Federal Reserve Bank of Minneapolis
Research Department

Measurement with Minimal Theory*

Ellen R. McGrattan

Working Paper 643

Revised October 2009

ABSTRACT

Applied macroeconomists interested in identifying the sources of business cycle fluctuations typically have no more than 40 or 50 years of data at a quarterly frequency. With sample sizes that small, identification may not be possible even if the analyst has a correctly-specified representation of the data. In this paper, I investigate whether small samples are indeed a problem for some commonly used statistical representations. I compare three—a vector autoregressive moving average (VARMA), an unrestricted state space, and a restricted state space—that are all consistent with the same prototype business cycle model. The statistical representations that I consider differ in the amount of a priori theory that is imposed but are all correctly specified. I find that the identifying assumptions of VARMAs and unrestricted state space representations are too minimal: the range of estimates for statistics of interest for business cycle researchers are so large as to be uninformative.

*McGrattan, Federal Reserve Bank of Minneapolis and University of Minnesota. I thank Elmar Mertens, Ed Prescott, and Warren Weber for their comments. Codes to replicate the results of this paper are available at my website. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. Introduction

Applied macroeconomists interested in identifying the sources of business cycle fluctuations typically have no more than 40 or 50 years of data at a quarterly frequency. With sample sizes that small, identification may not be possible even if the analyst has a correctly-specified representation of the data. In this paper, I investigate whether small samples are indeed a problem for some commonly used statistical representations applied to the same prototype business cycle model. The business cycle model is a prototype in the sense that many models, with various frictions and shocks, are observationally equivalent to it.

The statistical representations that I consider differ in the amount of theoretical detail that is imposed a priori, but all are correctly specified. In other words, if we had a sample of infinite length, all representations would correctly identify the sources of business cycles and the contributions of different shocks to economic fluctuations. I compare three representations: (*i*) a vector autoregressive moving average (VARMA), (*ii*) an unrestricted state space, and (*iii*) a restricted state space. All are consistent with the same prototype business cycle model, but the VARMA imposes few restrictions based on the underlying economic environment and the restricted state space imposes many. In particular, the VARMA representation is a system of equations in reduced form while the restricted state space representation uses specific details about the incentives and tradeoffs faced by economic agents in the theory.

I find that the identifying assumptions of the VARMA and unrestricted state space representations are too minimal to uncover statistics of interest for business cycle research with sample sizes used in practice. I demonstrate this by simulating 1000 datasets of length 200 quarters using the prototype business cycle model. For each dataset and each of the three statistical representations of the data, I apply the method of maximum likelihood to estimate parameters for that representation and then construct statistics of interest to business cycle analysts. The statistics include impulse responses, variance decompositions, and second moments of filtered data. For the VARMA and unrestricted state space representations, I find that many of the predictions are biased and have large standard errors. The errors are so large as to be uninformative.

Since the restricted state space representation relies on specific details of the economic environment, the maximum likelihood parameters are economically interpretable and can be constrained to lie in economically plausible ranges. In practice, business cycle researchers may put further constraints on the ranges of these parameters using independent micro or macroevidence. I also do this and compare results across experiments, varying constraints on the possible ranges of the maximum likelihood parameters. Interestingly, I find that the main results—which are the statistics of interest for business cycle analysts—are not sensitive to varying the constraints if they are confined to the economically plausible range.

In a related study, Kascha and Mertens (2009) compare the small sample performances of the VARMA and unrestricted state space representations with that of a structural vector autoregression (SVAR). (See Box A for some background on SVARs.) They do not consider restricted state space representations which impose much more theory. They find that the VARMA performs about as well as SVARs, and the state space representation performs slightly better than the SVARs. However, none of the representations they consider yield precise estimates for the statistics that these authors highlight.

In Section 2, I lay out the prototype business cycle model. Section 3 summarizes the three statistical representations. The method of maximum likelihood used to estimate parameters of the three representations is described in Section 4. In Section 5, I report on the business cycle statistics computed for each representation. Section 6 concludes.

