Review of Economic Dynamics **4,** 536–561 (2001)
doi:10.1006/redy.2000.0124, available online at http://www.idealibrary.com on IDEAL®

The Cyclical Behavior of Skill Acquisition[1]

David N. DeJong

Department of Economics, University of Pittsburgh, Pittsburgh, Pennsylvania 15260

and

Beth F. Ingram[2]

Department of Economics, University of Iowa, Iowa City, Iowa 52242

Received May 1, 1998; revised May 1, 2000; published online May 11, 2001

We use a business cycle model to analyze the general equilibrium implications of a representative agent's decision to devote time to skill acquisition activities, which are modeled as boosting subsequent labor productivity by increasing the stock of human capital. We use aggregate data on consumption, investment, and labor hours to estimate the parameters of the model, and then use the estimated model and the observed data to infer the aggregate behavior of skill acquisition activities. We find that these activities have important cyclical implications and are distinctly countercyclical; they also exhibit a systematic correspondence with college enrollment data. *Journal of Economic Literature* Classification Numbers: E32, J22, J24. © 2001 Academic Press

Key Words: on-the-job training; college enrollment; general-equilibrium model; full-information estimation.

Individuals routinely devote time to skill acquisition through participation in formal activities such as schooling, continuing education, and

[1] We gratefully acknowledge the financial support of National Science Foundation grants SBR 9422828 (to the University of Pittsburgh) and SBR 9422873 (to the University of Iowa). For helpful comments and suggestions, we thank two anonymous referees; Jeremy Greenwood, Gary Hansen, Greg Huffman, Ian King, Chuck Whiteman, Mark Wynne, seminar participants at McGill and Vanderbilt Universities and the Universities of Illinois, Iowa, Maryland, Pittsburgh, and Virginia; and participants at the 1997 Midwest Macroeconomics Meetings. The usual disclaimer applies.

[2] Corresponding author. E-mail: beth-ingram@uiowa.edu.

1094-2025/01 $35.00

training programs, and informal activities such as on-the-job training and professional activities pursued outside of the workplace. Many of these activities are not directly measurable; those that are measurable are often measured with considerable error and are rarely recorded on a time-series basis. For example, various surveys regarding on-the-job training activities provide snapshots of these activities at various points in time, but cannot be used to systematically assess cyclical behavior. These surveys clearly indicate the aggregate importance of on-the-job training and yet feature considerable measurement error (see, e.g., Barron, Berger, and Black, 1997). Regarding the importance of on-the-job training, the U.S. Congress Office of Technology Assessment (1988, table 3-26) estimated that firms invested between \$75 and \$205 billion in employee training in 1985; an investment of \$150 billion in that year would constitute roughly 4% of Gross National Product (GNP), 50% of that invested in plants and equipment, and 40% of total investment in human capital. Accounting for this large estimated range, the study notes for formal corporate training that "data in this large and important part of education are very poor." For informal training, the study notes that "the amount paid for training that does not occur in a formal setting is simply not known." Thus we lack comprehensive data on aggregate skill acquisition activities.

This lack of data is unfortunate; skill acquisition activities have clear macroeconomic implications and are themselves influenced by macroeconomic conditions. Regarding the former, by increasing the aggregate level of human capital and the marginal product of labor, decisions concerning human capital acquisition influence physical capital acquisition and labor/leisure choices of individuals, and thus the behavior of the aggregate economy. Regarding the latter, time spent pursuing skill acquisition activities entails an opportunity cost that varies with the business cycle. This opportunity cost appears to matter empirically. For example, college enrollments are an important component of skill acquisition activities for which we do have good time-series measures. These enrollments (as reported in the October supplement of the Current Population Survey) indicate a negative relationship between the growth rates of output and college enrollments in the U.S. The raw correlation between these series is −0.31 over the period 1970–1996. Moreover, Dellas and Sakellaris (1996) found college enrollments to be countercyclical using a probit analysis that controlled for a wide range of factors. According to their estimates, a 1% increase in the unemployment rate increased the average probability of college enrollment by 0.77 percentage points (with standard error 0.12) over the period 1968–1988; also, a 1% increase in GNP growth decreased the probability of enrollment by 0.16 percentage points (standard error 0.05) over this period.

This paper has two goals. The first is to explore business cycle implications of skill acquisition activities in a general equilibrium model. The second is to examine the cyclical behavior of skill acquisition activities implied by the model that we use. These goals are pursued jointly, using a business cycle model in which a representative agent endogenously allocates time among skill acquisition, leisure, and labor in separate production sectors (dedicated to consumption and investment goods). We estimate the parameters of the model by applying the Bayesian methods outlined by DeJong, Ingram, and Whiteman (2000a) to observable data on consumption expenditures, investment expenditures, and labor hours. Conditional on the observed data and the estimated model, we evaluate the importance of these activities in generating and propagating business cycle fluctuations and determine the implied behavior of skill acquisition activities (as well as the remaining unobservable variables included in the model) by backing out period-by-period values of the unobservable variables implied by the estimated model and the observed data.

In our model, skill acquisition activities influence macroeconomic activity by increasing the stock of human capital available for use in the separate production sectors in the following period. In parallel, time allocated to the production of investment goods increases the stock of physical capital available for use in the production sectors in the following period. Hence in the face of exogenous shocks, the agent may smooth her consumption profile by adjusting the extent of investment in both physical and human capital. Three types of exogenous shocks are included in the model: a total factor productivity shock, a productivity shock specific to the investment goods sector, and a shock to the skill acquisition sector.

An interesting result obtained from this analysis concerns the cyclical behavior of aggregate skill acquisition activities; conditional on the data and our estimated model, our measure of these activities is distinctly countercyclical. This result is consistent with the findings of Dellas and Sakellaris (1996) regarding college enrollments. Indeed, we compare our measure of skill acquisition activities with college enrollment data and find significant correspondence between these series. This result indicates that recessions are periods during which agents "retool" by investing in human capital. Coupled with the fact that investment in physical capital is clearly procyclical, this finding also indicates that agents treat investment in human and physical capital as substitutes in their efforts to smooth their consumption profiles.

The cyclical nature of skill acquisition activities is determined by two factors in this analysis: the equilibrium implications of the theoretical model and the empirical implications of the observed data. Regarding the theoretical model, *ceteris paribus*, positive productivity shocks have a negative effect on skill acquisition activities (because they increase the opportunity cost of not working), and positive skill acquisition shocks have a negative effect on output (because they increase the opportunity cost of not studying). Hence in the presence of uncorrelated shocks, the theoretical model entails countercyclical behavior in skill acquisition activities. Empirically, however, we find high positive correlation between the skill-acquisition shock and the two productivity shocks, suggesting that innovations that enhance productivity in the production sectors also enhance the productivity of skill acquisition activities. This correlation pattern conditions the model toward the prediction of procyclicality for the behavior of skill acquisition activities, because the direct effect of positive productivity shocks on output is of course positive, as is the direct effect of positive shocks in the skill acquisition sector on skill acquisition activities. However, our results indicate that the opportunity cost effects dominate the positive correlations between the shocks quite dramatically, thus our finding that skill acquisition activities behave countercyclically.

This finding demonstrates the influence that macroeconomic fluctuations exert on skill acquisition activities. Our analysis also illustrates the converse: skill acquisition activities provide an important source of, and explanation for, business cycle fluctuations. Specifically, despite the inclusion of standard sources of disturbances in the model [the total factor productivity (TFP) and investment shocks], our estimates assign clear importance to the role of study shocks in generating cyclical activity; the standard deviations estimated for the three shocks are similar, as are responses of the model variables to each of the shocks. Moreover, relative to standard real business cycle (RBC) models, we find that the behavior of the observable variables implied by our estimated model closely corresponds with the behavior of their empirical counterparts along several dimensions. In particular, our model has a strong internal propagation mechanism, generates an empirically relevant degree of labor hour volatility, and implies positive comovements between labor hours in the consumption goods and investment goods sectors. Finally, we find that the volatility of hours devoted to skill acquisition activities is higher than that of labor hours in the production sectors. Taken together, these findings indicate that skill acquisition activities play an important role in influencing observed patterns of business cycle activity.

