EC1010 Review

Paul Scanlon

February, 2010

Contents

Chapter 1

The Very Basics

"A large income is the best recipe for happiness I ever heard of."

- Jane Austen.

1.1 Measuring Output

Definition 1 (Nominal *GDP*) *Nominal GDP is the market value of final goods and services currently produced in the country in a given time period.*

Because its rise could be solely attributable to rising *prices*, nominal GDP doesn't really tell us much about welfare. For instance, a doubling of nominal GDP could be due to a doubling of *prices*. Yet we're primarily concerned with changes in *output*, not *prices*. Changes in output are reflected in changes in *real GDP*.

Definition 2 (Real GDP) *Real GDP measures the value of an economy's final production in a given period using the prices of some base year. As a result, changes in real GDP measure changes in quantities of output alone, and filters out the effects of price changes. Crucially, we measure real GDP using the same set of prices each year (i.e., using the prices from the "base year.")*

In other words, changes in real GDP measure changes in *physical* output produced within a country's borders. Our real concern, however, is not the level of real GDP, which measures aggregate output. Rather, we're concerned with real GDP per person or what's known as real GDP *per capita*; this is our measure of a country's *standard of living.* Because it is highly correlated with all other measures of well-being such as nutrition and life expectancy, real GDP per capita is a fairly reliable measure of welfare. That said, it says little about the *distribution* of wealth. Take, for instance, Bill Gates. If he walks into lecture, $\frac{Y}{L}$ will assuredly increase a million-fold. Yet this is hardly reflective of average welfare in the room. Moreover, GDP doesn't measure such factors as leisure time or environmental degradation.

Now picture this in your mind: a simple economy produces two goods in 2001 at \$10, so *nominal GDP* is \$20. The price of both goods increase to \$15 in 2002, so *nominal GDP* in 2002 is \$30. To evaluate *real GDP* we must use the *same* price level each year. So using 2001 prices (say) as our *base* year, we see that there has been *no change* in real *GDP* in 2002—it's still \$20. We have two goods each period, nothing "real" changes. Using prices from 2002 gives different values for real GDP. Yet the proportionate changes in real GDP—what we're frequently interested in—stay the same. Tables 1.1 and 1.2 give another example; for a one-good economy, Table 1.1 shows nominal GDP, while Table 1.2 presents real GDP in 1999 prices. When we compare GDPs across countries, we try to use the same prices across countries too. When we do this, such measurements are called PPP-adjusted GDPs.

Table 1.1: GDP at Current Prices

Year	***Quantity***	***Price***	***Nominal GDP***
1999	5	2	10
2000	5	4	20
2001	7	5	35

Table 1.2: Real GDP in 1999 Prices

Year	***Quantity***	***Price***	***Real GDP***
1999	5	**2**	10
2000	5	**2**	10
2001	7	**2**	14

1.2 Measurement

Moving on, we ahave the national accounts identity:

$$\boxed{\text{Income=Expenditure=Output}} \tag{1.1}$$

Imagine a small economy produces 20 cars a year, worth $20,000. So the value of *output* is $20,000. All those cars are sold, so *expenditure* is also $20,000. And all the producers receive that money, so *income* is $20,000.

All forms of expenditure end up as someone's income (so income=expenditure), so that's the first part. Furthermore, output produced is ultimately purchased and thus is equal to expenditure (so output=expenditure). In case you're wondering, what happens if some of the output is not purchased? In this case, the output is *still* counted as expenditure. In the national accounts that's treated as a purchase by the producer himself for his inventories (and this appears in *Investment*, as *inventory investment*). Using this convention, output is indeed equal to expenditure.

As another example, consider this: a carpenter buys wood from a forester for 20, and uses the wood to make a bookshelf, which he in turn sells for 50. There are three ways we can calculate (nominal) GDP. First, we can use the *expenditure approach*: this measures the amount of expenditure on final goods and services. In this simple example, this is simply the amount paid for the table i.e., 50. Second, we can use the *income approach*: this simply adds up the values of all incomes in the economy. In this example, the forester receives income of 20 and the carpenter receives (net) income of $50 - 20 = 30$. Hence the sum of incomes in the economy is 50. Third and finally, we can add up the *value added* by each worker in the economy: that is, what is the value each worker's contribution to the value of what they sell? In this example, the value added by the carpenter would be 20, while the value added by the carpenter would be 30. Again, the sum is 50.

Because it is the most common way to measure GDP, let's talk more about the *expenditure approach.* This entails adding up the values of final expenditure in the economy. Most importantly, we do not include *intermediate goods.* Adding up the components of final expenditure gives a measure of GDP, Y:

$$\boxed{Y = C + I + G + X - M} \tag{1.2}$$

Given that imports represent *foreign* production, we must subtract its value from domestic expenditure. This leaves us with the value of domestic expenditure that's domestically produced; i.e., GDP.

Consumption (C) C includes durable goods (e.g., cars, tvs), nondurable goods (e.g., food, clothing) and services, which are the largest component (items like education, healthcare, haircuts). It only includes the purchase of new goods, not the transfer of second-hand goods (this goes for all spending components. It's the largest component of GDP, currently constituting around 70% of GDP in the US.

Investment (I) This is the sum of nonresidential private investments (expenditures by firms on equipment and structures), residential investment (new homes) and changes in inventories (inventory investment). If firms place goods in inventories, this is regarded as a purchase by the firm from itself.

Government Spending (G) Sum of all purchases of goods and services by the government. It does not include government transfers (we are seeking to measure the value of *production.*) Investment by the government is part of G.[1]

Exports (X) Sale of domestically produced goods and services to the rest of the world.

Imports (M) The purchase of goods and services from the rest of the world.

Some Examples

1. An American firm, Boeing, sells a 747 to the American government - *this is part of G.*
2. Boeing sells a 747 to an American airline, US Airways - *this is part of Investment.*
3. Boeing sells a 747 to an American citizen - *part of Consumption.*
4. Boeing sells a 747 to an Ryanair - *in the U.S. accounts, this is an Export.*
5. Boeing builds a 747 to be sold next year - *this is part of inventory investment, I.*
6. A firm buys a new robot helping to automate production -*regarded as Investment.*
7. The Irish government purchases a foreign painting for an Irish museum - *included in G and M, so ultimately makes no difference to GDP (since it enters positively as G and negatively as M), as we'd predict. If a consumer purchased the painting it'd be recorded in C and M.*
8. The wage of an Irish immigration official - *this is part of G, since it's a service purchased by the Irish government.*

[1] In economic models—such as in the Solow model—it is conventional to assume G refers to government consumption, not investment.

1.3 Other Definitions

Definition 3 (Current Account) *This is the value of exports less the value of imports. It is also referred to as net exports.*

Definition 4 (Gross National Product, GNP) *Nominal GNP measures the value of output accruing to domestic factors of production. It is the value of output produced by citizens of a country regardless of where they are. (By contrast, GDP refers to output produced by people within a country's borders, regardless of their nationality.)*

Consider: if you travel to the U.S., and work there for the summer, then the value of output you produce there is part of U.S. GDP *and* Irish GNP. Because no production has taken place within Irish borders, this has no effect on Irish GDP.

Definition 5 (Potential Output) *Potential output, Y_n, is the level of output when the economy is working at a normal pace and is neither in boom nor recession.*

Definition 6 (Output Gap) *This is the difference between actual and potential output, $Y - Y_n$. So when the economy is in a boom the output gap is positive; and when it's in recession, the output gap is negative.*

1.4 The Price Level and Inflation

Definition 7 (Consumer Price Index (CPI)) *The CPI is a price index that measures the price of an average basket of goods in an economy. It is a weighted average of all the prices of consumer goods in the economy, weighted by the consumption shares of each good. One common measure of inflation is is the annual percentage change in the CPI.*

Definition 8 (Deflation) *This is defined as a sustained fall in the level of prices in an economy.*

Don't confuse *disinflation* and *deflation.* Deflation is when prices actually *fall.* Disinflation is a reduction in the rate of inflation, say from 15% to 10%. So, under *disinflation*, we can typically have prices still *rising.* In contrast, hyperinflation refers to extremely high rates of inflation. Think of it this way: hyperinflation is sprinting, inflation is running, disinflation is running at a *slower pace*, while deflation is going backwards.

Key Idea 1 (Policy Goals) *Good macroeconomic policy aims for low inflation, a zero output gap, and economic stability (as opposed to volatility). Demand side policies are aimed to increase aggregate demand, $Y = C + I + G + NX$, so as to bring the economy close to potential. Supply side policies (such as deregulation) are designed to increase potential output; i.e., to enhance the productive capacity of the country.*

1.4.1 Real Values

In economics, we are typically interested in the *real value* of something, not the nominal or monetary value. What we mean by the real value is the *purchasing power* value: how many *goods* can we buy with a given amount of money? For example, one can think of the real wage as the number of goods you can buy with that wage. Except in the case of interest rates, to obtain the real value we simply deflate by a price level such as the CPI. (i.e., the price of a normal basket of consumer goods). For instance, if you owe $1000 and the price level is $100, then your real debt burden is ten basket of goods. If the price level was raised to $1000, however, the real burden is now a mere one basket of goods. *And all that matters is the real burden in terms of purchasing power.* As another example, suppose I took a $10 loan from you fifty years ago. At that stage the purchasing power of the loan was relatively large, but given the substantial rise in the price level in the interim, the purchasing power of what I give back to you is pretty low, so I gain at your expense.

Because macroeconomics often entails analyzing growth rates, the following rule of thumb is useful:

Definition 9 (Rule of 70) *If a variable is growing at rate g% per year, then that variable will double in size after approximately $\frac{70}{g}$ years.*

Chapter 2

The Solow Model

"Is there some action a government of India could take that would lead the Indian economy to grow like Indonesia's or Egypt's? If so, what exactly? If not, what is it about "the nature of India" that makes it so? The consequences for human welfare involved in questions like these are simply staggering: Once one starts to think about them, it is hard to think about anything else."

- Robert Lucas, Nobel Laureate (1995).