2. The Prototype Business Cycle Model

I use a prototype growth model as the data generating process for this study. The model is a prototype in the sense that a large class of models, including those with various types of frictions and various sources of shocks, are equivalent to a growth model with time-varying *wedges* that distort the equilibrium decisions of agents operating in otherwise competitive markets. (See Chari et al. 2007.) These wedges are modeled like time-varying productivity, labor income taxes, and investment taxes. Since many models map into the same configuration of wedges, identifying one particular configuration does not uniquely identify a model; rather it identifies a whole class of models. Thus, the results are not specific to any one detailed economy.

Households in the economy maximize expected utility over per capita consumption c_t and per capita labor l_t,

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{\left(c_t \left(1 - l_t\right)^{\psi} \right)^{1-\sigma} - 1}{(1-\sigma)} \right] N_t \tag{2.1}$$

subject to the budget constraint and the capital accumulation law,

$$c_t + (1 + \tau_{xt})\, x_t = (1 - \tau_{lt})\, w_t l_t + r_t k_t + T_t \tag{2.2}$$

$$(1 + g_n)\, k_{t+1} = (1 - \delta)\, k_t + x_t \tag{2.3}$$

where k_t denotes the per capita capital stock, x_t per capita investment, w_t the wage rate, r_t the rental rate on capital, β the discount factor, δ the depreciation rate of capital, N_t the population with growth rate equal to $1+g_n$, and T_t the per capita lump-sum transfers. The series τ_{lt} and τ_{xt} are stochastic and stand in for time-varying distortions that affect the households' intratemporal and intertemporal decisions. Chari et al. (2007) refer to τ_{lt} as the *labor wedge* and τ_{xt} as the *investment wedge*.

The firms' production function is $F(K_t, Z_t L_t)$ where K and L are aggregate capital and labor inputs and Z_t is a labor-augmenting technology parameter which is assumed to be stochastic. Chari et al. (2007) call Z_t the *efficiency wedge* and demonstrate an equivalence between the prototype model with time-varying efficiency wedges and several detailed economies with underlying frictions that cause factor inputs to be used inefficiently. Here, I assume that the process for $\log Z_t$ is a unit-root with innovation $\log z_t$. The process for the exogenous state vector $s_t = [\log z_t, \tau_{lt}, \tau_{xt}]'$ is[1]

$$
\begin{aligned}
s_t &= P_0 + P s_{t-1} + Q\varepsilon_t \qquad (2.4) \\
&= \begin{bmatrix} g_z \\ (1-\rho_l)\,\tau_l \\ (1-\rho_x)\,\tau_x \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & \rho_l & 0 \\ 0 & 0 & \rho_x \end{bmatrix} s_{t-1} + \begin{bmatrix} \sigma_z & 0 & 0 \\ 0 & \sigma_l & 0 \\ 0 & 0 & \sigma_x \end{bmatrix} \varepsilon_t.
\end{aligned}
$$

where $\varepsilon_t = [\varepsilon_{zt}, \varepsilon_{lt}, \varepsilon_{xt}]'$ is the vector of shocks hitting the economy at date t.

Approximate equilibrium decision functions can be computed by log-linearizing the first-order conditions and applying standard methods. (See, for example, Uhlig 1999.) The equilibrium decision function for capital has the form

$$
\begin{aligned}
\log \hat{k}_{t+1} &= \gamma_k \log \hat{k}_t + \gamma_z \log z_t + \gamma_l \tau_{lt} + \gamma_x \tau_{xt} + \gamma_0 \\
&\equiv \gamma_k \log \hat{k}_t + \gamma_s' s_t + \gamma_0 \qquad (2.5)
\end{aligned}
$$

where $\hat{k}_t = k_t / Z_{t-1}$ is detrended capital. From the static first-order conditions, I also derive decision functions for output, investment, and labor which I use later, namely,

$$\log \hat{y}_t = \phi_{yk} \log \hat{k}_t + \phi_{ys}' s_t \qquad (2.6)$$

$$\log \hat{x}_t = \phi_{xk} \log \hat{k}_t + \phi_{xs}' s_t \qquad (2.7)$$

$$\log l_t = \phi_{lk} \log \hat{k}_t + \phi_{ls}' s_t \qquad (2.8)$$

[1] The assumption that the shocks are orthogonal is unrealistic for many actual economies, but adding correlations makes it even more difficult for atheoretical approaches.