Our model builds on recent business cycle analyses in which the production of consumption goods and investment goods occurs in separate sectors (e.g., Benhabib, Perli, and Sakellaris, 1997; Huffman and Wynne, 1998). The central feature of these models is that the economy is buffeted by both aggregate and sector-specific technology shocks, and some sort of friction prevents the agent from smoothly moving inputs between the sectors. As Huffman and Wynne noted, an unfortunate feature of multisector models is that the implied joint behavior of consumption and investment is often at odds with the behavior of comparable historically observed series: In U.S. data, consumption and investment are relatively highly correlated, whereas in multisector models, this correlation tends to be quite small, if not negative. The problem is that the friction that makes these models interesting also ensures that the sectors do not move together. As noted earlier, skill acquisition in our model is an activity that tends to be pursued during periods of low marginal productivity of labor. Agents move hours out of the investment sector and into skill acquisition during periods in which unfavorable technology shocks have pushed consumption below trend. Hence consumption and investment, as well as hours worked in these sectors, are more highly correlated in this model than in models in which leisure is the only alternative to labor in the use of time. Huffman and Wynne (1998) showed that introduction of an adjustment cost in the investment sector can induce a positive correlation between consumption and investment in a multisector setting; Einarsson and Marquis (1997) obtained a similar result by introducing human capital accumulation in a home-production model.

Our inclusion of skill acquisition activities in a business cycle model follows the work of Einarsson and Marquis (1997), Perli and Sakellaris (1998), and King and Sweetman (1988), who found that this modification is capable of resolving a long-standing empirical shortcoming of standard models: a too-low volatility of labor hours in standard models relative to U.S. data. Related efforts to improve on descriptions of labor hour volatility in RBC frameworks include the indivisible labor hours models of Hansen (1985) and Rogerson (1988); the home-production models of Benhabib, Rogerson, and Wright (1991) and Greenwood and Hercowitz (1991); the labor-hoarding models of Burnside, Eichenbaum, and Rebelo (1993) and Burnside and Eichenbaum (1997); and the fiscal disturbances model of McGrattan (1994).

We too study the empirical value added of incorporating skill acquisition activities in a business cycle framework, and our findings in this dimension complement previous work. But we go beyond this work by measuring the relative importance of innovations to the skill acquisition sector in generating business cycle fluctuations (compared to standard productivity shocks), and examining period-by-period movements in skill acquisition activities implied by the model that we estimate and the data used to obtain these estimates. The relative importance of innovations to the skill acquisition sector in generating aggregate fluctuations, the distinct countercyclicality of our measure of skill acquisition activities, and the correspondence of our measure of these activities with college enrollment data provide further support for the notion that skill acquisition activities are important components of aggregate business-cycle activity.

1. THE MODEL

Three goods are produced in the economy: a consumption good, a physical capital (or investment) good, and a human capital good. Physical capital allocated to the consumption sector, k_{ct-1}; labor allocated to the consumption sector, n_{ct}; and human capital, h_{t-1}, combine to produce the consumption good, c_t, through the following technology:

$$c_t \leq A_t k_{ct-1}^{\alpha}(h_{t-1}n_{ct})^{1-\alpha}, \tag{1}$$

where A_t represents a shock to total factor productivity. The same Cobb–Douglas technology governs the production of the investment good,

$$i_t \leq A_t A_{it} k_{it-1}^{\alpha}(h_{t-1}n_{it})^{1-\alpha}. \tag{2}$$

The shock A_t and the share parameter α are common to both production sectors, whereas the shock A_{it} is specific to the production of the investment good. Human capital augments the productivity of labor in each sector simultaneously, but with a lag. The allocation of physical capital between the sectors is fixed during the period, whereas labor is mobile. Physical capital depreciates at the constant rate δ. Hence investment increases the stock of physical capital according to

$$k_t - (1-\delta)k_{t-1} = i_t. \tag{3}$$

The representative agent has preferences over leisure and the single consumption good given by

$$E_0 \sum_{t=0}^{\infty} \beta^t \{\ln(c_t) + \ln(l_t)\}.$$

The agent is endowed with one unit of time, which is split between work in the consumption sector, work in the investment sector, skill acquisition, s_t, and leisure,

$$n_{ct} + n_{it} + s_t + l_t = 1. \tag{4}$$

The agent is also endowed with initial levels of consumption sector capital, k_{c0}; investment sector capital, k_{i0}; and human capital, h_0.

Skill acquisition activities augment the stock of human capital available for production as follows:

$$h_t = (1-\delta_h)h_{t-1} + A_{ht}g(s_t), \tag{5}$$

where δ_h is the rate at which human capital depreciates and A_{ht} is an exogenous shock that shifts the efficiency with which hours are transformed into human capital. Examples of a negative shock are the creation of a new computer operating system that is more difficult to learn than the previous system and a decrease in funding for government-sponsored training programs. A positive shock could be a technological improvement in employee training methods. The function $g(\cdot)$ controls the ability of the agent to transform time into human capital. We assume that $g(s_t) = (1/\theta)s_t^{\theta}$; the parameter θ may be either greater than 1 [so that $g(s)$ is convex] or less than 1 [so that $g(s)$ is concave]. Note that hours devoted to skill acquisition in the current period augment the stock of human capital available for use in production in the following period, enhancing the productivity of labor in each sector in the following period. The potential convexity of $g(s)$ does not introduce increasing returns to labor's input to production, because there remain decreasing returns to labor hours in the production functions and the opportunity cost of s is increasing in s.

Finally, the exogenous shocks are assumed to follow a vector autoregressive [VAR(1)] process with the constraint that the AR coefficient matrix be diagonal,

$$\begin{bmatrix} \ln A_t \\ \ln A_{it} \\ \ln A_{ht} \end{bmatrix} = \begin{bmatrix} \rho_A & 0 & 0 \\ 0 & \rho_i & 0 \\ 0 & 0 & \rho_h \end{bmatrix} \begin{bmatrix} \ln A_{t-1} \\ \ln A_{it-1} \\ \ln A_{ht-1} \end{bmatrix} + \begin{bmatrix} \varepsilon_{At} \\ \varepsilon_{it} \\ \varepsilon_{ht} \end{bmatrix}, \quad \varepsilon_t \sim N(0, \Sigma).$$

Nonzero off-diagonal coefficients are permitted for the covariance matrix Σ; that is, the shocks may be contemporaneously correlated.

Equilibrium for this economy is characterized by (1)–(5) and the following first-order conditions:

$$\frac{1}{1-\alpha}\frac{n_{ct}}{l_t c_t} = \frac{1}{c_t}, \tag{6}$$

$$\frac{n_{it}}{l_t i_t} = \beta E_t\left\{\frac{n_{it+1}}{l_{t+1} i_{t+1}}\left(1 - \delta + \alpha\frac{i_{t+1}}{k_{it}}\right)\right\}, \tag{7}$$

$$\frac{n_{it}}{l_t i_t} = \beta E_t\left\{\frac{n_{it+1}}{l_{t+1} i_{t+1}}\left(1 - \delta + \alpha\frac{i_{t+1}}{k_{it}}\frac{k_{it}/n_{it+1}}{k_{ct}/n_{ct+1}}\right)\right\}, \tag{8}$$

and

$$\frac{1}{l_t A_{ht} s_t^{\theta-1}} = \beta E_t\left\{\frac{1}{l_{t+1} A_{ht+1} s_{t+1}^{\theta-1}}\left[1 - \delta_h + \frac{n_{ct+1} + n_{it+1}}{h_t} A_{ht+1} s_{t+1}^{\theta-1}\right]\right\}. \tag{9}$$

Equation (6) is the condition governing intratemporal substitution; it implies a tight positive relationship between leisure and labor in the consumption sector. Hence intratemporal substitution between consumption and leisure is accomplished by adjusting labor hours in the consumption sector, thus inducing movements in hours in the investment and skill acquisition sectors. Equations (7)–(9) delineate the necessary conditions for intertemporal substitution. In (7) and (8), the marginal cost of physical capital is equated to the discounted expected marginal benefit of physical capital used in the investment sector and the consumption sector, respectively. In (9), the marginal cost of human capital is equated to the discounted expected marginal benefit of human capital, which is used simultaneously in both sectors. Hence intertemporal smoothing of consumption and leisure is accomplished through adjustment of the sector-specific physical capital stocks and the human capital stock.

