7 Competitive Equilibrium Business Cycles

In this chapter we explore the idea that macroeconomic fluctuations can for the most part be explained by the dynamic effects of shocks in a competitive economy. This line of inquiry is recent.[1] For most of the twentieth century, especially since the Great Depression, most macroeconomists have looked upon the sharp fluctuations in output and unemployment as prima facie evidence of major market imperfections and explored what these imperfections may be. In the last 15 years, however, some have argued that this is a misguided research strategy, since macroeconomic fluctuations can be explained without invoking imperfections. The first statement, that fluctuations can be explained as the realization through time of the set of transactions agreed upon in a complete market Arrow-Debreu economy, was presented by Black (1982). In a series of papers Prescott (1986, for example) has explored this idea further, by examining the dynamic effects of productivity shocks on the competitive equilibrium allocation, to conclude: "Economists have long been puzzled by the observations that, during peacetime, industrial market economies display recurrent, large fluctuations in output and employment over relatively short time periods.... These observations should not be puzzling, for they are what standard economic theory predicts."

As will become clear, we do not believe that the line of research that we present in this first chapter on aggregate fluctuations is likely to provide a satisfactory explanation of fluctuations. Nevertheless, this chapter is important for two reasons. The first is that it is a logical starting point to the study of fluctuations. To see why one may want to explore the role of imperfections in fluctuations, it is essential to understand what the dynamic effects of shocks would be in a competitive economy. The second is that productivity shocks, even if they do not account for all or even for most fluctuations, may nevertheless play a more important role in fluctuations than has been emphasized until now.

The chapter has three sections. In section 7.1 we extend the representative agent dynamic models of chapters 2 and 3 to allow for productivity shocks and uncertainty. Our emphasis in this section is on the dynamics of output and its components when productivity is random. We do not strive for generality, preferring where possible to obtain closed-form solutions that illustrate both the approach and the basic dynamic mechanisms at work.

In section 7.2 we relax the assumption of constant employment and examine the mechanisms that may be capable of generating the comovements in output and employment that characterize actual fluctuations. The challenge is to explain how technological shocks can generate persistent movements not only in output but also in employment. We indicate why we do not think that the challenge is met or is likely to be met by that approach to business cycles.

In the final section we explore the implications of decentralized markets for the dynamic effects of shocks on the economy. We analyze first a simple model of search to show how the heterogeneity of workers and jobs in the labor market modifies the effects of productivity shocks on employment, wages, and output. We then extend the model to allow for imperfect information about the realizations of current shocks and for two different types of shocks, real and monetary, an approach to fluctuations first explored by Lucas. Imperfect information modifies the dynamic effects of real shocks. It also implies a potential effect of monetary shocks which would not be present under perfect information. We conclude with an assessment of the competitive equilibrium approach.[2]

7.1 Productivity Shocks, Consumption, and Capital Accumulation

When studying the dynamics of capital accumulation in the models of chapters 2 and 3, we made counterfactual assumptions of certainty and perfect foresight. This was a necessary first step, but we now recognize the presence of uncertainty. The economy is constantly affected by the introduction of new technologies, by changes in tastes for specific goods, by changes in government policy, and so on. Most of these changes are neither perfectly predictable nor perfectly predicted by individuals and firms. The presence of uncertainty affects the behavior of agents. The shocks themselves lead them to constantly revise their optimal plans.

In this section we focus mostly on the dynamic effects of one type of shock, shocks to productivity. We concentrate on two main issues. First, we examine whether the propagation mechanism tends to amplify or to dampen the effects of shocks on aggregate output and whether productivity shocks can plausibly explain the pattern of serial correlation of output characterized in chapter 1. Second, we examine whether productivity shocks can explain the comovements of output and its components, also characterized in chapter 1. Finally, we take a brief detour and return to an issue that was central to many of our discussions in earlier chapters, that of the form of the golden rule under uncertainty.

We start by introducing shocks in both the Diamond and the Ramsey models—more specifically, in versions of those models for which explicit solutions can be obtained: this is a fairly substantial restriction.[3] We then turn to the additional dynamics introduced by inventory behavior.

Multiplicative Shocks in the Diamond and Ramsey Models

Productivity Shocks in the Diamond Model

Productivity shocks that affect output in the current period are likely to lead to increased consumption as well as to increased saving, and thus to increased capital accumulation. Because increased capital accumulation leads to higher output later, productivity shocks lead to a serially correlated response of output. This point can be made straightforwardly in a simple stochastic version of the Diamond overlapping generations model with capital.

