



Technology shocks and the business cycle: An empirical investigation[☆]

Peter N. Ireland*

*Department of Economics, Boston College, Carney Hall, 140 Commonwealth Avenue, Chestnut Hill,
MA 02467-3806, USA*

Received 1 October 1998; accepted 1 April 1999

Abstract

This paper uses maximum likelihood to estimate a prototypical real business cycle model under several different assumptions regarding the stochastic process governing technological change. The estimation approach permits the use of formal statistical hypothesis tests to help discriminate between competing specifications; the approach also allows each specification to be evaluated based on its ability to generate accurate out-of-sample forecasts. Overall, the results indicate that the data prefer a version of the model in which technology shocks are extremely persistent but still trend stationary.

© 2001 Elsevier Science B.V. All rights reserved.

JEL classification: E13; E32; E37

Keywords: Technology shocks; Business cycles

1. Introduction

This paper estimates a prototypical real business cycle model under several different assumptions regarding the stochastic process governing technological change. Its goal is to ascertain which specification fits the data best.

[☆]I would like to thank two anonymous referees for helpful comments and suggestions.

* Tel: 1-617-552-3687; fax: 1-617-552-2308.

E-mail address: peter.ireland@bc.edu (P.N. Ireland).

In two widely-cited survey articles, King et al. (1988a, b) demonstrate that technology shocks must exhibit a considerable amount of serial correlation if the real business cycle model is to account for key features of the US data. However, they leave open the question of whether technology shocks ought to be modeled as following a highly persistent, but trend stationary, process or as a random walk with drift. Hansen (1997) proposes an answer to this question by solving and simulating versions of his (1985) model of indivisible labor with both trend and difference stationary technology shocks. After evaluating each version based on its ability to match the relative volatilities of output, consumption, investment, and hours worked as measured in the data, Hansen concludes that the trend stationary specification is to be preferred.

As suggested above, this paper picks up where Hansen (1997) leaves off by actually estimating several versions of the real business cycle model via maximum likelihood. This alternative to the more conventional, moment-matching exercises performed by Hansen (1985, 1997) and many others, including Kydland and Prescott (1982), Prescott (1986); King et al. (1988a, b), offers several potential advantages.

First, these earlier, moment-matching exercises attempt to minimize the difference between the variances of two series: one measuring output, consumption, investment, or hours worked in the model and the other measuring the same variable in the data. The alternative, estimation exercises performed here, in contrast, attempt to minimize the variance of the difference between the same two series. As Smith and Zin (1997, p. 244) put it, ‘having a zero variance of the difference between two series is a much stricter criterion than having a zero difference of their variances’. In other words, the estimation exercises performed here demand far more from the real business cycle model than standard, moment-matching exercises; they may, therefore, provide additional information about the model’s ability to explain the data.

Second, some of the results derived here echo those obtained by King et al. (1988a); Hansen (1997) by showing that certain trade-offs emerge as the real business cycle model’s technology shock becomes more persistent: as the specification varies, the model’s performance improves along some dimensions but deteriorates along others. The estimation exercises performed here, like the moment-matching exercises performed in earlier studies, take these trade-offs into account. The estimation exercises, however, go a step further by providing a more rigorous and systematic way of choosing between alternative specifications: as soon as the model is estimated, formal statistical hypothesis tests may be used to discriminate between competing versions.

Third, and finally, the real business cycle model, once estimated, can be used to generate forecasts. Thus, unlike more conventional moment-matching exercises, the estimation exercises performed here allow different versions of the model to be evaluated based on their ability to track the data out of sample as well as within sample. As noted by Hamilton (1994, Chapter 15, Section 3),

forecasts generated by trend stationary models differ greatly from those generated by difference stationary models. The comparison of forecast accuracy allowed by the estimation approach taken here may therefore provide new and valuable information about which version of the model fits the data best.

In fact, these potential advantages are all realized here. The results presented below show that within sample, a formal hypothesis test rejects a popular version of the real business cycle model, in which technology shocks follow a trend stationary, first-order autoregressive process with persistence parameter $\rho = 0.95$, in favor of an alternative specification in which ρ lies much closer to one. Based only on within-sample performance, this alternative, trend stationary specification with ρ very close to one is difficult to distinguish from a third version of the model with difference stationary technology shocks. However, the trend stationary specification generates out-of-sample forecasts that are considerably more accurate than those from the difference stationary model. Thus, the results presented below provide sharp and decisive evidence in support of a version of the real business cycle model in which technology shocks are extremely persistent but still trend stationary.

2. A prototypical real business cycle model

The prototypical real business cycle model considered here is one with indivisible labor, first developed by Hansen (1985) and later used by Hansen (1997) in work that, as noted above, serves as the immediate predecessor to the present study. This popular model also ranks among those discussed in the famous survey articles by Prescott (1986); King et al. (1988a).