where $\hat{y}_t = y_t/Z_t$, $\hat{x}_t = x_t/Z_t$, and the coefficient vectors ϕ_{ys}, ϕ_{xs}, and ϕ_{ls} that multiply s_t in equations (2.6)-(2.8) are 3-dimensional. The coefficients in equations (2.5)-(2.8) are funcitons of the underlying parameters of preferences and technology that appear in the original household objective function (2.1) and constraints, (2.2)-(2.3).

2.1. Observables

In all representations later, I assume that the economic modeler has data on per capita output, labor, and investment. Because output and investment grow over time, the vector of observables is taken to be

$$Y_t = [\,\Delta \log y_t/l_t \quad \log l_t \quad \log x_t/y_t\,]' \,.$$

The elements of Y are: the growth rate of log labor productivity, the log of the labor input, and the log of the investment share.[2] All elements of Y are stationary.

For the prototype model, these observables can be written as functions of $S_t = [\log \hat{k}_t,$ $s_t,\ s_{t-1}, 1]'$. To see this, note that the change in log productivity is a function of the state today $(\log \hat{k}_t, s_t, 1)$ and the state yesterday $(\log \hat{k}_{t-1}, s_{t-1}, 1)$. The capital stock at the beginning of the last period $\log \hat{k}_{t-1}$ can be written in terms of $\log \hat{k}_t$ and s_{t-1} by (2.5). The other observables depend only on today's state $(\log \hat{k}_t, s_t, 1)$. Thus, all of the observables can be written as a function of S_t, which is a 8×1 vector.

3. Three Statistical Representations

I use the form of decision functions for the prototype model to motivate three different but related statistical representations of the economic time series.

[2] I have chosen variables that business cycle researchers typically do, but other variations that I tried—such as using output growth rather than labor productivity growth did not affect my results.

3.1. A Restricted State Space Representation

The state space representation for the prototype model has the form

$$\begin{aligned} S_{t+1} &= A(\Theta) S_t + B(\Theta) \varepsilon_{t+1}, \quad E\varepsilon_t \varepsilon_t' = I \\ Y_t &= C(\Theta) S_t \end{aligned} \tag{3.1}$$

where the parameter vector is

$$\Theta = [i, g_n, g_z, \delta, \theta, \psi, \sigma, \tau_l, \tau_x, \rho_l, \rho_x, \sigma_l, \sigma_x]'.$$

Here, i is the interest rate and is used to set the discount factor $\beta = \exp(g_z)^{\sigma}/(1+i)$. I use Θ to compute an equilibrium and then construct

$$A(\Theta) = \begin{bmatrix} \gamma_k & \gamma_s' & 0 & \gamma_0 \\ 0 & P & 0 & P_0 \\ 0 & I & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad B(\Theta) = \begin{bmatrix} 0 \\ Q \\ 0 \\ 0 \end{bmatrix} \tag{3.2}$$

$$C(\Theta) = \begin{bmatrix} (\phi_{yk} - \phi_{lk})(1 - 1/\gamma_k) & \phi_{lk} & \phi_{xk} - \phi_{yk} \\ \phi_{ys}' - \phi_{ls}' + \mathbf{1}' & \phi_{ls}' & \phi_{xs}' - \phi_{ys}' \\ -\phi_{ys}' + \phi_{ls}' + (\phi_{yk} - \phi_{lk}) \gamma_s'/\gamma_k & 0 & 0 \\ (\phi_{yk} - \phi_{lk}) \gamma_0/\gamma_k & \phi_{l0} & \phi_{x0} - \phi_{y0} \end{bmatrix}'$$

where $\mathbf{1}$ is a vector with 1 in the first element and zeros otherwise. Recall that P and Q are 3×3 matrices. Thus, $A(\Theta)$ is an 8×8 matrix, $B(\Theta)$ is a 6×3 matrix, and $C(\Theta)$ is a 3×6. Elements of these matrices are functions of coefficients in (2.4)-(2.8).

