Quantitative Macroeconomics: An Introduction

Dirk Krueger[1]
Department of Economics
University of Pennsylvania

April 2007

[1]The Author thanks Jesus Fernandez Villaverde for sharing much of his material on the same issue. This book is dedicated to my son Nicolo.

Contents

Para o ponto 3, ler apenas esta parte

I Motivation and Data **1**

1 Introduction **3**

1.1 The Questions 3
1.2 The Approach and the Structure of the Book 3

2 Basic Business Cycle Facts **5**

2.1 Decomposition of Growth Trend and Business Cycles 5
2.2 Basic Facts 15

II The Real Business Cycle (RBC) Model and Its Extensions **17**

3 Set-Up of the Basic Model **21**

3.1 Households 21
3.2 Firms 24
3.3 Aggregate Resource Constraint 27
3.4 Competitive Equilibrium 27
3.5 Characterization of Equilibrium 30

4 Social Planner Problem and Competitive Equilibrium **33**

4.1 The Social Planner Problem 33
4.2 Characterization of Solution 34
4.3 The Welfare Theorems 35
4.4 Appendix: More Rigorous Math 36

5 Steady State Analysis **39**

5.1 Characterization of the Steady State 39
5.2 Golden Rule and Modified Golden Rule 40

6 Dynamic Analysis **43**

6.1 An Example with Analytical Solution 44
6.2 Linearization of the Euler Equation 47
6.2.1 Preliminaries 47
6.2.2 Doing the Linearization 48

6.3 Analysis of the Results 54
6.3.1 Plotting the Policy Function 54
6.3.2 Impulse Response Functions 54
6.3.3 Simulations 54
6.4 Summary 56

7 A Note on Economic Growth **57**
7.1 Preliminary Assumptions and Definitions 57
7.2 Reformulation of Problem in Efficiency Units 58
7.3 Analysis 60
7.4 The Balanced Growth Path 60

8 Calibration **63**
8.1 Long Run Growth Rates 63
8.2 Capital and Labor Share 64
8.3 The Depreciation Rate 64
8.4 The Technology Constant A 65
8.5 Preference Parameters 65
8.6 Summary 66

9 Adding Labor Supply **67**
9.1 The Modified Social Planner Problem 67
9.2 Labor Lotteries 68
9.3 Analyzing the Model with Labor 69
9.4 A Note on Calibration 72
9.5 Intertemporal Substitution of Labor Supply: A Simple Example 73
9.6 A Remark on Decentralization 74

10 Stochastic Technology Shocks: The Full RBC Model **77**
10.1 The Basic Idea 77
10.2 Specifying a Process for Technology Shocks 78
10.3 Analysis 80
10.4 What are these Technology Shocks and How to Measure Them? . 82

III Evaluating the Model 85

11 Technology Shocks and Business Cycles **89**
11.1 Impulse Response Functions 89
11.2 Comparing Business Cycle Statistics of Model and Data 92
11.3 Counterfactual Experiments 94

IV Welfare and Policy Questions 95

12 The Cost of Business Cycles **99**

List of Tables

2.1 Cyclical Behavior of Real GDP, US, 1947-2004 15

8.1 Parameter Values 66

9.1 Labor and Consumption Allocations 74

11.1 Business Cycles: Data and Model 93

Preface

In these notes I will describe how to use standard neoclassical theory to explain business cycle fluctuations.

Part I

Motivation and Data

Chapter 1

Introduction

1.1 The Questions

Business cycles are both important and, despite a large amount of economic research, still incompletely understood. While we made progress since the following quote

> The modern world regards business cycles much as the ancient Egyptians regarded the overflowing of the Nile. The phenomenon recurs at intervals, it is of great importance to everyone, and natural causes of it are not in sight. (John Bates Clark, 1898)

there is still a lot that remains to be learned. In this class we will ask, and try to at least try to partially answer the following questions

- What are them empirical characteristics of business cycles?
- What brings business cycles about?
- What propagates them?
- Who is most affected and how large would be the welfare gains of eliminating them?
- What can economic policy, both fiscal and monetary policy do in order to soften or eliminate business cycles?
- Should the government try to do so?

