5. DYNAMIC EQUILIBRIUM MODELS III: BUSINESS-CYCLE MODELS

It is perfectly true, what philosophy says, that life must be understood backwards.

– Søren Kierkegaard (1813–1855)

In the growth models we studied in section 4 we focussed on steady-state or balanced growth and on what might determine long-run average rates of growth of consumption, output, wages, and so on. But of course macroeconomics also is concerned with the fact that we do not observe steady growth but instead see irregular cycles in most aggregates, perhaps around some trend.

To produce such cycles in the growth models we have studied we need some shocks (changes in some exogenous variable such as population or technical progress, though we would like to explain those variables ultimately too). Business cycles recur and hence we often think of those shocks as being random variables – each period we observe some draw from a probability density function. In business-cycle modelling, we usually add shocks to one of the growth models and then regard the transitional dynamics as a business cycle.

This section of the notes introduces tools we use in studying business-cycle models: dynamic programming, impulse response functions, calculating moments, and the generalized method of moments (GMM). In studying cycles the two workhorse models are the OLG model and the representative-agent model, just as was the case in studying growth. For simplicity, this section focuses on the representative-agent model. When we used this model to study growth in section 4, it had a constant saving rate (the Solow-Swan model). We'll begin this section by dropping that simplification, to make the model more realistic. Also, we'll set the model in discrete time, so it can be applied to data easily. Later on, we'll add shocks to the model and see whether it produces realistic cycles.

(a) Introduction to Dynamic Programming: The Neoclassical Growth Model

You can think of the neoclassical growth model two ways: it is the multi-period version of the two-period model in section 2, and it is the growth model of section 4 but with an explicit utility function instead of a fixed saving rate. I continuous time it is sometimes called the Ramsey-Cass-Koopmans model. We'll work with the discrete-time version.

Let us begin with an example with no random variables in it, to make the mathematics as simple as possible. Consider the problem of choosing a sequence of values for consumption $\{c_t\}$, to maximize

$$\sum_{t=0}^{T} \beta^t \ln(c_t),$$

subject to

$$c_t + k_{t+1} = y_t = k_t^{\alpha}.$$

You might recognize that this is simply a multi-period version of our two-period production economy from section 2. Like that example, this one has log utility, Cobb-Douglas production technology, and 100 percent depreciation of capital. The only way to have capital next period is to set some aside this period.

We could generalize the method we used in section 2 by forming a giant Lagrangean, and then maximizing. But this will give us a very large number of equations to solve when T is large. A better method is to use *dynamic programming*, which breaks this into a sequence of two-period problems we can solve easily.

We won't go into much mathematical detail here, but by solving this problem we will see the logic of dynamic programming, which extends to cases with an infinite horizon and with uncertainty. The first step in solving this problem will be to start at the last period, period T, and work backwards. Suppose that we enter period T with a capital stock k_T. What is the optimal thing to do, contingent on this stock? The utility function is additively separable over time, so we don't need to refer to earlier time periods to know how c_T affects utility. And utility is increasing, so of course the optimal decision is:

$$c_T = k_T^\alpha.$$

Live like there is no tomorrow. There is no reason to save, since we are in the last period.

The maximized value of the plan is just $\ln(c_T)$, which is equal to $\alpha \ln(k_T)$. Although this might not be immediately obvious, this is just an indirect utility function. In dynamic programming it is called the *value function*, and we denote it by V_T. It tells us the utility we get in period T from doing the optimal thing, starting with capital k_T.

At time T,

$$V_T(k_T) = u[f(k_T)] = \ln(k_T^\alpha].$$

By the chain rule of calculus, that means that

$$V_T'(k_T) = u'(c_T)f'(k_T),$$

which you can easily verify. This relationship is called the *envelope condition*, and we will use it again later.

This is all contingent on the investment decision at $T-1$. So the next step is to work backwards to the second-last period. Suppose that we enter this period with a capital stock of k_{T-1}. What is the best plan in this period? Now our problem is to maximize

$$\ln(c_{T-1}) + \beta\alpha \ln(k_T)$$

subject to

$$c_{T-1} + k_T = k_{T-1}^\alpha.$$

(We also can state the problem as maximizing $\ln(c_{T-1}) + \beta V_T(k_T)$.) But you should recognize that this is simply our two-period problem from section 2. Substituting the constraint in the objective – treating k_{T-1} as given – gives us a problem in c_{T-1}, and the first-order condition gives

$$c_{T-1} = \frac{k_{T-1}^{\alpha}}{1+\alpha\beta},$$

just as in section 2.

Once again we can find the value function, which now gives the value of the plan from $T-1$ onwards. It is

$$V_{T-1}(k_{T-1}) = \text{constant} + (\alpha + \alpha^2\beta)\ln(k_{T-1})$$

Everything will depend on the derivative of the value function, so we do not need to keep track of the constants.

Although you may not be happy to hear this, it now is time to consider the choice made at time $T-2$. Suppose that we enter this period with capital stock k_{T-2}. How should we divide output into consumption and investment? The problem now is to maximize

$$\ln(c_{T-2}) + \beta V_{T-1}(k_{T-1}),$$

subject to

$$c_{T-2} + k_{T-1} = k_{T-2}^{\alpha}.$$

Following the same steps gives us

$$c_{T-2} = \frac{1}{1+\alpha\beta+\alpha^2\beta^2} k_{T-2}^{\alpha}.$$

We also could find the value function, which once again summarizes the consequences of our decision for utility in future time periods.

Where is this leading? At each time period we are solving for the *policy function* that tells us the value of c_t given k_t. If we continued to work backwards, eventually, we would notice that the optimal policy function converged to

$$c_t = (1-\alpha\beta)k_t^{\alpha}.$$

Once we have solved for this function, we can study the model. Suppose we begin with k_0. We use the optimal policy to divide k_0^{α} into consumption and investment, and then the constraint gives us k_1, and so on. In this way, we construct sequences for consumption and investment.

The value function also converges in this example. Moroever, it satisfies the key equation of dynamic programming, called Bellman's equation:

$$V(k_t) = \max[u(c_t) + \beta V(k_{t+1})],$$

which is exactly the recursive relationship we've just been using, except that we've dropped the time subscripts.