Estimates $\hat{\Theta}$ are found by applying the method of maximum likelihood. The exact likelihood function is computed using a Kalman filter algorithm. (See, for example, Hamilton 1994.)

For the restricted state space representation, I consider three sets of restrictions on the parameter space. In what I refer to as the "loose constraints" case, I assume that the parameters in Θ can take on any value as long as an equilibrium can be computed. In what I refer to as the “modest constraints” case, I assume that the parameters in Θ are constrained to be economically plausible. Finally, I consider a “tight constraints” case with some parameters fixed during estimation. The parameters that are fixed are those that are least controversial for business cycle theorists. They are the interest rate i, the growth rates g_n and g_z, the depreciation rate δ, the capital share θ, and the mean tax rates τ_l and τ_x. In the tight-constraints case, I only estimate the parameters affecting key elasticities, namely, ψ and σ, and parameters affecting the stochastic processes for the shocks. There is no consensus on the values for these parameters.

3.2. An Unrestricted State Space Representation

In the restricted state space representation, all cross-equations restrictions are imposed. This necessitates making many assumptions about the economic environment. Suppose instead that I assume only that the state of the economy evolves according to (2.4) and (2.5), and that decisions take the form of (2.6)-(2.8).

In this case, I need not provide specific details of preferences and technologies. I do, however, need to impose some minimal restrictions that imply the parameters of the state space are identified. Let $\bar{S}_t = [\log \bar{k}_t, \bar{s}_t, \bar{s}_{t-1}]'$ where

$$\log \bar{k}_t = \left(\log \hat{k}_t - \log \hat{k}\right) / (\gamma_z \sigma_z)$$

$$\log \bar{z}_t = (\log z_t - \log z) / \sigma_z$$

$$\bar{\tau}_{lt} = (\tau_{lt} - \tau_l) / \sigma_l$$

$$\bar{\tau}_{xt} = (\tau_{xt} - \tau_x) / \sigma_x$$

and $\bar{s}_t = [\log \bar{z}_t, \bar{\tau}_{lt}, \bar{\tau}_{xt}]$. Then the unrestricted state space representation can be written

$$\bar{S}_{t+1} = A_u(\Gamma) \bar{S}_t + B_u \varepsilon_{t+1}, \quad E\varepsilon_t \varepsilon_t' = I$$

$$Y_t = C_u(\Gamma) \bar{S}_t \tag{3.3}$$

with

$$A_u(\Gamma) = \begin{bmatrix} \gamma_k & 1 & \tilde{\gamma}_l & \tilde{\gamma}_x & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \rho_l & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \rho_x & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}, \quad B_u = \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \tag{3.4}$$

and $C_u(\Gamma)$ unrestricted (except for zero coefficients on $\bar{s}_{t-1}$ in the second and third rows). The (1,3) element of $A_u(\Gamma)$ is $\tilde{\gamma}_l = \gamma_l\sigma_l/(\gamma_z\sigma_z)$. The (1,4) element is $\tilde{\gamma}_x = \gamma_x\sigma_x/(\gamma_z\sigma_z)$.

The vector to be estimated, Γ, is therefore given by

$$\Gamma = \left[\gamma_k, \tilde{\gamma}_l, \tilde{\gamma}_x, \rho_l, \rho_x, \operatorname{vec}(C_u)'\right]$$

where the $\operatorname{vec}(C_u)'$ includes only the elements that are not a priori set to 0. As in the case of the restricted state space representation, estimates are found by applying the method of maximum likelihood. From this, I get $\hat{\Gamma}$.

Proposition 1. The state space representation (3.3) is identified.