1.2 The Approach and the Structure of the Book

Our methodological approach will be to use economic theory and empirical data to answer these questions. We will proceed in four basic steps with our analysis. First we will document the stylized facts that characterize business cycles in modern societies. Using real data, mostly for the US where the data situation is most favorable we will first discuss how to separate business cycle fluctuations and economic growth from the data on economic activity, especially real gross domestic product. The method for doing so is called filtering. Our stylized facts will be quantitative in nature, that is, we will not be content with saying that the growth rate of real GDP goes up and down, but we want to quantify these fluctuations, we want to document how long a business cycle lasts, whether recessions and expansions last equally long, and how large and small growth rates of real GDP or deviations from the long run growth trend are. In a second step we will then construct a theoretical business cycle model that we will use to explain business cycles. We will build up this model up in several steps, starting as a benchmark with the neoclassical growth model. At each step we will evaluate how well the model does in explaining business cycles from a quantitative point. In the process we will also have to discuss how our model is best parameterized (a process we will either call calibration or estimation, depending on the exact procedure) and how it is solved (it will turn out that we will not always be able to deal with our model analytically, but sometimes will have to resort to numerical techniques to solve the model). Into the basic growth model we will first introduce technology shocks and endogenous labor supply, which leads us to the canonical Real Business Cycle Model. Further extensions will include capital adjustment costs, two sector models and sticky prices and monopolistic competition. Once the last two elements are incorporated into the model we have arrived at the New Keynesian business cycle model. In a third step we will then evaluate the ability of the different versions of the model to generate business cycles of realistic magnitude. Once the model(s) do a satisfactory job in explaining the data, we can go on and ask normative questions. The final fourth step of our analysis will first quantify how large the welfare costs of business cycles are and then analyze, within our models, how effective monetary and fiscal policy is to tame cyclical fluctuations of the economy.

Chapter 2

Basic Business Cycle Facts

In this chapter we want to accomplish two things. First we will discuss how to distill business cycle facts from the data. The main object of macroeconomists is aggregate economic activity, that is, total production in an economy. This is usually measured by real GDP or, if one is more interested in living standards of households, by real GDP per capita or worker. But plotting the time series of real GDP we see that not only does it fluctuate over time, but it also has a secular growth trend, that it, is goes up on average. For the study of business cycle we have to purge the data from this long run trend, that is, take it out of the data. The procedure for doing so is called filtering, and we will discuss how to best filter the data in order to divide the data into a long run growth trend and business cycle fluctuations.

Second, after having de-trended the data, we want to take the business cycle component of the data and document the main stylized facts of business cycles, that is, study what are the main characteristics of business cycles. We want to document the length of a typical business cycle, whether the business cycle is symmetric, the size of deviations from the long run growth trend, and the persistence of deviations from the long-run growth trend.

2.1 Decomposition of Growth Trend and Business Cycles

In Figure 2.1 we plot the natural logarithm of real GDP for the US from 1947 to 2004. The reason we start with US data, besides it representing the biggest economy in the world, is that the data situation for this country is quite favorable. There are no obvious trend breaks, due to, say, major wars, change in the geographic structure of the country (such as the German unification), and real data are available with consistent deflation for price changes over the entire sample period. In exercises you will have the opportunity to study your country of choice, but you should be warned already that for Germany consistent data for real GDP are available only for much shorter time periods (not least because of the German unification).

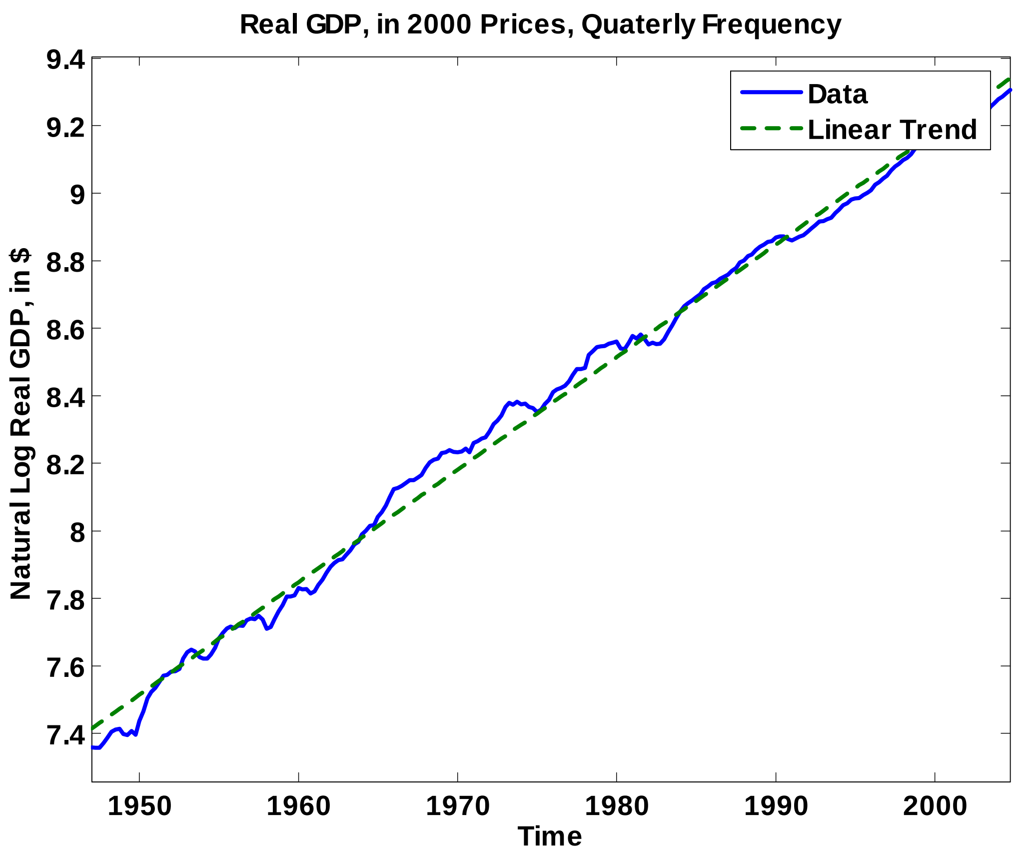

Figure 2.1: Natural Logarithm of real GDP for the US, in constant 2000 prices.