The behavioral implications of the model can perhaps be best understood by examining the agent's response in equilibrium to a positive movement in each of the exogenous shocks under the assumption that the shocks are uncorrelated with each other. Given logarithmic preferences, a positive TFP shock induces the agent to work harder, enjoy less leisure, and increase consumption in the period in which the shock occurs. The resulting decline in leisure generates a decline in hours used in the consumption goods sector. In addition, because the shock reduces the marginal cost of physical capital, the agent finds the accumulation of physical capital to be less expensive than the accumulation of human capital, and thus reduces the hours devoted to skill acquisition activities. A positive innovation in the investment-specific technology shock induces similar qualitative movements in hours, but generates a current-period decline in consumption. A positive shock in the study hours sector decreases the cost of accumulating human capital relative to physical capital; hence the agent shifts hours away from the investment goods sector to the skill acquisition sector. In addition, because the opportunity cost of leisure rises, the agent reduces leisure hours, and thus by implication, hours worked in the consumption goods sector also fall. To help offset the resulting output decline in the consumption goods sector, the agent moves physical capital from the investment goods sector to the consumption goods sector.

The behavior of the economy in periods following the realization of a shock depends on the autocorrelation properties of the shock process. Figure 1 illustrates the impulse response functions of the endogenous variables to a one-standard deviation increase in each shock, assuming that the shocks are independent AR(1) processes. Figure 1(a) illustrates responses to the TFP shock, 1(b) illustrates responses to the investment shock, and 1(c) illustrates responses to the skill acquisition shock. Because 1(a) and 1(b) differ only in terms of the response of consumption, we concentrate our discussion on 1(a) and 1(c). The parameters of the model were fixed at their posterior means in constructing these responses, with the exception of the covariance matrix for the shocks, which was restricted to be diagonal.

The most important feature of these responses is that (with the exception of the capital series) they are nonmonotonic. That is, peak responses of labor hours, study hours, and investment all occur with a one-period

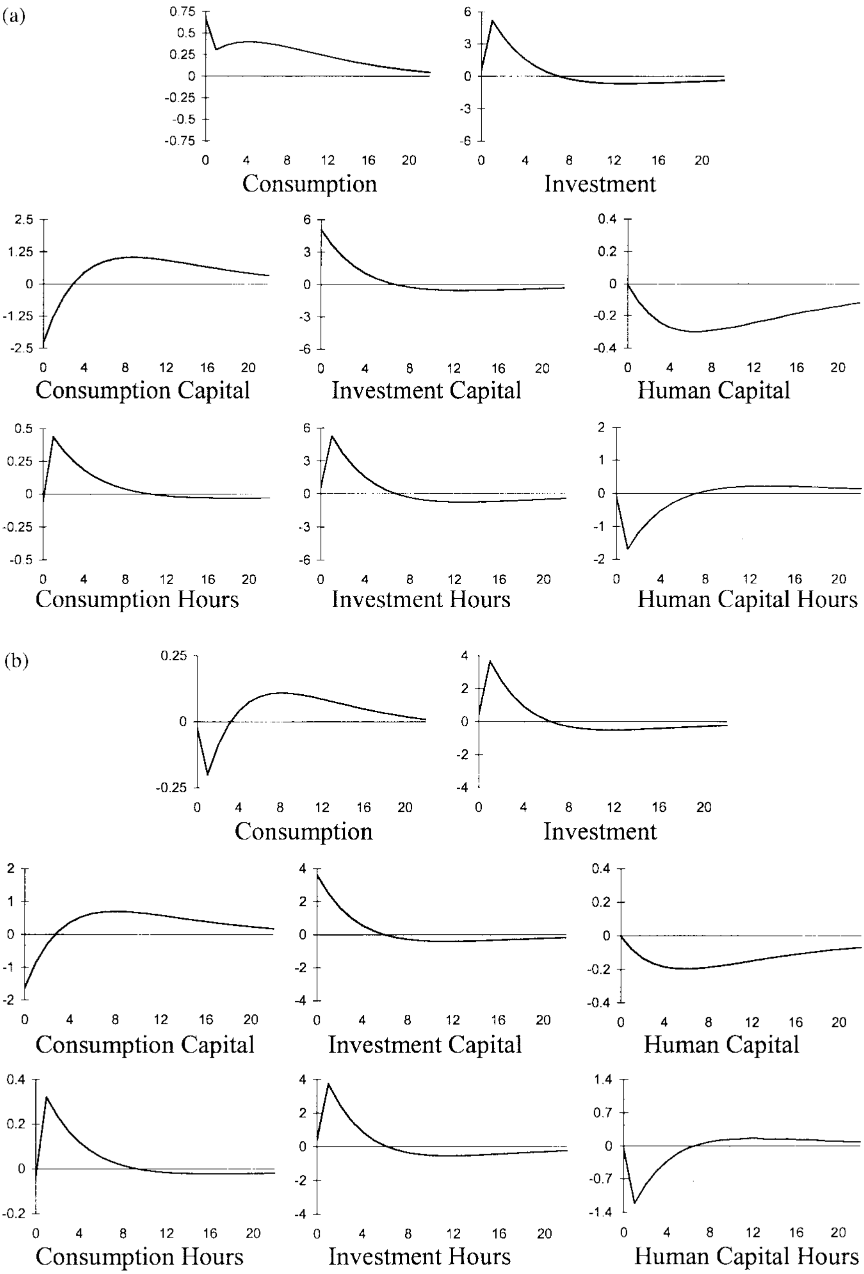

FIG. 1. Impulse response functions to (a) one-standard-deviation increase in TFP shock, (b) one-standard-deviation increase in investment shock and (c) one-standard-deviation increase in study hours shock.

(c)

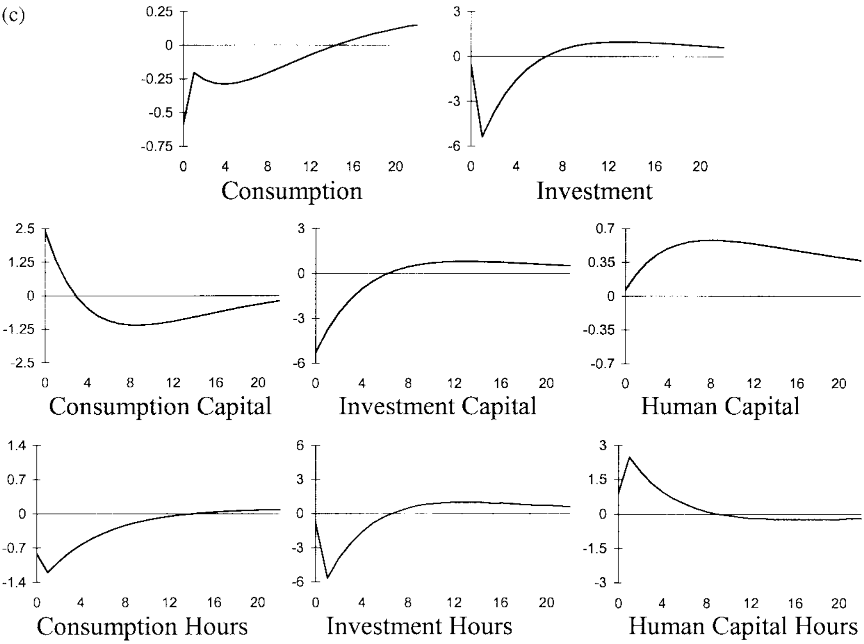

FIG. 1. *Continued*

lag, because of the sectoral rigidity in the stock of physical capital across production sectors. In addition, with the exception of consumption, all of the flow variables overshoot their steady-state values before returning to steady state. These patterns are indicative of the model's strong internal propagation mechanism, and contrast markedly with the pattern observed in standard RBC models.