2.1 Economic Growth

One of the most striking features of the world economy is the vast disparity in standards of living and rates of economic growth. For example, in 2000, real GDP per capita in the United States was more than fifty times that in Ethiopia. And over the period $1975-2003$, real GDP per capita in China grew at a rate of 7.6% annually, while, in Argentina, real GDP per capita grew at a rate of only .1%—seventy six times slower. Moreover, there are often vast reversals in prosperity over time. Argentina, Venezuela, Uruguay, Israel, and

South Africa were in the top 25 countries (by GDP per capita) in 1960, but none made it to the top 25 in 2000. From 1960 to 2000, the fastest growing country in the world was Taiwan, which grew at 6%. The slowest growing country was Zambia which grew at −1.8%. That is, people in Zambia were markedly worse off in 2000 than they were in 1960. The theory of economic growth seeks to address these issues and provide explanations.

The Solow model is a long-run model that seeks to explain why there are such vast income disparities and growth differences across the world. It describes the evolution of *potential* output or the productive capacities of countries over time. Because we assume an economy is always at potential in the long run, there are no recessions or booms in this analysis. Furthermore, there is no mention of nominal quantities such as money and prices since the *classical dichotomy* holds; namely, money has no effect on output in the long run and is thus irrelevant for explaining output differences. Over the long run, printing pieces of paper cannot generate increases in prosperity.

Before going on, it is important to note that we are not really interested in aggregate income/output, Y, in this world. What we are really concerned with is income *per capita*, $\frac{Y}{L}$—our conventional measure of *standard of living*. This shows how well each of us, on average, is doing.[1] Most importantly, I assume the population of size L is constant; that is, there is no population growth. As a result, when Y changes, $\frac{Y}{L}$ will also rise; and when K rises, $\frac{K}{L}$ rises.

The Aggregate Production Function

The *production function* for the economy is:

$$Y = AK^{\alpha}L^{1-\alpha}, \qquad 0 < \alpha < 1. \tag{2.1}$$

The production function is our *first key equation.* An important feature of this function is that it exhibits *decreasing returns to scale* to capital. What this means is that a second

[1] In the United States, and indeed most industrialized Western countries, this has been growing at an average rate of 1.8% since 1870 or about 2% since 1900.

laptop, say, will not give as much *bang per buck* as the first laptop, and so on. Formally, we say there's *diminishing marginal productivity* to capital accumulation. Keep in mind that we assume all the units of capital (say, laptops) are *the same.* So what I mean by increases in the capital stock are *more of the same* units.

Examining the production function more carefully, we see immediately that differences in GDP are attributable to differences in total factor productivity, A, differences in physical capital, K, and differences in population, L. The A term denotes *total factor productivity*; it is anything (in a broad sense) that, for a given K and L, leads to greater output. It's not simply technological advances, but also encompasses such factors as culture, climate, skills, health, work ethic, social capital (i.e., culture), human capital (education), institutions (the political system etc), and so on. Basically, A measures the efficiency with which K and L are combined. Total factor productivity rises if output per worker increases for any *given* K and L. It could be just people working more hours. For example, in a country with a persistently hot and inhospitable climate, then Y will likely be lower for a given level of capital per worker; hence A is relatively lower. Just imagine Trinity on a great day—no matter how great your PC is, you'll still get less done; that is, A will fall. Yet, for now, I assume A is a constant at any point in time.

To graph this, just let $\beta = AL^{1-\alpha}$ for a moment, and so the production function becomes

$$Y = \beta K^{\alpha}, \qquad 0 < \alpha < 1. \tag{2.2}$$

Using this formulation, we can treat β as a constant and just draw a graph of Y against K. Figure 1 shows the result.

The slope of the function subsides as we move outwards. As we shall see, this is what really frustrates things in this economy in the long run. Note crucially that an increase in A will raise Y for any given level of K. Graphically, the way to represent a rise in A is to shift the production function upwards. Figure 2 shows the result. For example, more *powerful* computers would be best represented as an increase in technology, A: just think

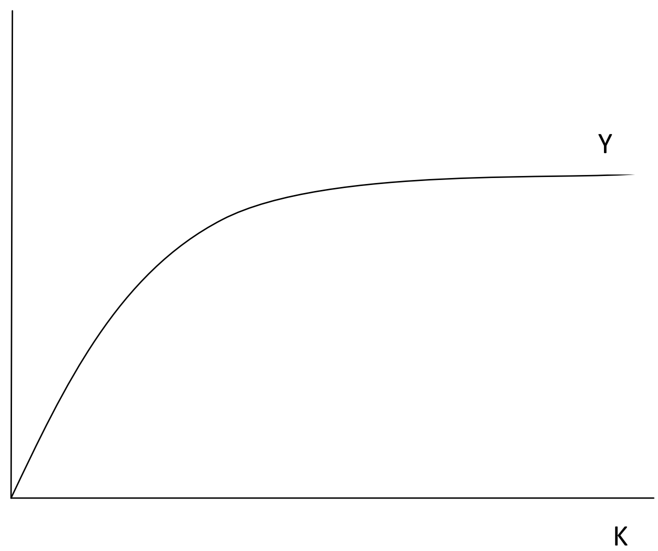

Figure 2.1: Basic Production Function. The first laptop is extremely useful; the second less so; the third hardly any use at all; and so on. Formally, the marginal product of capital (i.e., its usefulness), falls as K increases. This is called the diminishing marginal product of capital.

of each existing computer getting more power via new software.

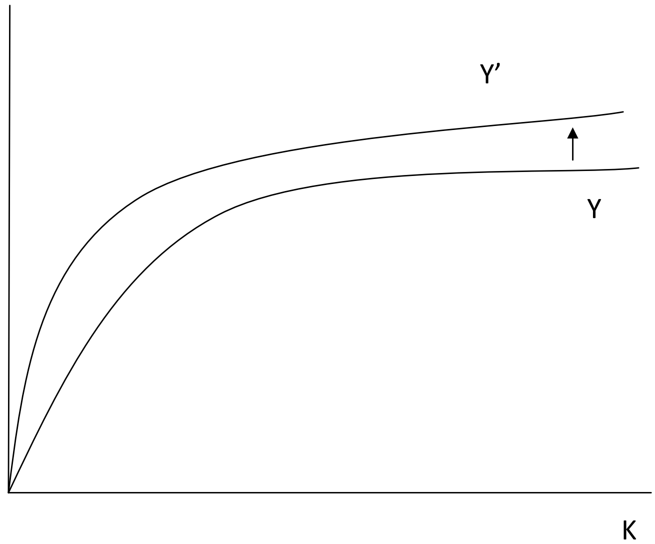

Figure 2.2: An Increase in A shifts the production function upwards. Such increases in A are a way of overcoming the curse of diminishing returns.

Growth in A

Because L is fixed, there are only two factors that increase Y, and hence the standard of living, $\frac{Y}{L}$: capital K and total factor productivity A. Because the marginal product of capital tends to zero, increments in capital are little use to us after a certain point. Therefore, in the long run, moving outwards to the right becomes futile, yielding only negligible increases in Y. For this reason, we need something else to generate *sustained* growth of Y and that something else (by a process of elimination) is productivity growth; i.e., growth in A. Without growth in A, we see that growth eventually stalls.[2] If we can't grow by moving out to the right (that is, by capital accumulation), then *the only way is up*; in the long run, the productivity factor, A, *must* grow to shift our function upwards and generate sustained growth. And that's basically what's happening in the industrialized economies.

[2]In reality, we *have* sustained growth in Y and $\frac{Y}{L}$; in developed economies such as the United States, for example, $\frac{Y}{L}$ increased tenfold in the U.S. from 1870 to 2000.

That's the only way we can have sustained rises in income per capita or standard of living. However, notice that capital accumulation is important at the *outset* of development. But once developed (i.e., when diminishing returns are prevalent), capital accumulation is less useful, and now it's productivity that drives sustained growth. Also, capital accumulation is the "easy" part; it's much harder to generate increases in A. Inspiration, not perspiration, is the key to sustained growth. Many countries such as Russia were remarkably successful in accumulating capital (through forced savings, generating so-called "Stalinist growth") and grew rapidly for a while, but couldn't subsequently generate increases in A.

Capital Accumulation

Central to the model is the evolution of the capital stock; once we know the level of capital, we can find every other variable. Thus, our *second key equation* describes the evolution of the capital stock or *capital stock accumulation.*

First, we assume the stock of *savings* in the economy are some constant fraction, $0 < s < 1$ of output Y. This leaves people with $(1-s)Y$ to consume. Savings includes private savings by individuals and public savings by the government; this sum is called *national savings.* For instance, other things constant, government deficits reduce national savings and hence the savings rate. In the model, savings are channeled into investment; i.e., to buy more laptops etc.[3] As a simple example of the saving/investment decision, consider corn: each period you can either eat the corn (i.e., consumption), or place the corn in the ground, and new corn will grow next period; i.e., investment. It is a similar process here. Now, denoting investment by I and aggregate savings by S, we have:

$$S = sY = I$$

where s denotes the national savings rate. That is, savings are used for investment. Figure 3 shows the *savings line.*

[3]Note the importance of *financial intermediation* or banks here. It's crucial to have a credible banking system so as to people can save and so as to channel those savings to the right investment projects.

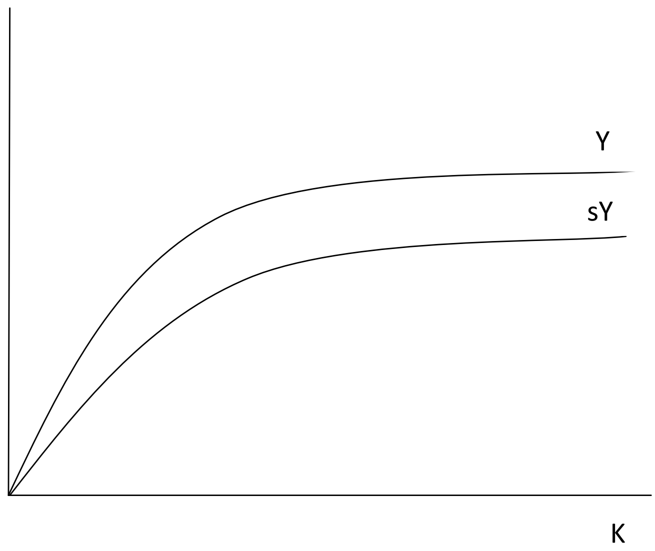

Figure 2.3: Savings is some fraction s of output. At each level of K, the savings line, sY, indicates what <u>adds to</u> the capital stock.

