Contents

1
Growth Facts

This material forms the empirical background for much of the theory that will follow. The facts developed in this chapter are what we're trying to explain with that theory. Additional facts will get introduced along the way, but these form the core for discussing economic growth over the long run. This chapter can be seen as an update on the original "Kaldor Facts" that motivated much of growth research for years.[1]

[1] Nicholas Kaldor. A model of economic growth. *The Economic Journal*, 67(268): 591–624, 1957

1.1 Growth rates and levels of GDP

The first thing is to simply look at how the (log) of real GDP per capita evolved over time, for a set of relatively developed economies. Figure 1.1 does this for the US, UK, Canada, Australia, and Mexico. The most important thing in this Figure is how boring it is. The growth rate was stable for each country.Moreover, that growth rate was very similar for all five. If you stare at the figure long enough, you can make the case that Mexico's growth rate became permanently lower around 1980 (or that Mexico's growth rate was somewhat higher than normal from 1965 to 1980).

Recall that the derivative with respect to time of something measured in logs is a growth rate. $\partial \ln y_t / \partial t \approx \Delta y_t / y_t$. Thus the slope of a line in Figure 1.1 is equal to the growth rate.

While the growth rates of these five countries are stable and quite similar, note that the level of GDP per capita was not similar. Mexico is demonstrably poorer than the others. The difference in log GDP per capita between Mexico and the US is about 1, which means that the ratio of GDP per capita between the US and Mexico is about 2.7 to 1, and has been for about sixty years. Mexico's living standards grew, but only enough to keep pace with the US. The other three countries in the figure are much closer to the US living standard, but again those gaps remained relatively stable over time.

Figure 1.1 and others use data on real GDP per capita from the Penn World Tables. This project attempts to measure living standards in different countries using a common set of prices, so that the measure of real GDP is comparable across countries. This is not a trivial task, nor is it foolproof, and there are issues with making these comparisons across countries.

Not every country is as boring as these five. Figure 1.1 plots (log) GDP per capita for the US (as a reference) and four new countries (Germany, Japan, South Korea, and China). For each of these four the locus of log GDP per capita was not linear, or often even close to

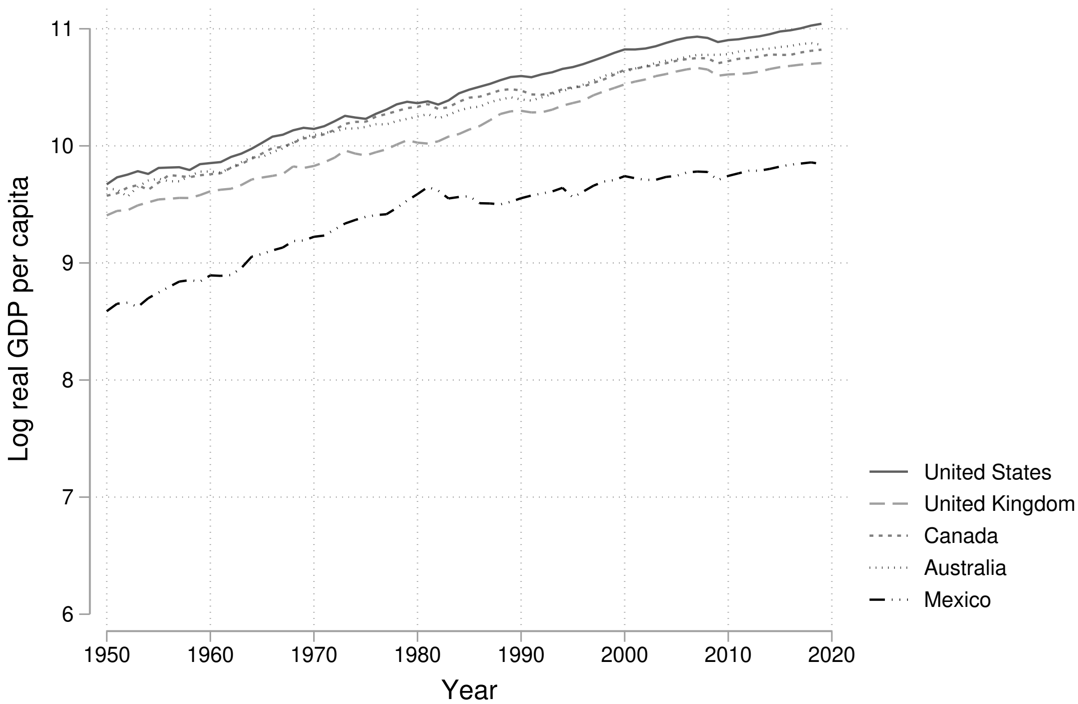

Figure 1.1: Log income per capita over time for select countries. Data is from the Penn World Tables 9.1.

linear. The slope of these lines changed over time, implying that the growth rate of GDP per capita changed as well. In almost every case, however, the lines look concave, meaning that the growth rate was falling over time.

Moreover, you can see that as GDP per capita in each country approaches the level of that in the US, the growth rate falls to match (roughly) the US growth rate. Germany, Japan, and now South Korea all seem to be asymptoting towards the level of living standards in the US. They started out much lower in 1950 and 1960, as you can see. In 1950 the gap in living standards between the US and Germany was about 2.7 to 1, between the US and Japan about 4.5 to 1, and between the US and South Korea about 25 to 1. But their high growth rates (as evidenced by the high slopes) allowed them to catch up to the US, at which point their growth rates slowed (lower slopes), and they are all now have living standards slightly lower than the US.

It is not imperative that the US be used as a reference point, although much of the economics literature implicitly does so. One of the reasons is that the US has good long-run data. We could use a different developed country as the reference points and all the facts in this chapter would still hold.

The experience of China is not as conclusive. In 1950 the gap in living standards with the US was about 43 to 1. Around 1980, the growth rate in China picked up, and since then it has been growing much faster than the other countries in the Figure. By 2015, this left the gap between the US and China at about 4.5 to 1. Based on the evidence of the other three countries, we might expect that in the

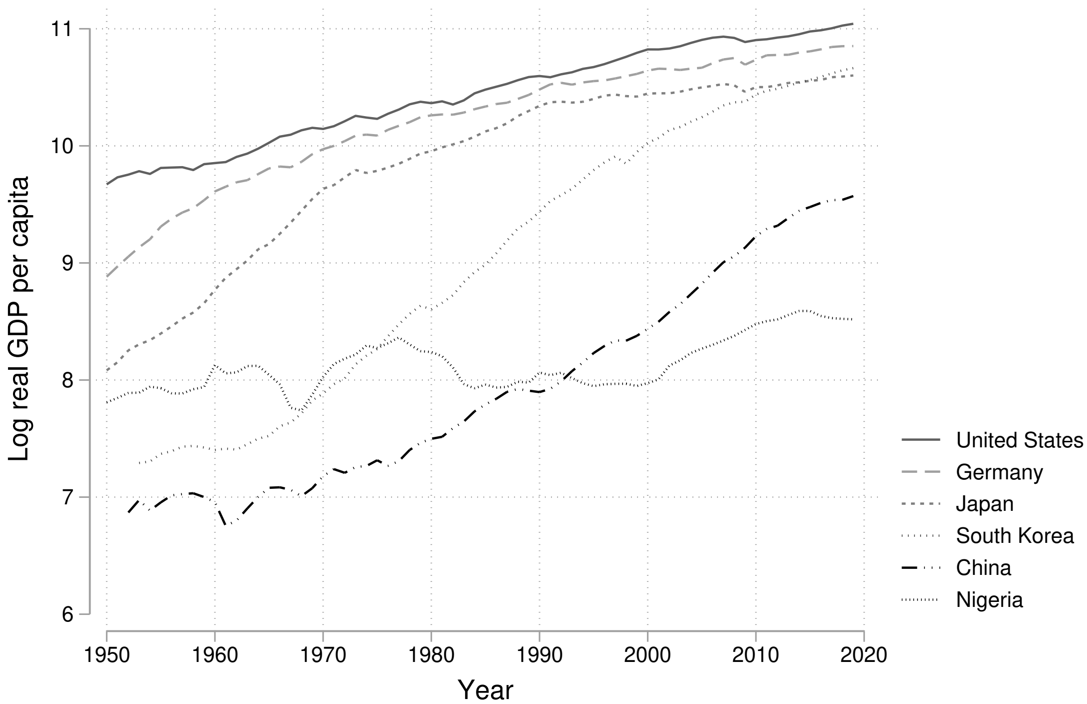

Figure 1.2: Log income per capita over time for select countries

future China will approach the living standard of the US, but that its growth rate will slow down so that it matches the US. We'll see in the following chapters why theory suggests that this is likely to happen. But we don't know this will happen for sure.

Nigeria also looks like an outlier here. In 1950 it had higher living standards than China and South Korea, but from that point forward the growth rate was close to zero, as shown by the line being close to flat. Only since 2000 does there appear to be some regular growth in GDP per capita. In this case we don't have anywhere close to enough evidence to decide whether Nigeria is entering a period in which it will catch up to the other countries, or whether it will continue to lag behind.

Nigeria aside, there is a tendency for countries to converge towards a common path for GDP per capita. This suggests that their growth rate is negatively related to the level of GDP per capita. It turns out to be helpful to look at this explicit relationship, as it helps illustrate the conditions where this convergence breaks down. Figure 1.1 plots the growth rate of GDP per capita, average over 10 years, against the level of output per capita at the beginning of a 10-year period, for six countries. It divides into two halves quite easily, with the dividing line being roughly a log GDP per capita of 8.5, which is roughly the living standard measured in Nigeria in 2015.

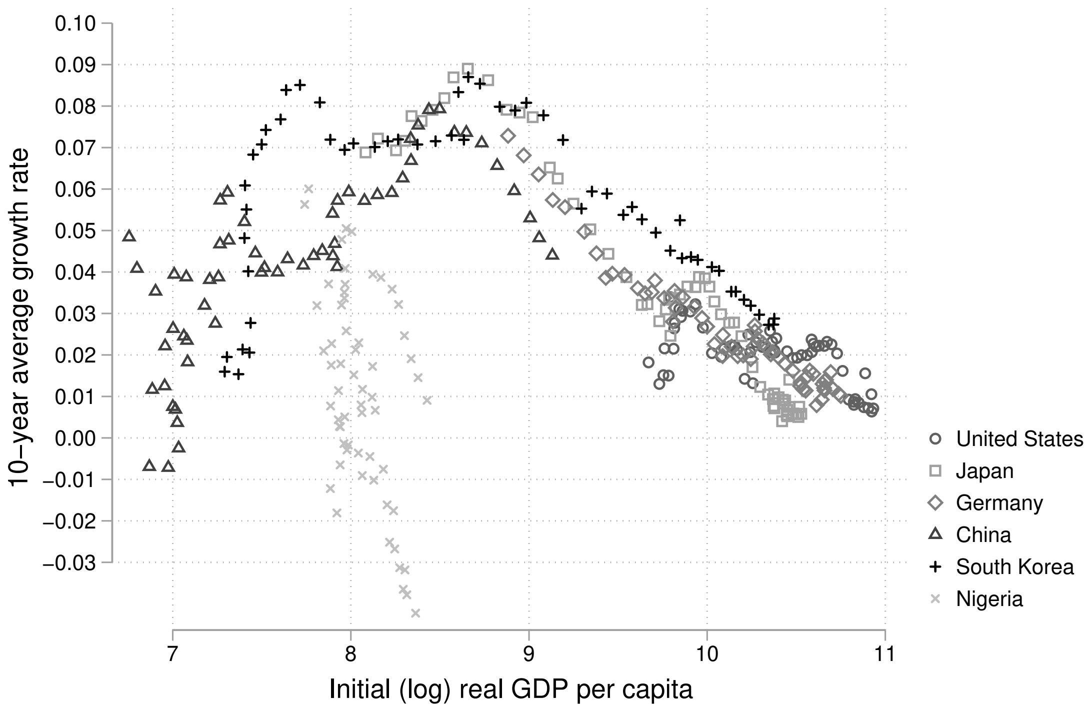

Figure 1.3: 10-year growth rates versus output per capita

To the right of that dividing line, there is a clear negative relationship of the growth rate and the level of GDP per capita. Given that economic growth pushes up the level of GDP per capita, by definition, this part of the Figure works a little like a slide. If you get on the slide with a GDP per capita around 8.5, then the growth rate is often around 6-8% per year, and GDP per capita rises rapidly. But as the living standard rises, growth slows down, until you get to the bottom of the slide at around (log) GDP per capita between 10 and 11, near the US level, and growth bottoms out at about 1-2% per year. The right-hand part of this Figure is the analogue of what Figure 1.1 showed for Japan, Germany, and South Korea converging to the US. The tight overlap of the data from different countries here indicates that there is a consistent process at work in all these countries, and at least suggests that we could use a similar model of economic growth to explain what was going on in all of them.

This follows, however, only to the right of that dividing line. For the data points with log GDP per capita less than 8.5 in Figure 1.1, there is no systematic relationship. If anything, it may appear that the growth rate gets lower as a country gets poorer. But overall it appears as if the growth rates in economies that are very poor tend to vary a lot, but that they do not systematically lead to convergence with rich countries. Nigeria is a good example, as you can see that the growth rate bounced between -3% and 5% throughout its recent history, leaving log real GDP per capita hovering around 8 the whole time. Once a country can lift its log GDP per capita above 8.5, then it appears to be on this slide or conveyor belt towards high living standards, but absent that there is no apparent logic to the growth rate of the economy. Note that this is not some iron law of nature, although if we plotted even more countries this same idea would come through as well.

There is a deep empirical literature on convergence across countries. Figure 1.1 captures a lot of the intuition in the results of that literature, without being systematic. What the Figure does not show is that if we looked at a much wider range of countries, we could find other "clubs" of countries that had negative relationships between growth rates and the level of GDP per capita, suggesting convergence with the club.

In terms of how we approach the study of economic growth, the division in Figure 1.1 is quite important. These notes, and this course, are focused on the explanation for the *right-hand* side of this Figure, with the implied convergence over time. The regularity of the negative relationship between growth rates and GDP per capita, along with several other regularities discussed below, will give us several very strong predictions about how economic growth works *in relatively developed economies*. Explaining what happens on the left-hand side of the Figure, among relatively poor economies, and what does or does not allow them to make the jump to the right-hand side, is beyond the scope of these notes and this course.

All those caveats aside, we can state some stylized facts for those developed economies that arise from the Figures already show. None of these should be taken as certainties or universal.

Fact 1.1 (Stable growth rates) *The long-run growth rates of developed economies appear to be stable in the long run.*

This fact comes from looking at Figure 1.1, in particular. In each case the growth rate (the slope of the plotted line) looks similar over many decades. Within Figure 1.1, the data for Germany and Japan corroborate this story, and South Korea appears headed towards a stable growth rate.

Fact 1.2 (Common growth rate) *The stable growth rates of developed economies appear to be the same in the long run, around 1.8% per year.*

This is a stronger statement that the first fact, and comes again from looking at the evidence in Figures 1.1 and 1.1. These economies not only appear to be headed towards stable growth rates, but given that the plotted trajectories of log GDP per capita are all parallel, this means they all have the same stable growth rate. If you calculate the actual slop of those lines, you get a number like 0.018, or 1.8% per year growth.

Fact 1.3 (Temporary growth differences) *Differences in the growth rate of GDP per capita between developed economies appear to be temporary.*

Fact 1.2 is a long-run fact, and as we saw in Figure 1.1 and Figure 1.1 the growth rates of countries can differ by quite a bit at any given point in time. But at least for developed countries the data suggest that those differences dissipate over time. In particular, the differences in growth rate appear to be driven by some countries being "behind" in living standards, and their high growth rate allows them to catch up, at which point their growth rate declines to match the other developed countries.

Fact 1.4 (Permanent level differences) *Even in the long run, there are permanent differences in the level of GDP per capita across countries.*

This fact is most evident by looking at Figure 1.1, and comparing countries to the US. While these countries all share a similar growth rate (i.e. the slope of the lines) they have different levels of GDP per capita (i.e. the intercepts). Thus there are persistent differences in living standards across countries even though their growth rates all converge to a similar rate.

The word "level" here is one that can get confusing. In Figure 1.1 we might say that the level of living standards is permanently higher in the US compared to the UK. This means the whole path of GDP per capita for the US lies above the UK. Another way of saying this is that the level of living standards is higher in the US than in the UK *at any given point in time*. But saying this does not imply a lack of growth in the UK, or higher growth in the US. As you can see in the Figure, they grow at the same rate, which keeps the level difference constant over time.

1.2 *Even longer-run evidence*

The evidence in this section is meant to complement that in the prior one. By using an alternate series measuring real GDP from Angus Maddison, we can trace back real GDP per capita even farther for a handful of countries. That data confirms that the facts just established above hold not just from 1950 to today, but from 1870 to today, a span of almost 150 years. Maddison's exact numbers for real GDP per capita do not match those used in the prior Figures, but that is not a failure of either data source. This only reflects differences in units.

Figure 1.2 plots the (log) real GDP per capita over time for four countries: US, UK, Germany, and Japan. For the US and UK, this shows that the Great Depression and World War II (roughly 1930-1950), while massive economic events, did not alter the long-run growth rate of either country. For the UK, there is some evidence that World War I, around 1915-1920, did push GDP per capita down

We could go back even further in time using Maddison's data. This would indicate that the stable growth rate of around 2% seen for developed countries only begins at some point in the early to mid-1800s, depending on the exact country. Prior to that the growth rate was probably closer to 0% per year, implying stagnant living standards. There is a rich literature encompassing economics and history studying the determinants of the jump from this "Malthusian" stagnation to sustained economic growth.

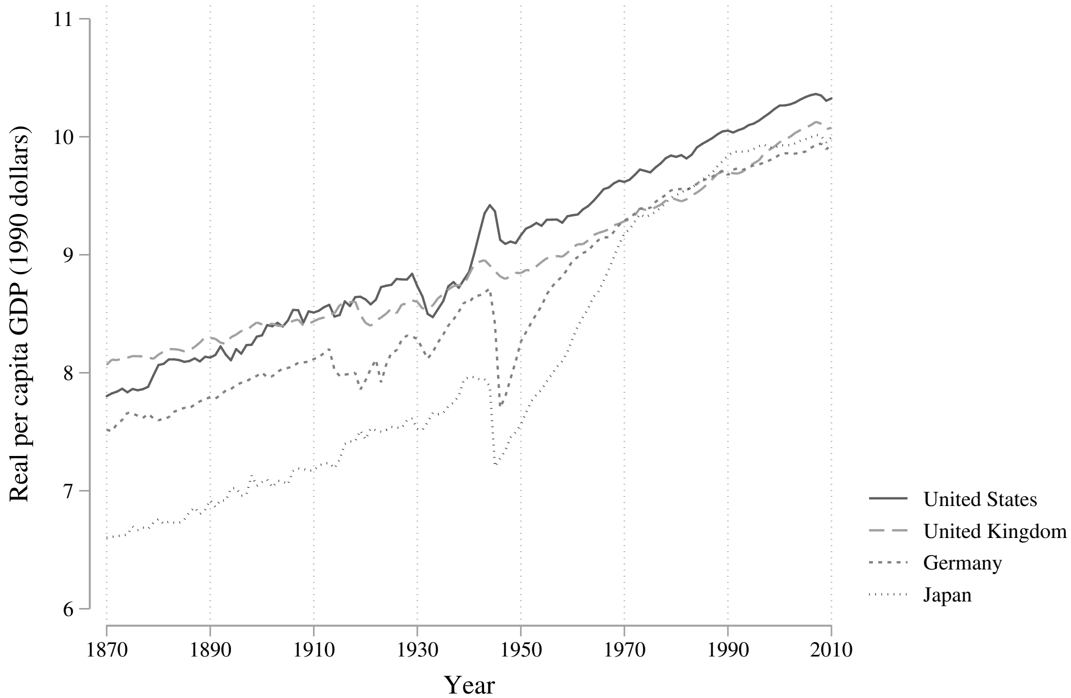

Figure 1.4: Real GDP per capita over time

permanently, but thereafter the growth rate was similar to before that war. For Germany, you can see as well that even with the disruptions caused by World War I and World War II, the economy eventually ended with the *level* of GDP per capita one might have extrapolated for it using the period from 1870 to 1910.

Japan had a different experience, which helps to illustrate that countries are not necessarily locked into a single path. Note that from 1870 to about 1940, Japan had a stable growth rate, but was much poorer than the other three countries. The gap between the US and Japan in 1930 was around 2.7 to 1. There was a significant drop in living standards because of World War II. But in the aftermath of that war, Japan grew very rapidly and did not return to the same level of GDP per capita one would have expected from 1870 to 1930. Unlike Germany, Japan changed the level of its living standards fundamentally in this period. It grew past the old level path of GDP per capita, and converged towards the richer countries like the US and UK.

The point is that you should not take Fact 1.1 about having a stable growth rate as implying that countries have stable levels of living standards. Japan illustrates that countries can, and occasionally do, shift from one level of living standards to another. Both before and after that shift they may end up with the same common, stable, growth rate. When we build the theory behind economic growth, we'll get some clues as to what might have changed in Japan after World War II to push up the level of living standards beyond the pre-war level. Unlike Germany, if we had extrapolated GDP per capita for Japan using data from 1870 to 1910 (or 1930) we'd have underestimated it from about 1960.

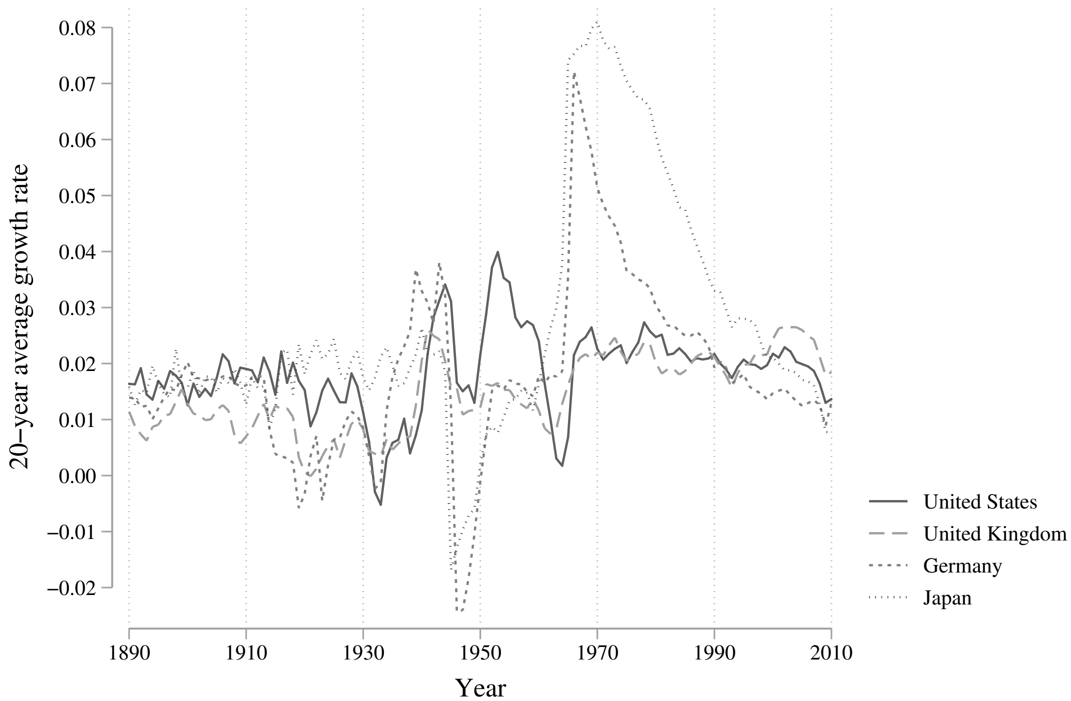

Figure 1.5: 10-year growth rates over time

We can gain some confirmation of Fact 1.3 from Figure 1.2. This plots 20-year average growth rates over time, to smooth out business cycle fluctuations. The point of the Figure is that while there were some significant spikes in the growth rates of these four countries (e.g. the Depression, after World War II), in each case the spikes dissipated over time, and the growth rate tended to return to its stable value. Note that this is true even for Japan, which shifted up the level of living standards. Nevertheless, its growth rate eventually dropped back to the roughly 2% per year value seen across almost all developed countries.

