Chapter 5
Business Cycles and Aggregate Labor Market Fluctuations

Finn E. Kydland

1. Introduction

Central to business cycle theory as well as to growth theory is the aggregate production function, which relates the nation's output of goods and services to the inputs of capital and labor. Of prime importance to business cycle theory is the behavior of the labor input. For growth, most of the output change is accounted for by changes in technology and in capital. In contrast, perhaps on the order of two-thirds of the business cycle is accounted for by movements in the labor input and one-third by changes in technology. Thus, most business cycle theorists agree that an understanding of aggregate labor market fluctuations is a prerequisite for understanding how business cycles propagate over time.

Table 5.1 lists statistics describing the cyclical behavior of key U.S. aggregates that are related to the labor input. The table includes measures of cyclical volatility, as well as correlations with cyclical real GNP, contemporaneously and at leads and lags of up to five quarters. The logarithms of the original series were detrended using the Hodrick-Prescott filter before the statistics were computed. (See Kydland and Prescott [1990] for details.) Some of the cyclical series are plotted against cyclical real GNP in figures 5.1–5.6.

Notable regularities related to the labor market are as follows:

1) Total hours, whether measured by the household or the establishment (payroll) survey, is almost as volatile as real GNP.

2) The household survey indicates that approximately two-thirds of the total-hours fluctuation is in the form of variation in employment and one-third is in hours per worker.

3) Total hours is highly procyclical, as indicated by the contemporaneous correlation coefficients with real GNP of nearly 0.9.

4) Total hours displays a slight phase shift in the direction of lagging the cycle, especially in the employment component. Hours per worker displays almost no phase shift.

Table 5.1
Cyclical Behavior of U.S. Labor Market Aggregates, 1954:I–1991:II

Variable	*Volatility (% SD)*	*Cross-Correlation of Real GNP with:*										
		$x(-5)$	$x(-4)$	$x(-3)$	$x(-2)$	$x(-1)$	x	$x(+1)$	$x(+2)$	$x(+3)$	$x(+4)$	$x(+5)$
Real Gross National Product	1.72	−.02	.16	.38	.63	.85		.85	.63	.38	.16	−.02
Hours (Household Survey)	1.49	−.10	.05	.25	.46	.70	.86	.85	.74	.58	.38	.17
Employment	1.09	−.17	−.03	.16	.38	.63	.83	.88	.80	.65	.46	.25
Hours per Worker	0.54	.07	.20	.36	.49	.64	.70	.58	.42	.28	.12	−.02
Hours (Establishment Survey)	1.66	−.23	−.07	.14	.39	.67	.88	.91	.80	.63	.42	.22
GNP/Hours (Household Survey)	0.87	.12	.23	.33	.47	.50	.51	.22	−.01	−.24	−.32	−.34
GNP/Hours (Establishment Survey)	0.82	.41	.47	.51	.53	.44	.32	−.06	−.30	−.47	−.50	−.49
Average Hourly Real Compensation (Business Sector)	0.93	.35	.39	.41	.43	.41	.35	.25	.16	.05	−.07	−.18
Real Employee Compensation (NIPA)/ Hours (Household Survey)	0.65	−.11	−.11	−.13	.06	.02	.10	.13	.14	.10	.08	.04
Real Employee Compensation (NIPA)	1.54	−.14	.00	.18	.41	.67	.88	.88	.76	.59	.38	.18
Employee Compensation (NIPA)/GNP		−.21	−.32	−.44	−.53	−.51	−.46	−.13	.13	.32	.39	.38

Source of basic data: Citicorp's Citibase data bank.

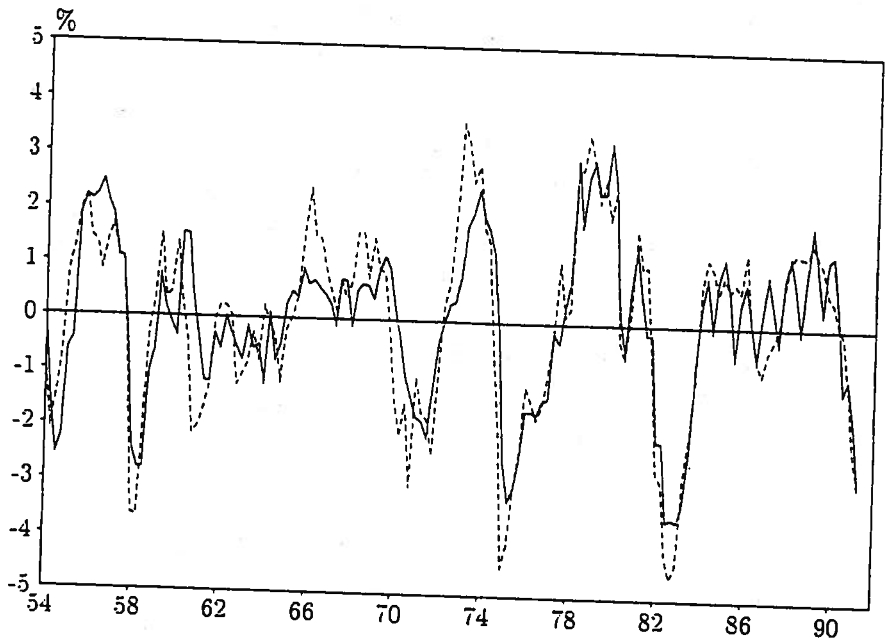

Figure 5.1 Total Hours (Household Survey) and Real GNP

Solid line shows hours and broken line shows real GNP. In Figures 5.1 through 5.6, data are quarterly from 1954:I to 1991:II and H-P filtered.

5) Average labor productivity is somewhat procyclical and leads the cycle. The degree of procyclicality is greater when output is divided by hours measured according to the household survey. The hours from the establishment survey indicate the longest lead: two to three quarters.

6) The statistics for average real hourly compensation in the business sector (which produces about 85 percent of GNP) are quite similar to those for productivity. If, on the other hand, we divide total employees' compensation from the national income accounts by total hours from either survey, series result whose correlations with real GNP are much lower.

7) Some writers have focused instead on the correlation of compensation (or productivity) with hours rather than with GNP (e.g., Christiano and Eichenbaum 1992). As a reflection mainly of the longer phase shift, the compensation series are less correlated contemporaneously with hours than with real GNP.

8) Real labor income is procyclical, but labor income as a fraction of GNP is countercyclical.

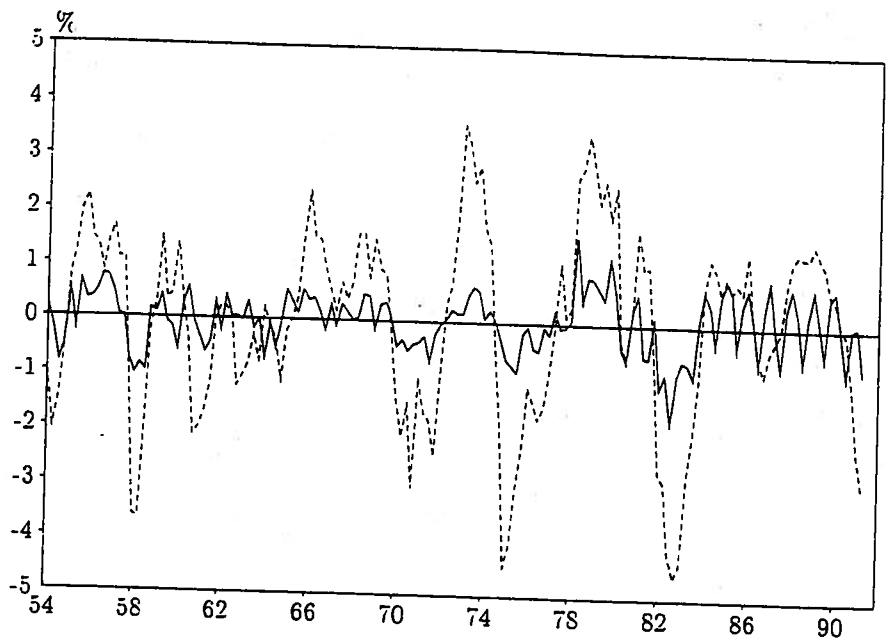

Figure 5.2 Hours per Worker and Real GNP
Solid line shows hours; broken line shows real GNP.

9) Over time, real hourly compensation has risen dramatically while hours worked per household has remained about constant. Cross-sectionally, however, there is a clear positive correlation between hours worked and the real wage. Moreover, the volatility of annual hours of work is much higher for wage earners in the two lowest quintiles than in the two highest (see Kydland 1984a; Ríos-Rull 1993b).

10) Benhabib, Rogerson, and Wright (1991) and Murphy, Shleifer, and Vishny (1989) argue that hours allocated to the production of consumption goods are procyclical. While direct observations based on a clear classification of the goods produced are not readily available, empirical evidence reported by Murphy, Shleifer, and Vishny points in that direction.