A short discussion of the data themselves. The frequency of the data is quarterly, that is, we have one observation fro real GDP in each quarter. However, the observation is for real GDP for the preceding twelve months, not just the last three months. In that way seasonal influences on GDP are controlled for. The base year for the data is 2000, that is, all numbers are in 2000 US dollars. In terms of units, the data are in billion US dollars (Milliarden). For example for 2004 real GDP is about

$$\exp(9.3) \approx \$10.000 \text{ Billion} = \$10 \text{ Trillion}$$

or about $\$36,000$ per capita. Finally, why would we plot the natural logarithm of the data, rather than the data themselves. Here is the reason: suppose an economic variable, say real GDP, denoted by Y, grows at a constant rate, say g, over time. Then we have

$$Y_t = (1+g)^t Y_0 \tag{2.1}$$

where Y_0 is real GDP at some starting date of the data, and Y_t is real GDP in period t. Now let us take logs (and whenever I say logs, I mean natural logs, that is, the logarithm with base e, where $e \approx 2.781$ is Euler's constant) of equation (2.1). This yields

$$\begin{aligned}
\log(Y_t) &= \log\left[(1+g)^t Y_0\right] \\
&= \log(Y_0) + \log\left[(1+g)^t\right] \\
&= \log(Y_0) + t * \log\left[1+g\right]
\end{aligned}$$

where we made use of some basic rules for logarithms.

What is important about this is that *if* a variable, say real GDP, grows at a constant rate g over time, then if we plot the logarithm of that variable it is exactly a straight line with intercept Y_0 and slope

$$\text{slope} = \log(1+g) \approx g$$

where the approximation in the last equation is quite accurate as long as g is not too large.[1] Similarly, we need not start at time $s = 0$. Suppose our data starts at an arbitrary date s (in the example $s = 1947$). Then if our data grows at a constant rate g, the figure for $\log(Y_t)$ is given by

$$\log(Y_t) = \log(Y_s) + (t-s)\log(1+g),$$

and if $s = 1947$, then

$$\log(Y_t) = \log(Y_{1947}) + (t-1947)\log(1+g)$$

Thus, if real GDP grew at a constant rate, a plot of the natural logarithm of real GDP should be straight line, with slope equal to the constant growth rate. Figure 2.1 shows that this is not too bad of a first approximation.

In this course, however, we are mostly interested in the deviations of actual real GDP from its long run growth trend. First we want to mention that the decision what part of the data is considered a growth phenomenon and what is considered a business cycle phenomenon is somewhat arbitrary. To quote Cooley and Prescott (1995)

---

[1] The fact that $\log(1+g) \approx g$ can be seen from the Taylor series expansion of $\log(1+g)$ around $g = 0$. This yields

$$\begin{aligned}
\log(1+g) &= \log(1) + \frac{g-0}{1} - \frac{1}{2}(g-0)^2 + \frac{1}{6}(g-0)^3 + \ldots \\
&= g - \frac{1}{2}g^2 + \frac{1}{6}g^3 + \ldots \\
&\approx g
\end{aligned}$$

because the terms where g is raised to a power are small relative to g, whenever g is not too big.

> Every researcher who has studied growth and/or business cycle fluctuations has faced the problem of how to represent those features of economic data that are associated with long-term growth and those that are associated with the business cycle - the deviation from the growth path. Kuznets, Mitchell and Burns and Mitchell [early papers on business cycles] all employed techniques (moving averages, piecewise trends etc.) that define the growth component of the data in order to study the fluctuations of variables around the long-run growth path defined by the growth component. Whatever choice one makes about this is somewhat arbitrary. There is no single correct way to represent these components. They ares imply different features of the same observed data.

Thus, while it is clear that "business cycle fluctuations" are the deviation of a key economic variable of interest (mostly real GDP) from a growth trend, what is unclear is how to model this growth trend. In the Figure above we made the choice of representing the long run growth trend as a function that grows at constant rate g over time. Consequently the business cycle component associated with this growth trend is given by

$$y_t = \log(Y_t) - \log(Y_t^{trend}) = \log\left(\frac{Y_t}{Y_t^{trend}}\right) = \log\left(\frac{Y_t - Y_t^{trend}}{Y_t^{trend}} + 1\right) \approx \frac{Y_t - Y_t^{trend}}{Y_t^{trend}}$$

By using logs, the deviation of the actual log real GDP from its trend roughly equals its percentage deviation from the long run growth trend. In Figure 2.2 we plot this deviation y_t from trend, when the trend is defined simply as a linear growth trend. From now on we will always use y_t to denote the business cycle component of real GDP.

