



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Review of Economic Dynamics 8 (2005) 805–828

Review of
**Economic
Dynamics**

www.elsevier.com/locate/red

Fluctuations in convex models of endogenous growth, II: Business cycle properties

Larry E. Jones^{a,b,*}, Rodolfo E. Manuelli^c, Henry E. Siu^d

^a *University of Minnesota*

^b *Federal Reserve Bank of Minneapolis*

^c *University of Wisconsin*

^d *University of British Columbia*

Received 10 December 2003

Available online 15 August 2005

Abstract

Using ideas from the endogenous growth literature, we present a model of the endogenous determination of productivity growth based on individual worker decisions about human capital investment. We calibrate a version of the model to match long run growth facts from the US and study the business cycle properties of this model. This approach offers improvements along several dimensions over standard exogenous growth methodologies. Most importantly, our stochastic endogenous growth model generates much greater serial correlation in output growth and labor supply volatility relative to its real business cycle counterpart. We conclude that using the extra discipline of reproducing the trend productivity growth features of the data endogenously constitutes an important missing component from the real business cycle approach.

© 2005 Elsevier Inc. All rights reserved.

JEL classification: E00; E01; E32

Keywords: Business cycle fluctuations; Human capital investment; Productivity growth

* Corresponding author.

E-mail address: lej@econ.umn.edu (L.E. Jones).

1. Introduction

The effects of aggregate shocks on economic performance is a topic that has been studied intensively in the real business cycle literature. Even though the real business cycle (RBC) program has been quite successful in accounting for the cyclical properties of post-war aggregate data (see Cooley, 1995, for an excellent review), it has some shortcomings. Those that we find particularly important are:

- (1) the “trend” in output and its components—and, hence, the methods used to remove the trend—have been taken as exogenous, and independent of the sources of fluctuations;
- (2) the models predict substantially lower variability of labor supply than that observed in the data unless utility is linear in leisure (or, equivalently, there is an indivisibility in labor supply and lotteries are introduced);
- (3) the models must resort to unobservable, or at least difficult to measure, costs of adjustment, margins of variation, or asymmetries to mimic the persistence of growth rates (see Cogley and Nason, 1995, for a discussion); and
- (4) since the properties of the models’ time series are not sensitive to the specification of the degree of intertemporal substitution, it is difficult to estimate this parameter in a way that is consistent with the theory, though this parameter is critical in understanding the economy’s response to exogenous shocks.

Our purpose in this paper is to present a class of convex endogenous growth models, and to analyze their performance in terms of both growth and business cycle criteria. The models we study have close analogs in the RBC literature, and hence are a natural first step when moving beyond the standard RBC model. In fact, we interpret the exogenous growth rate of productivity as an endogenous growth rate of human capital. This perspective allows us to compare the strengths of both classes of models using a large number of moments of the joint distribution of macroeconomic time series.

We find that this class of models offers improvements over simple RBC models on several dimensions. In particular:

- The endogenous growth model shows far more labor supply variability than the standard exogenous growth, RBC model. In fact, the endogenous growth version can match the observed patterns in the data for very reasonable parameter estimates. This finding highlights a key difference between the two classes of models. In the endogenous growth models that we study, human capital services and hours are jointly supplied to the market. Thus, cyclical fluctuations in labor supply are amplified by cyclical changes in the demand for human capital services.
- The endogenous growth model outperforms the RBC model in its predictions for the serial correlation properties of growth rates of both output and labor productivity. Hence, the models that we study embody important internal propagation mechanisms. This, as well as the rich dynamics displayed by the model, are due to the fact that both physical and human capital investments are generated by productive activity.

In addition, we find that the degree of intertemporal substitution and the relative depreciation rates of human and physical capital are two key parameters determining the overall performance of the endogenous growth models. Small differences in the intertemporal elasticity of substitution induce relatively large changes in the predicted variability of the physical capital investment–output ratio and hours worked. For our specifications, we find that the “best” fit is obtained for elasticities of substitution lower than one (the logarithmic case). In contrast, the degree of intertemporal substitution has a small effect in the exogenous growth model. Similarly, when the depreciation rate on human capital is close to that on physical capital, the improvements offered by the endogenous growth models are smallest.

The channel that is the source of this improved performance is the endogeneity of labor productivity growth through explicit human capital investments. Because human capital depreciates more slowly than physical capital, the response to a shock is drawn out over time with investments first made in plant and equipment and only subsequently in worker training. This is due to the relatively larger increase in the marginal product of capital goods whose depreciation rates are higher in periods when positive shocks occur.¹ We examine the realism of this channel by comparing the model implications for the human capital time series with that estimated by Ho and Jorgenson (2001) for the US. We find that, although this match is imperfect, both series show a lagged response to the cycle, with the contemporaneous cross correlation between output growth and human capital growth being small in both the model and the US data.

Throughout the paper, we consider only the simplest versions of both endogenous and exogenous growth models. In order to generate labor supply volatility close to that of the US data, the standard framework has been generalized to include indivisible labor (see Hansen, 1985); home production (Benhabib et al., 1991; and Greenwood and Hercowitz, 1991); cyclical factor utilization (Burnside and Eichenbaum, 1996; and King and Rebelo, 1999); and a separate, unshocked, sector producing human capital (Einarsson and Marquis, 1998, 1999). Variations of the basic setup designed to produce autocorrelated output growth include labor market search (Merz, 1995; and Andolfatto, 1996); cyclical capital utilization (Burnside and Eichenbaum, 1996); costs of adjustment (Cogley and Nason, 1995); extreme “time-to-build” restrictions (Christiano and Todd, 1996); learning-by-doing and externalities in aggregate labor productivity (Collard, 1999); and different technologies used to produce physical goods and human capital, as well the incidence of shocks across sectors (Perli and Sakellaris, 1998).

All of the models that have attempted to improve upon the standard RBC model on these two dimensions have resorted to various asymmetries in production functions that are difficult to document. These include different production functions in the home, human capital, and physical goods production sectors; in particular, they make strong assumptions about the capital–labor ratio and differential elasticities of substitution across sectors that are not backed by evidence. In addition, for the models to produce the desired results it is necessary to assume a particular pattern of incidence for technology shocks: in most

¹ Roughly speaking, since the $sF_{k_i} - \delta_{k_i}$ are approximately equalized across capital goods, k_i , a given change in the shock, s , has a larger effect on sF_{k_i} for those i for which δ_{k_i} is largest. Hence, these are the capital goods which are ‘adjusted’ most quickly.

models, the human capital (or the home production) sector is not subject to any shocks, since this facilitates substitution in and out of market work, increasing the volatility of measured hours. Finally, several of the models resort to difficult to measure adjustment costs. Our model contributes to this literature by showing that realistic values of labor supply volatility, autocorrelation in output growth, and a number of other second moments, can be attained without resorting to costs of adjustment, or asymmetries in production technologies and the incidence of shocks. The only asymmetry we consider is a difference in depreciation rates between human and physical capital.

In Section 2, we begin by laying out a general formulation of the class of models that we are interested in studying. In Section 3, we present the results of our benchmark parameterization of the endogenous growth model and compare this to its simple RBC counterpart. In Section 4, we discuss in detail the human capital investment channel that drives the differences between the endogenous and exogenous growth variants. In Section 5 we provide some sensitivity analysis of our main findings. Finally some concluding comments are offered in Section 6.

2. A model of business cycles with endogenous trend productivity

The model that we study features investment in both human and physical capital and a time stationary technology that is subjected to random shocks. It is a stochastic version of the convex models described in Jones and Manuelli (1990) and thus, the rate of productivity growth is endogenously determined. For simplicity, we treat the accumulation of human and physical capital symmetrically. The models that we study can be summarized as the solution to the following planner's problem:

$$\max E_t \left\{ \sum_t \beta^t u(c_t, \ell_t) \right\}, \quad (1)$$

subject to,

$$c_t + x_{ht} + x_{kt} \leq F(k_t, z_t, s_t),$$

$$z_t = n_t h_t,$$

$$k_{t+1} \leq (1 - \delta_k)k_t + x_{kt},$$

$$h_{t+1} \leq (1 - \delta_h)h_t + x_{ht},$$

$$\ell_t + n_t \leq 1,$$

with h_0 and k_0 given.