In the model, a representative consumer has preferences defined over consumption C_t and hours worked H_t during each period $t = 0, 1, 2, \dots$, as described by the expected utility function

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

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 is unemployed during any given period. If these individual consumers have utility functions that are additively separable in consumption and leisure, and if these individual consumers are able to fully insure themselves against the possibility of being unemployed, then their behavior in the aggregate can be summarized by considering the activities of a representative consumer with the utility function shown in (1).

During period t , 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

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

where $1 > \theta > 0$. The variable A_t in (2) represents a random shock to technology; its properties are described below.

The consumer divides output Y_t between consumption C_t and investment I_t , subject to the resource constraint

$$Y_t = C_t + I_t. \quad (3)$$

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

$$K_{t+1} = (1 - \delta)K_t + I_t, \quad (4)$$

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

Two general specifications for the law of motion governing the technology shock A_t are considered here. Under the first specification, deviations of the natural logarithm of A_t from a linear, deterministic trend follow a first-order autoregressive process, as in King et al. (1988a):

$$\ln(A_{t+1}) - (t + 1)\ln(\eta) = (1 - \rho)\ln(A) + \rho[\ln(A_t) - t\ln(\eta)] + \varepsilon_{t+1}, \quad (5)$$

where $A > 0$ and $1 > \rho > -1$. Under the second specification, the log of A_t follows a random walk with drift, as in King et al. (1988b):

$$\ln(A_{t+1}) = \ln(\eta) + \ln(A_t) + \varepsilon_{t+1}. \quad (6)$$

In both the cases, $\eta > 1$ measures the average rate of labor-augmenting technological progress and the serially uncorrelated innovation ε_{t+1} is normally distributed with mean zero and standard deviation $\sigma > 0$. In fact, (6) might be viewed as a special case of (5), in which the persistence parameter ρ is set equal to one. This apparent similarity, however, masks the key feature that distinguishes the two specifications: under (5), the technology shock is trend stationary, whereas under (6), the technology shock is difference stationary.

Equilibrium allocations for both versions of the model may be characterized by solving the representative consumer's problem: maximize, by choice of Y_t , C_t , I_t , H_t , and K_{t+1} for all $t = 0, 1, 2, \dots$ the utility function (1) subject to the constraints (2)–(4) and either (5) or (6) for all $t = 0, 1, 2, \dots$. This problem lacks a closed-form solution, but an approximate solution may be found, following King et al. (1988a, b), by defining a set of stationary, trendless variables

$$\begin{aligned} y_t &= Y_t/\eta^t, & c_t &= C_t/\eta^t, & i_t &= I_t/\eta^t, & h_t &= H_t, \\ k_t &= K_t/\eta^t, & \text{and } a_t &= A_t/\eta^t \end{aligned} \quad (7)$$

for the model with trend stationary technology shocks and another set of stationary, trendless variables

$$\begin{aligned} s_t^y &= Y_{t-1}/A_{t-1}, & g_t^y &= Y_t/Y_{t-1}, & r_t^c &= C_t/Y_t, & r_t^i &= I_t/Y_t, \\ h_t &= H_t, & s_t^k &= K_t/A_{t-1}, & \text{and} & s_t^a &= A_t/A_{t-1} \end{aligned} \tag{8}$$

for the model with difference stationary technology shocks. Eqs. (7) and (8) reveal that the exact form of the transformations needed to make each of the model's variables stationary depends on the process governing technological change. The form of these stationarity-inducing transformations also depends on the choice of utility and production functions in (1) and (2); in fact, King et al. (1988a) demonstrate that such stationarity-inducing transformations exist only when the utility and production functions are drawn from special classes that include the choices in (1) and (2).

The first-order conditions for the representative consumer's problem imply that in the absence of technology shocks, when $\varepsilon_t = 0$ for all $t = 0, 1, 2, \dots$, the economy converges to a steady state in which each of the stationary variables is constant, with

$$y_t = y, \quad c_t = c, \quad i_t = i, \quad h_t = h, \quad k_t = k \quad \text{and} \quad a_t = a$$

for all $t = 0, 1, 2, \dots$ in the model with trend stationary technology shocks and

$$\begin{aligned} s_t^y &= s^y, & g_t^y &= g^y, & r_t^c &= r^c, & r_t^i &= r^i, & h_t &= h, \\ s_t^k &= s^k, & \text{and} & & s_t^a &= s^a \end{aligned}$$

for all $t = 0, 1, 2, \dots$ in the model with difference stationary technology shocks. The consumer's first-order conditions can be log-linearized about this steady state; this log-linear approximation to the model has a solution of the form

$$s_{t+1} = \mathbf{A}s_t + \mathbf{B}\varepsilon_{t+1} \tag{9}$$

and

$$f_t = \mathbf{C}s_t \tag{10}$$

for all $t = 0, 1, 2, \dots$.