Examining the responses to the TFP shock, note the dramatic fall in study hours in the period following the shock. Given persistence in the TFP shock, then the marginal productivity of labor remains above steady-state in subsequent periods, so the opportunity cost of studying remains high. Hence the agent optimally allocates time to labor hours at the expense of study hours. In addition, higher income leads the agent to devote time to leisure (not pictured), further decreasing study hours; study hours remain below steady state for about eight quarters after the realization of a shock, responding to the positive shock in a countercyclical fashion.

The model also produces rich intertemporal patterns in the accumulation of the capital stocks. Because the TFP shock lowers the marginal cost of producing physical capital, the agent's initial response to the shock is to move physical capital out of the consumption goods sector into the investment goods sector. Moreover, the agent allows the stock of human capital to depreciate, given the increased opportunity cost of skill acquisition activities. The additional capital goods generated by the initial shifting of physical capital into the investment goods sector are funneled into the consumption goods sector, so that consumption sector capital is gradually replenished. At about the time that the level of capital in the consumption goods sector peaks, the hours devoted to skill acquisition activities rise above steady state, thus increasing the stock of human capital. Hence the mechanism by which the agent smooths the effect of the shock on consumption and leisure is by initially boosting the stock of physical capital in the investment goods sector, then shifting physical capital to the consumption goods sector, and finally replenishing human capital.

Consider now the response of the economy to a positive study shock, which represents a decline in the opportunity cost of skill acquisition activities. In this case, the agent moves hours out of leisure and labor hours and into skill acquisition activities. To maintain consumption at a reasonable level (albeit below steady state), the agent reallocates capital from the investment goods sector to the consumption goods sector. The agent then allows consumption sector physical capital to depreciate while replenishing investment sector physical capital and accumulating human capital. Investment remains below steady state for eight quarters after realization of the shock, whereas consumption remains below steady state for almost 16 quarters.

A positive study shock has some of the same characteristics as a negative TFP shock; in particular, both result in a decline in output and countercyclical movements in skill acquisition activities. But although these responses describe the model's behavior, they leave such issues as cyclicality of skill acquisition activities unresolved, because they do not take into account correlation patterns between the shocks. For example, if the TFP shock and the study hours shock were positively correlated historically, this would impose competing effects on the cyclicality of study hours; the relative strength of the competing effects would then determine this cyclicality.

2. EMPIRICAL IMPLEMENTATION OF THE MODEL

Our empirical analysis of the model begins with its formal estimation. This is accomplished using the full-information Bayesian procedure developed by DeJong, Ingram, and Whiteman (2000a); we provide details of the procedure in the Appendix. Briefly, the procedure involves linearizing the model to obtain a characterization of the likelihood function for the observed data, then combining the likelihood function with a prior distribution over the parameters of the model to obtain a posterior distribution. Once the model is estimated, the resulting posterior distribution can be used to derive posterior distributions of functions of interest. Here functions of interest include summary statistics for the implied behavior of the variables included in the model (both observable and unobservable). In particular, we examine correlation patterns among the exogenous shocks, the relative volatility and correlation patterns implied for observable series, and period-by-period values of the model's unobservable series, including skill acquisition activities.

As a first step in our empirical analysis, we specify a prior distribution for the parameters. The prior distribution comprises a series of independent normal distributions over all of the parameters excluding Σ. We attempted to center the normal distributions at parameter values considered standard in the RBC literature, and specified dispersions for these distributions that are sufficiently diffuse to allow the data to have a nontrivial influence on our results. The priors are depicted in Figure 2 (dashed lines). The mean rate of time discount, β, is set at 0.99, implying a steady-state yearly interest rate of approximately 4% because a period in our model corresponds to one quarter. Also, the rate of depreciation δ and the technology parameter α are set to standard values, $\delta = 0.02$ and $\alpha = 0.29$.

The parameters governing the accumulation of human capital are relatively new to this exercise and are more difficult to pin down. In steady state, the ratio of study hours to labor hours satisfies

$$\frac{\bar{s}}{\bar{n}} = \frac{\delta_h \theta}{(1/\beta) - 1 + \delta_h}.$$

We set the means of δ_h and θ so that this ratio would equal 0.3 when β was equal to its mean. According to the Bureau of Labor Statistics (BLS), approximately 7.15% of the population was enrolled in postsecondary education and 71% was employed at the end of our sample period. Assuming that these groups enjoy equal leisure time, a ratio of 0.3 is consistent with, for example, the employed population's devotion of an average of 5 on-the-job hours and 2 off-the-job hours per week to study time. It bears emphasizing here that n represents time devoted explicitly to production activities, not time spent on the job. Time devoted by workers to skill acquisition activities while on the job is reflected in s, not n.

Numbers such as these may seem high at first. However, data from the National Center for Education Statistics indicate that 51% of employed individuals participated in formal work-related education programs at the end of our sample period (60% provided by businesses, 20% by postsecondary schools, and 20% by government and other services). Moreover, these data do not include informal training activities, and thus they understate total on-the-job training activity.

These percentages are merely suggestive of aggregate study hours, so we specified considerable prior diffusion along this dimension. The dispersion chosen for θ (a standard deviation of 0.075) is sufficient to generously span

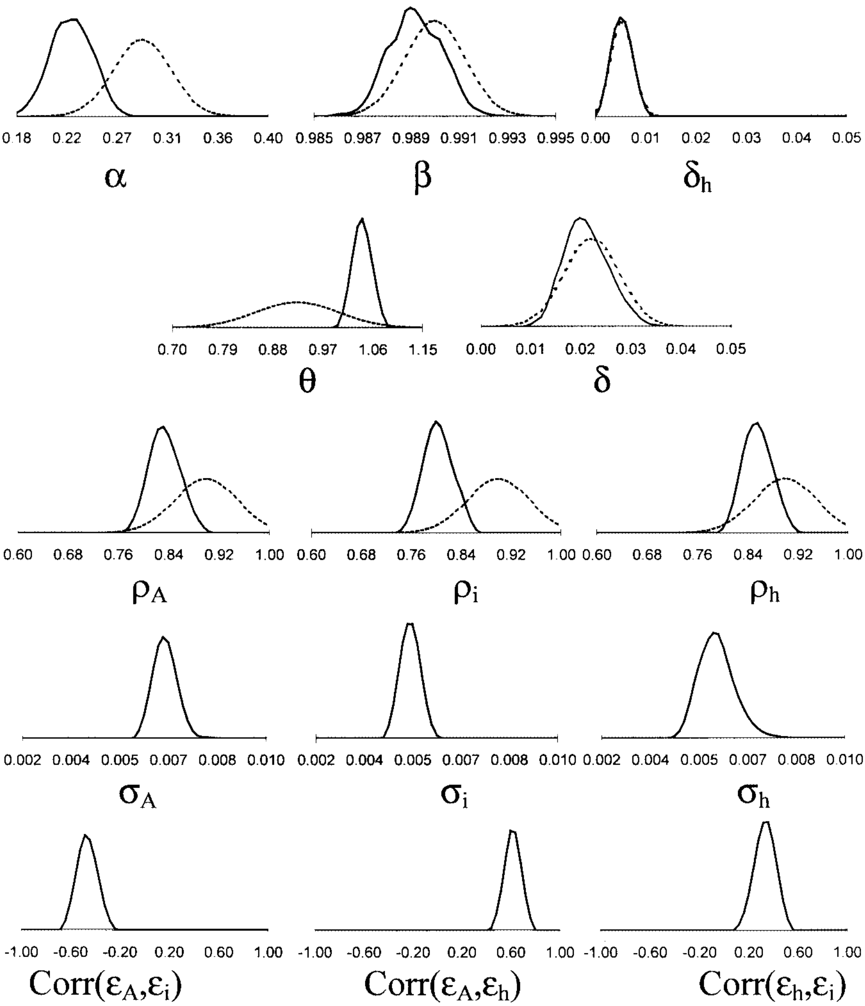

FIG. 2. Marginal prior (dashed lines) and posterior (solid lines).

both convex and concave specifications for the skill acquisition function $g(s)$. The dispersion chosen for δ_h is also generous (a standard deviation of 0.002); holding the remaining parameters at their prior means and $\bar{n}$ at 24 (the mean number of labor hours per week in our sample), this implies that as δ_h ranges two prior standard deviations above and below its prior mean, $\bar{s}$ has a range of approximately 2–12 hours per week. Of course, the remaining parameters can also vary, implying an even greater range for $\bar{s}$.