Here $\{s_t\}$ is a stochastic process which we assume is Markov with a time stationary transition probability function, c_t is consumption, x_{kt} is investment in physical capital, k_t is the stock of physical capital, x_{ht} is investment in human capital, h_t is stock of human capital, z_t is "effective labor," n_t is hours spent working, and ℓ_t is leisure. The depreciation rates on physical and human capital are given by δ_k and δ_h , respectively. If $F(\cdot)$ is concave and bounded below by a homogeneous of degree one function, it is possible to show that the competitive equilibrium allocation coincides with the solution to the planner's problem and, for some parameter values, displays income (and consumption) growth.

Thus, this is a fairly standard endogenous growth model in which effective labor is made up of a combination of hours and human capital supplied to the market. For this choice of functional forms, the model corresponds to the two capital goods version discussed in Jones et al. (1993). Other endogenous growth models are similar (e.g., Lucas, 1988), with the differences being slightly different choices of functional forms. From a formal point of view, our choice of a linear law of motion for both capital stocks amounts to an aggregation assumption: the technology used to produce investment in human capital goods (education, training, and health care, among others) is identical to the technology used to produce general output.² This allows us to focus on the properties introduced by endogenizing trend productivity growth, and the fact that $\delta_h < \delta_k$. Alternatives would be to assume that human capital investments are produced in a separate sector with perhaps a higher share going to labor, or assuming that $x_{k,t}$ and $c_t + x_{h,t}$ are complements in output. The first would be similar to modeling human capital accumulation as formal education with the extra share to labor arising because of the additional input into the process of student's time. The second corresponds to learning by doing at the individual level since the stock of human capital would then automatically increase as long as the household is working. These alternatives, while interesting, do not allow for aggregation to a one sector model.

The solution of models in this class does cause some technical problems. The natural choice of the state is the vector (k_t, h_t, s_t) . The problem that this poses is that both k_t and h_t are diverging to infinity (at least for versions that exhibit growth on average). To solve this problem, the key property that we exploit is that for models of this type to have a balanced growth path, both preferences and technology must be restricted in a specific way (see King et al., 1988; and Alvarez and Stokey, 1998). For our numerical strategy, it suffices that the model satisfies the following condition.

Condition 1 (*Preferences and Technology*).

(a) The instantaneous utility function satisfies:

$$u(c, \ell) = \begin{cases} v(\ell)c^{1-\sigma}/(1-\sigma) & \text{for } \sigma > 0 \text{ and } \sigma \neq 1, \\ \log(c) + v(\ell) & \text{for } \sigma = 1. \end{cases}$$

(b) F is concave and homogeneous of degree one in (k, z) .

These two restrictions, in turn, imply that knowledge of the current shock and the current *ratio* of human capital to physical capital (the two relevant pseudo-state variables) is sufficient to determine the optimal choices of employment, the growth rate of human capital, and next period's human to physical capital ratio. Given the state, the current stocks, and the current level of employment it is possible to determine consumption and future capital stocks using static first order conditions. This is the strategy that we will use. A solution to the model is a time stationary set of equations for the appropriately detrended, endogenous variables as functions of the time stationary state vector $(s, k/h)$. Thus, labor supply

² This is obviously an extreme assumption. However, we were unable to obtain reliable estimates of the physical capital–labor or physical capital–human capital ratios in specific activities like education and health.

is given by a function $n^* : S \times \mathfrak{R}_+ \rightarrow [0, 1]$ with $n_t = n^*(s_t, k_t/h_t)$, the new endogenous state is given by a function $\kappa^* : S \times \mathfrak{R}_+ \rightarrow [0, 1]$ with $k_{t+1}/h_{t+1} = \kappa^*(s_t, k_t/h_t)$, etc.

3. Implementing the model numerically

In this section, we study the properties of a calibrated version of the model outlined above. We also analyze a “related” exogenous growth model similar to those studied in the RBC literature. Throughout the analysis, we deviate from standard practice in the RBC literature and focus on annual frequencies. We do this because the principal mechanism that differs between this model and the RBC model is human capital investment (education, job training and retraining, etc.). Because of this, we want to be able to compare the predictions of the model directly with those for series of human capital that are available in the literature. Since those estimates are only available at annual frequencies, this is what we focus on in the paper.

3.1. Calibrating the model

For the purpose of model calibration and quantitative analysis, we have constructed annual US time series data for the period 1959–2000. Our data differ from those taken directly from government sources in order to align the empirical time series with their theoretical counterparts. Among others, differences stem from the decomposition of government spending into consumption and investment components, as well as the treatment of consumer durables. Details are contained in Appendix A.

For functional forms, we adopt the restrictions on preferences outlined above with $v(\ell) = \ell^{\psi(1-\sigma)}$ and assume that the production function is given by $F(k, nh, s) = sAk^\alpha(nh)^{1-\alpha}$. As can be seen, this is equivalent to $F(k, h, n, s) = sAk^\alpha h^{1-\alpha} n^{1-\alpha}$. Thus, there are increasing returns in the basic factors, (k, h, n) . Because of this, some care must be taken in its decentralization in a competitive setting. One interpretation, implicit above is that firms view k and $z = nh$ as the relevant factor inputs, confining the non-convexities to the household side of the model. Other methods are also possible. Looked at from the perspective of a planner’s problem then, we see that the total share of all capital is $\alpha + (1 - \alpha) = 1$. One of the defining differences, from a mathematical perspective, of this class of models is that the share of capital, interpreted in a broad sense, is large relative to the usual RBC model.

We set the share parameter for physical capital, α , equal to 0.36, hold β fixed at 0.95, and set δ_k to match the measured physical capital investment share of output (23.8%) found in the data (for the baseline calibration, $\delta_k = 0.102$).

There is considerable disagreement in the labor literature about the depreciation rate of human capital, δ_h . We choose a value of $\delta_h = 0.025$ which was obtained as an estimate of depreciation rates from examining life-cycle properties of wage data. Wages just before retirement age fall at roughly 1.5% per year which, assuming that human capital investment is close to zero near the end of the work life, gives us a lower bound on δ_h among those working. Adding to this the approximate loss to the economy of the human capital of retirees fixes the depreciation rate in the aggregate stock at approximately 2.2 to 3% per

year; we use 2.5% as an intermediate value. Haley (1976) estimates δ_h to be somewhere between 1 and 4%. Heckman (1976) obtains point estimates that are similar to Haley's, but are not significantly different from zero. Earlier work by Ben-Porath (1967) estimates δ_h close to 9%. The results in Jorgenson and Fraumeni (1989) are consistent with depreciation rates that range between 1 and 3%. Thus, the value that we use is similar to those obtained by other methodologies. In order to cover the range of estimates, we experiment with the value of δ_h in our section on sensitivity (see Section 5).

Finally, we choose the remaining parameters to match the average growth rate of US output over the 1959–2000 period of 1.77% per year, and average labor supply equal to 0.17 (see Jones et al., 1993). These two facts ($\gamma_y = 1.77\%$ and $n_{BGP} = 0.17$) pin down two of the three remaining parameters of the model: σ , ψ , and A . This leaves one degree of freedom in the choice of these parameters. One of the findings (discussed in detail in Section 5) is that overall variability in the endogenous growth model is extremely sensitive to the curvature parameter, σ . As a result, we choose σ to match the standard deviation of output growth ($\sigma(\gamma_y)$ below) found in the data. For the baseline calibration this requires $\sigma = 1.40$. Sensitivity analysis on this parameter is presented in Section 5.