1.3 *Consumption and investment rates*

The prior facts relate to the growth rate of GDP per capita, but do not speak to the composition of GDP. The next set of facts have to do with the share of GDP accounted for by spending on consumption (e.g. food, non-durable goods, personal services) and investment (e.g. capital goods, structures, and some durable goods).

Figure 1.3 shows the share of GDP accounted for by consumption for a set of countries with stable growth rates over a long period of time. This is about as boring as you'd expect in their case. Consumption shares ranged between 50% and 70%, but tended to remain stable over time. There are exceptions, as with Mexico and the United Kingdom's falling consumption shares from 1950 to 1970, or the slow rise of the consumption share in the United States over time. But in broad terms, the consumption share did not have a distinct trend.

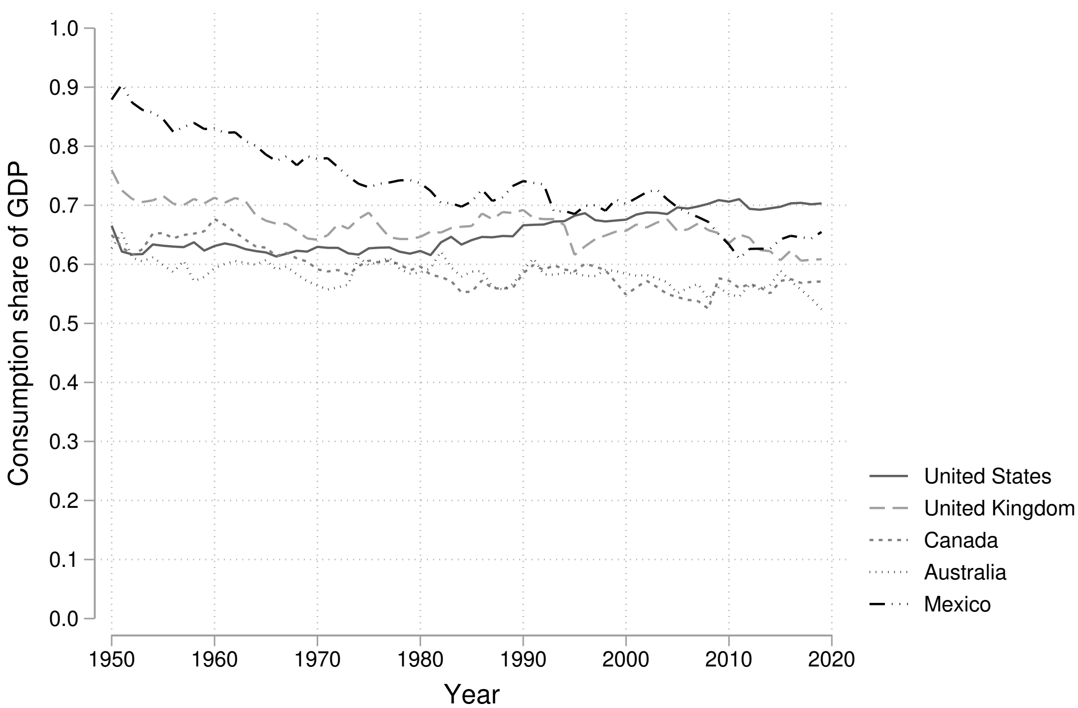

Figure 1.6: Consumption as share of GDP

There is more excitement with the second set of countries, in Figure 1.3. China stands out with a decline in the consumption share from 1950 to 2015, mirroring the rise in investment share over the same period. South Korea also saw a distinct drop in consumption share from 1950 to 1990, although at that point it stabilized around 50%. For Nigeria, the consumption share exhibited a lot of noise, which is probably due as much to a lack of data quality as it is to any real economic events. In the middle of the figure, at around 50-60% consumption shares, are Japan and Germany. Despite their rapid growth to catch up to the US, they did not have any distinct change in their consumption share over this period of time.

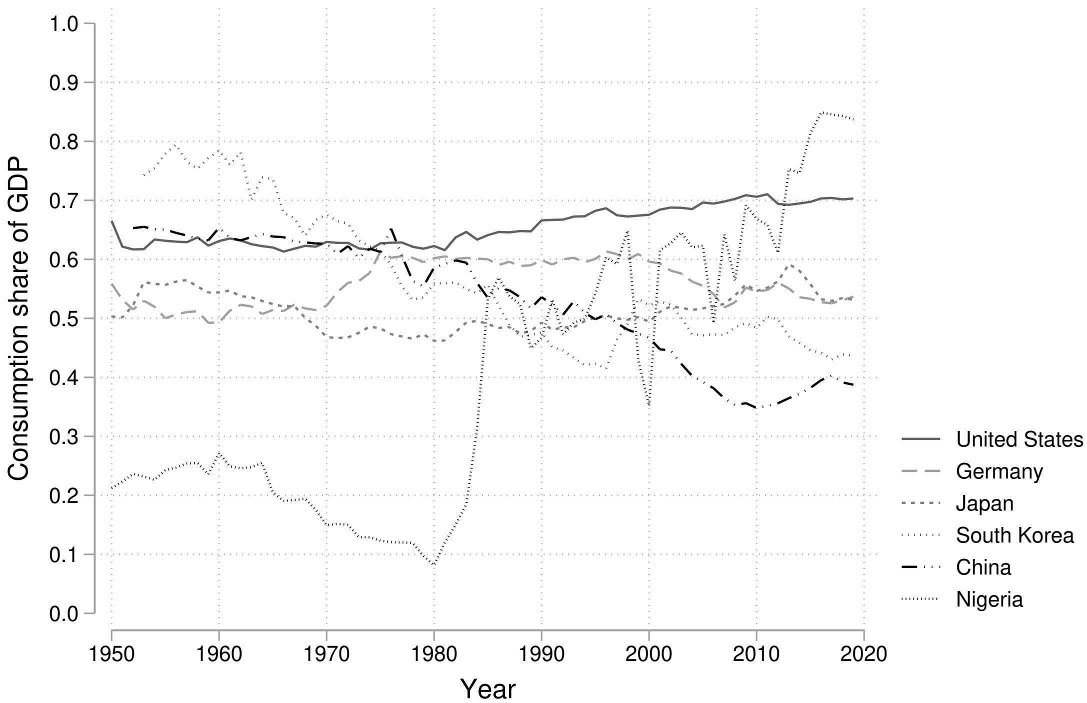

Figure 1.7: Consumption as share of GDP

We can do a similar study of investment shares of GDP. Figure 1.3 plots the investment shares for the countries with stable growth rates. The clear indication from the Figure is that these shares are stable as well, despite fluctuations. They also happen to be similar across the five countries, at a little more than 20% of GDP.

If we look instead at the set of countries that did not display such stable growth, there is again more variety in the data. Figure 1.3 plots the investment share in GDP for these countries, with the US again as a reference point. There is more dispersion here in the investment rate, with Germany, Japan, China, and South Korea all reaching investment shares above 30% at times. But even for these countries the investment shares all remain somewhat stable over time, with some evidence of a decline in recent decades for South Korea and Japan.

Nigeria appears to have experienced a significant drop in investment share around 1980, and then it remained at about 10% for a

China's evidence on investment and consumption shares do not seem to conform to patterns in most other countries. The growth rate of GDP per capita in China also appears as a clear outlier compared to the set of developed countries that are typically studied. Whether their economic growth is truly different than what we observed in the past is still open for debate. There are quetions about whether the data from China is reliable, and represents the actual facts on the ground.

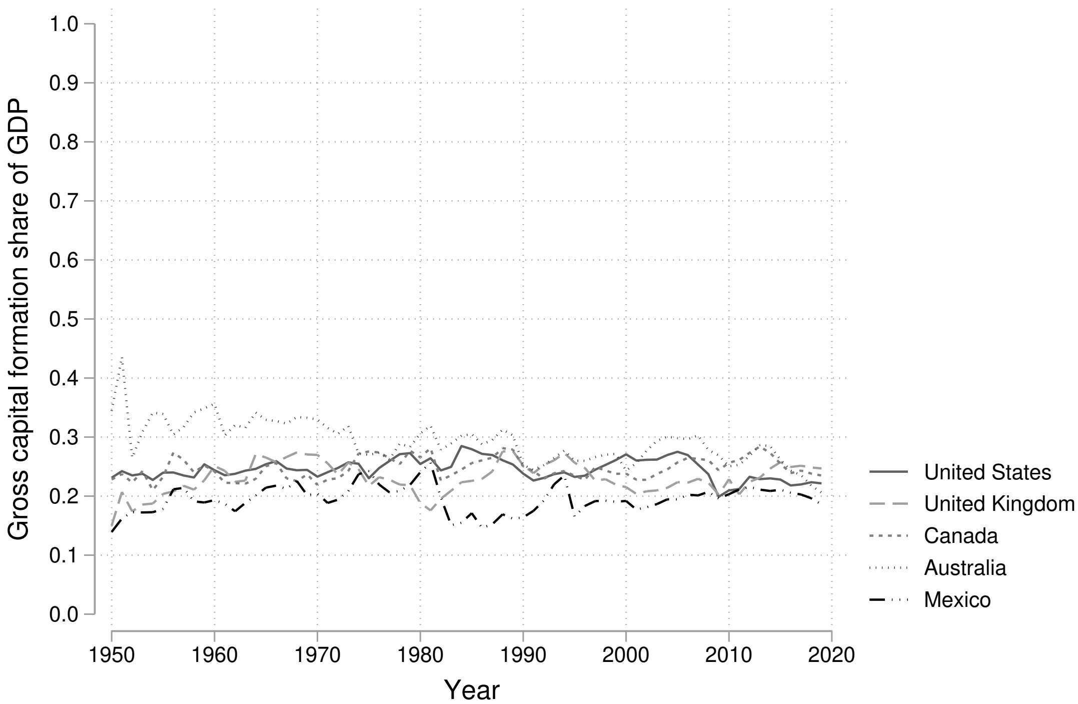

Figure 1.8: Investment as share of GDP

long period of time until a recent uptick. China is one country that seems to have a trend in the investment share. From around 20% in 1960, there was a steady increase to almost 50% by 2015.

What is not obvious from all these figures is whether there is any correlation of investment (or consumption) with the level of living standards. Figure 1.3 plots the investment share of GDP against the log of GDP per capita. To make this figure more readable, the over 9,000 actual data points representing investment shares and GDP per capita from each country/year observation were separated into 100 different bins, based on GDP per capita, and the data points each show the average investment share and GDP per capita for a bin.

Binning data as in Figure 1.3 is a convenient way of summarizing a large amount of observations that would otherwise make for a very messy figure. The binned data still capture the underlying relationship, and you could run a regression on the binned data and retrieve the same estimated correlation of investment shares and GDP per capita as you'd get using the full dataset.

There is a vague indication in Figure 1.3 that investment rates are higher when GDP per capita is higher, at least up to a log GDP per capita of 10. At that point, there is no obvious relationship. It is important to note that you cannot make any causal interpretation from Figure 1.3. That is, you cannot infer that if GDP per capita were to rise in a given country, then necessarily the investment share would rise. Nor can you infer that if the investment share were to rise, so would GDP per capita. All this data shows is that there is a very weak correlation of investment shares and GDP per capita for (relatively poor) countries. This correlation may be driven by some

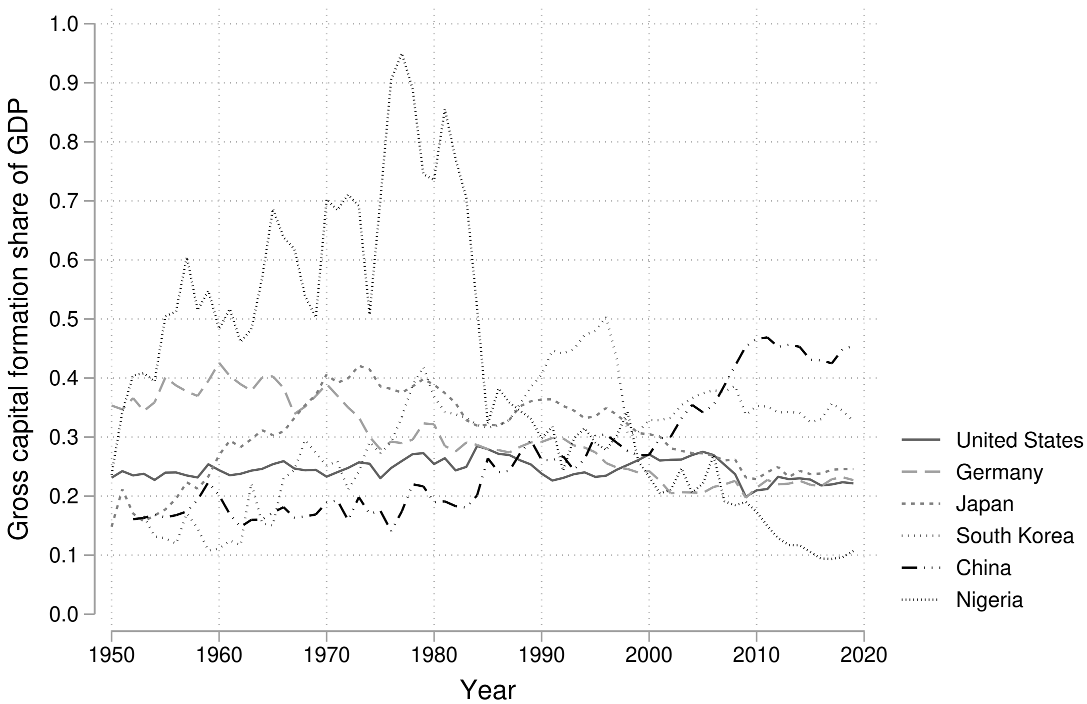

Figure 1.9: Investment as share of GDP

The investment share in Figure 1.3 is nominal investment over nominal GDP. But as the relative price of investment goods differs across countries, and tends to be higher in poor countries, the ratio of real investment to real GDP has a much stronger positive relationship with GDP per capita.

third factor we haven't considered.

We will establish a few significant facts from this section, although again you should be concious that these are not meant to be iron laws of economics.

Fact 1.5 (Stable investment and consumption shares) *The consumption and investment shares of GDP tend to be stable over long periods of time for most developed countries.*

There are obvious exceptions, which we can consider as we go forward. Also be aware that this fact on stable investment and consumption shares does not imply that these shares are the same across all countries.

Fact 1.6 (Investment share and GDP per capita) *There is a weak, positive, correlation of the investment share and the level of GDP per capita.*

As noted above, this is not a causal statement. It does mean that any theory of economic growth we develop should be capable of replicating this positive correlation. Our theory might imply that causality runs one way or the other, and that is something we might be able to establish empirically through some other method or data.

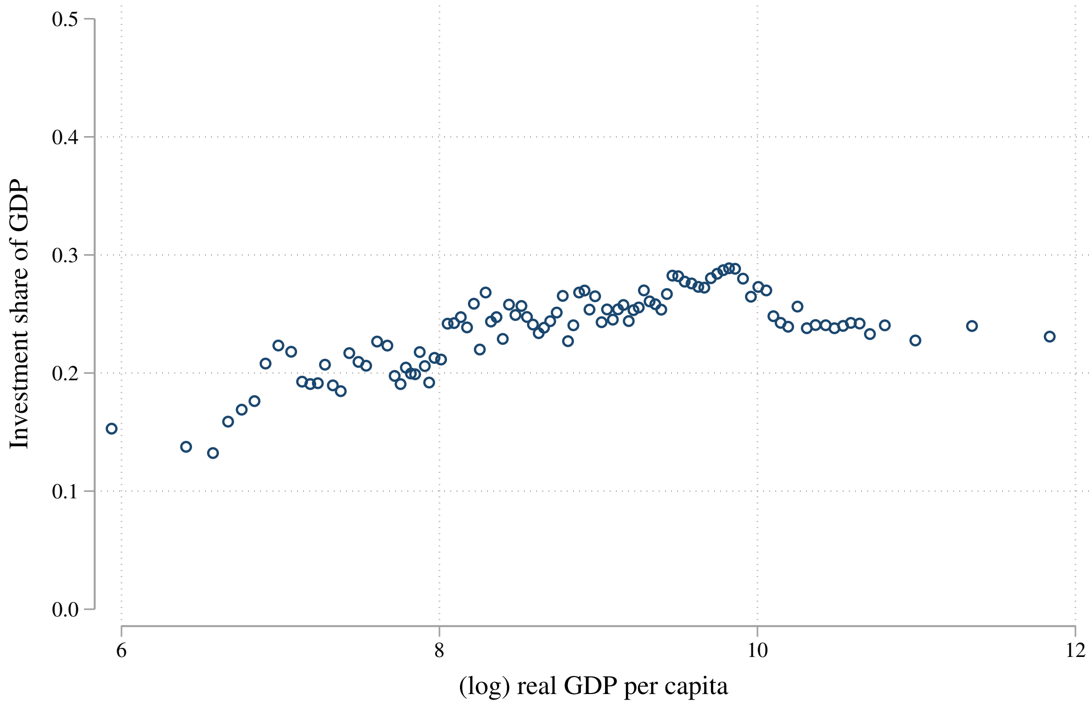

Figure 1.10: Investment shares and GDP per capita

But at a minimum our theory should capture this Fact in some manner.

1.4 Capital/output ratios

One of the reasons we are concerned with the investment share in GDP is that this investment spending represents the gross formation of new capital, where the "gross" refers to the fact that it doesn't account for the effect of depreciation on existing capital. That capital, which includes things like homes, office buildings, factories, equipment used by businesses, and intellectual property like software, is used to produce GDP. Hence investment spending, the capital stock, and GDP are all related. Much of the baseline theory of economic growth will deal with those relationships.

Prior to that, we can establish some facts about how capital and GDP are related. We'll look at capital/output ratios, which are simply the size of the capital stock divided by total GDP (output). The total stock of capital is measured here as described in the preliminary chapter, using a perpetual inventory method, meaning it is really a measure of the accumulated spending on investment, adjusted for depreciation.

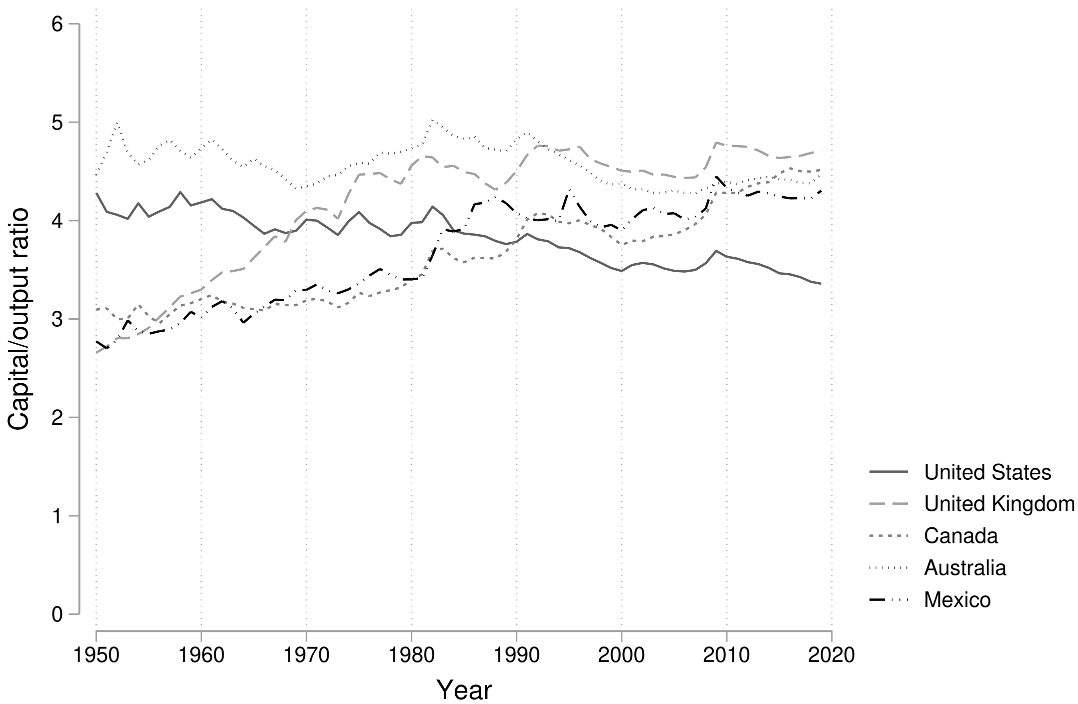

Figure 1.11: Capital/output ratio for select countries

What is the capital/output ratio measuring, leaving aside all of those caveats about how capital itself is measured? It would probably make more sense to think of the output/capital ratio, which would be a crude measure of the average product of capital, or how much GDP we were able to produce for each unit of capital we have. For reasons having to do with how we write down the theory mathematically, it has been more traditional to talk about the capital/output ratio, however. So the capital/output ratio is just something like the inverse of the average product of capital.

Don't confuse the average product of capital with the *return* on capital, which would be capturing something like the amount of output that the owner of a unit of capital gets paid for using that capital. That return could depend on the marginal product of capital, which *might* be related to the average product, but that relationship is not straightforward.

Figure 1.4 plots the capital/output ratios for the set of countries with stable growth rates. What this shows is that these ratios are somewhat stable over time, but obviously not constant. This implies that the average product of capital is somewhat stable over time as well, but not exactly constant. For the US, the capital/output ratio started around 2 and ended up closer to 3. In Canada and Australia, the ratios stayed more stable, but there is evidence they rose during the 2000s. Note that this means for all of these countries the average product of capital declined somewhat over time. In Mexico there was a shift up in the capital/output ratio around 1980, but stability before and after. Overall, however, these ratios do not have pronounced trends.

