Business Cycle Theory I (REAL)

I. Introduction

In this chapter we present the business cycle theory of Kydland and Prescott (1982), which has become known as Real Business Cycle theory. The "real" term was coined because (1) the model does not include money and because (2) shocks to the TFP component of the production function are the source of the business cycle that are emphasized in this research

The mathematics involved in solving out this model are far too advanced for this level. For this reason, we depart from our usual presentation of the solution methods. Instead, we present the basic elements of the business cycle model, provide intuition for the effect of productivity shocks on the economy, and then describe the calibration exercise with a summary of the findings of Kydland and Prescott (1982).

II. The Growth Model

The model that underlies real business cycle theory is just the growth model studied in Chapter X, with three notable differences. First, for the purpose of studying the business cycle, the empirical counterpart of the model economy is a quarter of a year. This different definition of the period length does not alter the model's structure in any way, but does affect some of the calibrated parameters. The second difference is that TFP in the production function is a random variable. In all the previous models to date, there were no stochastic elements. The models were thus deterministic. Deterministic models are what are called perfect-foresight models because they assume that the households and firms know the values of all the economic variables over their lives.

Models with random variables are called stochastic models. The inclusion of random elements means that agents are not all knowing, i.e., have perfect foresight of the future. The existence of stochastic elements means that agents must form an expectation of the value of the future variables, and make a decision based on those expectations.

Expected Utility

With the introduction of random variables, there is the question of what is the appropriate utility concept for households? The convention in the literature is to use the von Neumann and Morgenstern concept, which is just the sum of the utilities the household obtains for each possible state of nature weighted by the probability the household puts on that event happening.

To illustrate, consider the following example. There are two states of nature, we label good and bad. The probability assigned to the good state by the household is π and the probability assigned to the bad state is thus, 1-π. Let c(g) be the consumption that household chooses in the case of the good state and let c(b) be the household's consumption in the case he bad state is realized. Suppose the utlity function of the household is U(c). Then the household's expected utility is

(EU) $$\pi U(c(g)) + (1-\pi)U(c(b))$$

In the context of the growth model with the infinitely lived construct, the convention is to not write out the probability of each event. Instead, the notational convention is to write expected utility as $E_0 \sum_{t=0}^{\infty} \beta^t u(c_t, l_t)$. The subscript on the expectation operator refers to the period in which the expectation is formed. Here, it is the beginning period, time t=0.

III. Rational Expectations

Another issue we must deal with in economies with random variables is how expectations are formed. Specifically, what is the basis for the subjective probabilities the household places on each state of nature being realized? The convention in stochastic macro models is to impose the assumption that expectations are formed rationally.

Rational Expectations, as they are known in the literature, are attributed to John Muth who in 1960 introduced them the study of finance. Robert E. Lucas and Edward C. Prescott in 1970 introduced Rational Expectations to macroeconomics in their paper, Investment under Uncertainty.

The idea of rational expectations is rather simple. People use all available information they have to date, and then act in a manner that is consistent with those expectations. Formally, Rational Expectations is the conditional expectations operator. To illustrate, suppose we had a random variable, z, in each period. Then in period t, there would be some information available to the household, call it I_t. Then the rational expectation of z_{t+1} would be the expected value of z in period t+1 given the available information I_t.

To further illustrate the concept, suppose that a random variable, z, evolves according to the following law of motion

(AR1) $$z_{t+1} = .95 z_t + \varepsilon_{t+1}.$$

Such a process is called an auto regressive process of order 1. The order refers to the number of lagged values in the equation. If there were an additional term z_{t-1} on the right hand side of the equation, then the AR would be of order 2.

The above equation is stochastic in that ε is a random variable. As such, there is a probability distribution that governs its value, with a mean and variance. The convention is to assume that the random shock is drawn from the same distribution each period and that its value in period t+1 is independent of any past realized value. At time t, the value of A_t is given and hence part of the household's information set. The agent must form a rational expectation of z_{t+1}. The conditional expectation is found by applying the expectations operator to both sides of (AR1). Since the expectations operator is linear, we have

(CE) $E_t\{z_{t+1}\} = .95z_t + E\{\varepsilon_{t+1}\}$.