The prior over the parameters of the exogenous shock process is problematic; by assumption, these processes are unobserved. We posit a fair amount of persistence in the process by centering our prior for the AR coefficients at 0.9. Given the assumed prior standard deviation for these parameters (0.05), most of the prior mass falls between 0.8 and 1.0. Finally, we specify the usual noninformative prior over Σ, proportional to $\det(\Sigma)^{-(n+1)/2}$.

As noted, the second step in our empirical analysis involves coupling our prior distribution with the likelihood function induced by the linearized version of the model under the assumption of normality for the innovation vector ε_t. This yields the posterior distribution of the parameters, which is conditional on the model, our prior, and the dataset that we use. Our dataset comprises quarterly U.S. data from the first quarter of 1948 to the fourth quarter of 1995 for aggregate consumption, investment, and labor hours. Consumption is measured as real consumption of nondurables and services, investment is real gross fixed investment, and labor hours are average hours worked weekly from the household survey. All series are per capita, logged and detrended with Hodrick-Prescott (HP) filter and a parameter of 1600.

3. RESULTS

The empirical implications of the model of course depend on its parameterization. In this section, we first report the posterior distribution of the parameter vector of the model. We then discuss the implications carried by this distribution for the model's characterization of aggregate economic activity.

Summary statistics for the posterior distribution of the parameters are reported in Tables I and II; Table I lists the prior and posterior values of the means and standard deviations of the parameters, and Table II lists the posterior correlations among the parameters (recall that all prior correlations

TABLE I
Prior and Posterior Summary Statistics

	α	β	δ_h	θ	δ	ρ_A	ρ_i	ρ_h
Prior mean	0.290	0.990	0.005	0.925	0.020	0.900	0.900	0.900
Posterior mean	0.227	0.989	0.005	1.040	0.020	0.838	0.808	0.860
Prior S.D.	0.025	0.001	0.002	0.075	0.005	0.050	0.050	0.050
Posterior S.D.	0.019	0.001	0.002	0.013	0.004	0.023	0.023	0.022

	σ_A^2	σ_i^2	σ_h^2	$\mathrm{Corr}(A, I)$	$\mathrm{Corr}(A, h)$	$\mathrm{Corr}(h, I)$
Posterior mean	0.00671	0.00513	0.00581	−0.481	0.636	0.354
Posterior S.D.	0.00037	0.00027	0.00059	0.066	0.042	0.073

TABLE II
Posterior Correlation Matrix for Parameters

	α	β	δ_h	θ	δ	ρ_A	ρ_i	ρ_h	σ_A^2	σ_i^2	σ_h^2
α	1.00										
β	−0.31	1.00									
δ_h	0.01	0.12	1.00								
θ	0.34	0.20	0.58	1.00							
δ	−0.42	0.03	0.06	−0.03	1.00						
ρ_A	−0.01	−0.03	−0.03	0.17	0.20	1.00					
ρ_i	0.07	−0.12	0.07	0.17	0.13	0.86	1.00				
ρ_h	0.07	−0.01	−0.05	0.20	−0.13	0.85	0.87	1.00			
σ_A^2	−0.07	0.05	0.10	0.10	0.27	0.07	0.04	−0.03	1.00		
σ_i^2	0.04	0.05	0.10	0.05	−0.21	−0.17	−0.12	−0.06	0.23	1.00	
σ_h^2	−0.49	0.07	0.08	−0.59	0.17	−0.06	0.05	−0.08	0.19	0.10	1.00

are 0). Also, marginal prior and posterior distributions of the parameters are illustrated in Figure 2. Several features of our estimates are noteworthy. First, the posterior estimates governing the parameters of the exogenous process are tightly estimated relative to our priors; in addition, the process shows much less persistence than was assumed in the prior (with the posterior distributions of the ρ's shifted substantially to the left of our priors). These estimates imply that more of the persistence observed in the data emerges from features of the theoretical model than is typically the case in an RBC setting. This persistence arises in part from our empirical findings (low rates of depreciation and time discounting). But a more significant source of persistence is the addition of human capital in the model, which enhances the agent's ability to spread the impact of shocks over longer time horizons. Hence the model produces more of its persistence internally, rather than externally through the addition of ad hoc assumptions about the unobserved stochastic shocks.

Second, the posterior distribution of θ, which determines the concavity or convexity of the skill acquisition function, is tightly distributed in the convexity region: The data clearly indicate increasing returns to study time, despite the assignment of substantial prior weight to the concavity region. (Despite our results concerning θ, the model does not display increasing returns for parameter points in the posterior distribution.) This feature, coupled with the intersectoral immobility of the capital stock, enhances the nonmonotonicity of the impulse response functions of variables to shocks in this model.

Third, the posterior distributions of the contemporaneous correlation between the ε's indicate substantial correlation among the shocks. Consistent with a result found in a different setting (DeJong, Ingram, and Whiteman, 2000b), this particular dataset, in conjuction with the theoretical model, implies a negative correlation between the TFP and investment-specific shocks. (Here the posterior mean of this correlation coefficient is −0.48, which has a posterior standard deviation of 0.07.) Recall from Figure 1(b) that, *ceteris paribus*, the investment shock induces opposing movements in consumption and investment, whereas the TFP shock leads to comparable movements in consumption and investment. The negative correlation between the two shocks ensures that positive output movements are in general accompanied by positive movements in consumption, but may be associated with either rising or declining investment spending. This latter pattern of behavior is particularly important empirically. Although investment is clearly procyclical, its behavior is often distinct from that of output at business cycle turning points; the onset of recessions often coincides with positive investment growth, and the onset of recoveries often coincides with negative investment growth (see, e.g., DeJong, Ingram, and Whiteman, 2000b).

In contrast, our estimates indicate that TFP shocks are positively correlated with study shocks; the posterior mean of this correlation coefficient is 0.63, and its posterior standard deviation is 0.04. This implies that during a typical period in which the marginal productivity of labor is high; the marginal productivity of study hours is also high; thus the agent faces competing forces in determining how to allocate time among labor, leisure, and study hours. Viewed in isolation, this estimate renders the cyclicality of skill acquisition activities unclear. It also suggests that positive shocks to TFP have the additional positive effect of reducing the cost of adding to the stock of human capital; technological innovations seem to decrease the time cost associated with increasing the stock of human capital.

Finally, despite the assignment of zero prior correlation across parameters, there is substantial posterior correlation in many instances. Notably, the persistence parameters are highly positively correlated, indicating a lack of substitutability across alternative sources of persistence (i.e., the persistence in consumption, investment, and labor hours induced by the behavior of one shock is not ameliorated by the lack of persistence in another shock); the variances of the shock innovations are (moderately) positively correlated, indicating a similar lack of substitutability across sources of disturbances; and capital's share of output is highly positively correlated with θ and highly negatively correlated with β, δ, and σ_h^2.

To further explore the correspondence between the model and the data, we report posterior means and standard deviations of common summary statistics. Table III contains statistics on the behavior of the observable variables implied by theoretical model, and Table IV contains statistics calculated from observed data. The statistics on the observed data were generated using the posterior distribution of a six-lag VAR estimated using the HP-filtered data; the reported statistics are simple functions of the

TABLE III
Summary Statistics for the Model

x	$\sigma(x)^{*}100$	$\sigma(x)/\sigma(y)$	$\mathrm{Corr}(y(t), x(t-1))$	$\mathrm{Corr}(y(t), x(t))$
Output	1.77	1.00	0.95	1.00
	(0.214)	—	(1.18E−02)	—
Consumption	1.46	0.82	0.89	0.92
	(0.185)	(0.024)	(2.16E−02)	(1.56E−02)
Investment	5.38	3.07	0.73	0.79
	(0.574)	(0.074)	(2.70E−02)	(2.34E−02)
Labor hours	1.61	0.92	0.41	0.47
	(0.155)	(0.143)	(5.23E−02)	(5.05E−02)

Note. The posterior mean of indicated statistics with posterior standard deviation in parentheses. Standard deviations are reported as a percentage of own mean.