The figure shows that business cycles, so defined, are characterized by fairly substantial deviations from the long run growth trend. The deviations have magnitudes of up to 10%, and they are quite persistent: if in a given quarter real GDP is above trend, it looks as if it is more likely that real GDP is above trend in the next quarter as well. We will formalize this high degree of persistence below by defining and then computing autocorrelations of real GDP. Before doing this, however, observe that if we define the trend component simply as a linear trend, the figure suggest only three basic periods. From 1947 to 1965 real GDP was below trend, from 1966 to 1982 (or 1991) it was above trend, and then it fell below trend again. According to this, the postwar US economy only had two recessions and one expansion. This seems unreasonable and is due to the fact that by defining the trend they way we have we load everything in the data this is not growing at a constant rate into the business cycle. More medium term changes in the growth rate are attributed as business cycle fluctuations. While, as we argued above the division in trend and fluctuations is always somewhat arbitrary, most business cycle researchers and practitioners take the view that the growth trend should be defined more flexibly, so that the business cycle component only captures fluctuations at the three to eight year frequency.

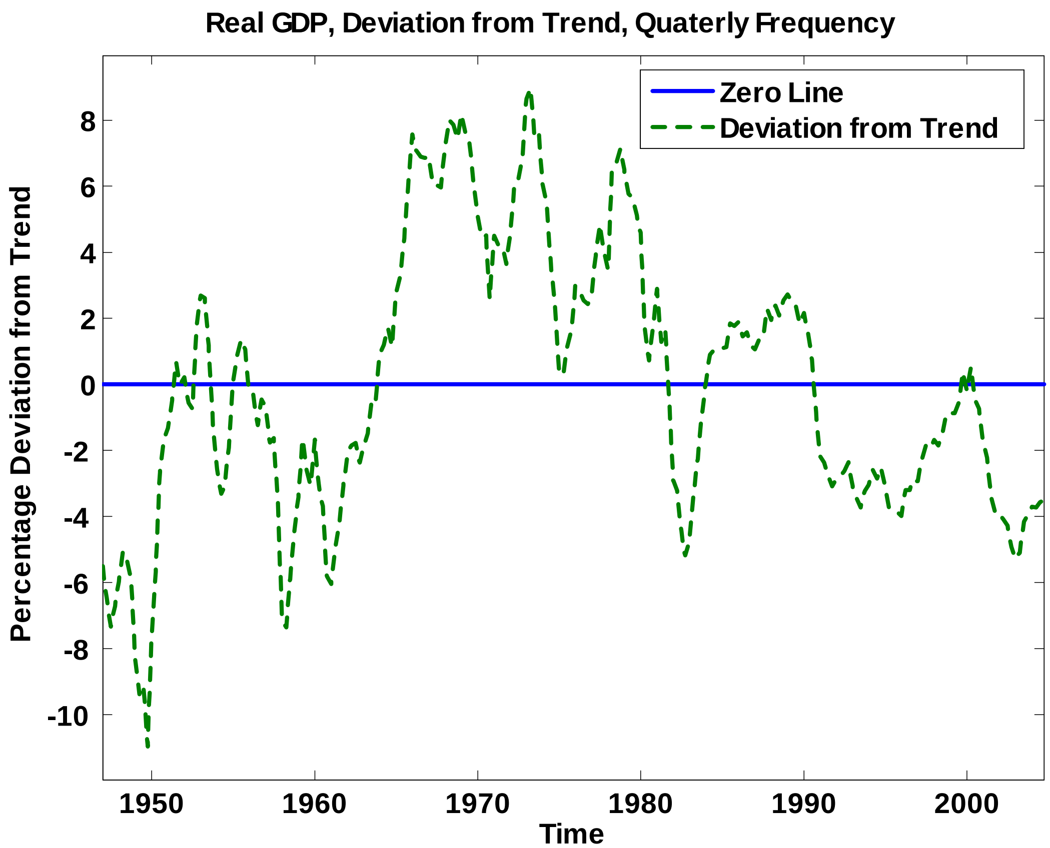

Figure 2.2: Deviations of Real GDP from Linear Trend, 1947-2004

So in practice most business cycle researchers measure business cycle fluctuations using one of three statistics: a) growth rates of real GDP, b) the cyclical component of Hodrick-Prescott filtered data, c) the data component of the appropriate frequency of a band pass filter. We will discuss the first two of these methods, and only briefly mention the third, because its understanding requires a discussion of spectral methods which you may know if you studied physics or a particular area of finance, but which I do not want to teach in this class.