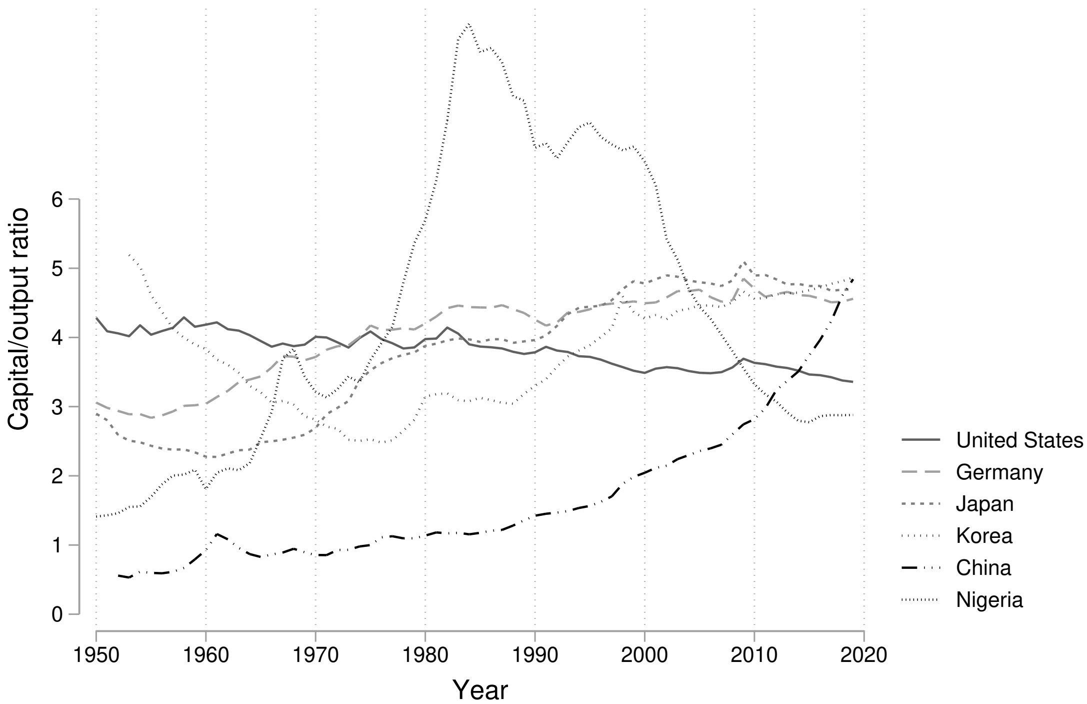

Figure 1.12: Capital/output ratio for select countries

For the other set of countries, there is again some more interesting action at work in Figure 1.4. Germany had a stable capital/output ratio from about 1970 forward, as did China, despite their rising investment share in GDP. In other words, their average product of capital did not change much over time. Japan's capital/output ratio rose, in contrast, from below 2 to over 4, before leveling off after 2000. Thus Japan's average product of capital *fell* for several decades. South Korea saw a similar rise in the capital/output ratio, although the rise seems to be slowing down in recent years. Nigeria had a significant swing up in the capital/output ratio, followed by just as dramatic a fall, leaving them with a capital/output ratio of just over 1 after the year 2000.

We can do something similar to what we did for investment ratios, and plot the capital/output ratio against the level of GDP per capita, as in Figure 1.4. Similar to the prior figure, there is a positive relationship between the capital/output ratio and living standards, except for at the very top end. The combination of Figure 1.3 and 1.4 makes some sense. As the capital stock depends on the amount of investment spending done, then we would expect that the capital stock should be larger (relative to GDP) if the amount of investment spending (relative to GDP) is larger. There is nothing deeper than that to learn from Figure 1.4. The same qualifier as before applies. Do not make any causal interpretations of the relationship in Figure 1.4. This only represents a correlation we should expect to explain in our theory.

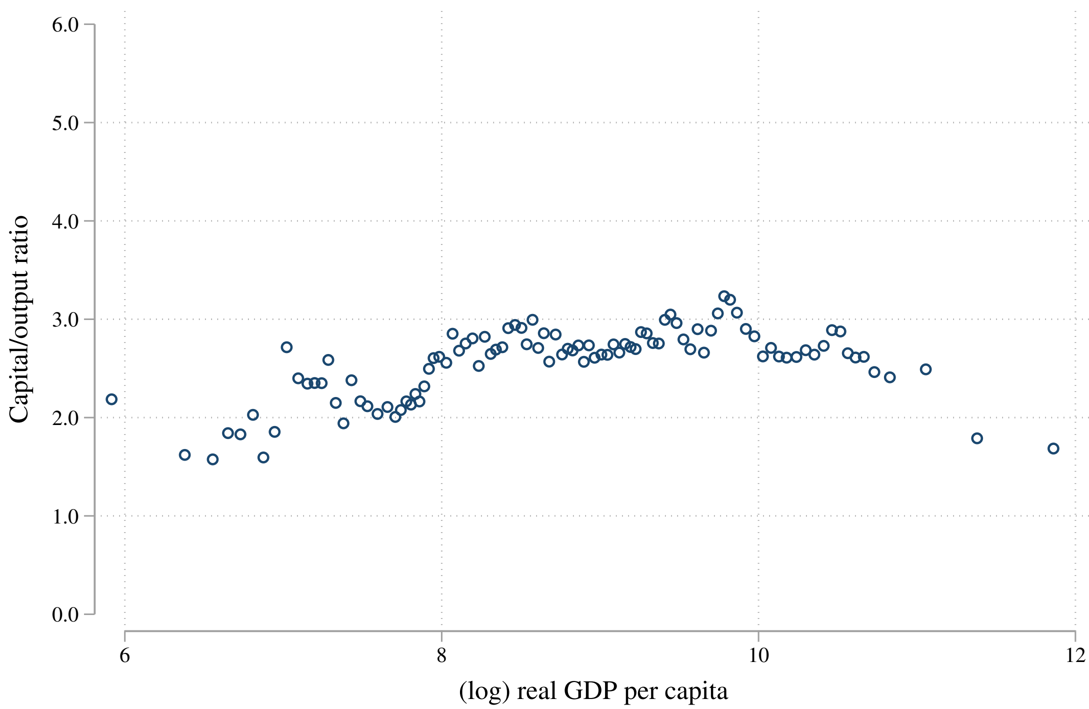

Figure 1.13: Capital/output ratio relative GDP per capita

To give some sense of the types of capital that have been folded into the aggregate capital stock, and establish that these too are stable with respect to output, Figure 1.4 plots the capital/output ratio for the US only. The data for this Figure come straight from the US National Accounts, and not from the Penn World Tables. Thus the actual capital/output ratios will differ.

The aggregate capital/output ratio is around 2.2-2.4 the whole time, but this Figure shows more detail about movement in the capital/output ratio because the vertical axis is compressed compared to Figure 1.4. Regardless, below that aggrete capital/output ratio is plotted the capital/output ratio for different types of capital. This is useful to establish first that the ratios were stable for almost all capital types (excluding intellectual property). Second, this is useful to show that structures, both residential and non-residential, are the dominant types of capital. Structures make up around 75-80% of all capital (by value) in any given year.

All this gives us enough information to establish two more facts

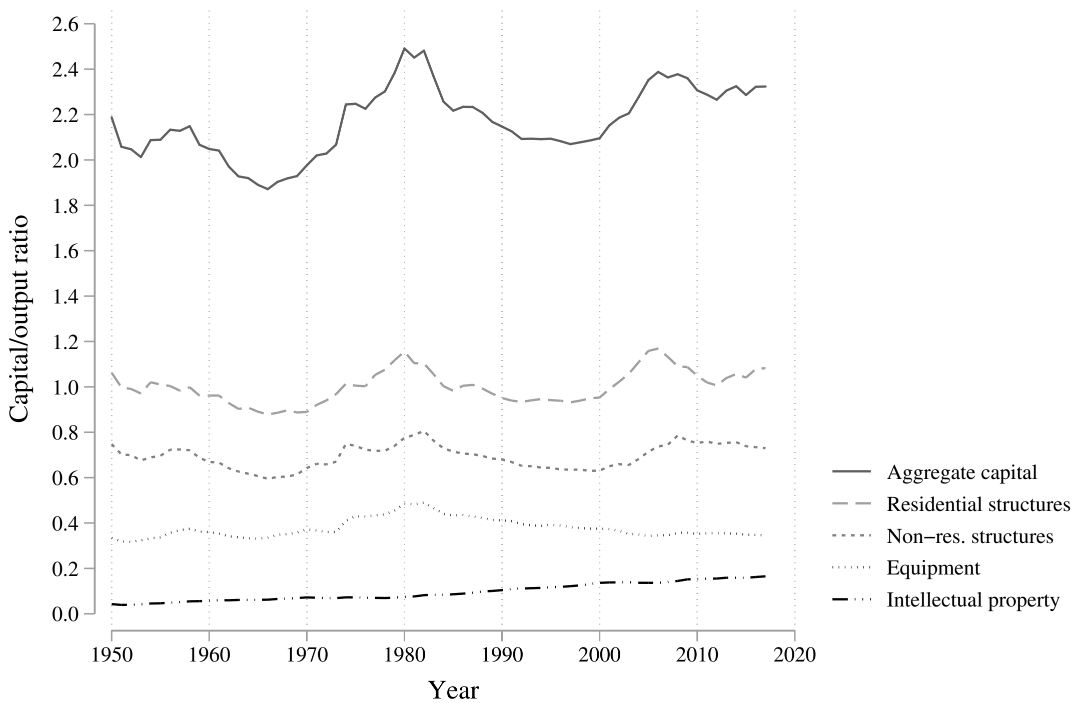

Figure 1.14: Capital/output ratio, by capital type, for the US

regarding economic growth.

Fact 1.7 (Stable capital/output ratios) *The capital/output ratio in developed countries is stable over long periods of time.*

As with several other facts, be careful in not over-interpreting this one. This fact says that for a given country, the capital/output ratio is stable over time. It does not say that capital/output ratios are identical across countries. As you saw in the Figures, these can differ.

Fact 1.8 (Capital/output ratio and GDP per capita) *There is a weak, positive, correlation of the capital/output ratios and the level of GDP per capita.*

1.5 Output and cost shares

The last set of facts we want to establish relate to shares of GDP, similar to what we looked at before. Consumption and investment, however, are part of the expenditure decomposition of GDP. Here we want to look instead at the income decomposition of GDP. Recall that this breaks down GDP into payments to labor (e.g. wages), payments to capital, and economic profits.

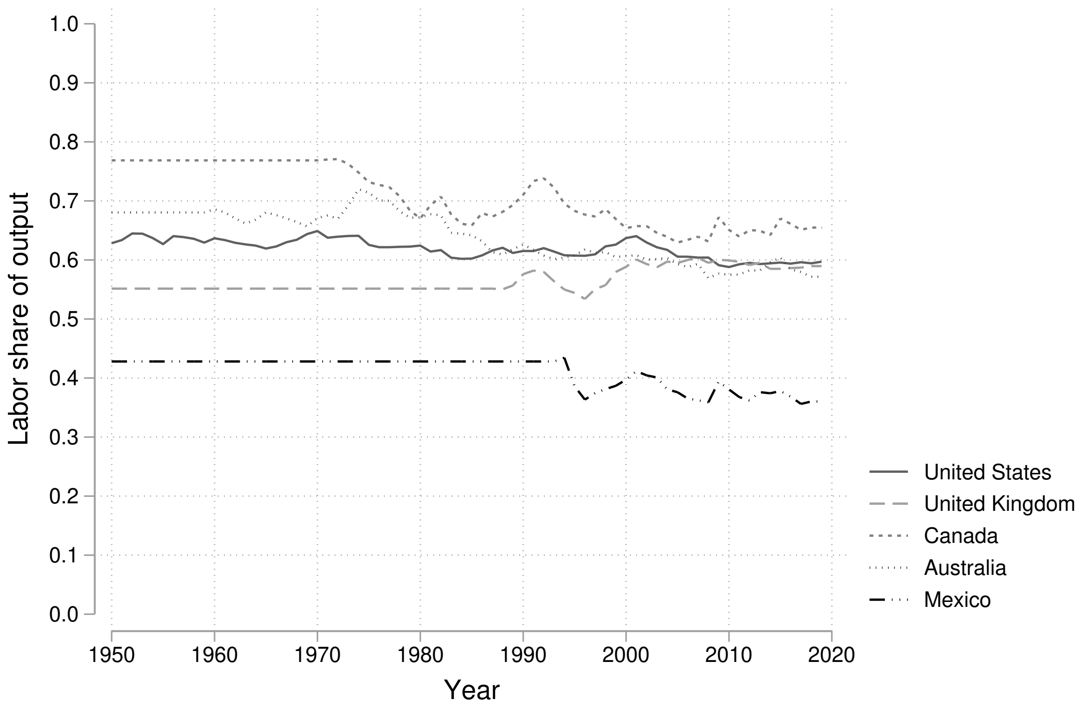

Figure 1.15: Labor share of GDP for selected countries

If we start with labor, Figure 1.5 plots the total payments to labor as a share of GDP. The data availability here is not great, so only the US has information running back to 1950. However, it appears to hold for all the countries in this Figure that the labor share of output is stable over long periods of time. There is a slight tendency for it to decline in Canada and Australia, but in general these appear to stay around 60% of output, except for Mexico, where the share is more like 40%.

One of the reasons Mexico may have a low reported labor share is that much of its labor income gets reported as part of operating surplus because workers run their own small businesses. Research has shown that if one makes adjustments for this, the labor share of countries like Mexico tend to get closer to developed country shares.

Turning to the other set of countries, there is a stimilar story. We again find that a relatively poor country, Nigeria, has a low reported labor share, like Mexico, at least until recently. But for the remainder, the labor share of output hovers around 60% over time. Unlike the prior sections where these countries often had more dispersed experiences, as far as the labor share is concerned they look more like the countries with stable growth rates.

As explained in the preliminaries, we cannot simply assume that one minus the labor share of output is equal to capital's share of output, because we do not have good measures of economic profits across countries. What we can do is look at different data on total *costs* instead, and look at what fractions of those costs are made up of labor payments and capital payments. What we lose here is that

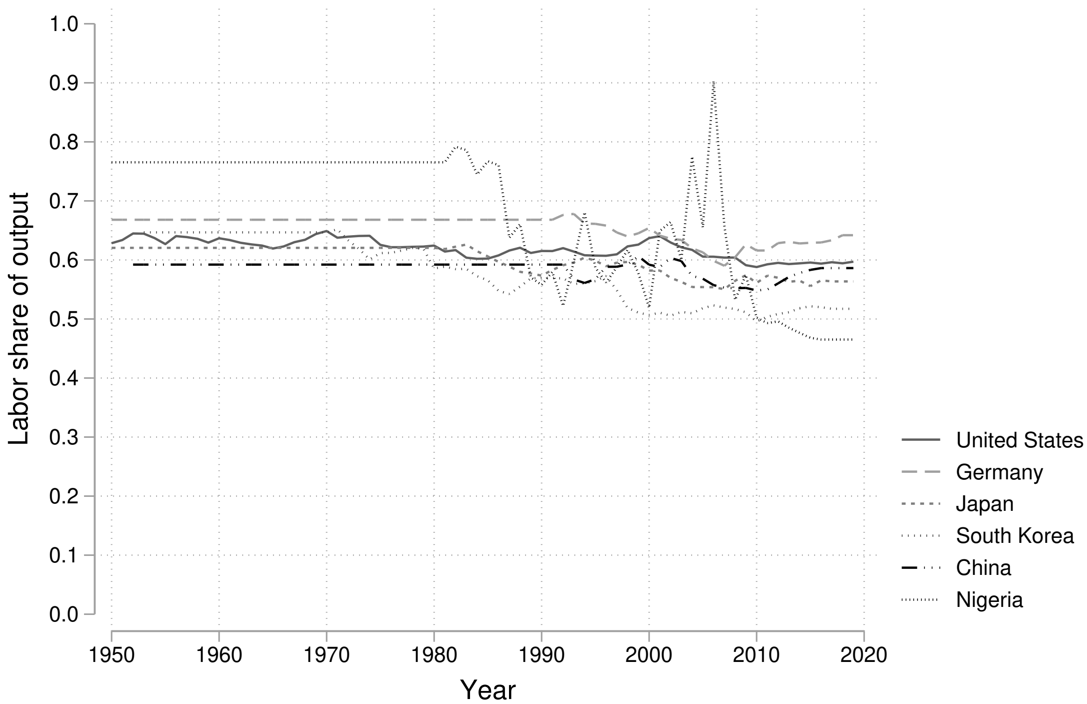

Figure 1.16: Labor share of GDP for selected countries

costs are not equal to GDP, so this won't give us capital's share of GDP. However, this is a small loss, because as we'll see when we turn to the theoretical work, the cost shares of capital and labor are going to be the more relevant pieces of information.

It is possible to infer capital's share of output if we're willing to make some assumptions about rates of return on capital. There is current research on this subject that establishes that the share of output paid to capital has fallen over time while the share going to economic profits has risen.

Figure 1.5 plots labor's share of total costs (i.e. total payments to labor and capital) for several countries. This data is drawn from a different data source, KLEMS, than the national accounts data, and so we are assuming that the relationships in this Figure apply to nations as a whole. This is probably a good assumption for the countries in this Figure, as their statistical agencies cooperate with KLEMS. The drawback of KLEMS is that it only goes back to 1980, and not to 1950.

Regardless, what the Figure shows is that labors share of costs are stable across time. And because this is constructed as a share of costs, we know then that *capital's* share of costs is stable across time for these countries as well. Those labor shares vary across countries, but tend to fall in the range of 50-70%. This is going to serve as our final fact of this chapter.

Fact 1.9 (Stable labor and capital cost shares) *For developed economies the shares of total costs accounted for by labor and capital are stable over*

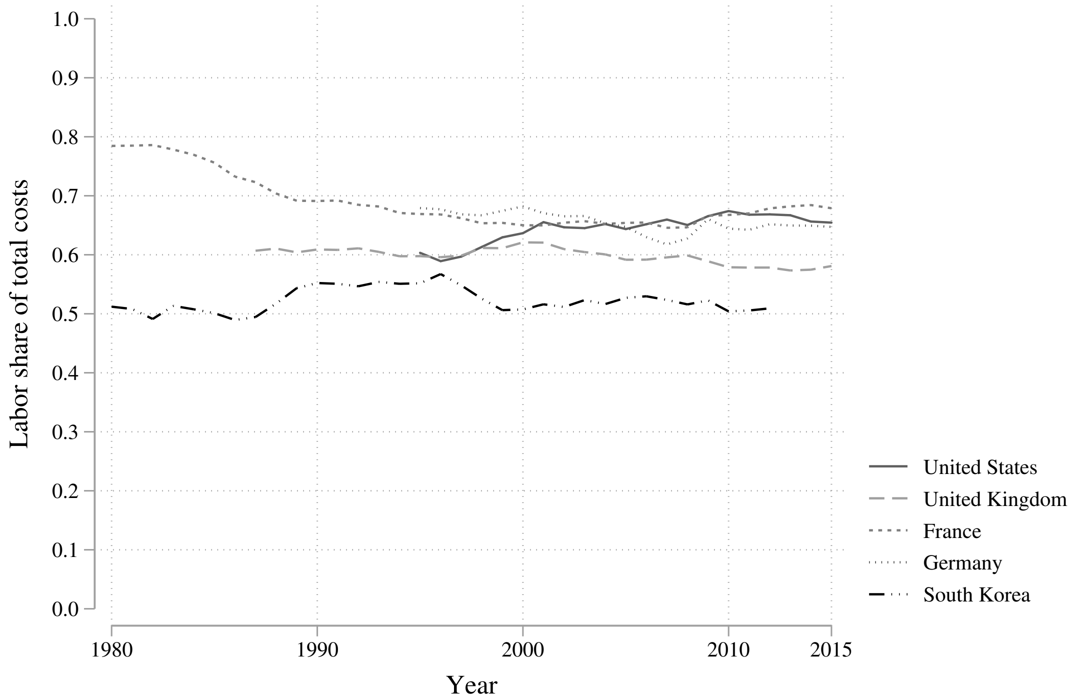

Figure 1.17: Labor share total costs for select countries

time.

As you have probably noted at this point, the overriding theme of the Facts has to do with stability. Stable growth rates, stable investment and consumption shares, stable capital/output ratios, and stable cost shares. We are going to reject certain theories of growth because they cannot match these facts, and we are going to build up our theory of growth with these facts in mind. In the next chapter they will allow us to make some strong statements about what drives economic growth.

2
Production

The facts regarding growth from the prior chapter indicated stability in several key variables over time. The stable facts of growth are often combined under a single umbrella term that is useful as a short-hand way of summarizing all those facts.

Definition 2.1 (Balanced growth path) *A balanced growth path (BGP) describes a situation where output (GDP) per capita, capital per capita, and consumption per capita grow at the same rate, and that rate is stable over time. A BGP, by implication, involves a constant capital/output ratio.*

The term "balanced growth path" is sometimes broadened to include the Kaldor Facts on labors share of output and the interest rate mentioned in the last chapter. I prefer to keep the definition of the term narrow, to avoid confusion.

This definition takes Fact 1.1 (the stability of the growth rate of GDP per capita), and combines that with several others. From Fact 1.7 we know that the capital/output ratio is stable over time, which implies that capital and output must grow at the same rate, and hence output per capita and capital per capita must grow at the same rate. Fact 1.5 said that the consumption share of GDP was constant over time, which means that consumption per capita must grow at the same rate as GDP per capita.

The term BGP straddles the boundary of theory and empirics. A BGP is a hypothetical construction in which the growth rate of GDP per capita, capital per capita, and consumption per capita are all exactly identical and perfectly constant over time. It therefore also has a perfectly constant capital/output ratio. No country even truly meets these conditions, but given what we saw in the prior chapter the BGP is a useful abstraction of what occurred in most developed nations over time. We therefore often say that a particular country appears to be "on" a BGP because its data roughly conforms to the definition.

Because we see that many economies tend to be "on" a BGP (e.g. the US or UK), when we write down a theory of economic growth we're going to insist it predicts a BGP will result. If it does, then we know it can replicate the broad features of economic growth we saw in the data. Given that, we'll be able to infer some things about how economic growth must work by seeing what assumptions are necessary for our theory to have a BGP as an outcome. As we also see that a number of economies are "off" a BGP (e.g. Japan or South Korea), we'll also want our theory to be able to describe growth in those situations, and why in those cases the countries all appear to move towards a BGP.

2.1 Accounting and notation

Before we get started theorizing, we need to establish some ideas about what we're accounting for in the sense of what GDP is counting and what GDP growth is measuring.

Definition 2.2 (Nominal GDP) *Nominal GDP is denoted*

$$PY = \sum_{j \in J} p_j c_j \tag{2.1}$$

where J are the set of final products or goods or services (or whatever) that are sold in the economy, p_j is the price of an individual product, and c_j are the real units of that product. P denotes a price index for GDP and Y denotes real GDP.