Proof. Applying the results of Wall (1984),[3] if (A_u^1, B_u^1, C_u^1) and (A_u^2, B_u^2, C_u^2) are observationally equivalent state space representations, then they are related by $A_u^2 = T^{-1}A_u^1T$, $B_u^2 = T^{-1}B_u^1$, and $C_u^2 = C_u^1T$. Identification obtains if the only matrix T satisfying these equations is $T = I$. It is simple algebra to show that this is the case for the unrestricted state space representation (3.3). ∎

It is useful to compare the matrices for the restricted state space representation in (3.2) and the unrestricted state space representation in (3.4). All coefficients in (3.2) are functions of the business cycle model's "deep structural" parameters Θ and must satisfy the cross-equation restrictions imposed by the theory. On the other hand, the only structure

[3] See Burmeister, Wall, and Hamilton (1986), Proposition 2.

imposed on coefficients of the unrestricted state space in (3.4) are zero restrictions. I am not imposing anything more.

3.3. A Vector Autoregression Moving Average Representation

Starting from the state space representation (3.1), the moving average for the prototype model with observations Y is easily derived by recursive substitution. In particular, it is given by

$$Y_t = CB\varepsilon_t + CAB\varepsilon_{t-1} + CA^2B\varepsilon_{t-2} + \ldots. \tag{3.5}$$

Assume that CB is invertible and let $e_t = CB\varepsilon_t$. Then I can rewrite (3.5) as

$$\begin{aligned} Y_t &= e_t + CAB\,(CB)^{-1}\,e_{t-1} + CA^2B\,(CB)^{-1}\,e_{t-2} + \ldots \\ &\equiv e_t + C_1e_t + C_2e_{t-2} + \ldots. \end{aligned}$$

Assuming the moving average is invertible, Y can also be represented as an infinite-order VAR,

$$Y_t = B_1Y_{t-1} + B_2Y_{t-2} + \ldots + e_t \tag{3.6}$$

where $B_j = C_j - B_1C_{j-1} - \ldots B_{j-1}C_1$.

Proposition 2. For the prototype economy, the implied VAR in (3.6) has the property that $M = B_jB_{j-1}^{-1}$ and therefore can be represented as a vector autoregressive moving average representation of order (1,1), namely,

$$Y_t = (B_1 + M)\,Y_{t-1} + e_t - Me_{t-1}, \quad Ee_te_t' = \Sigma \tag{3.7}$$

with $\Sigma = CBB'C'$.

Proof. See Chari et al. (2008). ■

Let Λ denote the vector of parameters to be estimated for the VARMA via maximum likelihood, which are all of the elements of matrices B_1, M, and Σ. If I allow these paramters to take on any values, it is possible that the system would be nonstationary or noninvertible. I reparameterize the VARMA as described in Ansley and Kohn (1986) to ensure stationarity and invertibility I also need to check that B_1 has nonzero elements and that $[B_1 + M, M]$ has full rank to ensure that the matrices are statistically identifiable. (See Hannan 1976.)

I now have three statistical representations that are consistent with the prototype model: the restricted state space representation that makes explicit use of the details of the underlying model and imposes these in cross-equation restrictions, the unrestricted state space representation which imposes zero restrictions on the state space but no cross-equation restrictions, and the VARMA(1,1) representation which uses only minimal information about the reduced form of the system. For all three, it is straightforward to apply the method of maximum likelihood.[4]

4. Setting up the laboratory

Before applying the estimation procedure, I first generate 1000 samples of data $\{Y_t\}$ using the prototype business cycle model. Each sample is 200 quarters in length, which is typical for actual applications. This is done by randomly drawing sequences for the shocks $\{varepsilon_t\}$. These shocks, along with an intial value of the state s_0, imply a sequence of exogenous states $\{s_t\}$ that statisfy (2.4). With an inital capital stock $\hat{k}_0$ and the sequence $\{s_t\}$, I can use (2.6)-(2.8) to generate data for the business cycle model. For each sample, the true parameters of the business cycle model are fixed and given by

$$\Theta = [.01, .0025, .005, .015, .33, 1.8, 1.0, .25, .0, .95, .95, 1, 1, 1]' .$$

Using the restricted state space representation, I apply the method of maximum likelihood

[4] Codes are available at my website. See Anderson et al. (1996) for more details on estimating dynamic linear economies.

to each of the 1000 samples. This procedure yields 1000 estimates $\hat{\Theta}$ of the parameter vector. For each estimate, I can construct the coefficients of the model's equilibrium equations (2.4)-(2.8). With numerical values for these coefficients, I can then construct the statistics that business cycle analysts care about which will be discussed later.