Our model puts restrictions on concepts that—although clearly identifiable from a theoretical point of view—are difficult to measure. Prime examples are consumption and investment in human capital. In the data, private expenditures on schooling and health (arguably investments in human capital) are assumed to be part of private consumption, and some forms of investment (e.g., training) are likely to remain unmeasured. In the model, those expenditures are more properly viewed as being investments in human capital. The resolution of this problem is not easy. As a first approximation, we assume that our empirical measure of consumption corresponds to the sum of consumption and investment in human capital in the model, $c + x_h$ (see Appendix A). Thus, it is the variable that enters the utility function, along with the level of investment (or spending) in human capital that coincides with measured consumption. Given our choice of parameters, it follows that on a non-stochastic balanced growth path, x_h/y is 26.2%.³ Lowering the calibrated value of δ_h lowers the balanced growth path value of x_h/y . It turns out that doing so improves our model's ability to generate labor supply volatility and autocorrelated output growth. This is discussed in Section 5.

Parameterizing the shock process is also not simple for this model since the stock of human capital is not directly measured. As a result, the standard Solow residual:

$$\begin{aligned} \log SR_t &\equiv \log y_t - \alpha \log k_t - (1 - \alpha) \log n_t \\ &= \log A + \log s_t + (1 - \alpha) \log h_t, \end{aligned} \quad (2)$$

is an agglomeration of the true shock and the stock of human capital. However, it is possible to identify the process $\{s_t\}$ by using the model's intertemporal first order conditions,

³ This is close to (though slightly larger than) the fraction we estimate in the data, properly accounting for both private and public expenditures on health care, training, and education (including expenditures such as books, computers, child care, students time, etc.). In earlier versions of the paper, we experimented with the possibility that some human capital investment (e.g., on-the-job training on the part of workers) goes unmeasured in the NIPA. This requires some adjustments to the way GDP accounting should be done. For small degrees of unmeasured investment, this made little difference in the results.

in addition to the accounting equation, (2). As an example, if $\delta_k = \delta_h$ optimality requires that the ratio of capital stocks, h_t/k_t , is identically $(1 - \alpha)/\alpha$. Hence:

$$y_t = s_t A \left[\frac{1 - \alpha}{\alpha} \right]^{1-\alpha} k_t n_t^{1-\alpha}.$$

Thus, given data on output, the capital stock, and hours, the time series of s_t can be directly identified up to the constant $A[(1 - \alpha)/\alpha]^{1-\alpha}$. In our calibration, $\delta_k > \delta_h$, so that the identification procedure is somewhat more involved, though conceptually the same as for the case with $\delta_k = \delta_h$. The details are contained in Appendix A.

We assume that the process, s_t , is given by:

$$s_t = \exp \left[\zeta_t - \frac{\sigma_\varepsilon^2}{2(1 - \rho^2)} \right],$$

with

$$\zeta_{t+1} = \rho \zeta_t + \varepsilon_{t+1},$$

where the ε 's are i.i.d., normal with mean zero and variance σ_ε^2 . It follows that $E(s_t) = 1$. Using the implied time series constructed from the US data, we estimate ρ and σ_ε to be 0.967 and 0.0135, respectively.⁴

The model that we study has a related exogenous growth version. More precisely, if h_t —our human capital variable—is assumed to grow exogenously at the rate \hat{y} , and further it is assumed that this growth in h has no resource cost, the technology becomes $F(k, nh, s; t) = s A_t k^\alpha n^{1-\alpha}$, with $A_t = A h_t^{1-\alpha} = A h_0^{1-\alpha} \hat{y}^{t(1-\alpha)}$. This is the RBC model that we study below. We calibrate the model in the same manner described above, with the exception of the curvature parameter, σ . It was not possible to find a value of σ to match $\sigma(\gamma_y)$, as was done in the endogenous growth model (this would require a value of σ significantly less than 1). Instead, we focus on the case of logarithmic preferences, $\sigma = 1$, considered widely in the RBC literature. The results presented are not sensitive to this assumption however, as shown in Section 5. Given this specification, s_t is calculated from Solow residual series implied by the US data. In this case, the estimated parameters are $\rho = 0.984$ and $\sigma_\varepsilon = 0.0092$.⁵

To solve both the endogenous and exogenous growth models, we use the method of parameterized value function iteration as discussed in Siu (1998). The method finds a linear combination of Chebychev polynomials to approximate the value function in the recursive representation of the model. This is done by iterating upon the contraction mapping over a discretization of the state space until the fixed point is found. See Appendix A for details.

⁴ Given these estimates, it is not obvious that the implied $\{s_t\}$ series is a realization of a stationary process. Since our non-stochastic model has the balanced growth property, our assumptions force the shock to explain the productivity slowdown that started in the mid-seventies. This implies that, in our sample, the estimated shocks are decreasing from a peak of 1.10 (10% above average) in 1965. Though feasible, modeling the productivity slowdown is beyond the scope of this paper. For alternative explanations, see Greenwood and Yorukoglu (1997), Caselli (1999) and Manuelli (2000).

⁵ The difference in the estimated standard deviation here relative to the endogenous growth specification is due to the fact that the latter allows the growth rate to covary with the shock, while the former fixes it at 1.77%.

We used simulated time series with $T = 10,000$ periods in order to calculate estimates of key population moments of the models' endogenous variables. To facilitate comparisons throughout the paper, the same realization of $\{s_t\}$ was used for all cases with the same parameters for the stochastic process.

3.2. Results

As indicated above, our purpose is to assess the performance of the endogenous growth model in terms of its ability to replicate the distribution of observed variables. We also study the differences in amplification and propagation mechanisms between the endogenous and the exogenous growth models. We concentrate on three dimensions of that distribution: standard deviations, autocorrelations and cross-correlations. In terms of first moments, we find that the quantitative difference between the mean values of simulated data and their calibrated, non-stochastic balanced growth path values to be very small. This finding is robust to variation in all of the model parameters, except when the variance of the exogenous shock process becomes unrealistically large. See Jones et al. (2005) for discussion.

Table 1 presents the results for the standard deviation of: the growth rates of output, γ_y , and labor productivity, $\gamma_{y/n}$; the standard deviation of the investment–output ratio,⁶ x_k/y ; and the coefficient of variation in hours worked, n . This is done for both the endogenous growth model and its exogenous growth counterpart; the table also presents the corresponding values for the US data. Throughout, we limit our attention to situations in which the model economy is on a stochastic balanced growth path. Because of this, the theory predicts that the variables listed above all have stationary distributions and hence no detrending is necessary. However, since it is not obvious that this holds for the data, and to facilitate comparisons with the standard approach, we also include the related calculations for H-P filtered data in Appendix B.

As can be seen from the table, the calibrated endogenous growth model performs well at matching the volatilities observed in the data. (Recall that the value of σ was chosen to match $\sigma(\gamma_y)$ found in the data, given the calibration of the remaining parameters.) For ex-

Table 1
Volatilities in the models and the data

	$\sigma(\gamma_y)$	$\sigma(\gamma_{y/n})$	$\sigma(x_k/y)$	$\sigma(n)/E(n)$
<i>(h, k) model</i>	0.0192	0.0123	0.0134	0.0338
RBC model	0.0121	0.0080	0.0064	0.0070
<i>US data</i>	<i>0.0192</i>	<i>0.0112</i>	<i>0.0139</i>	<i>0.0481</i>

Notes. The column labeled $\sigma(j)$ gives the standard deviation of j , where j corresponds to the growth rate of output (γ_y), the growth rate of labor productivity ($\gamma_{y/n}$), the investment–output ratio (x_k/y), and labor (n). The term $E(n)$ denotes mean number of hours. The rows correspond to simulated moments for the endogenous growth model, the exogenous growth model and the time series data from the US, respectively.

⁶ We report only the standard deviation of the investment–output ratio since, in the model, $y = c + x_h + x_k$, and it follows that $\sigma(x_k/y) = \sigma((c + x_h)/y)$. Hence, by construction, the variability of the consumption–output ratio and the investment–output ratio coincide.

ample, the standard deviation of the investment–output ratio is virtually the same as the US value, the coefficient of variation of hours worked is 71% of the US value, and the standard deviation of the growth rate of labor productivity exceeds the US value by 10%.⁷ In contrast, the RBC model shows too little volatility in all categories. The difference in hours worked is the most dramatic, with the model displaying only 15% of the variability found in the data.⁸ This confirms the well-known result that RBC models generate little amplification of business cycle shocks. In sum, the endogenous growth model studied here performs better than its exogenous growth counterpart on these dimensions, particularly in the volatility of hours worked.