GDP measures *final* goods and services, which are just those used to consume or used to produce in the future. This excludes intermediate goods (e.g. energy, raw materials) that are purchased and then turned into final goods or services. Gross output is a different concept that is the sum of GDP and spending on intermediate goods and services.

From this you can see that real GDP, Y, is something that is implicit. We cannot measure it directly, the way we can with nominal GDP.

Definition 2.3 (Real GDP growth) *Holding prices constant, the change over time in real GDP is*

$$dY = \sum_{j \in J} \frac{p_j}{P} dc_j.$$

where dc_j is the change in real purchases of each product. This means the growth rate of real GDP is

$$g_Y = \frac{dY}{Y} = \sum_{j \in J} \frac{p_j c_j}{PY} \frac{dc_j}{c_j}. \tag{2.2}$$

The ratios $p_j c_j / PY$ are nominal expenditure shares and dc_j / c_j is the growth rate of real purchases of each product.

That's a lot. It just says that the growth rate of real GDP is a combination of growth in real consumption or purchases of each product in the economy, weighted by their expenditure share. We're not assuming that GDP is just a single monolithic "thing", it's made up of multiple goods and services. We will work almost exclusively with things like g_Y and talk about GDP growth as one thing, but you should keep in mind that this encompasses plenty of changes within it. Saying that g_Y is equal to 0.02 does not mean each and every product grew at 2%. Some will be higher, some lower.

GDP is an accounting concept, and the two primary ways of adding up GDP have to reconcile. You are probably familiar with the expenditure side from intermediate macro classes, which say something like $Y = C + I + G + (X - M)$.

Definition 2.4 (Expenditure approach to GDP) *The expenditure approach to GDP classifies the J transactions by the type of product, and in accounting terms is*

$$PY + \sum_{j \in M} p_j c_j = \sum_{j \in C} p_j c_j + \sum_{j \in I} p_j c_j + \sum_{j \in G} p_j c_j + \sum_{j \in X} p_j c_j \tag{2.3}$$

where M (imports), C (consumption), I (new capital goods), G (government purchases), and X (exports) are sets of products that are each a subset of the J overall set of products.

Notice that I put imports on the "left". That is because the combination of GDP (gross *domestic* product) and imports are the combined set of products that are available for purchase in the economy. The expenditures on those products are classified either as consumption, new capital goods, government, or exports. The choice of this classification is arbitrary (e.g. education is consumption).

From the income side every transaction is income to someone or some firm (who then pays workers and owners of capital)

Economic profits are payments for transactions over and above their cost of production in the same period of time. They might reflect market power (e.g. a monopolist) or might reflect payments for fixed costs.

Definition 2.5 (Income approach to GDP) *The income approach to GDP classifies total GDP as*

$$PY = \Pi + W + RK \tag{2.4}$$

where Π *are economic profits (not accounting ones), W is the total wage bill, and RK are payments to owners of capital.*

The income approach does not assign each transaction to one of these three, but breaks each transaction (notionally) down into three parts.

Those income classifications are used to write things in terms of shares as in

Measuring W is not too hard, but assigning the remainder $PY - W$ to either capital or economic profits is notoriously hard, as firms do not tend to think of their own income from this perspective.

Definition 2.6 (GDP Shares) *The income shares of GDP are*

- *Labor:* $s_L = W/PY$
- *Capital:* $s_K = RK/PY$
- *Profits:* $s_\Pi = \Pi/PY$

Note that these shares are out of nominal GDP.

Last, let's clean up and be clear on some notation regarding growth rates and changes.

See A.1 if this is not familiar.

Definition 2.7 (Notation for growth and change) *For any variable X*

1. *dX is the differential in X, or the change over time. dX is shorthand for $\partial X/\partial t dt$.*

2. *g_X is the growth rate of X and is defined as $g_X = dX/X$.*

3. *$d \ln X \approx g_X$.*

These definitions are all in continuous time, in the sense that I'm taking derivatives/differentials. But I'll freely use approximations of these to discrete time concepts. For example $dX \approx X_{t+1} - X_t$ and $g_X \approx (X_{t+1} - X_t)/X_t$. When we get to specific discrete time models we will be more precise about things.

2.2 *Capital accumulation*

One part of GDP is investment purchases, which are *new* capital goods. Thus investment purchases add to the capital stock. At the same time, we think/know/assume that a certain amount of the existing capital stock depreciates every period. It is common to assert that this takes place at a constant rate (e.g. 5% of the capital stock depreciates every year), and this rate is denoted by δ.

We can put together a differential equation for the capital stock using all this information as

$$dK = I - \delta K \tag{2.5}$$

which says that the change in the capital stock (dK) is equal to the amount of investment done at a given point in time minus the amount of the capital stock that depreciates. This equation is very sloppy with time subscripts, which is in part due to the fact that we're trying to combine a stock (capital) with a flow (investment). A more tedious but more accurate way to describe this capital accumulation would be

$$dK(t, t+1) = I(t) - \delta K(t) \tag{2.6}$$

where $dK(t, t+1)$ is the change in the capital stock from period t to period $t+1$, say January 1st 2018 to January 1st 2019. $I(t)$ is the investment purchases done during period t, say during 2018. $\delta K(t)$ is the amount of the capital stock at time t, say on January 1st 2018, that depreciated from t to $t+1$, say during 2018. As I said, tracking the timing is tedious.

The notation for derivatives and growth rates throughout the book is explained in A.1.

In the notation regarding logs and growth rates from above, if we divide both sides of (2.5) by K, then we can say that the growth rate of the capital stock is

$$g_K = \frac{I}{K} - \delta \tag{2.7}$$

2.3 *Production*

To begin the theory, we need a description of how GDP is produced. This is where we have to introduce the first assumptions about how the economy works, although thanks to some recent work these assumptions are less demanding than one might guess.[1]

[1] David Baqaee and Emmanuel Farhi. A Short Note on Aggregating Productivity. NBER Working Papers 25688, National Bureau of Economic Research, Inc, March 2019

Assumption 2.1 (Production of output) *Production has the following features:*

- *There are two factors of production, capital (K) and labor (L) which are used as part of the production process.*
- *There are an arbitrary number of individual production units (e.g. establishments or firms) which use capital and labor, and may also use intermediate inputs purchased from other production units.*
- *The cost functions (i.e. the cost to the unit to produce one unit of output) of those production units are constant returns to scale with respect to capital and labor.*
- *The total capital used by the production units equals the aggregate capital stock, $\sum_i K_i = K$, and the total labor used by production units equals the supply of workers, $\sum_i L_i = L$.*

It is worth noting what is *not* assumed here. We do not have to assume production units produce homogenous goods. We do not have to assume competition, much less perfect competition, across production units. Those units could be monopolists, or some of them could be part of an oligopoly, while others may in fact be competing fiercely with one another. We do not have to assume that those units face the same cost of labor or capital, either. In short, there can be any arbitrary markup, or wedge between the price and marginal cost for each production unit.

The assumption that there is just one type of capital, and one type of labor, is restrictive. It is also not necessary, but allowing for multiple types of either or both in the analysis would add tedious algebra at this point without much insight. See A.10

Given the assumptions, it can be shown that

$$d \ln Y = \epsilon_K d \ln K + \epsilon_L d \ln L + \epsilon_L d \ln A \tag{2.8}$$

where $d \ln x$ refers to the change in (log) of variable x, and ϵ_K and ϵ_L are the elasticity of output with respect to capital and labor, respectively. A here refers to the level of total factor productivity (or productivity, for short). The change in log productivity can be broken down further into terms related to unit-specific technological change and the allocation of factors across those units, but for the moment there is no need to specify that breakdown.

The equation for $d \ln Y$ seems simple, but showing that it in fact holds for an economy with the few assumptions shown above is anything but trivial.

Equation (2.8) simply says that the change in (log) output - which note is equivalent to a percent change for small enough changes - is a combination of the percent change in the capital stock (weighted by ϵ_K), the percent change in labor (weighted by ϵ_L), and the percent change in productivity.[2] The fact that $d \ln A$ is multiplied by ϵ_L implies that productivity changes are also weighted by the elasticity with respect to labor, or that productivity is "labor-augmenting"..

[2] H. Uzawa. Neutral Inventions and the Stability of Growth Equilibrium. *The Review of Economic Studies*, 28(2):117–124, 02 1961; and Charles I. Jones and Dean Scrimgeour. A new proof of Uzawa's steady-state growth theorem. *The Review of Economics and Statistics*, 90(1):180–182, 2008

This does not mean productivity growth can only be labor-augmenting, or that it has to be labor-augmenting. The Uzawa Theorem says that we need to be able to express productivity growth as labor-augmenting for our model to have a BGP. But that does not preclude capital-augmenting or disembodied productivity growth. See A.11

The above equation holds for any changes in K, L, or A, but we are particularly interested in changes over time. If we simply divide both sides of equation (2.8) by dt, then we will have an expression for the change in log output for a change in time, or in other words, the growth rate. Using the notation introduced in the preliminary chapter,

$$g_Y = \epsilon_K g_K + \epsilon_L g_L + \epsilon_L g_A \tag{2.9}$$

The elasticities will be relevant for our analysis of growth, and they have some properties that are very useful. I'm going to state these properties as assumptions, because in these notes I will not be showing how to derive these properties, but they are implications of the assumptions about production.

Assumption 2.2 (Factor elasticities) *Factor elasticities have the following properties:*

- $\epsilon_K + \epsilon_L = 1$, *which is a result of the constant returns to scale assumption for individual units of production.*
- ϵ_K *is equal to input/output weighted share of costs accounted for by capital, and* ϵ_L *is equal to the input/output weighted share of costs accounted for by labor*

The first property simply says that if both capital and labor grow at say, 10%, then output grows at 10% as well (ignoring TFP growth). Thus there are neither increasing returns to capital and labor at the aggregate level (e.g. output grows by more than 10% if factors grow by 10%) or decreasing returns at the aggregate level (e.g. output grows by less than 10% if factors grow by 10%).

The crude cost share of labor reported in Fact 1.9 is constant over time, but that is not necessarily the right measure of ϵ_L. A.7 describes how to construct the proper input/output weighted cost share for labor, and shows that it is also roughly constant and similar in magnitude to the crude cost share.

Neither of the properties in Assumption 2.2 imply that the elasticities must be constant over time, or insensitive to the amount of capital and labor supplied in the economy. However, if we calculate those input/output weighted shares of costs, they are roughly constant over time. The fact that these are constant over time will allow us to make several strong assertions about how economic growth works, so let's set this off as a separate assumption.

Assumption 2.3 (Stable factor elasticities) *The factor elasticities* ϵ_K *and* ϵ_L *are stable over time.*

Given all these assumptions, we can do some manipulation of (2.9) to arrive at an equation for what drives growth in output per capita (as opposed to total output). Subtract $\epsilon_K g_Y$ from both sides of that equation, and you have

$$(1-\epsilon_K)g_Y = \epsilon_K(g_K - g_Y) + \epsilon_L g_L + \epsilon_L g_A. \tag{2.10}$$

Divide both sides of this by $(1-\epsilon_K)$ and apply the assumption that $\epsilon_K + \epsilon_L = 1$, and you arrive at

$$g_Y = \frac{\epsilon_K}{\epsilon_L}(g_K - g_Y) + g_L + g_A. \tag{2.11}$$

We can manipulate this equation to the following,

$$g_y = \frac{\epsilon_K}{\epsilon_L} g_{K/Y} + g_A, \tag{2.12}$$

where $g_{K/Y} = g_K - g_Y$ is the growth rate of the capital/output ratio, to be clear. This shows that the growth rate of GDP per capita depends on a combination of the growth rate of the capital/output ratio and the growth rate of productivity. This relationship holds whether an economy is on a BGP or not. It is simply a re-arrangement of the production relationship we started with in (2.9).

There is one interesting note to make about the role of capital accumulation and economic growth. Equation (2.12) shows that the growth rate depends on how fast the capital stock grows *relative* to output, $g_{K/Y}$. This is because capital is itself a function of how much we produce. We would expect the capital stock to increase along with output no matter what. This expression indicates that growth in the capital stock only contributes to growth in output per capita to the extent that it grows *faster* than output.

It is plausible to connect the elasticities ϵ_L and ϵ_K to the income shares s_L and s_K, but it is important to remember that these are distinct concepts and there is no strict relationship of the shares and elasticities except under very restrictive assumptions about competition.

A.9 describes how shares and elasticities are related for a given firm.

2.4 *Balanced and transitional growth*

The above section provided a description of the growth rate of output per capita, g_y, that is valid no matter whether an economy is on or off a BGP. To proceed we're going to look first at what this production structure implies about growth for countries that are in fact on a BGP, and then we'll turn to countries that are not.

We can step in immediately here and draw one immediate conclusion about growth on a BGP.

Conclusion 2.1 (Source of growth on a BGP) *Given that the capital/output ratio is constant on a BGP ($g_{K/Y} = 0$), it is the case that the* *growth rate of output per capita on a BGP depends only on productivity growth, $g_y^{BGP} = g_A$.*

Growth is denoted as g_y^{BGP} to make clear that this is the growth rate along a BGP. The proposition does not say that the growth rate is always equal to g_A, as economies may not be on a BGP (e.g. South Korea or Japan). A second aside is that this means $g_Y^{BGP} = g_A + g_L$, or that growth in GDP on a BGP is the sum of productivity and population growth.

This is strong statement about what drives growth in the long run. The growth rate on a balanced growth path is proportional to total factor productivity growth, and does not depend on the particulars of how capital accumulates relative to output. When you look back at Figure 1.1, the stable growth you see for developed economies is a result of growth in productivity. This doesn't mean that capital accumulation doesn't occur (capital must be growing at the same rate as output), or that capital accumulation doesn't matter for the *level* of living standards (which we'll get to), only that it doesn't matter for the growth rate of output per capita in the long run.

This result has a corollary implied by the stability of cost shares and the growth rate.

Conclusion 2.2 (Stable productivity growth on a BGP) *Given that the cost share of labor is constant over time, and g_y^{BGP} is constant on a BGP, it must be the case that the growth rate of productivity, g_A, is also constant on a BGP.*

There is no way to test this proposition. Productivity cannot be observed directly. We could assume that production happens according to equation (2.9), and use data on capital, labor, and cost shares to back out a measure of g_A. But then we'd be assuming this proposition was true.

When we turn to studying productivity growth in more detail, this proposition will be an important fact that we will try to match with the theory. It will lead to some interesting implications for what drives growth in productivity in the long run. For the moment just keep in mind that the stability of growth, the capital/output ratio, and the cost shares impy that productivity growth must be stable as well.

Knowing how growth works on a BGP, we can say something about the variation we see in growth rates in the data.

Conclusion 2.3 (Growth differences) *Given that g_A appears similar across countries, observed differences in growth rates must be due to differences in $g_{K/Y}$.*

This comes from fact 1.2, which implies that g_A is similar across developed countries. Knowing that, the conclusion follows from examining equation (2.12). Variation in growth rates is due to some countries being "off" their BGP ($g_{K/Y} \neq 0$) while some are "on" their BGP ($g_{K/Y} = 0$). Since $g_{K/Y}$ plays such an important role, we're going to give it a specific name.

Definition 2.8 (Transitional growth) *Transitional growth is equal to the growth due to changes in the capital/output ratio, $g_{K/Y}$.*

2.5 *Accounting for growth*

We know from 2.12 that growth is either via $g_{K/Y}$ or g_A. How important are each of these components in the data? Figure 2.5 plots the growth rate of GDP per capita, g_y, as well as the growth rate of g_A over time for a set of relatively rich countries (GDP per capita above arond $3,000). The growth rates of both are 10-year forward looking averages (e.g. the data for 1980 is the average growth rate between 1980 and 1990). This figure also "bins" up data from all the countries, so what you are seeing in any given year is something like the average outcome across the countries included.

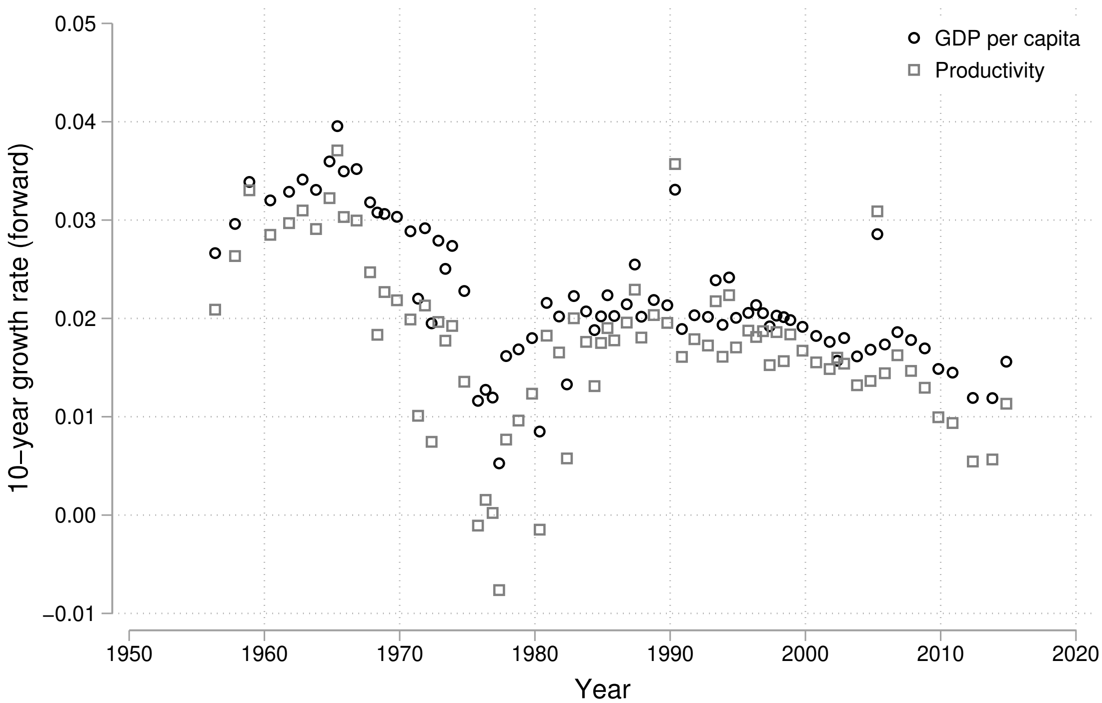

Figure 2.1: Comparison of g_y and g_A in relatively rich countries over time

The most relevant thing here is that the two series track one another closely. That is, most of growth in GDP per capita can be accounted for by growth in productivity. We know that's true along a BGP, but this figure also indicates that the contribution of transitory growth when countries are off the BGP tends to be small. The size of $\frac{\epsilon_K}{\epsilon_L} g_{K/Y}$ is just the gap between the two different series.

You can see that more clearly in Figure 2.5, which plots $\frac{\epsilon_K}{\epsilon_L} g_{K/Y}$ in the same manner, meaning 10-year growth rates and binned up across countries. It's not that transitory growth is zero, at times it averaged over 1% as many countries were converging to new BGP's

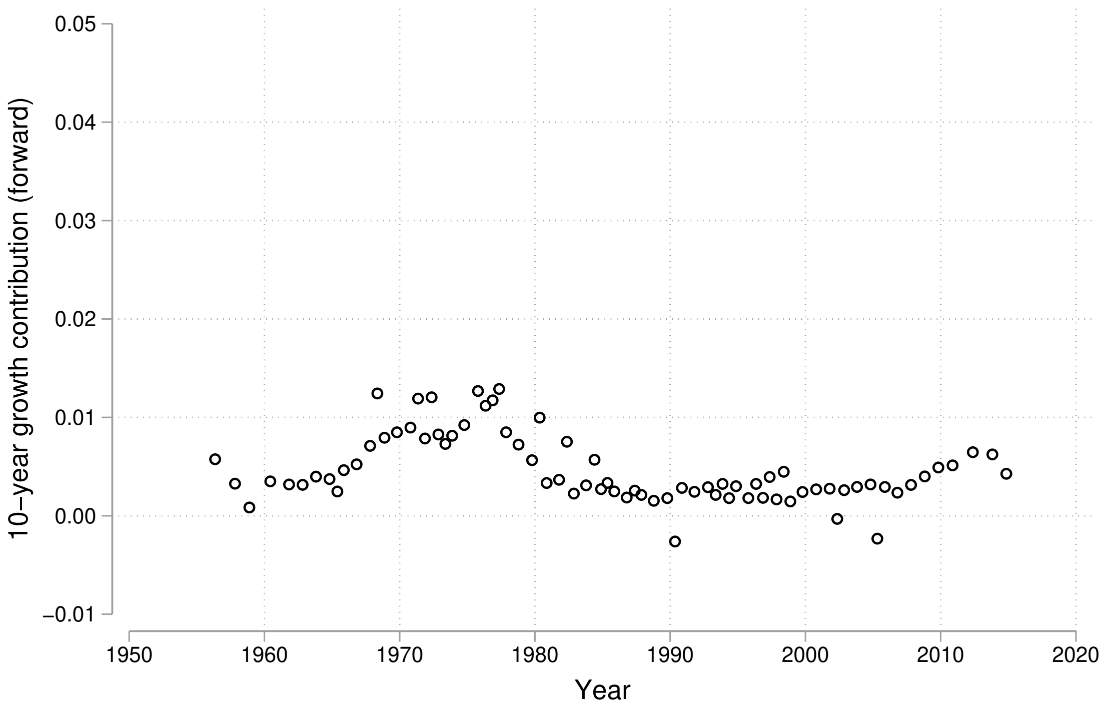

Figure 2.2: Contribution of $\epsilon_K/\epsilon_L g_{K/Y}$ in relatively rich countries over time

in the 1970s, and it remains positive on average even out to the 2010's, indicating some shifting between BGPs. But just in terms of size it isn't large, leading us to this fact,

Fact 2.1 (Importance of prouctivity growth) *Most g_y is due to productivity growth, g_A, and transitional growth due to $g_{K/Y}$ tends to be small.*

This doesn't mean capital growth is unimportant. What these figures don't show is that capital accumulation remains important for assuring stability *around* the BGP. No, capital accumulation doesn't drive long-run growth, but it does help ensure we stay on or close to the BGP.

3
The Solow model

The facts regarding growth from the prior chapters indicated stability in several key variables over time. We were already able to conclude several things about the BGP just from the structure of production, but those did not provide any information about why economies tend towards a BGP. The Solow model is a theoretical structure that describes that tendency towards stability.[1]

[1] Robert M. Solow. A contribution to the theory of economic growth. *The Quarterly Journal of Economics*, 70(1):pp. 65–94, 1956

3.1 Transitional growth

The essential point of the Solow model is that transitional growth is temporary, and that it drives the economy towards the BGP. To see this, let's first subsitute in what we know about the growth rate of the capital/output ratio

$$g_{K/Y} = g_K - g_Y = \epsilon_L g_K - \epsilon_L g_L - \epsilon_L g_A, \tag{3.1}$$

where the second equality follows from the definition of g_y in Equation (2.9), and the assumption about the sum of the elasticities.

