

Graduate Macro Theory II: Stylized Business Cycle Facts and the Quantitative Performance of the RBC Model

Eric Sims
University of Notre Dame

Spring 2016

1 Introduction

This note describes some basic facts about US business cycles and then examines how well the basic RBC model can fit those facts.

2 Stylized US Business Cycle Facts

The convention within the literature is to look at HP detrended data to examine the business cycle. We focus on deviations from the smooth HP trend. There are, of course, alternative ways in which one can look at the cyclical component of the data (e.g. first differences, the Band Pass filter, linear detrending, etc.).

“Business cycle moments” focus primarily on a variety of second moments. In particular, the standard deviation of an HP filtered series is referred to as its volatility. We are also interested in looking at a series’ cyclical component, which is defined as its contemporaneous correlation with GDP. We call the first order autocorrelation of a series a measure of its persistence, and we also look at how strongly a series is correlated with output led or lagged a number of periods, so as to say something about which series are “lagging indicators” and which series are “leading indicators”.

The series we are most interested in looking at are the same endogenous variables that come out of a simple real business cycle model – output, consumption, investment, total hours worked, the real wage, and the real interest rate. In addition, we will look at average labor productivity (the ratio of output to total hours worked), the price level, and total factor productivity (TFP). The price level is not in the model as we have thus far specified it, but can easily be added. TFP is the empirical counterpart of the driving force A_t in the model. We measure it as output minus share weighted inputs:

$$\ln \hat{A}_t = \ln Y_t - \alpha \ln K_t - (1 - \alpha) \ln N_t \tag{1}$$

Constructing this series requires an empirical measure of the capital stock. In practice, this is hard to measure and most existing capital stock series are only available at an annual frequency. Typically the way in which people measure the capital stock is by using the “perpetual inventory” method. This method essentially takes data on investment, an initial capital stock, and an estimate of the rate of depreciation δ and constructs a series using the accumulation equation for capital: $K_{t+1} = I_t + (1 - \delta)K_t$.

All series (with the exception of TFP, the price level, and the interest rate) are expressed in per capita terms after dividing by the civilian non-institutionalized population aged 16 and over. All series are also in real terms (except the price level) and in logs (with the exception of the interest rate). The measure of GDP is the real GDP from the BEA accounts. The measure of consumption is the sum of non-durable and services consumption, also from the BEA accounts. Investment is measured as total private fixed investment plus consumption expenditures on durable goods (durable goods should be thought of as investment because they provide benefits in the future, just like new physical capital does). Consumption and investment data are from the BEA. Total hours is measured as total hours in the non-farm business sector, available from the BLS. Productivity is output per hour in the non-farm business sector, also from the BLS. Wages are measured as real compensation per hour in the non-farm business sector, from the BLS. The price level is measured as the implicit price deflator for GDP, from the BEA. The nominal interest rate is measured as the three month Treasury Bill rate. The real interest rate is approximately $r_t = i_t - E_t\pi_{t+1}$. I construct a measure of the ex-post real interest rate using the measured T-bill rate at time t and actual inflation from t to $t + 1$, where inflation is measured from the GDP deflator. This is ex-post because actual will not equal expected inflation in general.

All series are HP filtered. The data are from 1948 quarter 1 to 2010 quarter 3. The selected moments are shown below.

Series	Std. Dev.	Rel. Std. Dev.	Corr w/ y_t	Autocorr	Corr w/ Y_{t-4}	Corr w/ Y_{t+4}
Output	0.017	1	1.00	0.85	0.07	0.11
Consumption	0.009	0.53	0.76	0.79	0.07	0.22
Investment	0.047	2.76	0.79	0.87	-0.10	0.26
Hours	0.019	1.12	0.88	0.90	0.29	-0.03
Productivity	0.011	0.65	0.42	0.72	-0.50	0.35
Wage	0.009	0.53	0.10	0.73	-0.10	0.10
1 + Interest Rate	0.004	0.24	0.00	0.42	0.27	-0.25
Price Level	0.009	0.53	-0.13	0.91	0.09	-0.41
TFP	0.012	0.71	0.76	0.75	-0.34	0.34

The first column with numbers gives the standard deviation; the second column gives the standard deviation relative to the standard deviation of output. This can be interpreted as a

measure of relative volatility. We observe that consumption is significantly “smoother” than output; in contrast investment is more than 2.5 times more volatile than output. Hours are about as volatile as output. Productivity, the real interest rate, and real wages are significantly smoother than output. TFP is about two-thirds as volatile as output.

