject the null of stability for the six structural parameters, indicating that these parameters have remained stable over time, despite important changes in monetary and fiscal policies.

4.4. Forecast Accuracy

Table 5 reports on the accuracy of the hybrid model's out-of-sample forecasts. As noted above, the model has 21 estimated parameters: the six structural parameters from the real business cycle model and the 15 parameters from **D** and **V** that describe the behavior of the residuals. An unconstrained, first-order VAR for the logs of output, consumption, and hours worked with a constant and a linear trend for each variable also has 21 estimated parameters. Thus, the table compares the root mean squared forecast errors from the hybrid model with those from the unconstrained VAR.

To create these statistics, both models are estimated with data from 1960:1 through 1984:4 and used to generate out-of-sample forecasts one through four quarters ahead. Next, the sample is extended to 1985:1, and additional forecasts are generated using the updated estimates. Continuing in this way yields series of one-quarter-ahead forecasts running from 1985:1 through 1997:3, series of two-quarters-ahead forecasts running from 1985:2 through 1997:3, series of three-quarters-ahead forecasts running from 1985:3 through 1997:3, and series of four-quarters-ahead forecasts running from 1985:4 through 1997:3, all of which can be compared to the actual data that were realized over those periods.

The results indicate that in nearly every case, forecasts from the hybrid model outperform those from the unconstrained VAR. To determine whether any of these differences are statistically significant, table 5 also reports a statistic that is used by Diebold and Mariano (1995) to test the null hypothesis of equal forecast accuracy across two models. Let $\{e_t^h\}_{t=1}^T$ denote a series of k-step-ahead forecast errors from the hybrid model, let $\{e_t^u\}_{t=1}^T$ denote the corresponding forecast errors from the unconstrained VAR, and construct a sequence $\{d_t\}_{t=1}^T$ of loss differentials using $d_t = (e_t^u)^2 - (e_t^h)^2$ for all $t = 1, 2, ..., T$. Diebold and Mariano show that the test statistic

$$S = d/\sigma_d$$

is asymptotically distributed as a standard normal random variable, where d is the sample mean of $\{d_t\}_{t=1}^T$ and where σ_d, the standard error of d, can be estimated using formulas given in their paper, under the null hypothesis of equal forecast accuracy: $d = 0$.

In table 5, positive values of S indicate cases where the hybrid model's forecasts outperform the VAR's, while negative values of S indicate cases where the opposite is true. In fact, tests of the null hypothesis $S = 0$ against the alternative $S > 0$ often reject the null of equal accuracy of the models' forecasts. In no case can the null $S = 0$ be rejected in favor of the alternative $S < 0$. Overall, therefore, the hybrid model delivers forecasts that are superior to those from the unconstrained VAR.

5. Conclusion

This paper adds to a prototypical dynamic, stochastic, general equilibrium model—Hansen's (1985) real business cycle model with indivisible labor—a vector of residuals that follows a first-order autoregressive process. The result is a hybrid model that exploits the power of detailed economic theory but remains flexible enough to be taken to the data: the model can be estimated via maximum likelihood and, once estimated, can be used to perform statistical hypothesis tests as well as to generate out-of-sample forecasts.

Some of the results presented above echo the well-known successes and shortcomings of the real business cycle model as documented by Watson (1993), Cooley and Prescott (1995), and Rotemberg and Woodford (1996), among others. The results show, for example, that technology shocks do a better job in explaining the behavior of output and its components than they do in explaining the behavior of aggregate hours worked. In addition, technology shocks account for much of the variability in output that occurs over very short and very long horizons, but are less successful in accounting for output variation at business cycle frequencies. And finally, estimates of the model reveal that the statistical uncertainty surrounding Kydland and Prescott's (1982) finding that the real business cycle model explains a substantial fraction of the output variation in the US data is large, though not as large as first suggested by Eichenbaum (1991).

Other results, however, illuminate aspects of the real business cycle model that are less widely appreciated. A statistical test performed above fails to reject the hypothesis that the real business cycle model's parameters have remained stable over time, despite important changes in monetary and fiscal policies. This result confirms one of the model's greatest strengths: it identifies parameters that are truly structural. Furthermore, the hybrid model developed here—which takes the real business cycle model as its starting point—delivers out-of-sample forecasts that outperform those from an unconstrained VAR.

As the surveys in Cooley's (1995) volume make clear, work on dynamic, stochastic, general equilibrium theory has now moved well beyond its real business cycle origins to consider the effects of monetary and fiscal policy shocks, household production, imperfectly competitive market structures, and numerous other extensions. The method developed here can be applied more generally to take these extended models to the data and to assess their explanatory power, both within-sample and out-of-sample. Doing so remains a task for future research.

6. References

Altug, Sumru. "Time-to-Build and Aggregate Fluctuations: Some New Evidence." *International Economic Review* 30 (November 1989): 889-920.

Andrews, Donald W.K. and Ray C. Fair. "Inference in Nonlinear Econometric Models with Structural Change." *Review of Economic Studies* 55 (October 1988): 615-639.

