Business Cycles: Empirical Evidence[1]

We now examine the empirical evidence on the short run determinates of macroeconomic performance. Our focus is on the effectiveness of monetary and fiscal policies. The econometrics of the short run are more complex than those of the long run.

Some Technical Background

A key tool for this analysis is the *impulse response function.* In these notes, we will first develop IRFs from a theoretical model. We will then discuss IRFs obtained through econometric estimation. After comparing these results, we will begin the process of developing the main business cycle model of the course, the New Keynesian Model. Consider the following a preview. You are not expected to appreciate where these equations come from yet. We will soon develop them in considerable detail.

The model may be written as just two equations. The first is known as *aggregate demand*:

$$\tilde{Y}_t = g_t - b(m-1)(\pi_t - \bar{\pi}) \tag{1}$$

Here, $\tilde{Y}_t$ is the "output gap," the difference between what output is and what it would be if unemployment equaled its natural rate. π_t is inflation. $b(m-1) < 0$, implying an inverse relationship between output and inflation. $\bar{\pi}$ is a target inflation rate set by the Central Bank (2% in the United States).

Note that the aggregate demand curve is downward sloping. Most market demand curves are downward sloping because consumers demand more as the price falls. This intuition does not extend to the aggregate demand curve. Instead, we will see that inflation increases real interest rates. Higher real interest rates then induce households to substitute from consumption to saving, which lowers output in the short run. Again, this is just a preview.

The second equation is *aggregate supply*:

$$\Delta\pi_t = \bar{v}\tilde{Y}_t + u_t \tag{2}$$

Here $\bar{v} > 0$, implying a positive relationship between output and inflation. The model includes two random shocks, one to demand (g_t), and one to supply (u_t). These capture other

[1]These are undergraduate lecture notes. They do not represent academic work. Expect typos, sloppy formatting, and occasional (possibly stupefying) errors.

exogenous factors that may affect aggregate supply or aggregate demand. On average, they equal zero.

Once, again the intuition from micro (where most supply curves are upward sloping) does not extend to this aggregate supply curve. We will see that to make this curve upward sloping, we need to make assumptions that deviate from the baseline supply and demand framework that is taught in basic microeconomics.

Conceptually, suppose that the economy is at its steady state and that all of the shocks (g_t and u_t in our New Keynesian model, for example) have equaled zero forever. Now suppose that a one-time shock hits the economy. After this, no shock ever hits the economy again (the shocks again equal zero).

Impulse Response Functions result from both theoretical and empirical work. The process is similar in both cases.

To calculate an IRF, we do the following:

1. Select numerical values for the model's parameters and insert them into (1)-(2) (for m, b. etc.). This is known as *calibration*. In serious work, it is important to justify a specific calibration. This may be done with econometric estimation or by using microeconomic data to obtain values.[2]

2. Solve for the model's steady state. Here, this consists of setting all shocks equal to their mean of zero and dropping all time subscripts. Note that (2) yields $Y = \tilde{Y} = 0$. Inserting this into (1) then yields $\pi = \bar{\pi}$.

3. For an IRF to a demand shock, set $g_t = 10\%$ (or any other value, it doesn't really matter), and $u_t = 0$. Solve for $\tilde{Y}_t$ and π_t or, more generally, the model's endogenous values.

4. For all periods beyond t, set both shocks equal to zero. Continue to solve for the model's endogenous values.

5. For an IRF to a supply shock, instead set $u_t = 1\%$ and $g_t = 0$ in step #2.

I set $\bar{\pi} = 0.02$, $m = 2$, $\bar{b} = 1$, and $\bar{v} = \frac{1}{3}$. The following charts shows the IRF for a demand shock (possibly to interest rates):

[2] In practice, popular calibrations are often widely cited. The researcher thus does not typically have to build a new calibration from the ground up. Another approach is to simulate the model for a range of calibrations. The latter approach is often taken when a good calibration is difficult to obtain or where different calibrations yield interesting differences.

