Chapter 9

Business Cycles

In this chapter we explore the causes of business cycles. Briefly, business cycles are the recurring fluctuations that occur in real GDP over time. For further descriptions of business cycles, refer to Barro's Chapter 9. Here, we concentrate on explaining business cycles. We begin with an overview of potential explanations. Then we work out a real business cycle model in detail.

While there are many different theories of business cycles, they share some properties. There is always a driving force behind economic fluctuations, some sort of shock or disturbance that is the original cause of the cycle. In addition, most theories build on a propagation mechanism that amplifies shocks. Unless the disturbances are already big enough by themselves to account for the fluctuations, there has to be some propagation mechanism that translates small, short-lived shocks into large, persistent economic fluctuations.

We will start our search for the cause of business cycles in Section 9.1 by listing a number of possible shocks and propagation mechanisms. Competing theories of the business cycle differ in which shocks and mechanisms they emphasize. In Section 9.2 we will concentrate on the real business cycle model, which is a straightforward extension of the market-clearing models that we developed in earlier chapters. Section 9.3 presents simulations for our real business cycle model and assesses the success of the model in matching real-world fluctuations.

9.1 Shocks and Propagation Mechanisms

Among the many shocks and disturbances that are present in an economy, only a few have received special attention in research on business cycles. Here are some of the more important candidates:

- **Technology shocks:** Real-world production functions change over time. New technologies like computers or robots alter the production process and raise overall productivity. Sometimes, production facilities break down or do not work as expected, so productivity falls. This technological change is not always smooth; it often comes in the form of shocks.

- **Weather shocks and natural disasters:** Many industries like agriculture or tourism are weather-dependent. Rainfall and sunshine influence the output of these sectors, so the weather is a a potential source of fluctuations. This is also true for disasters like earthquakes or landslides. El Niño is a shock of this kind that received a lot of attention lately. We can regard these kinds of shocks as one type of technology shock. Weather changes the production function for wheat, and an earthquake that wiped out, say, Silicon Valley, would change the production function for computers.

- **Monetary shocks:** We saw in Chapter 8 on inflation that there are real effects of monetary policy. Therefore random changes to money supply or interest rates are a potential source of fluctuations as well.

- **Political shocks:** The government influences the economy both directly through government enterprises and indirectly through regulation. Changes in tax laws, antitrust regulation, government expenditure and so on are a potential source of disruption in the economy.

- **Taste shocks:** Finally, it is also conceivable that shifts in preferences cause fluctuations. Fashion and fads change rapidly, and they may cause fluctuations in areas like the apparel, music, or movie industries.

While the shocks just mentioned are present to some degree in every economy, they are probably not large enough to serve as a direct explanation of business cycles. For example, in the United States real GDP fell by 2.8% between October 1981 and 1982. It is hard to imagine any shock that caused a direct output loss of almost 3% of GDP within only a year, and if there was one, we would probably be aware of it. It appears more likely that there are mechanisms present in the economy that amplify shocks and propagate them through time. Here are some candidates:

- **Intertemporal substitution:** Shocks that have a negative impact on productivity lower the marginal return to labor and other factors of production. If marginal products fall, consumer's might prefer to work less and consume leisure instead. Labor input would fall, which amplifies the negative impact on output. At the same time, since consumers prefer a smooth consumption profile they might prefer to lower savings for some time when a shock hits. On an aggregate level, this leads to lower investment and a lower capital stock in the future. Therefore a short-lived shock may have an impact in the future as well.

- **Sticky prices:** Market economies react to changes with price adjustments. For example, a negative productivity shock lowers the marginal product of labor, so that the real wage would have to move downward to adjust labor demand and supply. But if wages are inflexible for some reason, the adjustment cannot take place. The result is unemployment and an output loss that is larger than the direct effect of the shock. Similar effects arise if goods prices are sticky.

- **Frictions in financial sector:** Even small shocks can force the firms the are hit directly into bankruptcy. This will affect other firms and banks that lent money to the now bankrupt firms. Often additional firms have to declare bankruptcy, and sometimes even banks fail. Bank failures affect all creditors and debtors and therefore can have large economic consequences. Serious economic crises are often accompanied and amplified by series of bank failures. Examples are the great depression and the current Asian crisis.

Business cycle models can be broadly subdivided into two categories. Some theories regard cycles as a failure of the economic system. Because of frictions or imperfections of the market mechanism, the economy experiences depressions and fails to achieve the efficient level of output and employment. Models of this kind often rely on financial frictions, sticky prices, or other adjustment failures as the propagation mechanism. Both technology shocks and monetary shocks are considered to be important sources of fluctuations. The Keynesian model of output determination[1] falls into this category.