At various stages of the recent development of business cycle theory, some of these cyclical patterns have been regarded as deviations from existing theory. An application of real business cycle theory has been to address the question, How much of postwar business cycles would have remained if technology shocks were the only source of fluctuations? Major deviations along dimensions central to this question obviously could reduce one's confidence in the quantitative answer obtained. Through the interaction of theory and measurement, the deviations

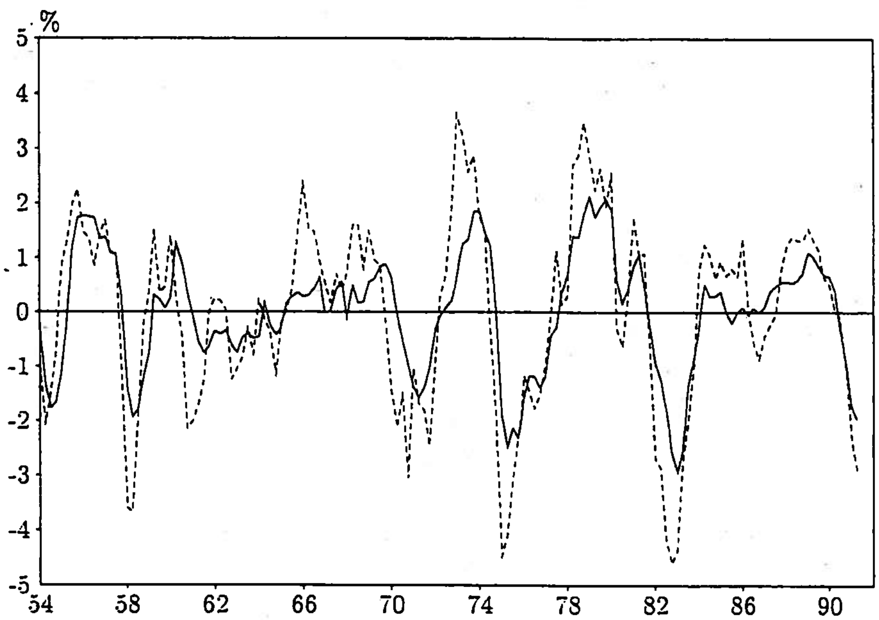

Figure 5.3 Total Employment and Real GNP
Solid line shows employment; broken line shows real GNP.

or anomalies relative to theory have led to stronger theory as well as to better measurements.

This chapter has two main objectives. The first is to give examples of the perceived deviations relative to theory, especially those related to labor market fluctuations, and of how researchers have attempted to resolve them. In the process, it will become clear that some of the proposed modifications still leave open important theoretical and measurement issues. The second objective is to present in detail an example of a model environment that is reasonably rich in its description of the labor market. It will incorporate movements of labor inputs in the forms of hours per worker as well as employment—both the intensive and the extensive margins.

In the next section, we present as a benchmark the standard neoclassical stochastic growth model, extended to include an explicit role for time allocation. It can be regarded as the starting point for the purpose of addressing business cycle questions. Then we review some of the developments in theory and measurement that have been motivated by perceived deviations from established theory. One such development is consideration of the use of nonmarket time in the household, possibly jointly with other inputs, to produce nonmarket goods. This is the subject of Section 3. Section 4 considers the fact that the work force consists of workers with a wide range of skills, whose behavior over the cycle differs substantially.

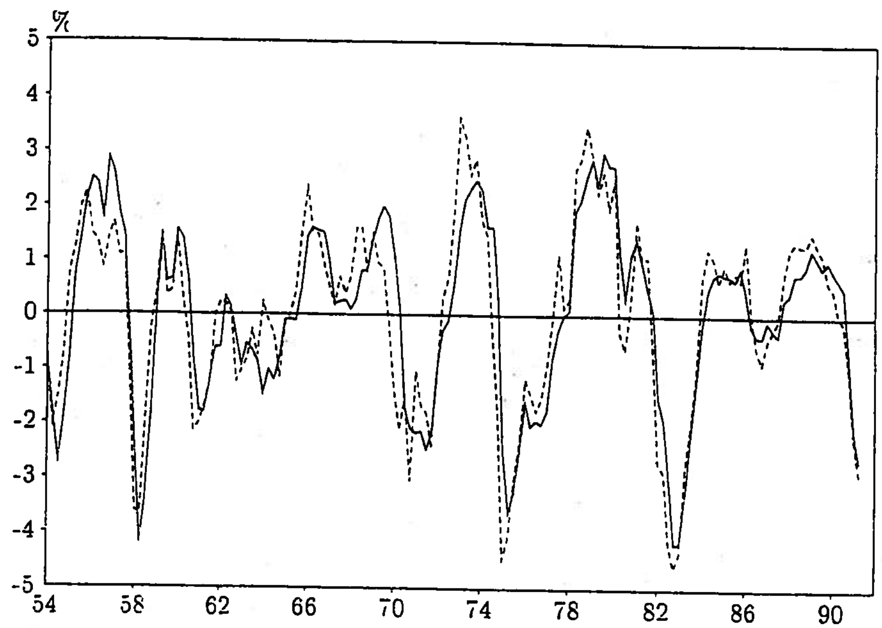

Figure 5.4 Total Hours (Establishment Survey) and Real GNP
Solid line shows hours; broken line shows real GNP.

This issue is discussed both from a modeling standpoint and from the perspective of measuring the labor input in aggregate production. The model formulations described in Sections 3 and 4 represent, with today's methods, tractable extensions of basic neoclassical theory.

Section 5 deals with the implications for the business cycle of the fact that labor input changes take the forms of both hours-per-worker and employment changes. The significance of introducing the employment margin became clear from the important paper by Hansen (1985) based on the theoretical insight of Rogerson (1984, 1988). The methodological foundation permitting the introduction of both margins has been developed only recently. A fundamentally new issue in this context is what shape the production function should take. In the business sector, the change of output associated with a given change of total hours in a given period surely is different when the change is in the number of hours a plant is being used rather than in the number of workers operating the plant.

In this chapter. several ways are presented in which the roles of market and nonmarket time for business cycles have been modeled. Section 6 provides a comparison of four of these in terms of the main business cycle characteristics. Section 7 contains an example of how one can extend one of these model economies (the one presented in section 5) to incorporate a new feature, in this case, learning by doing.

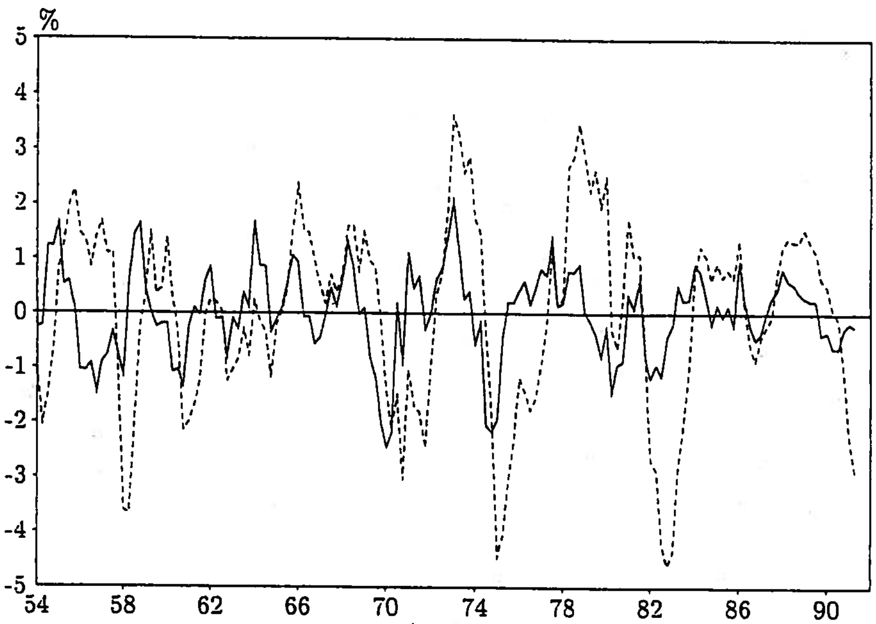

Figure 5.5 Average Productivity (Establishment Survey) and Real GNP
Solid line shows productivity; broken line shows real GNP.

Finally, in the last section we attempt an assessment of where we stand, particularly regarding the labor market's contribution to the propagation of shocks.