The next column gives the contemporaneous correlation with output. We observe that consumption, investment, hours, and TFP are all strongly correlated with output (and indeed also with each other). Productivity is somewhat more weakly correlated with output; and indeed this correlation masks a large sign shift in the correlations from strongly positive prior to the mid-1980s to weakly negative thereafter. Real wages are weakly correlated with GDP; although this probably significantly understates the true cyclical nature of wages because of a selection bias.¹ The real interest rate is essentially acyclical (that correlation will typically go slightly negative if I use ex-ante expectations of inflation as opposed to ex-post inflation). The price level is mildly countercyclical.

We see that almost all series are strongly persistent in the sense of having a large first order autocorrelation coefficient. The least persistent series is the real interest rate, but this autocorrelation is still 0.42. We observe that hours are a lagging indicator in the sense that the correlation with output lagged one year is quite positive. The real interest rate is negatively correlated with output led four quarters.

3 The Basic RBC Model and Calibration

The basic RBC model can be characterized by the first order conditions of the decentralized model, as described in class. These first order conditions are:

$$u'(C_t) = \beta E_t (u'(C_{t+1})(A_{t+1}F_K(K_{t+1}, N_{t+1}) + (1 - \delta))) \quad (2)$$

$$v'(1 - N_t) = u'(C_t)A_tF_N(K_t, N_t) \quad (3)$$

$$K_{t+1} = A_tF(K_t, N_t) - C_t + (1 - \delta)K_t \quad (4)$$

$$\ln A_t = \rho \ln A_{t-1} + \varepsilon_t \quad (5)$$

$$Y_t = A_tF(K_t, N_t) \quad (6)$$

$$Y_t = C_t + I_t \quad (7)$$

$$u'(C_t) = \beta E_t u'(C_{t+1})(1 + r_t) \quad (8)$$

$$w_t = A_tF_N(K_t, N_t) \quad (9)$$

$$R_t = A_tF_K(K_t, N_t) \quad (10)$$

I use the functional form assumptions that $u(C_t) = \ln C_t$, $v(1 - N_t) = \theta \ln(1 - N_t)$, and

¹Aggregate measures of wages are essentially averages of wages across workers. In a recession, job loss is mostly concentrated among low wage workers. Through a composition effect this tends to work to make the average wage rise when output is low because the composition of workers is shifting towards relatively more high wage earners. If one attempts a correction for this composition bias real wages actually look significantly more procyclical.

$F(K_t, N_t) = K_t^\alpha N_t^{1-\alpha}$. These first order conditions can then be re-written imposing these function forms to get:

$$\frac{1}{C_t} = \beta E_t \left(\frac{1}{C_{t+1}} (\alpha A_{t+1} K_{t+1}^{\alpha-1} N_{t+1}^{1-\alpha} + (1 - \delta)) \right) \quad (11)$$

$$\frac{\theta}{1 - N_t} = \frac{1}{C_t} (1 - \alpha) A_t K_t^\alpha N_t^{-\alpha} \quad (12)$$