Figure 2.3 plots growth rates of real GDP for the US. Remember that even though the data frequency is quarterly, these are growth rates for yearly real GDP. The average growth rate over the sample period is $3,3\%$. As a side remark, about one third of this growth is due to population and thus labor force growth, and two thirds are due to higher real GDP per capita. Note that if the log of

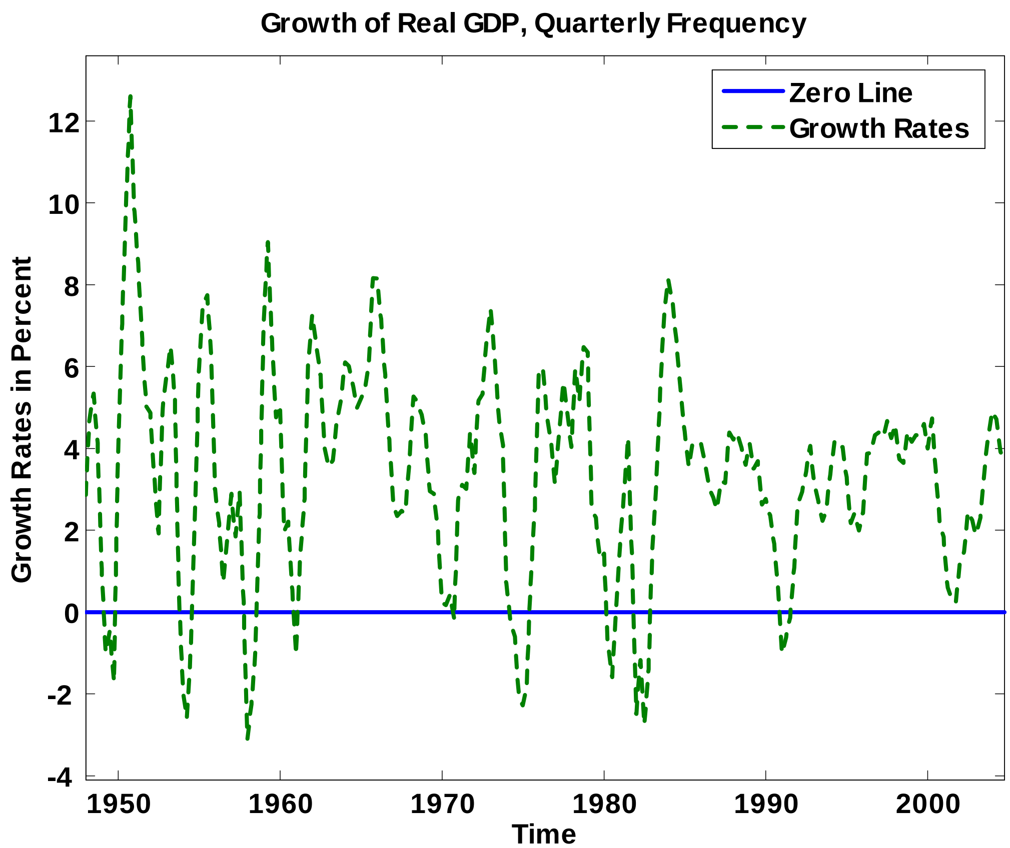

Figure 2.3: Growth Rates of Real GDP for the US, 1947-2004

real GDP would really follow exactly a linear trend, then

$$\log(Y_t) = \alpha + g(t - 1947)$$

and thus the growth rates would be given by

$$\begin{aligned} g_Y(t-1,t) &= \frac{Y_t - Y_{t-1}}{Y_{t-1}} \approx \log\left(\frac{Y_t - Y_{t-1}}{Y_{t-1}} + 1\right) = \log\left(\frac{Y_t}{Y_{t-1}}\right) = \log(Y_t) - \log(Y_{t-1}) \\ &= \alpha + g(t-1947) - \alpha - g(t-1-1947) = g \end{aligned}$$

that is, then the plot above would look like a straight line equal to the average growth rate of $3,3\%$. Note however that since the actual data do not follow this constant growth path exactly, plotting growth rates and plotting the residuals of a linear egression does *not* result in the same plot shifted by $3,3\%$ upward (compare Figure 2.3 with the trend line at $3,3\%$ and Figure 2.2). Also note that, when dealing with $y_t = \log(Y_t)$, computing the growth rate

$$g_Y(t-1,t) = \log(Y_t) - \log(Y_{t-1}) = y_t - y_{t-1}$$

amounts to plotting the data in deviations from its value in the previous quarter. Thus effectively all variations of the data longer than one quarter are filtered out by this procedure, leaving only the very highest frequency fluctuations behind. This is why the plot looks very "jumpy", and observations in successive quarters not very correlated. We will document this fact more precisely below.