On the other hand, there is a class of models that regards business cycles as the optimal reaction of the economy to unavoidable shocks. Shocks are propagated through intertemporal substitution within an efficient market mechanism. Technology shocks are considered to be the main cause of economic fluctuations. Models of this kind are often referred to as *real business cycle* models.[2]

We can be fairly certain that there is some truth to both views of economic fluctuations. Major economic breakdowns like the great depression or the recent Asian crisis appear to be closely connected to disruptions in the financial sector. Bank failures and financial instability played an important role in both cases.

On the other hand, most business cycles are far less severe than the great depression or the Asian crisis. In the entire post-war history of the United States and the Western European countries there is not a single depression that caused an output loss similar to the one suffered between 1929 and 1933. The question is whether normal business cycles are caused by the same kind of frictions that caused the great depression. The Keynesian model with its emphasis on slow adjustments and sticky prices supports this view. Real business cycle theorists argue that breakdowns like the great depression are a phenomenon distinct from usual business cycles, and that usual cycles can be explained as the optimal reaction of an efficient market system to economic shocks.

[1]See Barro, Chapter 20.

[2]The term derives from the fact that shocks in real business cycle theory are real, as opposed to monetary, and that sluggish nominal adjustment does not play a role as a propagation mechanism.

In this chapter, we will primarily look for explanations for normal-scale business cycles, like those experienced in the United States since World War II. How can we determine whether such cycles are small-scale failures of the economic system rather than simply the markets' efficient reactions to shocks? A natural way to answer this question is to build a number of model economies that include alternative propagation mechanisms, expose the model economies to shocks, and see whether the outcomes look like real-world business cycles. This is exactly the road that has been taken by real business cycle theorists. They have taken standard equilibrium models as a point of departure and exposed them to productivity shocks. As it turns out, models of this kind are quite successful at explaining real-world business cycles. We will now take a closer look at such a real business cycle model.

9.2 A Real Business Cycle Model

Real business cycle models are straightforward extensions of equilibrium models of the kind that we use throughout this course. In most cases, the models feature infinitely lived consumers, and business cycles are generated by random disturbances to production possibilities. Unfortunately, solving that kind of model is difficult. Often no explicit solution is available, so numerical approximations have to be used. To keep the presentation tractable, in this chapter we will use a simpler framework in which people live for two periods only. The model does not fit the facts as well as a full-scale real business cycle model, but it serves its purpose as a simple illustration of the main ideas of real business cycle theory.

In the model world there is a sequence of overlapping generations. Each period a new generation of consumers is born, and each consumer lives for two periods. We will sometimes refer to the periods as years, and for simplicity we assume that exactly one consumer is born each year. People work in the first period when they are young. In the second period they are retired and live on savings. Throughout the model, superscripts refer to the year when a person was born, while subscripts refer to the actual year. For example, c_t^t is the period-t consumption of a consumer who was born in year t, so such a consumer would be young in period t. Similarly, c_{t+1}^t is the consumption of the same consumer in period $t+1$, when he is old. The consumers do not care about leisure. A consumer born in year t has the following utility function:

$$u(c_t^t, c_{t+1}^t) = \ln(c_t^t) + \ln(c_{t+1}^t).$$

We could introduce a discount factor, but for simplicity we assume that the consumers value both periods equally. Note that at each point of time there are exactly two people around: one who was just born and is young, and another who was born the year before and is now retired. In each period the young person supplies one unit of labor and receives wage income w_t. The labor supply is fixed, since consumers do not care about leisure. The wage income can be used as savings k_t and as consumption c_t^t. The budget constraint of a young worker is:

$$c_t^t + k_t = w_t,$$

i.e., consumption plus savings equals income from labor. In period $t+1$ the consumer born in t is old and retired. The old consumer lends his savings k_t to the firm. The firm uses the savings as capital and pays return r_{t+1} to the old consumer. A fraction δ of the capital wears out while being used for production and is not returned to the consumer. δ is a number between zero and one, and is referred to as the *depreciation rate*. The budget constraint for the retirement period is:

$$c_{t+1}^t = (1 - \delta + r_{t+1})k_t,$$

i.e., consumption equals the return from savings.