Population is assumed to be constant. People live for two periods, each person supplying one unit of labor inelastically in the first period. For convenience, the size of each generation is normalized to one. The production function is assumed to be Cobb-Douglas, with production given by

$$Y_t = U_t K_t^a N_t^{1-a} = U_t K_t^a, \tag{1}$$

where U_t is the level of productivity (a random variable with properties to be defined later), K_t is capital, and L_t is labor. The second equality results from the normalization of the labor force. Y_t is gross output, that is, output that includes the capital stock left after production. Equivalently, Y_t is net output, and capital depreciates fully after one period.[4]

An individual born at time t supplies one unit of labor and earns a wage ω_t. She then chooses consumption at time t to maximize her expected utility:

$$\ln C_{1t} + (1+\theta)^{-1} E[\ln(C_{2t+1}) \mid t]$$

subject to the budget constraint

$$C_{2t+1} = (1 + r_t)(\omega_t - C_{1t}).$$

The assumption of logarithmic utility yields a simple solution to the maximization problem. First-period consumption and savings are proportional to wage income and independent of the interest rate. The independence of savings from the interest rate, which is stochastic here, is what makes the model so tractable. Letting S_t denote the savings of the young, and noting that $S_t = \omega_t - C_{1t}$,

$$S_t = \frac{\omega_t}{2+\theta}.$$

Thus, by using the fact that $\omega_t = (1-a)U_t K_t^a$, and that $K_{t+1} = S_t$, we obtain a stochastic difference equation for the capital stock:

$$K_{t+1} = \frac{(1-a)U_t K_t^a}{2+\theta}. \tag{2}$$

The capital stock today determines labor income, which in turn determines saving and the capital stock in the next period.[5]

Equation (2) is linear in logarithms. We denote logarithms by lowercase letters:

$$k_{t+1} = b + ak_t + u_t, \tag{3}$$

where

$$b \equiv \ln\left(\frac{1-a}{2+\theta}\right).$$

By taking logarithms in (1) and replacing capital by its expression from (3), we obtain a dynamic equation for output:

$$y_t = ab + ay_{t-1} + u_t. \tag{4}$$

The logarithm of output follows a first-order difference equation, with forcing term u_t; the first-order autoregressive coefficient a is equal to the share of capital in output. This shows the contribution of capital accumulation to the persistence of the effects of technological shocks on output. We now explore this contribution further.

Suppose first—and counterfactually—that u_t is a white noise random variable so that the logarithm of productivity is equal to a constant plus white noise. In this case output follows a first-order autoregressive process, with a serial correlation coefficient equal to a: a positive productivity shock leads to an increase in consumption and savings and to higher levels of capital and output. Over time, capital and output return to their initial levels.

The empirical counterpart of a is difficult to determine, since this is a gross production function for a model with a unit period of about 30 years. But consideration of the economics of the mechanism specified here suggests it can account for relatively little serial correlation. Suppose that a shock raises GNP this year by 1% and that saving increases by as much as 0.5% of GNP—a very high saving response. With the real return on capital equal to about 10%, the extra saving would increase GNP in the following year by 0.05%. Capital accumulation clearly cannot account for much of the serial correlation in output.

The assumption of white noise productivity, or even of white noise productivity around a deterministic trend, is, however, not particularly appealing. New techniques, once introduced, should be available forever. Suppose that innovations in productivity have permanent effects on the level of productivity, that is, the process for productivity has a unit root.[6] A stochastic process that may more accurately describe productivity growth is

$$u_t = g + u_{t-1} + \varepsilon_t$$

where ε_t is white noise. Productivity follows a random walk with drift g so that it grows on average at rate g.[7] Put another way, productivity growth is a white noise process. Given this process for productivity, we replace it in (4), and output becomes

$$\Delta y_t = g + a\Delta y_{t-1} + \varepsilon_t. \tag{5}$$

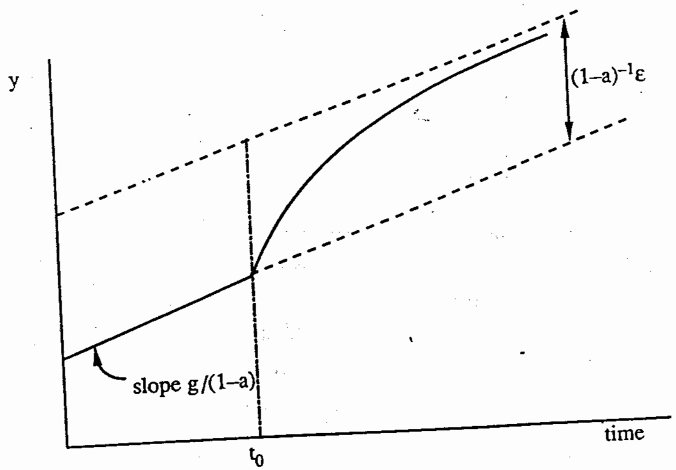

Figure 7.1
Effects of an increase in ε at t_0 on output

As a result of capital accumulation, *output growth* is serially correlated, with serial correlation coefficient a. The effects of a positive productivity shock, an increase in ε, on the level of y are given in figure 7.1. Output increases above its initial growth path at time zero, and increases further through time at a rate faster than g. Eventually, the rate of growth returns to g, but the level of output is higher than it would otherwise have been.