Figure 1: IRF to Demand Shock in the Simple NK Model

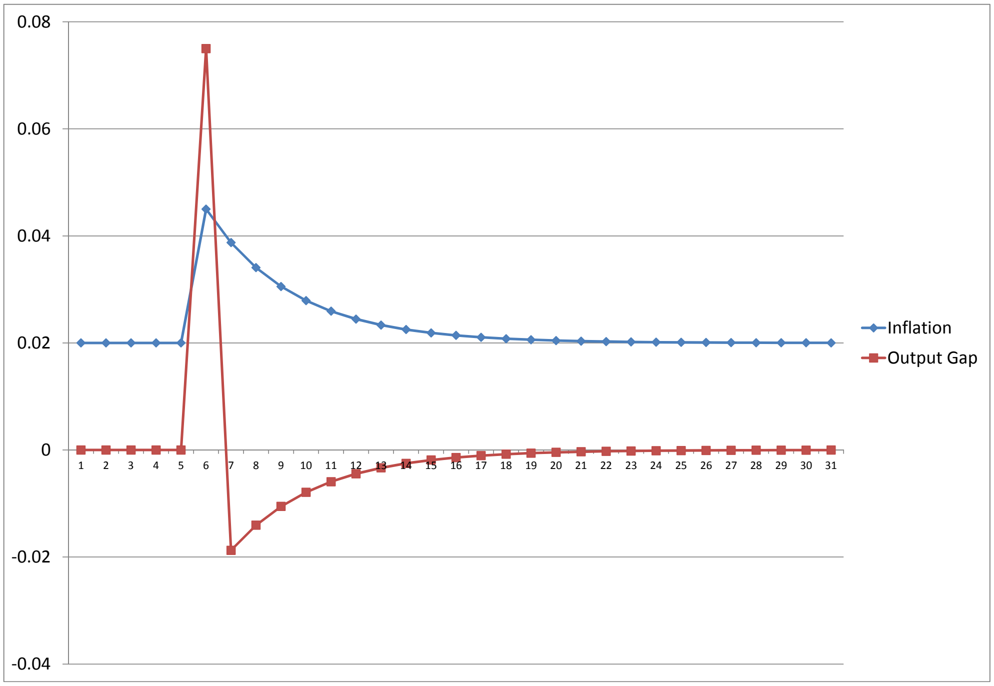

The model thus predicts that a one-time shock to demand increases output temporarily. The output gap then turns negative as it converges toward zero. Inflation jumps up and then converges to its target, $\bar{\pi}$.

Conceptually, the IRF isolates the effect of the shock to aggregate demand. Obviously, such a shock does not happen in isolation in reality. But isolating the effect of a demand shock is important. It helps us understand whether fiscal or monetary polices that change aggregate demand have beneficial effects. This is crucial to evaluating many of the polices which have been pursued in response to the recent recession. Note that the IRF does not just identify the instantaneous effects, but instead identifies the entire time path. This is important because most policies take time to yield large effects.

The next figure shows the IRF for a supply shock:

Here, inflation increases above its target and output falls below potential output.

IRFs can also be estimated from the actual data. The econometrics behind this are beyond the scope of the class.[3] We will, however, compare these estimated IRFs with those from various

[3] They are covered in ECO 341 and widely used in senior theses in macroeconomics.

Figure 2: IRF to Supply Shock in the Simple NK Model

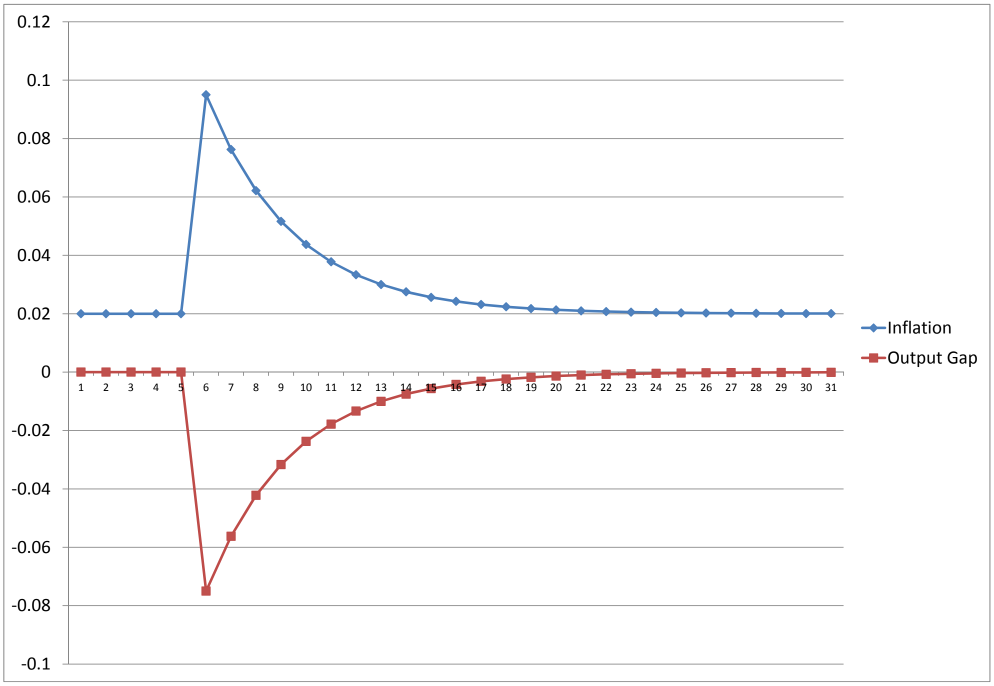

theoretical models. This is another common way to test theoretical models.