Similarly, I can apply the method of maximum likelihood in the case of the other two statistical representations. For the unrestricted state space, the procedure yields estimates for $\hat{\Gamma}$ and, in turn, for $A_u(\hat{\Gamma})$ and $C_u(\hat{\Gamma})$ of (3.3). For the VARMA(1,1), the procedure yields estimates for $\hat{\Lambda}$ and, in turn, $\hat{B}_1$, $\hat{M}$, and $\hat{\Sigma}$ of (3.7). As before, once I have numerical values for the coefficients in these equations, I can construct the statistics of interest for business cycle analysts.

For the restricted state space representation, three levels of constraints on the parameter vector are investigated. Recall that the only restriction in the "loose constraints" case is that an equilibrium exists. In the "modest constraints" case, I assume that the parameter constraints are

$$\begin{aligned}&[.0075, 0, .0025, 0, .25, .01, .01, .15, -.1, -1, -1, 0, 0, 0]\\ &\qquad < \hat{\Theta} < [.0125, .0075, .0075, .025, .45, 10, 10, .35, .1, 1, 1, 10, 10, 10]\,. \quad (4.1)\end{aligned}$$

This implies an annual rate of interest between 3 and 5 percent; an annual growth rate of population between 0 and 3 percent; an annual growth rate of technology between 1 and 3 percent; an annual depreciation rate between 0 and 10 percent; a capital share between 25 and 45 percent; ψ and σ between 0.01 and 10; the mean labor wedge between 0.25 and 0.35; the mean investment wedge between -0.1 and 0.1; serial correlation coefficients between -1 and 1; and standard deviations of the shocks between 0 and 10 percent. In the "tight constraints" case, I fix the interest rate, the growth rates, the depreciation rate, the capital share, and the means of the tax rates during estimation and use bounds in (4.1) for the other parameters.

5. Business cycle statistics

Statistics of interest for business cycle analysts include impulse response functions, variance decompositions, autocorrelations, and cross-correlations. In this section, I use the three representations (3.1), (3.3), and (3.7) to construct these statistics.

The first set of statistics are impulse responses of the three observables—growth in labor productivity, the log of labor, and the log of the investment share—to 1 percent shocks in each of the three shocks in ε_t. Here, I report the responses of productivity, labor, and investment in the period of impact of the shock. In the restricted state space representation, the impact of the shock is summarized by the elements of CB. Similarly, the impact responses are summarized by $C_u B_u$ for the unrestricted state space representation. For the VARMA, one needs additional information to identify CB from the variance-covariance $\Sigma = (CB)(CB)'$. A typical assumption made in the literature to identify the first column of CB is to assume that demand shocks have no long run effect on labor productivity. This assumption allows me to infer the first column of CB. (See Chari et al. 2008.) However, it does not imply anything for the relative impacts of ε_{lt} and ε_{xt}. Since these are not identifiable, they are not reported.

The impact coefficients of the impulse responses are reported in Table 1. The first row shows the true value of each statistic. For example, in the model, productivity rises by 0.58 percent in response to a 1 percent increase in ε_{zt}, labor rises by 0.27 percent, and the investment share rises by 0.88 percent. Responses to shocks in ε_{lt} are shown in the middle three columns, and responses to shocks in ε_{xt} are shown in the last three columns.