2. Basic Business Cycle Framework

Neoclassical growth theory has become the dominant theoretical framework in quantitative business cycle theory, as well as in most other areas of aggregate economics. It represents an environment that includes household and business sectors, and, for some questions, a government sector as well. The simplest growth model ignores time allocation decisions (see Stokey and Lucas with Prescott [1989, ch. 2] or Section 2 of Chapter 1 of this volume). A version that still is simple, but contains enough ingredients potentially to address business cycle questions, is as follows. The economy is inhabited by a large number of identical households, whose preferences are represented by a utility function:

$$E\sum_{t=0}^{\infty}\beta^t u(c_t, \ell_t),$$

where c_t is consumption, ℓ_t is time spent in nonmarket activity, or leisure for short, and β is the subjective discount factor. The production technology uses as

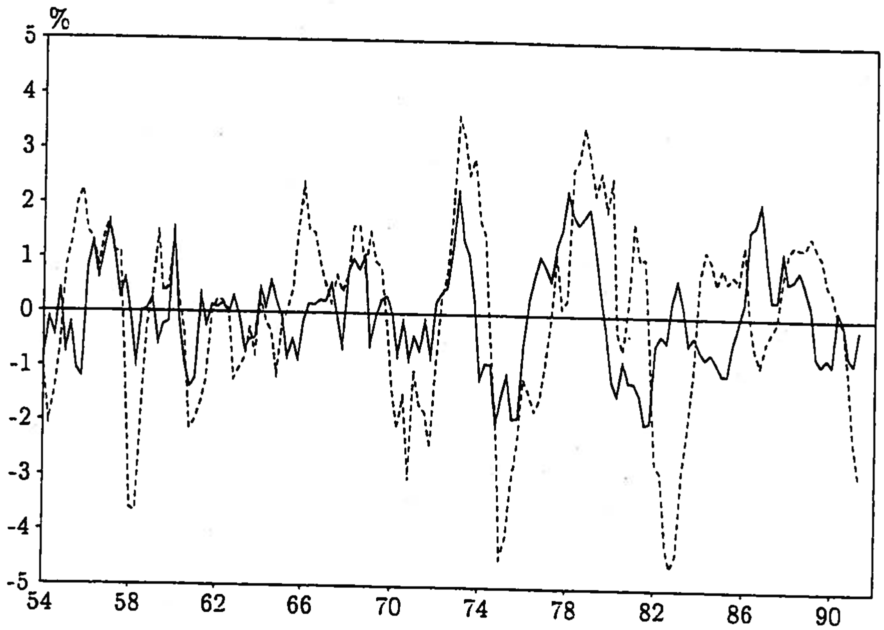

Figure 5.6 Average Hourly Real Compensation and Real GNP
Solid line shows hours; broken line shows real GNP.

inputs capital, k_t, and labor, h_t. There is perfect substitution in production between investment, x_t, and consumption. The constraints on the uses of output and time are

$$c_t + x_t \leq z_t f(h_t, k_t),$$

and

$$h_t + \ell_t \leq 1,$$

where, for simplicity, total discretionary time (net of sleep and personal care) is normalized to one. Laws of motion for the capital stock and technology are

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

$$z_{t+1} = \rho z_t + \epsilon_{t+1}, \tag{1}$$

where ϵ_{t+1} is a random disturbance with positive mean.

This framework departs from the simplest neoclassical growth framework in two ways. Leisure is included in the utility function, a feature from which models designed to address growth questions usually abstract. The emphasis on the time allocation decision distinguishes business cycle theory from growth theory. Another extension is the inclusion of stochastic technology shocks, which have been considered in the theoretical growth literature by Brock and Mirman (1972) and by Danthine and Donaldson (1981). With these features included, one could use the model to address, for example, questions about the role of technology shocks. Following Solow (1957), the z's can be measured as the residual in output variation after the capital and labor inputs have been accounted for. With the Cobb-Douglas production function,

$$f(h_t, k_t) = h_t^\theta k_t^{1-\theta} \tag{2}$$

one can write

$$\log z_t = \log \text{GNP}_t - \theta \log h_t - (1-\theta) \log k_t.$$

The value of θ corresponds to the average labor share in GNP. By studying the resulting series of z's, one can characterize statistically their persistence, as reflected in the parameter ρ in (1), as well as the volatility of the innovations, ϵ.

With explicit forms for the u and f functions and numerical values for the parameters of these functions and of the laws of motion, one can compute the solution in the form of decision rules for the variables c_t, n_t, and x_t. These decision rules, along with the laws of motion for the state variables, k_t and z_t, and the stochastic specification of the random shocks, can be used to perform computational experiments with the aim of yielding quantitative answers to business cycle questions.

A standard utility function is

$$u(c_t, \ell_t) = (c_t^\alpha \ell_t^{1-\alpha})^{1-\sigma}/(1-\sigma). \tag{3}$$

Here, the elasticity of substitution between consumption and leisure is one. In the general class of constant-elasticity-of-substitution (CES) functions, this is the only value consistent with the observation that in spite of a large increase in the average real wage over the past few decades, there has been little change in long-run hours per household in market activity. In a model of this type, this fraction of long-run time spent in market work typically turns out to be close to the value of the parameter α. Thus, with measurement of this fraction from data on individuals or households, its average value implies a value for α. Such time allocation measurements were reported by Ghez and Becker (1975), who, when defining the total discretionary time available for market and nonmarket activity, were careful to measure and to exclude time devoted to sleep and personal care.

Determining values of θ and α as well as those of the elasticities of substitution in the utility and production functions are examples of calibrating the model economy. The curvature parameter, σ, is harder to quantify with confidence. Studies of attitudes towards risk may suggest a reasonable range for this parameter.

Business cycle theory organizes quarterly national income and product accounts (NIPA) data. With this period length, however, it makes a difference that building new factories takes much longer than one quarter. Accordingly, Kydland and Prescott (1982) assume that the construction of productive capital in the business sector takes J quarters, where J on the average may be 3 or 4, but with resources used throughout the construction period. The law of motion for the productive capital stock then is

$$k_{t+1} = (1 - \delta)k_t = s_{1t}, \tag{4}$$

where the notation is to let s_{jt}, $j = 1, \ldots, J$ be capital (in units of finished capital) that is j periods from completion. Thus,

$$s_{j,t+1} = s_{j+1,t}, \qquad j = 1, 2, \ldots, J - 1. \tag{5}$$

The amount of resources used at each stage when building one unit of new productive capital is φ_j. Total investment, the sum of fixed investment and inventory investment, then is

$$x_t = \sum_{j=1}^{J} \varphi_J s_{jt} + y_{t+1} - y_t, \tag{6}$$

where y_t is the inventory stock at the beginning of period t. Including inventories is another way of extending the standard growth model. In a period with high productivity, for instance, people may wish to smooth consumption and carry into the subsequent quarter some finished goods in the form of inventories. Moreover, as motivated in Kydland and Prescott (1982), the inventory stock may be treated as an input in aggregate production. A specification of the resource constraint then is

$$c_t + x_t \leq [(1 - \gamma)(z_t h_t^{\theta} k_t^{1-\theta})^{-\nu} + \gamma y_t^{-\nu}]^{-1/\nu}.$$

With these features added, the model environment accounts quite well for the key properties of postwar U.S. business cycles, including relative volatility of investment and consumption, the procyclicality of most aggregates, and the contemporaneously uncorrelated capital stock. The model yields the preliminary estimate that technology shocks account for more than half of postwar U.S. business cycles. This estimate follows from computational experiments that use as an input the volatility of Solow residuals obtained for the U.S. economy; it is based on the fraction of U.S. output volatility implied by the model economy. This finding is supported by the model behavior of other aggregates, such as relative consumption and investment fluctuations. The key deviation relative to theory is that in this simple model with everyone working the same number of hours, the percentage standard deviation of the hours is substantially smaller than that of the model's real GNP.

3. Household Production

The realization that the empirical procyclical volatility of hours may be a problem for a general equilibrium theory of the cycle dates back at least to Lucas and Rapping (1969). Confronting this issue, they were led to the question, Are there reasons to substitute intertemporally, not captured by the standard specification of the household problem, that give rise to greater procyclical hours volatility? Lucas and Rapping suggest the theoretical possibility that future utility may depend, in part, directly on this period's choice of hours of work.

Kydland and Prescott (1982) make this idea operational and represent preferences in such a way that current utility is a function of a weighted average of current and past choices of nonmarket time:

$$u[c_t, \mu(L)\ell_t] = \frac{1}{1-\sigma}[c_t^{\alpha}(\sum_{i=0}^{\infty}\mu_i \ell_{t-i})^{1-\alpha}]^{1-\sigma}.$$

With weights summing to one, as can be assumed without loss of generality, their choice of parameter values was to let as much as one-half of the weight fall on current leisure ($\mu_0 = 0.5$), with the remainder spread over the past with geometrically declining weights. Thus, continuing with this numerical example, if the weights decline by 10 percent per quarter, then $\mu_1 = 0.05$, $\mu_2 = 0.045$, and so on. With that specification, the dependence of utility on current and past leisure choices is characterized by two parameters, μ_0 and ν, where ν is the rate of decline of the weights, that is, $\mu_{i+1} = (1 - \nu)\mu_i$ for all $i \geq 1$.