Empirical Evidence on Monetary Policy

We begin with the estimation of IRFs for monetary policy. The following IRF is for a random 75 basis point (0.75%) increase to the Federal Funds Rate. It is representative of similar studies:[4]

Note the following:

i. GDP begins to decrease. The peak effect occurs after about 9 quarters where GDP falls by about 1%, a significant effect.

ii. Inflation (the slope of the price impulse response function) picks up after about 3 quarters, and the cumulative decline in the price level is about 3%.

To compare our theoretical and empirical IRFs, we need to mentally flip one of them so that they compare either two increases to demand or decreases to demand. Fortunately, IRFs

[4]Taken from Gail, J. 1992. "How Well Does the IS-ML Model Fit Postwar U.S. data?" *Quarterly Journal of Economics*, vol. 107: 709-738.

Figure 3: Estimated Impulse Response Functions

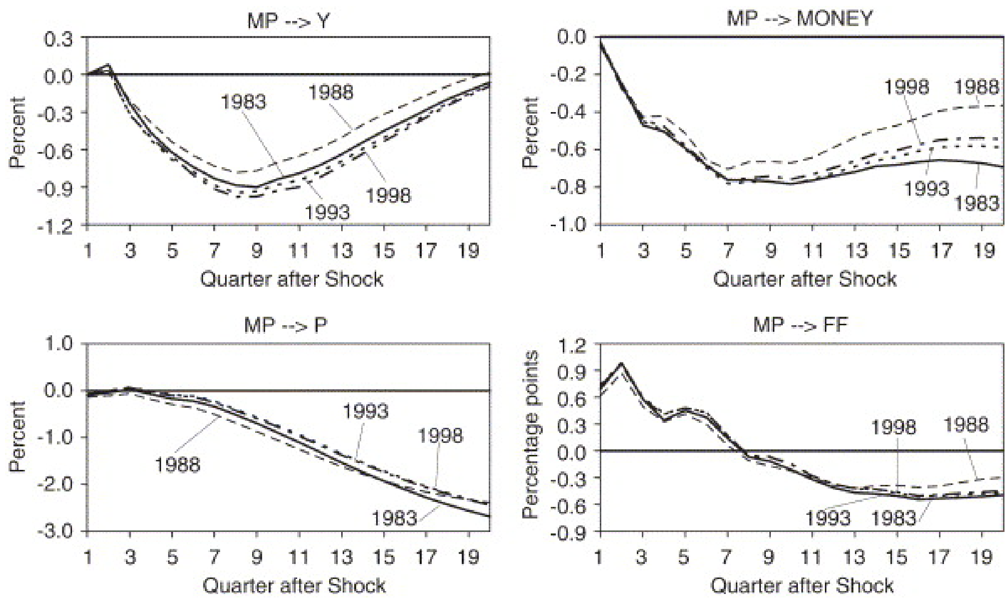

are symmetric in that the response to positive shocks is simply the opposite of the response to the negative shocks. I argue the following:

1. Qualitatively (direction only, not considering magnitudes), the data support the New Keynesian model's prediction that demand shocks affect output. This is in contrast to many classical type models that predict that there is no such effect. This type of result is a major reason why I have chosen the New Keynesian Model as this course's main model of the business cycle.

2. Quantitatively, our New Keynesian model does not do a great job. In particular, the model predicts too little *persistence*. The impact of the demand shock is felt right away and then turns negative. The data, however, show a "hump-shaped" pattern where the effects of the demand shock build up over time and peak after 1-2 years.

Improving the model's empirical fit, using appealing assumptions, is the subject of much recent work in New Keynesian macroeconomics. This is largely beyond the scope of this class.[5]

To conclude the discussion on IRFS, take note that we are comparing two different types. One comes from our theoretical model that we will soon develop. The second is taken from econometrics. The idea going forward is that we compare the two of them to help judge the empirical relevance of the theoretical model.

[5]ECO 318 spends some time on this issue, but it is not a major emphasis of that course.