Table 2 reports the autocorrelation properties of the simulated data as well as the US data. There are several interesting results. First, the endogenous growth model generates considerable persistence in output growth. Quantitatively, the endogenous growth model can account for about 50% of the degree of persistence in growth rates.⁹ It is shown in Section 5 that the degree of first order autocorrelation increases with the intertemporal elasticity of substitution, but this dependence is not quantitatively significant. In the exogenous growth model, in all cases considered, the autocorrelation of output growth is close to zero. This, of course, is just another instance of the well-known failure of RBC models to display realistic propagation mechanisms (see Cogley and Nason, 1995).

Second, in the endogenous growth model the autocorrelation of hours worked matches closely that of the exogenous shock process; in fact, for the case with $\delta_k = \delta_h$, these would match exactly (since k/h is constant). For our baseline calibration, the model overstates the measured autocorrelation of hours by 6%. The exogenous growth model displays values for $\rho(n)$ lower than that of the US data, and for some specifications, comes close to matching this dimension of the data. Third, the baseline endogenous growth model overstates the autocorrelation of labor productivity growth, while the RBC model understates it for all cases considered; but in both cases, the differences are small. For both specifications—

Table 2
Autocorrelations in the models and the data

	$\rho(\gamma_y)$	$\rho(\gamma_y/n)$	$\rho(n)$
(h, k) model	0.1890	0.3308	0.9745
RBC model	0.0320	0.1958	0.7945
US data	0.4049	0.2746	0.9172

Notes. The column labeled $\rho(j)$ gives the first-order autocorrelation coefficient of j , where j corresponds to the growth rate of output (γ_y), the growth rate of labor productivity (γ_y/n), and the number of hours worked (n). The rows correspond to simulated moments for the endogenous growth model, the exogenous growth model and the time series data from the US, respectively.

⁷ This is true irrespective of the calibrated balanced growth value of labor supply, and hence, the calibrated values of ψ and σ . In experiments, we have calibrated the model using $n_{BGP} = 0.3$, a value close to those used in the RBC literature. None of the substantive results presented in this paper are sensitive to this choice.

⁸ The smallest value of σ for which we were able to obtain accurate solutions was 0.65. This case constitutes the “best fit” for the RBC model, with $\sigma(\gamma_y)$ still only 83%, and $\sigma(n)/E(n)$ only 28%, of the US values.

⁹ The value of 0.4049 for the US data is quite sensitive to the time period chosen. For example, for the period 1959 to 1996 it is 0.3278, much closer to that from the (h, k) model.

endogenous and exogenous growth—the autocorrelation of growth rates of the endogenous variables depends on the value of σ , the intertemporal elasticity of substitution. This dependence is documented in detail in Section 5.

There are two sources for the improvement in matching the first order autocorrelation of output growth over what is seen in RBC models. If $\delta_k = \delta_h$, the class of models here behaves just like a one sector growth model, but with a high total capital share (of one, overall). Thus, the overall share of capital being higher is one source of difference. In addition, in the calculations presented above, $\delta_k > \delta_h$. As it turns out this assumption itself gives rise to interesting and rich model dynamics in response to business cycle shocks. This is because with $\delta_k > \delta_h$, it is optimal to follow a positive shock by first investing in physical capital, and investing in human capital with a delay. Roughly, first plant and equipment is installed, and, second, workers are trained. This property of the model is discussed in detail below. By analyzing a version of the model in which $\delta_k = \delta_h$, it is possible to estimate the quantitative importance of these two qualitative changes from the RBC model. As it turns out, this is about 50% for each of these two components. This is discussed in more detail in Section 4 below.

The last statistic we present is the cross correlation, at several leads and lags, of key variables in the simulated and US data. Table 3 displays the cross-correlation of output growth, labor productivity growth, hours worked, and the physical capital investment–output ratio, with the growth rate of output. Overall, neither model is a complete success at matching the US data. As indicated above, the endogenous growth model generates significant persistence in output growth at the 1-year lag before falling at lag 2. Though this is an imperfect match with the US data (too little persistence at lag 1, slightly too much at lag 2), it is a marked improvement over the RBC model. In the US data, productivity growth leads output growth, while hours worked lags it. The endogenous growth model is able to match

Table 3
Cross-correlations in the models and the data

Cross-correlation		$j = -2$	$j = -1$	$j = 0$	$j = 1$	$j = 2$
$\gamma_{y,t}, \gamma_{y,t+j}$	(h, k) model	0.0969	0.1890	1.0000	0.1890	0.0969
	RBC model	0.0264	0.0320	1.0000	0.0320	0.0264
	US data	0.0549	0.4049	1.0000	0.4049	0.0549
$\gamma_{y,t}, n_{t+j}$	(h, k) model	0.2061	0.2117	0.4495	0.4646	0.4485
	RBC model	0.0121	0.0233	0.6344	0.5064	0.4053
	US data	-0.0628	-0.0116	0.3357	0.4413	0.3551
$\gamma_{y,t}, \gamma_{y/n,t+j}$	(h, k) model	0.1367	0.2225	0.9745	0.2545	0.1979
	RBC model	0.0299	0.0376	0.9793	0.1617	0.1295
	US data	0.3947	0.5578	0.2413	-0.1706	0.2316
$\gamma_{y,t}, (x_k/y)_{t+j}$	(h, k) model	0.0336	0.0443	0.9516	-0.0189	-0.0139
	RBC model	0.0123	0.0234	0.6341	0.5062	0.4053
	US data	0.0370	0.2850	0.6339	0.5928	0.3163

Notes. The variable labeled $\gamma(i)$ is the growth rate of i , where i corresponds to output (y) and labor productivity (y/n); n indicates hours worked, and x_k/y is the physical capital investment–output ratio. The table displays the cross-correlation of each of these variables with output growth at leads and lags of $j = -2, \dots, +2$. The rows correspond to simulated moments for the endogenous growth model, the exogenous growth model and the time series data from the US, respectively.

the lagging structure for hours worked, while the RBC model better captures the near zero leading correlations. Neither model is able to match the cross-correlation pattern of labor productivity and output growth; this is not surprising given that both models are driven by a single technology shock. Finally, the RBC model provides a better match for investment's share of output. The reason for the endogenous growth model's poor performance on this dimension is not surprising, given the model's dynamics explained in the next section.

To summarize, we find that the simple endogenous growth model performs quite well at matching key statistics of the US time series at annual frequencies, and outperforms the analogous exogenous growth version. The endogenous growth model displays greater variability in hours worked, despite the absence of an unshocked sector which "competes" with the market sector for labor resources (e.g., home production and human capital production through formal training) and utility which is linear in leisure. Similarly, it generates autocorrelated output growth without relying on hard-to-measure adjustment costs, asymmetries, and other more extreme assumptions.

4. The dynamics of a response to a shock—the human capital investment channel

The model that we analyzed in the previous section shows that endogenous growth models generate increased labor supply volatility and display the kind of internal propagation mechanisms that are necessary to match US observations on persistence. In this section we detail the mechanisms in the model that generate these properties.

There are two basic differences between the model that we present and the standard RBC model. The first is that capital's share overall (i.e., the combined shares of physical and human capital) is larger in endogenous growth models, which has important consequences for the response to a shock. This can be seen even in non-stochastic versions of the models. In particular, how do these two classes of models respond to a surprise increase in capital stocks when the system is on its balanced growth path? As it turns out, they respond very differently. In the exogenous growth model, such an increase from the balanced growth path level of k causes an immediate increase in output followed by a dampened return, in levels, to the balanced growth path. Thus, the initial increase is followed by a series of relative decreases. From a mechanical point of view, this seems to be the basic force driving the lack of persistence in the RBC model: every increase in the growth rate of output (triggered by a higher than average productivity shock) is followed by a series of lower than average growth rates (as output and capital track back down to the balanced growth path). In contrast, in the endogenous growth model a surprise increase in capital stocks (for simplicity assume that both k and h are increased proportionally) triggers a new trend line in levels, but with the same growth rate as before.¹⁰ Thus, there is no dampened return to the original time path, as seen in the exogenous growth model.