Kydland (1984a) interprets this utility function as a stand-in for household production, in which part of nonmarket time is used to accumulate household capital, which yields utility in the future. Examples of such capital may be quality of children, health, and perhaps the quality of the residence and other durable household property. The relatively large weight, α_0, on current nonmarket time then reflects the notion that a substantial portion of nonmarket time yields immediate pleasure. The remainder represents an investment in a form of household capital, which depreciates at a rate of ν. This analog of the specification above to the household production idea is exact when the two uses of nonmarket time are in fixed proportions and leisure and the durable home good are perfect substitutes in preferences. These are conservative assumptions. Relaxing them presumably would make market hours more responsive to changes in market opportunities.

With this feature included, not only does the volatility of hours increase relative to those of productivity and output, but technology shocks are also more potent in generating overall business cycle volatility. Referring to those findings, Kydland (1984b) concludes: "Using a standard time-separable utility function, about two-thirds of the fluctuations in the data were accounted for. If households are assumed to value leisure more if they have consumed less leisure in the past, the growth model explained nearly all."

This preliminary statement was not based on direct measurements of the volatility of the technology shocks. A more precise estimate uses Prescott's (1986) measurements. They are based on Solow's (1957) method for measuring technological change as the residual after the inputs have been accounted for. Subsequently, the resulting estimate of the volatility of the Solow residual has been used in computational experiments with a variety of model economies. The statistical properties of these residuals indicate that they are highly persistent—have high serial correlation. On the basis of these estimates, the "two-thirds" in the above quotation instead would have been 55 percent.

The fact remains that the quantitative importance of household capital formed by past nonmarket time can make a substantial difference to the estimate of the role of technology shocks. An attempt at assessing independently the magnitude of this form of household capital is made in Hotz, Kydland, and Sedlacek (1988). Using annual panel data for 482 men who in the first year of the twelve-year sample period were between the ages of 23 and 52, they estimate the parameters characterizing the role of household capital for life cycle behavior, taking into account differences in age, number of children, and other demographic factors. The estimates are consistent with the parameter values for μ and ν used by Kydland and Prescott (1982). It is probably fair to say, however, that this feature of household production has not been verified sufficiently by measurements to be regarded as necessary for a reliable estimate of the role of technology shocks for the cycle.

This formulation of time as an input into producing a form of household capital is simple and abstracts from the possibility that market-produced goods may be required as a joint input. The general idea, however, that attention to household production is important for understanding labor market fluctuations is an appealing one. It has been pursued in greater detail in two recent papers, by Benhabib, Rogerson, and Wright (1991) and by Greenwood and Hercowitz (1991). Both these papers consider the use in the household of physical capital (residential housing and/or consumer durables) that, along with nonmarket time, can be used to produce consumption goods. Greenwood and Hercowitz focus on the joint pattern of capital accumulation in the business and household sectors. Although this question has indirect implications for the labor input in market production, we shall leave a discussion of that topic for another chapter. Benhabib, Rogerson, and Wright, on the other hand, address issues that have a more direct bearing on the labor market. For example, they are motivated partly by the impression that fact (10) on the list in Section 1 represents a deviation from standard business cycle theory.

A key feature in the Benhabib, Rogerson, and Wright (BRW) model is the inclusion of a commodity made in the home using time and capital as inputs in a way analogous to the production of the market good. This home-produced consumption good is an imperfect substitute for market goods. Home production is a function of technology shocks in a manner analogous to that for the business sector.

The utility function is similar to (3), except that the variable c_t is replaced by a CES aggregator function representing a composite consumption good, which depends on c_{mt} and c_{nt}, where the subscripts m and n stand for market and nonmarket, respectively. Leisure in the utility function is net of time allocated to market and nonmarket production: $\ell_t = 1 - h_{mt} - h_{nt}$. Investment goods are produced in the market sector only. Capital can be moved between the two sectors. In practice, this reallocation takes place in the form of new investment. The home and business technologies both are Cobb-Douglas, with share parameters calibrated separately. The laws of motion for the technology shocks in the two sectors are identical, including a serial correlation coefficient of 0.95.

In addition to the motivation already mentioned, Benhabib, Rogerson, and Wright (1991) refer to measurements indicating that the fraction of nonmarket time devoted to production in the household is large. An interesting question, then, is whether household production may interact with market production in such a way that, combined, technology shocks to market and household production account for a considerably larger fraction of the business cycle than do those of the market sector alone. The answer hinges on three parameters—an elasticity of substitution in preferences, the standard deviation of the home technology shock, and its correlation with the business one. Indeed, a main contribution of the article is to demonstrate this fact. Since measurements of these parameters are either lacking or rudimentary at best, it underlines the importance of such measurements for a reliable answer.

Among those three parameters, the key one is the elasticity of substitution in preferences between market- and home-produced consumption, which BRW set equal to 5. This figure is based, in part, on estimates in Eichenbaum and Hansen (1990), according to which there is little statistical evidence against the hypothesis of perfect substitution elasticity between nondurables and durables. This empirical result can be interpreted as having a bearing on the model at hand. The estimate, however, is hard to reconcile with the observation that over time the price of durables relative to nondurables and services has fallen while the expenditure share has remained roughly constant. This fact would suggest an elasticity much closer to one than to infinity. Thus, one may doubt whether the elasticity in the BRW model will hold up under empirical scrutiny. Clearly, it plays a significant role for the model properties.

Other new parameters in the home-production specification are the standard deviation of the innovation to home technology and its contemporaneous correlation coefficient with the innovations in business sector technology. Again, good measurements upon which to base the values are not available. It seems much less likely, however, that the findings hinge upon future measurements of these parameters. For one thing, the authors make a good case for their reasonableness. Also, the theoretical findings appear to be not nearly as sensitive to variations within a moderate range of these parameters.

Benhabib, Rogerson, and Wright find that for their economy in comparison with the standard growth model, the volatility of output rises from 1.29 to 1.71 percent, which is very close to that observed for the postwar U.S. economy. The volatility of hours in relation to that of GNP, $\text{std}(h_m)/\text{std(GNP)}$, rises from 0.50 to 0.75, where std is short for percentage standard deviation.

The correlation in the model between real GNP and hours spent producing consumption goods in the market sector is 0.10. This magnitude may strike the reader as quite disappointing. One contribution of the article, however, is to show that this correlation can be turned from a large negative value to this slightly positive value simply through the introduction of household production. The simplicity of the model environment in other respects accounts for the negative correlation in the benchmark model. There are several reasons. Most important perhaps is the omission of inventories. Changes in business inventories have been procyclical and highly volatile, and a large part of those changes in every quarter has been in consumption goods. For instance, if inventory changes were divided between consumption and investment goods in the same proportion as are their average fractions of GNP (about three-fourths and one-fourth, respectively), then a standard business cycle model without explicit household production, such as that described at the end of Section 2, would imply a positive correlation between real GNP and the hours spent producing consumption goods. With the introduction of the BRW household production function in that environment, the correlation would presumably be substantially higher than the BRW model's 0.10. A numerical comparison is discussed in Section 6.

Another issue is whether the statistics that serve as a basis for fact (10) in Section 1 include consumer durables. Empirically, this aggregate shares many of the properties of business investment: it is highly volatile and strongly procyclical. Most model economies abstract from consumer durables and, one may argue, cannot hope to produce very procyclical hours in the consumption sector. In the BRW economy, consumer durables are, to a larger extent, the empirical counterpart to household capital, yet the hours spent producing them are not allocated to the consumption sector. For the BRW benchmark parameter values, household investment is strongly countercyclical. This fact leaves some doubt as to how much has been resolved with regard to accounting for the procyclical hours in the consumption sector.

We have discussed two approaches to modeling household production, each of which may have an important bearing on labor market fluctuations. The first emphasizes the use of nonmarket time to accumulate a durable, which is not necessarily tangible, in the home sector. The other approach is to think of nonmarket time as being combined with tangible market-produced durables to produce another consumption good. In either case, if these features can be shown to be of quantitative importance, they will help to account for a considerable part of output and, in particular, hours variability. Both cases share the characteristic, however, that the underpinnings in the form of measurements are still shaky.

Another model motivated by home production, or by the interaction between home and market production, is presented in Cho and Cooley (1994). Their idea is that a fixed cost is associated with each day when people work. This cost can be motivated, in part, by the notion that some home production, such as child care, needs to be replaced. Moreover, Cho and Cooley assume an externality in the sense that this cost depends on the aggregate number of workers. They then show that introducing this feature potentially can lead to a substantial increase in the volatility of hours relative to that of productivity. Strict calibration of the model to micro observations gives less encouraging results, however, and the authors conclude that the evident deviation shows that some important feature still is missing from their model environment.