$$K_{t+1} = A_t K_t^\alpha N_t^{1-\alpha} - C_t + (1 - \delta) K_t \quad (13)$$

$$\ln A_t = \rho \ln A_{t-1} + \varepsilon_t \quad (14)$$

$$Y_t = A_t K_t^\alpha N_t^{1-\alpha} \quad (15)$$

$$Y_t = C_t + I_t \quad (16)$$

$$\frac{1}{C_t} = \beta E_t \frac{1}{C_{t+1}} (1 + r_t) \quad (17)$$

$$w_t = (1 - \alpha) A_t K_t^\alpha N_t^{-\alpha} \quad (18)$$

$$R_t = \alpha A_t K_t^{\alpha-1} N_t^{1-\alpha} \quad (19)$$

We can solve for the steady state of the model as we have before. In the non-stochastic steady state we set A_t equal to its non-stochastic mean, which is 1 (0 in the log). Let variables without time subscripts denote steady states. In macroeconomics “calibration” is a particular way to choose parameter values. The gist of the calibration approach is to pick parameters so that the steady state of the model matches certain long run (i.e. average) features of the data. One of the novelties of the RBC approach (and with its calibration) was that it took a model that was (a) designed to explain the “long run” and (b) pick parameters designed to explain the “long run” but then (c) used that model and those parameters to explain the short run. You can think about the long run properties of the data as “identifying” the parameters of the model while short run moments are “overidentifying” restrictions which we can use to “test” the model. I put all these terms in quotations marks because we are not formally doing any of these things, but loosely this is what we are doing when we calibrate the model.

Let’s begin by coming up with a value for β . Go to (17), the Euler equation for risk free bonds. Evaluated in steady state, it implies:

$$\beta = \frac{1}{1 + r} \quad (20)$$

In the data the average real interest rate on riskless debt can be measures as $r = i - \pi$, where i is a “safe” nominal interest rate and π is the inflation rate. This implies an average real interest rate of something on the order of two percent (at an annual frequency). We can use that to back out the required β to make the above hold. This implies a value of $\beta = 0.995$. I’m going to round that down and set $\beta = 0.99$.

Given the function form assumptions, (18)-(19) yield:

$$wN + RK = Y$$

In other words, total output/income is the sum of payments to factors. Using the expression for the real wage, we get:

$$wN = (1 - \alpha)Y \Rightarrow \frac{wN}{Y} = (1 - \alpha) \quad (21)$$

In other words, we can measure α by looking at the fraction of total income that gets paid out in the form of wages. Doing this yields a value of $\alpha = 0.33$ or so.

Next, look at the accumulation equation for capital. This reveals that, in steady state:

$$\frac{I}{N} = \delta \frac{K}{N} \quad (22)$$

Steady state output is:

$$\frac{Y}{N} = \left(\frac{K}{N} \right)^\alpha \quad (23)$$

Hence the steady state investment-output ratio can be written:

$$\frac{I}{Y} = \delta \left(\frac{K}{N} \right)^{1-\alpha} \quad (24)$$

The steady state capital-labor ratio is:

$$\frac{K}{N} = \left(\frac{\alpha}{\frac{1}{\beta} - (1 - \delta)} \right)^{\frac{1}{1-\alpha}} \quad (25)$$

Hence, we can write the steady state investment-output ratio as:

$$\frac{I}{Y} = \frac{\alpha\delta}{\frac{1}{\beta} - (1 - \delta)} \quad (26)$$

The average investment-output ratio (depending on exactly how you measure things) is roughly 20%. I'll round this up to 22.5%, which given my values of α and β gives me a value of $\delta = 0.02$ give or take.

Using our values for α , β , and δ , this implies that the capital-labor ratio should be about 35.

Now solve the accumulation equation for steady state consumption per worker:

$$\frac{C}{N} = \left(\frac{K}{N} \right)^\alpha - \delta \frac{K}{N} \quad (27)$$

Now solve the first order condition for labor supply for another expression for consumption to labor:

$$\frac{C}{N} = \frac{1}{\theta} \frac{1-N}{N} (1-\alpha) \left(\frac{K}{N}\right)^\alpha \quad (28)$$

Equate these:

$$\frac{1}{\theta} \frac{1-N}{N} (1-\alpha) \left(\frac{K}{N}\right)^\alpha = \left(\frac{K}{N}\right)^\alpha - \delta \frac{K}{N} \quad (29)$$

Solve for θ , taking N as given:

$$\theta = \frac{\frac{1-N}{N} (1-\alpha) \left(\frac{K}{N}\right)^\alpha}{\left(\frac{K}{N}\right)^\alpha - \delta \frac{K}{N}} \quad (30)$$

In the data, people work about one-third of their time endowment (this may be a bit of an overstatement, but it's simple enough). So set $N = 0.33$ and solve for θ . I get 1.75.