To proceed we need to draw in what we know about the growth rate of the capital stock from Equation (2.7). Using that, we can say that

$$g_K = \frac{I}{Y}\frac{Y}{K} - \delta, \tag{3.2}$$

where the first fraction on the right is the investment share of GDP and the second is the (inverse of) the capital/output ratio. We know from Fact 1.5 that the investment share is stable over time. Let's introduce a little new notation here

Traditionally, the variable s was used for I/Y, and was understood to mean "savings". I don't like this terminology, as "savings" has more common usages that do not mean anything like the investment share of GDP. Thus the interpretation of s here as "share of GDP".

$$s_I \equiv \frac{I}{Y} \tag{3.3}$$

which should be read as the share (s) of GDP that goes to investment (I). Given the data, we're going to assume that s_I is a constant, at least for the time being.

Combine the common terms involving g_K and the definition of s_I with Equation (3.1) and we arrive at

$$g_{K/Y} = \epsilon_L \left(s_I \frac{Y}{K} - \delta - g_L - g_A \right). \tag{3.4}$$

When the capital/output ratio is high, then the growth rate of capital is low because the absolute amount of investment spending done ($s_I Y$) is small relative to the amount of capital that is depreciating. When the capital/output ratio is small, the growth rate is high because the investment being done is large relative to depreciation. Beyond that, the growth rate of the capital/output ratio depends negatively on g_L and g_A because these things create growth in Y over and above the growth attributable to additional K.

This gives us the ability to draw several conclusions.

Conclusion 3.1 (Inverse relationship of growth and capital/output) *Transitional growth, $g_{K/Y}$, is inversely related to the size of the capital/output ratio, K/Y.*

We just established above the logic of why this conclusion holds. This follows from examining the equation (3.4). Furthermore, if you take the limit of this growth rate as K/Y goes to zero, then $g_{K/Y}$ goes to infinity. We're not particularly concerned with this as a real-life case, but it is important because it establishes that $g_{K/Y}$ is definitely *positive* when the capital/output ratio is small. If you take the limit of the growth rate as K/Y goes to infinity (again, not a real-life case but useful to illustrate things) then $g_{K/Y} \approx -\epsilon_L(\delta + g_L + g_A) < 0$, a negative growth rate. This leads to a new conclusion.

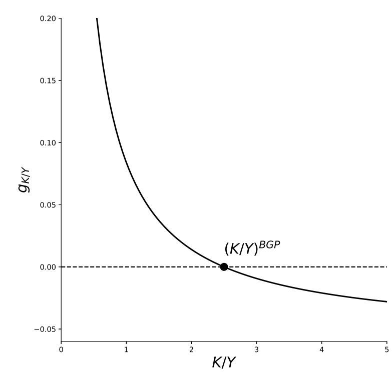

Figure 3.1: The inverse relationship of growth in the capital/output ratio and the level of the capital/output ratio.

Conclusion 3.2 (Existence of a BGP) *There must be a level of the capital/output ratio, call it $(K/Y)^{BGP}$, at which $g_{K/Y} = 0$.*

Because as we just established $g_{K/Y}$ smoothly decreases with K/Y, and eventually turns negative, it has to be the case that at some point $g_{K/Y} = 0$. This is most obvious when looking at Figure 3.1.

Variation in growth rates is due to variation in $g_{K/Y}$, and variation in $g_{K/Y}$ is due to variation in K/Y itself. But if you examine Figure 3.1 or equation (3.4) you can see that $g_{K/Y}$ will evolve over time as the capital/output ratio changes. We can make the following conclusion.

As we'll see in the next section, it is really the gap between K/Y and $(K/Y)^{BGP}$ that dictates $g_{K/Y}$.

Conclusion 3.3 (Transitional growth is temporary) *The capital/output ratio is stable around $(K/Y)^{BGP}$, meaning that no matter where the K/Y ratio starts, it will always end up at $(K/Y)^{BGP}$, and $g_{K/Y}$ will always end up at zero in the long run.*

This statement is a direct result of the negative relationship between $g_{K/Y}$ and the level of K/Y from conclusion 3.1. That negative relationship ensures that K/Y is always being pushed towards the BGP level. When $K/Y < (K/Y)^{BGP}$, $g_{K/Y} > 0$, and hence the capital output ratio moves closer to to $(K/Y)^{BGP}$, and then the process continues until $g_{K/Y} \Rightarrow 0$. You can run the same logic in reverse if the capital/output ratio is above the BGP level.

Conclusion 3.3 explains fact 1.3, which said that growth rate differences were temporary. This showed up most clearly when looking at countries like Japan, South Korea, and Germany, who each had very high growth rates for several decades but that high growth rate eventually dissipated as they reached their BGP. What conclusion 2.3 tells us was that their abnormally high growth rates were due to transitional growth, $g_{K/Y}$. And conclusion 3.1 suggests that this high transitional growth was because those countries had capital/output ratios below their balanced growth path value of $(K/Y)^{BGP}$.

3.2 *The level of the BGP*

What is apparent is that the level of K/Y is important in determining the size of transitional growth, and how much it differs from the BGP value of $(K/Y)^{BGP}$. It is important to keep in mind that the value of $(K/Y)^{BGP}$ need not be identical across countries, as was seen in Figures 1.4 and 1.4. Transitional growth arises in a particular country because its actual capital/output ratio differs from its own, specific $(K/Y)^{BGP}$.

So what determines that specific BGP capital/output ratio? Go back to equation (3.4), set $g_{K/Y} = 0$, and solve.

Conclusion 3.4 (Level of capital/output on a BGP) *On a balanced growth path, the level of the capital/output ratio is*

$$\left(\frac{K}{Y}\right)^{BGP} = \frac{s_I}{g_A + g_L + \delta}. \tag{3.5}$$

The capital/output ratio on a BGP depends on how fast capital accumulates (the investment share) relative to how fast capital depreciates (δ) combined with how fast output grows (TFP and population growth). The intuition behind the investment share is straightforward. If an economy commits a greater share of its output to building capital goods, it follows that the stock of capital will be large relative to output. And the fact that the capital/output ratio depends negatively on the depreciation rate makes sense as well. The faster that capital breaks down, the smaller will be the capital stock relative to output. Population growth, g_L, is negatively related to

Putting numbers on this equation, s_I is around 0.20 for developed countries, and g_A is around 0.018, based on conclusion 2.1. If depreciation runs around 0.05, and the population growth rate is around 0.01, then the capital/output ratio on the BGP should be about 2.56, which is a decent approximation of what was shown in the last chapter.

the capital/output ratio because expanding the number of people (workers) expands output without expanding the capital stock, and hence it makes output larger relative to the capital stock.

What is not as obvious is why the growth rate of productivity has a negative effect on the capital/output ratio, but this is similar to the effect of population. Higher productivity makes output higher without affecting the size of the capital stock, so faster productivity growth decreases the capital/output ratio.

Regardless, transitional growth is related to how far away the actual capital/output ratio is from the (country-specific) $(K/Y)^{BGP}$ described in equation (3.5). We know from conclusion 3.1 that the farther K/Y is below $(K/Y)^{BGP}$, the larger is transitional growth $g_{K/Y}$. What this doesn't tell us is *why* there are countries that have a capital/output ratio below their BGP level, but we will take this up in the next section.

And symmetrically, the farther K/Y is above $(K/Y)^{BGP}$, the smaller is transitional growth.

Before moving on, it is worth establishing the close connection of the capital/output ratio and the level of output per capita on a BGP. Intuitively, it makes sense that the level of output per capita should depend positively on the capital/output ratio, but we can make the connection exact.

Go back to equation (2.12), and multiply both sides through by dt to recover

$$d \ln y = \frac{\epsilon_K}{\epsilon_L} d \ln K/Y + d \ln A, \tag{3.6}$$

which is just a manipulation of the original equation (2.8) describing production of GDP. Integrate both sides of this equation, and you get

$$\ln y = \frac{\epsilon_K}{\epsilon_L} \ln K/Y + \ln A + C, \tag{3.7}$$

This integration relies on the crucial assumption 2.3 that ϵ_K and ϵ_L are constant, which is consistent with the data. But if those elasticities changed as K/Y or A changed, then this would not be as simple a process.

where C absorbs all the constants of integration that arise. This equation is effectively a production function, in log form. It holds for the economy whether it is on a BGP or not. But we can use it to describe the path that output per capita follows on a BGP.

The capital/output ratio is constant on a BGP, and described by (3.5). We also know that on a BGP productivity grows at a constant rate, meaning that $\ln A$ is a function of time. Put that all together and we have

Constant growth of A can be written mathmetically as $\ln A(t) = \ln A(0) + g_A t$, where $A(0)$ is some arbitrary baseline level of productivity.

$$\ln y(t)^{BGP} = \frac{\epsilon_K}{\epsilon_L} \ln \left(\frac{s_I}{g_A + g_L + \delta} \right) + \ln A(0) + g_A t. \tag{3.8}$$

Note that I've indicated this is $\ln y(t)^{BGP}$, as output per capita on a BGP is a function of time, given that productivity is growing over time. I've also set the value of C to zero, without loss of generality, as we can adjust the baseline level of productivity, $A(0)$ to incorporate that.

Equation (3.8) tells us that a country with a higher BGP capital/output ratio should have a higher level of output per capita, holding productivity constant. This conforms to the rough positive relationship between capital/output and log GDP per capita demonstrated in Figure 1.4 and Fact 1.8, although that data does not meet the "holding productivity constant" standard. As the capital/output ratio depends directly on the investment share in GDP, this also means our description of a BGP conforms to the rough positive relationship of investment share and log GDP per capita in Figure 1.3 and Fact 1.6.

3.3 *The Golden Rule*

The BGP's imply something about consumption.[2] No matter what

$$c(t)^{BGP} = (1 - s_I)y(t)^{BGP}. \tag{3.9}$$

[2] Edmund Phelps. The golden rule of accumulation: A fable for growthmen. *The American Economic Review*, 51(4): 638–643, 1961

Given 3.8 you can see there is a trade-off here. The higher is s_I, the higher is GDP per capita along the BGP. But the higher is s_I, the smaller fraction of that gets to be consumed. From the consumption perspective, it is not ideal to maximize s_I (e.g. set it equal to one).

We can set a useful benchmark here, which is finding the s_I that maximizes $c(t)^{BGP}$. Take the derivative of consumption with respect to s_I and set equal to zero, and you get

$$-y(t)^{BGP} + (1 - s_I)\frac{\partial y(t)^{BGP}}{\partial s_I} = 0 \tag{3.10}$$

which you can resolve to

$$\frac{\partial y(t)^{BGP}}{\partial s_I}\frac{s_I}{y(t)^{BGP}} = \frac{s_I}{1 - s_I}. \tag{3.11}$$

In other words, the elasticity of GDP per capita along the BGP with respect to s_I is going to tell us about how to set s_I. If GDP per capita is very sensitive to s_I, then it pays to have a high s_I, and vice versa. From 3.8 we can read off this elasticity quite easily which gives us

Conclusion 3.5 (The Golden Rule) *The "Golden Rule" level of s_I that maximizes c_t^{BGP} is*

$$s_I^{GR} = \epsilon_K. \tag{3.12}$$

There is nothing about this that is "optimal" in the sense of maximizing utility or anything. This is purely a mental benchmark.

This implies that the thing that maximizes consumption along a BGP is spending the fraction of income equal to the elasticity of output with respect to capital on building new capital. It ties the importance of capital to the amount that is invested in new capital.

From Figure 1.5 we know that labor's share of costs is between 0.5 and 0.7, which should be an estimate of ϵ_L, and that makes ϵ_K somewhere between 0.3 and 0.5. From Figure 1.3 is looks like observed s_I

You'll see in Chapter 7 that the only viable accumulation rate is *below* s_I^{GR} once we allow for individuals to optimize their consumption/accumulation choice. It happens because people are impatient.

are all below 0.3, so in that sense the data tell us that economies are below the Golden Rule accumulation rate.

3.4 *Transitional growth*

Last, compare this expression for the level of output per capita along a BGP to the expression for actual output per capita from (3.7), presented again here with explicit time subscripts,

$$\ln y(t) = \frac{\epsilon_K}{\epsilon_L} \ln\left(\frac{K(t)}{Y(t)}\right) + \ln A(0) + g_A t. \tag{3.13}$$

If you compare equation (3.13) to (3.8), you'll see the only difference is the capital/output ratio. We can use this comparison to make one other conclusion.

Conclusion 3.6 (Transitional growth and output per capita) *Transitional growth, $g_{K/Y}$, depends positively on the difference $\ln y(t)^{BGP} - \ln y(t)$.*

The farther actual output per capita is below its balanced growth path, the faster a country grows. This follows from the fact that $\ln y(t)^{BGP}$ and $\ln y(t)$ only differ because of the capital/output ratio, and the farther actual $K(t)/Y(t)$ is below $(K/Y)^{BGP}$, the higher is $g_{K/Y}$. This isn't an absolute statement about the level of output per capita and the growth rate. The conclusion tells us that transitional growth is associated with countries that are poor *relative to their own balanced growth path*.

The conclusion also means that the farther actual output per capita is above its balanced growth path, the slower it grows.

The prior sections established that transitional growth arises because the capital/output ratio isn't equal to the balanced growth path value, but that doesn't explain why this might occur. Given that the natural dynamics of the economy always push the capital/output ratio towards the BGP level, you might expect that all countries should be near that value at all times. But we can consider several different situations where something exogenous happens that creates a gap between $K(t)/Y(t)$ and $(K/Y)^{BGP}$, and hence between $\ln y(t)$ and $\ln y(t)^{BGP}$.

Germany is a classic example. $K(t)/Y(t)$ fell because of World War II, while $s_I/(g_A + g_L + \delta)$ remained fundamentally the same. This in turn meant that actual output per capita, $\ln y(t)$, was less than the BGP level of output per capita, $\ln y(t)^{BGP}$. The result was the miracle growth of the 1950s and 1960s. In terms of output per capita, Figure 3.2 shows what happened, and you can compare that to the actual data for Germany in Figure 1.1. It obviously isn't a perfect match, but the broad outline of the model is correct. Note that this drop in $K(t)/Y(t)$ was completely exogenous from the perspective of the Solow model; nothing in our equations could have predicted World

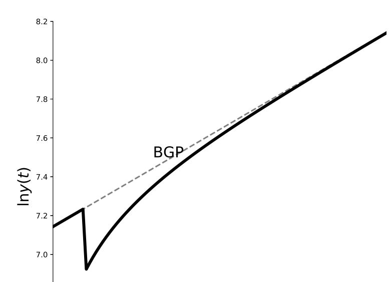

Figure 3.2: An example of a shock to the K/Y ratio that occurs in 1960. The dark line traces the actual path of output per capita.

War II. But the model does allow us to explain what happened in Germany given that exogenous shock happening.

A different case to consider is South Korea after 1960. There, it doesn't appear that there was a distinct fall in $K(t)/Y(t)$, as we do not see a distinct fall in $\ln y(t)$. Rather, there must have been an increase in $s_I/(g_A + g_L + \delta)$. The BGP itself moved, not the actual capital/output ratio, and South Korea was in a position where $\ln y(t) < \ln y(t)^{BGP}$, and so it had the rapid growth observed in the 1970s, 1980s, and into the 1990s. Figure 3.3 plots the stylized path, which again isn't an exact match, but gets the right idea.

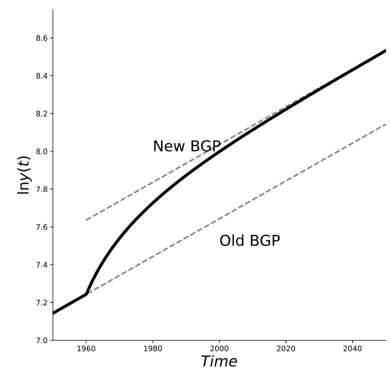

Figure 3.3: An example of a shift in the BGP that occurs in 1960. The dark line traces the actual path of output per capita.

Again, the South Korean case arose because of an exogenous shock, this time to s_I or g_L (or possibly $A(0)$ as we'll discuss below). Nothing in the model could have predicted how or why South Korea enacted policies or changed behaviors around 1960 to create that shock to the BGP capital/output ratio. But conditional on it occurring, the Solow model can explain what happened afterwards.

Those examples are useful because we see the whole evolution of transitional growth, from large values of $g_{K/Y}$ right after the exogenous shocks, and then the drop of $g_{K/Y}$ as the economies approach their balanced growth path. A different kind of example is China, where we see the distinct change around 1980, and an acceleration of the growth rate. The Solow model tells us that their transitional growth will be high for a while, but that eventually growth will slow down towards g_A as their capital/output ratio reaches the BGP level. But nothing about the Solow model can explain the political and historical reasons Deng Xiaoping initiated market reforms in 1980. And because China appears to be in the middle of this transitional growth, it may well turn out that we are wrong about what happens in the future. We're assuming it will follow a path similar to those of Germany, South Korea, or Japan.

It is possible to derive an exact prediction for $\ln y(t)$ over time using the Solow model. It is somewhat tedious to work out, and is best implemented on a computer, not pencil and paper. See A.15.

The logic behind a shock to baseline productivity, $A(0)$, is similar to the examples given above, but it takes a little extra work to see why. If $A(0)$ rises exogenously, we know that nothing happens to $(K/Y)^{BGP}$. And if you look at equations (3.13) to (3.8) you'll see that they both appear to shift up the same amount.

A different question is what happens when g_A changes. This becomes involved, as it changes the BGP capital/output ratio and the growth rate.

But actual output per capita doesn't jump by this full amount because the productivity shock also *lowers* the actual capital/output ratio. The capital stock itself is unchanged, but output goes up due to higher productivity, and thus $K(t)/Y(t)$ falls. The net effect of the rise in productivity and the drop in the capital/output ratio will be for output per capita to rise, but not enough to reach the new BGP.

To see this, go all the way back to to the expression for changes in output per capita in equation (3.6), and manipulate that as follows.

You can get this same logic from looking back at the original equation (2.8). This already shows us that the net effect of a productivity change on output is $\epsilon_L d\ln A$.

$$
\begin{aligned}
d\ln y &= \frac{\epsilon_K}{\epsilon_L} d\ln K/Y + d\ln A \\
&= \frac{\epsilon_K}{\epsilon_L}\left(d\ln K - d\ln Y\right) + d\ln A \\
&= \frac{\epsilon_K}{\epsilon_L}\left(d\ln K - \epsilon_K d\ln K - \epsilon_L d\ln L - \epsilon_L d\ln A\right) + d\ln A.
\end{aligned}
$$

You can see here the two effects of a shock to productivity. First, inside the parentheses is the negative effect on the capital/output ratio. But notice that when you multiply this out you get $-\epsilon_K d\ln A$. Second, at the end you get the positive productivity effect, and notice that this is simply $d\ln A$. So the raw productivity effect is larger in size than the negative effect working through capital/output, and on net the change in output per capita is $\epsilon_L d\ln A$.

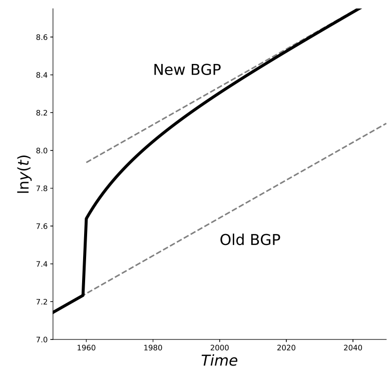

Figure 3.4: An example of a positive shock to $A(0)$ in 1960. The dark line traces the actual path of output per capita.

Figure 3.4 shows what this looks like over time. At the time of the shock output per capita jumps partway up towards the new BGP, and then after this follows a normal, slow transition path. The growth rate of output per capita in that period during the shock to $A(0)$ would be massive, and then the growth rate would taper off towards the same g_A that existed before.

Going back to the example of South Korea or China, this could be an additional explanation for their rapid growth. Rather than a jump in investment share or a drop in population growth (although both occurred in those countries), there could have been a significant jump in baseline productivity. One might argue that the market reforms of China were really a productivity increase, perhaps. But again, we'd have to investigate the actual history and data from these countries to decide what drove their rapid transitional growth. The Solow model tells us where to look for this evidence, but doesn't tell us why things changed in the first place.

There is some interesting evidence from East Asia that suggests productivity improvements were *not* the major factor in their rapid growth.

3.5 *Implications*

Given that we observed countries on balanced growth paths, we were able to make some powerful conclusions about what drives growth in the long run. In particular, conclusions 2.1 and 2.2 tell us that the growth rate depends on productivity growth, g_A, in the long-run, and not on something like the investment share in physical capital. This will motivate the deeper study of productivity growth going forward.

The Solow model gave us a framework for understanding why countries end up on a BGP at all. The way that capital accumulates ensures that the capital/output ratio always returns to some BGP value eventually, which by necessity drives the growth rate back to g_A. This allowed it to make sense of the fact that differences in growth rates across countries were due to transitional growth, $g_{K/Y}$, and are temporary. The additional value of the model was that it gave us clues for why some countries found themselves "below" their balanced growth path at some point, and why there was any difference in growth rates to begin with.

4
Investment in capital

One of the key facts regarding growth was that the investment share of GDP, s_I, was stable over time. The consumption share of GDP, which we could call s_C, was also stable over time. These facts were central to the understanding of balanced growth; the *definition* of balanced growth included consumption growth matching GDP growth, implying s_C was constant. The stability of the capital/output ratio was dependent on a stable value for s_I.

But those shares represent choices about how we allocate spending in the economy. There is nothing that forces these shares to be stable. What we need to understand, then, is why the households and firms in the economy *chose* to keep s_C and s_I stable over time. In some sense, this is just a standard expenditure problem. But the choice on how much to spend on consumption versus investment is not quite the same as deciding how much to spend on different types of consumption (e.g. shirts versus food), as investment spending buys products that will deliver consumption value in the *future*. The investment products can be used to produce consumption goods and services (e.g. a factory that packages food), or are themselves durable consumption goods (e.g. a house).

This means there is an inherent inter-temporal aspect to the choice about how spending on consumption versus investment. We thus have to think harder about how to measure the relative price of consumption versus investment. There is the actual cost of the investment goods themselves today which matter, but then we have to account for the fact that the investment goods will then generate consumption goods (or can be used for consumption themselves) in the future. That return on investment spending is going to be a key component of the decision to buy investment goods. And because of that return, investment goods are going to look relatively "cheap" in terms of pure consumption goods. At the same time, the return we get from investment goods comes in the future, not today, and we probably want to take into account that time delay.

Beyond the relative price, we cannot just focus on current income when thinking about the investment/consumption decision. If we know that income will be larger tomorrow regardless of what we do, then we might not worry as much about buying investment goods that deliver consumption value in the future. We can complicate that more by realizing that income tomorrow probably depends on how much investment spending we do today, because that investment spending determines the capital stock we have. Regardless, we'd expect that the investment/consumption decision depends on our projection of future income, not just the income we have today.