Assuming that the mean of ε is zero, the conditional expectation of z_{t+1} is just .95 z_t. This is the rational expectation of z.

The rational expectation (i.e., the conditional expectation) is very different from the unconditional expectation. To calculate the unconditional expectation we iterate backwards on the term A_t on the left hand side of (CE). Given that (AR1) must apply in all periods, we can rewrite the equation as

$$z_{t+1} = .95[.95z_{t-1} + \varepsilon_t] + \varepsilon_{t+1} = .95^2 z_{t-1} + .95\varepsilon_t + \varepsilon_{t+1}.$$

If we continue this backward iteration forever, we arrive at

$$z_{t+1} = \varepsilon_{t+1} + .95\varepsilon_t + .95^2 \varepsilon_{t-1} + .95^3 \varepsilon_{t-2} \ldots\ldots$$

Now if we apply the expectations operator to both sides, and maintain our assumption that the mean of the error term is zero each period, then we find the unconditional expectation of z_{t+1}=0.

That is to say that if you did not have any information at time t, (i.e., the value of z_t), then your best guess for z_{t+1} is zero.

IV. Understanding the Effects of Productivity Shocks

We have used throughout this book the unit elasticity utility function, $u(c_t, l_t) = \ln c_t + \alpha \ln l_t$. The justification for this functional form is the observation that since 1950 the real wage has shown constant growth whereas average hours worked has shown no secular trend. With the unit elasticity substitution, the substitution effect offsets the income effect so that there is no price effect.

This is actually a long-run result, and implicitly involves a permanent increase in the price of leisure. If there is a temporary change in the price of leisure, then leisure and work hours will change in the short-run. To see this, consider the two key household optimizing conditions

(H1) $$\frac{1-h_t}{\alpha c_t} = \frac{1}{w_t}$$

(H2) $$\beta E_t\left(\frac{c_t}{c_{t+1}}\right) = \frac{1}{1+i_t}$$

Using (H1) for both period t and period t+1 and substituting into (H2) we arrive at

$$\beta E_t\left((1+i_t)\frac{w_t}{w_{t+1}}\frac{l_t}{l_{t+1}}\right) = 1.$$

Now if there is a positive productivity shock today, then w_t will increase. Being temporary w_{t+1}

either does not change or increases less than w_t. If the interest rate is not much changed (or increases), then it follows that the ratio of l_t/l_{t+1} must fall, which is the same as saying that work hours increase today. Notice that the increase in the real wage were permanent so that $w_t=w_{t+1}$, there would not be any change in leisure.

V. Calibration- Estimating Productivity shocks

Although the business cycle is the center of the analysis, the calibration of the model is essentially the same as before. Namely, the preferences and technology parameters are calibrated so that the deterministic model's steady state- match the long-run observations of the United States economy over the 20th century. In doing this, the quantitative exercise of the real business cycle school is to examine the short-run implications of a model restricted to match the long-run behavior of the US economy.

There are a couple of noted changes in the calibrated technology and preference parameters on account that we are analyzing the economy at the much higher frequency of 3 month intervals. Since output is a flow variable, its value over a quarter of a year is one fourth its value over the entire year. This means that the observation for k/y must be multiplied by four for the purpose of calibrating the value of δ, around .025. With this much value for k/y and much lower calibrated value of δ, the imputed real rate of interest and rental rate of capital are lowered. This implies a much higher value for the subjective time discount factor, β around .99. These are the only changes in Step 4 of the calibration.

To complete the calibration, the process for the technology shocks needs to be estimated. Kydland and Prescott's strategy was to undertake a Solow Growth Accounting exercise so as to impute a time series for Solow Residual's $\{A_t\}$. Then having this "data", they estimated an (AR1) process of the form.