As we noted earlier, the coefficient a is, in practice, likely to be small, and therefore capital accumulation does not amplify the dynamics of productivity shocks significantly *through this particular mechanism*.[8] Also virtually any stochastic process that is desired for output can be obtained by appropriately specifying the corresponding stochastic process for the productivity shock. Thus independent information on the stochastic process for productivity shocks is needed to assess their role in economic fluctuations.

Productivity Shocks in the Ramsey Model

Very similar results obtain in the corresponding version of the Ramsey model with infinitely long-lived individuals. Consider the following modification of the model of chapter 2; in it we again rely on logarithmic specifications of utility and production functions to obtain explicit solutions.

There is no population growth, and the population size is normalized to one. Individuals live forever, supplying one unit of labor in each period and maximizing at time t:

$$\sum_{i=0}^{\infty} (1+\theta)^{-i} E[\ln C_{t+i} | t]$$

subject to

$$K_{t+i+1} + C_{t+i} = Y_{t+i} \equiv U_{t+i} K_{t+i}^{a}.$$

As before, U, a random variable, is the level of productivity, and the production function represents gross output. From the first-order conditions we obtain[9]

$$\left(\frac{1}{C_t}\right) = (1+\theta)^{-1} E\left[\frac{aU_{t+1}K_{t+1}^{a-1}}{C_{t+1}} \middle| t\right].$$

An educated guess at the solution to this difference equation is that $C_t = \beta U_t K_t^a$ for a value of β to be determined. If we replace it in the above expression and use the budget constraint, the guess proves to be right and the appropriate value of β to be $1 - [a/(1+\theta)]$.

The logarithmic assumptions therefore deliver the very strong result that consumption and saving are proportional to income, independent of the stochastic properties of U.[10] The implied process for capital is

$$K_{t+1} = S_t = \left(\frac{a}{1+\theta}\right) U_t K_t^a, \tag{6}$$

or, using logarithms,

$$k_{t+1} = b + a k_t + u_t, \tag{7}$$

where

$$b \equiv \ln\left(\frac{a}{1+\theta}\right).$$

The results are therefore similar to those of the Diamond model. The only difference is in the constant term, which determines the average level of capital in the economy.

Multisector Extension

Long and Plosser (1983) have developed a multisector extension of the Ramsey model. Each sector is affected by its own productivity shock and uses as inputs both labor and the outputs of other sectors, with a one-period production lag. The dynamics of the system are richer because the deterministic part of the system now has as many roots as the number of sectors, in constrast to the dynamic system given by equation (7) which has only one root.[11] Using an aggregate empirical input-output table to calibrate their model, they show that even under the assumption that productivity shocks are independent both across time and across sectors, the outputs of individual sectors may exhibit both serial and cross correlations. The average pairwise cross correlation across sectors is around 20%, and the average first-order serial correlation (for annual data) around 30%.

A Detour: The Modified Golden Rule under Uncertainty

Having introduced uncertainty in the Ramsey model, we can return to the form of the modified golden rule, this time under uncertainty. What is the relation between the *marginal product of capital* and the *discount rate* in this stochastic extension of the Ramsey model?

Assume, for simplicity, that the stochastic process for productivity is either stationary or has a unit root and zero drift [i.e., if productivity follows the process given by equation (5), g is equal to zero]. Then under certainty the modified golden rule in this economy, without population or productivity growth, would be that the gross marginal product of capital, which

would also be equal to the gross interest rate, is equal to one plus the subjective discount rate. This can be seen, for instance, by setting U in (6) equal to a constant and solving for the real interest rate with $K_{t+1} = K_t$, which is the steady state condition.