VAR parameters and were generated using standard numerical integration techniques.

Because the theoretical model was estimated using observed consumption, investment, and labor hours data, we do not wish to overemphasize the importance of the correspondence between the model and these data. However, the model's performance along these dimensions lends credibility to our inferences concerning the behavior of the unobservable variables reported later.

Beginning with the first column of Table III and Table IV, we note that the posterior mean of the standard deviation of output in the model (1.77%) is slightly higher than that found in the data (1.35%). However, as the second column indicates, the standard deviations of consumption, investment, and labor hours relative to output accord well with the relative standard deviations found in the data. (All standard deviations are reported as percentages of own means.) In addition, with the exception of labor hours, the model delivers roughly the correct degree of correlation between these series and output. Hence this model produces more rea-

TABLE IV
Summary Statistics for Data

x	$\sigma(x)^{*}100$	$\sigma(x)/\sigma(y)$	$\mathrm{Corr}(y(t), x(t-1))$	$\mathrm{Corr}(y(t), x(t))$
Output	1.35	1.00	0.85	1.00
	(0.039)		(0.014)	—
Consumption	0.86	0.64	0.79	0.86
	(0.031)	(0.016)	(0.021)	(0.013)
Investment	5.13	3.79	0.86	0.89
	(0.134)	(0.090)	(0.013)	(0.011)
Labor hours	1.48	1.10	0.81	0.83
	(0.046)	(0.036)	(0.018)	(0.017)

Note. Data are logged and detrended.

sonable correlation patterns than standard models. In one standard RBC model, Huffman and Wynne (1999) reported a negative correlation between consumption and output (−0.584) and a near-perfect correlation between investment and output (0.983). Adding adjustment costs in the investment sector reduces the investment output correlation (0.976) and changes the sign of the consumption output correlation (0.870). Christiano and Fisher (1998) also generated positive comovement by introducing habit persistence in the preference function of the agents.

The correlation of labor hours with output is lower than it should be (0.47 in the model compared to 0.81 in the data), but is nevertheless higher than the near-zero correlation implied by most RBC models. Finally, although not reported in Tables III and IV, the correlation between consumption and investment, as well as hours worked in these sectors, is of interest. As noted in the introduction, RBC models in which consumption and investment goods are produced in separate sectors typically imply a negative correlation between these variables, whereas these series are positively correlated in the data (in our dataset, the correlation is 0.59). Our model does not match the data perfectly along this dimension (the correlation that we obtain is 0.38), but including skill acquisition activities in this multisector framework clearly yields an improved fit along this dimension.

Conditional on the model and its parameterization, we next conduct inferences concerning the behavior of consumption sector labor hours, investment sector labor hours, and study hours. (Recall that leisure hours are proportional to consumption sector labor hours, and thus statistics on this series are not reported.) The procedure for measuring these series is as follows. Given a particular parameter setting, we use a smoothing algorithm to extract the unobserved series from the observed consumption, investment, and aggregate labor hours series. We then compute summary statistics for these series, and assign these statistics the posterior weight received by the corresponding parameter setting. Repeating this process for a large number of parameter drawings yields approximations of the posterior distributions of the summary statistics of the unobserved series. These posteriors are summarized in Table V.

Regarding the labor hours series, investment goods hours are far more volatile than consumption goods hours (with the standard deviation nearly four times greater for the former) and are much more closely correlated with output (0.67 versus 0.22). Appropriate sector-specific labor hours data are not generally available. Huffman and Wynne (1999) created one such series using U.S. quarterly data in which labor hours were allocated to the consumption and investment sectors based on the disposition of final output in the sector in which the labor is used. (Mining, construction, manufacturing, transportation, public utilities, and wholesale trade compose the investment sector, and finance, insurance, real estate, retail trade and ser-

TABLE V
Summary Statistics for Measured Series

	S.D.	Corr(Output)
Consumption hours	0.014	0.22
	(0.001)	(0.07)
Investment hours	0.053	0.67
	(0.005)	(0.03)
Skill acquisition hours	0.097	−0.36
	(0.028)	(0.058)

Note. Entries represent the posterior mean of the statistics, with posterior standard deviation in parentheses.

vices compose the consumption sector. Because all economic activity in a sector is allocated to either consumption or investment independent of the percentage of final output that flows to each activity, the correlations and standard deviations will tend to be biased toward equality.) Based on the Household Survey data for the period 1976–1994, the correlation between investment sector labor hours and output is 0.72, and the correlation between consumption sector labor hours and output is 0.64: Investment hours are twice as volatile as consumption hours (a relative standard deviation of 2.9 for investment compared to 1.4 for consumption). Hence our model produces less correlation between consumption labor hours and output and less volatility in consumption hours than is exhibited in this dataset.

Given that investment is far more volatile than consumption, both in the data and in the model (recall Table III), the relative volatility of the two labor hours series is not surprising. But the relatively high correlation with output exhibited by investment goods hours is surprising, because in the model investment itself is more weakly correlated with output than consumption is. This surprising result reflects the wedge between labor and output in the two production sectors driven by the behavior of the exogenous shocks.

The behavior of skill acquisition activities is quite distinct from the other hours series. First, the series is extremely volatile; its standard deviation is 0.097, nearly seven times higher than consumption goods hours and two times higher than investment goods hours. (Because all standard deviations are reported as percentages of own means, this comparison overstates the absolute volatility of study hours relative to labor hours. Calculated using posterior means of the parameters, the steady-state ratio of study hours to labor hours is approximately 1:3, so the overstatement is by a factor of around 3.) Hence the model and data imply that agents actively adjust their behavior along this dimension in response to shocks. Also, the series is distinctly countercyclical; its correlation with output is −0.36, with a posterior standard deviation of 0.058. As noted in the introduction, Dellas and Sakellaris (1996) obtained a similar result in a dataset that includes only 18 to 22 year-olds. Betts and McFarland (1995) measured the correlation between unemployment and enrollment at public 2-year colleges and found a negative correlation between unemployment and enrollment. Finally, Sakellaris and Spilimbergo (1999) found a countercyclical pattern of enrollment in U.S. universities by foreign students arriving from OECD countries. This finding is obtained despite the fact that, as noted earlier, the study hours shock is positively correlated with the TFP and investment goods shocks according to our estimates. Clearly, the behavior of skill acquisition activities is dominated by opportunity cost considerations: agents actively adjust their behavior along this dimension, diverting time resources away from skill acquisition activities when labor productivity is relatively high. By implication, agents treat investment in human and physical capital as substitutes, investing procyclically in physical capital and countercyclically in human capital.

Figure 3 presents the entire time paths of the three hours series generated by the model under our posterior-mean specification. Specifically, the series were constructed by fixing the parameters of the model at their posterior means, and using a smoothing algorithm to extract values of the unobserved series at each point in time from observed values of consumption, investment, and aggregate labor hours at each point in time. National Bureau of Economic Research (NBER) dates are included as reference points; the series are measured in terms of percentage deviations away from steady-state values. Both the volatility and countercyclicality of the study hours series are clearly evident in the figure: deviations from steady-state of 15% or greater are not uncommon for this series, with the sharpest spikes

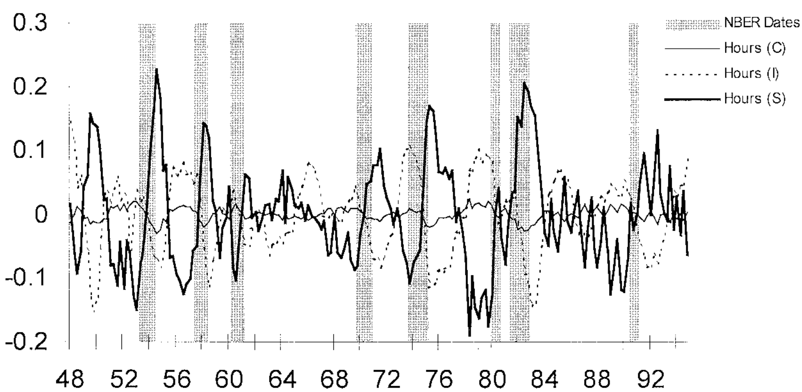

FIG. 3. Hours series implied by the model. Series are measured form the mean as a percentage of the standard deviation of the series and are calculated at the mean of the posterior distribution of the parameters. Shaded bars are NBER-dated recessions.