In the next three rows, I report statistics based on the restricted state space representation with varying degrees of tightness in the constraints imposed during maximum likelihood estimation. The last two rows are the results for the unrestricted state space representation and the VARMA(1,1) representation. In each case, the first number displayed is the mean estimate of the statistic averaged over the 1000 datasets. The second number displayed below in parentheses is the root mean square error which is defined as

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(\hat{\zeta}_i - \zeta\right)^2}.$$

In this formula, $\hat{\zeta}_i$ is the ith estimate of the statistic, $i = 1, \ldots, N$, and ζ is the true value. If there is no bias due to small samples, then ζ is equal to the mean of the estimates $\hat{\zeta}_i$, $i = 1, \ldots, N$, and the RMSE is equal to the standard deviation.

It is clear from Table 1 that the differences in results for the restricted state space and the other two representations are large. Consider first the means of the estimates. There is little bias in the estimates for the restricted state space. This is especially true when tight constraints are used during maximum likelihood estimation. However, even in the case of modest constraints, the means of the estimates are very close to the true values shown in the first row. For the unrestricted state space representation and the VARMA, the biases are large. For example, all of the predicted responses following a technology shock are significantly below the actual responses. In the case of the shocks to the labor and investment wedges for the unrestricted state space model, large biases are also evident.

Next consider the root mean squared errors that appear in parentheses below the means. As I remove restrictions, these errors grow large. Compare, for example, the errors of the restricted state space representation with tight constraints with those of the VARMA in columns one through three. The errors are so large in the latter case that the impulse response predictions range from large negatives to large positives. In other words, the VARMA predictions are uninformative. Similarly, the unrestricted state space has large root mean squared errors for all of the statistics reported in Table 1 and, like the VARMA, is therefore uninformative about impulse responses.

To generate tight predictions, we need to impose the cross-equation restrictions as long as the parameter estimates are restricted to lie in the economically plausible range. When I allow all of the parameters to be completely free for the restricted state space representation, I find that the root mean squared errors do get significantly larger. For example, one can see a significant difference in the responses of labor and the investment share.

The next statistics that I consider are variance decompositions. For a general state space system of the form

$$S_{t+1} = AS_t + Be_{t+1} \tag{5.1}$$

$$Y_t = CS_t \tag{5.2}$$

with $Ee_te_t' = \Sigma$, the *population variances* of the observables in Y are the diagonal elements of the matrix V, where

$$V = AVA' + BLL'B'.$$

where L is a lower triangular matrix that satisfies $LL' = \Sigma$. In this case, the (i,i) element of V is the variance of the ith variable in Y. The *variance decomposition* summarizes the contribution of the variances due to each of the shocks in e_t. To be specific, let V_j be the contribution of the variance of Y due to shock j. This is given by

$$V_j = AV_jA' + BL\Phi_jL'B'$$

where Φ_j is a matrix with the same dimensions as Σ and one nonzero element, element (j,j), that is equal to 1. In this case, the (i,i) element of V_j is the variance of the ith variable in Y which is due to the jth shock. Note that $V = \sum_j V_j$.

In the case of the VARMA(1,1), I can rewrite the system in (3.7) in the form of the state space above, namely

$$S_{t+1} = \begin{bmatrix} B_1 + M & I \\ 0 & 0 \end{bmatrix} S_t + \begin{bmatrix} I \\ -M \end{bmatrix} e_{t+1}$$

$$Y_t = [I \quad 0]\, S_t$$

where the coefficients on S_t and e_{t+1} can be mapped to A, B, and C in (5.1)-(5.2).

In Table 2, I report the predictions of the population variance decompositions. The ordering of results in Table 2 is the same as in Table 1, with the most restrictive appearing first and the least restrictive appearing last. Comparing the means of the statistics with the actual values, we see that the biases are not as large for the variance decompositions as they were for the impact coefficients of the impulse responses, with some exceptions for the unrestricted state space. For example, the results show some bias for the decompositions of labor and investment shares. However, in terms of the root mean squared errors, the results for the unrestricted state space and VARMA representations again show that the predictions are not informative. In effect, the range of variances is close to everything in 0 to 100 percent.