The second difference is due to the fact that $\delta_k > \delta_h$. Because of this, the endogenous growth model has the property that shocks have different effects on the accumulation of human and physical capital. That is, a positive technology shock results not only in a

¹⁰ This follows directly from the homogeneity properties of the model.

contemporaneous increase in physical capital investment, x_k , but also an increase in x_k relative to x_h . This is because any given change in the shock has a larger impact on the rate of return to physical capital investment than it does on the rate of return to human capital investment (whenever the system is near the balanced growth path).

To see this, assume that initially the economy is operating along its balanced growth path with $s_t = 1$. From the non-stochastic version of the no arbitrage equation it follows that:

$$\tilde{F}_k(t) - \delta_k = \tilde{F}_h(t) - \delta_h,$$

or,

$$\tilde{F}_k(t) - \tilde{F}_h(t) = \delta_k - \delta_h > 0,$$

where $\tilde{F} = Ak^\alpha(nh)^{1-\alpha}$. That is, $\tilde{F}_k(t) > \tilde{F}_h(t)$ along the non-stochastic balanced growth path. In the stochastic version of the model, the relevant version of the no arbitrage condition is:

$$E_t \{ u_c(t+1) [s_{t+1} \tilde{F}_k(t+1) - \delta_k - (s_{t+1} \tilde{F}_h(t+1) - \delta_h)] \} = 0.$$

But since, $\tilde{F}_k(t) > \tilde{F}_h(t)$ near the balanced growth path and, following a positive shock to s_t , s_{t+1} is likely to be greater than one, the solution cannot keep $h_{t+1}/k_{t+1} = h_t/k_t$. If it did, it follows that:

$$s_{t+1} \tilde{F}_k(t+1) - \delta_k - (s_{t+1} \tilde{F}_h(t+1) - \delta_h) > 0,$$

violating the no-arbitrage condition. This is simply because a given change in s has a larger effect on the rate of return to k . This is because a larger portion of it is subject to the shock (the depreciation component is not subject to the shock). Because of this asymmetry, it is necessary to follow a positive shock to s with a policy that lowers the rate of return on k (and the opposite for a negative shock to s). This is accomplished by directing relatively more investment to k and relatively less to h . Thus, the optimal policy is such that $h_{t+1}/k_{t+1} < h_t/k_t$ and this, in turn, results in relatively low human capital investment in the period of a positive shock.¹¹ Of course, as the shock trends back to its average value, these investment incentives eventually turn the opposite way (in the calibrated model presented in the previous section, this occurs after one period).

This property has a simple interpretation. When times are good ($s > 1$), firms optimally choose to increase their stocks of physical capital first. This is followed, with a short lag, with investment in human capital. Each of these requires additional hours and hence the response in hours worked is sustained. An alternative, more realistic interpretation of this timing could be that in response to a positive shock, firms first increase the value of their physical capital stock by upgrading the quality of their equipment. The subsequent increase in human capital would be, under this interpretation, retraining the work force to work with the upgraded physical capital stock.

¹¹ Of course the explanation is only approximate since the theory only restricts the integral, and not each term, to be zero. However, it captures the right effects. Also, if labor is inelastically supplied and utility is linear in consumption, the reader can verify that this argument is exactly correct. Thus, it is a useful intuitive explanation for the effects seen in the model.




Fig. 1. Impulse response functions—investment.

In Fig. 1, we show the response of the two forms of investment to a positive one standard deviation shock to productivity which “hits” the economy in period 6. The data displayed in Fig. 1 correspond to percentage deviations from the original balanced growth path (i.e. the path that would have obtained in the absence of shocks). Figure 1 illustrates the arguments sketched above. During the period of the shock there is a large response—over 5% above trend—of investment in physical capital. As can be seen, this falls quickly to 1.5% within one period. However, in keeping with the discussion above, this investment boom does not die out as the shock does. Rather, the system moves to a new trend line, with the same growth rate but a higher overall level (i.e. a new balanced growth path). Thus, even after 40 periods, investment in physical capital is more than 3% above the previous trend growth path.

Although the long run effects of the productivity shock on x_h are the same as those to x_k , the timing of the response of is quite different. During the period of the shock human capital investment actually falls by 1% as x_k is increased. (Whether or not this occurs depends on the magnitude of the difference between δ_h and δ_k .) In the period after the shock, however, x_h rises to 3% above trend growth, and in line with the dynamics of physical capital investment, human capital investment settles to a level permanently higher than the original trend line.

Figure 2 displays the impulse response functions for other variables in the model: output, consumption (defined as $c + x_h$), hours and investment (in physical capital). The two most interesting features are the delayed response of measured consumption to a shock, and the relatively long lasting increase in hours. As indicated above, the behavior of consumption is driven not only by intertemporal substitution effects, as in the standard exogenous growth model, but also by the delayed response of x_h (included in our measure of con-




Fig. 2. Impulse responses—various.

sumption) to a technology shock. Similarly, the response of hours worked does not peak in the period of the shock. For this specification, hours are highest in the period following the shock. Moreover, the effects of the shock are relatively long lasting: after 10 periods, hours are more than 0.58% above normal, and after 20 periods they are still 0.4% above trend. This contributes to the high first-order autocorrelation in output growth displayed by the model. One period of higher than average output growth (due to the productivity innovation) is followed by a second period of higher than average output growth, coinciding with the peak response of hours; this is followed by a gradual transition to the new balanced growth path.

To understand why the endogenous growth model generates these dynamic responses, it is useful to note that the model differs from the standard RBC framework in two important ways. First, it is a model of two capital stocks with different depreciation rates. Second, taken together, these two types of capital have a large share in output (one, in the endogenous growth model). Each of these features plays a role in the qualitative differences relative to the RBC model. To better understand this, we examined a two capital stock neo-classical model along the lines of the standard RBC model. When the depreciation rates of the two capital goods are equal, these models behave much like the standard one capital stock, RBC model. In response to a shock, investment in both types of capital increases equally along with labor supply. This response is quickly dampened as the system returns to the balanced growth path.

When the depreciation rates of the two capital goods are significantly different however, the time path of the response varies from this typical pattern. First, in the period of the shock there is increased investment in the capital good with the higher depreciation rate. This is in keeping with the intuition about equalizing rates of returns given above. Because

of this difference in the timing of investment, this change moves the ratio of the two stocks off its stationary level. In the period after the shock, investment in the high depreciation capital good falls, much like in a one capital good model; however, investment in the low depreciation capital good increases. This helps to bring the stocks back into alignment as the system returns to the balanced growth path. Because of this second adjustment, the persistence of the response to the shock is drawn out longer than in standard neoclassical models. This response has an analog in the time path of labor supply as well. Indeed, when the difference in depreciation rates is large enough, the labor supply response peaks in the period after the shock, not the first as in the standard RBC model.

These statements about the timing of response are true roughly independent of the size of the shares of the capital goods. However, as the overall share gets larger, the size of the effects gets larger. This increases the persistence of the response (and the autocorrelation in γ_y) even more and dramatically increases the volatility of labor supply as a result. For example, in a version of the model where the combined share of the two capital stocks is 0.36, but the depreciation rates are unequal, a 1% shock to productivity causes a 0.4% increase in labor supply at the peak (which occurs in the period of the shock and persists into the subsequent periods), much the same as in standard RBC models.¹² On the other hand, with parameterizations similar to the one used here for the endogenous growth model, there is an 0.8% increase in labor at the peak, and this occurs in the period after the shock (with a 0.6% in the first and third periods). As a rough summary then, the difference in depreciation rates give rise to the persistence that we find in output growth and labor supply, while the increase in capital's share makes labor supply more volatile. (Details on these experiments are available upon request from the authors.)