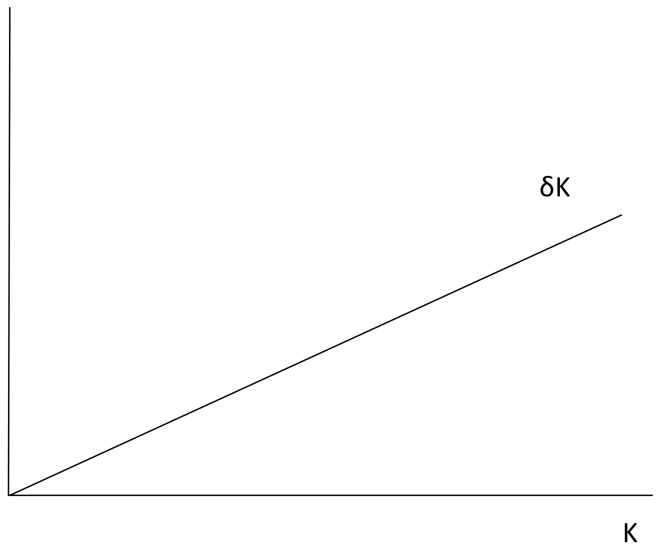

Figure 2.4: At each level of K, the depreciation line indicates how much capital depreciates in the economy. In other words, this tells us what <u>leaves</u> the capital stock.

Therefore, by raising investment, savings leads to an increase in the capital stock. But unfortunately a fraction $0 < \delta < 1$ of that stock depreciates (or *melts*) each period. Figure 4 shows the *depreciation line.* Overall, the equation for the evolution of the aggregate capital stock, K, is:

$$\Delta K = I - D = sY - \delta K \tag{2.3}$$

where D denotes depreciation, assumed to be some fraction δ of the capital stock, K. So the net increase in the capital stock is investment minus depreciation. Think of it this way: what determines the increase in *you* is the amount of calories you consume, minus the number of calories you burn up. It is precisely the same idea here.[4]

At the heart of the model is the tension between investment and depreciation. It is useful to think of these as a tug of war between the "good" and "bad" force. The tension between total depreciation and investment determines the evolution of the capital stock and ultimately the standard of living. But there is bad news. Depreciation is linear in capital. By contrast, the gains to saving are high initially, but fall as diminishing returns to capital set in. As a result, we must reach a point where the depreciation force just offsets the gains to investment and, assuming A is fixed, there is no further growth K or Y. This point is called the *steady state* level of capital, K^*. Associated with this level of capital, is the steady state level of output, Y^*.

The Transitional Path to Steady State

Figure 5 shows the complete Solow model. The steady state capital stock is at the intersection of the savings and depreciation lines. When the economy is the the left of the steady state, investment exceeds depreciation. As a result, the capital stock K must increase. And given that K is increasing as it moves towards steady state, we know from the production function, $Y = AK^{\alpha}L^{1-\alpha}$, that Y is necessarily increasing too. When the

[4]People attempting to lose weight decrease their food intake (i.e., their "investment") and engage in exercise (i.e., trying to create more "depreciation.")

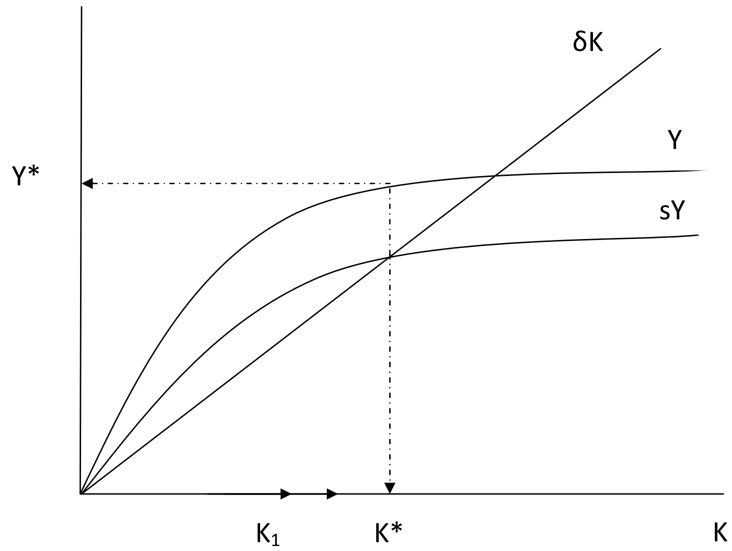

Figure 2.5: The Complete Solow Model. Once the savings/investment function is above the depreciation line (as at point K_1), the capital stock increases. In contrast, when the depreciation line is above the savings/investment function, the capital stock falls. Point is, the economy will always gravitate to steady state. At the steady state, the two forces just offset each other and the capital stock stalls. It can take an economy at least a decade to go from K_1 to K^*.

savings/investment function and the depreciation line intersect, then total depreciation and investment just offset each other. Thereafter, absent growth in A, there is no further growth in K or Y. The economy is now at *steady state* and in its *golden years.* At this point, growth in A must drive sustained rises in the standard of living.[5]

To summarize, we have two key equations that describe the evolution of the economy. They are:

$$\boxed{Y = AK^{\alpha}L^{1-\alpha}} \tag{2.4}$$

and

$$\boxed{\Delta K = sY - \delta K} \tag{2.5}$$

So when the savings line is above the depreciation line, mathematically we have, $sY > \delta K$, and hence from Equation (2.5), $\Delta K > 0$. Then from Equation (2.4), output will also increase. At the steady state level of capital $\Delta K = 0$, and hence from Equation (2.4), output, Y, will also stop growing.

Key Idea 2 *If the savings/investment function is above the depreciation line, then the capital stock and output are rising. The opposite applies if the savings/investment function is below the depreciation line.*

The Steady State

It's important to bear in mind that steady state Y and hence $\frac{Y}{L}$ (i.e., welfare/standard of living) is constant and determined by the exogenous parameters of the model: s, δ, A. While increases in s and A lead to higher standard of living, increases in δ lowers it. Increases in s and A ultimately result in more savings and investment, and so lead to a higher steady state capital stock.[6]On the other hand, a higher rate of depreciation causes a greater losses in capital stock, which ultimately lowers the steady state level of capital. Because countries across the world have different values of s, δ, and A, they

[5]You could continue to increase the savings rate, but the resulting increases in output would be tiny.

[6]Note that a higher A leads to a higher $Y = AK^{\alpha}L^{1-\alpha}$ and so to higher stock of savings, sY.

converge to different steady states. Also, just because an economy is in equilibrium, it does not mean the equilibrium is good; rather, it could settle at a point of impoverishment. Once a country reaches its steady state, we have no further growth in K or Y, without growth in A. So, without A growing, we have no *sustained* growth in output per worker in steady state. All of the growth we see in the model is thus *transitional* and occurs on the *transition* to steady state. To ensure continually rising living standards, A must rise. Such continuous rises in A are attributable to technological progress. While a better political system would lead to a once-off increase in A, it is unlikely that it would lead to a continuous rise in A over time. For this reason, it is technological progress that ultimately drives growth in developed economies. And because technological advances can be "shared" across advanced countries, the model predicts advanced economics should grow at about the same rates—a prediction that is confirmed in the data.

Key Idea 3 *When an economy reaches steady state (or its golden years), there is no further growth in Y or $\frac{Y}{L}$ (without growth in A), since the forces of depreciation just offset the forces of investment. What determines the steady state level of Y are the "fundamentals": s, δ, and A. To generate sustained growth in Y, A must rise continuously.*

Convergence

Each country grows towards its own steady state determined by its fundamentals, A, s, and δ. So clearly, countries with *different* parameters converge to *different* steady states. From the basic Solow diagram in Figure 5, we can see that countries further below their steady states grow faster since $\Delta K > 0$. The economy, in this sense, is analogous to a *spring*—the further the economy is "stretched" from its steady state point, the faster it will grow. Fundamentally, the reason for this is that there countries have not yet encountered severe diminishing returns. For them, the *marginal product of capital* is relatively high, leading to large increases in output.[7] Think of *running* from Pearce St. to the Arts Block;

[7] The marginal product of capital is the increase in output from adding one more unit of capital.

you start off pretty quickly, but this is unsustainable, and your speed slows down as you approach the Arts Building. It's the same intuition for the growth of nations towards their respective steady states.

But wait. What about Africa, is this continent now poised for rapid growth? Not at all. Africa would *not* be expected to grow fast under this rationale, since it has a low steady state due to poor fundamentals (mainly, low A and low s). It's only countries with good fundamentals that are *starting off* poor that would be expected to grow faster. What determines growth is the distance between where they are now and their steady state (i.e., where they are going.) This is called *conditional convergence* or *catch-up growth.* Examples of such rapid growth include China and the other Asian "tiger" economies like Taiwan and South Korea. Think of it like this: if your micro result was pretty poor, but you now start working harder (and so improving your "fundamentals") you have the potential from rapid improvement for the macro test. By contrast, those who performed well in micro have little potential for growth in their mark. More significantly, we don't expect any improvement from those who performed poorly in micro, and take no measures to improve.

Key Idea 4 *The further the economy is from its steady state, the faster it will grow.*

Definition 10 (Conditional Convergence) *Poor countries that are heading to a good steady state can experience rapid growth. Namely, such countries have not yet hit severe diminishing returns, and so output grows quickly as they head towards their steady state.*

2.1.1 Savings and Development

One central insight from the Solow model is its prediction for what happens when the savings rate changes. So suppose the savings rate, s, changes due to a rise in personal savings. The investment/saving function shifts up, and we initially grow towards a new high steady state. As we move towards the new steady state, K and hence Y rise. In the new steady state K and Y are both higher and remain constant at a higher rate.