In (9) and (10), the vectors s_t and f_t keep track of percentage deviations of each stationary variable from its steady-state level, with

$$s_t = [\ln(k_t/k)\ln(a_t/a)]'$$

and

$$f_t = [\ln(y_t/y)\ln(c_t/c) \ln(i_t/i) \ln(h_t/h)]' \tag{11}$$

for the model with trend stationary technology shocks and

$$s_t = [\ln(s_t^k/s^k)\ln(s_t^y/s^y) \ln(s_t^a/s^a)]'$$

and

$$\mathbf{f}_t = [\ln(g_t^y/g^y)\ln(r_t^c/r^c) \ln(r_t^i/r^i) \ln(h_t/h)]' \quad (12)$$

for the model with difference stationary technology shocks. In both cases, the matrices \mathbf{A} , \mathbf{B} , and \mathbf{C} have elements that depend on the parameters describing tastes and technologies: β , γ , θ , δ , η , A , ρ , and σ for the model with trend stationary technology shocks and β , γ , θ , δ , η , and σ for the model with difference stationary technology shocks.

3. Estimation strategy

Equations (9) and (10) are in state-space form: for both versions of the model, the components of \mathbf{s}_t can be treated as unobservable state variables, while the components of \mathbf{f}_t can be measured in the data. Hence, methods described by Hamilton (1994, Chapter 13) may be applied to obtain maximum likelihood estimates of the real business cycle model's parameters under each assumption regarding the stationarity properties of the technology shocks.

As emphasized by Ingram et al. (1994), however, the assumption that just one shock — the technology shock — drives all aggregate fluctuations makes the real business cycle model singular: the model predicts that certain linear combinations of the observable variables in \mathbf{f}_t will be deterministic. If, in the data, these exact linear relationships do not hold, then any attempt to estimate the model via maximum likelihood will fail.

To overcome this obstacle, consider augmenting each equation in (10) with a serially correlated residual, so that the empirical model consists of (9),

$$\mathbf{f}_t = \mathbf{C}\mathbf{s}_t + \mathbf{u}_t \quad (13)$$

and

$$\mathbf{u}_{t+1} = \mathbf{D}\mathbf{u}_t + \boldsymbol{\xi}_{t+1} \quad (14)$$

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

This strategy of adding serially correlated residuals to the observation equation shown earlier as (10) is also used by Altug (1989); McGrattan (1994); Hall (1996); McGrattan et al. (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 and by assuming that these measurement errors are uncorrelated across variables, so that the matrices \mathbf{D} and \mathbf{V} are diagonal. Here, no such restrictions are imposed; instead, the residuals in \mathbf{u}_t are allowed to follow an unconstrained, first-order vector

autoregression. Thus, the residuals may still pick up measurement errors in each variable, but they might 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.

Of course, the assumption that ξ_{t+1} and ε_{t+1} are orthogonal is a strong one. It may not be satisfied if, for instance, the residuals in u_t pick up the effects of monetary and fiscal policy variables, which are not explicitly considered in the basic real business cycle model but which may be correlated with the real business cycle model's technology shock. With the addition of these residuals, however, maximum likelihood estimation no longer requires the technology shock to account for every aspect of the data. The hope, as always, is that the extra identifying assumption will hold as an approximation, so that the estimation approach will provide useful and informative results.

In the data, consumption C_t is measured by real personal consumption expenditures in chained 1992 dollars, investment I_t is measured by real gross private domestic investment, also in chained 1992 dollars, and output is measured by the sum $Y_t = C_t + I_t$. Hours worked H_t is measured by hours of wage and salary workers on private, nonfarm payrolls. Each series is expressed in per-capita terms using the civilian, noninstitutional population, age 16 and over; each series, except for population, is seasonally adjusted. The data for consumption, investment, and population come from the Federal Reserve Bank of St. Louis' FRED database; the data for hours worked come from the Bureau of Labor Statistics' Establishment Survey. All series are quarterly and run from 1960:1 through 1998:1.

It would, in principle at least, be feasible to bring additional data to bear in assessing the real business cycle model's empirical performance. The model, for instance, also makes predictions about the behavior of real wages and interest rates over the business cycle. Results obtained by Mankiw et al. (1985), however, suggest that representative agent models (at least those with time-separable utility functions, like the one considered here) have great difficulty in explaining the cyclical relationships between aggregate prices (such as wages and interest rates) and quantities (such as consumption and hours worked) that can be found in the data. For this reason, perhaps, most studies that work with real business cycle models, including those by Kydland and Prescott (1982); Hansen (1985, 1997); Prescott (1986); King et al. (1988a, b), focus mainly on the behavior of quantities; the present study follows in this tradition.