The household born in period t maximizes utility subject to the budget constraints, and takes prices as given:

$$\max_{c_t^t, c_{t+1}^t, k_t} \left\{ \ln(c_t^t) + \ln(c_{t+1}^t) \right\}, \text{ subject to:}$$
$$c_t^t + k_t = w_t, \text{ and:}$$
$$c_{t+1}^t = (1 - \delta + r_{t+1})k_t.$$

We can use the constraints to eliminate consumption and write this as:

$$\max_{k_t} \{ \ln(w_t - k_t) + \ln((1 - \delta + r_{t+1})k_t) \}.$$

This is similar to the problem of the consumer in the two-period credit market economy that we discussed in Section 3.2. From here on we will drop the practice of denoting optimal choices by superscripted stars, since the notation is already complicated as it is. The first-order condition with respect to k_t is:

$$0 = -\frac{1}{w_t - k_t} + \frac{1 - \delta + r_{t+1}}{(1 - \delta + r_{t+1})k_t}.$$

Solving this for k_t yields:

$$k_t = \frac{w_t}{2}. \tag{9.1}$$

Thus, regardless of the future return on capital, the young consumer will save half of his labor income. Again, this derives from the fact that wealth and substitution effects cancel under logarithmic preferences. This feature is is helpful in our setup. Since there will be productivity shocks in our economy and r_{t+1} depends on such shocks, the consumer might not know r_{t+1} in advance. Normally we would have to account for this uncertainty explicitly, which is relatively hard to do. In the case of logarithmic utility, the consumer does not care about r_{t+1} anyway, so we do not have to account for uncertainty.

Apart from the consumers, the economy contains a single competitive firm that produces output using capital k_{t-1} and labor l_t. Labor is supplied by the young consumer, while the supply of capital derives from the savings of the old consumer.The rental rate for capital is r_t, and the real wage is denoted w_t. The production function has constant returns to scale and is of the Cobb-Douglas form:

$$f(l_t, k_{t-1}) = A_t l_t^{\alpha} k_{t-1}^{1-\alpha}.$$

Here α is a constant between zero and one, while A_t is a productivity parameter. A_t is the source of shocks in this economy. We will assume that A_t is subject to random variations and trace out how the economy reacts to changes in A_t. The profit-maximization problem of the firm in year t is:

$$\max_{l_t, k_{t-1}} \left\{ A_t l_t^{\alpha} k_{t-1}^{1-\alpha} - w_t l_t - r_t k_{t-1} \right\}.$$

The first-order conditions with respect to l_t and k_{t-1} are:

$$\text{(FOC } l_t) \qquad A_t \alpha l_t^{\alpha-1} k_{t-1}^{1-\alpha} - w_t = 0; \text{ and:}$$

$$\text{(FOC } k_{t-1}) \qquad A_t (1-\alpha) l_t^{\alpha} k_{t-1}^{-\alpha} - r_t = 0.$$

Using the fact that the young worker supplies exactly one unit of labor, $l_t = 1$, we can use these first-order conditions to solve for the wage and return on capital as a function of capital k_{t-1}:

$$w_t = A_t \alpha k_{t-1}^{1-\alpha}; \text{ and:} \tag{9.2}$$

$$r_t = A_t (1-\alpha) k_{t-1}^{-\alpha}. \tag{9.3}$$

Since the production function has constant returns, the firm does not make any profits in equilibrium. We could verify that by plugging our results for w_t and r_t back into the firm's problem. Note that the wage is proportional to the productivity parameter A_t. Since A_t is the source of shocks, we can conclude that wages are procyclical: when A_t receives a positive shock, wages go up. Empirical evidence suggests that wages in the real world are procyclical as well.

To close the model, we have to specify the market-clearing constraints for goods, labor, and capital. At time t the constraint for clearing the goods market is:

$$c_t^t + c_t^{t-1} + k_t = A_t l_t^{\alpha} k_{t-1}^{1-\alpha} + (1-\delta) k_{t-1}.$$

On the left hand side are goods that are used: consumption c_t^t of the currently young consumer, consumption c_t^{t-1} of the retired consumer who was born in $t-1$, and savings k_t of the young consumer. On the right hand side are all goods that are available: current production and what is left of the capital stock after depreciation.

The constraint for clearing the labor market is $l_t = 1$, since young consumers always supply one unit of labor. To clear the capital market clearing we require that capital supplied by the old consumer be equal to the capital demanded by the firm. To save on notation, we use the same symbol k_{t-1} both for capital supplied and demanded. Therefore the market-clearing for the capital market is already incorporated into the model and does not need to be written down explicitly.

In summary, the economy is described by: the consumer's problem, the firm's problem, market-clearing conditions, and a random sequence of productivity parameters $\{A_t\}_{t=1}^{\infty}$. We assume that in the very first period there is already an old person around, who somehow fell from the sky and is endowed with some capital k_0.