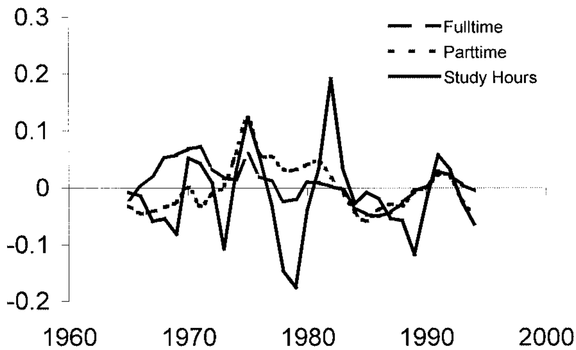

FIG. 4. Hours series. U.S. College enrollment data (full-time and part-time), measured in October and study hours data derived from theoretical model. All data are HP-filtered with a smoothing parameter of 1600.

in the series often closely coinciding with business cycle turning points. Between business cycle peaks, the series shows a general tendency to follow a gradual downward trend; once the next peak is reached, the series then tends to follow a steep upward trend (covering a total range of as much as 35 percentage points in the span of a few quarters). This pattern seems to be well underway in the latest set of observations.

We conclude our analysis by attempting to assess the coherence of our measure of skill acquisition activities with the actual behavior of these activities. This is impossible to do directly, because, as noted at the outset, we lack a comprehensive aggregate measure of these activities. The data we use for this purpose are full- and part-time per capita college enrollments, which are reported on an annual basis in the fall by the National Center for Education Statistics. (Source: Digest of Education Statistics, 1997, Table 172: Total Fall Enrollment in Institutions of Higher Education.) HP-filtered values of these measures, along with fall-quarter values of our skill acquisition measure, are illustrated in Figure 4. The three series exhibit substantial correspondence, with correlations between our measure and movements in full- and part-time enrollments of 0.35 and 0.31. Our measure is more volatile than the enrollment series; the standard deviation of our measure is 0.076, compared to 0.033 and 0.042 for full- and part-time enrollments. But this finding is not surprising, because college enrollments are probably relatively unresponsive to business cycle conditions as compared to other types of skill acquisition activities, such as employer-sponsored training or retraining programs. This caveat serves to emphasize that this comparison merely suggests the coherence of our measure with overall skill acquisition activity. Nevertheless, the comparison suggests that along this additional dimension, our model once again seems to have empirical relevance.

Each of these results is conditional on the data we have brought to bear in this exercise, as well as the model we have specified. Although not perfect, our model's overall performance in characterizing observable activity is quite good, even along dimensions that have traditionally been troublesome for RBC models. Moreover, the strong qualitative performance of the model is robust to the parameter uncertainty reflected in the parameter estimates we have obtained. Thus the clear linkages between aggregate economic activity and skill acquisition activities apparent in our results seem to us to have a high degree of empirical relevance, and further study of these linkages, facilitated by more direct measures of skill acquisition activities, seems warranted.

4. CONCLUSIONS

We have attempted to shed light on the interaction between aggregate economic activity and the pursuit of skill acquisition activities by individuals. Because aggregate skill acquisition activities, broadly defined, are unmeasurable, we have sought to infer their behavior using a formally estimated RBC model and measurable data on consumption, investment, and labor hours. Conditional on the model and the observable data, probabilistic measures of the unobservable variables included in the model are readily available. Statistical uncertainty associated with our parameterization of the model may be translated directly into uncertainty associated with our measures of these variables.

Our findings suggest the existence of important linkages between aggregate economic activity and skill acquisition activities. Innovations in skill acquisition technology appear to have important business cycle implications and have nontrivial interactions with conventionally measured innovations (e.g., TFP shocks). Moreover, agents appear to actively adjust their skill acquisition activities in response to exogenous shocks to achieve utility maximization objectives. An implication of this behavior is that skill acquisition activities, according to our measures, are distinctly countercyclical.

Our finding of countercyclicality is perhaps not surprising: It is certainly consistent with the behavior of college enrollments and accords with intuition (i.e, individuals pursue skill acquisition activities when their opportunity cost is relatively low). Alternatively, our finding that these activities have such a strong cyclical component—they are highly variable relative to labor and leisure hours, and moderately negatively correlated with output—is perhaps surprising; for example, college enrollments are relatively smooth. In response to this observation, we note that measurable components of skill acquisition activities, such as college enrollments, are arguably the least responsive to business cycle conditions: For instance, because of the skilled-wage premium, it is likely that the demand for higher education is relatively inelastic to opportunity cost considerations. However, unmeasurable activity, such as employee retraining necessitated by the adoption of a new computer system, is likely to be far more elastic. Thus our results may not be surprising after all.

This discussion serves to emphasize that as long as we lack comprehensive data on skill acquisition activities, the behavior and economic implications of these activities will be debatable. Having contributed to this debate, we hope to see progress toward attaining more direct measures of these activities. This would surely enhance our understanding of business cycle activity.

APPENDIX: ESTIMATION METHODOLOGY

Here we briefly outline our estimation methodology. (For further details, see DeJong, Ingram, and Whiteman, 2000a.) The methodology involves coupling the likelihood function associated with the theoretical model with a prior distribution specified over the parameters of the model to obtain a posterior distribution of these parameters. Posterior distributions of functions of these parameters are then approximated numerically.

The model that we estimate is the 12 equation system given by (1)–(9) and the three-equation VAR of the shocks. We begin by using standard log-linearization techniques to obtain a first-order system of the form

$$x_t = Fx_{t-1} + G\varepsilon_t, \tag{A1}$$

where x_t is a 12×1 vector of the variables of the model expressed as logged deviations from their steady-state values, $\varepsilon_t = (\varepsilon_{At}\varepsilon_{it}\varepsilon_{ht})'$ is the vector of VAR disturbances, and the elements of F and G are functions of the model's deep parameters, hereafter denoted by the vector μ. Although the system is of dimension 12, it is stochastically singular because there are only three random shocks. Therefore, the model carries nontrivial predictions for any three of the variables. We focus on its implications for consumption, investment, and labor hours, collected in the 3×1 vector X_t, and append to (A1) the observation equation

$$X_t = H'x_t, \tag{A2}$$

where the 3×12 matrix H' maps x_t into X_t.

We assume that the VAR disturbances ε_t are normally distributed, and thus obtain a likelihood function associated with the observer system (A1)–(A2). We analyze the likelihood function from a Bayesian perspective; the data are treated as fixed, and the parameters are treated as random variables. Denoting the sample of observations on the observed variables X_t by X, write the posterior distribution of interest as $P(\mu \mid X)$. This posterior is proportional to the product of the likelihood for the data given the parameters, the prior for the parameters, and the startup values of the data,

$$P(\mu \mid X) \propto L(X \mid \mu)P(\mu). \tag{A3}$$

(The factor of proportionality is the marginal distribution of the data.)