From this discussion it is clear that human capital investment is the key feature of the endogenous growth model, responsible for the improvements relative to the RBC approach. This naturally raises the question of whether this effect has a counterpart in the data. Several series for the human capital stock in the US have been developed which can be compared with the series generated by the model (for examples, see Jorgenson and Fraumeni, 1989; and Ho and Jorgenson, 2001). Of particular interest is the timing of human capital growth relative to the booms and busts of the business cycle, i.e. the cross-correlations with output growth. This is presented, for both the model and the data constructed by Ho and Jorgenson (2001), in Table 4.¹³

In the data, there is virtually no contemporaneous correlation in output growth and human capital growth. At lag 1, the correlation is (mildly) negative. However, at lags of 2 and 3 years, the correlations are positive. This is consistent with lower than average human capital investment in the period of business cycle booms, and higher than average human capital investment 1 and 2 years following the boom. Although an imperfect match quan-

¹² Note that with a total share of the two capital goods less than one, this version of the model is no longer an endogenous growth one.

¹³ We use the human capital stock data for 1959 through 1999. These estimates are based on changes in the composition of years of education and experience over time within the workforce relative to their base year. They do not adjust for changes in the 'quality' of a given number of years of education due to the changes in subjects taught, etc., over the years. This data is available as an updated table from <http://post.economics.harvard.edu/faculty/jorgenson/papers/papers.html>.

Table 4

Cross-correlation of output and human capital growth in the models and the data

$corr(\gamma_{yt}, \gamma_{ht+j})$	$j = 0$	$j = 1$	$j = 2$	$j = 3$
(h, k) model	0.1989	0.1548	0.4773	0.4597
<i>US data</i>	0.0208	-0.0547	0.1241	0.2196

Notes. Each column gives the cross-correlation coefficient of output growth (γ_y) with the j th lag of human capital growth (γ_h). The rows correspond to simulated moments for the endogenous growth model and the time series data from the US, respectively.

titatively, the endogenous growth model displays a qualitatively similar pattern: a lower cross-correlation at lag 1 relative to lag 0, and positive cross-correlations at lags of 2 and 3 years.

We view this as a promising first step—in understanding business cycle dynamics, and the accumulation of human capital over the cycle. Given the imperfection in the data and the fairly simple model of human capital formation studied here, the results suggest a positive payoff to researching this channel for internal propagation further.

5. Sensitivity analysis

In this section, we present a sensitivity analysis of our results to the assumptions we have made. Of all of these, there are two key factors in determining the quantitative properties of the model. These are the degree of risk aversion in the utility function, and the difference in depreciation rates between physical and human capital. Because of this, we present two tables, showing how the results change when these two parameters are moved away from our benchmark values holding all other parameters fixed.

Table 5 shows how extreme the dependence of volatilities on utility curvature is in the endogenous growth model. For example, increasing σ from 0.9 to 2.00 decreases the variability in hours worked by a factor of 11 in the endogenous growth model. Though not as extreme, similar results hold for the other measures; for example, the standard deviation

Table 5

Sensitivity of the models to the curvature parameter σ

σ	(h, k) model				RBC model			
	$\sigma(\gamma_y)$	$\sigma(x_k/y)$	$\sigma(n)/E(n)$	$\rho(\gamma_y)$	$\sigma(\gamma_y)$	$\sigma(x_k/y)$	$\sigma(n)/E(n)$	$\rho(\gamma_y)$
0.9	0.0483	0.0450	0.1857	0.2965	0.0127	0.0074	0.0080	0.0243
1.0	0.0315	0.0258	0.0990	0.2491	0.0121	0.0064	0.0070	0.0320
1.2	0.0225	0.0166	0.0510	0.2062	0.0115	0.0052	0.0056	0.0416
1.4	0.0192	0.0134	0.0338	0.1890	0.0111	0.0043	0.0047	0.0477
1.6	0.0175	0.0117	0.0249	0.1818	0.0108	0.0037	0.0040	0.0523
1.8	0.0165	0.0107	0.0194	0.1789	0.0105	0.0032	0.0035	0.0561
2.0	0.0157	0.0099	0.0155	0.1780	0.0104	0.0029	0.0032	0.0594
<i>US data</i>	0.0192	0.0139	0.0481	0.4049	0.0192	0.0139	0.0481	0.4049

Notes. The values of key second moments are presented as σ varies across rows. The definitions of statistics and variables are contained in previous tables. This is shown for the endogenous growth model (on the left), the exogenous growth model (on the right), and the time series data from the US.

of output growth falls by a factor of 3 over this range. That the curvature of utility has such a large effect is somewhat surprising, since for most models of exogenous growth, this parameter plays only a small role.¹⁴ This can be seen in the rightmost columns of Table 5. As σ increases from 0.9 to 2.00, the variability in hours worked falls by only a factor of 2.5, while the standard deviation of output growth falls by only 22%.

As seen in the table, for values of σ approximately 1.4 and greater, the sensitivity in the endogenous growth model drops dramatically. Thus, although we choose $\sigma = 1.4$ as our base case to match $\sigma(\gamma_y)$ exactly, any value of σ greater than or equal to about 1.3 would work more or less as well on this statistic. For the most part, this is true for the other measures of volatility. The exception to this is the coefficient of variation in hours worked, where the quantitative reduction as a function of σ is still fairly large. The best match with the data occurs at $\sigma = 1.4$ for all of the volatility measures, except $\sigma(n)/E(n)$ where $\sigma = 1.3$ is the best. Some, but not all, of the autocorrelations generated by the model are also sensitive to σ . In particular, $\rho(\gamma_y)$ varies significantly as σ is changed, and is too low for all values of σ (though much higher than in the corresponding RBC model).

In sum, the endogenous growth model does quite well at values of σ in the 1.3 to 1.6 range, displaying too much volatility at lower values and too little volatility at higher values. Moreover, throughout this range of parameters, the endogenous growth model studied here performs better than its exogenous growth counterpart. This is particularly evident in the volatility of the growth rates and hours worked.

Table 6 reports the same set of statistics when δ_h is adjusted, holding all other parameters fixed. As with changes in σ , the key model statistics are also quite sensitive to this parameter. In keeping with the discussion above, the volatility of hours falls from 0.0498 to 0.0117 as δ_h increases from 1 to 10.2% per year. Thus, although the model works well at matching this statistic at our benchmark parameterization, this is sensitive to the value of δ_h that we have chosen. However, for the 1 to 4% range found in the empirical literature, the model performs well at matching the US data, and outperforms the standard RBC

Table 6
Sensitivity of the models to the depreciation rate δ_h

δ_h	$\sigma(\gamma_y)$	$\sigma(\gamma_y/n)$	$\sigma(x_k/y)$	$\sigma(n)/E(n)$	$\rho(\gamma_y/n)$	$\rho(n)$	$\rho(\gamma_y)$
0.01	0.0230	0.0138	0.0212	0.0498	0.3671	0.9801	0.2339
0.025	0.0192	0.0123	0.0134	0.0338	0.3308	0.9745	0.1890
0.04	0.0172	0.0114	0.0089	0.0252	0.3060	0.9686	0.1600
0.06	0.0154	0.0107	0.0050	0.0184	0.2871	0.9609	0.1401
0.08	0.0144	0.0104	0.0028	0.0146	0.2784	0.9541	0.1350
0.102(= δ_k)	0.0133	0.0100	0.0017	0.0117	0.2781	0.9486	0.1418
US data	0.0192	0.0112	0.0139	0.0481	0.2746	0.9172	0.4049

Notes. The values of key second moments are presented as δ_h varies across rows. The definitions of statistics and variables are contained in previous tables. This is shown for the endogenous growth model and the time series data from the US.

¹⁴ This may explain why the RBC literature has typically not explored the effects of alternative values of the elasticity of substitution. For a good survey, see Cooley (1995).

model. As with the experiments varying σ , the autocorrelation properties of the model are less sensitive to the value of δ_h than are the volatility statistics.

6. Conclusion

In this paper we have taken a preliminary look at a class of stochastic endogenous growth models. This class of models show an improvement—over exogenous growth models—in accounting for the first-order serial correlation of the growth rate of output and the measured variability of hours per worker. This is true despite the fact that the models that we study do not rely on adjustment costs, or asymmetries in either the technologies of production across sectors, or the incidence of shocks. We find these results encouraging. They suggest that the class of endogenous growth models has the potential to deliver the kind of amplification and internal propagation mechanisms that the real business cycle literature has been searching for.