That is all the parameters of the model that do not govern the stochastic process for A_t , which we turn to next.

4 How Well Can the Model Fit the Data?

We need a systematic way to calibrate the parameters governing the process for TFP. I begin by linearly detrending my measure of log TFP. I need to linearly detrend because it is (implicitly) assumed in the model that a linear, deterministic trend drives the observed trends in the actual data (though we ignore the small variations to the first order conditions that are necessary to be consistent with that).

As such, I regress my empirical measure of TFP on a constant and a linear time trend:

$$\ln \widehat{A}_t = \phi_0 + \phi_1 t_t + u_t \quad (31)$$

I get the following estimates: $\phi_0 = 0.147$ and $\phi_1 = 0.003$. This basically can be interpreted that TFP grows, on average, at about 0.012, or 1.2 percent, per year. Then I take the measured residual, \widehat{u}_t , which can be interpreted as the detrended TFP series, and estimate an AR(1) process:

$$\widehat{u}_t = \rho \widehat{u}_{t-1} + e_t \quad (32)$$

I get the following estimates: $\rho = 0.974$ and the standard deviation of the residual of 0.009.

I use these parameters to solve the model. Below is a table of moments from the model. These moments are HP filtered to facilitate comparison with the moments from the data:

Series	Std. Dev.	Rel. Std. Dev.	Corr w/ y_t	Autocorr	Corr w/ Y_{t-4}	Corr w/ Y_{t+4}
Output	0.017	1.00	1.00	0.72	0.13	0.13
Consumption	0.006	0.35	0.95	0.78	0.34	-0.03
Investment	0.056	3.29	0.99	0.71	0.04	0.20
Hours	0.007	0.41	0.98	0.71	0.00	-0.23
Productivity	0.010	0.59	0.99	0.74	-0.50	0.06
Wage	0.010	0.59	0.99	0.74	0.74	0.06
1+Interest Rate	0.001	0.03	0.96	0.71	-0.05	0.26
TFP	0.012	0.69	0.99	0.72	0.11	0.15

Let’s take a moment to compare the numbers in this table (from model-generated data) to the numbers from the US data. Let’s begin by focusing on areas where the model does well. We see that the model does a pretty good job at matching the volatilities of output, consumption, and investment – in particular, consumption is significantly smoother than output, and investment is significantly more volatile than output. The model does a good job of matching the volatilities of labor productivity and TFP. The model also does a good job at matching the own autocorrelations – the series are all persistent with first order autocorrelation coefficients typically in the neighborhood of 0.75. Lastly, the model captures the fact that most quantity series (consumption, investment, hours, productivity, and TFP) are quite procyclical (high contemporaneous correlations with output), though these correlations are too high in the model relative to the data.

Now let’s move on to where the model does less well. Whereas it does pretty well with “quantity” moments (see above), it does much less well with prices. In particular, the model does not generate enough volatility of interest rates (relative standard deviation of 0.03 in model vs. 0.24 in the data). Further, it generates wages and real interest rates that are *far* too procyclical relative to the data. In the data, wages are very modestly procyclical and real interest rates are acyclical or countercyclical, depending on how you measure them. There is some evidence that aggregate wage data understates the procyclicality of wages due to a composition bias (Solon, Barsky, and Parker, 1994), so the wage cyclicality can potentially be reconciled. It’s much harder to deal with the interest rate cyclicality. The model does not do great at the dynamic correlations. One particular area of failure is the fact that real interest rates *positively* lead output in the model, whereas they *negatively* lead output in the data (King and Watson, 1988). Finally, another failure of the model is that it does not generate sufficient volatility of hours – in the data, hours are actually slightly more volatile than output (this relative volatility has risen over time), but in the model hours are about half as volatile as output.

Some of these deficiencies are easier to deal with than others. For example, we can employ a version of Hansen (1985) and Rogerson (1988) “indivisible labor model” to get what essentially amounts to infinitely elastic labor supply. This will work to raise the volatility of hours and lower the cyclicality of wages. We could also add money to the model in a way that doesn’t change

any of the above results, but which makes the model capable of matching the countercyclicity of the price level. In addition, we can add shocks to things other than technology – things like government spending shocks or distortionary tax shocks (McGrattan, 1994). These can work to lower the contemporaneous correlations with output, which are too high in the data.