We can build up a model of the choice over how much to consume and invest, similar to how we built up a model of how GDP was produced. Mapping that model to the data, we'll be able to draw some conclusions about how the preferences for consumption work at the aggregate level. Given those, we will be able to build up an explanation for why the investment and consumption shares appear stable in the long run.

4.1 The return on capital

Before we build any more theoretical structure, let's establish some basic facts surrounding the return on capital and the relationship of that return to the capital/output ratio.

We run into an immediate problem here, similar to the one in trying to measure real GDP growth. We do not observe the actual physical output of each unit of capital, and so we need some way to infer the return on capital from things we do observe. We'll use the financial returns from owning different types of capital, making the assumption that people would only own this capital if it had some real return, and that over the long run the financial returns will be informative about the real.

Financial returns are in nominal terms, so they are deflated by the consumer price index to recover the financial return in terms of consumption goods.

Evidence from over a long period of time on the rate of return to different types of capital is now available for a set of advanced economies.[1] Figure 4.1 plots the return to housing, equities, and government bonds for the United States. The returns are the average of the yearly returns for five-year windows.

[1] Òscar Jordà, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M Taylor. The Rate of Return on Everything, 1870–2015*. *The Quarterly Journal of Economics*, 04 2019

There are certainly fluctuations in the rates of return over time, and if the data were plotted on an annual basis, the fluctuations would be more pronounced. However, despite the fluctuations, there does not appear to be a distinct trend in the return of any capital type over time. The rate of return on housing over the whole course of the data centers around 5-6%, the average for equities is around 8%, and the return on safe assets is around 2.5% over this whole period.

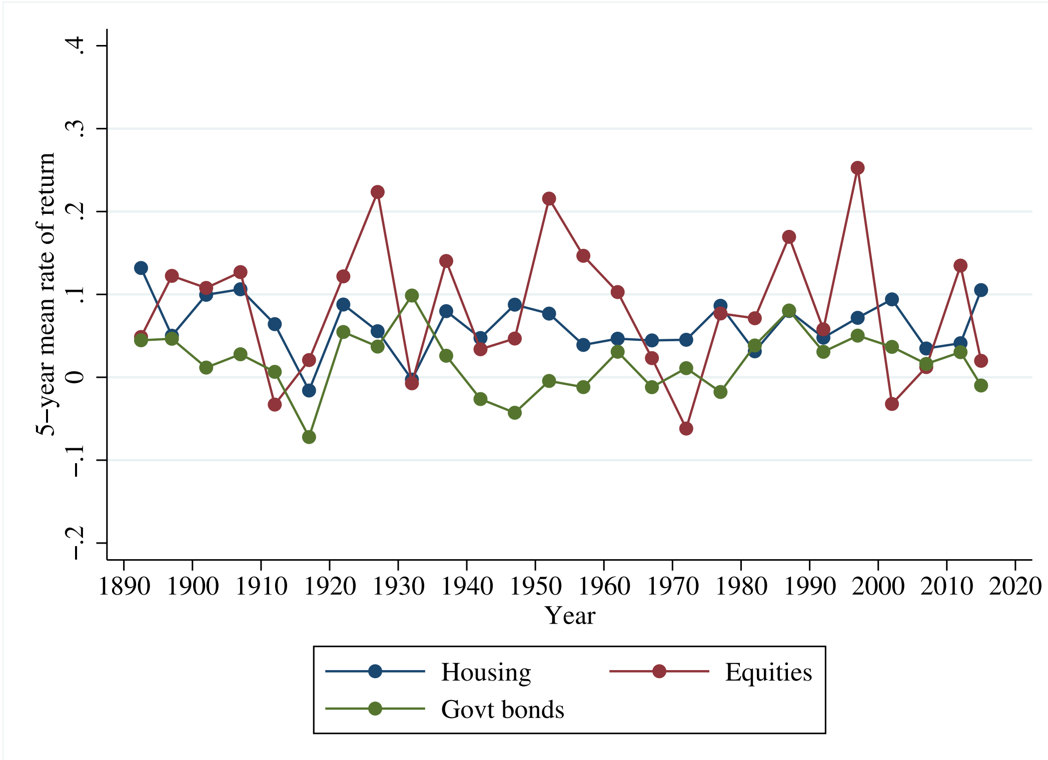

Figure 4.1: Rate of return on housing, equities, and government bonds, USA, 1980-2016. The average return from each 5-year period (e.g. 1970-1974) is shown.

The fluctuations for the United Kingdom, in Figure 4.1, are more dramatic than in the United States. But a similar pattern emerges that the rate of return on equities and housing do not have a distinct trend over time. The averages for the three types of capital are not that much different for the U.K., either, at around 6% for equities, 5% for housing, and 1.7% for safe assets.

Similar plots for other developed countries can be done, and the overall story remains the same: fluctuations in the rate of return on housing and equities, but no overall trend.

We can use the patterns in Figures 4.1 and 4.1 to establish another broad fact.

Fact 4.1 (Stable rates of return) *The rates of return on different types of capital (housing, equities, government bonds) are stable in the long run.*

This stability in rates of return gives us our first clue about why investment shares were stable over time. We think that the rate of return is an important element of why people purchase investment goods in the first place. If the rate of return were rising steadily over time, we might have expected even more investment purchases (and an increase in s_I), or perhaps even less given that purchasing fewer could still yield a big return. Regardless, with a stable return, it makes some intuitive sense that s_I would be stable as well.

Despite the facts on financial rates of return, we'd still like some assurance that we can connect this to the evidence we have on capital

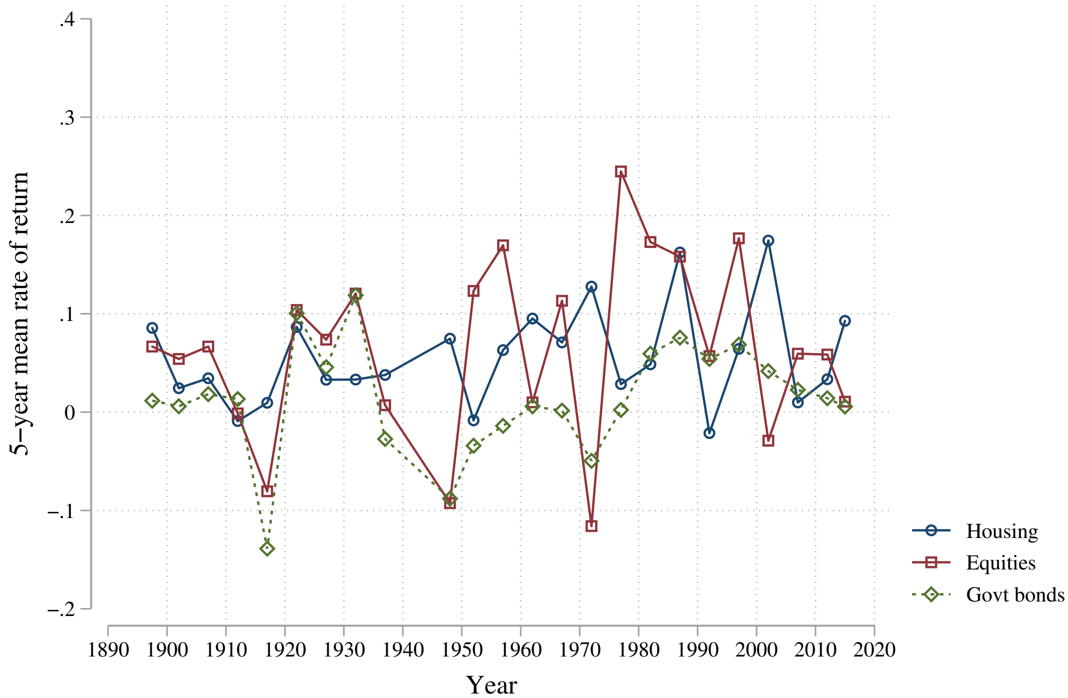

Figure 4.2: Rate of return on housing, equities, and government bonds, United Kingdom, 1980-2016. The average return from each 5-year period (e.g. 1970-1974) is shown.

and output. The rate of return on capital should depend positively on its marginal product, and in the theoretical setting we laid out in Chapter 3 the marginal product of capital is proportional to the average product of capital. That means the rate of return should be *negatively* related to the capital/output ratio, K/Y, which is just the inverse of the average product of capital.

This is a consequence of the diminishing marginal product of capital, which was reflected in the fact that ϵ_K was less than one.

Using all 16 countries that Jordà et al. (2019) provided data for, Figure 4.1 summarizes the relationship of the rate of return on equities with the capital/output ratio. The figure shows the relationship net of both country and year effects, meaning that the relationship is not driven by differences in the average returns of countries or years. Regardless, this is not very promising, although there is a slight negative relationship. There is a lot of noise in the return on equities relative to the size of the capital/output ratio.

Controlling for the country and year fixed effects in this way essentially "de-means" the data for each country and year.

However, if we examine Figure 4.1, then the negative relationship of the rate of return on housing and the capital/output ratio is quite strong. Again, this figure shows the data net of country and year effects. This gives us a little more hope that the rate of return data are telling us something real about the return on capital we're using in the theoretical model. Furthermore, recall that housing (and structures) make up the vast majority of the capital stock in nearly all

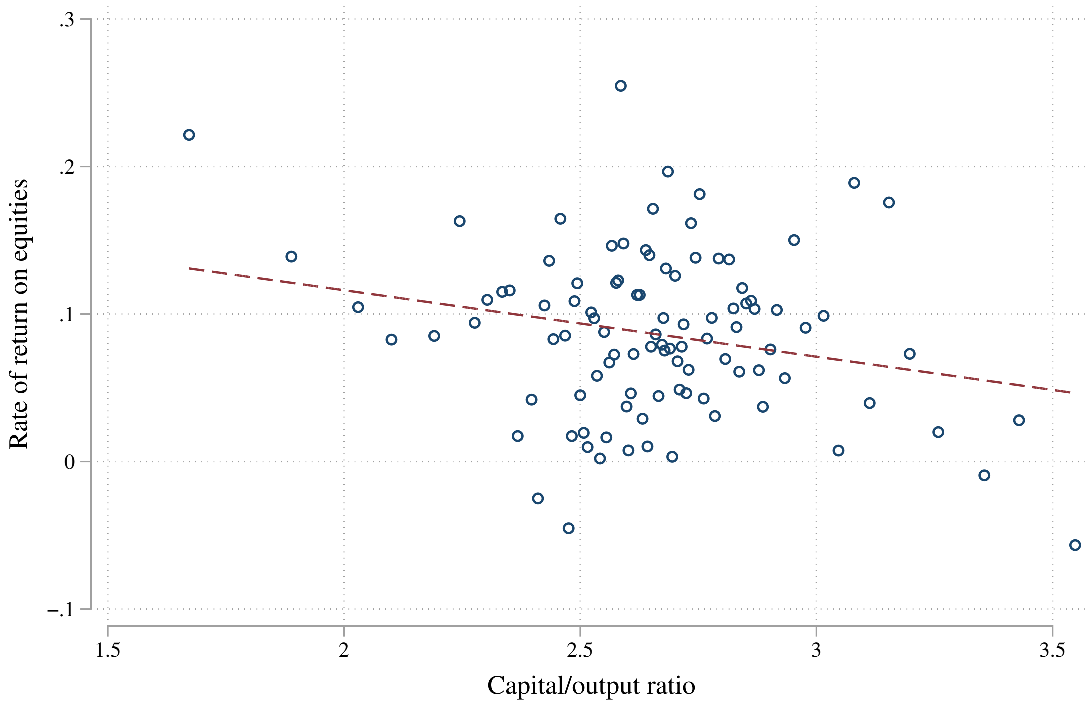

Figure 4.3: Relationship of rate of return on equities to the aggregate capital/output ratio for 16 countries 1950-2015. This is the residual after removing both country and year fixed effects.

countries, and so the return on housing is quite close to the overall return on capital, of which equities are but a small fraction.

Despite the noise in the figures, we're going to establish this as a further fact about growth.

Fact 4.2 (Capital/output and returns) *The rate of return on capital has a (loose) negative relationship to the capital/output ratio.*

This fact is what we're going to use to link our theoretical findings about the capital/output ratio to the consumption/investment decision, which depends on the rate of return on capital. It will help us establish why s_I is stable in the long run. One very important caveat about Fact 4.2 is that it is not a *causal* claim, only an observation about a correlation. That is, the fact suggests that capital/output ratios are informative about returns on capital, but does not prove that if the capital/output ratio were to rise (for example), that the return on capital must fall.

All that said, from our model of production we have established how rates of return and the capital/output ratio could be related. From 2.6 we had that the share of capital income in GDP was $s_K = RK/Y$, which if we manipulate implies the following:

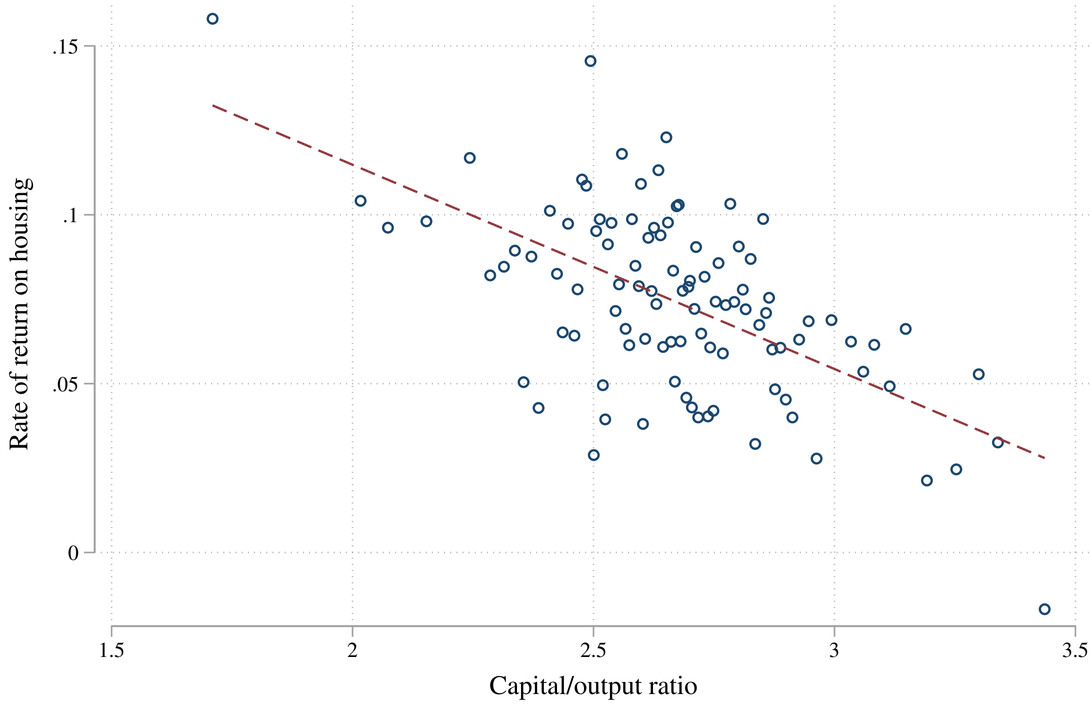

Figure 4.4: Relationship of rate of return on housing to the aggregate capital/output ratio for 16 countries 1950-2015. This is the residual after removing both country and year fixed effects.

Assumption 4.1 (Gross rate of return) *The **gross** rate of return on capital is*

$$R = s_K \frac{Y}{K} \tag{4.1}$$

From an individual perspective what matters is the net rate of return, which allows for the fact that in using capital (or renting it to someone else who uses it) some of that capital depreciates at the rate δ.

Assumption 4.2 (Net rate of return) *The **net** rate of return on capital is*

$$r = \frac{s_K Y - \delta K}{K} = s_K \frac{Y}{K} - \delta. \tag{4.2}$$

These definitions draw a distinct relationship between rates of return and K/Y. The reason the data above is only indicative of these relationships is that the data are about market returns on financial assets. While in principle those financial assets represent ownership stakes in real capital (e.g. shares of a company represent ownership of that company's capital stock) the exact nature of ownership is not always obvious (e.g. investor protections and rights vary by type of share) and there is substantial risk and uncertainty due to trading in those financial assets. It is probably best to think of these definitions of R

and r, and their relationship with K/Y, as representing the baseline or underlying rates of return that all these other considerations build off of.

4.2 *The capital accumulation choice*

We can establish some conditions on how s_I can act that will support the stability of the Solow model and ensure that s_I will be constant in the long run. That will help point us towards the nature of the full decision process that we will go into in the next chapter.

That decision process is based on people's choice of what to consume versus what they choose to accumulate, in essence meaning they pick s_C and s_I. But the decision problem we'll solve is most readily expressed in terms of the growth rate of consumption, so let's connect that here to see why we care about it.

It's the case that $s_I = 1 - s_C$, and $s_C = C/Y$, so

$$ds_I = -ds_C = -(dC/Y - s_C dY/Y) = -s_C(g_C - g_Y) = s_C(g_Y - g_C). \tag{4.3}$$

Any change in s_I can be connected via this to the growth rate of consumption relative to the growth rate of GDP. Here it helps if you flip around to the income defintion of GDP, and then g_Y means the growth rate of income. Therefore the change in s_I will be positive if consumption grows slower than income, and the change in s_I will be negative if consumption grows faster than income.

Along a balanced growth path we want s_I to be constant, so we want $ds_I = 0$, which means something that should not be surprising, that $g_Y = g_C$, or consumption grows just as fast as income. That's already something we established was part of a BGP. So in that sense saying s_I is constant on a BGP was already a necessary condition of saying that $g_Y = g_C$. Or if you like, these are just two ways of saying the same thing.

But that still leaves the question of the choice of g_C that individuals make. And there are a few specific questions we can think about.

1. Under what conditions does it make sense that when K/Y is at steady state they'd *choose* to set $g_C = g_Y$ and have s_I be constant?
2. Under what conditions does it make sense that When K/Y is not at steady state, they *choose* to set g_C (and/or s_I) such that the economy moves back towards steady state?
3. Under what conditions does it make sense that they choose $s_I > 0$ at all?

For the rest of this chapter we will talk about the first question, and next chapter we'll take on the second and third, which is far more complicated. Remember, we see in the data that economies work in a way to support the BGP. We're not asking if people will act this way. We're asking why they act this way.

4.3 *The consumption decision*

We need some additional structure to think about the consumption decision and assess whether it supports a balanced growth path. As mentioned at the opening of this chapter, the consumption/investment choice is inter-temporal in nature, and so looks a little different than a static choice problem over how much pizza and beer to purchase. In the case of the consumption/investment decision, we think of people getting utility from consumption goods today and from consumption goods tomorrow (or next quarter or nexdt year). Investment goods purchased today are only valuable to a person in the sense that they provide consumption goods tomorrow. So our inter-temporal choice problem is really about choosing how much to consume today and tomorrow, where investment spending today is just the mechanism that people use to "buy" some extra consumption tomorrow. For a more concrete example, think of investment spending on housing, meaning new construction. I might spend money today building a house - buying an investment good - for the purpose of enjoying the consumption value of living in that house tomorrow.

The choice problem should thus concern itself with consumption today and tomorrow, and we can back out the implications for investment spending. As with any typical choice problem, we'd expect that people would act until there was no way to increase their utility. Let's think about a household who has consumption of c_1 today, and consumption of c_2 tomorrow

How would their utility change if they made a tiny change, dc, to consumption today? From today's perspective, they would experience a change of $U'(c_1)dc_1$ to their utility, where $U'(c_1)$ is the marginal utility of consumption today. If $dc_1 > 0$, then their utility would go up, and if $dc_1 < 0$ their utility would go down. $U'(c_1)$ captures by how much their utility would change in either direction.

If they did make this dc_1 change to consumption today, what happens tomorrow? Imagine that $dc_1 < 0$, so that they spent less today. They could take that money and buy investment goods, and this would give them $-dc_1(1+r)$ in consumption tomorrow. That $(1+r)$ is the return on investment goods, and the evidence presented in the prior chapter gives us an idea of how big of a return that might be. The negative sign on $-dc_1(1+r)$ reflects the fact that lowering consumption today $dc_1 < 0$ means we have more consumption tomorrow $-dc_1 > 0$.

What is the effect on utility of having that additional $dc_1(1+r)$ to consume? That depends on the marginal utility of consumption *tomorrow*, which is $U'(c_2)/(1+\theta)$. The $U'(c_2)$ term is similar to the one today, but this marginal utility is divided by $(1+\theta)$. This term reflects the fact that we expect people to discount the future, perhaps because they are impatient, so that the marginal utility tomorrow is scaled down by the factor $(1+\theta)$.

A value of θ of, say, 0.05, would imply that utility tomorrow would be $1/(1.05) = 0.952$ of the utility today for the same amount of consumption.

We can look at the total effect on utility of this proposed change, which given our assumption that the person chose c_1 and c_2 to maximize their utility, should mean it is impossible to increase their utility, or that

$$U'(c_1)dc_1 - \frac{U'(c_2)}{1+\theta}dc_1(1+r) = 0.$$

This can get re-arranged to

$$\frac{U'(c_2)}{U'(c_1)} = \frac{1+\theta}{1+r}, \tag{4.4}$$

which looks a lot like the standard condition that the ratio of marginal utilities is equal to the ratio of "prices". Here, the relative price of consumption tomorrow to consumption today is $(1+\theta)/(1+r)$. If $\theta > r$, then tomorrow is "expensive" relative to today, because the discounting of utility tomorrow is high and/or the rate of return on investment goods is low. In this case, we'd get that $U'(c_2) > U'(c_1)$, which would imply that $c_2 < c_1$, or that the person would consume more today than tomorrow. You can work through the reverse logic when $\theta < r$, and see that this implies $c_1 < c_2$, as today looks expensive relative to tomorrow given a low discount rate and/or a high rate of return.

This doesn't mean they don't buy investment goods at all. The actual purchase of investment goods would depend on how big their income today was relative to their choice of c_1.

4.4 *Consumption along a BGP*

It's possible to do a surprising amount of analysis with this simple relationship. Recall that what we're trying to do is establish why s_C and s_I are constant, which is equivalent to explaining why the growth rate of consumption is constant (and equal to the growth rate of output). Denoting g_c as the growth rate of consumption *per capita*, we could write $c_2 = (1+g_c)c_1$.