McGrattan (1994); Hall (1996); McGrattan et al. (1997) exploit data on the capital stock, as well as data on investment, to estimate versions of the real business cycle model. However, published data on the capital stock are constructed from data on investment using capital accumulation equations such as (4); see, for example, Katz and Herman (1997). Thus, as noted by McGrattan (1994), the use of data on both investment and capital to estimate a model that includes the capital accumulation equation (4) serves only to uncover the value

of the depreciation parameter δ that was used to construct the capital stock series in the first place; the very small standard errors that Hall (1996) and McGrattan et al. (1997) attach to their estimates of δ undoubtedly reflect this fact. Here, therefore, a simpler and more direct approach is taken: data on capital are not considered explicitly, but the depreciation rate δ is fixed prior to estimation at 0.025, the value originally used by Hansen (1985) and a value that approximates those used in the construction of the published capital stock series (again, see Katz and Herman, 1997).

Finally, less aggregated measures of consumption and investment might be used in place of the broad measures described above. In fact, preliminary attempts to estimate the model considered here used data on the consumption of nondurables and services alone, since this measure seems to correspond more closely to the variable C_t that appears in the model. These preliminary attempts at estimation failed, however, due to the fact that when consumption is measured using nondurables and services alone, trends appear in the consumption–output and investment–output ratios, at least over the sample period considered here. Since the theoretical model cannot account of the presence of trends in these ratios, the maximum likelihood procedure attempted to attribute the trends to nonstationary behavior in the residuals \mathbf{u}_t . Thus, to facilitate estimation and, in particular, to avoid the appearance of explosive roots in the matrix \mathbf{D} , broader measures of consumption and investment are used instead.

The definitions of the variables in (7) and (11) indicate that deviations of the log levels of output, consumption, and investment from a common deterministic trend are stationary under the assumption that technology shocks are governed by (5); the definitions of the variables in (8) and (12) indicate that the logs of output growth, the consumption–output ratio, and the investment–output ratio are stationary under the assumption that technology shocks are governed by (6). Thus, as noted above, each version of the model imposes its own stationarity-inducing transformation on the data; the series need not be filtered in any other way.

Moreover, since the resource constraint (3) holds in both the model and the data, only the series for Y_t , C_t , and H_t need to be used in estimating the model; the series for I_t is redundant. For the purposes of estimation, therefore, \mathbf{f}_t reduces to a 3×1 vector, with

$$\mathbf{f}_t = [\ln(y_t/y) \ln(c_t/c) \ln(h_t/h)]'$$

for the model with trend stationary technology shocks and

$$\mathbf{f}_t = [\ln(g_t^y/g^y) \ln(r_t^c/r^c) \ln(h_t/h)]'$$

for the model with difference stationary technology shocks. In both cases, \mathbf{u}_t and ξ_{t+1} become 3×1 vectors that can be written as

$$\mathbf{u}_t = [u_{yt} \ u_{ct} \ u_{ht}]'$$

and

$$\xi_{t+1} = [\xi_{yt+1} \ \xi_{ct+1} \ \xi_{ht+1}]',$$

while \mathbf{D} and \mathbf{V} become 3×3 matrices that 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 each version of the model, the parameters describing tastes and technologies are constrained to satisfy the theoretical restrictions noted in Section 2, above. In addition, the eigenvalues of \mathbf{D} are constrained to lie inside the unit circle, making the residuals in \mathbf{u}_t stationary. Finally, the covariance matrix \mathbf{V} is constrained to be positive definite.

4. Results

Table 1 presents maximum likelihood estimates of the parameters of the augmented model consisting of (9), (13), and (14) under several different assumptions regarding the process governing the behavior of the technology shocks. In each case, preliminary attempts to estimate the model led to unreasonably low values for the discount factor β ; here, as in Altug (1989), more sensible results obtain when this parameter is held fixed during estimation. Accordingly, the table displays results when β is set equal to 0.99 and when, as noted above, δ is set equal to 0.025. The standard errors, also shown 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 second column of Table 1 shows the results when the technology shocks are assumed to follow the trend stationary process (5); the third column shows the results when the technology shocks are assumed to follow the difference stationary process (6). The maximum likelihood estimates of γ , θ , η , and σ are quite similar across the two specifications; in each case, they are reasonable and precise.