4. High- and Low-Wage Earners

Theory

The model environments discussed so far assume that all workers are homogeneous. If there are substantial differences in cyclical behavior across demographic groups, then this assumption could bias considerably the estimate of the role of technology shocks. As fact (9) in Section 1 indicates, an example of such a difference is the greater hours volatility of the low-wage earners as compared with high-wage workers.

A simple way to introduce heterogeneity in this class of economies is to divide the model population into groups according to skills. Kydland (1984a) considers two equal-sized groups, where the first is better skilled for market production than the second. The resource constraint then can be written as

$$c_{1t} + c_{2t} + x_t \leq z_t f(h_t^e, k_t),$$

where c_{1t} and c_{2t} are consumption by the high- and low-skilled workers, respectively, and $h_t^e = \omega h_{1t} + h_{2t}$ is total labor input measured in quality-weighted units. If we divided the work force in two according to skills and used average per-person labor compensation to compute this weight, the numbers in Kydland (1984a) or in Ríos-Rull (1993b) suggest a value for ω of 2 or higher. The equilibria studied are those corresponding to the Pareto problem of maximizing the weighted utilities of the two groups (see Negishi 1960). The weights are calibrated to yield average hours spent in market activity comparable to those in the U.S. data when workers are divided into two similar-sized groups according to human capital.

The associated equilibrium is such that average consumption less labor income is greater for the low-wage earners than for those with high wages. Steady-state aggregate consumption has to satisfy the constraint

$$c_1 + c_2 = w_1 h_1 + w_2 h_2 + b_1 + b_2$$

where b_i stands for net nonlabor income for workers of type i. It consists of net capital income as well as any net transfers, τ_i, from the other skill group, that is, $\tau_1 = -\tau_2$. Thus, total steady-state nonlabor income, $b_1 + b_2$, is simply the real interest rate, r, multiplied by the capital stocks. For each skill group, b_i is defined so that $c_i = w_i h_i + b_i$. For the equilibriums reported in Kydland (1984a), the steady-state magnitude of b_2 exceeds that of b_1 by nearly 4 percent of GNP. Given what we know about relative capital income for the two groups, this means that some of b_2 has to be a transfer from the high-wage to the low-wage earners. In view of the amounts of such transfers that take place through the government as well as within the household, this magnitude does not appear unreasonable.

The paper compares the case in which ω, the wage of the skilled workers relative to that of the unskilled, is constant with the case in which this relative wage is allowed to move countercyclically by a small amount, say, with a standard deviation of .25 percent. A finding is that in the latter case, the standard deviation of aggregate unweighted hours rises by more than 20 percent relative to that of productivity. On the other hand, the fraction of output volatility accounted for by Solow residuals declines by about 10 percent. The model's cyclical relation between the relative wages of the skill groups is consistent with that reported by Reder (1962), although it would be interesting to have this empirical regularity investigated again using more recent, perhaps higher-frequency, data. Intuitively, it seems reasonable that the high-skilled workers are more adaptable in recessions, but that the skills of some, such as certain engineers, become obsolete in periods of rapid technological advance. There are, of course, numerous microstudies of the interaction in production of such categories as white- and blue-collar workers or workers with different levels of training.

This model economy introduces heterogeneity in a way that makes it tractable within a framework with infinitely lived agents. It illustrates a channel through which skill differences may have a bearing on the role of technology shocks for the cycle in general and for the implied volatility of hours of work in particular. A sharper assessment of this importance will depend on measurements such as those suggested in the preceding paragraph. Moreover, since the equilibria studied require transfers from the skilled to the unskilled of particular magnitudes, the reliability of the findings may depend on the presence of similar magnitudes in the actual economy. Although sizable transfers clearly do take place, their exact quantities are not easy to determine for the appropriate classification of people.

We have described an environment with the population divided into two different infinitely lived groups. It abstracts from life cycle behavior, for instance. Such behavior for mortal consumers can be built into an overlapping-generations framework. Until recently, however, it was difficult to see how one could calibrate such models while at the same time maintaining computational tractablity.

Economists' perspective on the feasibility of using aggregate equilibrium models with life cycle behavior now has changed, in part as a consequence of research that develops further the quantitative-theoretic approach pioneered for such models by Auerbach and Kotlikoff (1987).

In the past few years, Ríos-Rull has led the way in developing and using overlapping generations models in order to obtain quantitative answers to a variety of questions. Of particular interest in our context is his paper (1993b) on the interaction between household production and the choice of whether or not to become better skilled for market production. His paper is motivated to a large extent by fact (9) in Section 1. The driving forces are the presence of a home-produced good with poor market substitutes, and the possibility of choosing whether or not to acquire skills through schooling. In this model economy, meaningful heterogeneity arises even though everyone is born alike. The model accounts well for some of the key movements both cross-sectionally and secularly. Cyclically, however, a remaining discrepancy is that in contrast to the U.S. data, in the model, the hours volatility of the unskilled workers does not exceed that of the skilled.

Measurement

An alternative to modeling explicitly the heterogeneity of workers in terms of skills for market production is to take account of these differences in the measurements to which models are compared. Given the central role played by the production function for aggregate theory in general and for business cycle theory in particular, an important question is, How reliable are the available measurements of the labor input? For output and its components, the principles behind the measurements are those pioneered by Kuznets (1946a) and Stone (1947) for national income and products accounts. According to these principles, steady-state or base-year prices are used to weight the different goods being summed up to form the aggregate real quantities. A similar approach is used for the capital stock. The difficulties for capital are perhaps even more severe, as the capital controversy between the two Cambridges illustrated. It is clear, however, that while Cambridge, England, was right in theory, Cambridge, U.S., prevailed in practice. The capital stock measurements have contributed to the important developments and insights in growth theory in the past thirty or forty years (see Solow 1970). In contrast, the same NIPA principles typically have not been applied to the measurement of the labor input. Standard practice is to give equal weight to the hours of all workers, including people with dramatically different stocks of human capital. If the cyclical behavior of these workers differs widely, then the standard procedure of simply adding up the hours may produce a poor measure of the labor input.

From the viewpoint of a theory in which the production function is a central feature, it is natural to think of the labor input in efficiency units. One would then like to weight the hours of different individuals by their relative base-year prices in the same way that other NIPA quantity data are constructed. The urgency of this task is demonstrated in Kydland (1984a). Using data from the Panel Study of Income Dynamics (PSID) on about 1000 men over the age of 30, which presumably is the least volatile major category of the labor force, and dividing the subjects into five nearly equal groups according to years of schooling, he estimates that over the eleven-year period of the sample, those with the least formal education changed their annual hours on the average by about 100 hours more for each percentage point change in the unemployment rate than did those with the most formal education.

A more detailed study of this measurement issue is reported in Kydland and Prescott (1993). Using a sample from the PSID of nearly 5,000 people consisting of all major demographic groups, they compared the cyclical behavior of two alternative measures of the labor input as follows. Let N_t be the number of people in the population in year t, and let h_{it} be person i's hours of work in that year. The standard measure is simply to add up, in each period, the hours across all workers: $H_t = \sum_i h_{it}$. Another measure is to multiply the hours of each individual by relative human capital weights that do not change cyclically: $L_t = \sum_i \phi_i h_{it}$. For the sample period, there was little secular change in average real compensation per hour. Therefore, a fixed relative weight for person i was constructed by dividing his or her total real labor earnings over all the years by total hours worked in those same years; that is, $\phi_i = \sum_t e_{it} / \sum_t h_{it}$, where e_{it} is real labor earnings of individual i in year t, and the summations are over all the years of the sample period for which observations for that person were available. This measure of the worker's "normal" efficiency is used in every period as the stand-in for his or her relative efficiency in market production.

The finding is that if the sample were representative for the entire population, the standard measure of labor input would overstate the labor input volatility by about 40 percent. This is a large number from the standpoint of business cycle theory. Another finding is that the real hourly compensation of the quality-adjusted labor input is more procyclical than the corresponding average compensation per unweighted hour.

5. Hours versus Employment Behavior

Indivisible Labor

An important development in the understanding of hours volatility was the article by Hansen (1985). In the models discussed so far, all the variability in hours takes place in the form of changes in hours per worker. Hansen went to the opposite extreme. In his environment, all the labor-input volatility takes the form of employment changes. There is a fixed cost of working, with the implication that everyone works either zero hours or some positive number h_1.

As an illustration, assume that the utility function is logarithmic (corresponding to $\sigma = 1$ above):

$$u(c_t, \ell_t) = \log c_t + \alpha \log \ell_t$$

To get around the nonconvexity implied by the binary choice of hours of work, assume instead that individuals choose the probability π_t of working. In other words, a contract to work h_1 hours with probability π_t and 0 hours with probability $1 - \pi_t$ is traded between workers and firms. This means that workers get paid whether they work or not. (Hansen discusses in an appendix the interpretation in terms of insurance).