While the popular discussion mostly uses growth rates to talk about the state of the business cycle, academic economists tend to separate growth and cycle components of the data by applying a filter to the data. In fact, specifying the deterministic constant growth trend above and interpreting the deviations as cycle was nothing else than applying one such, fairly heuristic filter to the data. One filter that has enjoyed widespread popularity is the so-called Hodrick-Prescott filter, or HP-filter, for short. The goal of the filter is as before : specify a growth trend such that the deviations from that trend can be interpreted as business cycle fluctuations. Let us describe this filter in more detail and try to interpret what it does. As always we want to decompose the raw data, $\log(Y_t)$ into a growth trend $y_t^{trend} = \log(Y_t^{trend})$ and a cyclical component $y_t = \log(Y_t^{cycle})$ such that

$$\begin{aligned}\log(Y_t) &= \log(Y_t^{trend}) + \log(Y_t^{cycle}) \\ y_t &= \log(Y_t) - y_t^{trend}\end{aligned}$$

Of course the key question is how to pick y_t and y_t^{trend} from the data? The HP-filter proposes to make this decomposition by solving the following minimization problem

$$\min_{\{y_t, y_t^{trend}\}} \sum_{t=1}^{T} (y_t)^2 + \lambda \sum_{t=1}^{T} \left[(y_{t+1}^{trend} - y_t^{trend}) - (y_t^{trend} - y_{t-1}^{trend}) \right]^2 \tag{2.2}$$

subject to

$$y_t + y_t^{trend} = \log(Y_t) \tag{2.3}$$

where T is the last period of the data. Note that we are given the data $\{\log(Y_t)\}_{t=0}^{T}$, so the right hand side of 2.3 is a known and given number, for each time period. Implicit in this minimization problem is the following trade-off in choosing the trend. We may want the trend component to be a smooth function, but we also may want to make the trend component track the actual data to some degree, in order to capture also some fluctuations in the data that are of lower frequency than business cycles. These two considerations are traded off by the parameter λ. If λ is big, we want to make the terms

$$\left[(y_{t+1}^{trend} - y_t^{trend}) - (y_t^{trend} - y_{t-1}^{trend}) \right]^2$$

small. But the term

$$\begin{aligned} & (y_{t+1}^{trend} - y_t^{trend}) - (y_t^{trend} - y_{t-1}^{trend}) \\ = & \left[\log(Y_{t+1}^{trend}) - \log(Y_t^{trend})\right] - \left[\log(Y_t^{trend}) - \log(Y_{t-1}^{trend})\right] \\ = & g_{Y^{trend}}(t, t+1) - g_{Y^{trend}}(t-1, t) \end{aligned}$$

is nothing else but the change in the growth rate of the trend component. Thus a high λ makes it optimal to have a trend component with fairly constant slope. In the extreme as $\lambda \to \infty$, the weight on the second term is so big that it is optimal to set this term to 0 for all time periods, that is,

$$\begin{aligned} g_{Y^{trend}}(t, t+1) - g_{Y^{trend}}(t-1, t) &= 0 \\ g_{Y^{trend}}(t, t+1) &= g_{Y^{trend}}(t-1, t) = g \end{aligned}$$

and thus

$$\begin{aligned} y_t^{trend} - y_{t-1}^{trend} &= g \\ y_t^{trend} &= y_{t-1}^{trend} + g \end{aligned}$$

for all time periods t. But this is nothing else but our constant growth linear trend that we started with. This is, the HP-filter has the linear trend as a special case.

Now consider the other extreme, in which we value a lot the ability of the trend to follow the real data. Suppose we set $\lambda = 0$, then the objective function to minimize becomes

$$\min_{\{y_t, y_t^{trend}\}} \sum_{t=1}^{T} (y_t)^2 \text{ subject to } y_t + y_t^{trend} = \log(Y_t)$$

or substituting in for y_t

$$\min_{\{y_t^{trend}\}} \sum_{t=1}^{T} \left(\log(Y_t) - y_t^{trend}\right)^2$$

and the solution evidently is

$$y_t^{trend} = \log(Y_t)$$

that is, the trend is equal to the actual data series and the deviations from the trend, our business cycle fluctuations, are identically equal to zero. These extremes show that we want to pick a λ bigger than zero (otherwise there are no business cycle fluctuations and all of the data are due to the long run trend) and smaller than ∞ (so that the trend picks up some medium run variation in addition to long run growth and does not leave everything but the longest run movements to the fluctuations part).

Thus which λ to choose must be guided by our objective of filtering out business cycle fluctuations, that is, fluctuations in the data with frequency of between three to five years. Which choice of λ accomplishes this depends crucially on the frequency of the data; for quarterly data a value of $\lambda = 1600$ is commonly used, which loads into the trend component fluctuations that occur at frequencies of roughly eight years or longer. Note that for any $\lambda \in (0, \infty)$ it is not completely straightforward to solve the minimization problem associated with the HP-filter. But luckily there exists pre-programed computer code in just about any software package to do this.

Figure 2.4 shows the trend component of the data, derived from the HP-filter with a smoothing parameter $\lambda = 1600$. That is, the figure plots y_t^{trend} against time. We observe that, in contrast to a simply constant growth trend the HP-trend captures some of the medium frequency variation of the data, which was the goal of applying the HP-filter in the first place. But the HP trend component is still much smoother than the data themselves.