The idea behind this is straightforward. The value of the optimal plan at time t involves maximizing today's utility plus the discounted sum of utility from all future time periods. The underlying principle is sometimes called the *principle of optimality*. It means simply that if we have a function to maximize with respect to each of $\{c_0, c_1, c_2, \ldots, c_T\}$ we can maximize it with respect to c_T, then substitute that value into the function and maximize with respect to c_{T-1}, and so on.

If we begin with Bellman's equation, we can state our problem as:

$$V(k_t) = \max[\ln(c_t) + \beta V(k_{t+1})]$$

subject to

$$k_{t+1} = k_t^\alpha - c_t.$$

In this problem k_t is called the *state variable*, and the second equation gives its law of motion. The state variable tells us everything we need to know about initial conditions carried into each period.

We have already found the optimal policy function for this problem, but let us briefly consider this formal statement of the problem so that we can see how to find answers in other problems. The first-order condition (found by differentiating with respect to c_t) is in general

$$u'(c_t) + \beta V'(k_{t+1}) \frac{dk_{t+1}}{dc_t} = 0,$$

where we have simply used the chain rule of calculus. Now it is time to use the envelope condition, which states that

$$V' = u' \cdot f'.$$

The combination of these two equations gives:

$$u'(c_t) = \beta f'(k_{t+1}) u'(c_{t+1}),$$

which is the now painfully familiar Euler equation.

In our problem, notice that $f'(k_t) = \alpha k_t^{\alpha-1} = \pi = \delta + r_{t+1} = 1 + r_{t+1}$. And with log utility the result is:

$$\frac{1}{c_t} = \beta(1 + r_{t+1}) \frac{1}{c_{t+1}}.$$

We've now proved that this applies to adjacent time periods in the multi-period case.

Suppose that instead of figuring out the optimal policy by working backwards and looking for convergence we instead guessed the form of the policy and then tried to work it out exactly. This is analogous to the method of undetermined coefficients in section 3. In our problem, suppose that we guess that the best policy is to consume a constant fraction, say γ, of output each period and to save and invest the rest.

Substituting our guess into the Euler equation gives:

$$\frac{1}{\gamma k_t^\alpha} = \frac{\beta \alpha k_{t+1}^{\alpha-1}}{\gamma k_{t+1}^\alpha}.$$

Thus

$$\frac{1}{k_t^\alpha} = \frac{\beta\alpha}{k_{t+1}}.$$

But remember that

$$k_{t+1} = k_t^\alpha - c_t = k_t^\alpha(1-\gamma),$$

and so $\gamma = 1 - \alpha\beta$.

The guess-and-verify method is certainly easy. And we also could guess and verify the value function instead of the policy function (also called the decision rule). The other advantage is that it can be used in problems with an infinite horizon, where we have no terminal condition at T from which to begin.

Usually exogenous state variables are modelled as *Markov* processes, a glamourous way of saying that the forecasts of future values depend only on the current value (they follow first-order autoregressions). That is a way of keeping the state vector from becoming too large. Recall the rule-of-thumb in the method of undetermined coefficients in section 3: the guess has one lag less than the number of lags in the process for the exogenous variable. The same rule applies here, so if a state variable depends on its previous value then the guess includes the current value.

Exercise: Consider the problem of consumption and saving out of initial wealth. The goal is to maximize

$$\sum_{t=0}^{\infty} \beta^t \ln(c_t),$$

subject to

$$w_t = R(w_{t-1} - c_{t-1}),$$

where w_t is wealth, R is the gross interest rate, and w_0 is given. See if you can show that the optimal policy is $c_t = (1-\beta)w_t$ using the guess-and-verify method. •

To make this a growth model, let us give it an infinite horizon, and also introduce some labour-augmenting productivity growth. Households maximize:

$$\mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t \ln(c_t),$$

subject to

$$c_t + k_{t+1} = y_t = k_t^\alpha[(1+\lambda)^t n_t]^{1-\alpha},$$

where there is no population growth and $n_t = 1$. This special case, with log utility, Cobb-Douglas production, and 100% depreciation, is called the Brock-Mirman model. We could study questions in growth economics with this model, just as in section 4; the exercises invite you to do this.

We've now studied several different growth models. The following table is designed to make sure that you don't confuse them:

Growth Models

	OLG	Solow-Swan	*AK*	Brock-Mirman
horizon:	finite	infinite	infinite	infinite
time:	discrete	continuous	continuous	discrete
saving:	u-based	constant	constant	u-based
growth:	exogenous	exogenous	endogenous	exogenous

The entries in the AK column correspond to the simple, endogenous growth model we studied at the end of section 4, which was a modified Solow-Swan model. Please note that we also could construct AK versions of the OLG model or of the Brock-Mirman model.

Finally, you should be aware that very few dynamic programmes can be solved analytically like the examples we've seen here. Almost always the solution is approximated and/or found using numerical methods.

(b) Application: Real Business Cycles

To make the growth model into a business-cycle model, we simply add some shocks. We'll begin with a single shock – to technolgy – and consider other shocks later. With this single shock the model is sometimes called a *real business cycle* model. We'll learn about the predictions of that model and, more generally, what steps to follow in studying any business-cycle model.

The distinguishing feature of rbc models is that they attach an important role to technology shocks in explaining cyclical fluctuations. In fact, early versions contained only technology shocks. This contrasted sharply with some traditional macroeconomics, which attributed some part of cycles to monetary and fiscal policies. However, recent business-cycle models contain both types of shocks, so this dichotomy is not really sharp.

The other distinguishing feature of the models is that they explain cycles using the stochastic growth model (multi-period competitive equilibrium model). This seems a logical place to begin, because that model is used in studying growth, asset prices, and so on.

Suppose that our problem is now to maximize

$$\mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t \ln(c_t),$$

subject to

$$c_t + k_{t+1} = y_t = z_t k_t^{\alpha},$$

where $z_t = exp(\epsilon_t)$, with ϵ_t iid and unpredictable, with mean zero and variance σ^2. The random variable z_t is called a technology shock or a productivity shock.

Now the value function must satisfy,

$$V(k_t, z_t) = \max[\ln(c_t) + \beta \mathrm{E}_t V(k_{t+1}, z_{t+1})].$$

There are now two state variables, one endogenous (k) and one exogenous (z). The expectation on the right-hand side depends on forecasts of both of them. You should be able to write down the Euler equation directly. In fact, in this case it turns out that the optimal policy is simply

$$c_t = (1 - \alpha\beta) z_t k_t^\alpha,$$

which you can verify if you wish.