To illustrate, let us rewrite the production function as

$$y_t = e^{z_t} k_t^{\theta} [(1+\gamma)^t h_t]^{1-\theta}$$

Where z is given by (AR1)

Then the Solow residual is

(SR) $$e^{z_t} (1+\gamma)^{t(1-\theta)} = \frac{y_t}{k_t^{1/3} h_t^{2/3}} .$$

Using quarterly data from the US economy on GDP, the capital stock and total work hours one can calculate the time series for TFP, i.e., Solow Residuals. Next, if we divide both sides of (SR) by $(1+\gamma)^{t(1-\theta)}$ and take the log of both sides, we can get the time series for $\{ z_t\}$. We then use the time series for $\{z_t\}$ and estimate the (AR1) equation. This gives us an estimate for the parameter, ρ, of .95.

The calibration is not complete because what is critically important for the theory is the magnitude of the shocks. This is determined by the variance of the shock, ε. To assign the value to the variance, we compute the error term in each period, namely, the difference between the actual and predicted values for z. This gives us a time series for the error terms, $\{\varepsilon_t\}$. Using these numbers, we then calculate the mean of the errors and the variance. With these parameter values in hand, Step 4 of the calibration procedure is complete, and the theory can be tested.

To do this, Kydland and Prescott fed in a sequence of TFPs that are based on the estimated AR1 process above with ρ=.95 and with computer generated random numbers that are based on the variance and means of the error terms. With these realized productivity shocks they then solve out the actual path for the model economy. In this way they simulated the US economy.

Now, it is quite possible that he computer random number generate produces a series of good shocks in each period, or a series of bad shocks every period. To eliminate the effect of a particular draw on the conclusions, Kydland and Prescott simulated the economy over 20 times. In this way, they generate a sampling distribution of statistics. This requires that for each simulation, the HP filter is applied in order to find the trend and deviation components. Next, the volatility for each simulation is determined.[1] The average of these volatilities is then computed. This is the relevant statistic to be compared with the volatility computed for the actual US economy. In comparing comovements as predicted by the model and as displayed by the actual US economy, again we compute for each simulation the correlation coefficients, then we take the average across the simulations. This average comovement is compared to the correlation coefficient computed using the actual quarterly data for the US economy.

VI. Assessing the Real Business Cycle Theory.

How much of the volatility in US GDP can be accounted by productivity shocks? The following table reprinted from a 1992 *Federal Reserve Bank of Minneapolis Quarterly Review* article written by Gary Hansen and Randy Wright compares the model's predictions with the data both in terms of volatility and comovements. The period of comparison is 1947:Q1-1992:Q3.

[1] Recall, to calculate the measure of volatility, we first express the deviation as a percentage of the trend, namely, define xt = dt/τt. Next, we take the standard deviation of the (percentage) deviation from trend. This is the measure of volatility.

Relative to the tables in Chapter X on the Business Cycle Facts, the Hansen and Wright table presents the business cycle properties of the model and data in a slightly different way and considers fewer dimensions of the data. Importantly, except for the volatility of output, the volatility of all other variables are expressed relative to the volatility of output. Additionally, the table includes a measure of productivity defined as output divided by hours. This variable is indexed by the letter ω.

	% s.d. of output σ_y	Consumption σ_c / σ_y	Investment σ_c / σ_y	hours σ_h / σ_y	Productivity σ_ω / σ_y	σ_h / σ_ω	Corr(h,ω)
US Time series	1.92	.45	2.78	.87	.51	1.76	-.07
Model	1.30	.31	3.15	.49	.53	.94	.93

Note: the last four columns is an average of the statistic implied by the Household Survey data and the Establishmnet Survey data.