In particular, our model implies that labor is allocated not just to the production of physical capital and consumption goods, but also to the production of human capital. The RBC model with exogenous trend growth obviously abstracts from this last role for labor input. As a result, it ignores the fluctuations in labor input that are due to the cyclical changes in the demand for human capital. Our analysis shows that these fluctuations in human capital investment amplify the variability in hours worked; this amplification mechanism is strengthened as the overall share of physical and human capital in production increases. In our endogenous growth model, this capital share is large, resulting in labor supply volatility close to that observed in the US data.

A second important finding stems from the fact that the depreciation rate on physical capital is larger than that on human capital. This single asymmetry imparts rich dynamics in the model's response to cyclical shocks. In particular, the demand for the two capital services are asynchronised. Investment in physical capital responds contemporaneously to a productivity shock, with human capital investment responding in the period after. This has the effect of drawing out the dynamic response of both labor and output. Not surprisingly, the endogenous growth model displays output growth with significant first-order autocorrelation. RBC models which abstract from this consideration lack this internal propagation mechanism.

Finally, we find that in contrast to exogenous growth models, the economies that we study imply that the intertemporal elasticity of substitution plays a major role in determining the second moment properties of macroeconomic time series. On the basis of our specification, we find that an intertemporal elasticity of substitution of approximately 0.7 does best at replicating the moments in the data. Moreover, from the perspective of the model, there is a substantial difference between 0.7 and 1.0 (logarithmic preferences), which is the most common specification used in the real business cycle literature.

There are several dimensions in which the model is found lacking. Foremost of these is that, although an improvement over the standard RBC model, there is still too little autocorrelation in output growth in the model relative to the data. A second weakness is that, although there are qualitative similarities in the overall pattern of cross correlations of output growth and human capital growth, this match is far from perfect.

It seems to us that the next step is to carefully explore the effects of generalizing the model. This includes generalizing both the details of the market production technologies, and the human capital production technologies. Some of this work has been done, and we plan to use more evidence to recover the differences in factor intensities across sectors.

Acknowledgments

We thank Craig Burnside, Larry Christiano, Ellen McGrattan, Jim Nason and Ed Prescott for their help, and the National Science Foundation for financial support.

Appendix A. Numerical methods used for solving the model

A.1. Constructing the US time series

Because we are interested in the study of simple, one-sector, closed economy models without government spending, we construct US time series data as follows. Our empirical measure of consumption is the sum of consumer expenditures on nondurable goods and services (both taken from the NIPA), the imputed service flow from the stock of consumer durables, and government expenditures on nondurables and services. The imputation of the service flow from consumer durables follows the procedure of Brayton and Mauskopf (1985), using the NIPA and the BEA's Fixed Asset/Wealth data. Government consumption is the NIPA's measure of government expenditures, less government investment on fixed assets (i.e., physical capital). Government investment on physical capital is calculated from the BEA data as the first difference in the net stock of government fixed assets (non-residential equipment, software, structures, and residential structures), plus consumption of fixed capital on existing stocks.

Physical capital investment is measured as the sum of private domestic investment on fixed assets, consumer expenditures on durable goods, net exports, inventory investment (all taken from the NIPA), and the estimated flow of government investment on fixed assets. Output is the sum of measured consumption and measured physical capital investment. Physical capital is measured as the net stock of private and government fixed assets plus the net stock of consumer durables, taken from the BEA data. All of the above data are measured in fixed-weight (1996) dollars.

The hours worked series is constructed from CES data as the total number of employees, multiplied by the average weekly hours of production workers (both from non-farm payrolls). Because the average weekly hours series covers only 1964–2000, these data were extrapolated back to 1959 using a regression on a constant and a quartic time trend. Data are expressed in per capita terms using the Census Bureau's civilian, non-institutional population aged 16 years and over.

A.2. Estimating the shocks from time series

For the RBC model, it is possible to back out the productivity shock series from the data using the model's production function. In particular, assume that we have observations on $\{y_t, k_t, n_t\}$, and write the production function as:

$$y_t = \exp(\zeta_t) A k_t^\alpha n_t^{1-\alpha}.$$

Given a calibrated value for α , dividing through by $k_t^\alpha n_t^{1-\alpha}$ allows us to identify the series, $\exp(\zeta_t)A$. The value of A and the series $\{\zeta_t\}$ are separately identified by using the restriction that $E(\zeta_t) = 0$.

For the endogenous growth model, it is not possible to identify the productivity shock series directly from the data and the production function. This is because of the additional unobserved variable, $\{h_t\}$. To identify the $\{\zeta_t\}$ series we need to make use of the model's decision rules. In this model:

$$y_t = \exp(\zeta_t) A \kappa_t^\alpha (h_t n_t)^{1-\alpha} = \exp(\zeta_t) A \kappa_t^\alpha h_t n_t^{1-\alpha},$$

where $\kappa_t = k_t/h_t$. Given the stationarity properties of the model, κ_{t+1} is a non-linear function of the state, (κ_t, ζ_t) . In particular, this decision rule can be derived from the no-arbitrage condition:

$$E_t \left\{ u_c(t+1) \left[\delta_h - \delta_k + \frac{y_{t+1}}{k_{t+1}} (\alpha - (1-\alpha)\kappa_{t+1}) \right] \right\} = 0.$$

Given data on $\{y_t, k_t, n_t\}$, calibrated values for model parameters and this decision rule, it is possible to derive the relevant unknowns, A , $\{\zeta_t\}$, and $\{h_t\}$.

Begin with an initial value for the human to physical capital ratio, κ_0 . Given the observation for k_0 , we get $h_0 = k_0/\kappa_0$. The decision rule allows us to obtain each subsequent value for κ_{t+1} and $h_{t+1} = k_{t+1}/\kappa_{t+1}$. With this, we are able to identify A and $\{\zeta_t\}$ as described above for the RBC model.

In the special case in which $\delta_h = \delta_k$, this is particularly simple. This is because the decision rule is a constant function, $\kappa_t = (1-\alpha)/\alpha$ for all t . Observation on $\{y_t, k_t, n_t\}$ allow us to immediately identify the shock process. For the case in which $\delta_h \neq \delta_k$, we back out the shocks iteratively. We begin with an initial guess of the parameters of the shock process, ρ and σ_ε^2 . Using these, we solve the model to generate the decision rule for κ_{t+1} . This allows us to back out the shock process $\{\zeta_t\}$ as described above. We recompute the parameters, ρ and σ_ε^2 , using these estimated shocks and re-solve the model with the updated parameter estimates. This procedure is repeated until the $\{\zeta_t\}$ converge.

A.3. Numerical solution method

This appendix outlines the numerical method used to solve the endogenous growth models studied in this paper. The method is an extension of the general class of projection methods developed in Judd (1992). For further details, as well as the description of the method's use in solving the standard real business cycle model, see Siu (1998).

Let $\kappa = k/h$, $\tilde{c} = c/h$, and $\eta = h'/h$. The method begins by specifying the approximation to the value function to be of the form:

$$v(k, \zeta) = \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} a_{ij} T_i(\varphi(\kappa)) T_j(\varphi(\zeta)),$$

where, T_i is the i th order Chebychev polynomial, and $\varphi: \mathfrak{R} \rightarrow [-1, 1]$ is a linear transformation mapping the bounded, ergodic capture regions of κ and z into the Chebychev polynomial's domain of definition.