Bencivenga, Valerie R. "An Econometric Study of Hours and Output Variation with Preference Shocks." *International Economic Review* 33 (May 1992): 449-471.

Blanchard, Olivier Jean and Charles M. Kahn. "The Solution of Linear Difference Models Under Rational Expectations." *Econometrica* 48 (July 1980): 1305-1311.

Chow, Gregory C. and Yum K. Kwan. "How the Basic RBC Model Fails to Explain US Time Series." *Journal of Monetary Economics* 41 (April 1998): 301-318.

Christiano, Lawrence J. "Why Does Inventory Investment Fluctuate So Much?" *Journal of Monetary Economics* 21 (March/May 1988): 247-280.

Cooley, Thomas F., ed. *Frontiers of Business Cycle Research.* Princeton: Princeton University Press, 1995.

Cooley, Thomas F. and Edward C. Prescott. "Economic Growth and Business Cycles." In Thomas F. Cooley, ed. *Frontiers of Business Cycle Research.* Princeton: Princeton University Press, 1995.

DeJong, David N., Beth F. Ingram, and Charles H. Whiteman. "A Bayesian Approach to Dynamic Macroeconomics." Manuscript. Pittsburgh: University of Pittsburgh, Department of Economics, August 1997*a*.

DeJong, David N., Beth F. Ingram, and Charles H. Whiteman. "Keynesian Impulses Versus Solow Residuals: Identifying Sources of Business Cycle Fluctuations." Manuscript. Pittsburgh: University of Pittsburgh, Department of Economics, August 1997*b*.

Diebold, Francis X. and Roberto S. Mariano. "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics* 13 (July 1995): 253-263.

Eichenbaum, Martin. "Real Business-Cycle Theory: Wisdom or Whimsy?" *Journal of Economic Dynamics and Control* 15 (October 1991): 607-626.

Hall, George J. "Overtime, Effort, and the Propagation of Business Cycle Shocks." *Journal of Monetary Economics* 38 (August 1996): 139-160.

Hamilton, James D. *Time Series Analysis.* Princeton: Princeton University Press, 1994.

Hansen, Gary D. "Indivisible Labor and the Business Cycle." *Journal of Monetary Economics* 16 (November 1985): 309-327.

Ingram, Beth Fisher, Narayana R. Kocherlakota, and N.E. Savin. "Explaining Business Cycles: A Multiple-Shock Approach." *Journal of Monetary Economics* 34 (December 1994): 415-428.

Ireland, Peter N. "A Small, Structural, Quarterly Model for Monetary Policy Evaluation." *Carnegie-Rochester Conference Series on Public Policy* 47 (December 1997): 83-108.

Kim, Jinill. "Monetary Policy in a Stochastic Equilibrium Model with Real and Nominal Rigidities." Manuscript. New Haven: Yale University, Department of Economics, December 1995.

Kydland, Finn E. and Edward C. Prescott. "Time To Build and Aggregate Fluctuations." *Econometrica* 50 (November 1982): 1345-1370.

McGrattan, Ellen R. "The Macroeconomic Effects of Distortionary Taxation." *Journal of Monetary Economics* 33 (June 1994): 573-601.

McGrattan, Ellen R., Richard Rogerson, and Randall Wright. "An Equilibrium Model of the Business Cycle with Household Production and Fiscal Policy." *International Economic Review* 38 (May 1997): 267-290.

Perron, Pierre. "The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis." *Econometrica* 57 (November 1989): 1361-1401.

Rappoport, Peter and Lucrezia Reichlin. "Segmented Trends and Non-Stationary Time Series." *Economic Journal* 99 (Conference Supplement 1989): 168-177.

Rogerson, Richard. "Indivisible Labor, Lotteries and Equilibrium." *Journal of Monetary Economics* 21 (January 1988): 3-16.

Rotemberg, Julio J. and Michael Woodford. "Real-Business-Cycle Models and the Forecastable Movements in Output, Hours, and Consumption." *American Economic Review* 86 (March 1996): 71-89.

Runkle, David E. "Vector Autoregressions and Reality." *Journal of Business and Economic Statistics* 5 (October 1987): 437-442.

Sargent, Thomas J. "Two Models of Measurements and the Investment Accelerator." *Journal of Political Economy* 97 (April 1989): 251-287.

Sims, Christopher A. "Macroeconomics and Reality." *Econometrica* 48 (January 1980): 1-48.

Stock, James H. and Mark W. Watson. "Evidence on Structural Instability in Macroeconomic Time Series Relations." *Journal of Business and Economic Statistics* 14 (January 1996): 11-30.

Watson, Mark W. "Measures of Fit for Calibrated Models." *Journal of Political Economy* 101 (December 1993): 1011-1041.