The estimate of $\rho = 0.9983$ in the trend stationary case is very close to the setting $\rho = 1$ that is implicitly assumed under the difference stationary

Table 1
Parameter estimates and standard errors

Parameter	Trend stationary model	Difference stationary model	Trend stationary model with $\rho = 0.95$
	Estimate (std error)	Estimate (std error)	Estimate (std error)
γ	0.0044 (0.0001)	0.0044 (0.0002)	0.0044 (0.0001)
θ	0.2352 (0.0055)	0.2361 (0.0056)	0.2363 (0.0068)
η	1.0039 (0.0005)	1.0035 (0.0013)	1.0038 (0.0002)
A	10.5610 (0.6973)	—	10.9678 (0.3997)
ρ	0.9983 (0.0024)	—	—
σ	0.0064 (0.0004)	0.0063 (0.0007)	0.0046 (0.0003)
d_{yy}	1.1923 (0.1065)	0.6744 (0.0794)	1.1075 (0.0876)
d_{yc}	0.5648 (0.1535)	0.1377 (0.0496)	0.3275 (0.1427)
d_{yh}	-0.5500 (0.1268)	-0.0141 (0.0159)	-0.3063 (0.1084)
d_{cy}	0.0956 (0.0637)	-0.4137 (0.0635)	0.0475 (0.0568)
d_{cc}	1.1463 (0.0670)	0.8657 (0.0373)	1.0904 (0.0610)
d_{ch}	-0.1893 (0.0747)	-0.0190 (0.0105)	-0.1100 (0.0568)
d_{hy}	0.4120 (0.1057)	0.7934 (0.1021)	0.3207 (0.1011)
d_{hc}	0.4839 (0.1507)	-0.0595 (0.0375)	0.2778 (0.1477)
d_{hh}	0.3346 (0.0940)	1.0205 (0.0139)	0.5665 (0.0735)
v_y	0.0063 (0.0008)	0.0062 (0.0009)	0.0080 (0.0008)
v_c	0.0057 (0.0006)	0.0052 (0.0004)	0.0066 (0.0005)
v_h	0.0006 (0.0014)	0.0024 (0.0016)	0.0018 (0.0009)
v_{yc}	0.00001987 (0.00000617)	-0.00001268 (0.00000691)	0.00003749 (0.00000687)
v_{yh}	0.00000250 (0.00000274)	0.00000934 (0.00000616)	0.00001121 (0.00000435)
v_{ch}	0.00000327 (0.00000311)	-0.00000176 (0.00000513)	0.00001145 (0.00000416)
L^*	1704.9130	1706.2601	1684.3039

Note: L^* denotes the maximized value of the log likelihood function.

specification, indicating that the data prefer versions of the model that feature extremely persistent technology shocks. This finding is reinforced by the results displayed in the fourth column of Table 1, which shows maximum likelihood estimates of the trend stationary specification when ρ is constrained to equal 0.95, the value recommended by Hansen (1985, 1997). Under the null hypothesis that $\rho = 0.95$, the likelihood ratio statistic

$$LR = 2(L^u - L^c)$$

is asymptotically distributed as a chi-square random variable with one degree of freedom, where L^u is the maximized value of the unconstrained log likelihood function and L^c is the maximized value of the constrained log likelihood function. Table 1 indicates that $L^u = 1704.9130$ and $L^c = 1684.3039$, implying that $LR = 41.2182$. Since the 0.1 percent critical value for LR is 10.8, the null hypothesis is easily rejected by the data.

Some insight into the rejection of the hypothesis that $\rho = 0.95$ can be found in Table 2. There, the augmented model consisting of (9), (13), and (14) is used, together with the maximum likelihood estimates shown in Table 1, to decompose the k -step-ahead forecast error variances of the log levels of output, consumption, investment, and hours worked into two orthogonal components: the first attributable to the real business cycle model's technology shock and the second attributable to the residuals in u_t . Each figure in Table 2 measures the percentage of the forecast error variance that is due to the technology shock. Thus, like the R^2 statistic from a regression model, these figures measure the extent to which the behavior of each variable is explained by the real business cycle model rather than by the serially correlated residuals. Like the R^2 from a regression, these figures must be used with some caution in choosing between models: if in truth, technology shocks play absolutely no role in generating actual business cycle fluctuations in the US economy, then the true model ought to assign zero weight to technology shocks in explaining the data. But with this caveat in mind, the figures in Table 2 can be used to get an intuitive feel for how well each version of the real business cycle model does in accounting for observed movements in the data.

In fact, the results shown in Table 2 echo those presented by King et al. (1988a); Hansen (1997), who find that a trade-off emerges as technology shocks in the real business cycle model become more persistent: as ρ increases, technology shocks account for a larger fraction of the volatility in output and consumption and a smaller fraction of the volatility in investment and hours worked. Hansen (1997) considers this trade-off by comparing the relative volatilities of output, consumption, investment, and hours worked as predicted by the model to the same statistics as they appear in the data and concludes that the version of the model with $\rho = 0.95$ works best. The figures in Table 2, however, suggest that the models with $\rho = 0.9983$ and $\rho = 1$ do much better than the model with $\rho = 0.95$ in accounting for the behavior of output and consumption and, at the same time, fare only slightly worse in accounting for the behavior of investment and hours worked. On balance, therefore, the hypothesis test, which examines this trade-off in a formal way, favors the model with more persistent technology shocks.