The $N^2 \times 1$ coefficient vector, $a = \{a_{ij}\}_{i,j=1}^N$, characterizing the approximation, v , is chosen to solve the following system of N^2 equations:

$$\iint R(\kappa, \zeta; \hat{v}) w_{ij} \partial \varphi(\kappa) \partial \varphi(\zeta) = 0, \quad i, j = 1, \dots, N, \tag{3}$$

where

$$R(\kappa, \zeta; v) \equiv v(\kappa, \zeta) - \max_{n, \eta, \kappa' \in \Gamma(\kappa, \zeta)} \left[\frac{[c(1-n)^\psi]^{1-\sigma}}{1-\sigma} + \beta \eta^{1-\sigma} \int v(\kappa', \zeta') P(\zeta', d\zeta') \right],$$

and

$$\Gamma(\kappa, \zeta) = \{n, \eta, \kappa' \mid 0 \leq n \leq 1, 0 \leq \tilde{c} \leq \exp(\zeta) A \kappa^\alpha n^{1-\alpha} + (1 - \delta_k) \kappa + 1 - \delta_h - \eta(1 + \kappa')\}.$$

We select the N^2 functions $\{w_{ij}\}$ to be:

$$w_{ij} = \frac{T_{i-1}(\varphi(\kappa))}{(1 - \varphi(\kappa)^2)^{1/2}} \frac{T_{j-1}(\varphi(\zeta))}{(1 - \varphi(\zeta)^2)^{1/2}}.$$

To compute the double integral above, we use M^2 -point Gauss–Chebychev quadrature integration (with $M > N$). In the special case where $\delta_h = \delta_k$, $\kappa \equiv \alpha/(1 - \alpha)$ so that the dimension of the state space is reduced in half and the dimension of the unknown coefficient vector a is reduced by a factor of N .

Due to the orthogonality conditions possessed by Chebychev polynomials, we are able to rewrite the system of equations (3) as the following matrix expression:

$$X'(Xa - Y_a) = 0,$$

where X is an appropriately defined matrix of Chebychev polynomials and Y_a is a vector consisting of non-linear functions of the unknown coefficients a . To solve this matrix problem easily, we implement the following iterative procedure. Given an initial guess for the coefficient vector, a_0 , the vector Y_{a_0} is computed. A new guess, a_1 , is computed as:

$$a_1 = (X'X)^{-1} X'Y_{a_0}.$$

This procedure is repeated until we converge upon the limit point, a^* .

Table A.1
Volatilities in the models and the data

	$\sigma(y)$	$\sigma(y/n)$	$\sigma(c + x_h)$	$\sigma(x_k)$	$\sigma(n)$
<i>(h, k)</i> model	0.0209	0.0123	0.0224	0.0638	0.0090
RBC model	0.0132	0.0095	0.0088	0.0300	0.0046
<i>US data</i>	0.0237	0.0122	0.0201	0.0690	0.0211

Notes. The variable labeled $\sigma(j)$ gives the standard deviation of j , where j corresponds to (filtered, logged) output (y), labor productivity (y/n), measured consumption ($c + x_h$), investment (x_k), and labor (n). The rows correspond to simulated moments for the endogenous growth model, the exogenous growth model and the time series data from the US, respectively.

Appendix B. H-P filtered second moments

In this appendix, we present volatility statistics for key variables from the models and the US data. This is done in Table A.1. All series have been filtered using the Hodrick-Prescott filter with a smoothing parameter value of 400. In terms of matching the data, the endogenous growth model outperforms the RBC model on all dimensions. It is interesting to note that the endogenous growth model does a good job at matching the volatility of filtered output, despite the fact that the curvature parameter, σ , was chosen to match the standard deviation of output growth. Finally, the calibrated model does a worse job of matching the US data in terms of filtered per capita hours as opposed to unfiltered hours. This is, in part, due to the fact that the endogenous growth model produces a highly auto-correlated hours series; from Table 2, we see that $\rho(n) = 0.9745$. The H-P filter removes much of the lower frequency variability generated by the model. Nonetheless, filtered hours in the endogenous growth model are twice as volatile as that of the RBC model.

References

- Andolfatto, D., 1996. Business cycles and labor-market search. *American Economic Review* 86 (1), 112–132.
- Alvarez, F., Stokey, N.L., 1998. Dynamic programming with homogeneous functions. *Journal of Economic Theory* 82, 167–189.
- Benhabib, J., Rogerson, R., Wright, R., 1991. Homework in macroeconomics: Household production and aggregate fluctuations. *Journal of Political Economy* 99 (6), 1166–1187.
- Ben-Porath, Y., 1967. The production of human capital and the life cycle of earnings. *Journal of Political Economy* 75 (4-1), 352–365.
- Brayton, F., Mankiw, E., 1985. The Federal Reserve Board MPS quarterly econometric model of the US economy. *Economic Modelling* 2 (3), 170–292.
- Burnside, C., Eichenbaum, M., 1996. Factor hoarding and the propagation of business-cycle shocks. *American Economic Review* 86 (5), 1154–1174.
- Caselli, F., 1999. Technological Revolutions. *American Economic Review* 89 (1), 78–102.
- Christiano, L.J., Todd, R.M., 1996. Time to plan and aggregate fluctuations. *Federal Reserve Bank of Minneapolis Quarterly Review* (Winter), 14–27.
- Cogley, T., Nason, J.M., 1995. Output dynamics in real-business-cycle models. *American Economic Review* 85 (3), 492–511.
- Collard, F., 1999. Spectral and persistence properties of cyclical growth. *Journal of Economic Dynamics and Control* 23, 463–488.
- Cooley, T.F. (Ed.), 1995. *Frontiers of Business Cycle Research*. Princeton Univ. Press, Princeton.

- Einarsson, T., Marquis, M.H., 1998. An RBC model with growth: The role of human capital. *Journal of Economics and Business* 50, 431–444.
- Einarsson, T., Marquis, M.H., 1999. Formal training, on-the-job training and the allocation of time. *Journal of Macroeconomics* 21 (3), 423–442.
- Greenwood, J., Hercowitz, Z., 1991. The allocation of capital and time over the business cycle. *Journal of Political Economy* 99 (6), 1188–1214.
- Greenwood, J., Yorukoglu, M., 1997. 1974. *Carnegie–Rochester Conference Series on Public Policy* 46, 49–95.
- Haley, W.J., 1976. Estimation of the earnings profile from optimal human capital accumulation. *Econometrica* 44 (6), 1223–1238.
- Hansen, G.D., 1985. Indivisible labor and the business cycle. *Journal of Monetary Economics* 16, 309–327.
- Heckman, J.J., 1976. A life-cycle model of earnings, learning and consumption. *Journal of Political Economy* 84 (4-2), S11–S44.
- Ho, M.S., Jorgenson, D.W., 2001. The quality of the US workforce: 1948–1995. Working paper.
- Jones, L.E., Manuelli, R.E., 1990. A convex model of equilibrium growth: Theory and policy implications. *Journal of Political Economy* 98 (5), 1008–1038.
- Jones, L.E., Manuelli, R.E., Rossi, P.E., 1993. Optimal taxation in models of endogenous growth. *Journal of Political Economy* 101 (3), 485–517.
- Jones, L.E., Manuelli, R.E., Siu, H.E., Stacchetti, E., 2005. Fluctuations in convex models of endogenous growth, I: Growth effects. Working paper.
- Jorgenson, D.W., Fraumeni, B.M., 1989. The accumulation of human and nonhuman capital, 1948–84. In: Lipsey, R.E., Tice, H.S. (Eds.), *The Measurement of Saving, Investment and Wealth*. In: *NBER Studies in Income and Wealth*, vol. 52. Univ. of Chicago Press, Chicago.
- Judd, K.L., 1992. Projection methods for solving aggregate growth models. *Journal of Economic Theory* 58, 410–452.
- King, R.G., Rebelo, S.T., 1999. Resuscitating real business cycles. In: Taylor, J., Woodford, M., (Eds.), *The Handbook of Macroeconomics*. In press.
- King, R.G., Plosser, C.I., Rebelo, S.T., 1988. Production, growth and business cycles, I: The basic neoclassical model. *Journal of Monetary Economics* 21, 195–232.
- Lucas Jr., R.E., 1988. On the mechanics of economic development. *Journal of Monetary Economics* 22, 3–42.
- Manuelli, R.E., 2000. Technological revolutions and the labor market. Working paper. University of Wisconsin, February.
- Merz, M., 1995. Search in the labor market and the real business cycle. *Journal of Monetary Economics* 36, 269–300.
- Perli, R., Sakellaris, P., 1998. Human capital formation and business cycle persistence. *Journal of Monetary Economics* 42, 67–92.
- Siu, H.E., 1998. Parameterized value function iteration. Working paper. Northwestern University, October.