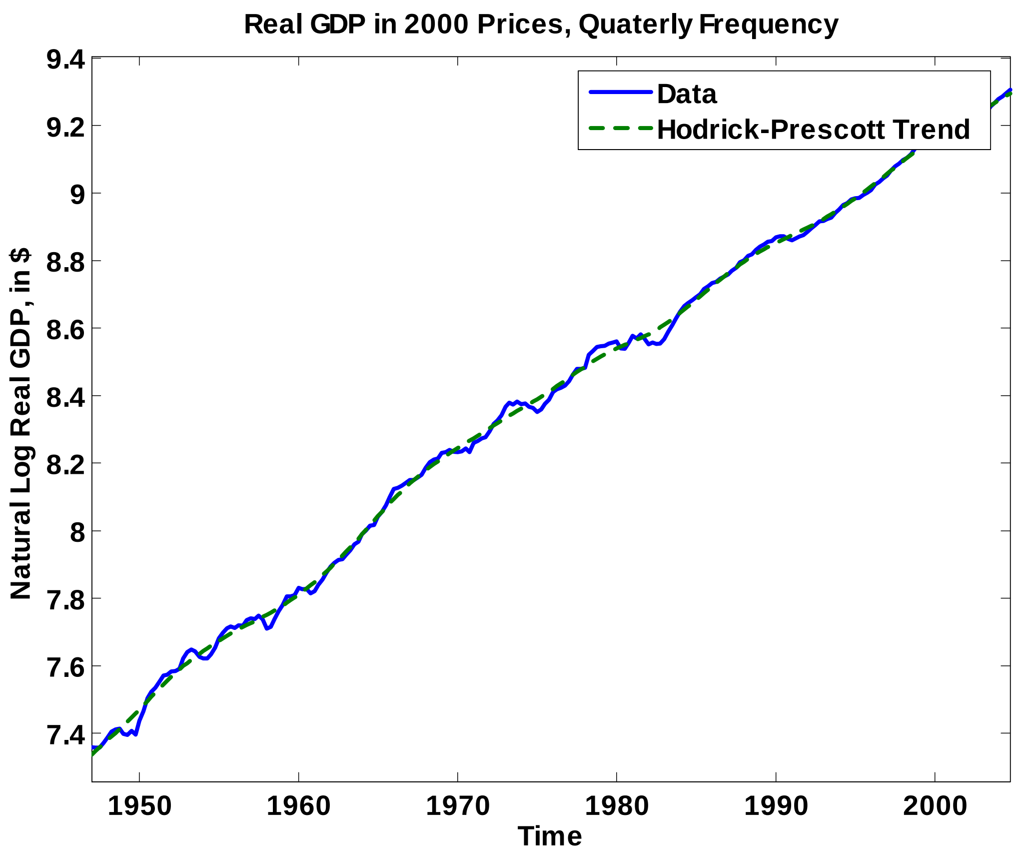

Figure 2.4: Trend Component of HP-Filtered Real GDP for the US, 1947-2004

Our true object of interest is the cyclical component that comes out of the

HP filtering Figure 2.5 displays the business cycle component of the HP-filtered data, y_t. As in Figures 2.3 and 2.2 this figure shows that the cyclical variation in real GDP can be sizeable, up to $4-6\%$ in both directions from trend. Clearly visible are the mid-seventies and early eighties recessions, both partially due to the two oil price shocks, the recession of the early 90's that cost George W. Bush's dad his job and the fairly mild (by historical standard) recent recession in 2001.