It is difficult to distinguish between the trend stationary model with $\rho = 0.9983$ and the difference stationary model based on any of the figures shown in Tables 1 and 2. However, the augmented model consisting of (9), (13), and (14) can also be used to generate forecasts; the estimation approach used here allows the two specifications to be evaluated and compared based on their out-of-sample performance as well as their within-sample performance.

Thus, Table 3 reports the root mean squared forecast errors from the trend and difference stationary models. To produce the figures in this table, both models are estimated with data from 1960:1 through 1984:4 and used to generate out-of-sample forecasts for the log levels of output, consumption, investment,

Table 2
 Percentage of forecast error variance explained by technology

Quarters ahead	Trend stationary model	Difference stationary model	Trend stationary model with $\rho = 0.95$
<i>A. Output</i>			
1	60.9427	60.1338	47.8317
2	48.6111	44.7982	39.7814
3	40.0576	36.9341	33.6542
4	34.4170	32.9874	29.0584
8	26.1010	31.3967	19.5286
12	26.1539	35.0690	16.2368
16	28.4028	38.5311	15.0314
20	31.0773	41.6112	14.5562
40	41.1620	48.6653	14.0063
<i>B. Consumption</i>			
1	33.2848	28.6219	10.6072
2	30.6460	26.4294	10.6110
3	28.7690	25.1757	10.6211
4	27.5244	24.7238	10.6498
8	26.6037	27.3889	10.9989
12	28.9950	32.3640	11.6177
16	32.5376	37.2127	12.3051
20	36.2201	41.4784	12.9171
40	49.7109	50.0345	14.0387
<i>C. Investment</i>			
1	41.3051	43.7179	51.6493
2	36.7016	34.7512	45.5309
3	31.7833	29.2363	39.8427
4	27.7310	26.2138	34.9469
8	20.3320	24.8931	23.1121
12	19.5338	27.4457	18.6439
16	20.2759	29.1517	17.0578
20	21.1227	30.4416	16.4679
40	23.7625	33.8904	16.0061
<i>D. Hours worked</i>			
1	97.7231	68.3838	90.2711
2	51.4909	30.2780	63.5337
3	26.7930	16.4517	41.7110
4	16.1011	10.8196	27.9847
8	5.3144	5.0394	9.0844
12	3.4434	3.9656	5.0603
16	2.8541	3.5179	3.7582
20	2.5952	3.2060	3.2390
40	2.0964	2.3401	2.7614

and hours worked, with one-step-ahead forecasts being for 1985:1, two-step-ahead forecasts being for 1985:2, and so on. Next, the models are reestimated with data from 1960:1 through 1985:1 and used to generate one-step-ahead forecasts for 1985:2, two-step-ahead forecasts for 1985:3, and so on. Continuing in this manner yields series of one-step-ahead forecasts running from 1985:1 through 1998:1, series of two-step-ahead forecasts running from 1985:2 through 1998:1, and so on, all of which can be compared to the actual data that were realized over those periods; the exercise is designed to assess each model's real-time forecasting ability.

Table 3 shows that the model with trend stationary technology shocks produces more accurate forecasts of output and its components while the model with difference stationary technology shocks produces more accurate forecasts of hours worked. To determine whether any of these differences are significant, the table also reports a statistic that Diebold and Mariano (1995) use to test the null hypothesis of equal forecast accuracy across two models. Let $\{d_t\}_{t=1}^T$ denote a series of loss differentials, constructed as $d_t = (e_t^n)^2 - (e_t^s)^2$ for all $t = 0, 1, 2, \dots$, where $\{e_t^n\}_{t=1}^T$ is a series of k -step-ahead forecast errors from the model with difference stationary technology shocks and $\{e_t^s\}_{t=1}^T$ is a series of k -step-ahead forecast errors from the model with trend stationary technology shocks. Let d denote the sample mean of $\{d_t\}_{t=1}^T$. Diebold and Mariano show that under the null hypothesis of equal forecast accuracy, $d = 0$, the test statistic

$$S = d/\sigma_d$$

is asymptotically distributed as a standard normal random variable, where σ_d , the standard error of d , can be estimated using formulas reported in their paper.