Using this, we can manipulate the marginal utility of consumption tomorrow. So long as g_c isn't very big, then

$$U'\left((1+g_c)c_1\right) \approx U'(c_1) + U''(c_1)c_1 g_c$$

where $U''(c_1)$ is the second derivative of utility, and it tells us how much the *marginal* utility of consumption changes when consumption changes. Diminishing marginal utility implies that $U''(c_1) < 0$. If we put this approximation to work in equation (4.4), we get

This approximation is from a Taylor series expansion of marginal utility around the point c_1. See A.2 for a refresher on Taylor expansions.

$$1 + \frac{U''(c_1)c_1}{U'(c_1)} g_c = \frac{1+\theta}{1+r}.$$

To go further, we're going to manipulate the "price" ratio as well. For small values of θ and r, it holds that

$$\frac{1+\theta}{1+r} \approx 1 + \theta - r.$$

Put all this together and we get

$$g_c = (r - \theta)\left[-\frac{U'(c_1)}{U''(c_1)c_1}\right]. \tag{4.5}$$

This expression looks ugly, in particular the part in the brackets, but that is going to have a straightforward interpretation. Given that $U''(c_1) < 0$, that bracketed term is positive. Before focusing on that, consider what this equation tells us. It says that the growth rate of consumption depends on the relative size of r and θ. When $r > \theta$, $g_c > 0$, which is just a restatement of what we said above. The opposite case, with $r < \theta$, implies that $g_c < 0$. The bracketed term just scales how big of an effect the difference in rates of return and discount rates have on consumption growth.

To go forward, let's get some idea of what that bracketed term involves.

Definition 4.1 (Intertemporal elasticity of substitution) *The intertemporal elasticity of substitution (IES) is defined as*

$$IES \equiv \frac{-U'(c)}{U''(c)c} = \frac{-d\ln c}{d\ln U'(c)}$$

and measures how sensitive consumption is to a change in the marginal utility of consumption.

The IES captures how willing people are to change their consumption behavior in response to a change in marginal utility. The change in marginal utility, based on the choice problem outlined above, would depend on the change in relative price, so the IES tells us how sensitive consumption behavior is to a change in the price of consumption today versus consumption tomorrow.

4.5 *Conditions on stability*

Now, under what conditions does it make sense that people choose to set $g_c = g_y$ along a BGP? What's necessary for this all to hold together?

Conclusion 4.1 (Constant IES) *Given that the rate of return on investment goods, r, is stable on a BGP, that the growth rate of consumption, g_c, is stable on a BGP, and the assumption that the discount rate θ is constant, it must be that the inter-temporal elasticity of substitution (IES) is stable along a BGP.*

This comes from examining equation (4.5), and applying what we know from the data about rates of return and consumption growth. The one additional assumption is that the discount rate θ, as a fundamental preference parameter, is constant over time. We're assuming that people are consistently impatient. The combination of these facts and assumptions means that the IES must be stable along a BGP as well. Apparently the willingness to adjust consumption behavior in response to relative prices did not change as the absolute size of consumption grew over time. A constant IES is what rationalizes the consumption decision with the BGP.

We can go even further than this, and establish something that will seem entirely too exact to describe the real world.

Conclusion 4.2 (Balanced growth preferences) *Let the IES be equal to a constant, so that $-U'(c)/U''(c)c = 1/\sigma$. Then the utility function over consumption in any given period must be of the form*

$$U(c) = \frac{c^{1-\sigma}}{1-\sigma}. \tag{4.6}$$

Proving this conclusion holds is a little tedious mathematically, but the intuition is not too hard. A constant IES means that no matter how big consumption is, every time consumption grows by 1%, the marginal utility falls by some constant percent. That means marginal utility has to be a power function of consumption, as in $U'(c) = c^{-\sigma}$. Integrate both sides of this and you get the preferences in the conclusion.

These preferences are referred to as Constant Relative Risk Aversion (CRRA) preferences as well. This form of preferences also implies that risk aversion does not change as the absolute size of c changes. See A.17.

These are called "balanced growth preferences" because they are the only form for preferences that necessarily will match the consumption behavior we see. If you change something fundamental about the structure of the economy here, say by adding something different about how people discount time, then you can break the need to have this exact form of preferences. But these are the standard in a model of the economy because they support the stability of the BGP in a typical setting.

We can go further in establishing conclusions about the rate of return along a balanced growth path.

Conclusion 4.3 (Rate of return and time discounting) *Given that there is constant positive growth in consumption along a BGP, it must be that the rate of return r^{BGP} is greater than the time discount rate θ.*

This is easy to establish from equation (4.5), which shows that the sign of consumption growth depends on the size of r relative to θ. Positive consumption growth requires $r > \theta$. Beyond that, we know that the growth rate of consumption along a BGP must be $g_c = g_y$, as this ensures that the share of consumption in GDP, s_c, remains constant as well. From conclusion 2.1 that $g_y = g_A$, we can put the following conclusion together.

Conclusion 4.4 (Rate of return on the BGP) *Given that s_C is constant along a BGP then the rate of return along a BGP is $r^{BGP} = \theta + \sigma g_A$.*

The conclusion implies that $r^{BGP} > g_A$ as we'll assume $\sigma \geq 1$. Thomas Piketty used the idea of $r > g$ to explain rising inequality. His definitions of r and g are not quite the same as those used here, so there isn't an obvious conclusion to draw about inequality from the condition here.

This gives us a way of understanding the forces that determine the rate of return. First, the rate of return depends on the discount rate θ, which recall captures how impatient we are. To get people to buy investment goods at all, the return on them must be high enough to overcome their natural impatience. So the higher is θ, the higher the rate of return on the BGP will be.

Second, the growth rate g_A matters for the rate of return. This determines how fast the economy grows. The rate of return has to be at a level that will convince people to buy investment goods such that consumption growth is the same as output growth. The faster the economy grows, the faster consumption has to grow, and this will happen if the rate of return is higher.

Finally, the inter-temporal elasticity of substitution matters negatively (recall that the IES is $1/\sigma$). When the IES goes *down*, and σ is higher, then the rate of return is higher. A low IES means that people are not willing to substitute consumption tomorrow for consumption today. The only way to induce people to have consumption growing at a rate equal to the growth rate of output is for the rate of return to be relatively high. In situations where the IES is high already (and σ is small), then people are easy to convince, and the rate of return does not have to be as high.

Conclusion 4.4 is probably best seen as a way of mentally accounting for the rate of return, rather than a hard and fast quantitative rule. Two of the elements in it, θ and σ, are preference parameters that we cannot observe directly. Most important, the setting here has ignored risk in investment, which would presumably add an additional element to the rate of return along a balanced growth path.

Caveats aside, if θ is around 0.05, and some alternative estimates of σ suggest it is close to 2, and we know that g_A is around 0.018, then this implies $r^{BGP} = 0.086$, which is in the right ballpark for the observed rates of return.

4.6 *Reconciling with the production side*

We already had a definition of r in 4.2, though, and that depended on K/Y. And we already know from 3.4 that K/Y along a BGP is determined by $s_I/(\delta + g_A + g_L)$. Now we've got another defintion of r along a BGP, and somehow they all have to make sense together.

The consumption problem sort of "wins" here, in that it establishes what must be true such that individuals make a choice over consumption will act in a way that $g_c = g_y$, which in turn ensures that s_I is constant. So from 4.4 we know that

$$r^{BGP} = \theta + \sigma g_A \tag{4.7}$$

which given the definition of r must mean that

$$(K/Y)^{BGP} = \frac{s_K}{\theta + \sigma g_A + \delta}. \tag{4.8}$$

At the same time, it has to be from the production side that $(K/Y)^{BGP} = s_I/(\delta + g_A + g_L)$. This tells us that

Conclusion 4.5 (Savings rate along a BGP) *The savings rate s_I along a BGP (but not necessarily off the BGP) must be*

$$s_I^{BGP} = s_K \frac{\delta + g_A + g_L}{\theta + \sigma g_A + \delta} \tag{4.9}$$

The savings rate that supports the BGP - the savings rate that is consistent with people choosing to set $g_c = g_y$ - depends on the share of output that goes to capital, s_K. That makes sense. If capital earns a lot it makes sense for people choosing consumption to invest a lot in accumulating it because they'll be able to consume the big return in the future.

It's also true that if $\theta > g_L + (1 - \sigma)g_A$, then it must be that $s_I^{BGP} < s_K$, or that the economy along a BGP puts less effort into accumulating capital than it's share of income. In that sense you could say the economy "under-invests". Essentially, this is because individuals are impatient and don't care enough about the far future to invest at a rate such that $s_I^{BGP} = s_K$.

In the next chapters we'll see that this condition must hold in any rational consumption decision.

We have a lot of information about how any BGP must operate. Recall that s_I^{BGP} dictates the *level* of GDP per capita, even though it does not influence the *growth rate* of GDP per capita along the BGP. This level depends in part on the preferences of individuals about time - θ - and how willing they are to substitute consumption over time - σ. We've replaced the static s_I in the Solow model with a "deeper" set of parameters in some sense.

What we don't have from all this work is any sense of what happens to s_I *off* the BGP. We really should be talking about s_{It}, the capital accumulation rate at any given point in time, and there is nothing we've established so far that tells us exactly what this will be. Most important, there is nothing about what we've established so far that ensures that s_{It} will act in a way that supports the stability of the economy. What we need is that s_{It} acts to ensure that K/Y always collapses towards a steady state. The next chapter will lay out the larger consumption problem behind section 4.3 and that will let us establish that in fact there is a reasonable model of individual choice that leads to stability.

5
Consumption decisions

The stability of the the investment share s_I (and of consumption s_C) is a core feature of the BGP. The purpose of this chapter is to evaluate whether a reasonable model of individual decision-making about consumption and investment decisions is consistent with a relationship of s_I and K/Y that supports stability. The unsurprising answer is yes, it is, and the structure of that decision forms the basis for most models of macroeconomic behavior.

We're going to start with a simple two-period problem and use discrete time concepts. We need to build the tools of discrete time optimization because ultimately we're going to add uncertainty, and this is easier to conceive and model in terms of discrete steps (e.g. things happen on specific dates). Once we do that, we'll come back and see how this looks in a continuous time model, which is what we used to do the intuitive work in the last chapter. The economics are all the same, there is just a difference in how the math works, and sometimes one or the other is useful.

5.1 Simple consumption problem principles

All of the important principles of the consumption problem can be learned looking at a two-period model of consumption. Our decision-maker is alive for two periods and has some stock of assets a_1 that they start with, and they earn some income w_1 in the first period. They have to choose how much to consume in period 1, c_1. They earn a rate of return on their assets of r. This gives us the following dynamic budget constraint

$$a_2 = w_1 + (1+r)a_1 - c_1. \tag{5.1}$$

In this sense the constraint is that you have $w_1 + (1+r)a_1$ to deal with, and are choosing between c_1 and a_2.

There is a notational convention here, which is that you earn a return on your assets coming into the period, a_1, and end up with the assets a_2 that you take to period 2. This is purely a choice in how to denote things.

In period 2 they have

$$0 = w_2 + (1+r)a_2 - c_2, \tag{5.2}$$

where I've written that in an odd way. The assertion here is that $a_3 = 0$, which makes sense because this decision-maker is only alive for two periods. There is no reason for them to accumulate assets for the future once they get to the second period.

The period 2 choice of consumption is thus trivial,

$$c_2 = w_2 + (1+r)a_2. \tag{5.3}$$

We can combine the two constraints into a *lifetime budget constraint* of

$$\frac{c_2}{1+r} - \frac{w_2}{1+r} = w_1 + (1+r)a_1 - c_1 \tag{5.4}$$

or

$$c_1 + \frac{c_2}{1+r} = (1+r)a_1 + w_1 + \frac{w_2}{1+r}. \tag{5.5}$$

The left hand side is the present discounted value of consumption, and the right hand side is the present discounted value of income and assets, or lifetime wealth. We could easily assert that this person has $a_1 = 0$ and the constraint is not really any different.

What does this person care about? They care about consumption in the two periods, and we'll assert they have liftime utility of

$$V = U(c_1) + \beta U(c_2), \tag{5.6}$$

where the time preference rate β measures how much you care about period 2 compared to period 1. In this finite time setting it doesn't matter whether β is bigger or less than one. In problems with infinite time horizons it will have to be that $\beta < 1$. This lifetime utility has two properties that we'll use throughout the analysis of consumption.

The time preference rate β is related to the discount rate θ via $\beta = 1/(1+\theta)$. They are measuring the same thing, just using different notation.

Assumption 5.1 (Additive separability) *The lifetime utility function V is additively seperable if the marginal utility of consumption at time t is independent of consumption at time s.*

For our simple example, note that if you want the marginal lifetime utility of c_1, that is $\partial V/\partial c_1 = U'(c_1)$, and this does not depend on the size of c_2. The simpler way to see this is that we've literally *added* the utilities $U(c_1)$ and $U(c_2)$, and that in neither case is it $U(c_1, c_2)$. Additive separability might arguably be wrong - think of durable goods - but it delivers a problem that is tractable and which can match many pieces of data, if not all.

Assumption 5.2 (Concave utility) *The per-period utility function $U(c)$ is concave with respect to c, meaning $U'(c) > 0$ and $U''(c) < 0$.*

This concavity assumption means that $U(c)$ has diminished marginal utility, which is the kind of thing we'd assume about almost any kind of good, like pizza, beer, or in this case, total consumption.

The $U(c)$ function is sometimes referred to as the "felicity" function to distinguish it from V, lifetime utility.

It means that it doesn't make sense to load up all your consumption in period 1 or period 2, but rather that you'd like to spread it out.

The BGP preferences from last chapter, $U(c) = c^{1-\sigma}/(1-\sigma)$ are concave because $U'(c) = c^{-\sigma} > 0$ and $U''(c) = -\sigma c^{-\sigma-1} < 0$.

The problem for this person is to maximize lifetime utility - choose c_1 and c_2 - subject to the lifetime budget constraint. With two periods you can "plug and chug" to get an answer on this, but we will set this up more formally to help illustrate how this works in more complex settings.

We're going to add one additional assumption:

Assumption 5.3 (Perfect loan market) *Individuals can borrow and lend/save at the rate r with no limit.*

I've spoken about the setup so far as if individuals always save, which would mean $c_1 < w_1 + a_1(1+r)$, but this assumption means we're not holding anyone to this requirement. That's a big part of the question for us. If people are free to set c_1 higher (borrow) or lower (lend/save) than $w_1 + a_1(1+r)$, then why would it be the case that in the long run they tend to set c_1 lower? Perfect loan markets are an obvious stark assumption, and it's quite possible to add frictions here to make this problem more realistic. There could be borrowing limits, or different interest rates, etc.

The terms "lend", "save", "invest", and "accumulate capital" are all synonyms in these settings. It usually just depends on whether we are talking about individuals (lend/save/invest) or the economy (accumulate capital).

Set up the Lagrangian for this problem as follows:

$$\mathcal{L} = U(c_1) + \beta U(c_2) + \lambda \left(a_1(1+r) + w_1 + \frac{w_2}{1+r} - c_1 - \frac{c_2}{1+r} \right) \quad (5.7)$$

and maximize with respect to c_1, c_2, and λ. That parameter λ is the Lagrange multiplier, and in this problem the interpretation is worth thinking about. As with any multiplier, it tells us the marginal value of relaxing the constraint. In this case, the value is utility, so λ is the marginal utility of lifetime wealth. It translates dollars of lifetime wealth into units of utility.

The first order conditions are:

$$U'(c_1) - \lambda = 0 \quad (5.8)$$

$$\beta U'(c_2) - \frac{\lambda}{1+r} = 0 \quad (5.9)$$

$$a_1(1+r) + w_1 + \frac{w_2}{1+r} - c_1 - \frac{c_2}{1+r} = 0 \quad (5.10)$$

The first condition says that the optimal solution is to set the marginal utility of consumption in period 1 equal to the marginal utility of lifetime wealth. That kind of makes sense. If the marginal utility in period 1 was higher, then it seems like you could lower lifetime wealth by one dollar, and lose only λ in utility, but you'd pick up $U'(c_1)$ in additional utility, and you'd win. So it only makes sense for $U'(c_1) = \lambda$.

The second condition is the same, but we have to translate all this through time. Write that condition this way

$$\lambda = (1+r)\beta U'(c_2). \tag{5.11}$$

This says that the marginal utility of lifetime wealth should be equal to the, well, marginal utility of consumption in period 2. But if you gave up a dollar today (and lost λ in utility) you'd get back $1+r$ dollars tomorrow, and each of those $1+r$ dollars tomorrow would be worth $\beta U'(c_2)$ in marginal utility.

Put those first two conditions together and they say

Conclusion 5.1 (Consumption Euler Equation) *The consumption Euler equation relates the marginal utility of consumption in two periods to the rate of return and the time preference rate:*

$$\frac{U'(c_1)}{U'(c_2)} = \beta(1+r) \tag{5.12}$$

This is the same structure we used in the last chapter to come up with the growth rate of consumption along the balanced growth path. The intuition is identical.

We should really write $\frac{U'(c_1)}{\beta U'(c_2)} = 1+r$ as this is the ratio of marginal utilities equal to the marginal rate of transformation, but convention is to put the β on the right-hand side.

That's it, that's the consumption solution. To get a firm answer for the size of c_1 and c_2 you'd solve the Euler equation along with the last first-order condition - the budget - and get an answer. But as we saw in the last chapter, the key element here for us is the Euler equation which relates consumption between two periods to the time preference rate, β and the rate of return r. This Euler equation implicitly tells us g_c between periods 1 and 2.

There is another important feature to note of this consumption problem

Conclusion 5.2 (Consumption and lifetime wealth) *Given the assumption of perfect loan markets, it is the case that the consumption Euler equation does not depend on either the size of lifetime wealth or on the distribution of lifetime wealth between initial assets and income in different periods.*

Because they can move money around at will, individuals don't care *when* they get their money, their choice about consumption via the Euler equation only depends on the rate r and the time preference β.

If we use the same approximation as in the last chapter, that

$$U'(c_2) = U'\left((1+g_c)c_1\right) \approx U'(c_1) + U''(c_1)c_1 g_c$$

then the Euler equation becomes

$$g_c = \frac{U'(c_1)}{U''(c_1)c_1}\left(\frac{1}{\beta(1+r)} - 1\right) \tag{5.13}$$

which if you plug in that $\beta = 1/(1+\theta)$ you'll get that $(1+\theta)/(1+r) - 1 \approx \theta - r$ and so we're back to

$$g_c = (r-\theta)\left(-\frac{U'(c_1)}{U''(c_1)c_1}\right) \tag{5.14}$$

and the growth rate of consumption depends on the difference in r and the discount rate, modified by the intertemporal elasticity of substitution. All we've done is show that the simple derivation in the last chapter is the solution to a formal two-period problem of optimal consumption.

5.2 *Savings and the rate of return*

Even with the approximation for g_c this still doesn't quite tell us anything about s_I, which remember is what we're examining to assure that capital accumulation supports the stability of the BGP we see in the data. We already know from last chapter that a valid and necesary form for the utility function is $U(c) = c^{1-\sigma}/(1-\sigma)$, so it is obvious to work with that. In that case the Euler equation becomes

$$\frac{c_1^{-\sigma}}{c_2^{-\sigma}} = \beta(1+r) \tag{5.15}$$

or

$$\frac{c_2}{c_1} = [\beta(1+r)]^{1/\sigma} \equiv 1 + g_c. \tag{5.16}$$

Like our more approximate answer above, this says that the growth rate of consumption (or the ratio) depends on the relative time preference rate and rate of return ($\beta(1+r)$), and how sensitive people are is determined by the IES, $1/\sigma$. If $\beta(1+r) > 1$, then it must be that $c_2 > c_1$, or consumption grows. If $\beta(1+r) < 1$, then $c_2 < c_1$ and consumption shrinks. But changes in r will change growth in consumption depending on how big $1/\sigma$ is. People with low elasticities will not respond much, and people with high elasticities will respond strongly to r.

Use that Euler equation in the budget constraint (the last FOC from the Lagrangian, technically) to plug in for

$$a_1(1+r) + w_1 + \frac{w_2}{1+r} - c_1 - \frac{1+g_c}{1+r}c_1 = 0 \tag{5.17}$$

and this resolves to

$$c_1 = \frac{1}{1+\frac{1+g_c}{1+r}}\left(a_1(1+r) + w_1 + \frac{w_2}{1+r}\right). \tag{5.18}$$

The thing in parentheses is total lifetime wealth, and note that c_1 doesn't depend on when wealth arrives, just how big it is, as mentioned in 5.2.

The initial fraction is the share of total lifetime wealth that the person consumes in period 1. This fraction depends on the optimal growth rate of consumption, and the higher is that growth rate the lower is this share, which makes sense. If you want fast growth of consumption, you have to start relatively small. That extra $1 + r$ in the ratio is capturing the fact that it's not just whether you want consumption to grow fast, but whether you want consumption to grow fast *relative* to how fast you could grow consumption by saving. If you want to grow consumption faster than r, you need to push down current consumption a lot. If you are fine letting consumption grow slower than r, you can actually consume a lot today, because r will ensure you get a good amount tomorrow as well.

There are multiple ways that the rate of return influences the consumption choice.

If you have a few hours, go ahead and try and take the derivative of s_I with respect to $1 + r$ and see that there is no simple solution.

- Substitution effect. Think of this as the impact of a change in r on g_c, which is purely about the ratio of "prices" in the Euler equation. An increase in r raises the price of consumption in period 1 because you are foregoing future consumption. You substitute away from c_1 and hence g_c goes up in response. You save more in response.

- Income effect. Think of this as the effect of a change in r acting on the $1 + r$ in ratio $(1 + g_c)/(1 + r)$. An increase in r makes you feel richer, and expands your options. This might lead to a higher consumption c_1, and therefore might lead you to save *less*. If $\sigma > 1$, which is what we tend to assume, then this income effect "wins" versus the substitution effect, and an increase in r tends to raise consumption and lower savings. A high value of σ means that g_c does not respond much to changes in r, leading to the income effect dominating.

- Wealth effect. The last effect comes from lifetime wealth. If r goes up, then this *lowers* the present discounted value of w_2, as it comes in the future. Wanting to smooth consumption, you need to consume less, and end up saving more for this future that is less wealthy.

I was loose about the term "savings" there, because we have to be careful. If we want to think about how s_I works, we need to remind ourselves that s_I is a macro/production side concept dealing with the ratio of capital accumulation to GDP. In our consumption model the various individual terms do not necessarily map directly to this. We're going to hold off on linking things back to s_I for now, as this mapping gets easier once we expand out and incorporate this into the wider production side. Right now, our consumers don't perceive any notion of the macroeconomy - they take r and such as given - and just make decisions.

One useful notion, though, is whether people choose to have consumption bigger or smaller than their initial period resources. In period 1 they have $(1+r)a_1 + w_1$ available to them. An interesting question is whether they choose to consume more or less than this particular amount. It will be the case that $c_1 < (1+r)a_1 + w_1$ if

$$\frac{1}{1+\frac{1+g_c}{1+r}}\left(a_1(1+r) + w_1 + \frac{w_2}{1+r}\right) < (1+r)a_1 + w_1 \qquad (5.19)$$

or if

$$\frac{w_2}{(1+r)a_1 + w_1} < 1 + g_c. \qquad (5.20)$$

People will save some of their initial period resources if the ratio of second period exogenous income (w_2 just happens to them without any choice on savings) is sufficiently small compared to total first period resources. How small that has to be depends on preferences embedded in g_c. If people want high growth in consumption, this is easier to meet as a condition, and people naturally save some of their first period resources. If g_c is low, however, it is more likely that they "eat" some first period resources and just wait for w_2 to arrive in the second period.