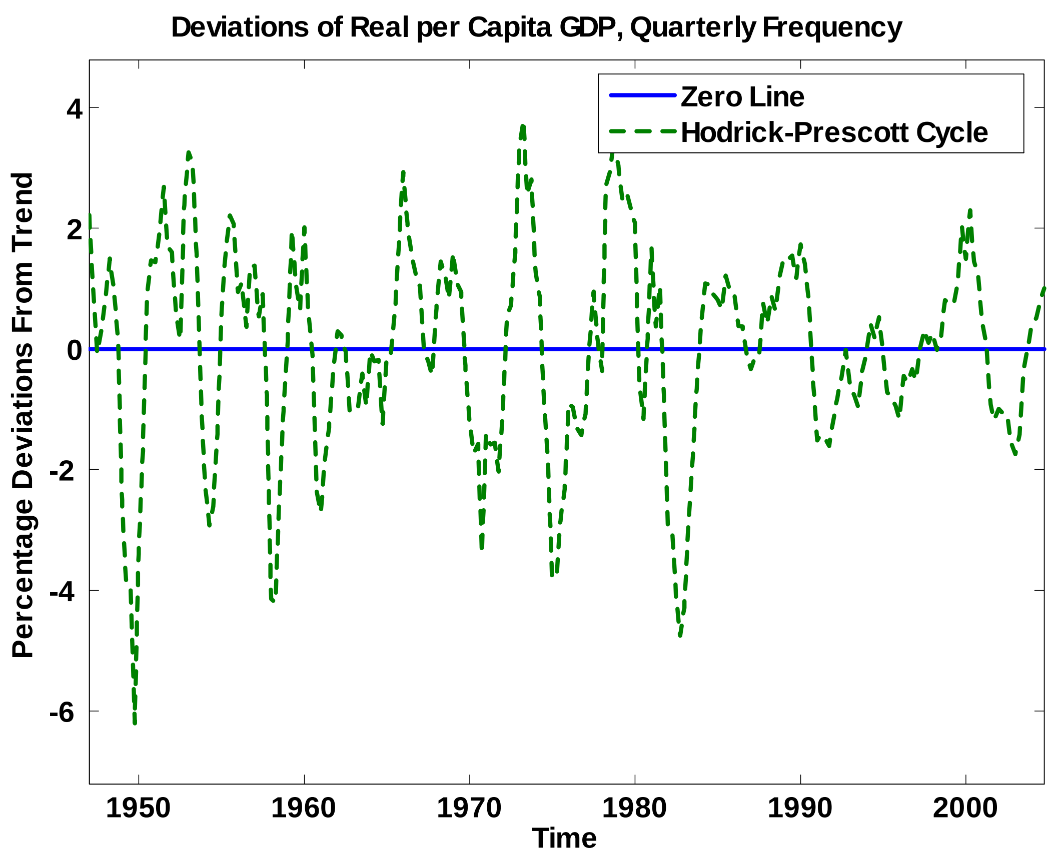

Figure 2.5: Cyclical Component of HP-Filtered Real GDP for the US, 1947-2004

But now we want to proceed with a more systematic collection of business cycle facts. We first focus on our main variable of interest, real GDP. In later chapters, once we enrich our benchmark model with labor supply and other realistic features, we will augment these facts with facts about other variables of interest.

Variable	Mean	St. Dv.	A(1)	A(2)	A(3)	A(4)	A(5)	min	max	%(>0)	%(>0,033)
Growth Rate	3.3%	$2,6\%$	$0,83$	$0,54$	$0,21$	$-0,09$	$-0,20$	$-3,1\%$	$12,6\%$	$88,6\%$	$54,0\%$
HP Filter	0%	1.7%	$0,84$	$0,60$	$0,32$	$0,08$	$-0,10$	$-6,2\%$	$3,8\%$	$53,9\%$	N/A

Table 2.1: Cyclical Behavior of Real GDP, US, 1947-2004

2.2 Basic Facts

Now that we have discussed how to be measure business cycle facts, let us document the main regularities of business cycles. Sometimes the resulting facts are called stylized facts, that is, facts one gets from (sophisticatedly) eyeballing the data. These facts will be the targets of comparison for our quantitative models to be constructed in the next part of this class. The goal of the models is to generate business cycles of realistic magnitude, and to explain what brings them about. In order to do so, we need empirical benchmark facts.

Table 2.1 summarizes the main stylized facts or quarterly real GDP for the US between 1947 and 2004, both when using growth rates and when using the cyclical component of the HP-filtered series. The mean and standard deviation of a time series $\{x_t\}_{t=0}^T$ are defined as

$$
\begin{aligned}
mean(x) &= \frac{1}{T}\sum_{t=0}^{T} x_t \\
std(x) &= \left(\frac{1}{T}\sum_{t=0}^{T}(x_t - mean(x))^2\right)^{\frac{1}{2}}
\end{aligned}
$$

The autocorrelations are defined as follows

$$
A(i) = corr(x_t, x_{t-1}) = \frac{\frac{1}{T-i}\sum_t (x_t - mean(x))(x_{t-i} - mean(x))}{std(x) * std(x)}
$$

and measure how persistent a time series is. For time series with high first order autocorrelation tomorrow's value is likely to be of similar magnitude as today's value. Finally the table gives the maximum and minimum of the time series and the fraction of observations above zero, and, for the growth rate, the fraction of observations bigger than its mean, $3,3\%$.

We make the following observations. First, besides the fact that the mean growth rate is $3,3\%$ whereas the cyclical component of the HP-filtered series has a mean of 0, the main stylized facts derived from taking growth rates and HP-filtering are about the same. They are:

1. Real GDP has a volatility of about 2% around trend, or more concretely, $1,7\%$ when considering the HP-filtered series.

2. The cyclical component of real GDP is highly persistent (that is, positive deviations are followed with high likelihood with positive deviations). The autocorrelation declines with the order and turns negative for the fifth order (that is if real GDP is above trend this quarter, it is more likely to be below than above trend in five quarters from now).

3. Positive deviations from trend are more likely than negative deviations from trend. This suggests that recessions are short but sharp, whereas expansions are long but gradual.

4. It is rare that the growth rate of real GDP actually becomes negative, at least for the US between 1947 and 2004.

It is an instructive exercise (that you will do with Philip's help) to carry out the same empirical exercise described in these notes for an alternative country of your choice. All you need is a sufficiently long time series for real GDP for a country (preferably seasonably adjusted), a little knowledge of MATLAB (or some equivalent software package) and a pre-programed HP filter subroutine (which for MATLAB I will give you). But now we want to start constructing our theoretical model that we will use to explain existence and magnitude of the business cycles documented above.

Part II

The Real Business Cycle (RBC) Model and Its Extensions