To see what holds here, consider the logarithm of the marginal product of capital (MPK), which is given by

$$\ln(\text{MPK}_t) = \ln a + u_t + (a - 1)k_t. \tag{8}$$

By solving for k_{t+1} as a function of past u's in (7), replacing k_t in (8), and rearranging, we get

$$\ln(\text{MPK}_t) = \ln(1+\theta) + (1 - aL)^{-1}(u_t - u_{t-1}), \tag{9}$$

where L is the lag operator. Thus, by taking unconditional expectations on both sides,[12] we get

$$E[\ln(\text{MPK})] = \ln(1+\theta).$$

In this version of the Ramsey model with uncertainty, the modified golden rule is that the expectation of the logarithm of the marginal product is equal to the logarithm of one plus the discount rate. By Jensen's inequality, this implies that under uncertainty the expectation of the marginal product exceeds one plus the discount rate.

If we are ready to specify the process for productivity, we can go further and get an explicit solution for the expected value of the marginal product. If, for example, productivity growth is white noise, with innovations e_t normally distributed with variance σ^2, we have from equation (9),

$$E[\ln(\text{MPK})] = \ln(1+\theta)$$

and

$$V[\ln(\text{MPK})] = \left(\frac{1}{1-a^2}\right)\sigma^2.$$

Thus

$$E[\text{MPK}] = (1+\theta)\exp\left[\frac{\sigma^2}{2(1-a^2)}\right].$$

We can also derive the relation between the *riskless rate* and the discount rate. Although there is no riskless asset in this economy, we can find out what its equilibrium rate of return would be, were we to introduce one at the margin. The rate of return on the riskless asset in period t, R_t, would have to satisfy the first-order conditions of consumers, namely,

$$\frac{1}{C_t} = R_t(1+\theta)^{-1}E\left[\left(\frac{1}{C_{t+1}}\right)\bigg| t\right],$$

or, upon rearranging,

$$\ln(R_t) = \ln(1+\theta) - \ln\left(E\left[\frac{C_t}{C_{t+1}}\bigg| t\right]\right).$$

If we assume that productivity growth is white noise, with innovations e_t normally distributed with variance σ^2, the logarithm of consumption follows

$$c_{t+1} - c_t = a(c_t - c_{t-1}) + e_t$$

so that

$$\ln(R_t) = \ln(1+\theta) + a(c_t - c_{t-1}) - \left(\frac{\sigma^2}{2}\right).$$

The riskless rate is therefore not constant. At time t it depends on the past change in consumption. Because changes in consumption are positively correlated, an increase in consumption in period t implies an expected increase in consumption from t to $t+1$ and thus a higher riskless rate. The riskless rate is a decreasing function of the uncertainty associated with changes in consumption. An intuitive explanation, based on our analysis of precautionary saving in the previous chapter, is as follows: an increase in uncertainty makes people more prudent, so that, at the same real rate, they would consume less today and more tomorrow. To reestablish equilibrium, the real rate must decrease, leading to an offsetting increase in consumption today and a decrease tomorrow.

The last step is to compute the unconditional expected value of the riskless rate. From above,

$$E[\ln(R_t)] = \ln(1+\theta) - \left(\frac{\sigma^2}{2}\right),$$

$$V[\ln(R_t)] = \frac{a^2\sigma^2}{(1-a^2)}.$$

The variance in the logarithm of the riskless rate depends on the variance in the expected rate of change of consumption. From the above equations we get

$$E[R] = (1+\theta)\exp\left\{\left(\frac{\sigma^2}{2}\right)\left[\frac{a^2}{1-a^2} - 1\right]\right\}.$$

The effect of uncertainty is ambiguous. The riskless rate may be greater or smaller than the subjective discount rate. The closer consumption growth is to white noise, the more likely it is that the real rate is less than the subjective discount rate.[13]

We have shown that in a dynamically efficient economy such as the Ramsey economy, there is, once we allow for uncertainty, no simple relation between the marginal product of capital, the riskless rate, and the subjective discount rate. The reason for deriving those results was, however, our interest in a different but closely related question, that of dynamic efficiency in an arbitrary economy. In an economy where people have finite horizons and where there is uncertainty, what rate should we look at to tell if the economy is dynamically efficient? Empirically, the average real rate of interest on (relatively) safe Treasury bills has averaged less than 1% over the past 60 years, while the return on stocks has exceeded the rate of growth. Which is the relevant rate of return to compare with the rate of growth in assessing whether the economy has overaccumulated capital? The preceding example suggests that neither of the two rates is likely to be appropriate and that there may not be a simple answer to that question. The question of how to assess empirically whether an economy is dynamically efficient is still quite open. Abel et al. (1986) derive a sufficient condition for efficiency,[14] that the level of investment must be less in every year than the share of profits. (Under certainty the dynamic efficiency can be stated as the condition that investment must be less or equal to the share of profits. This was first shown by Phelps 1961.) They show that this condition has been satisfied in every year since 1929 and therefore, they reason, the U.S. economy is very likely to be dynamically efficient.