At an even deeper level, people have criticized RBC models because they don't seem particularly realistic. To generate fluctuations that resemble those in the US, one needs large, high frequency variation in A_t . No other shock (e.g. government spending, preferences, etc.), within the confines of the model, can be the main driving force behind the data, although, as mentioned in the preceding paragraph, adding other shocks can make the overall correlations with output “look better”. The reason is because hours and consumption are highly positively correlated in the data (correlation between HP filtered series of 0.78). Combine (3) with (9), and use our baseline functional form assumptions. You get:

$$\frac{\theta}{1 - N_t} = \frac{1}{C_t} (1 - \alpha) A_t K_t^\alpha N_t^{-\alpha}$$

Let's log-linearize this:

$$\begin{aligned} \ln \theta - \ln(1 - N_t) &= -\ln C_t + \ln(1 - \alpha) + \ln A_t + \alpha \ln K_t - \alpha \ln N_t \\ \left(\frac{N^*}{1 - N^*} \right) \tilde{N}_t &= -\tilde{C}_t + \tilde{A}_t + \alpha \tilde{K}_t - \alpha \tilde{N}_t \end{aligned}$$

Again define $\gamma = \frac{N^*}{1 - N^*}$ and we can write this as:

$$\tilde{N}_t = \left(\frac{1}{\gamma + \alpha} \right) \left(-\tilde{C}_t + \tilde{A}_t + \alpha \tilde{K}_t \right)$$

If neither K nor A move, then it must be the case that N and C move in opposite directions. Since shocks to K don't seem like a very plausible explanation, we are left with shocks to A . These shocks must be the main driving force behind the data, otherwise consumption and hours will not be correlated strongly enough.

People have said that assuming business cycles result from exogenous technological progress (or “worse”, regress) is unappealing (Summers, 1986). To generate recessions, one needs A to decline. What does this even mean, especially for a modern economy? If A is moving around so much, why don't we read about it in newspapers? Rather, we typically read about output and employment declining with some speculation as to why, but that speculation rarely if ever mentions fluctuations in “technological possibilities” or something similar.

Hence, critics of the real business cycle model are uncomfortable with the facts that it is (a) driven by technology “shocks” and that (b) these shocks must be large and sometimes negative. Hence, much of business cycle research since the 1980s has been involved in modifying the basic model to (a) allow other shocks to “matter” in a way that they can't in the basic model (e.g. monetary policy) and (b) generating better and more realistic mechanisms for the model to take

“small” shocks (as opposed to large) and produce relatively large business cycles.

The real business cycle model has a fairly weak *amplification* mechanism and an even weaker *propagation* mechanism. *Amplification* refers to a model’s ability to have output react by substantially more than the exogenous shock – i.e. to take “small” shocks and produce “large” fluctuations. The only amplification mechanism here is labor supply, and it is fairly weak. The volatility of output in the model is only about 40 percent bigger than the volatility of TFP that is fed into the model. This is not trivial, but it means that a substantial amount of output volatility is in essence “assumed” by the shock – if output were simply equal to the productivity shock, we would generate about 70 percent of the output volatility that the model does.

Propagation refers to a model’s ability to make shocks have persistent effects. The only propagation mechanism in the model is capital accumulation, but this is weak as well. We can see this by noting that output is only a little more persistent (as measured by the first order autocorrelation coefficient) than is the assumed driving force, TFP. If I set $\rho = 0.0$ (so I assume no persistence) I actually get both output and TFP being mildly negatively autocorrelated (-0.07 for both). The bottom line is that you don’t get any more persistence in output than you assume in TFP in this model.

Much of the rest of what we do in this class will involve adding amplification and propagation mechanisms and re-tooling the model in other ways such that non-technology shocks (e.g. different sources of “demand”) can have large real effects.