Given a sequence of productivity parameters $\{A_t\}_{t=1}^{\infty}$, an equilibrium for this economy is an allocation $\{c_t^t, c_t^{t-1}, k_{t-1}, l_t\}_{t=1}^{\infty}$ and a set of prices $\{r_t, w_t\}_{t=1}^{\infty}$ such that:

- Given prices, the allocation $\{c_t^t, c_t^{t-1}, k_{t-1}, l_t\}_{t=1}^{\infty}$ gives the optimal choices by consumers and firms; and
- All markets clear.

We now have all pieces together that are needed to analyze business cycles in this economy. When we combine the optimal choice of savings of the young consumer (9.1) with the expression for the wage rate in equation (9.2), we get:

$$k_t = \frac{1}{2} A_t \alpha k_{t-1}^{1-\alpha}. \tag{9.4}$$

This equation shows how a shock is propagated through time in this economy. Shocks to A_t have a direct influence on k_t, the capital that is going to be used for production in the next period. This implies that a shock that hits today will lead to lower output in the future as well. The cause of this is that the young consumer divides his income equally between consumption and savings. By lowering savings in response to a shock, the consumer smoothes consumption. It is optimal for the consumer to distribute the effect of a shock among both periods of his life. Therefore a single shock can cause a cycle that extends over a number of periods.

Next, we want to look at how aggregate consumption and investment react to a shock. In the real world, aggregate investment is much more volatile than aggregate consumption (see Barro's Figure 1.10). We want to check whether this is also true in our model. First, we need to define what is meant by aggregate consumption and investment. We can rearrange the market-clearing constraint for the goods market to get:

$$c_t^t + c_t^{t-1} + k_t - (1-\delta)k_{t-1} = A_t l_t^{\alpha} k_{t-1}^{1-\alpha}.$$

On the right-hand side is output in year t, which we are going to call Y_t. Output is the sum of aggregate consumption and investment. Aggregate consumption C_t is the sum of the consumption of the old and the young person, while aggregate investment I_t is the difference between the capital stock in the next period and the undepreciated capital in this period[3]:

$$\underbrace{c_t^t + c_t^{t-1}}_{C_t} + \underbrace{k_t - (1-\delta)k_{t-1}}_{I_t} = \underbrace{A_t l_t^{\alpha} k_{t-1}^{1-\alpha}}_{Y_t}.$$

Consumption can be computed as the difference between output and investment. Using equation (9.4) for k_t yields:

$$\begin{aligned} C_t = Y_t - I_t &= A_t k_{t-1}^{1-\alpha} + (1-\delta)k_{t-1} - k_t \\ &= A_t k_{t-1}^{1-\alpha} + (1-\delta)k_{t-1} - \frac{1}{2}A_t \alpha k_{t-1}^{1-\alpha} \\ &= \left(1 - \frac{1}{2}\alpha\right) A_t k_{t-1}^{1-\alpha} + (1-\delta)k_{t-1}. \end{aligned} \tag{9.5}$$

Aggregate investment can be computed as output minus aggregate consumption. Using equation (9.5) for aggregate consumption yields:

$$\begin{aligned} I_t = Y_t - C_t &= A_t k_{t-1}^{1-\alpha} - \left(1 - \frac{1}{2}\alpha\right) A_t k_{t-1}^{1-\alpha} - (1-\delta)k_{t-1} \\ &= \frac{1}{2}A_t \alpha k_{t-1}^{1-\alpha} - (1-\delta)k_{t-1}. \end{aligned} \tag{9.6}$$

We are interested in how C_t and I_t react to changes in the technology parameter A_t. We will look at relative changes first. The *elasticity* of a variable x with respect to another variable y is defined the percentage change in x in response to a one percent increase in y. Mathematically, elasticities can be computed as $\frac{\partial x}{\partial y}\frac{y}{x}$. Using this formula, the elasticity of consumption with respect to A_t is:

$$\frac{\partial C_t}{\partial A_t}\frac{A_t}{C_t} = \frac{\left(1 - \frac{1}{2}\alpha\right) A_t k_{t-1}^{1-\alpha}}{\left(1 - \frac{1}{2}\alpha\right) A_t k_{t-1}^{1-\alpha} + (1-\delta)k_{t-1}} < 1,$$

and for investment we get:

$$\frac{\partial I_t}{\partial A_t}\frac{A_t}{I_t} = \frac{\frac{1}{2}A_t \alpha k_{t-1}^{1-\alpha}}{\frac{1}{2}A_t \alpha k_{t-1}^{1-\alpha} - (1-\delta)k_{t-1}} > 1.$$

It turns out that the relative change in investment is larger. A one-percent increase in A_t leads to an increase of more than one percent in investment and less than one percent in consumption. Investment is more volatile in response to technology shocks, just as real-world investment is. Of course, to compare the exact size of the effects we would have to specify the parameters, like α and δ, and to measure the other variables, like k_t.