The fundamental equation in the capital stock now is a stochastic difference equation:

$$\ln(k_{t+1}) = \ln(\beta\alpha) + \alpha \ln(k_t) + \epsilon_t.$$

This is a first-order autoregression, with a root of α.

This model does not have growth, the cycles are not very persistent (since α is less than 0.5), and employment does not fluctuate. Nevertheless, we can use it to see how we might study cycles in models without these defects.

Exercise: Find the variance of log output, the variance of log consumption, and their covariance.

Exercise: Find the covariance between the log real wage and log output.

To calculate moments, such as variances, remember that we know the moments of the exogenous shock: ϵ_t has mean zero and variance σ^2. So if we can write each endogenous variable – capital, output, the real wage – in terms of current and lagged shocks then we can calculate their moments.

To do this, we simply solve the relevant difference equation backwards. For example, ignoring the constant term (which will not affect the variance), we know that the capital stock follows:

$$\ln(k_{t+1}) = \alpha \ln(k_t) + \epsilon_t.$$

If we use repeated substitution backwards we can write this in terms of current and past shocks:

$$\ln(k_{t+1}) = \epsilon_t + \alpha\epsilon_{t-1} + \alpha^2\epsilon_{t-2} + \ldots$$

Thus

$$\mathrm{var}[\ln(k_{t+1})] = \sigma^2(1 + \alpha^2 + \alpha^4 + \ldots) = \frac{\sigma^2}{1 - \alpha^2}.$$

A useful tool that saves some time in this sort of calculation is the *lag operator*, denoted L. This provides a shorthand for keeping track of lags and leads. Multiplying a variable by L lags it once, and multiplying it by L^i lags it i times. So the difference equation in capital (again without the constant term) is:

$$\ln(k_{t+1}) = \alpha L \ln(k_{t+1}) + \epsilon_t.$$

L can be treated just like a constant, so

$$\ln(k_{t+1})(1-\alpha L) = \epsilon_t,$$

and

$$\ln(k_{t+1}) = \frac{\epsilon_t}{1-\alpha L} = \epsilon_t(1+\alpha L+\alpha^2 L^2 + ...),$$

which is our earlier result.

Exercise: Graph the impulse response function of log output in response to a log productivity shock, ϵ_t.

An impulse response function is simply the graph of what happens to y_t, y_{t+1}, y_{t+2}, and so on, in response to a single shock ϵ_t. It has y on the vertical axis and time after the shock on the horizontal axis.

You can see that the properties of output and consumption depend on the properties of the shocks, as well as on the parameters of the model. How do we know the properties of the shocks? One possibility is to measure them using Solow residuals. Such measurements usually yield quite persistent technology shocks.

Exercise: Suppose that we model ϵ_t as

$$\epsilon_t = 0.8\epsilon_{t-1} + \nu_t$$

where ν_t is white noise. What are the moments and time series properties of output now?

In other cases, we may want to treat shocks as genuinely unobservable. But then how can we make predictions with or test the business-cycle model? The answer is that we can restrict the stochastic process followed by the shock, and then estimate its parameters. This sounds complicated, but an example will convince you that it is straightforward. Suppose that we assume that ϵ_t follows a first-order autoregression with parameter ρ. Then,

$$\ln(y_t) = (\alpha+\rho)\ln(y_{t-1}) - \alpha\rho\ln(y_{t-2}) + \nu_t.$$

We can estimate α and ρ easily from this equation, using the log of output and its lagged values.

The estimation here is by ordinary least squares. But we also could estimate the parameters by the method of moments. When we calculate the moments of output and consumption, we find that they depend on the unknown parameters. That gives us identification: we can estimate the parameters by matching the model's moments with the moments we find in historical data.

Recall from section 1 that to construct moments in the data we need variables that are stationary. What *detrending* method should we use? Probably the most natural method is to use the model as a guide to detrending. For example, suppose we keep a trend in labour-augmenting technology, denoted $(1+\lambda)^t$ in the production function. Then the model will produce cycles around a balanced growth path with growth rate λ. We could estimate this common trend in time series for output, consumption, and investment, and then isolae the remaining cycles. If instead we modelled the exogenous technological change as a random walk, say, then the endogenous variables would inherit that pattern, and we would detrend by taking differences.

How successful is the basic real business cycle model? Let's look at three of its predictions. First, it succeeds in reproducing the fact (seen in section 1) that investment is more volatile than output, and that output in turn is more volatile than consumption. Our Brock-Mirman special case doesn't have this feature, but versions with more realistic depreciation rates do. Second, the model does not have much propagation, as we've seen. Most of dynamics in output come from the shocks rather than from capital accumulation, which is the propagation mechanism. Third, the model produces a strongly procyclical real wage, whereas for most countries the real wage is at most mildly procyclical.

From these comparisons between theory and evidence, researchers have been led to reformulate the model in a variety of ways. For example, we can add variable labour supply and a second shock – to government spending – to the model. Fiscal shocks have several effects. First, an increase in g raises the interest rate. This is the price signal by which the government commands a greater share of resources. Second, with elastic labour supply, the same shock lowers the real wage. The mechanism is that greater government spending makes people poorer because taxes rise too, and so they increase their labour supply. With productivity shocks, the real wage is procyclical, as the labour demand curve shifts back and forth. But with government spending shocks the real wage may be countercyclical, because an increase in labour supply lowers real wages but raises output. A mixture of the two shocks may give an acyclical real wage, which is realistic empirically. Exercise 15 asks you to work out this reformulation explicitly.

By now you should see that there are several steps involved in most business-cycle modelling. First, one writes down a model, which is the hard part. Second, one solves it. Our example here has an analytical solution, but most models don't and so the solution must be approximated. Third, one assigns values to the parameters, perhaps by econometric estimation or by calibration. Fourth, one calculates predictions or implications from the model – such as moments, impulse response functions, or sample paths – and compares these to historical evidence. Fifth, one reformulates the model.