Calculating the likelihood function for the observer system (A1)–(A2) is straightforward using the Kalman filter. Then the mapping from the underlying economic parameters μ to the observer system parameters (the mapping implicit in the linearization procedure discussed earlier) can be used to complete the specification of the likelihood function $L(X \mid \mu)$. For the observer system, the Kalman filter equations are

$$\begin{aligned} x_{t\mid t} &= x_{t\mid t-1} + P_{t\mid t-1}H(H'P_{t\mid t-1}H)^{-1}(X_t - H'x_{t\mid t-1}), \\ x_{t+1\mid t} &= Fx_{t\mid t}, \\ P_{t\mid t} &= P_{t\mid t-1} - P_{t\mid t-1}H(H'P_{t\mid t-1}H)^{-1}H'P_{t\mid t-1}, \end{aligned}$$

and

$$P_{t+1\mid t} = FP_{t\mid t}F' + Q,$$

where $x_{t\mid t}$ is the optimal estimate of the (unobserved) state vector (the solution to the "signal extraction" problem) at time t given data through time t ($X_s, s = 1, \dots, t$), $x_{t+1\mid t}$ is the corresponding one-step-ahead predictor, $P_{t\mid t}$ is the covariance matrix of the error in estimating the state, $E(x_t - x_{t\mid t})(x_t - x_{t\mid t})'$, $P_{t+1\mid t}$ is the corresponding covariance matrix of the one-step-ahead predictor, $Q = GEv_tv_t'G'$, and the recursions are initialized by $x_{1\mid 0} = 0$ and $\operatorname{vec}(P_{1\mid 0}) = (I - F \otimes F)^{-1}\operatorname{vec}(Q)$, the values from the unconditional distribution.

Conditional on an initial observation X_0, the likelihood can be built up using the so-called "prediction error decomposition." In particular, define the prediction error as the error in predicting the observables one step ahead: $u_t = X_t - H'x_{t\mid t-1}$. Then the likelihood is given by

$$\log L \propto -\tfrac{1}{2}\sum_{t=1}^{T}\log\left|P_{t\mid t-1}\right| - \tfrac{1}{2}\sum_{t=1}^{T}u_tP_{t\mid t-1}^{-1}u_t.$$

Thus, given data and a candidate parameter value μ, the model can be transformed into the observer system (A1)–(A2), the Kalman filter applied, and the value of the likelihood function calculated.

Our interest is in conducting posterior inference for general functions of μ, $g(\mu)$. Examples of functions considered in the paper include the parameters themselves, impulse response functions, and date-by-date values of the unobservable variables of the model implied by the observed data and the parameterized model. The expected value of $g(\mu)$ under the posterior is given by

$$E[g(\mu)] = \frac{\int g(\mu)P(\mu \mid X)d\mu}{\int P(\mu \mid X)d\mu}. \tag{A4}$$

In general, the integrals in (A4) cannot be evaluated analytically; instead, they must be approximated using numerical integration techniques. Ideally, this is done by generating an artificial sample $\{\mu_k\}$ for $k = 1, \ldots, n$ directly from the posterior density (A3), and approximating (A4) by calculating the average value of $g(\mu)$ obtained using these drawings. Unfortunately, because the likelihood function in this case is in the form of an observer system, it is not possible to generate parameter drawings from its associated posterior distribution directly. Instead, we proceed by generating an artificial sample from a different distribution from which it is possible to sample, and assign weights to the elements of the sample so that they can be thought of as originating from the posterior distribution of interest. This technique is known as importance sampling. The distribution used to obtain drawings of μ is known as the importance density, denoted by $I(\mu)$. Given an artificial sample, (A4) is approximated by calculating the weighted average

$$\bar{g}_n = \sum_{i=1}^{n} g(\mu_i)w(\mu_i) \Big/ \sum_{i=1}^{n} w(\mu_i), \tag{A5}$$

where the weight function $w(\mu_i) = P(\mu_i \mid X)I(\mu_i); I(\mu_i)$ appears in the denominator of $w(\mu_i)$ to offset the direct influence that $I(\mu)$ has in obtaining the particular drawing μ_i. Given that the support of $I(\mu)$ includes that of $P(\mu \mid X)$, $\bar{g}_n$ converges almost surely to $E[g(\mu)]$, as long as $E[g(\mu)]$ exists and is finite.

To obtain the results reported in the paper, we specified a multivariate t distribution for $I(\mu)$ to ensure that its support included that of $P(\mu \mid X)$. The mean and covariance matrix of $I(\mu)$ that we ultimately used resulted from a sequence of preliminary runs. Initially, we assigned the prior mean and covariance matrix to $I(\mu)$, and over 10,000 drawings computed first-pass approximations of the posterior mean and covariance matrix of μ using (A5). However, very few of the drawings obtained in this manner were assigned appreciable weight by the posterior distribution, so we relocated $I(\mu)$ at our first-pass approximations and obtained second-pass approximations using 10,000 more drawings. After several rounds, our moment calculations converged (subject to numerical sampling error) to those used in deriving the results presented in the paper, which are based on 90,000 drawings. Of these drawings, that which was assigned the greatest weight received less than 3% of the total assigned weight; hence we are confident that our results closely approximate the actual posterior calculations that we seek.

REFERENCES

Barron, J. M., Berger, M. C., and Black, D. A. (1997). “How Well Do We Measure Training?,” *Journal of Labor Economics* **15**, 507–528.

Benhabib, J., Perli, R., and Sakellaris, P. (1997). “Persistence of Business Cycles in Multisector RBC Models,” unpublished manuscript, University of Pennsylvania.

Benhabib, J., Rogerson, R., and Wright, R. (1991). “Homework in Macroeconomics: Household Production and Aggregate Fluctuations,” *Journal of Political Economy* **99**, 1166–1187.

Betts, J., and McFarland, L. (1995). “Safe Port in a Storm—The Impact of Labor-Market Conditions on Community College Enrollments,” *Journal of Human Resources* **30**, 741–765.

Burnside, C., and Eichenbaum, M. (1997). “Factor-Hoarding and the Propagation of Business-Cycle Shocks,” *American Economic Review* **86**, 1154–1174.

Burnside, C., Eichenbaum, M., and Rebelo, S. (1993). “Labor Hoarding and the Business Cycle,” *Journal of Political Eocnomy* **101**, 245–273.

Christiano, L., and Fisher, J. (1998). “Stock Market and Investment Good Prices: Implications for Macroeconomics,” unpublished manuscript, Federal Reserve Bank of Chicago.

Congress of the United States, Office of Technology Assessment (1988). *Technology and the American Economic Transition: Choices for the Future.* Washington, D.C.: U.S. Government Printing Office.

DeJong, D., Ingram, B., and Whiteman, C. (2000a). “A Bayesian Approach to Dynamic Macroeconomics,” *Journal of Econometrics*, **98**, 203–223.

DeJong, D., Ingram, B., and Whiteman, C. (2000b). “Keynesian Impulses vs. Solow Residuals: Identifying Sources of Business Cycle Fluctuations,” *Journal of Applied Econometrics*, **15**, 311–329.

Dellas, H., and Sakellaris, P. (1996). “On the Cyclicality of the Demand for Education: Theory and Evidence,” unpublished manuscript, University of Maryland.

Einarsson, T., and Marquis, M. H. (1997). “Home Production with Endogenous Growth,” *Journal of Monetary Economics* **39**, 551–569.

Greenwood, J., and Hercowitz, Z. (1991). “The Allocation of Capital and Time over the Business Cycle,” *Journal of Political Economy* **99**, 1188–1214.

Hansen, G. (1985). “Indivisible Labor and the Business Cycle,” *Journal of Monetary Economics* **16**, 309–327.

Huffman, G., and Wynne, M. (1999). “The Role of Intratemporal Adjustment Costs in a Multisector Economy,” *Journal of Monetary Economics* **43**, 317–350.

King, I., and Sweetman, A. (1998). “Procyclical Retooling,” unpublished manuscript, University of Victoria.

McGrattan, E. (1992). “The Macroeconomic Effects of Distortionary Taxation,” unpublished manuscript, Federal Reserve Bank of Minneapolis.

Perli, R., and Sakellaris, P. (1998). “Human Capital Formation and Business Cycle Persistence,” *Journal of Monetary Economics* **42**, 67–92.

Rogerson, R. (1988). “Indivisible Labor, Lotteries and Equilibrium,” *Journal of Monetary Economics* **21**, 3–16.

Sakellaris, P., and Spilimbergo, A. (1999). “Business Cycles and Investment in Human Capital: International Evidence on Higher Education,” unpublished manuscript, University of Maryland.