Significantly, there is no effect on the long-run growth of K and Y. It was only along the *transition path* to steady state that we had a spurt of growth. But that growth effect is ultimately transitory and leads us merely to a higher *level* of Y in the long run. What is striking is we don't get a permanently higher growth rate of Y. Underlying this important result is the principle of *diminishing returns*. More savings just means more investment, but this is not of *that* much use to us. Because all the antecedents to the Solow model predicted that the long-run growth of Y was *increasing* in the savings rate, Solow himself described this as a "real shocker" Finally, keep in mind that the time to complete the transitional path to a new steady state is long in this model—probably decades—so we can actually have *transitional growth effects* for quite a while. Finally, empirical work shows that the elasticity of output with respect to the saving rate is about .5. That is, a 10% increase in the saving rate raises the level of output per person in the long run by about 5%, only a modest increase.

To understand these effects, consider body weight. One often reads something to the effect that eating an extra bar of chocolate a day will cause you to put on, say, 2 pounds a week. Just think about that for a moment. After a year you'll put on 104 pounds, and after two years you'll put on 208 pounds, and so on. Fortunately, this is self-evidently *wrong*; when you increase your consumption rate of chocolate, there cannot realistically be a *long-run growth effect*. Most likely, you will grow for a while in size—i.e., a *transitory growth effect*, to be sure— but you will asymptote eventually to some greater weight/*level*. This prediction seems more reasonable. In other words, there's a transitory growth effect and long-run effect on your weight level. In particular, you wont keep growing *forever*. Of course, this is analogous to what happens when the savings rate increases in the Solow model.

2.1.2 Examples

Getting back to macro, here are three examples of changes to the model. In all cases, there are three mains steps to figuring out how Y (and hence $\frac{Y}{L}$) changes. First, you figure out which curve shifts. Second, figure out where the new intersection of the savings and depreciation line is. This will nail down the new steady state capital stock. Third, from the new steady state capital stock, go up to the production function to find the new steady state level of Y.

Example 1: An Fall in the Savings Rate

Suppose we are at the steady state K^* and there is a sudden and permanent fall in the savings rate from s to s' time t. This is shown in Figure 6. This change causes the savings line to shift downwards. We start off at K^*. Just after the change, then, we now have the depreciation line above the new savings line. As a result, the capital stock falls, and will continue to fall until the depreciation line intersections with the new savings line. This point of intersection is the new steady state K', where the associated output level is Y'.

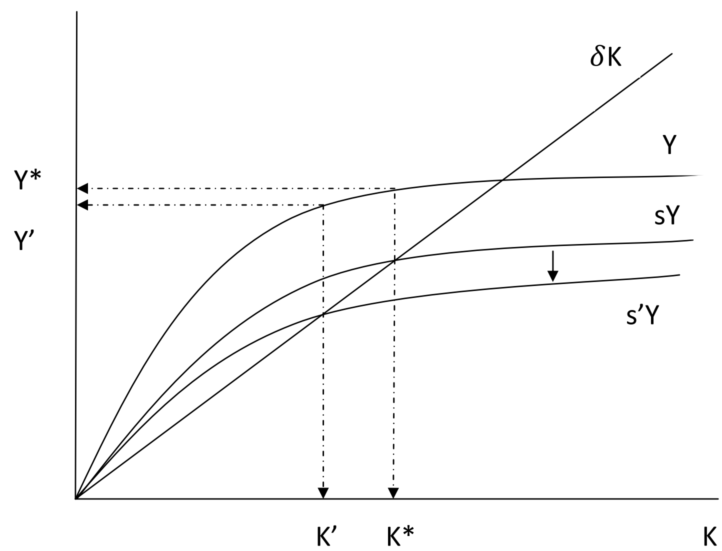

Figure 2.6: A fall in the savings rate from s to s' in the Solow Model.

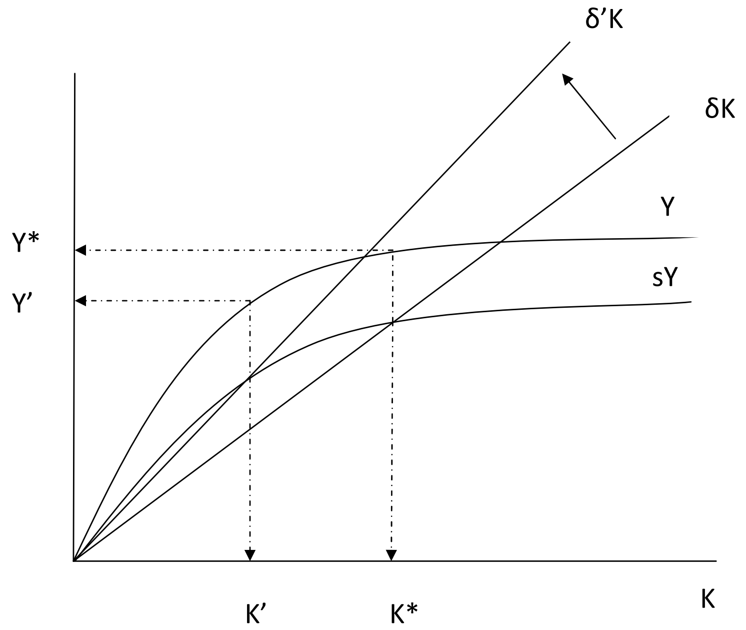

Figure 2.7: An Increase in Depreciation from δ to δ' in the Solow Model.

Example 2: A Rise in the Depreciation Rate

Figure 7 illustrates a rise in the depreciation rate. A rise in the rate of depreciation from δ to δ' causes an increase in the slope of the depreciation line. As a result, the line pivots upwards. The intersection of this line with the the savings line determines the new steady state level of capital, K', which is lower than before. The associated steady state output Y' is also lower.

Example 3: A Rise in Total Factor Productivity, A

A rise in A has two effects. First, it causes the production function to shift upwards: A rise in A causes $Y = AK^{\alpha}L^{1-\alpha}$ to rise for any given level of K i.e., the production function shifts up. Second, because savings is a constant fraction of Y, the savings stock is now a constant fraction of a *higher* Y. That is, savings moves from being sY to sY'. As a result, the savings line shifts upwards. The intersection of this new savings line with the depreciation line pins down the new steady state capital stock, K'. As noted already, it is such increases in A that ultimately cause sustained increases in living standards in developed economies.

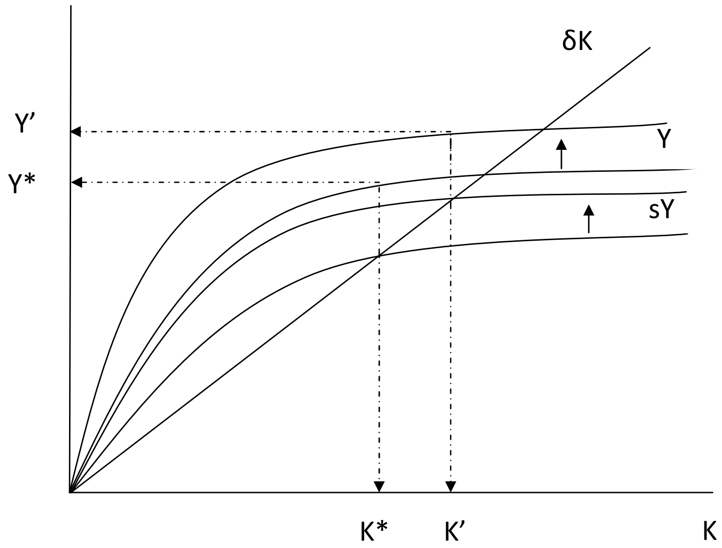

Figure 2.8: An increase in A. Ultimately, it is continual increases in A that sustain increases in living standards in developed countries. Observe how an increase in A induces an increase in K too.

Chapter 3

Further Growth Theory

"Little else is requisite to carry a state to the highest degree of opulence from the lowest barbarism but peace, easy taxes, and a tolerable administration of justice: all the rest being brought about by the natural course of things."

- Adam Smith.

3.1 Growth Accounting

Recall our production function from Chapter 2:

$$Y = AK^{\alpha}L^{1-\alpha}.$$

The objective of *growth accounting* is to quantify the contributions of A, K, and L to growth in Y. In particular, it enables us to deduce the contribution of total factor productivity A to growth in Y. Denoting the growth rate of variable X by g_X, a little manipulation of the production function—don't worry about the details—gives:

$$g_Y = g_A + \alpha g_K + (1 - \alpha) g_L. \tag{3.1}$$

That is, the growth rate of real GDP is equal to the growth rate of total factor productivity (TFP) plus a weighted average of the growth rates of labor supply and the capital stock (where the weights indicate how important these latter factors are in the production process). Now, you might wonder: how do we measure total factor productivity growth g_A? Isolating the TFP growth rate g_A gives

$$g_A = g_Y - \alpha g_K - (1-\alpha) g_L. \tag{3.2}$$

Hence, given we have data on α, g_Y and g_K, we can deduce what g_A is. Because we are inferring g_A from what's left after accounting for the contributions of labour and capital to growth, g_A is often referred to as the *Solow residual.* It is the growth unaccounted for by factor (i.e., capital and labor) accumulation. For example, when we go to U.S. data we find that g_A accounts for a large part of output growth, g_Y. By contrast, when growth accounting is done for "catch-up" countries—such as the "Asian tigers" in the eighties and China today—we typically find that the contribution from g_K is the most important determinant of g_Y. How come? Since these countries haven't yet hit diminishing returns, capital accumulation is especially productive in developing economies. Yet because they've already done considerable capital accumulation, g_A is more important as a source of growth in developed economies.

Our expression g_A gives total factor productivity growth's contribution to output growth. It is the growth that can't be explained by either changes in capital or labour. For instance, if we had output growth, without *any* change in inputs, from above we would have $g_Y = g_A$, so all output growth in this case would be attributable to *total factor productivity.* So somehow, we must be utilizing our inputs in a different manner. In particular, if g_A rises, then we must have been using our inputs more *efficiently.* It is not necessarily technology, but is *anything that somehow changes the relationship between our inputs.* As an extreme example, imagine a corrupt government comes to power and expropriates *any* output produced by firms and workers. In this case, people go to work each day, but don't produce *anything* (why should they?) and the factories simply come to a halt. Therefore, in this case, we clearly have an enormous fall in output Y, but we certainly have no change in the capital stock K or the number of workers L. All the change in Y stems from a fall in A. (In fact, this little story describes pretty well a lot of developing countries.)