(c) Introduction to GMM

As mentioned earlier, dynamic programmes usually can't be solved analytically, so numerical approximations are used. However, there is an alternative way to learn about some of the implications of a model that sometimes can tell us almost as much, without requiring as much information. In most models we can study some implications without completely solving the model or even specifying what the shocks are. The implications usually are first-order conditions. An example is the usual Euler equation, which holds even if we do not know how the productivity shock evolves. We can use that equation to estimate parameters and also to test the model, using the generalized method of moments (GMM). You'll certainly come across GMM if you do any reading in applied macroeconomics.

Here is the idea. With log utility and a representative agent the Euler equation for aggregate consumption is:

$$\frac{1}{c_t} = E_t \beta (1 + r_{t+1}) \frac{1}{c_{t+1}}.$$

Remember from introductory statistics that the E operator simply calculates averages. So the theory predicts that the average of the expression on the right is equal to the expression on the left. We can calculate that average, and set the value of β so the two sides match as closely as possible, say by minimizing the squared errors just as in ordinary least squares. The resulting value is our estimate, $\hat{\beta}$.

The same idea can be used to test theories and see where they are successful and where they are not. Euler equations are necessary conditions, so while satisfying them doesn't mean a model is a good guide in all respect, rejecting them is enough to send us back to the drawing board.

You may recall from section 2 that the Euler equation holds for returns on many assets, so our $\hat{\beta}$ should match up the two sides whatever return series we use. The extent to which it does this serves as a test of the theory.

A second way to test the theory is to use the law of iterated expectations. Notice that the Euler equation uses E_t, the conditional expectation operator. That means that the forecasted value of the right-hand side, conditional on any information at time t, should equal the left-hand side. So if we try forecasting the right-hand side, using a regression and several different regressors, then in each case the fitted value should equal the left-hand side. Again this provides a test.

To formalize all this define an error term:

$$e_{t+1} = \beta (1 + r_{t+1}) \frac{1}{c_{t+1}} - \frac{1}{c_t}.$$

Then this error must satisfy these conditions:

$$\mathrm{E} e_{t+1} = 0$$
$$\mathrm{E}(e_{t+1} \cdot x_t) = 0.$$

because of rational expectations, where x_t is any variable known at time t. These variables sometimes are called instruments. We can use the corresponding sample versions to estimate the parameter β;

$$\sum_{t=0}^{T-1} e_{t+1} = 0$$
$$\sum_{t=0}^{T-1} (e_{t+1} \cdot x_t) = 0.$$

Imagine simply varying β until the first condition is satisfied. Then the second condition can be used as a test (based on overidentification). The same value of β should satisfy this equation too.

When there is overidentification – more moment-condition equations than parameters to estimate – then the moment conditions usually can't be satisfied exactly. In that case, how close they are to zero provides a test of the model. To measure 'closeness' we need a standard error. In this course we won't study how to calculate this. But from introductory statistics, you should be familiar with the idea that a sample moment – such as $\bar{x}$ – has a standard error.

Remember that both c_t and r_t are endogenous, so this is like instrumental variables estimation. The actual estimation works by minimizing a quadratic form in the moment conditions – trying to make them as close to zero as possible. The key feature is that we don't need to completely solve the model in order to estimate some of its parameters and test some elements of it. That can speed up reformulating the model. This method of investigation does not tell us time paths for the variables or allow us to forecast them. But by testing necessary conditions it can usually tell us something about improving the model.

A classic application of GMM is to the equity premium puzzle: linking both bond and equity returns to aggregate consumption. The readings in section 2 expand on this example.

A second good example of a study which examines Euler equations is Hall's (1980) paper on labour supply (see also the comment by Barro). Hall noted that the real wage is roughly acyclical and yet labour supply is procyclical. Moreover, estimates of the wage elasticity of labour supply usually are quite small. How can these facts be reconciled with labour supply models? In other words, if there are not large temporary movements in real wages then what causes movements in labour supply?

Hall suggested a different price signal which might influence labour supply: the interest rate. If you look at a standard two-period model (as in section 2) you will see that an increase in the interest rate reduces current consumption and increases labour supply.

In models with shocks to productivity remember that a shock shifts the labour demand curve but also changes the interest rate, which may thus change labour supply. Incidentally, this shows a problem with dividing shocks into two categories called 'supply' and 'demand' – some shocks affect both curves.

Exercise: For a utility function of your choice, derive an Euler equation linking labour supply in successive periods. Discuss how this could be estimated. Discuss the effect of the interest rate.

There is a potential problem with this channel too though. If you chose a logarithmic utility function, separable in consumption and leisure, then one of the Euler equations is:

$$c_t = (1 - n_t)w_t.$$

If w_t is acyclical then c_t and n_t should move inversely (because both consumption and leisure are normal goods with this utility function). In fact, they seem to be positively correlated in business cycles.

A good example of a study which estimates and tests a complete set of Euler equations for a household planning problem is the one by Mankiw, Rotemberg, and Summers (1985). They find that over-identifying restrictions are rejected; in other words, there does not seem to be a set of preference parameters consistent with all the Euler equations and instruments x_t they examine.

In the next section of the course, we'll see that many studies of aggregate consumption focus on finding an Euler equation that holds up in data.

Further Reading

On dynamic programming, a good place to begin is with a Ph.D.-level theory book, such as David Kreps's *A Course in Microeconomic Theory* (1990), appendix two, or Andreu Mas-Colell, Michael Whinston, and Jerry Green's *Microeconomic Theory* (1995), appendix M.N. I next recommend chapter 1 of Thomas Sargent's *Dynamic Macroeconomic Theory* (1987). Two fine books not written specifically for economists are by Ronald Howard *Dynamic Programming and Markov Processes* (1960) and Sheldon Ross *Introduction to Stochastic Dynamic Prgramming* (1983).

A fine survey on real business cycles is by George Stadler 'Real business cycles,' *Journal of Economic Literature* (1994) 1750-1783. Two equally readable but opposite perspectives are found in the *Journal of Economic Perspectives* (1989): Charles Plosser's 'Understanding real business cycles' 51-77 and Greg Mankiw's 'Real business cycles: a new Keynesian perspective' 79-90. For textbook discussions, see David Romer's *Advanced Macroeconomics* (1996) chapter 4 and Olivier Blanchard and Stanley Fischer's *Lectures on Macroeconomics* (1989) chapter 7. An excellent example of constructing business cycle models is provided by Robert King, Charles Plosser, and Sergio Rebelo in 'Production, growth, and business cycles: I the basic neoclassical model' *Journal of Monetary Economics* (1988) 195-232. For Solow's perspective on the use of the neoclassical growth model as a model of cycles, see his essay 'Growth theory and after' *American Economic Review* (1988) 307-317.