[3] More precisely, I_t in the model is *gross investment*, which includes replacement of depreciated capital. The net difference between capital tomorrow and today $k_t - k_{t-1}$ is referred to as *net investment*.

If we look at absolute changes instead of relative changes, the results are less satisfactory. The absolute change is higher in consumption than in investment, while in the real world it is the other way around. This failure of the model derives from the fact that people are too short-lived. In real business cycle models, the smaller variations in consumption relative to investment result from consumers trying to smooth their consumption. In our model, the possibilities for smoothing are rather limited. The old person has no more time left and therefore cannot smooth at all, while the young person has only one more year to go. Therefore a comparatively large fraction of the shock shows up in consumption. In more-advanced real business cycle models with infinitely lived consumers, the absolute changes in consumption are much smaller than the absolute changes in investment.

9.3 Simulations

We can get an even better impression of the business cycle in our model by simulating the economy. This means that we specify all parameters, start at some initial capital stock, and generate a series of random shocks. We can use the solutions to the model to compute consumption, investment, output, and the capital stock in the economy for any number of periods. Then we can compare the results to real-world business cycles.

There are only two parameters to be specified in the model, α and δ. Our choices are $\alpha = .7$ and $\delta = .05$. The choice for α matches the labor share in the economy to real world data[4], while the value for δ is an estimate of the actual average depreciation rate in an industrialized economy. The initial capital stock k_1 was set to .22. The productivity parameter was generated by:

$$A_t = \bar{A} + \epsilon_t.$$

Here $\bar{A}$ is the average level of productivity, while the ϵ_t are random shocks. We set $\bar{A} = 1$. The ϵ_t where generated by a computer to be independent over time and uniformly distributed on the interval $[-.1,\ .1]$. Thus the shocks can change productivity by up to ten percent upward or downward.

Figure 9.1 shows the reactions to a single productivity shock of five percent. That is, in the first period A_t is equal to its average, $A_1 = 1$. In the second period the shock hits, $A_2 = 1.05$. From then on, A_t is back to one and stays there. We can see that even this single shock has an impact that can be felt for a long period of time. Figure 9.1 shows the absolute deviations of consumption, investment, and capital from their average values. It takes about eight periods until all variables are back to their average. In the second period, when the shock takes place, both consumption and investment are up. In period 3 the capital stock is higher because of the higher investment in period 2. At the same time, investment falls. Consumption is higher than average because the capital stock is

[4] The *labor share* in an economy is defined to be total wages as a fraction of output. See Chapter 11 to see why α is equal to the labor share.

higher, even though productivity is back to normal again. From then on, all variables slowly return to their average values. Note that from period 4 on no one is alive anymore who was present when the shock took place. The higher investment in the period of the shock has increased the capital stock, and the effects of that can be felt for a long time. Thus even a single shock has long-run effects, and investment goes through a full cycle in response to this shock.

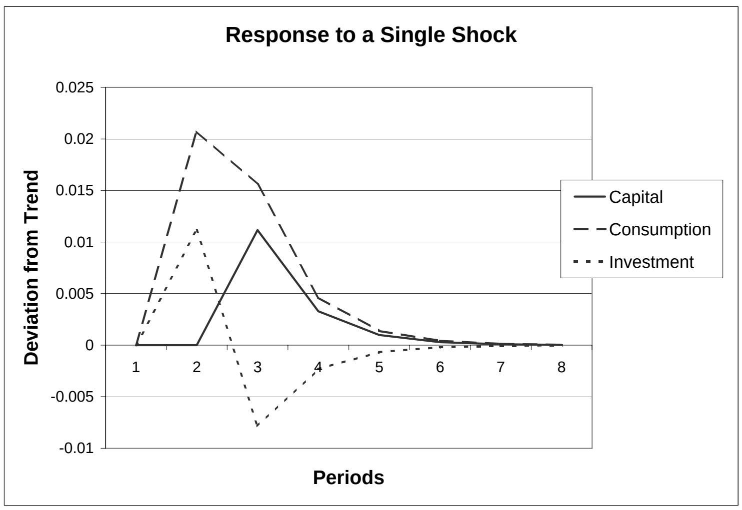

Figure 9.1: Response to a Five-Percent Productivity Shock

Figure 9.2 shows the same information as Figure 9.1, but variables are divided by their mean so that we can see the relative changes. Investment is by far the most volatile series. Compared to investment, the changes in capital and consumption are hardly visible.