3.1.1 TFP, Incentives, and Growth

A widely cited paper in this area notes that *"the dominant cause of large international dispersion in levels of output per worker is due to differences in total factor productivity."* Determining what causes the low levels of TFP in developing countries is difficult, but the two main contenders are geography and institutions.[1] Yet there is a growing consensus that poor governments are the main cause of low income levels. Probably the most convincing evidence in favour of the government's role in development is that of Korea. After World War 2, the South become capitalist, while the North become communist. In 1953, when they were divided, GDP per capita in both the North *and* the South was $770. But by 1998—45 years of this "experiment"—North Korea had an income per capita of $1183, while South Korea had an income per capital of $13317. In addition, by 2002 life expectancy in the North was 60, while it was 73 in the South.

Characteristics of a *good* economic environment—leading to a high A—are the preservation of property rights, good bureaucracy, sound financial institutions, lack of corruption, high degree of political/economic freedom, openness to trade, lack of ethnic conflict, and good law enforcement. Yet governments can adversely affect all of these factors. For instance, if there are no property rights (i.e., legal claims to what I produce), I have no incentive to innovate. If an economy is closed to trade, domestic firms have no incentive to remain competitive (via lower costs, product development etc), and can't benefit from the adaption of new technologies from abroad. And if there's reams of bureaucracy, firms have little incentive to start any kind of enterprize and invest (and recall how investment

[1] The term *institutions* is a broad one, referring broadly to the economic environment established by the government.

determined income levels in *Solow*).[2] If people can't use their education productively due to the aforementioned issues, then they have no *incentive* to get an education.

Geography

The view that it's bad geography that stifles growth has become less prominent as the importance of good government and *economic freedom* have become more apparent. Moreover, there are many successful countries such as Hong Kong located in climates not typically perceived as conducive to growth. For example, Las Vegas, located in a desert, is currently the fastest growing city in the U.S. It seems that if the economic environment is favorable, growth can take place almost *anywhere*.

The Industrial Revolution

The commencement of sustained growth in real GDP per capita, $\frac{Y}{L}$, began in Britain around 1760. Before that, the world was generally described by the *Malthusian model*. According to this model, rises in GDP Y induced rises in population L, thereby depressing real GDP per capita $\frac{Y}{L}$ again. So, over most of historical experience (that is, the period prior to 1760), the standard of living of living remained relatively stable. Only after that, did Britain break from its so-called Malthusian trap. You might ask, why did this happen? After 1760 increases in output were so large that they were able to overwhelm any attendant increase in L; overall, therefore, real GDP per capita, $\frac{Y}{L}$ rose. In addition, as incomes continued to rise, fertility rates declined. There are two reasons for this. First, as females' wages rose, the opportunity cost of having children also rose. As a result, instead of having large families, many women decided to have fewer children and spend more

[2] Hernando de Soto, in his revealing book *"The Mystery of Capital"*, reports that in Lima, Peru, it took 289 days for a team of people working six hours a day to meet the regulations required to legally open a small business producing garments. In New York City, it took a day. (In an earlier book, "The Other Path," he revealed that along the way, ten bribes were solicited, and it was necessary to pay two of the requested bribes in order to get permission to operate legally.)

time in the workplace. Second, in developing countries, children are frequently viewed as potential workers—for the family farm, say—and as carers who will tend to parents in their old age. Yet as economies develop, economic activity becomes less focussed on agriculture and more geared towards manufacturing and services, thus reducing the need for children.

3.1.2 New Growth Theory

The emphasis in the Solow model was capital accumulation. But when all is said and done, the model explains *everything but the reason for long-run growth itself*—total factor productivity growth. Herein lies a serious problem: the main reason for long run growth is almost incidental to the mechanics of the model, it is just *assumed* at the end. To redress this, the *New Growth Theory* seeks to explain where the A comes from and how it evolves. Because the New Growth Theory is largely concerned with growth in developed economies, it takes the view that A evolves as a result of technological progress (and so ignores other factors such as improvements in institutions, climate, and so on.) Across developed countries, most rises in A are primarily due to technological progress.

Therefore, in this framework, we think of A as ideas and inventions. One important implication of the theory is that we must give firms *incentives* to undertake research and improve technology—and that incentive is profit. One way of ensuring firms that they can reap profits from innovating is to grant them *temporary monopoly power* via intellectual property rights; i.e., patents and copyrights. This way, they can recover the costs of innovation/research, and also make profit. If they didn't have such guarantees, they wouldn't embark in research in the first instance. For instance, in the U.S., patents—which confer property rights on innovations, and hence guarantee a firm monopoly power—last for twenty years.

According to this theory, the government has a crucial role to play in maintaining property rights and a stable economic/political environment conducive to innovation and entrepreneurship. In addition, by giving, say, tax subsidies to those involved in scientific research, the government can actively promote research and development. Another insight from New Growth Theory is that a higher population can be beneficial for growth: when there are more people, there are more creators/researchers, and, through positive externalities from research, their innovations can eventually benefit everyone.

Chapter 4

Savings and Investment

4.1 Interest Rates and Real Values

Interest rates are almost always quoted in monetary terms—what we call the *nominal* interest rate. But remember what I said earlier about *real* values? In finance, what we are really concerned with here are *real* returns: the return in terms of *purchasing power.* For example, if I made a real return of 3%, say, then I'd be able to purchase 3% more goods and service with the money that was returned to me. The real return gives the percentage increase in purchasing power. Mathematically, it is given by the nominal interest rate less the rate of inflation; intuitively think of it as the monetary return less what is "eaten away" by inflation.

In fact, lots of savers lost out in the $70's$ since they bought bonds and only after did they witness a resurgence in inflation, thereby diminishing their *real* returns. In other words, they lent out money, but the purchasing power of what they lent was much greater than what they got received back. For instance, if I lent \$10 to someone a hundred years ago and I got a mere \$12 back today, then this has hardly any *purchasing power* compared to what I lent out, given the enormous price level increase in the interim.

Definition 11 (Real Interest Rate) *The real rate of return r is given by $r = i - \pi$,* *where i denotes the nominal interest rate and π inflation. The real rate of return indicates the percentage change in purchasing power from making a loan.*

Because the real interest rate represents the ultimate reward and ultimate burden of payment, it is the real interest rate which determines savings and investment decisions.

4.2 Determination of the Equilibrium Real Rates: Loanable Funds Model

What I mean by investment is demand for funds for building new factories, machines, building homes (i.e., residential investment), and so on.[1] Investment demand comes from firms and residential investment. When real interest rates fall, investment rises. Why? Basically, the real interest rate represents the cost of loanable funds and when that cost falls, funds become cheaper, making investment projects more attractive. For instance, a firm will typically have a rule whereby they should invest in all projects that satisfy the condition $r^p > r$; i.e., where the expected rate of return on the investment project, r^p, is greater than the real interest rate. So when the real interest rate falls more projects satisfy this condition, causing the demand for investment funds to rise. Often we write the investment curve—representing investment demand—as

$$I(r) = a - br,$$

where $b > 0$ mediates the sensitivity of investment to interest rate changes, while a captures exogenous influences on investment demand. An example of an increase in a would be a rise in investment demand due to a new technology which firms wish to invest in. Most importantly, such an investment demand arises not from a change in interest rates, but from a factor unrelated to interest rates. Graphically, we would represent this

[1]It is common in the media to refer to savers as investors. However, in our model, these are two distinct agents. One is the source of supply of loanable funds (the savers); another is the source of demand for those funds (the investors). Investment demand here strictly refers to demand for investment projects.

by a shift outwards in the investment curve; this way, for any given level of the interest rate, investment demand rises.

Turning to the other side of the market, when the return to saving is higher, people save more. There are three sources of savings: households, governments, and foreign savings. The sum of private and government savings is called *national savings.* We write the savings curve

$$S(r) = c + dr,$$

where $d > 0$. That is, savings are increasing in the interest rate: as the interest rate rises, savings becomes more attractive. Empirically, savings are fairly irresponsive to changes in interest rates, which translates into a low d and an inelastic savings curve. As with investment, c embodies exogenous changes in savings. For instance, if people become concerned about the future, people might save more; i.e., there would a rise in what are called *precautionary savings.* Again, this change is independent of the interest rate. Graphically, we would represent it as a shift outwards in the savings curve; i.e., the level of savings would rise for each level of the interest rate.

4.3 Savings/Investment Diagrams and Examples

Reflecting the relationships above, we draw an upwardly sloping savings curve and a downward sloping investment curve. Think of the savings curve as simply a *supply* curve, with the investment curve being the source of *demand* for funds.[2] The slopes of each curve determine elasticities. As we shall see in the examples, the slopes play an important role in the determination of equilibrium interest rates. Implicitly, we are assuming that all savers and all those with investment projects meet together at some auction—our loanable funds market—and bid over funds. Although this is unrealistic, it is a useful way to analyze the basic forces determining the real interest rate in an economy.

[2]Conveniently, we denote the Savings curve with an S, the standard way of indicating a supply curve.

Here's how it works. When the level of investment demand exceeds the level of savings, those with the most profitable investment projects outbid those with less profitable ones. As a result, there is upward pressure on interest rates. Moreover, it will ultimately be those with the most profitable projects who manage to get funds. When the price of loanable funds rises, those with the least profitable projects are priced out of the market. In a sort of Darwinian "survival of the fittest", the market leads to an efficient outcome, and allocates funds to the most worthy investment projects. This way, the market system or "invisible hand" allocates funds in the most efficient way.

Finally, keep in mind that I am assuming output is always at potential. To see what this means, consider a closed economy with no government; in this case we will always have $Y_n = C + I$, where Y_n denotes potential output. In this analysis, therefore, what changes behind the scenes is the *distribution* of output between consumption and investment. For example, suppose consumers decide to consumer more, causing saving to rise. What happens? Well, mathematically, since $Y_n = C + I$, we know that I must fall. Central to this adjustment is the interest rate. As we shall see in the examples, when savings fall, the interest rate rises, which in turn reduces investment. As such, one can view this analysis as a way of determining how output is distributed across the economy. In particular, consumption is the flip-side to savings, so when savings fall, consumption rises; and because savings equals investment, investment also falls.