On GMM, a fine survey of the econometrics and some applications is given by Masao Ogaki, in "Generalized method of moments: econometric applications," chapter 17 in the *Handbook of Statistics* (1993) volume 11. Hall's study of first-order conditions is "Labor supply and aggregate fluctuations," *Carnegie Rochester Conference Series on Public Policy* (1980), 12, 7-38. A study of several Euler equations is undertaken by Gregory Mankiw, Julio Rotemberg, and Lawrence Summers, "Intertemporal substitution in macroeconomics," *Quarterly Journal of Economics*, (1985) especially 225-232.

Exercises

1. Consider the following problem of labour supply:

$$\max_{\{c_1,n_1,c_2,n_2\}} \ln(c_1) + \ln(1-n_1) + \beta \ln(c_2) + \beta \ln(1-n_2)$$

The intertemporal budget constraint is

$$c_1(1+r) + c_2 = w_1 n_1 (1+r) + w_2 n_2$$

and the household takes w_1, w_2, and r as given.

(a) Solve for c_1, n_1, c_2, and n_2 in terms of prices.

(b) Find the response of labour supply in the first period to an increase in the real interest rate.

2. Consider a two-period world in which consumers maximize their objective function

$$U = ln(c_1) + \beta ln(c_2) - \frac{n_1^\gamma}{\gamma} - \frac{\beta n_2^\gamma}{\gamma}$$

subject to the budget constraint,

$$(w_1 n_1 - p_1 c_1)(1+i_1) + w_2 n_2 - p_2 c_2 = 0$$

where c is goods consumption, n is the supply of labour, p is the price of goods and w is the wage. Subscripts 1 and 2 denote periods 1 and 2 respectively.

(a) Set up and solve the Lagrangean, choosing consumption and labour supply in each period.

(b) Derive expressions for relative consumption and labour supply.

(c) How does the relative supply of labour in period 1 respond to a temporary increase in wages? How does that response itself change when γ decreases?

(d) How does the relative supply of labour in period 1 respond to a permanent wage increase? (To keep things simple, consider the case of a multiplicative increase – *e.g.* a doubling of the wage rate)

(e) Derive expressions for the individual demands for consumption in each period and for the supply of labour in each period. Show that the labour demand functions satisfy the results in parts (c) and (d).

3. This question concerns the interpretation of co-movements in aggregate consumption, hours, and real interest rates. Suppose that the period utility function is:

$$u(c_t, n_t) = \frac{c_t^{1-\alpha}}{1-\alpha} + \frac{(1-n_t)^{1-\alpha}}{1-\alpha}; \quad \alpha > 0$$

where c_t is consumption, n_t is hours worked. Suppose that a representative consumer maximizes

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

with $\beta \in (0, 1)$, subject to a budget constraint

$$a_{t+1} = (1 + r_t)a_t + w_t n_t - c_t$$

for all $t > 0$. There is no uncertainty and the real wage w and real interest rate r are taken as given by the consumer.

(a) Write down the the Euler equations describing the tradeoffs between c_{t+1} and c_t, n_{t+1} and n_t, and c_t and n_t respectively.

(b) Show that for these preferences the elasticity of intertemporal substitution in consumption is equal to the inverse of the coefficient of relative risk aversion.

(c) Show that the response of hours is greater to a temporary increase in the real wage than to a permanent one. What is the economic intuition into this result?

(d) Suppose that real wages are basically acyclical. Show that the model then implies that consumption and hours should move inversely in response to changes in interest rates. Is this result observed empirically?

4. Consider a two-period model of household budgetting:

$$\max \ \ln(c_t - n_t^{\mu}) + \beta \ln(c_{t+1} - n_{t+1}^{\mu})$$

subject to

$$w_t n_t + \frac{w_{t+1} n_{t+1}}{1 + r} = c_t + \frac{c_{t+1}}{1 + r}.$$

(a) Find the three Euler equations.

(b) This period utility function is non-separable between consumption and leisure. We have seen that one problem with models with separable utility is that if wages are acyclical then they predict c and n should move inversely, which they do not do empirically. Can this reformulation avoid that difficulty?

5. Suppose there are two countries, indexed by i. In country i the utility of a representative agent is given by:

$$U_i = \sum_{t=0}^{T} \mathrm{E}_0 \beta^t u(c_{it}),$$

with

$$u(c_{it}) = \ln(c_{it}).$$

In each country there is an endowment of a non-storable commodity. The endowment in country i is given by

$$y_{it} = z_{it}, \qquad z_{it} \sim iid, \quad \text{corr}(z_{1t}, z_{2t}) = 0.$$

(a) Solve for the consumption allocation in a competitive equilibrium, by solving a social welfare maximization problem with equal weights on the two countries.

(b) What is the correlation between c_{1t} and c_{2t}?

(c) Now let us modify the environment in two ways. First, suppose that

$$y_{it} = n_{it}^{\alpha} z_{it},$$

where n_{it} is labour supply in country i. The exogenous component of output, z, is still *iid* and has zero correlation across countries. Second,

$$u = \ln(c_{it} - n_{it}).$$

Again solving a Pareto planning problem, find the optimal n_{it}.

(d) Does this modification affect the prediction for the cross-country, consumption correlation?

Answer

(a) The allocation is

$$c_1 = c_2 = \frac{z_1 + z_2}{2}.$$

(b) The correlation is one.

(c) From the first-order conditions we see that:

$$n_i = (\alpha z_i)^{\frac{1}{1-\alpha}}.$$

(d) The planner equates marginal utilities of consumption, so

$$\frac{1}{c_1 - n_1} = \frac{1}{c_2 - n_2}.$$

Thus when z_1 rises, and n_1 rises (shown in part (c)) so does c_1. This breaks the perfect correlation between the two consumptions, which we do not observe empirically. For more details, see Devereux *et al*, *Journal of International Money and Finance* (1992).