In Table 3, positive values of S indicate cases where $d > 0$, so that the model with trend stationary technology shocks produces more accurate forecasts; negative values of S indicate cases where $d < 0$, so that the model with difference stationary technology shocks produces more accurate forecasts. In many cases, the null hypothesis that $d = 0$ can be rejected in favor of the alternative that $d > 0$. In no case, however, can the null hypothesis that $d = 0$ be rejected in favor of the alternative that $d < 0$. Based on these results, one can say that the model with stationary technology shocks clearly outperforms the model with difference stationary technology shocks in terms of its ability to generate accurate out-of-sample forecasts.

Before closing, it is useful to note that the state-space model consisting of (9), (13), and (14) can be used, together with methods outlined by Hamilton (1994, Chapter 13, Section 6), to produce 'smoothed' estimates of the variable A_t that appears in (2) and measures technological change in the US economy. These methods use data from the entire sample period, 1960:1–1998:1, to estimate A_t at each point in time. Since the procedure relies on all of the model's equations, including (2)–(4), (5) or (6), and the first-order conditions from the representative

Table 3
Forecast accuracy, 1985:1–1998:1

<i>A. Output</i>				
Quarters ahead	1	2	3	4
RMSE — Trend stationary model (%)	0.7064	1.2815	1.8584	2.3799
RMSE — Difference stationary model (%)	0.7957	1.5315	2.2726	2.9268
S	3.0956***	2.6180***	2.2093**	1.8770*
Quarters ahead	8	12	16	20
RMSE — Trend stationary model (%)	3.8887	4.3281	4.0854	4.2786
RMSE — Difference stationary model (%)	4.5993	5.1995	5.3350	5.8871
S	1.1455	1.0149	0.9627	0.9621
<i>B. Consumption</i>				
Quarters ahead	1	2	3	4
RMSE — Trend stationary model (%)	0.5836	0.8769	1.1633	1.4877
RMSE — Difference stationary model (%)	0.6470	1.0655	1.4963	1.9206
S	3.5149***	2.9601***	2.4737**	2.0742**
Quarters ahead	8	12	16	20
RMSE — Trend stationary model (%)	2.3895	2.6854	2.4478	2.4068
RMSE — Difference stationary model (%)	3.1118	3.6914	3.8467	4.2401
S	1.4810	1.3144	1.0870	1.0018
<i>C. Investment</i>				
Quarters ahead	1	2	3	4
RMSE — trend stationary model (%)	2.7377	4.5257	6.2331	7.6935
RMSE — Difference stationary model (%)	2.8647	4.9628	6.9526	8.8007
S	1.4204	1.7484*	1.6818*	1.7749*
Quarters ahead	8	12	16	20
RMSE — trend stationary model (%)	11.8548	12.8781	12.7833	13.9378
RMSE — Difference stationary model (%)	12.9259	13.9251	14.2611	15.3499
S	1.0138	1.0097	0.9847	0.8569
<i>D. Hours worked</i>				
Quarters ahead	1	2	3	4
RMSE — trend stationary model (%)	0.4111	0.9433	1.5568	2.2065
RMSE — Difference stationary model (%)	0.3936	0.9042	1.5028	2.1344
S	-0.6723	-0.5801	-0.5553	-0.5743
Quarters ahead	8	12	16	20
RMSE — trend stationary model (%)	4.8112	6.5748	7.2683	7.4894
RMSE — Difference stationary model (%)	4.3288	5.6329	6.3095	6.6080
S	-1.5623	-1.4982	-1.2213	-1.0288

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

consumer's problem, these estimates may be referred to as 'full-model' estimates of A_t .

The full-model estimates of A_t generated from the specification with $\rho = 0.9983$ are shown in Fig. 1; these estimates look very similar to those




Fig. 1. Alternative estimates of technological change.

generated from the alternative specifications with $\rho = 0.95$ and $\rho = 1$. Hansen's (1997) simulations also reveal that the variable A_t behaves quite similarly across specifications. Thus, both here and in Hansen (1997), the various specifications appear to differ more in their implications for the way in which the endogenous variables respond to shocks than in their implications for behavior of the shocks themselves.

Fig. 1 compares the full-model estimates described above to the estimates of A_t that are constructed in a more familiar way. Prescott (1986) follows Solow (1957) in measuring the productivity variable by taking logs on both sides of (2) and rearranging to obtain

$$\ln(A_t) = \left(\frac{1}{1 - \theta} \right) \ln(Y_t) - \left(\frac{\theta}{1 - \theta} \right) \ln(K_t) - \ln(H_t).$$

To form these 'Solow-Prescott' estimates of A_t , data on investment are used to construct a series for the capital stock from (4) with $\delta = 0.025$. The parameter θ is set equal to 0.2352, the estimated value from the trend stationary specification with $\rho = 0.9983$. In contrast to the full-model estimates, the Solow-Prescott estimates are based only on (2) and (4).