4.4 Graphs

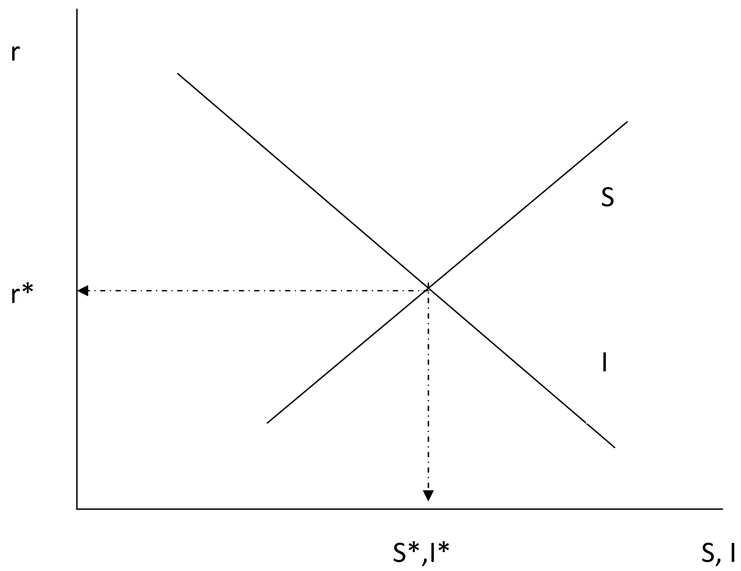

Figure 4.1: Financial market equilibrium. At equilibrium, savings equal investment. When the interest rate is above the equilibrium rate, savings supply exceeds investment demand. This surplus causes the price of loanable funds—i.e., the interest rate—to fall. As a result, two things happen. First, because saving is becoming less attractive, some savers decide to leave the market. Second, as the cost of funds falls, some investors are seduced into the market. The fall in savings and rise in investment demand causes the initial differential between savings and investment to fall. This process continues until both supply and demand for funds are equal—which occurs at the intersection of the curves. The opposite forces dominate if the economy starts below the equilibrium interest rate. In contrast to the Solow model, we assume these adjustments occur almost instantaneously.

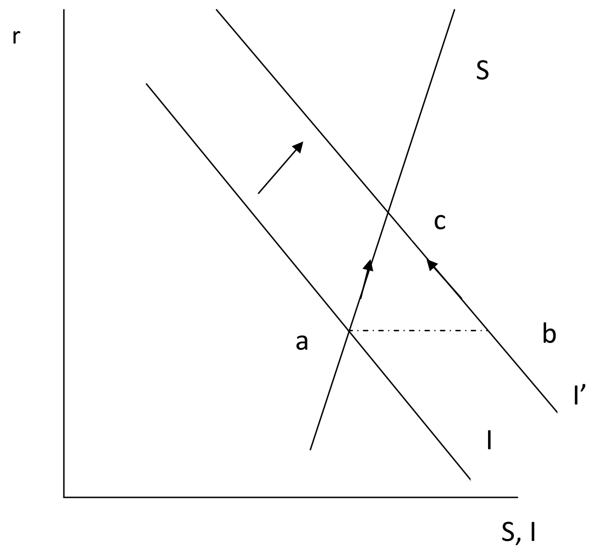

Figure 4.2: A rise in Investment Demand. We start off at point a. In particular, savings are initially at point a. When investment demand rises, investment demand immediately rises to point b: The horizontal distance between a and b measures the rise in investment demand. Because investment demand now exceeds savings supply, the financial market is in disequilibrium. At this point, the investors with the most profitable projects bid up the interest rate. Two things now happen. First, as the interest rate rises, savings rise. As we move from a to c, higher interest rates draw more savers into the market. Second, as we move from b to c, the interest rate is rising, causing some investors to leave the market, resulting in a decline in investment demand. This process continues until equilibrium is reached at point c. Investment and savings are both higher, but note that the rise in actual investment is less than the initial rise in investment demand. This is because the higher interest rates have caused some investors to leave the market.

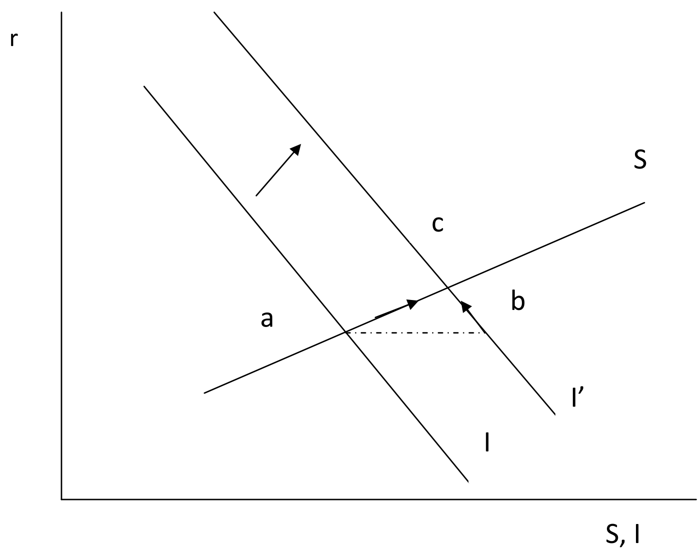

Figure 4.3: A More Elastic Savings Response when Investment Demand Rises. This represents the same increase in investment demand as in Figure 4.2. However, in this case, savers are highly responsive to rising interest rates. As a result, savings increases a lot when the interest rate rises. In equilibrium, both savings and investment are higher, while the interest rate is lower. The more elastic the savings response, the greater the rise in actual investment in equilibrium. This is because savers readily supply more funds as the interest rate rises, and this increase in supply can finance more investment.

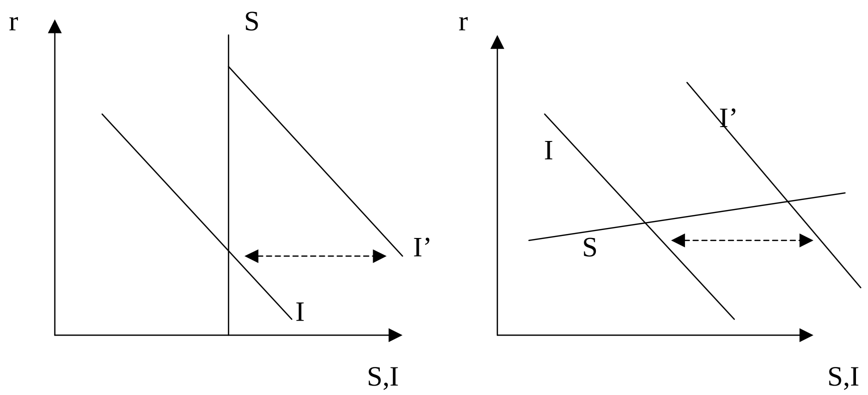

Figure 4.4: Rise in Investment Demand with Inelastic Savings Supply. (Note that an Inelastic curve looks like an I.) On the left, we assume savers are irresponsive to changes in interest rates. Despite higher interest rates, they always save the same amount. In this case, investment can't increase (since its pinned down by the fixed stock of savings), so an increase in the demand for funds is reflected wholly in a rise in price of funds; i.e., a rise in r. That is, the interest rate rises a lot to "choke off" investment demand. The burden of adjustment falls entirely on price; i.e., the interest rate. In this case, the investors with the least profitable projects are driven out of the market. For a given increase in investment demand, we get two markedly different equilibria.

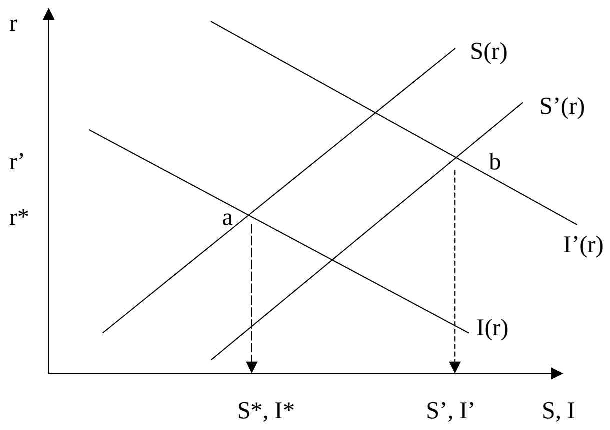

Figure 4.5: An rise in investment demand and an increase in capital inflows. In the U.S. in the mid-nineties, investment demand rose due to companies seeking to invest in information technology. In addition, during the Asian financial crisis, there was a surge of capital inflows to the country (since the U.S. is always seen as a "safe-haven" in times of global economic turmoil.) Whether the interest rate rises or falls depends on the magnitudes of the respective shifts; here the interest rate rises from r^* to r'. Because of the capital inflows, interest rates don't rise as much: the capital inflows neutralize some of the upward pressure on the interest rate.

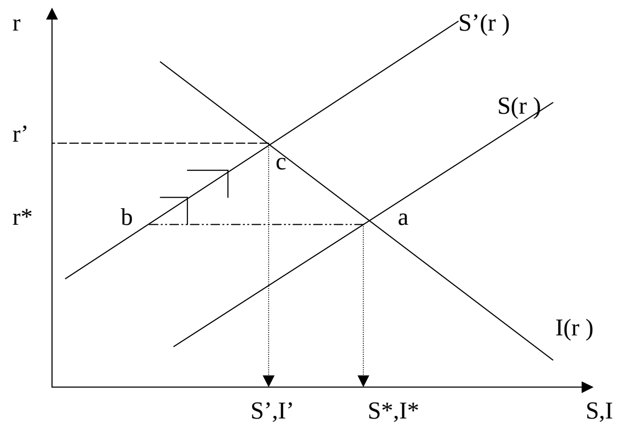

Figure 4.6: Crowding Out: A fall in government savings—due to permanent tax cuts, thus reducing $T - G$ (the budget balance)—leads to a shift inwards of the saving curve (we initially move from a to b), with the distance between a and b indicating the actual fall in savings. However, in response to rising domestic real rates, savings rise as the economy moves from b to c. In addition, because of the rising interest rate, investment demand falls as we move from a to c. Overall, because of the induced rise in savings, actual investment does *not* fall by the full extent of the fall in savings. Nonetheless, investment still falls from I^* to I', reflecting the notorious crowding out effect of government expenditure. We know from the Solow model that a lower investment rate ultimately reduces the standard of living in an economy.