Table 1. Full Sample Estimates and Standard Errors

Parameter	Estimate	Standard Error
γ	0.0045	0.0001
θ	0.2342	0.0046
η	1.0039	0.0005
A	6.0952	0.5927
ρ	0.9983	0.0024
σ	0.0050	0.0003
d_{yy}	1.1961	0.1042
d_{yc}	0.5693	0.1434
d_{yh}	-0.5656	0.1288
d_{cy}	0.0983	0.0620
d_{cc}	1.1497	0.0656
d_{ch}	-0.1977	0.0757
d_{hy}	0.4177	0.1035
d_{hc}	0.4773	0.1395
d_{hh}	0.3297	0.0964
v_y	0.0060	0.0009
v_c	0.0057	0.0006
v_h	0.0006	0.0016
v_{yc}	0.00001813	0.00000613
v_{yh}	0.00000160	0.00000288
v_{ch}	0.00000323	0.00000308

Table 2. Forecast Error Variance Decompositions

A. Output

Quarters Ahead	Percentage of Variance Due to Technology	Standard Error
1	63.7468	7.4107
4	36.2954	6.1128
8	27.8473	6.6408
12	28.3861	7.7473
20	34.3150	9.5320
40	45.4542	12.2212
∞	84.9800	19.8384

B. Consumption

Quarters Ahead	Percentage of Variance Due to Technology	Standard Error
1	34.1186	5.8732
4	28.2569	5.7496
8	27.4942	7.1697
12	30.3149	8.7008
20	38.4756	11.0472
40	52.9727	14.4447
∞	90.3517	14.2374

C. Investment

Quarters Ahead	Percent of Variance Due to Technology	Standard Error
1	42.7775	6.4653
4	29.4801	5.0084
8	22.1535	5.0172
12	21.8195	5.6030
20	23.9082	6.2863
40	27.0359	7.0614
∞	51.3788	30.2073

D. Hours Worked

Quarters Ahead	Percent of Variance Due to Technology	Standard Error
1	97.8137	11.7038
4	17.1219	4.0571
8	5.6879	1.4949
12	3.7631	1.1457
20	2.9524	1.0593
40	2.4723	1.1227
∞	2.2882	1.3138

Table 3. Subsample Estimates and Standard Errors

Parameter	Subsample 1 Estimate	Std Error	Subsample 2 Estimate	Std Error
γ	0.0046	0.0001	0.0043	0.0002
θ	0.2307	0.0048	0.2369	0.0048
η	1.0046	0.0005	1.0033	0.0004
A	6.3576	0.2582	6.4266	0.3356
ρ	0.9927	0.0100	0.9718	0.0412
σ	0.0040	0.0012	0.0033	0.0012
d_{yy}	1.1590	0.1615	0.8793	0.2499
d_{yc}	0.3328	0.1831	0.6900	0.2246
d_{yh}	-0.4850	0.1322	-0.3867	0.1772
d_{cy}	0.1265	0.0826	-0.0090	0.0987
d_{cc}	1.0413	0.0860	1.1849	0.1075
d_{ch}	-0.2224	0.0711	-0.1349	0.0882
d_{hy}	0.3732	0.1544	0.2228	0.1534
d_{hc}	0.2130	0.1690	0.3993	0.1185
d_{hh}	0.4863	0.1315	0.5499	0.1470
v_y	0.0073	0.0022	0.0070	0.0007
v_c	0.0062	0.0009	0.0061	0.0010
v_h	0.0035	0.0027	0.0023	0.0005
v_{yc}	0.00002664	0.00001400	0.00002858	0.00000990
v_{yh}	0.00000881	0.00001954	0.00001590	0.00000153
v_{ch}	0.00000461	0.00000934	0.00001093	0.00000478

Table 4. Tests for Parameter Stability

Stability of all 21 Estimated Parameters:	$W = 40.2731^{***}$
Stability of the 6 Structural Parameters:	$W = 9.0709$
Stability of the Remaining 15 Parameters:	$W = 27.6233^{**}$

Note: ** and *** denote significance at the 5% and 1% levels.

Table 5. Forecast Accuracy, 1985:1-1997:3

A. Output

Quarters Ahead	1	2	3	4
RMSE - Hybrid Model (%)	0.6373	1.1883	1.7046	2.1873
RMSE - VAR (%)	0.8488	1.4912	1.9987	2.3778
S	3.5418***	2.2876**	1.5419	0.7620

B. Consumption

Quarters Ahead	1	2	3	4
RMSE - Hybrid Model (%)	0.5743	0.8559	1.1190	1.4495
RMSE - VAR (%)	0.5895	0.8749	1.1138	1.3571
S	0.5257	0.3124	−0.0443	−0.5028

C. Investment

Quarters Ahead	1	2	3	4
RMSE - Hybrid Model (%)	2.7228	4.4673	5.9743	7.1188
RMSE - VAR (%)	3.6074	5.9378	7.6938	9.0156
S	3.3774***	2.2188**	1.7465*	1.5859

D. Hours Worked

Quarters Ahead	1	2	3	4
RMSE - Hybrid Model (%)	0.4192	0.9138	1.4840	2.1113
RMSE - VAR (%)	0.7323	1.4194	2.0750	2.6915
S	4.5319***	2.7056**	2.1005**	1.7485*

Note: *, **, and *** denote significance at the 10%, 5%, and 1% levels.