There are a number of key points to be taken from the table. First, returning to the question of the calibration, productivity shocks can account for roughly 2/3rds of the volatility of output. This is actually quite amazing, given that there are clearly other shocks in the real world that cause output to deviate from its long-run trend. In terms of consumption and investment, the theory is judged to be a relative success. The model predicts a little too much by way of consumption smoothing, and hence for that reason, a little too much by way of investment volatility.

Where the model is deemed to be less successful relates to the predictions with respect to hours and productivity. First, hours in the model are not volatile enough relative to the data. This is the main reason why the model does not account for 100 percent of the output volatility.

Productivity shocks do not bring about enough of an increase in hours in the model relative to the data, so that the model does not generate a larger deviation in output.

For the purpose of understanding this issue as well as the greatest failure of the model, the correlation between hours and productivity, it is useful to show the implications of the RBC theory using the labor market diagram. Labor demand is just the Marginal product of labor from the firm's profit maximization problem. A positive TFP shock shifts the labor demand curve to the right, implying both a rise in the real wage and an increase in hours. Notice, importantly, the the size of the hour increase depends on the steepness of the labor supply curve, with a larger response being associated with a flatter curve. This is shown in Figure 1. Here we have drawn an increase in the Labor Demand curve associated with a positive TFP shock and compare the effect on the equilibrium hours for two labor supply curves that differ in their steepness.

Figure 1: Change in Hours Depends on Slope of Labor Supply

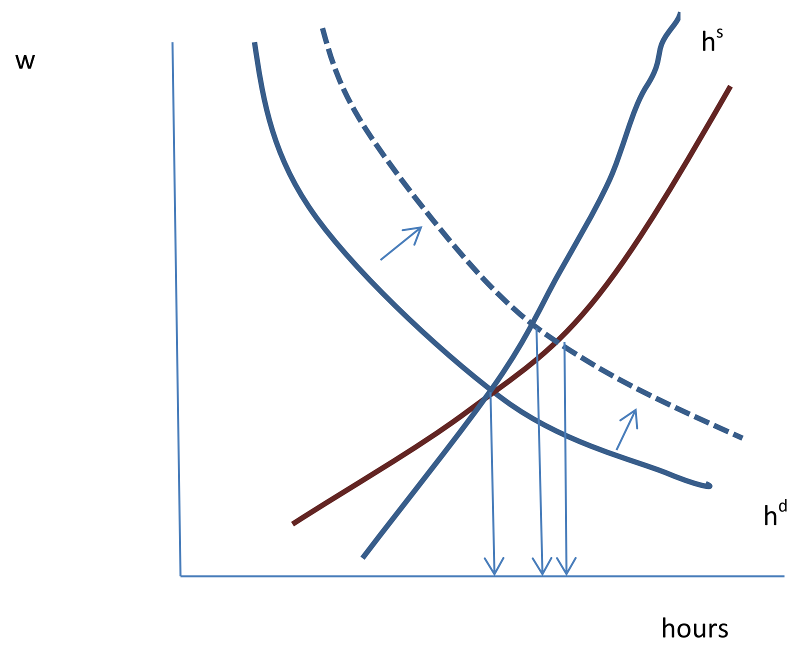

Recall that the labor supply is upward sloping because the higher wage is viewed as a temporary phenomenon. The slope then reflects the degree which the household is willing to intertemporally substitute between today and tomorrow. Intuitively, this determined by the curvature, or concavity of the utility function, which in this case has been specified as the log. Log utility although concave, simply implies too small of an intertemporal substitution effect, and for this reason, hours respond by far less in the model compared to the data.

Although predicting less volatility in hours than observed in actuality, the model is still successful along this dimension. Much harder to reconcile is the models' prediction for the correlation of hours and productivity. In the data, it is essentially zero; in the data it is close to 1.