6. Assume that a period utility function is given by:

$$u(c) = \frac{c^{1-\alpha}}{1-\alpha},$$

with $\alpha > 0$. Then consider the problem:

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

with $0 < \beta < 1$, subject to

$$A_{t+1} = R_{t+1} \cdot (A_t - c_t),$$

and given A_0. Assume that R_t is independently and identically distributed, with $ER_t^{1-\alpha} < 1/\beta$.

(a) Write Bellman's equation for this problem.

(b) Write the Euler equation for this problem.

(c) Show, using the guess-and-verify method, that the optimal policy function takes the form $c_t = \lambda A_t$ and give an explicit formula for λ.

Answer

(a)

$$V(A_t) = max\left[\frac{c_t^{1-\alpha}}{1-\alpha} + \beta E_t V(R_{t+1} \cdot (A_t - c_t))\right]$$

Notice that R is not a state variable because it is iid.

(b)

$$c_t^{-\alpha} = E_t \beta R_{t+1} c_{t+1}^{-\alpha}$$

(c) We show this by showing that our guess of this form yields a solution:

$$1 = E\beta R^{1-\alpha}(1-\lambda)^{-\alpha}$$

Solve this for λ. Incidentally, you then can see the effect of interest-rate volatility on consumption.

7. Imagine an economy with two-period overlapping generations, and no population growth. Young agents supply labour inelastically, receive a wage w_t, consume c_{1t}, and save s_t. When old, they consume c_{2t+1} financed from their savings. The economy is competitive. Period utility is logarithmic, and there is a discount factor β in the two-period utility function.

The production function is:

$$y_t = A_t k_t^{\alpha} n_t^{1-\alpha},$$

where

$$A_t = exp(\lambda t + \epsilon_t)$$

and ϵ_t is $iin(0, \sigma^2)$. The market clearing condition is

$$s_t = k_{t+1}.$$

(a) Find the law of motion for $\ln(k_t)$.

(b) What is the correlation between $\ln(w_t)$ and $\ln(y_t)$?

(c) What are some of the empirical failures of business cycle models driven by productivity shocks?

Answer

(a)

$$\ln(k_t) = \ln[(1-\alpha)\beta/(1+\beta)] + \alpha \ln(k_{t-1}) + \lambda t + \epsilon_t$$

(b) The correlation is unity.

(d) (See Plosser and Mankiw on the reading list.)

8. Imagine a model economy with a representative agent, with utility function:

$$U = \sum_{t=0}^{\infty} \beta^t log(c_t),$$

where β is a discount factor and c_t is consumption. There is no uncertainty. Time is discrete. The production function uses capital, k_t, and labour, n_t, to produce output, y_t:

$$y_t = k_t^{\alpha}[(1+\lambda)^t n_t]^{1-\alpha},$$

where $\lambda > 1$ is an exogenous rate of labour-augmenting technical progress. There is no population growth, and $n_t = 1$. Capital depreciates fully each period, so that

$$k_{t+1} = y_t - c_t.$$

The equilibrium is competitive, and k_0 is given.

(a) State Bellman's equation for this problem.

(b) Hence state the Euler equation linking consumption in adjacent time periods.

(c) Use the guess-and-verify method to show that the optimal investment rule is of the form

$$k_{t+1} = \omega y_t$$

and find the value of ω.

(d) Solve for the real wage in terms of the state variable. What is the growth rate of the real wage?

(e) Suppose that λ falls from 1.05 to 1.03, say in 1973. What would be the effect on measured growth in total factor productivity?

Answer (a)

$$V(k_t) = max[\ln(c_t) + \beta V(k_{t+1}]$$

subject to

$$k_{t+1} = k_t^\alpha [(1+\lambda)^t n_t]^{1-\alpha} - c_t$$

and given k_0.

(b) From the envelope condition:

$$\frac{1}{c_t} = \frac{\beta(1 + r_{t+1})}{c_{t+1}}$$

where $1 + r_{t+1} = \alpha y_{t+1}/k_{t+1}$.

(c) $\omega = \alpha\beta$

(d)

$$w = (1-\alpha)k_t^\alpha [(1+\lambda)^t]^{1-\alpha}$$

You can show that k_t grows at rate λ, and so does w.

(e) From growth accounting (using the continuous-time approximation), the Solow residual is

$$(1-\alpha)\lambda.$$

Say α is between 0.3 and 0.4, then the decline in TFP growth is between 0.014 and 0.012.

9. The first equilibrium business-cycle models featured only one random exogenous variable, a shock to total factor productivity. This question examines how adding a random fiscal policy might correct some of the unrealistic predictions of the model.

Imagine a household planning problem of maximizing

$$EU = E_0 \sum_{t=0}^{\infty} \beta^t u_t,$$

where in each period t

$$u_t = \ln(c_t - \mu g_t) + \ln(1 - n_t).$$

Here c is consumption, n is labour supply, and g is government spending. The parameter μ describes the direct impact of public spending on utility: if $\mu = 0$ there is no impact while if $\mu = 1$ government spending is a perfect substitute for private spending. Labour income in period t is $w_t n_t$, and there is a lump-sum tax. Government spending is taken as exogenous by the household.

(a) What are some of the unrealistic predictions of real business cycle models based on technology shocks only?

(b) In the model described above, find the Euler equation linking c_t, g_t, n_t, and w_t. Can this model be consistent with procyclical consumption and employment and an acyclical real wage?

(c) Find the Euler equation linking the marginal utilities of goods in adjacent time periods. Do you think that this model can resolve any of the empirical shortcomings of the consumption-based CAPM?

Answer (a) Shocks may not actually measure technology; difficult to see how technology gets worse or is predictable; the model predicts procyclical real wages and may have difficulty with hours or employment as well; the model does not have much propagation so that income dynamics are much like shock dynamics. See Romer, pp. 186-189.

(b)

$$\frac{c_t - \mu g_t}{1 - n_t} = w_t$$

Yes, it can provided g is procyclical enough. That seems a bit unlikely. Government spending would have to be at least as procyclical as consumption, so that $c_t - \mu g_t$ would be countercyclical.

(c) Perhaps. We now have an additional source of volatility in the intertemporal marginal rate of subsitution, which might lead to more volatile interest rates and a higher equity premium. But, recall from part (b) we needed g to be correlated with c to solve one problem. In this problem a positive correlation will reduce the volatility of $c - \mu g$, so it is less likely that marginal utility will be much more volatile.