4.5 Other Factors Affecting Interest Rates

It is best to consider the rate we have derived as the risk-free real rate of interest. This is the real interest rate upon which all other interest rates are based. To see reasons why rates might diverge from this, we must go beyond the simple loanable funds model. In that model, we simply assumed all savers and investors met together and bid over funds. Yet in reality there are a variety of mechanisms in an economy, through which savers and investors can interact. One obvious example is a bank; savers leave their money in banks, and banks lend that money out to people who borrow for investment projects. Another common example of this interaction is a firm issuing a bond. This is essentially a case where a saver gives the firm money, and the firm issues the saver—called the bondholder—with an IOU. This process of issuing bonds is the way governments borrow. A less obvious example of savers and investors interacting is the stockmarket. Savers give new firms funds—which they use to finance a new business—and the firm issues them with shares. These shares confer part ownership of the firm on the savers (who have now become *shareholders*).

The rates of return on these various means of savings can differ quite substantially. But the crucial point is, the rates of returns from these assets are all benchmarked on the rate we derived in the loanable funds model. Below I list reasons for why interest rates may differ from the rate derived in the loanable funds model.

1. *Inflation Premium.* In reality, almost all rates quoted are nominal, not real rates. To ensure savers earn the real return in the loanable funds model—often called the "required real return"—we must add an *inflation premium* to it to compensate for inflation eating away at returns. For example, suppose the required real rate is 4 percent. If people expect inflation to be 5 percent, then the saver will demand a nominal rate of 9 percent. This way, the real rate will be $r = i - \pi = 9 - 5 = 4$ percent as required. As you might infer, in economies where people expect inflation to be high, nominal interest rates are also high. This positive relationship between

nominal interest rates and the inflation is called the *Fisher effect.*

2. *Risk Premium.* In the case of bonds, a risk premium compensates for risk of default. There are a variety of so-called rating agencies—such as Moody's—which rate the quality of bonds from various companies and governments. The lowest quality bond is called a "junk bond"; because of the associated risk, they typically pay the highest returns. In the case of stocks, a risk premium is required, because their returns are so volatile.

3. *Liquidity premium.* This is to compensate bond-holders for the fact that they might not be able to resell the bond to someone else. Bonds issued by the U.S. government are the most liquid in the world, and are easy to sell; there would be little, if any, liquidity premium attached to them.

4. *Term premium.* The longer the period for which a bond is issued, the higher the interest rate it will command. This is because there is greater risk involved in such a loan. The *yield curve* illustrates graphically the interest rates (or yields) offered on bonds of all maturities. For reasons mentioned, it almost always slopes upwards, indicating higher interest rates for loans of longer duration.

To sum up, higher nominal interest rates are not necessarily good. This underlies a general economic principle: *there's no free lunch.* High interest rates merely compensate savers for greater risk. In reality, all assets are available to savers, so returns must adjust to make all assets equally attractive.

4.5.1 Bond and Stock Returns

How are returns generated in practice? In the case of a bank, it is straightforward: you just place 80 in a savings account and get back 100, say, a year later. First I will discuss bond returns, and then stock returns.

Bond Returns

To see what I mean by a bond return, say you buy a bond for, say, 80, and it specifies it will return 100 in a years time. Therefore, the return on this bond is

$$r = \frac{100 - 80}{80} = \frac{20}{80} = 25\%.$$

Notice that the cheaper the bond, the greater the return.[3] There is an important inverse relationship between bond prices and interest rates/returns on bonds: *when bond prices rise, their return—or the interest rate they pay—falls.* For example, suppose the bond's price rises to 90. Then the return or interest rate from buying the bond will be

$$r = \frac{100 - 90}{90} = \frac{10}{90} \approx 11\% < 25\%.$$

Conversely, if the return or interest rate on the bond rises, the price must have fallen. Namely, the price falls *to generate* the higher return.

Key Idea 5 *There is an inverse relationship between bond prices and the interest rate/return they pay.*

Apart from the risk of default, bonds are risky if you wish to resell them. Imagine, for instance, I buy a bond from you: I give you 80 and you promise to pay me back 100 next year. Now suppose that 6 months later it transpires that I want cash quickly, so I decide to sell my bond to someone else. However, suppose that for some reason, everyone believes you are heading towards bankruptcy. In this case, people will be willing to pay me very little for the bond you issued me; namely, the risk of default has risen. In this case, someone might only be willing to pay me only 50, say, for the bond—in which case I make a capital *loss.*

The example above deals with a 1-year bond. In practice, many bonds are issued for longer than one year. Suppose now that I buy a bond from you for 60, and the bond

[3] You can also find the return by solving the equation $80(1 + r) = 100$ for r.

specifies it will pay back 100 in *two* years time. In this case I would make a total return of

$$r = \frac{100 - 60}{60} = \frac{40}{60} \approx 67\%.$$

This is the aggregate return you would get over two years. However, the return on bonds are usually quoted in terms of the interest rate you would get *annually*. In this case, to get a rough estimate of what this would be, you would divide the aggregate return by 2, giving around 33%.[4]

Stock Returns

Let me give you an example. Imagine you purchase a stock for 50. Next period the stock price rises to 60 and you get a dividend—i.e., a share of firm profits distributed to you—of 15. In this case, you make a capital gain of $60 - 50 = 10$ and a dividend of 15. Therefore, your return is

$$r = \frac{15 + 10}{50} = \frac{25}{50} = 50\%.$$

For a given dividend and capital gain, a higher share price translates into lower returns. In reality, stock returns are quite volatile, with savers frequently making large capital losses on stocks, when stock prices fall. In addition, because future stock prices and dividends are frequently uncertain, a saver likely does not even know what the return will be. As already noted, it is because of this uncertainty that stock returns are, on average, relatively high.

In practice most of the discussions of the stock market revolve around the prices of *existing* shares. In this case, however, when one purchases shares, you are simply buying

[4] Technically, this is incorrect, and only gives a rough estimate. To find the equivalent annual return, you would solve the equation $60(1 + r)(1 + r) = 100$, which gives $r \approx 30\%$ a year.

them off an existing shareholder.[5] Stock prices change frequently, and for a variety of reasons. Crucially, what determines a stock's price are the forces of supply and demand. At any given time, the supply of stocks issued by a firm is fixed. As a result, when demand changes, the stock price changes. And why would demand fall? Well, demand for a stock falls when people believe the company is question is facing difficulty (and vice versa.)[6] In such a circumstance, it is natural to expect the forecasts of future profits for the firm would be low, which in turn would lower expected dividends. This would reduce expected returns from the stock and hence reduce demand for it, causing its price to fall.[7]

Which brings me to my next point. The famous *efficient markets hypothesis* says that all publicly available information about a firm's prospects are almost instantly impounded into its share price. Underlying this important dynamic are people in financial markets who continually scrutinize firms' prospects, and sell or buy shares once news about future profitability is released. For instance, if news is just released suggesting a firm's future profitability will rise, people in financial markets will almost instantly buy the firm's shares, causing its share price to rise. Noting our formula above, this makes it quite difficult for the average person to make above-average returns from stocks; namely, the shares of firms with good prospects will have high share prices, tending to offset the fact that future dividends will be high; this in turn leads to average returns.

[5] In this case, investment would not increase when shares are purchased. You would be saving, but the seller would be dissaving, so there would be no change in aggregate savings.

[6] Discussion of such share price movements dominate much of media coverage of financial issues. For example, on 23 February 2010 one article at the Bloomberg website indicates "Stocks in Europe fell after German business confidence unexpectedly declined for the first time in 11 months." Lower business confidence indicates profits will likely be lower in the future; this in turn means dividends will be lower, and this expectation reduces the value of existing shares.

[7] Much of the discussion of the stockmarket in the media refers to the price changes of *existing* shares. Sometimes an established firm may issue new shares, thereby "diluting" the existing shareholders. Because firms typically use the funds from the new share offering for investment, this should lead to higher future profits. And because higher profits lead to higher dividends, this dilution of existing shareholders is not as bad as it seems.

4.6 The Maths of Savings/Investment

4.6.1 International Exchange

In 2004, the U.S. ran a current account deficit of −\$600bn. As a result, \$600bn left the U.S. and entered foreign hands. But what do the foreigners do with those dollars? Well, they buy U.S. assets—assets like real estate, stocks, bonds, and so on. Analogous to a boomerang, the money comes right back again. When the money comes back to the financial markets in the U.S., we call these funds *international savings*, which are denoted S_I. So in 2004, international savings in the U.S. economy were \$600bn. It follows that

$$S_I = -NX.$$

Crucially, note that *the counterpart to an exchange of goods is an exchange of assets.* Current account deficits aren't a free lunch.

Another way to think of this is as follows. Because the U.S. is running a current account deficit of \$600bn, it is living beyond its means by \$600bn. For this reason, the country must borrow savings from abroad. And since its absorbing \$600bn more than its producing, it must be borrowing \$600bn; and, as a country, it is borrowing from foreign savers. Those borrowings are called *international savings* and we again have $S_I = -NX = \$600$.

A current account deficit means a country is *borrowing* from abroad and accumulating foreign debt/liabilities. No matter how we look at it, *the counterpart to an exchange of goods is a capital flow.* Hence we often refer to S_I as *net foreign borrowing* from savers abroad. By contrast, a country can also lend funds abroad. In any case, the sum of all a country's past net foreign borrowing is called its *net foreign assets.* The U.S., which has been running persistent current account deficits for over twenty years, has experienced a dramatic fall in its net foreign assets, which are now decidedly negative. In fact, the U.S. is now the world's largest debtor.