Figure 2: Positive Productivity Shock with Increase in Labor Supply

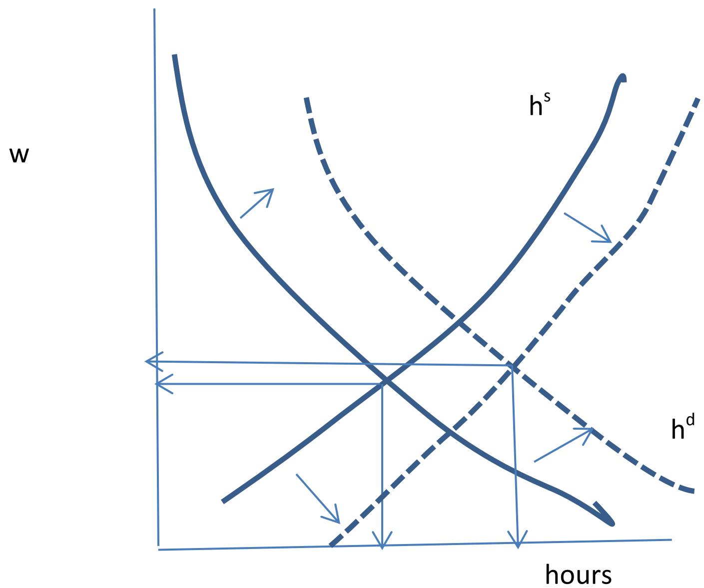

The reason why the model predicts a correlation close to one is evident from the labor diagram above. With the increase in labor demand, hours and wages move together. Since productivity is just the wage divided by the labor share, the hours and productivity move more or less the same. What would have to happen to get (1) hours to increase with productivity shocks and (2) no change in the wage? One way to achieve (1) and (2) is to have the labor supply increase at the same time as labor demand increases. Note that this shift would not only solve the correlation issue but generate a bigger response in hours.

IV. Extensions

A number of extensions to the Real Business Cycles were put forward in an effort to show that the theory could be reconciled with the volatility of work hours over the cycle and the correlation between hours and labor productivity.

Non-Separable Leisure

This is a fix to the first deficiency of the standard RBC model, namely, its inability to generate enough volatility in work hours. It was explored in a paper by Hotz, Kydland and Sedlacek (1988). Recall from Figure 1 that greater volatility is implied if the today's supply of labor responds more to a higher wage, i.e., the labor supply curve is flatter, or the elasticity of substitution is higher.

With non-separable leisure, today's utility does not depend just on today's leisure but a weighted sum of the household's past leisure choices. Specifically, current period utility is

(N-SUtility) $$u(c_t, L_t) = \ln c_t + \alpha \ln L_t .$$

The key difference is that L_t is the weighted sum of this period's and all past period's leisure, namely,

$$L_t = \eta_0 l_t + \eta_1 l_{t-1} + \eta_2 l_{t-2} + \eta_3 l_{t-3} + ...$$

One interpretation of this type of preference is that people are willing to work a lot of hours in this quarter if they had enjoyed a nice vacation 3 or 6 months earlier. It is the idea of your university giving you a Spring break- with a week off, you will be better able to focus and study longer hours as finals approach.

Hotz, Kydland and Sedlacek (1988) impose the following restriction on the weights so that

$$\eta_{i+1} = \rho \eta_i$$

The implication of this weighting system is that the weights decline geometrically by the factor ρ each period.

With respect to the calibrations, non-seperable leisure introduces two new parameters whose values must be restricted in step 4 of the calibration procedure, namely, η_0 and ρ. The authors use employment histories for white male head of households taken from the Panel Study of Income Dynamics Using to restrict the value of these two parameters. The estimated parameters are found to be $\eta_0 = .35$ and ρ=.90.

Home Production

Benahabib, Rogerson and Wright (1991) considered the role of home production in the context of the business cycle. Home production is the goods and services that are done by households outside the market. These goods and services do not receive any compensation and are generally excluded from the NIPA. Examples are dinners you prepare at home, laundry, cleaning, yardwork, driving yourself and family members to their jobs and activities.