Individuals are identical ex ante, but the ex post outcome in every period depends on the lottery. Expected utility is

$$\begin{aligned} U(c_t, \ell_t) &= \pi_t[\log c_t + \alpha \log(1 - h_1)] + (1 - \pi_t)(\log c_t + \alpha \log 1) \\ &= \log c_t + \alpha \pi_t \log(1 - h_1). \end{aligned}$$

Per capita hours worked are simply $h_t = \pi_t h_1 = 1 - \ell_t$, implying that $\pi_t = (1 - \ell_t)/h_1$. Substituting this expression for π_t into the utility function, we obtain the representative individual's utility function:

$$U(c_t, \ell_t) = \log c_t - \frac{\alpha \log(1 - h_1)}{h_1} \ell_t + \text{constant}.$$

In other words, the planner's utility function is linear in ℓ_t. Thus, the startling finding is that the intertemporal elasticity of substitution in the aggregate can be very large even though, as a property of each individual's utility function, this elasticity has the much smaller value associated with the logarithmic utility function. On the basis of this model economy, Hansen found that Solow residuals could produce business cycles even more volatile than those observed in the postwar U.S. economy.

With the extreme assumption that the employment margin is where all the hours variability takes place, the implied estimate naturally overstates the role of technology shocks for the cycle. An economy that permits variation along both margins—employment and hours per worker—presumably would yield an estimate somewhere in between those of Hansen's model and those of a model with only hours-per-worker variation. Such an estimate is provided in Kydland and Prescott (1991).

Two Margins

The goal here is to construct a business cycle model in which there is variation in labor input along both the hours-per-worker margin and the employment margin. In order to provide a credible estimate of the role of technology shocks, this model ought to mimic to a reasonable degree facts (2) and (4) in the list in Section 1.

In this economy, the obvious analogue of the standard production function is $z_t f(h_t n_t, k_t)$, where n_t is the number of workers and h_t is hours per worker. This production function implies that the marginal product of labor input is the same no matter which of the two forms the change takes. A better assumption is that a fixed number of workers are assigned to each machine or, more generally, to each unit of capital input. Adding workers to a fixed stock of capital then reduces the marginal product in the usual way, while letting the existing workers operate the machines longer hours would, to a reasonable approximation, increase output in the same proportion; the production function is $z_t h_t f(n_t, k_t)$.

Another issue is how to deal with the labor indivisibility analogous to that in Hansen's economy. The economy still is inhabited by a large number of ex ante identical individuals, although some will not work ex post in every period. Some preliminary insight can be gained from a related one-period example from Hornstein and Prescott (1993).

Each agent is endowed with $\bar{k} > 0$ units of capital. Preferences with respect to consumption-work pairs, (c, h), are represented by their expected utility, $E[u(c, h)]$, where h is the fraction of time allocated to market activity. For simplicity, we assume that $s = (c, h, k)$ is a member of $S = C \times H \times K$, where C, H, and K are finite sets. In practice, these sets could be constructed as a grid of values in the relevant range for each of the variables. For each individual, the commodity bundle is interpreted as a contract that obliges him or her to provide k units of capital and h units of time, for which he or she receives c units of the consumption good. The probability of an event $s = (c, h, k)$ is x_s.

In the business sector, add the finite set N, and let $A = H \times K \times N$ with elements of the type $a = (h, k, n)$. The choice is how many plants z_a to operate for h hours using k units of capital and n workers. An allocation satisfies the resource constraints if

$$\sum_s c x_s - \sum_a h f(n, k) z_a \leq 0,$$

$$-\sum_s k x_s + \sum_a k z_a \leq 0.$$

and

$$-\sum_{c,k} x_s + \sum_{k,n} n z_a \leq 0 \quad \text{for all } h \in H.$$

The first constraint says that the amount consumed is less than or equal to the quantity produced. According to the second constraint, the quantity of capital used in production cannot exceed the quantity available. The last constraints (one for each value of h) say that the number of people working in plants that are operated h hours does not exceed the number of people working h hours.

For this economy, as shown by Hornstein and Prescott (1993), the competitive equilibrium can be obtained by solving a stand-in Pareto problem. This problem is a linear programming (LP) problem with the x_s as variables:

$$\max_{x\geq 0} \sum_s u(c,h)x_s$$

$$\text{s.t.} \quad \sum_s x_s = 1$$

$$\sum_s [c - f(1,k)h]x_s \leq 0$$

$$\sum_s kx_s \leq \bar{k}.$$

A general property of the solution to an LP problem with three constraints is that at most three variables are positive. That is, there are no more than three combinations of $s = (c, h, k)$ such that $x_s > 0$.

Now consider the utility and production functions given by (3) and (2), respectively, with standard parameter values. It turns out that when the grids of the points in S are made successively finer, the solutions to the corresponding planner's problems tend to cluster in such a way that at least two of the points that receive positive x_s get closer and closer. As Hornstein and Prescott (1993) show, this pattern reflects the property that when the sets C, H, and K contain infinitely many points (S is a subset of $\mathbb{R}^3_+$), then the solution to the LP problem implies mass on either two points or only one point depending on the parameter values for the utility and production functions.

When the equilibrium consumption vector places mass on only one point ($x_s = 1$ for some $s = s_1$), it is of the form $s_1 = (c_1, h_1, k_1)$. Since $h_1 > 0$, everyone works the same number of hours. When there is mass on two points, s_0 and s_1, then the value of h_0 in s_0 is zero. Thus, some fraction of people work h_1 hours and receive consumption c_1, while for everyone else h_0 is zero and consumption is c_0.

Business Cycle Model

We shall now embed an analogous structure within a fully dynamic business cycle model. This model will be calibrated to correspond to that with mass on two points. The variable n_t will be the fraction of people who work in period t. A person working h hours and using k units of capital produces $zhk^{1-\theta}$ units of some intermediate good. This good, along with inventory services, y, is an input to a CES production function.

For this economy, the aggregate resource constraint in period t is

$$c_t + x_t + m_t \leq [(1-\gamma)(z_t h_t n_t^{\theta} k_t^{1-\theta})^{-\nu} + \gamma y_t^{-\nu}]^{-1/\nu} \tag{7}$$

where m_t is the aggregate cost of moving people between the market and nonmarket sectors. This cost will be approximated by a quadratic function, $m_t =$

$\mu(n_t - n_{t-1})^2$. As suggested by Kydland and Prescott (1991), this specification is a stand-in for an environment in which there is a distribution of moving costs across the population, and those with the smallest cost are moved first. The moving-cost distribution is independent over time. This formulation gives rise to labor hoarding in this economy.

The cost of getting to work every day may also play a role. Most of that cost probably is in the form of time that neither is enjoyed as leisure nor contributes as an input in the production of goods. Such a cost is allowed for in the original model. Although it affects the calibration somewhat, it makes little difference to the cyclical properties, and we ignore it here.

As in Section 2, the inventory stock is included as an input. This assumption is made partly for analytic reasons. One can then ignoore the nonnegativity constraint for inventories and use a linear-quadratic economy. The assumption that larger inventories economize on the other two inputs can be justified in several ways. For example, by making longer production runs and thus holding larger inventories on the average, firms reduce equipment downtime associated with shifting from producing one good to producing another. For this economy, the observed procyclical behavior of the aggregate inventory stock is mimicked reasonably well.

The remainder of the model specification is analogous to that in Section 2. The laws of motion for finished and unfinished capital stocks are given by (4) and (5), and total investment is given by (6). Finally, we use the law of motion given by (1) for the technology level.

An implication analogous to that in Hornstein and Prescott (1993) is that the equilibrium can be computed by solving a social planner's problem:

$$\max E \sum_{t=0}^{\infty} \beta^t [n_t u(c_{1t}, 1 - h_t) + (1 - n_t) u(c_{0t}, 1)]$$

$$\text{s.t.} \quad c_t = n_t c_{1t} + (1 - n_t) c_{0t}$$

and to the constraints just mentioned. The utility function, u, is the standard one given by (3), where the elasticity of substitution already has been calibrated to equal one for reasons discussed in Section 2.

Steady State and Calibration

The steady state for this economy is its deterministic rest point, that is, the point resulting when the variance of the shock is zero. The steady state is important for two reasons. First, since this highly nonlinear model will be represented by a quadratic approximation, the steady state represents the point about which this approximation is made. More important, however, the properties of the steady state for the model economy correspond to analogous long-run relations in the actual economy that in many cases can be measured with high signal-to-noise ratios and are used in the calibration.