The two measures of technological change display similar movements in Fig. 1. The Solow-Prescott estimates are more volatile, however. This explains

why the estimates of σ shown in Table 1 are smaller than the values of σ selected by Prescott (1986); Hansen (1985, 1997), and others, who calibrate this parameter based on the Solow–Prescott measure of A_t .

5. Conclusion

As noted above, in the introduction, the estimation exercises performed here demand far more from the real business cycle model than the conventional, moment-matching exercises used throughout the existing literature. These exercises also allow for a more formal assessment of the trade-offs that arise when, for example, an increase in the persistence of technology shocks helps the model to explain the behavior of output and consumption even as it hurts the model's ability to explain the behavior of investment and hours worked. Here, a formal statistical hypothesis test weighs both sides of this trade-off and decisively rejects a version of the model with $\rho = 0.95$ in favor of an alternative with ρ much closer to one.

Once estimated, the real business cycle model can also be used for forecasting. Thus, the exercises performed here allow competing specifications to be evaluated based on their ability to track the data, not just within sample, but out of sample as well. Here, in fact, the extra information made available by a comparison of forecast accuracy proves especially valuable in discriminating between two versions of the model. Within sample, the model with trend stationary technology shocks and ρ very close to one is difficult to distinguish from the model with difference stationary technology shocks; any approach to model evaluation that relies solely on within-sample performance will be of little help in choosing between these two. Out of sample, however, the trend stationary model generates forecasts that are considerably more accurate than those from the difference stationary model; the comparison of forecast accuracy makes it easy to choose in favor of the trend stationary model.

Overall, therefore, the results obtained here indicate that the data prefer a version of the real business cycle model in which technology shocks are extremely persistent but still trend stationary. More generally, the analysis shows how maximum likelihood estimation, when applied to analyze dynamic, stochastic, general equilibrium models of aggregate fluctuations, can yield sharp and meaningful results.

References

- Altug, S., 1989. Time-to-build and aggregate fluctuations: some new evidence. *International Economic Review* 30, 889–920.
- Diebold, F.X., Mariano, R.S., 1995. Comparing predictive accuracy. *Journal of Business and Economic Statistics* 13, 253–263.

- Hall, G.J., 1996. Overtime, effort, and the propagation of business cycle shocks. *Journal of Monetary Economics* 38, 139–160.
- Hamilton, J.D., 1994. *Time Series Analysis*. Princeton University Press, Princeton.
- Hansen, G.D., 1985. Indivisible labor and the business cycle. *Journal of Monetary Economics* 16, 309–327.
- Hansen, G.D., 1997. Technical progress and aggregate fluctuations. *Journal of Economic Dynamics and Control* 21, 1005–1023.
- Ingram, B.F., Kocherlakota, N.R., Savin, N.E., 1994. Explaining business cycles: a multiple-shock approach. *Journal of Monetary Economics* 34, 415–428.
- Katz, A.J., Herman, S.W., 1997. Improved estimates of fixed reproducible tangible wealth, 1929–95. *Survey of Current Business* 77, 69–92.
- King, R.G., Plosser, C.I., Rebelo, S.T., 1988a. Production, growth and business cycles: I. the basic neoclassical model. *Journal of Monetary Economics* 21, 195–232.
- King, R.G., Plosser, C.I., Rebelo, S.T., 1988b. Production, growth and business cycles: II. new directions. *Journal of Monetary Economics* 21, 309–341.
- Kydland, F.E., Prescott, E.C., 1982. Time to build and aggregate fluctuations. *Econometrica* 50, 1345–1370.
- Mankiw, N.G., Rotemberg, J.J., Summers, L.H., 1985. Intertemporal substitution in macroeconomics. *Quarterly Journal of Economics* 100, 225–251.
- McGrattan, E.R., 1994. The macroeconomic effects of distortionary taxation. *Journal of Monetary Economics* 33, 573–601.
- McGrattan, E.R., Rogerson, R., Wright, R., 1997. An equilibrium model of the business cycle with household production and fiscal policy. *International Economic Review* 38, 267–290.
- Prescott, E.C., 1986. *Theory Ahead of Business-Cycle Measurement*. Carnegie-Rochester Conference Series on Public Policy, vol. 25, pp. 11–44.
- Rogerson, R., 1988. Indivisible labor, lotteries and equilibrium. *Journal of Monetary Economics* 21, 3–16.
- Sargent, T.J., 1989. Two models of measurements and the investment accelerator. *Journal of Political Economy* 97, 251–287.
- Smith, G.W., Zin, S.E., 1997. *Real Business-Cycle Realizations*. Carnegie-Rochester Conference Series on Public Policy, vol. 47, pp. 243–280.
- Solow, R.M., 1957. Technical change and the aggregate production function. *Review of Economics and Statistics* 39, 312–320.