By looking at a single shock, we were able to examine the propagation mechanism in isolation and to get an impression of the relative volatility of consumption and investment. But if we want to compare the model outcomes to real-world business cycles, we need to generate a whole series of shocks. Figure 9.3 shows such a simulation for our model economy. The combined effects of many shocks cause an outcome that looks similar to real-world business cycles. There are booms and depressions, the cycles vary in length within a certain interval, and investment is more volatile than consumption.

Our simple business cycle model is quite successful in emulating a number of business-cycle facts. Shape, length, and amplitude of business cycles are comparable to real-world data, investment is relatively more volatile than consumption, and the wage is procyclical. More-advanced real business cycle models are even better in matching the facts. By introducing variable labor supply we can generate procyclical employment. Using infinitely lived consumers would get the absolute changes in consumption and investment right.

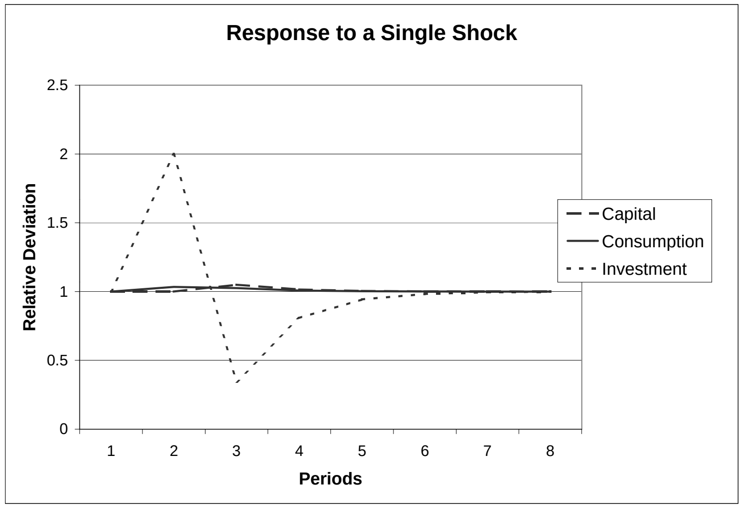

Figure 9.2: Relative Changes in Response to a Five-Percent Productivity Shock

State-of-the-art real business cycle models match most business cycle facts, and when fed with measured productivity shocks, they generate cycles that explain about 70% of the size of actual business cycles.

This success has led some researchers to the conclusion that business cycles are exactly what standard economic theory predicts. In the presence of shocks to production possibilities, optimal adjustments of households and firms within an efficient market system generate just the pattern of fluctuations that is observed in the real world. From this perspective, business cycles are no miracle at all. We would be surprised if there were no business cycles!

Even though technology shocks combined with efficient markets appear to provide a convincing explanation for business cycles, it cannot be ruled out that other shocks or propagation mechanisms also play a role. After all, real business cycle theory does not account for 100% of the amplitude of actual business cycles, so there have to be other factors as well. Other types of shocks can be analyzed within the real business cycle framework. There are also a number of models that emphasize other propagation mechanisms. The Keynesian model of output determination is the most prominent example[5] but models that combine monetary shocks with frictions in the financial sector have also received a lot of attention lately. However, so far none of these models matches the ability of the real business cycle model to mimic actual economic fluctuations.

[5]See Barro's Chapter 20 for a discussion of that model. For a critique of the real business cycle model from a Keynesian perspective, try the following article: Mankiw, N. Gregory. 1989. "Real Business Cycles: A New Keynesian Perspective". *Journal of Economic Perspectives* 3(3): 79-90.