The Maths of Savings/Investment: The Baseline Case

Bear with me for a moment. First, I will deal with the most simple example: a closed economy with no government. Now recall the one way of calculating GDP is to add the components of final expenditure: consumption and investment. Therefore

$$Y = C + I.$$

But recall that the value of expenditure is equal to the value of income; every expenditure is someone's income. Now ask yourself: what do people do with income? Well, they can either consume it or save it. So we have

$$Y = C + S_p.$$

Now equating both expressions for Y gives

$$C + I = C + S \Rightarrow S_p = I$$

$$\boxed{S_p = I} \tag{4.1}$$

That is, private savings equal investment. In other words, the stock of savings is used for investment. Finally, it is common to divide by Y and write this in terms of rates:

$$\frac{S_p}{Y} = \frac{I}{Y}$$

That is, the *savings rate* equals the *investment rate.*

Incorporating Government and Net Exports

Adding up the components of expenditure, we have

$$Y = C + I + G + X - M.$$

This tells us where the money comes from. And by the national accounts identity, this is equal to *income*. And what do people do with income? They either consume, save, or pay taxes. So we have

$$Y = C + S_p + T.$$

Now equating both expressions for Y yields:

$$C + I + G + X - M = C + S_p + T$$

$$\Rightarrow I + G + NX = S_p + T$$

$$\Rightarrow I = S_p + (T - G) - NX$$

and noting that $T - G = S_g$ (where S_g denote government savings) and $-NX = S_I$ this yields

$$\boxed{I = S_p + S_g + S_I} \tag{4.2}$$

This says that the flows *out* of the financial markets (i.e., for investment) are equal to the flows *into* the financial markets (private, government, and international savings.)

4.7 Private Savings and the Permanent Income Hypothesis

So far, we haven't discussed rigorously the issue of what determines savings by the private sector. In other words, what motivates people to save? For analyzing savings decisions, we consider the *permanent income hypothesis.* Essentially this states that a person's consumption level today depends on *all* lifetime income—and not necessarily *current income.* According to the theory, a person figures out what his lifetime wealth is, and then simply consumes a fraction of that each period (with the fraction depending on how many periods the person expects to live.) Think about that for a moment. That means, for example, that if my only income is 100 euros in ten years time, I will borrow today and for the next nine years, and consume 10 each period. Consumption-wise, *nothing* should change in 10 years time. Overall, my consumption pattern is independent of my income stream.

To see one implication, say you win a prize of 1000 euros, and so your income this period increases. If people base consumption on lifetime income, then winning the prize this year should have relatively little effect on consumption this year. To the extent that the prize changes lifetime income, it does affect consumption to some extent. That is, people will consume a little of it this year, but will *smooth* the rest of it over their lifetime. Now ask yourself: would your consumption response be the same if you received a *permanent* income rise of 1000 (due say to promotion)? Hardly. Because lifetime income goes up a lot, with a *permanent* change in income arising from, say, promotion, consumption rises by 1000 *each period*.

Underlying the permanent income hypothesis is the basic idea of diminishing marginal utility to consumption. Just imagine listening to a song or eating food: the last unit of consumption is never quite as good as the first. Point is, the "bang per buck" falls as consumption of a good rises. Precisely because of diminishing marginal utility to consumption in any *given* period, to maximize utility, consumers spread consumption as thinly as possible across *all* periods. Forget about macro for a moment, and think of it like this: like the way people spread butter over bread, it's optimal to spread consumption over a number of periods. And for the same reason too: would it make sense to put too much butter on part of the sandwich, leaving other parts dry? Of course not.

Chapter 5

The Money Supply and Inflation

"I don't care too much for money, money can't buy me love."

The Beatles

The Equation of Exchange

Over the period 1987 – 1996, Brazil had an annual inflation rate of 656%; yet by 2005, the inflation rate had been reduced to to 5.9%. This chapter addresses the question: what causes such large differentials in inflation? The *quantity theory of money* gives us the answer, but first we must discuss the *equation of exchange.*

Let me start with an example. Imagine, in a world of 15 people, I give one person a $10 bill in return for some new good they produced. In turn, that person engages in some exchange—of another newly produced good—with some other member of the population, and gives them the bill; for instance, they might purchase a newly made sweater from them. In turn, the sweater-maker then engages in trade with some other member, and this trading continues until everyone has had the bill. At the end of all the trades, the following equation—the *equation of exchange*—must be satisfied

$$MV = PY,$$

where PY is the monetary value of all the transactions: the price of each good multiplied by the number of goods. Here it is \$150. M is just the \$10 bill. Now V is what we call the *velocity* of that bill. It indicates how many times it goes around, or how many times it has been "recycled."[1] In our homely example, the velocity is just given by the number of people. Note that the quantity equation, $MV = PY$, *has* to hold. The money value of the bill (\$10) multiplied by the number of people (15) has to equal the value of total transactions. When dealing with the equation of exchange, we assume the value of total transactions is simply *nominal GDP*.[2] Recall that nominal GDP is the monetary value of new goods produced in an economy. In a more general setting, where there are different prices for different goods, P refers to the *price level*—e.g., the CPU—which is the average price of a good in the economy. Empirically, we measure velocity, the speed at which money "gets around," by $V = \frac{PY}{M}$ (which equals 15 in the example.)

The Long Run and the Quantity Theory

We assume the *quantity theory of money* holds in the long run. Underlying the theory are two assumptions. First, the theory assumes the velocity of money V is *constant.* Second, it assumes that money has no effect on output Y. This is called *money neutrality* or the *classical dichotomy.* To understand this, satisfy yourself that increasing bits of paper cannot increase output in the long run. See Table 5.1 showing vast dispersion in the rate of inflation (broadly reflecting money growth) and *no corresponding relationship* to unemployment (broadly reflecting *real economic activity*). Over the long run, there's no inverse relationship between unemployment and inflation. We can see above that unemployment is approximately the same over the long run, but inflation varies a lot. This is consistent with money neutrality.

Combining these assumptions with the *equation of exchange*, $MV = PY$, we see that

[1] You can think of it as measuring *how much work* the bill is doing.

[2] We are implicitly assuming all the trades were for *newly produced goods,* and not second-hand goods. The exchange of second-hand goods does not reflect new production and so are not part of GDP.

Table 5.1: The Classical Dichotomy: United States, 1954-2004

Year	***Unemployment Rate***	***Inflation Rate***
1954	5.6	-.37
1964	5.2	1.198
1974	5.64	12.1
1984	7.5	4.0
1994	6.1	2.6
2004	5.6	2.4

rise in the money supply causes a proportional rise in the price level. That is, $\uparrow M \Rightarrow \uparrow P$. For example, a doubling of the money supply merely leads to a doubling of the price level, with output unaffected. As a result, nominal and real variables are *independent* in the long run. Conveniently, this enables us to talk about and analyze real variables without worrying about how high the price level or inflation is. Indeed, this is what we did in the Solow model.

Yet, in the short run, money *can* affect output. As a result, the classical dichotomy breaks down. To see why, notice from the equation of exchange, $MV = PY$, that if prices are fixed and velocity is constant, a rise in M leads to a rise in Y.

5.0.1 Inflation Rates

Now, what implications does the quantity theory have for inflation, the rate of change of the price level? Noting that the growth rate of a product is approximately equal to the product of growth rates, take growth rates of the *equation of exchange* to get:[3]

$$g_M + g_V = g_P + g_Y,$$

[3]To obtain the growth rate of a *product*, we *add* the individual growth rates; i.e., if $Y = XZ$ then $g_Y = g_X + g_Z$.

where g_X denotes the growth rate of variable X. As already noted, the quantity theory assumes that the velocity of money is stable or constant, $g_V = 0$. Then, from above we have

$$g_M = g_P + g_Y \Rightarrow g_P = g_M - g_Y$$

In the long run, any level of money growth exceeding the level of output growth manifests itself as inflation. For example, suppose money growth is 12% and output growth is 10%. Then we need 10% money growth just to be able to buy the new output. For this reason, output growth (i.e., more *supply*) alleviates inflationary pressure. Or think of output growth as *absorbing* or "soaking up" some of the money growth. But any increase in money growth beyond the level of 10% leads to a higher level of price growth (i.e., inflation). So in the case of 12% money growth, 10% is used to buy the new goods (since they've increased in quantity by 10%), while the other "excess" 2% rears its head as a higher inflation rate.[4]

Now, what happens if money growth, g_M, increases? From above, we have:

$$g_P = g_M - \underbrace{g_Y}_{\text{not affected by changes in } g_m}$$

Because money growth g_M does not affect output growth in the long run, higher money growth leads to a higher rate of inflation; $\uparrow g_M \Rightarrow \uparrow g_P$. So, in the long run, $\frac{\partial g_P}{\partial g_M} = 1$; i.e., there is a $1 - 1$ relationship between changes in inflation and money growth; both move in concert—"*sustained inflation is always and everywhere a monetary phenomenon.*"

Now, you might wonder: why do governments increase the level of money growth beyond the level of output growth if it's inflationary? Well, because of weak tax systems and underdeveloped financial markets, governments in developing countries often *have* to finance a significant share of their expenditure by resorting to the printing press. This is

[4] With high levels of inflation and relatively low levels of output growth—i.e., $g_Y \approx 0$—it is common to write the relationship above as $g_P \approx g_M$.

called *debt monetization* or *seignorage* and is *the* unambiguous, underlying cause of high inflation or hyperinflation. High inflation rates, therefore, are always and everywhere a *fiscal* phenomenon; namely, they are an inevitable consequence of an unsustainable level of government expenditure. So although a hyperinflation is inherently a *monetary* phenomenon, the root cause of the problem lies in the fiscal domain. This is called *fiscal dominance*; the demands on the fiscal side frequently override concerns about rising inflation. Fiscal reform/restraint is *the* way to halt high inflation. Of course, financing through inflation is not a free lunch. When prices rise, the purchasing power of peoples' money falls. As such, inflation acts like a tax, and deprives people of purchasing power; this is the "inflation tax," often described as the cruellest tax of all.

Key Idea 6 (The Quantity Theory) *In the long run, changes in money growth (all else constant) induce proportional changes in the rate of inflation.*