Some relations do not require much analysis of the model. Examples are NIPA relations for the model environment. Without loss of generality, we choose units such that steady-state output is one. Steady-state consumption and investment shares of GNP are set to 0.75 and 0.25, respectively. For the United States in the postwar period, the inventory stock has been about a quarter of annual GNP. Thus, we set $y = 1$. Steady-state n corresponds to the long-run fraction of the working-age population who actually work and is taken to be 0.75, while h, the steady-state fraction of time spent working, conditional on being in the market sector, is 0.40. As an average of the entire population of the model economy, then, the time spent in market activity is 0.30, or just over thirty hours per week. This is a standard magnitude for this relation and in line with the measurements by Ghez and Becker (1975).

The elasticities of substitution between consumption and leisure in utility and between capital and labor in production have been discussed already. Both equal one. There is less clear-cut evidence on which to base the value of the elasticity $1/(1+\nu)$ between inventories and the composite input. It is probably quite small, and ν is therefore probably substantially larger than zero. We choose $\nu = 3$. If the question dictates it, one should of course investigate the robustness of the answer to this choice.

A value for J of 3 (quarters) is a reasonable compromise. Some capital, of course, takes more time, and some less, to build. There is little evidence that the average time to build varies over the cycle. We assume that the resources needed are used up evenly throughout the construction period, that is, $\phi_j = 1/J$ for all j. The evidence is that the yearly depreciation rate is in the range of 8–10 percent. Since we assume no growth, we shall use the upper end of this range and assume that $\delta = 0.025$. This value, along with an investment share of output of 0.25, corresponds to a yearly capital/output ratio of 2.5 ($k = 10$). Also, with no growth, the steady-state real interest rate, r, equals $(1-\beta)/\beta$. A value for r of 0.01 per quarter implies that β is approximately 0.99.

Before we consider the remaining parameter values, we need to derive the steady-state implications of equilibrium behavior for the model environment. For this purpose it is convenient to work with the decentralized problems of the household and of the firm separately. (For a discussion of decentralization of the standard growth model, see Chapter 1 of this volume.) We think of firms as being owned by the households, and the input factors as being rented or hired from these same households. For either problem, we initially take hours per period, h, as given. The remaining decision variables for the firm, then, are n, k, and y, and those for the household are c_0, c_1, and n. In the end, we determine h from the equilibrium condition that the marginal product of working h hours equals the negative of the ratio of marginal utilities with respect to hours and consumption.

The Firm's Problem

The firm is endowed with a technology whereby it uses labor, capital, and inventories as inputs to produce output of goods and services. Defining q_k and q_y to be the rental prices of capital and inventories, respectively, and $w_h = wh$ to be a worker's real earnings per period conditional on working h hours, the firm maximizes in every period

$$F(zhn^{\theta}k^{1-\theta}, y) - q_k k - q_y y - w_h n.$$

In the steady state, the equilibrium q_y equals r and, with no additional time to build (that is, with $J = 1$), the rental price of capital would be $r + \delta$. For multiple-period construction ($J > 1$), however, the real price, p_k, of newly produced capital exceeds one because resources are tied up during the construction period. Defining the prices of s_j, the capital being built, to be p_j, for $j = 1, \ldots, J-1$, we must have $p_{J-1} = \phi_J$. The other prices are determined recursively as

$$p_{j-1} = (1+r)p_j + \phi_j, \qquad j = 2, \ldots, J-1.$$

The equilibrium steady-state price of a unit of productive capital, then, is

$$p_k = \sum_{j=1}^{J} \phi_j (1+r)^{j-1},$$

implying a steady-state rental price of $q_k = (r+\delta)p_k$.

Units in which to measure output, such that its steady-state quantity is one, are chosen by selecting the average z appropriately. To turn to the inventory decision, the condition $F_y = q_y$ yields

$$\gamma = q_y y^{\nu+1}$$

Similarly, from the condition $F_k = q_k$ one obtains

$$1 - \theta = q_k k/(1 - \gamma y^{-\nu}) = q_k k/(1 - ry).$$

That is, the parameter $1 - \theta$ equals the capital share of income net of the income share of the inventory input. Thus, both γ and θ are quantified from relations between variables or parameters whose values we already have determined. In particular, γ equals 0.01 (implying that 1 percent of the model's national income can be attributed to inventories), and θ is approximately equal to 0.64. Finally, the wage rate w, which is a parameter of the household's problem, is implied by $w_h = wh = F_n$.

The Household's Problem

The household's problem treats the capital income parametrically. Steady-state net capital income is

$$b = q_y y + (q_k - \delta)k,$$

which also can be written as

$$b = ry + rp_k k + rp_1 \delta k + rp_2 \delta k,$$

that is, the interest rate times each of the values of the four capital stocks. Given this steady-state net capital income, the household maximizes discounted utility subject to an infinite-period budget constraint. The resulting values of the variables c_{0t}, c_{1t}, n_t and h_t clearly are date independent. Consequently, we can drop the time subscripts. The steady-state problem of the household then can be written as

$$\max\ (1-n)u(c_0, 1) + n \cdot u(c_1, 1-h)$$

$$\text{s.t.}\quad (1-n)c_0 + nc_1 \le whn + b.$$

Maximization yields first-order conditions with respect to the variables c_0, c_1, and n. Moreover, hours per worker, h, has to satisfy the condition $-u_h/u_{c_1} = F_h$. These four conditions and the budget constraint determine the Lagrange multiplier along with four additional unknowns. These four will be α and σ from the utility function, and c_0 and c_1. The resulting values are $\alpha = 0.29$, $\sigma = 2.41$, $c_0 = 0.57$, and $c_1 = 0.81$. We note that in the steady state, those who work consume about 40 percent more than do those who are not in the market sector.

The value of σ warrants a comment. This value is larger than the value of 2.0 used in Kydland and Prescott (1991) and results mainly from a lower calibrated value of h, namely, 0.40 rather than 0.44. With a total time allocation of about 100 hours per week, the value of 0.44 probably was a little too high. It may be easier to think about σ in relation to the empirical finance literature if we multiply $1 - \sigma$ by α, thus obtaining the overall exponent on c in the utility function. This exponent (whose value here is -0.4) should be comparable conceptually to what is used in finance studies that abstract from the time allocation decision, so that the implied degree of relative risk aversion is in the ballpark of what those studies find.

6. Cyclical Properties of Model Economies

The purpose of this section is to compare cyclical properties of four of the economies that we have discussed:

1) a homogeneous-worker economy similar to that in Kydland and Prescott (1982), but with standard utility function;

2) as in economy 1, but with part of nonmarket time used to produce a durable household good;

3) as in economy 1, but including a household technology for using capital and labor as inputs to producing consumption goods (similar to the BRW model); and

4) as in economy 1, but with two margins for changing the labor input, as described in Section 5.

All four environments include inventories in the same ratio to GNP. It takes three quarters to build new productive capital. Other sources of calibration that are common to these economies also are assigned the same values. These magnitudes are presented and motivated in the preceding section.

The differences in calibration across economies are as follows. In economies 1, 2, and 3, the fraction of time devoted to market activity is 0.3, as in section 5, but all in the form of h, since by assumption n is one. In economy 2, the magnitude of μ_0 is set equal to 0.60, which gives slightly more weight to current leisure in the utility function than in Kydland and Prescott (1982). The depreciation rate η for household capital equals 0.10. In economy 3, the parameters of the aggregator function for consumption in the utility function and those of the household technology are assigned the same values as in the BRW model. In other respects, the economy is analogous to economy 1. For example, it includes the same curvature parameter σ, which is greater than that used by Benhabib, Rogerson, and Wright (1991), who employ a logarithmic utility function.

The statistics on which we focus, in addition to output and its two main components, are those corresponding to the aggregates listed in Table 5.1. They are summarized in Table 5.2, borrowing the format in BRW. The notation h_c represents the hours spent producing consumption goods in the market economy, while c_m denotes consumption goods produced in the market economy. This distinction is relevant only for economy 3.

In the simplest version of the growth model, as modified in Section 2, the standard deviation of cyclical output is 1.25 percent. Introducing household capital produced solely by leisure raises the figure to 1.39 percent. The increase in hours volatility is substantially greater, however, while productivity volatility is lower, so that for economy 2, hours volatility actually is larger than that for productivity.

The household technology shock evidently has the potential to account for a substantial fraction of the business cycle. The comparison of economy 3 with economy 1, where the introduction of the household technology is the only difference, indicates a rise in output volatility from 1.25 to 1.60 percent. Moreover, productivity becomes substantially less correlated with the cycle.