Intuitively, the addition of home production allows households to substitute between market and non-market activities. Specifically, when there is a positive shock to TFP in the production of market goods and services, households will find it optimal to work more in the market and less at home. Essentially, when there is a positive shock to market production and a higher wage is offered, you start eating out more as well as having your clothes laundered at the dry cleaner.

Government Spending

Christiano and Eichenbaum (1992) proposed that government spending shocks could explain why labor is more volatile than in the standard model and why the correlation between hours and labor productivity is close to zero. The idea is to formulate a way by which such shocks would lead to a shift in the labor supply curve. In the Christiano and Eichenbaum model, government expenditures are paid for by lump-sum taxes. Thus, a higher shock would imply larger lump-sum taxes in the current period. This would reduce the household's wealth, and as we saw in our model of labor supply decision, this makes the household choose less leisure, i.e., supply more work hours.

The law of motion for government expenditures was specified as

(G Shocks) $\ln g_{t+1} = (1-\lambda)\ln(\overline{g}) + \lambda \ln g_t + \mu_t$

In (G Shocks), $\bar{g}$ denotes the steady-state level of government spending. The variable μ_t represents the shock to government spending. It is assumed to be independently and identically distributed with mean zero and with variance σ_μ. Importantly, the shock to government expenditures is assumed to be independent of the technology shock.

In this model, the government is assumed to run a balanced period each period. Expenditures are paid for by lump-sum taxes.

The inclusion of government expenditures in the model has some important implications for step 3 of the calibration procedure. Recall, in the Solow model and Neoclassical Growth Model without government, we moved government consumption in the NIPA into the model consumption category, and we moved government investment in to the investment category. Now, we no longer would make those adjustments.

This different reorganization would change some of the parameter values from the standard RBC model, but these are not particularly important in the context of this discussion. Important to determining the model's implications for business cycle is the assignment of the parameters in the (G Shock) Equation. This is done by using Ordinary Least Squares to estimate the law of motion for (G Shocks). Importantly, the value of $\bar{g}$ is calibrated to the long-run ratio of government expenditures to GDP, which is roughly 20%. Christiano and Eichenbaum (1992) report estimates for λ=.96 and σ_μ= .021.

The government spending shocks and TFP shocks are then fed into the model, and the equilibrium path is computed. This is done 200 times, and then the average volatility and correlations are compared to the data.

	% s.d. of output σ_y	Consumption σ_c / σ_y	Investment σ_c / σ_y	hours σ_h / σ_y	Productivity σ_ω / σ_y	σ_h / σ_ω	Corr(h,ω)
US Time Series	1.92	.45	2.78	.87	.51	1.76	-.07
Model	1.30	.31	3.15	.49	.53	.94	.93
Non-Separable Leisure	1.51	.29	3.23	.65	.40	1.63	.80
Government Spending	1.24	.54	3.08	.55	.61	.90	.49
Home Production	1.71	.51	2.73	.75	.39	1.92	.49

VI. Conclusion

In this chapter we have studied the Real Business Cycle Theory paradigm initiated by Finn Kydland and Ed Prescott and extended by many other economists. The findng of this research agenda is that productivity shocks can account for a large portion of the business cycle, although not all.

Despite this success, there are many economists who do not appreciate the contributions of this literature. As we shall discuss in the next chapter, these economists do not like a number of aspects pertaining to the Real Business Cycle Theory such as the lack of any frictions. The RBC theory assumes that all markets clear. Relatedly, opponents of this theory criticize it for the implication that recessions are just bought of laziness, that anyone not working out there is voluntarily unemployed. Finally, there is no room for policy, particularly monetary policy. In part this reflects that only real prices matter and that the equilibrium is Pareto Optimal. In the next chapter, we will turn to the prototypical model that supposedly avoids these failures- the so called New Keynesian Dynamic Stochastic General Equilibrium model.

Assignment:

1. Flip of coin- AR(1)
2. (i) Calculate Solow Residuals for US economy (Quarterly Data?), Capital, GNP, and Total work hours. (2) HP filter and (3) simple linear AR1 regression