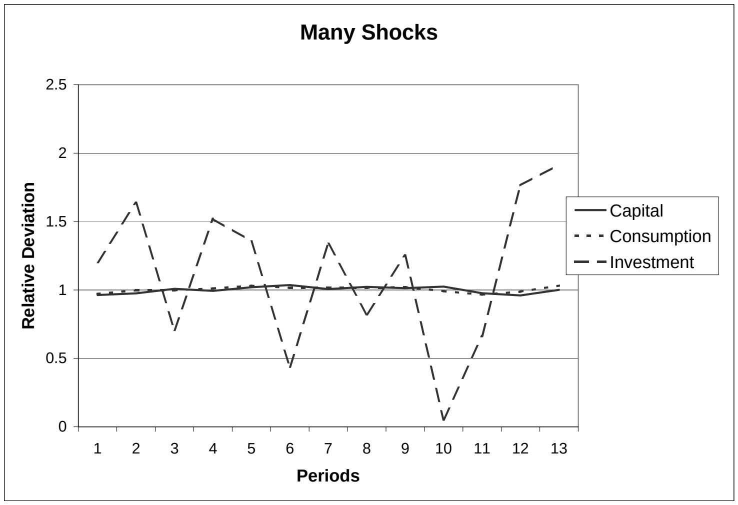

Figure 9.3: Capital, Consumption, and Investment with Many Shocks

Variable	Definition
c_t^t	Consumption of generation t when young
c_{t+1}^t	Consumption of generation t when old
$u(\cdot)$	Utility function
k_t	Capital saved in t and used in $t+1$
l_t	Labor
w_t	Wage
r_t	Rental rate of capital
$f(\cdot)$	Production parameter
A_t	Productivity parameter
α	Parameter in the production function
δ	Depreciation rate
C_t	Aggregate consumption
I_t	Aggregate investment
Y_t	Aggregate output
ϵ	Random shock

Table 9.1: Notation for Chapter 9

Exercises

The following exercises make up a project that can be done in groups or individually.

Exercise 9.1 (Moderate)
As the word "cycle" indicates, for a long time economists thought of business cycles as regular, recurrent events. The length and severity of business cycles was thought to be mostly constant. For example, the typical length of one full cycle (from boom through recession back to boom) was supposed to be between four and seven years. In this question you will examine the actual business cycles of a country of your choice and examine whether they seem to follow a regular pattern.

The first thing to do is to get the necessary data. Business cycles are roughly defined as deviations of real GDP from trend. Therefore you will need to acquire data on real GDP for some country. A good source is the Penn World Tables, a set of standardized measures of economic activity for most countries in the world. You can access the World Tables through a website at the University of Toronto. The address is:

`http://arcadia.chass.utoronto.ca/pwt/`

Once you are there, select "Alphabetical List of Topics", then "Real GDP per capita in constant dollars using chain index", then click on the country of your choice (not the United States), then use the "Submit Query" button to get the data. Load the data into a spreadsheet, and you are ready to go.

The first step is to compute the trend component of GDP. Good methods for computing the trend of a time series require a relatively high amount of complicated computations. Therefore we will offer you an ad hoc, quick-and-dirty method of computing the trend. Once we get to the business cycles, it turns out that this method works sufficiently well for our purposes. We will use GDP_t to denote real GDP at time t. The computation of the trend proceeds in steps:

- Compute the growth rate of GDP for each year. In terms of your spreadsheet, let us assume that column A is year and column B is real GDP. The first year is in row one. Now you can put the growth rates into column C. Put the growth rate from year 1 to 2 into cell C1, and so on.
- From now on, we are going to apply a method called *exponential smoothing* to get smooth versions of our data. Assume you want to get a smooth version of a times series x_t. Let us call the smooth version $\hat{x}_t$. Basically, the $\hat{x}_t$ are computed as a forecast based on past observations of x_t. The first $\hat{x}_t$ is set equal to the first x_t: $\hat{x}_1 = x_1$. From then on, the forecasts for the next period are computed as an average of the last forecast and the actual value: $\hat{x}_{t+1} = \beta\hat{x}_t + (1-\beta)x_t$, where β is a number between

zero and one. If you plug this formula recursively into itself, you will see that each $\hat{x}_t$ is a weighted average of past x_t.

Let us now put a smooth growth rate into column D. Since $\hat{x}_1 = x_1$, the first smooth value is equal to the original value: D1=C1. For the next value, we apply the smoothing formula. We recommend that you set β to .5: D2=.5*D1+.5*C1. In the same way, you can get the other smoothed growth rates. For future reference, We will call the smooth growth rates $\hat{g}_t$.

- In the next step, we are going to apply the same method to real GDP, but additionally we will use the smooth growth rates we just computed. This smooth real GDP is the trend we are looking for, and we will place it in column E. As before, in the first year the smooth version is identical to the original one: $\text{Trend}_1 = \text{GDP}_1$, thus E1=B1. From then on, we get the trend in the next period by averaging between the trend and the actual value (as before), but also applying the smooth growth rates we just computed. If we do not do that, our trend will always underestimate GDP. From year two on the formula is therefore:

$$\text{Trend}_{t+1} = (1 + \hat{g}_t)(0.5)\text{Trend}_t + (0.5)\text{GDP}_t.$$

In terms of the spreadsheet, this translates into E2=(1+D1)*(.5*E1+.5*B1), and so on.