The third set of statistics are very common in the real business cycle literature that typically reports statistics for filtered time series using the method of Hodrick and Prescott (1997). Specifically, for each statistical representation and each set of parameter estimates, I simulate 500 time series for output, labor, and investment of length 200. In each case, the output and investment data are filtered because they are nonstationary. I then take averages of standard deviations, autocorrelations, and cross-correlations over the 500 simulations. This is done for each representation and for each of the 1000 maximum likelihood parameter vectors.

The implied statistics are reported in Table 3. Notice that the bias and root mean squared errors of the predictions are small for all representations. For example, in all cases, the distribution of cross-correlations of output and labor has a mean of 0.89 and the largest root mean square error is 0.02. Perhaps this is not too surprising given that we do not need all of the details of a model to get an accurate prediction for unconditional moments.

The final set of statistics is related to those reported in Table 2. In Table 4, I report the variance decompositions for the filtered data. This is a similar exercise to that done in Table 2 but is included for easy comparison to estimates in the business cycle literature. As before, the range of predictions for the unrestricted state space and the VARMA representations is so large that they are uninformative. In the restricted state space model, the estimates for the technology shock are very informative. This is true even for labor and investment, whose variation depends little on technology shocks. The restricted state space estimates for the labor shock imply that it contributes significantly to all three variables. The restricted state space estimates for the investment shock are the least informative, but still imply that ε_x has a big effect on investment.

6. Conclusion

In this paper, I conduct a simple small-sample study. I ask how much can business cycle theorists learn from actual time series if they impose very little theory when applying their statistical methods. The answer is very little.

References

Anderson, Evan W., Lars P. Hansen, Ellen R. McGrattan, and Thomas J. Sargent. 1996. "Mechanics of Forming and Estimating Dynamic Linear Economies," in H. Amman, D. Kendrick, and J. Rust, eds., *Handbook of Computational Economics*, North-Holland.

Ansley, Craig F. and Robert Kohn. 1986. "A Note on Reparameterizing a Vector Autoregressive Moving Average Model to Enforce Stationarity," *Journal of Statistical Computational Simulations*, 24: 99–106.

Blanchard, Olivier J., and Danny Quah. 1989. "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review*, 79: 655–673.

Burmeister, Edwin, Kent D. Wall, and James D. Hamilton. 1986. "Estimation of Unobserved Expected Monthly Inflation Using Kalman Filtering," *Journal of Business and Economic Statistics*, 4(2): 147–160.

Chari V. V., Patrick J. Kehoe, Ellen R. McGrattan. 2007. "Business Cycle Accounting," *Econometrica*, 75(3): 781–836.

Chari V. V., Patrick J. Kehoe, Ellen R. McGrattan. 2008. "Are Structural VARs with Long-Run Restrictions Useful in Developing Business Cycle Theory?" *Journal of Monetary Economics*, 55(8): 1337–1352.

Galí, Jordi, and Pau Rabanal. 2005. "Technology Shocks and Aggregate Fluctuations: How Well Does the RBC Model Fit Postwar U.S. Data?" *NBER Macroeconomics Annual 2004*, 19: 225–288.

Hamilton, James D. 1994. *Time Series Analysis* (Princeton, N.J.: Princeton University Press).

Hannan, E. J. 1976. "The Identification and Parameterization of Armax and State Space Forms," *Econometrica*, 44(4): 713–723.

Hodrick, Edward C., and Robert J. Hodrick. 1997. "Postwar U.S. Business Cycles: An Empirical Investigation." *Journal of Money, Credit and Banking,* 29(1): 1–16.

Kascha, Christian, and Karel Mertens. 2009. "Business Cycle Analysis and VARMA Models," *Journal of Economic Dynamics and Control,* 33: 267–282.

Uhlig, Harald. 1999. "A Toolkit for Analysing Nonlinear Dynamic Stochastic Models

Easily," in R. Marimon and A. Scott, eds., *Computational Methods for the Study of Dynamic Economies* (Oxford, U.K.: Oxford University Press), 30–61.

Wall, Kent D. 1984. "Identification Theory for Varying Coefficient Regression Models," paper presented at the Sixth Conference of the Society of Economic Dynamics and Control.