The introduction of a distinction between employment and hours-per-worker variation, along with the modified production function in (7), raises the standard deviation of output from 1.25 to 1.55 percent. The latter figure was produced with

Table 5.2
Statistical Properties of Model Economies

	$x =$:					
	c_m	i	h_m	n	$GNP/h_m n$	h_c
	Model Economy 1: std(GNP) = 1.25					
std(x)/std(GNP)	.40	2.49	.41		.60	
Corr.(x,GNP)	.97	0.95	.99		.99	
	Model Economy 2: std(GNP) = 1.39					
std(x)/std(GNP)	.37	2.57	.53		.49	
Corr.(x,GNP)	.95	0.95	.98		.98	
	Model Economy 3: std(GNP) = 1.60					
std(x)/std(GNP)	.66	2.59	.69		.46	.82
Corr.(x,GNP)	.73	0.90	.91		.79	.48
	Model Economy 4: std(GNP) = 1.55					
std(x)/std(GNP)	.43	2.61	.20	.46	.47	.28
Corr.(x,GNP)	.98	0.95	.75	.86	.97	.17

the same value of the standard deviation of innovations to technology as in the other experiments. Allowing for variable capacity utilization, however, means that the standard expression for determining the Solow residuals no longer is theoretically correct. A way of checking the size of the bias is to use the standard method in the model economy to see if the variance estimate is different from the variance of ϵ (0.0076^2) used as input to the experiments. The resulting bias suggests that the estimate of the standard deviation for economy 4 should be reduced from 1.55 to 1.49.

For economies 3 and 4, we have computed the statistics for h_c, hours devoted to the production of consumption goods. This variable, which in part motivated the BRW model, no longer has a straightforward definition because of the presence of inventory changes. A considerable fraction of these changes presumably are in the form of consumption goods. The assumption made in Table 5.2 is that in every period the fraction of inventory change that is in the form of consumption goods is the same as that in final sales. This is probably a conservative assumption.

Then even economy 4 implies procyclical h_c, indeed with a greater correlation coefficient with cyclical GNP than in the BRW model. But for the modified BRW economy, our economy 3, this correlation coefficient is as high as 0.48. Had the model economy included market-produced consumer durables in a way implying that they were procyclical as in the data, then an even larger correlation coefficient presumably would result. Thus, it seems safe to say that fact 10 in the list in Section 1 no longer can be regarded as a fact from which the theory deviates.

7. On-the-Job Learning

In constructing a model environment with heterogeneous workers, Kydland (1984a) assumes that the division of human capital between the two groups is given. That assumption precludes consideration of issues that relate to the timing of the accumulation of human capital over the cycle. As Mincer (1962, S73) concludes, "Investment in on-the-job training is a very large component of total investment in education in the United States economy." Human capital of this form thus is large enough that abstracting from its accumulation when evaluating the role of technology shocks, one risks omitting a potentially important propagation mechanism. One may guess a priori that introducing on-the-job training will change the cyclical properties of several aggregates, perhaps of labor input and productivity variables in particular. The main question, however, is to what extent the estimate of the cyclical role of technology shocks is affected.

An example of a tractable specification is to assume that workers enter the labor force at the lowest efficiency level and accumulate skills through the process of learning for I periods. Let e_{it}, $i = 0, \ldots, I$, be the number of workers at efficiency level i at time t, where e_{0t} represents the bottom of the skill distribution. Consider the following laws of motion:

$$e_{i+1,t+1} = (1-\eta)e_{it}, \qquad i = 0, \ldots, I-2,$$

and

$$e_{I,t+1} = (1-\eta)(e_{I-1,t} + e_{It}).$$

In other words, a fraction η of the workers at each level lose their previously accumulated skills or "die." In the steady state, a corresponding number reenter at the inexperienced level. The total number of workers in period t is $n_t = \sum_{i=0}^{I} e_{it}$. If the relative efficiencies are $\pi_0 < \pi_1 < \cdots < \pi_I$, where we normalize π_0 to one, then the corresponding quality-adjusted number is $e_t = \sum_{i=0}^{I} \pi_i e_{it}$. This variable replaces n_t in the production function.

The rest of the model is as in Section 5. Indeed, that economy is a special case (for $\Delta\pi = 0$) of the one considered here. With on-the-job learning, I state variables are added. With the computational method used, computer time increases a little, but there is no practical difficulty in setting up the computational experiments.

Assume that the absolute increments to π_i are equal at all stages, that is. $\Delta\pi_i = \pi_i - \pi_{i-1}$ are the same for all i. This means, of course, that the percentage increases get smaller at each higher stage. We choose $I = 8$ and $\Delta\pi = 0.05$, so that the most highly skilled workers are 40 percent more productive than those just entering the market sector. This is a compromise. Measurements probably would indicate steeper growth of efficiency at the initial stages and flatter growth at the later ones, with growth of some magnitude continuing after two years. The attrition rate, η, is set equal to 0.08 per quarter. Consequently, in the steady state, about half of the model's working population is in the highest-earning group.

The comovements of the various aggregates with GNP and most of the relative volatilities are quite similar to those for the case of $\Delta\pi = 0$. The main difference is that the standard deviation of output drops by 0.10; in other words, Solow residuals account for a slightly smaller fraction of the business cycle.

It has been suggested that with human capital, different measurements are needed for the Solow residuals. This is not necessarily so. The situation is analogous to that in Kydland and Prescott (1991), where the authors permit variation in the number of hours a plant is operated, while the measurements of Solow residuals do not assume this. The magnitudes of the technology shocks going into the model are known. One can then measure the shocks in the model in the same way that they are measured in the data and estimate the magnitude of the bias. In the Kydland and Prescott (1991) study, this procedure led to a slight reduction in the estimate of the fraction of the output variance accounted for.

8. Conclusion

This chapter has presented variants of what can be regarded as the dominant framework of shared knowledge in aggregate economics. It is a framework within which one can organize and interpret NIPA data. The particular choice of model environment within this framework, of course, depends on the question to be addressed. The question of the role of shocks to aggregate production technology for the business cycle has received considerable attention in the past ten years. In this chapter we have focused on the extent to which the estimate of this role depends on the model specification as it relates to the labor market in particular. To some extent, the different environments represent a progression over time in our understanding of the role of the labor input.

As we have seen, in spite of using an identical stochastic process for the impulse—the technology shock in the market sector—in each of the economies, the resulting volatility of GNP across models can be quite different. In other words, the roles of the propagation mechanisms are of central importance. In choosing models to consider, we have focused on the extent to which they represent different specifications of features that affect aggregate behavior as reflected more or less directly in the labor market.

In the initial development and use of this framework, some features of the workings of the labor market in the U.S. data, especially the volatility of aggregate hours of work and the correlation between hours and productivity, were regarded as important deviations relative to theory. As theory and measurements have progressed, however, the status of these features as deviations has diminished. Better abstractions have been developed, for instance, to indicate that a great deal of aggregate intertemporal substitution of hours is what the theory predicts. From a measurement standpoint, evidence suggests that the volatility of the labor input, which one would like to measure by weighting the hours of different workers according to their normal efficiency, is considerably less than the unweighted hours variability. The high correlation between hours and productivity, of course, is to be expected in environments with only technology shocks as a source of impulse. As illustrated in the model with shocks to household production added, the presence of other impulses will reduce that correlation. This has also been demonstrated with government shocks as the additional impulse (Christiano and Eichenbaum 1992).

Among other things, we have discussed ways in which the propagation of shocks via the labor market is affected through interaction of the business and household production. It is probably fair to say that we know mainly about the *potential* for household production to play a sizable role. A clearer answer about its actual role, however, will have to await measurements that have not yet been carried out. This is an important area of future research. Another area is consideration of whether the findings using environments with adjustment along both the intensive and extensive margins are affected by the degree of insurance assumed in those models.

Many recent contributions to the understanding of the labor market and the cycle have been omitted from this overview. For example, while Hansen (1985) shows that intertemporal substitution in the aggregate may be much larger than that reflected in individuals' preferences, Smith (1989) finds a tendency in the same direction due to asymmetric information between workers and firms about the workers' skills. We did not focus on the countercyclical labor share of national income observed in the data. Ways of accounting for this fact are studied in Danthine and Donaldson (1990), who use a contracting set-up, and in Gomme and Greenwood (1993).

It may be surprising to some that we make few references to micro labor studies, given that such studies are potential sources of calibration. The main reason is that much of that literature has been occupied by the goal of measuring such things as supply and demand elasticities for labor. With modern general equilibrium language, measurements of such elasticities do not map naturally into model parameters. Moreover, to the extent that one can interpret low elasticities as evidence of limited willingness, according to individual preferences, to substitute intertemporally, the insight from Hansen's (1985) economy suggests that this has little or no relevance to aggregate questions.

We have already listed some interesting measurement issues that remain for future research. On the theory side, many features of the labor market have received little attention and also represent interesting research areas for the future. Examples are the role of the differences of skills across workers for market production, the role of variation in capacity utilization and its implications for the aggregate production function, and the role of less than perfect insurance for workers against shocks.

Note

The National Science Foundation and the Federal Reserve Bank of Cleveland has provided research support. Christian Zimmermann assisted with the computational experiments.