This completes the computation of the trend. Plot a graph of GDP and its trend. If the trend does not follow GDP closely, something is wrong. (Document your work, providing spreadsheet formulas, etc.)

Exercise 9.2 (Moderate)
Now we want to see the cyclical component of GDP. This is simply the difference between GDP and its trend. Because we are interested in relative changes, as opposed to absolute changes, it is better to use log-differences instead of absolute differences. Compute the cyclical component as $\ln(\text{GDP}) - \ln(\text{Trend})$. Plot the cyclical component. You will see the business cycles for which we have been looking. (Document your work, providing spreadsheet formulas, etc.)

Exercise 9.3 (Easy)
Now we will examine the cycles more closely. Define "peak" by a year when the cyclical component is higher than in the two preceding and following years. Define "cycle" as the time between two peaks. How many cycles do you observe? What is the average length of the cycle? How long do the shortest and the longest cycles last? Do the cycles look similar in terms of severity (amplitude), duration, and general shape? (Document your work, providing spreadsheet formulas, etc.)

Exercise 9.4 (Moderate)
Having seen a real cycle, the next step is to create one in a model world. It turns out that doing so is relatively hard in a model with infinitely lived agents. There we have to deal with uncertainty, which is fun to do, but it is not that easy as far as the math is concerned.

Therefore our model world will have people living for only one period. In fact, there is just one person each period, but this person has a child that is around in the next period, and so on. The person, let us call her Jill, cares about consumption c_t and the bequest of capital k_{t+1} she makes to her child, also named Jill. The utility function is:

$$\ln(c_t) + A\ln(k_{t+1}),$$

where $A > 0$ is a parameter. Jill uses the capital she got from her mother to produce consumption c_t and investment i_t, according to the resource constraint:

$$c_t + i_t = \sqrt{Bk_t} + \epsilon_t,$$

where $B > 0$ is a parameter, and ϵ_t a random shock to the production function. The shock takes different values in different periods. Jill knows ϵ_t once she is born, so for her it is just a constant. The capital that is left to Jill the daughter is determined by:

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

where the parameter δ, the depreciation rate, is a number between zero and one. This just means that capital tomorrow is what is left over today after depreciation, plus investment.

Compute Jill's decision of consumption and investment as a function of the parameters k_t, and ϵ_t.

Exercise 9.5 (Moderate)
If we want to examine the behavior of this model relative to the real world, the next step would be to set the parameters in a way that matches certain features of the real world. Since that is a complicated task, we will give some values to you. B is a scale parameter and does not affect the qualitative behavior of the model. Therefore we set it to $B = .1$. δ is the depreciation rate, for which a realistic value is $\delta = .05$. A determines the relative size of c_t and k_t in equilibrium. A rough approximation is $A = 4$. Using these parameters, compare the reactions of c_t and i_t to changes in ϵ_t. (Use calculus.)

Exercise 9.6 (Moderate)
In the last step, you will simulate business cycles in the model economy. All you need to know is the capital k_1 at the beginning of time and the random shocks ϵ_t. As a starting capital, use $k_1 = 3.7$. You can generate the random shocks with the random number generator in your spreadsheet. In Excel, just type "=RAND()", and you will get a uniformly distributed random variable between zero and one. Generate 50 such random numbers, and use your formulas for c_t and i_t and the equation for capital in the next period, $k_{t+1} = (1-\delta)k_t + i_t$, to simulate the economy. Plot consumption and investment (on a single graph). How does the volatility of the two series compare? Plot a graph of GDP, that is, consumption plus investment. How do the business cycles you see compare with the ones you found in the real world? You don't need to compute the length of each cycle, but try to make some concrete comparisons.

Exercise 9.7 (Easy)
Read the following article: Plosser, Charles. 1989. "Understanding Real Business Cycles". *Journal of Economic Perspectives* 3(3): 51-78. Plosser is one of the pioneers of real business cycle theory. What you have done in the previous exercises is very similar to what Plosser does in his article. His economy is a little more realistic, and he gets his shocks from the real world, instead of having the computer draw random numbers, but the basic idea is the same.

Describe the real business cycle research program in no more than two paragraphs. What question is the theory trying to answer? What is the approach to answering the question?

Exercise 9.8 (Moderate)
What does Plosser's model imply for government policy? Specifically, can the government influence the economy, and is government intervention called for?