Table 1: Empirical Estimates of the Government Spending Multiplier

Source	Estimate
Evans (1969)	> 2
Barro (1981), Hall (1986, 2009), Barro and Redlick (2011)	0.6-1
Rotemberg and Woodford (1992)	1.25
Ramey and Shapiro (1998), Edelberg, Eichenbaum, and Fisher (1999), Eichenbaum and Fisher (2005), Cavallo (2005)	0.6-1.2
Blanchard and Perotti (2002)	0.9-1.29
Mountford and Uhling (2009)	0.65
Romer and Bernstein (2009)	1.57
Cogan, Cwik, Taylor, and Wieland (2010)	0.64
Ramey (2011)	0.6-1.2
Fisher and Peters (2010)	1.5
Auerbach and Gorodnichenko (2011)	-0.3-3.6
Gordon and Krenn (2010)	1.8

Empirical Evidence on Fiscal Policy

The following table shows various estimates for the government spending multiplier.[6]

Clearly, there are a wide range of estimates on the government spending multiplier. There are two sources of these large differences.

i. Technical differences including differences in data and differences in econometric techniques.

ii. The size of the government spending multiplier surely depends on the underlying macroeconomic conditions. Ilzetzki and Mendoza (2012) examine these differences.[7]. They find:

a. In developed economies, the multiplier is positive both on impact and at peak. In developing economies, however, the multiplier is zero or negative.

b. Economies with flexible exchange rates (like the United States) have lower multipliers than countries with fixed exchange rates.

c. Countries that are open to trade have lower, often negative, multipliers.

d. Countries with debt to GDP ratios above 60% have negative multipliers.

[6]Complied by: Ramey, V. 2011. "Can Government Purchases Stimulate the Economy?" *Journal of Economic literature*, 49(3): 673-685. Ramey (2011) also provides more more detailed references for each study.

[7]See Ilzetski, E. and R. Mendoza. 2012. "How Big (Small?) are Fiscal Multipliers?" *Working Paper*

Other work generally finds that fiscal policy is more effective when the economy is in a recession as opposed to an expansion.[8]

Collectively, all these results imply that we don't really know what effect short run fiscal stimulus has on GDP and unemployment. This is currently the source of much ongoing research, including by a colleague here at Bates.

Chahrour, Schmitt-Grohe, and Uribe (2010) examine the tax multiplier.[9] They find that different studies show that reducing taxes by 1% of GDP increases GDP by between 1% and 3%. This does not imply that reducing marginal income tax rates by 1% will yield this effect. Rather the share of taxes to GDP (about $\frac{1}{4}$ for all levels of government in the U.S.) must fall by 1%

Laffer Curves

Consider the following argument: There are always two tax rates that yield the same level of revenue. This is not controversial. Suppose for example, that the labor tax rate is set to zero. It follows that tax revenue will equal zero. Now suppose that this tax rate equals 100%. Such a tax rate will also yield no tax revenue because rational agents will choose to supply no labor at this rate.[10] This argument yields a Laffer Curve, a relationship between tax rates and tax revenues:

[8] See, for example, Auerbach and Gorodnichenko (2011).

[9] See. Charour, R., Schmitt-Grohe, S., and M. Uribe. 2010. "A Model-Based Evaluation of the Debate on the Size of the Tax Multiplier." *NBER Working Paper 16169*.

[10] Or more likely, they will report no income at this rate.

Graph: Laffer Curve

There are a pair of important implications from the Laffer Curve:

i. If the economy lies to the right of the peak of the Laffer Curve, then lowering taxes will raise tax revenue.

ii. If the economy is to the left of the peak of the Laffer Curve, then cutting taxes by 1% (relative to their previous level) will reduce tax revenue by less than 1%. The difference is known as the self-financing component of tax cuts.

Neither of these results are controversial. The tax rates that maximize tax revenue and the self-financing rate of tax cuts, however, are. Trabandt and Uhlig (2011) estimate these values.[11] They find the following:

i. For the United States, the labor tax rate that maximizes tax revenue is between 52% and 72%, well below even the highest current marginal tax rate. For taxes on capital, this figure is between 60% and 64%.

ii. The self-financing component of a labor tax cut in the U.S. is between 21% and 49%. For a capital tax cut, it is between 45% and 60%.

iii. Because Europe generally has higher tax rates than the U.S., labor tax cuts are more self financing. The authors estimate 35% for Ireland, 50% for Germany, 62% for France, and 83% for Sweden.

[11] Trabandt, M. and H. Uhlig. 2011. "The Laffer Curve Revisited." *Journal of Monetary Economics*, 58(4): 305-327.

Keep in mind that it is not a goal of policy makers to maximize tax revenue. These results, along with related work, shows that there is little doubt that tax cuts in the United States do not pay for themselves.

The Laffer Curve is also an important part of the 1980s classic *Ferris Bueller's Day Off*. If you have not seen this film, do so immediately. It WILL be on the final exam.