10. This question examines growth theory with an endogenous saving rate. Let time be discrete. Suppose that the production function is:

$$Y_t = K_t^{\alpha}[N_t(1+\lambda)^t)]^{1-\alpha},$$

where λ is a constant rate of labour-augmenting technical progress.

Suppose that we imagine a representative household which maximizes

$$EU = E_t \sum_{j=0}^{\infty} \beta^j \ln(C_{t+j}),$$

where β is a discount factor. The population N is constant and labour is supplied inelastically. Markets are competitive. The depreciation rate is 100%.

(a) Write this as a dynamic programming problem.

(b) Use the guess-and-verify method to find the optimal plan for consumption.

(c) On a steady-state growth path, find the wage rate and the real interest rate as functions of the parameters.

(d) Suppose that we wish to use growth theory to predict the duration of a recovery from a war, in which half the capital stock is destroyed. An economist argues as follows: "The Solow-Swan model overestimates the recovery time because it does not allow for the fact that households increase their saving rate in response to high interest rates." Is this economist's point important theoretically or empirically?

Answer

(a)

$$V(K_t) = max[\ln(C_t) + E_t\beta V(K_{t+1}],$$

subject to

$$K_{t+1} = K_t^{\alpha}[1+\lambda)^t]^{1-\alpha} - C_t.$$

(b) It is easy to see from the Euler equation that the usual Brock-Mirman result holds:

$$C_t = (1-\beta\alpha)Y_t.$$

(c) On a steady-state growth path the gross interest rate is

$$1 + r = \frac{1+\lambda}{\beta} = (1+\lambda)(1+\theta).$$

The wage grows at gross rate λ.

(d) In the example is this question there is log utility so the savings rate does not depend on the interest rate. So that is one theoretical example in which the economist is incorrect. Moreover, the empirical work on saving and interest rates (admittedly not collected during recovery from war) suggests the effect is quite small.

11. Suppose that a competitive economy has the following technology, in discrete time:

$$y_t = z_t k_t^{\alpha}(h_t n_t)^{1-\alpha},$$

where z is a non-negative random variable. Suppose that in equilibrium $n_t = 1$ and

$$h_t = \gamma k_t^{\gamma},$$

due to an external effect of capital accumulation on human capital.

(a) For what parameter values does this model produce endogenous growth? Assume this condition holds from now on.

(b) There is 100 percent depreciation and log utility with discount factor β. If $\alpha = 0.4$ and $\beta = 0.95$ find the law of motion for the capital stock k_{t+1} as a function of z_t and k_t.

(c) Suppose that there is some persistence in the technology shock:

$$\ln(z_t) = (1-\rho) + \rho \ln(z_{t-1}) + \epsilon_t,$$

where ϵ_t has mean zero. What is the average growth rate of output?

Answer

(a) $\gamma = 1$ ($\alpha = 1$ and $\gamma = 0$ does not seem realistic)

(b)

$$k_{t+1} = \beta\alpha y_t = 0.38 z_t k_t$$

(c) The average value of the log shock is 1. So

$$E\Delta \ln(k_{t+1}) = \ln(\beta\alpha) + 1 = 0.03.$$

12. A country produces a nonstorable consumption good with this technology:

$$y_t = z_t n_t^{\frac{1}{2}},$$

where y is output, n is hours worked and z is a positive, iid, random productivity measure. All households have the following preferences:

$$EU = \sum_{t=1}^{\infty} E_0 \beta^t \ln(c_t - \gamma n_t),$$

where c is consumption and β is a discount factor. Markets are competitive.

(a) An econometrician does not know the production function. Describe how the econometrician could use observations on consumption, hours, real interest rates, and real wages to estimate β and γ.

(b) Suppose that $\gamma = 0.5$. Solve for labour supply and consumption as functions of productivity.

(c) Do you think the responses of consumption and labour supply to a productivity shock would be any differerent if this were an open economy?

Answer

(a) Use GMM to estimate the three Euler equations, using various instruments.

(b) Substitute the constraint in the objective, and choose n to maximize:

$$\ln(zn^{\frac{1}{2}} - 0.5n)$$

which gives

$$n = c = z^2.$$

The second-order condition also holds.

13. Suppose you are asked this question: Are business cycles due to productivity shocks? This series of exercises looks at how an answer to this general question might be approached.

(a) First, how can technology or productivity shocks be measured? Do you think Solow residuals can be regarded as exogenous changes in aggregate productivity? To answer these questions, try sampling some of the review articles in the suggestions for further reading.

(b) Suppose that you decide to measure technology shocks using actual Solow residuals:

$$A_t = \frac{Y_t}{K_t^{\alpha} N_t^{1-\alpha}}.$$

What value of α should you use in your measurements?

(c) Suppose that you selected $\alpha = 0.3$. You then decide to use this form:

$$A_t = A_{t-1}^{\rho} exp(\epsilon_t),$$

where ϵ_t is iid with mean zero and variance σ^2. How could you find values for ρ and σ^2?

(d) Suppose that you found $\rho = 0.9$ and $\sigma^2 = 1.5$. Now that you have described shocks, you need to describe how the economy propagates them. Suppose you adopt the Brock-Mirman model with log utility and inelastic labour supply. Describe how to use GMM to test your assumption of log utility, using the Euler equation for consumption growth. Describe how to estimate β, the discount factor, using data on aggregate consumption and asset returns.

(e) Suppose that you find $\beta = 0.9$. Graph the impulse response function for $\ln(A)$ and for the model's predictions for $\ln(Y)$, in response to a 'surprise' of $\epsilon_t = 1$. Does this real business cycle model include a significant propagation mechanism?

(f) Find the variance of log output predicted by the model.

(g) Could you use this model to compare predicted and actual paths for output, just as you would in testing an econometric model?

(h) A student in Econ 852 writes: "Most economists construct confidence intervals around predicted moments or sample paths. The failure of real business cycle modelers to do this makes their findings impossible to interpret." Discuss how you could allow for sampling variability in the predictions of this business-cycle model.

14. Consider the Brock-Mirman environment where a representative agent seeks to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t [\ln(c_t) + \ln(1 - n_t)]$$

subject to

$$c_t + k_{t+1} = y_t$$

and k_0 given.

(a) Consider the basic case $y_t = k_t^\alpha n_t^{1-\alpha}$. Use dynamic programming to solve for the decision rules for consumption, hours worked and investment.