Additive Shocks in a Linear Ramsey Model

The assumptions made thus far in this section imply that the saving rate is constant and independent of the process followed by productivity. To obtain some insight into dynamics when the saving rate can vary, we develop another simplified case. We continue to use the Ramsey model but assume that productivity shocks are additive rather than multiplicative and that the real interest rate is constant. As we shall see, this affects not only the short-run but also the long-run effects of productivity shocks.

Consider an economy with a constant population whose size is normalized at one. Individuals are infinitely long lived and maximize at time t:

$$\sum_{i=0}^{\infty}(1+\theta)^{-i}E[C_{t+i} - bC_{t+i}^2|t], \qquad b > 0$$

subject to

$$K_{t+i+1} + C_{t+i} = Y_{t+i} \equiv (1 + r)K_{t+i} + U_{t+i}.$$

Instantaneous utility is quadratic. Production requires only capital and takes place under constant returns to scale, with r being the net marginal product of capital. (An alternative interpretation, similar to a model developed in chapter 2, is that this is a small economy that can borrow and lend at the world interest rate r.) U_t is now additive and therefore does not affect the marginal product of capital.[15] It is best thought of as an exogenous endowment each period, as manna from heaven, that follows a given stochastic process. We will refer to it as the endowment. Together these assumptions yield certainty equivalence.

Recalling the results of chapter 2, an economy with infinitely long lived individuals will, if r is different from θ, be either accumulating or decumulating capital forever. Thus, to obtain a well-behaved stochastic steady state, we assume in what follows that θ is equal to r, that is, the discount rate is equal to the marginal product of capital. Given this assumption, the first-order conditions imply that

$$C_t = E[C_{t+i}|t], \qquad \text{for all } i > 0.$$

Consumers choose to have constant expected consumption. The highest feasible expected level of consumption that satisfies the intertemporal budget constraint is given by[16]

$$C_t = r\left\{K_t + (1 + r)^{-1} \sum_{i=0}^{\infty} (1 + r)^{-i}E[U_{t+i}|t]\right\}. \tag{10}$$

The basic feature of this economy can be understood by examining the coefficient associated with U_t in equation (10). That coefficient is $r/(1 + r)$. This means that the consumer treats U_t as an annuity, from which he draws down the interest for consumption each period. Thus, if movements in U_t are transitory, there is consumption smoothing in this version of the Ramsey model. Consumption depends on the current capital stock and the expected present discounted value of endowments.

Replacing this expression for consumption in the budget constraint gives the behavior of capital and thus also (from the production function) that of output:

$$K_{t+1} = K_t + \left\{U_t - r(1 + r)^{-1} \sum_{i=0}^{\infty} (1 + r)^{-i}E[U_{t+i}|t]\right\}. \tag{11}$$

The change in the capital stock is equal to the difference between the current endowment and the expected present discounted value of current and future endowments.

We can obtain further insight by solving for a specific endowment process. For a change we consider the following IMA(1, 1) process:

$$U_t = U_{t-1} + e_t - ae_{t-1}, \qquad |a| < 1.$$

This is the simplest process that allows for both permanent and transitory effects of shocks on endowments. The long-run effect of a shock e on U is equal to $1 - a$. Thus, if $a = 1$, the process reduces to white noise, with shocks having purely transitory effects. For a between 0 and 1, shocks have a permanent effect, though smaller than their initial effect. If a is equal to 0, the process reduces to a random walk. For a between 0 and -1, the permanent effect exceeds the initial effect.

By solving for expectations in (11), we find the impact of a given shock e_t on the change in the capital stock:

$$K_{t+1} - K_t = \left(\frac{a}{1 + r}\right)e_t. \tag{12}$$

Thus a shock, e, increases capital accumulation only to the extent that its effect on productivity is *transitory*. If its effect is permanent, if productivity follows a random walk, the increase in output is matched by an equal increase in consumption, resulting in an unchanged capital stock. If, however, its effect is partly transitory (a between 0 and 1), consumption increases less than one for one with output, resulting in an increase in capital accumulation. This ensures that the effects of the productivity increase on consumption are spread over time.

In this case, in contrast to the previous examples, even purely transitory shocks ($a = 1$) have a permanent effect on capital accumulation and output. This occurs because the model, by assuming that the interest rate is constant and independent of the level of capital, implicitly assumes constant returns to scale given labor, an assumption we have made for convenience of argument and without explicit justification.[17]

The first two examples emphasized the fact that positive technological shocks generally lead to more capital accumulation, which then amplifies the initial effect of these shocks. This example, by contrast, emphasizes the potential role of consumption smoothing in determining the serial correlation of output.[18]