5 A Hidden Calibration: Labor Supply

In the preference specification I used above, the first order condition for labor supply can be written:

$$\theta \frac{1}{1 - N_t} = \frac{1}{C_t} w_t \quad (33)$$

Log-linearizing about the steady state yields:

$$\gamma \tilde{N}_t = -\tilde{C}_t + \tilde{w}_t \quad (34)$$

Here $\gamma = \frac{N^*}{1 - N^*}$. The Frisch labor supply elasticity is defined as the derivative of log hours to the log wage, holding the marginal utility of wealth fixed (which is here equal to the marginal utility of consumption). Since $\tilde{N}_t = \ln N_t - \ln N^*$, the Frisch elasticity can be seen to be $\frac{1}{\gamma}$ here:

$$\left. \frac{d\tilde{N}_t}{d\tilde{w}_t} \right|_{\tilde{\lambda}_t} = \frac{1}{\gamma}$$

If we want to target $N^* = 1/3$, this implies that the Frisch labor supply elasticity is about 2. This is much larger than most microeconomists would tell you that elasticity is. This elasticity is important in the model, because it in essence governs how “flat” the labor supply curve is – the flatter the labor supply curve is, the more fluctuations in labor demand (driven by productivity shocks) affect labor hours and hence output.

An alternative way to specify preferences is not over leisure but rather over labor. Consider the following preference specification:

$$u(C_t) - v(N_t) = \ln C_t - \theta \frac{N_t^{1+\chi}}{1+\chi}$$

As long as $\chi \geq 0$, we have that utility is (i) increasing and (ii) concave in leisure (where leisure is defined as $L_t = 1 - N_t$). The first order condition for labor supply can now be written:

$$\theta N_t^\chi = \frac{1}{C_t} w_t \quad (35)$$

Log-linearizing about the steady state, we get:

$$\chi \tilde{N}_t = -\tilde{C}_t + \tilde{w}_t \quad (36)$$

This means that the Frisch elasticity here is $\frac{1}{\chi}$. What's nice about this is that the Frisch elasticity here is controlled by this one parameter; in the case above, the Frisch elasticity depends on N^* , which is in turn a function of many parameters. I can still calibrate θ to generate a target level of N^* here, but that will not affect the Frisch elasticity, which is solely governed by the parameter χ . I can write the first order condition for labor supply with these preferences evaluated in steady state as:

$$\theta N^{1+\chi} = \frac{N}{C} (1 - \alpha) \left(\frac{K}{N} \right)^\alpha$$

Given that I know $\frac{C}{N}$, if I have a target level of N in steady state I can solve for θ to hit that as:

$$\theta = \frac{1}{N^{1+\chi}} \frac{(1 - \alpha) \left(\frac{K}{N} \right)^\alpha}{\left(\frac{K}{N} \right)^\alpha - \delta \frac{K}{N}} \quad (37)$$

Suppose I work with my existing calibration, but set $\chi = 1$. I need a value of $\theta = 7.883$ for this to be consistent with steady state hours of 0.33. I re-solve the model and produce volatilities of the main variables of interest.

Series	Std. Dev.	Rel. Std. Dev.
Output	0.015	1.00
Consumption	0.006	0.40
Investment	0.049	3.26
Hours	0.005	0.33
Productivity	0.010	0.66
Wage	0.010	0.55
1+Interest Rate	0.001	0.03
TFP	0.012	0.80

In this specification of the model there is even less output volatility (0.015 instead of 0.017). This means that there is less amplification. We also observe a worse relative volatility of hours to output here than before (0.3 versus 0.4); both are still far too low relative to the data but now things are worse. The other moments (correlations, autocorrelations) are not substantially affected by this change.

There are two important take-aways here. First, functional form assumptions may contain “hidden calibrations” – with the log-log utility specification, targeting steady state hours of 0.33 implied a relatively large Frisch labor supply elasticity without us really knowing it. Second, a chief problem in the basic RBC model (and really, in almost all macro models) is generating sufficiently volatility of labor input to output (which is equivalent to saying it’s difficult to get much amplification). One way to get more amplification is to assume more elastic labor supply (in this preference specification this could be accomplished via lower values of χ).