(b) Consider exogenous labor-augmenting technical change $y_t = k_t^\alpha (n_t \lambda^t)^{1-\alpha}$. Use dynamic programming to solve for the decision rules for consumption, hours worked and investment.

(c) Consider endogenous growth (learning by doing) $y_t = k_t^\alpha (n_t h_t)^{1-\alpha}$ where $h_t = k_t$ in equilibrium. Use dynamic programming to solve for the decision rules for consumption, hours worked and investment.

Note that in each case you can guess $k_{t+1} = \gamma y_t$ and therefore $c_t = (1-\gamma) y_t$.

15. Consider a model of cycles driven by shocks to government spending, which are financed by lump-sum taxation. A representative agent seeks to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t [\ln(c_t) + \ln(1-n_t)]$$

subject to

$$c_t + k_{t+1} + g_t = y_t$$

and k_0 given. Suppose $y_t = k_t^\alpha n_t^{1-\alpha}$. Also suppose that

$$g_t = \theta_t y_t$$

and that θ_t is a random variable. However, this is an externality: firms ignore the effect of their investment decisions on government spending.

(a) Show that

$$c_t = (1 - \beta\alpha - \theta_t) y_t$$

and that

$$n_t = \frac{1-\alpha}{2 - \alpha - \beta\alpha - \theta_t}.$$

(b) Why does a shock to θ_t increase labour supply? What is the model's prediction for the cyclicality of the real wage?

16. This question uses the discrete-time optimal growth model to study the productivity slowdown. Suppose that the production function is:

$$Y_t = K_t^\alpha [N_t (1+\lambda)^t]^{1-\alpha},$$

the depreciation rate is 100 percent, and utility is logarithmic and time-separable with discount factor β. The economy is competitive. The population grows at a constant rate η so that $N_{t+1} = (1+\eta)N_t$.

(a) Solve for a competitive equilibrium as a dynamic programming problem.

(b) Suppose that the index t counts years, $\alpha = 0.33$, $\beta = 0.9$, $\eta = 0.01$ and $\lambda = 0.04$ up to 1973, while from 1974 on $\lambda = 0.02$. Describe the behaviour of real wages in response to the decline in λ (including the transition if possible).

(c) Discuss alternatives to a decline in λ as possible explanations for the measured productivity slowdown.

Answer (a) Combining the Euler equation and the constraint using the guess-and-verify method gives:

$$K_{t+1} = \beta\alpha Y_t.$$

Then C, w, r, and Y follow.

(b) The fundamentel equation in efficiency units of capital is:

$$k_{t+1} = \frac{\beta\alpha}{(1+\eta)(1+\lambda)} k_t^{.}33$$

so

$$k^* = \Big(\frac{\beta\alpha}{(1+\eta)(1+\lambda)}\Big)^{\frac{3}{2}}$$

Next,

$$w_t = (1-\alpha)k_t^{\alpha}(1+\lambda^t).$$

At the initial values $k^* = .150349$ and at the final values $k^* = .15479$. Of course the growth rate of the real wage drops from 4 percent to 2 percent. How does the transition look?

(c) [Discussion of measurement questions and of oil price shocks. Also, what is the evidence on real wage growth?]

17. Imagine a country of small producers, each operating this technology:

$$y_t = z_t n_t^{\frac{1}{2}},$$

where n is their labour input and z is a technology shock. The output they produce is not storable. Their utility function is:

$$EU = E_0 \sum_{t=0}^{\infty} \ln(c_t - n_t).$$

The government collects a share τ_t of output as tax revenue each period, leaving a share $1-\tau_t$ to the producers.

(a) Find optimal consumption and labour supply.

(b) Does this economy exhibit Ricardian equivalence?

(c) Does this model produce realistic cycles in consumption, output, and employment in response to a technology shock?

Answer (a) This becomes a static problem of choosing n to maximize

$$\ln(z(1-\tau)n^{\frac{1}{2}} - n),$$

which gives

$$n_t = 0.25[z_t(1-\tau_t)]^2.$$

so that

$$c_t = 0.5[z_t(1-\tau_t)]^2,$$

(b) No. Changes in tax timing (with the present value of taxes unchanged) affect labour supply and consumption and the real interest rate. See if you can explain the effect of a tax deferral on the real interest rate.

(c) The model produces some realistic facts: consumption and employment are procyclical; consumption is less volatile than output. The persistence comes entirely from the shock though, as there is no propagation mechanism in the model.

18. In this question, we'll see whether a dynamic model with linear technology and quadratic utility can produce business cycles. Suppose that the utility function is

$$E_0 \sum_{t=0}^{\infty} \beta^t (c_t - \theta c_t^2)$$

and that the capital stock accumulates according to

$$k_{t+1} = Ak_t + u_t - c_t,$$

where $A = 1/\beta$. The economy is competitive. The variable u_t is a shock, which follows a first-order autoregression:

$$u_t = \rho u_{t-1} + \epsilon_t,$$

where $0 < \rho < 1$ and ϵ_t is white noise. Thus $E_t u_{t+1} = \rho u_t$.

(a) Why are we able to first solve for quantities without studying prices at the same time? What environment might produce the linear production function: $y_t = Ak_t + u_t$?

(b) Set this up as a dynamic programming problem and solve for the optimal decision rules. (Hint: Guess that $c_t = ak_t + bu_t$ and solve for the coefficients a and b.)

(c) How is consumption related to current and past technology shocks? To see whether the model propagates shocks, graph the impulse response function for consumption, for the special case with $\rho = 0$.

Answer. (a) The sceond welfare theorem. The endogenous growth model with learning-by-doing.

(b) The state is (k_t, u_t). With $A = 1/\beta$ the Euler equation is: $c_t = E_t c_{t+1}$. The decision rules are:

$$c_t = (A-1)k_t + \frac{r}{1+r-\rho}u_t,$$

and

$$k_{t+1} = k_t + \frac{1-\rho}{1+r-\rho}u_t.$$

(c)

$$c_t = bu_t + a(1-b)u_{t-1} + a(1-b)u_{t-2} + ...$$

so using the definitions and $\rho = 0$ shows that all the coefficients are $r/(1+r)$. Thus each shock has a small permament effect, as befits their permanent effect on capital and then the permanent income hypothesis.