This helps explain some of the relationship between r and g_A we saw in the last chapter. To support a BGP r had to be equal to $\theta + \sigma g_A$. The value of g_A is going to dictate how big w_2 will be in relationship to period 1 resources; if the economy grows fast so does w_2. If growth is fast enough this might convince people to just eat some first period resources and not accumulate anything at all. Just wait for economic growth to make you rich. So the only way to support a BGP is for the rate of return to be high enough that their preferences in g_c create a big growth rate of consumption, which can only happen if they don't eat those resources. A stable BGP depends on the production side yielding a high enough rate of return to convince everyone to keep accumulating even as the economy grows anyway because of g_A.

5.3 *Arbitrary time periods*

The game now is to extend this consumption problem out to an arbitrary number of periods, T, which could include up to an infinite number of periods. If T is the end point, then the last period with wages or consumption is $T-1$. We have to do a little work her to write down the appropriate budget constraints, but once we've done that the essential nature of the problem remains identical.

Don't get weird about infinite time periods. You can think of people who care about their kids (and hence their kids kids, etc.) or you can think of people who, conditional on being alive, have a non-zero chance of living another day, meaning there is no upper limit to lifespan, even if the probability of living gets quite low.

Definition 5.1 (Dynamic budget constraint) *The dynamic budget constraint is*

$$a_{t+1} = w_t + (1 + r_t)a_t - c_t, \tag{5.21}$$

holds for every period $t = 0, ..., T - 1$

We will add to this an assumption that governs the individual problem

Assumption 5.4 (Terminal asset condition) *It must be that*

$$a_T \geq 0 \tag{5.22}$$

Those two constraints can be combined to say the following

If we had uncertainty about life expectancy, then we could allow for some people to die "early" compared to expectations, and we could then add an insurance market which allowed people to insure against this so that in expectation $a_T = 0$ and in aggregate it held exactly. We could also allow for bequests, so that people wanted to leave $a_T > 0$ to their kids, and then we'd think about utility for a family which would go out to infinity.

Conclusion 5.3 (Lifetime budget constraint) *Given the dynamic budget constraint and the terminal asset condition, it must be that*

$$a_0 + \sum_{t=0}^{T-1} w_t \Pi_{s=0}^{t}(1 + r_s)^{-1} \geq \sum_{t=0}^{T-1} c_t \Pi_{s=0}^{t}(1 + r_s)^{-1} \tag{5.23}$$

or the PDV of lifetime wealth must be greater than or equal to the PDV of lifetime consumption.

This thing is kind of ugly because of the combined discount factors. If we assume that all $r_s = r$, then this becomes simpler as in

$$a_0 + \sum_{t=0}^{T-1} \frac{w_t}{(1 + r)^t} \geq \sum_{t=0}^{T-1} \frac{c_t}{(1 + r)^t}. \tag{5.24}$$

Lifetime utility is now

$$V = \sum_{t=0}^{T} \beta^t U(c_t). \tag{5.25}$$

There's a few ways to solve this overall, including "plug and chug" with the budget cosntraint and lifetime utility, although that gets rough. Let's stick with the Lagrangian, as again that sets us up to think about more complex problems. But we have to be a little more careful here as we do not have a lifetime budget constraint, just the dynamic constraint. That means there are really T constraints, one for each period, so we have

$$\mathcal{L} = \sum_{t=0}^{T} \beta^t U(c_t) + \lambda_t \left(w_t + (1 + r_t)a_t - c_t - a_{t+1} \right). \tag{5.26}$$

Taking first-order conditions gets a little trickier here. We want to look at the FOC for a single given period, t. That means we're taking FOC with respect to c_t, a_{t+1}, and λ_t:

$$\beta^t U'(c_t) - \lambda_t = 0 \tag{5.27}$$

$$-\lambda_t + (1 + r_{t+1})\lambda_{t+1} = 0 \tag{5.28}$$

$$w_t + (1 + r_t)a_t - c_t - a_{t+1} = 0 \tag{5.29}$$

Pay attention to the middle one, because a_{t+1} shows up in two separate dynamic budget constraints. It's the link over time, and so increasing assets lowers utility today (by λ_t) and raises it tomorrow (by $(1 + r_{t+1}\lambda_{t+1})$. We also know that if we did these FOC with respect to period $t+1$, we'd get a FOC for consumption that looks like

$$\beta^{t+1}U'(c_{t+1}) - \lambda_{t+1} = 0. \tag{5.30}$$

We can put all of this together into a solution. First, the condition on assets means that

$$1 + r_{t+1} = \frac{\lambda_t}{\lambda_{t+1}} \tag{5.31}$$

and therefore if the rate of return is positive then it means the marginal utility of wealth must be going *down* over time. That makes some sense. Wealth in the future should be worth a little less than wealth today in utility terms, because it involves waiting. The rate r_{t+1} is the rate at which the market will transfer utility over time.

Take the conditions on c_t and c_{t+1} together to get

$$\frac{U'(c_t)}{\beta U'(c_{t+1})} = \frac{\lambda_t}{\lambda_{t+1}} \tag{5.32}$$

which says that the ratio of marginal utilities must be equal to the ratio of marginal utilities of wealth. That also should make sense. This is saying the value of consumption in utility terms between periods needs to match the ratio of marginal utility of wealth. If it didn't, you could move consumption around to take advantage of the difference. This is probably the more accurate way to think about the ratio of marginal utilities equal to the ratio of *prices*, and the prices here are marginal utilities of *wealth* in the two periods.

We know the ratio of marginal utilities of wealth is bigger than one if $r_{t+1} > 0$. And thus the ratio of marginal utilities must be bigger than one, or MU in t must be bigger than MU in $t+1$. But because of the β term, this doesn't mean $U'(c_t)$ is necessarily bigger than $U'(c_{t+1})$. It's going to depend on how big β is, or how much you care about period $t+1$ compared to t.

Put all of this together into the Euler equation. In 5.1 we already solved for this for a two-period problem, but being pedantic lets write down

Conclusion 5.4 (Consumption Euler Equation) *The consumption Euler equation relates the marginal utility of consumption in two periods to the rate of return and the time preference rate:*

$$\frac{U'(c_t)}{U'(c_{t+1})} = \beta(1 + r_{t+1}) \tag{5.33}$$

The fact that the Euler equation doesn't depend on the size or timing of lifetime wealth ties this to Friedman's "Permanent Income Hypothesis" and Modigliani's "Lifecycle Hypothesis". In the PIH changes in income *today* will have a small impact on consumption *today* because they don't change the Euler equation. In the LCH your choice of consumption isn't bound to when you earn money, so you borrow while young and live off savings when old.

The same proposition on timing of wealth holds as in 5.2. This Euler equation does not depend on the timing or size of lifetime wealth.

With BGP preferences the consumption Euler equation is

$$\frac{c_{t+1}}{c_t} = [\beta(1+r_{t+1})]^{1/\sigma} \equiv 1 + g_{c,t+1} \tag{5.34}$$

where the growth rate of consumption is now specific to period t because the rate of return is specific to period $t+1$.

The remaining part of this problem is to choose the actual amount of initial consumption, c_0, to start with, such that your choice of c_0 and the time path of c_t implied by the Euler equation ensures that you stay within the lifetime budget constraint.

In theory one could work this out via pencil and paper, given the string of rates of return, wages, and initial assets. In practice this is tedious and hard, and we can use a computer to solve this out. Simplifying things via things like a constant rate of return or constant growth rate of the wage will make things much easier.

5.4 *Infinite periods*

Let's extend this to infinite periods. For lifetime utility, there is no issue, this just becomes

$$V = \sum_{t=0}^{\infty} \beta^t U(c_t). \tag{5.35}$$

For the dynamic budget constraints we continue to have that

$$a_{t+1} = w_t + (1+r_t)a_t - c_t, \tag{5.36}$$

at any given time t.

But now we don't have $a_T \geq 0$, because there is no T at which the problem stops. The equivalent condition in infinite time is

Assumption 5.5 (No Ponzi Condition) *The "No Ponzi Game" condition requires that:*

$$\lim_{t\to\infty} a_t \Pi_{s=0}^{t}(1+r_s)^{-1} \geq 0. \tag{5.37}$$

This condition says that the present discounted value of assets has to remain positive or zero as time goes to infinity. It's not that one cannot accumulate debt at any point during the infinite number of periods, it's that the present discount value has to vanish over time. The big mess of $\Pi_{s=0}^{t}(1+r_s)^{-1}$ is the series of market discount rates applied to the value of a_t. If $r_s = r$ for all periods, then this is the simpler term $a_t/(1+r)^t$.

Charles Ponzi got temporarily rich through a chain letter scheme in the 1920's. The modern equivalent in Bernie Madoff, and there are occassional instances of these all over the world.

I'm being loose here, but for all intents and purposes you can just take the same lifetime budget constraint with T periods and take the limit at $T \to \infty$ to get that

$$a_0 + \sum_{t=0}^{\infty} w_t \Pi_{s=0}^{t} (1 + r_s)^{-1} \geq \sum_{t=0}^{\infty} c_t \Pi_{s=0}^{t} (1 + r_s)^{-1}. \tag{5.38}$$

Everything works out the same as with a large T problem, you just need to assert that the left-hand side of the constraint here adds up to something finite.

If you assume $r_s = r$, then lifetime wealth will be finite so long as w_t doesn't grow too fast. if $w_t = w_0(1 + g_w)^t$, you'd need to have $g_w < r$ for finite wealth.

5.5 *Continuous time*

The continuous time version of this problem looks similar in most respects. Lifetime utility is now

$$V = \int_0^{\infty} U(c) e^{-\theta t} dt \tag{5.39}$$

so that consumption is continuously discounted at the rate θ. I've discarded any t subscript or notation because it kind of gets annoying, but you can understand that c is in fact $c(t)$. The dynamic budget constraint is

$$da = ra + w - c \tag{5.40}$$

where da is the instantaneous change in a, and it depends on the flow of income, $ra + w$, and the choice of instantaneous consumption, c. Again, I dropped the t. Don't get too hung up on "instantaneous", just think of doing discrete time in very, very, small increments. In addition to this we need a No-Ponzi requirement to keep accumulation sensible, and this is

$$\lim_{t \to \infty} a(t) e^{-\int_0^t r(s) ds} \geq 0. \tag{5.41}$$

The exponential term is the continuous equivalent of the string of discount rates run from time 0 to time t. I'm specific about $a(t)$ here because this limit is keyed off of the time. All this says is that the PDV of assets has to remain non-negative as time goes to infinity. Same logic as before.

The continuous time equivalent problem is solved using a Hamiltonian, rather than a Lagrangian. They are not quite the same, as the Hamiltonian is about the current value of the objective (utility) and not lifetime value. It's going to tell us what you have to do any any given moment t to optimize.

This is a *current* value Hamiltonian, so it's the utility value of $H(t)$ from the perspective of time zero. You can write down and solve a *present* value Hamiltonian and get the same answers.

$$H(t) = e^{-\theta t} u(c) + \lambda (ra + w - c). \tag{5.42}$$

The Hamiltonian is thus a function of t. The first term is the utility value of consumption at some time t, which note depends on the discount rate.

There are a standard set of conditions we evaluate H with. The value of λ is again a marginal value of relaxing the constraint.

1. Maximize with respect to control variable, c.

$$\frac{\partial H}{\partial c} = e^{-\theta t} u'(c) dc - \lambda = 0 \tag{5.43}$$

   Notice that this is the same construct as what we got with a Lagrangian, as it relates the MU of consumption to the marginal utility of wealth.

2. Evaluate the change in the value of the multiplier due to a change in assets

$$d\lambda = -\frac{\partial H}{\partial a} = -\lambda r \tag{5.44}$$

   where that negative sign shows up because we're varying a while holding H constant. This is again showing us that the rate of return tells us something about how the multiplier changes, as this implies $r = -d\lambda/\lambda$, or the rate of return is the growth rate of the multiplier.

3. Evaluate the change in the state variable and recover the dynamic constraint,

$$da = \frac{\partial H}{\partial \lambda} = ra + w - c. \tag{5.45}$$

You solve these together. From the first condition you can get another expression for $d\lambda$

$$-\theta e^{-\theta t} u'(c) + e^{-\theta t} u''(c) dc - d\lambda = 0 \tag{5.46}$$

and use with the second condition to get

$$-\theta e^{-\theta t} u'(c) + e^{-\theta t} u''(c) dc = -\lambda r \tag{5.47}$$

and using the first condition again we get

$$-\theta e^{-\theta t} u'(c) + e^{-\theta t} u''(c) dc = -e^{-\theta t} u'(c) r \tag{5.48}$$

which solves down to

$$u''(c) dc = \theta u'(c) - u'(c) r \tag{5.49}$$

and to

$$\frac{dc}{c} g_c = (r - \theta) \left[-\frac{u'(c)}{u''(c)c} \right] \tag{5.50}$$

and this should all look familiar now. The left side is the growth rate of consumption. The right side is $r - \theta$ times the IES. This is the formal derivation of the intuitive version of the consumption problem from the last chapter. You'd combine this with the No-Ponzi condition to get a solution for consumption levels at any given point. The same rules apply here, which is that solving for the actual value of consumption will be tedious and we tend to rely on numerical answers from the computer.

6
Exchange Economies

The consumption problem we set up is not a complete model in the sense that there is a price - r - that is just taken as given. That might make sense for the individual (they don't control the price) but we don't have any way to determine where r came from. We also don't have any way to know where the string of w earnings come from, but often we are willing to take that as a given endowment, especially as for our consumer here there is not choice over how much to work, so w doesn't really constitute the price of their time (it could and we could model that).

One way we could endogenize r and make that price an equilibrium of supply and demand is to merge this consumer problem with the production side from the Solow model. There the rate of return is determined by K/Y, and the value of K/Y would depend on their choices. That will be the neo-classicial model.

But before we get to the neo-classicial model there are other ways to think about endogenizing r, meaning other ways to think about there being a supply and demand for assets. These are useful in other contexts and come up often in macro research. In these settings the focus is typically more on asset prices or other aspects of the economy and they are less concerned with growth or stability (or they take those as givens). We often refer to these settings as *endowment economies* because the stream of income w and initial assets a_0 are all just taken as given. They could even have uncertainty around them, but nothing the individual does (or the group of individuals does) will influence the pattern of that income, unlike the neo-classical model.

In these economies people still want to borrow and save, potentially. But now the only way to borrow and save is to find another person who wants to do the opposite. In the neo-classical model you'll be able to borrow or save with yourself in the sense that your savings (anything you don't consume) will necessarily still be there tomorrow in the form of capital. Here, anything you don't eat disappears, so you have to rely on others to alter your consumption path.

So now we have to find a way to construct an equilibrium where everyone is acting according to their consumption problem maximization rules - the Euler, etc. - but in which the net decisions are such that there is no net borrowing or lending in the economy in any given period. You and I can trade, but we can't trade off against the future.

Because these economies explicitly have different people or types of people we often refer to them as *exchange economies* and they are nothing more than the typical two-person markets you set up in intermediate or talked about in terms of Edgeworth boxes and the welfare theorems.

The last thing about these economies that we'll add is that we will be more explicit about the kind of equilibrium we'll set up. That is, we need to define the situation where supply and demand equalize. Supply and demand of what? It depends on how we allow these people to trade. The different setups are useful to know because they come up in other contexts, including the neo-classical. In one case we presume that at period zero people can make any kind of trades about any future periods they want, and that those trades are binding and impossible to cancel. These are called "Arrow-Debreu assets" and this is an "Arrow-Debreu equilibrium" where the supply and demand of those assets is in equilibrium and we will see how those prices dictate the path of consumption.

A different concept is that the only assets that can be traded are one-period assets, meaning I can only write a contract to borrow or lend to you today with a payoff in the next period. It's still perfectly enforced, but I can't look ahead, and at time zero I cannot make all the possible trades I want to make. This is called a "sequential equilibrium" and it is implicitly what we looked at last chapter. A key point is that the sequential and Arrow-Debreu equilibrium's give you the same answer in our narrow consumption problem, which is useful because it means we can leverage that to solve problems later in one form or the other. It's often easier to think about Arrow-Debreu equilibria (all trades take place at once) rather than worrying about sequential trades. That will fail at some point when we add uncertainty (you can't know what to trade at time zero if you don't know what will happen), but even then there are cases where the AD still holds up.

6.1 An Arrow-Debreu Equilibrium

The first step is that we can re-cast the consumption problem in terms of prices, not returns. Let lifetime utility be over a set of J goods, and there is a preference weight of β_j on each one (note not raised to a power or anything) and the utility function $U(c_j)$. So goods differ because β_j might be higher or lower, but the marginal utility of each one works in the same way, with $U(c_j)$ having normal properties.

$$V = \sum_{j=0}^{J} \beta_j U(c_j) \tag{6.1}$$

The only change is the lifetime budget constraint will be written in terms of

$$\sum_{j=0}^{J} p_j w_j \geq \sum_{j=0}^{J} p_j c_j. \tag{6.2}$$

This person has an endowment of w_j of each good, and that is their total income or wealth, and they want to spend that on the various consumption goods. This is nothing more than a slightly complicated intermediate micro problem.

You can set up the Lagrangian and get a relationship like this for any two goods i and j

$$\frac{\beta_i U'(c_i)}{\beta_j U'(c_j)} = \frac{p_i}{p_j}. \tag{6.3}$$

The ratio of marginal utilities should equal the ratio of prices. If we use our normal assumption about utility with CRRA then it will be the case that

$$\frac{c_j}{c_i} = \left(\frac{\beta_j}{\beta_i}\frac{p_i}{p_j}\right)^{1/\sigma}. \tag{6.4}$$

Relative consumption of the two goods depends on their relative price, and the sensitivity to the ratio of prices depends on σ. $1/\sigma$ represents their willingness to substitute between goods.

Let's say we have two people A and B, the only two people in the economy. They each have first-order conditions of

$$\frac{c_j^A}{c_i^A} = \left(\frac{\beta_j}{\beta_i}\frac{p_i}{p_j}\right)^{1/\sigma} \tag{6.5}$$

and

$$\frac{c_j^B}{c_i^B} = \left(\frac{\beta_j}{\beta_i}\frac{p_i}{p_j}\right)^{1/\sigma}. \tag{6.6}$$

We could let their β values differ too, but that just gets tedious to keep track of without any real intuition.

It's also got to be the case that their consumption totals for a given good j have to add up to the endowment,

Definition 6.1 (Endowments) *In an endowment economy total consumption of individuals of good j has to add up to the total endowment of j*

$$c_j^A + c_j^B = w_j^A + w_j^B. \tag{6.7}$$

Using the first-order conditions it would be that

$$\left(\frac{\beta_j}{\beta_i}\frac{p_i}{p_j}\right)^{1/\sigma}(c_i^A + c_i^B) = w_j^A + w_j^B, \tag{6.8}$$

and note that because of the endowments this must be

$$\left(\frac{\beta_j}{\beta_i}\frac{p_i}{p_j}\right)^{1/\sigma}(w_i^A + w_i^B) = w_j^A + w_j^B \tag{6.9}$$

or that

Conclusion 6.1 (Relative price) *The relative price of two goods i and j is determined by*

$$\frac{p_i}{p_j} = \frac{\beta_i}{\beta_j}\left(\frac{w_j^A + w_j^B}{w_i^A + w_i^B}\right)^{\sigma}. \tag{6.10}$$

The equilibrium price of i to j depends on a few things. First, it depends on the relative preference for i and j. If you care more about i it will have a higher price because A and B will bid up the price on it. It also depends on the relative supply, given the endowment. If there is more j that will drive down the relative price of j, and vice versa. All this equation says is that supply and demand influence the price. The degree to which the supply changes the price depends on σ, how willing customers are to substitute. We can write down a relative price ratio like this for any two goods.

There should not be anything mysterious about this problem, it just has a lot of products to think about. The key thing we want to see is that we can frame the inter-temporal consumption problem in just this way. It's just a change to think about products as "consumption in period j" rather than pizza or beer or whatever.

The only mental change we have to make is that the *order* of the J goods is important. If J are time periods then we can't just arbitrarily reorder the numbering. In a regular static problem pizza could be good 1 and beer good 2, but we just as easily let beer be good 1 and pizza good 2. Here, there is something specific about goods i and j and we can't just call them j and i.

Just like the pizza/beer problem, our individual is making a static decision over a set of ... well, securities or assets rather than produts. When they pay p_j they are buying an asset that will allow them to consume one unit of consumption in period j. They have an endowment now not of pizza and beer *today*, but some expected income in period j of w_j (they could consume this if they wanted), and they can sell an asset to someone else promising them one unit of consumption in period j based on that income.

Definition 6.2 (Arrow-Debreu asset) *As Arrow-Debreu asset i is an asset that sells for price p_i in period zero, and which delivers one dollar (or one unit of consumption) in period i in the future.*

The AD asset is a concept that turns an inter-temporal good into something that is bought and sold *today*. This isn't entirely hypothetical, of course, and this is just what most financial products are. You pay a price today to deliver some money (which can be used to buy consumption) in the future. A complete set of Arrow-Debreu securities or assets means that an individual can move consumption around from any period i to any period j. This makes a full set of Arrow-Debreu securities equivalent to assuming that one can borrow and lend at will. That need not mean that the interest rate for all that borrowing and lending is the same for all periods, just that one can move money around.

If we think about these as assets, then we can think about their rate of return. Think about the annualized rate of return, r_i, you get on an asset i that is sold for p_i in period zero, and that pays off one dollar in period i.

Definition 6.3 (Return on AD assets) *The annualized rate of return on an AD asset that pays off in period i is defined by*

$$p_i(1+r_i)^i = 1, \tag{6.11}$$

and note that $p_0 = 1$ and $r_0 = 0$ by definition.

Now, let's be more specific about some things and see how this relates to our overall consumption problem from before. Keep in mind that the order of j matters, as these are time periods. Let's say that the preference weights are $\beta_j = \beta^j$. Let's say for simplicity that $w_j^A + w_j^B = (1+g)^{j-i}(w_i^A + w_i^B)$ or that the total endowment grows at the rate g every period. This means that the endowment in period i from the perspective of period zero is $w_i^A + w_i^B = (1+g)^i(w_0^A + w_0^B)$.

Together this means that if we compare the price ratio in i to zero, p_i/p_0, using what we know this is

$$\frac{1}{(1+r_i)^i} = \frac{\beta^i}{1}\left(\frac{1}{(1+g)^i}\right)^{\sigma}. \tag{6.12}$$

Raise this to the $1/i$ power on both sides, and flip over, and you get

$$1+r_i = \frac{1}{\beta}\left((1+g)\right)^{\sigma}, \tag{6.13}$$

or that the annualized rate of return you'd get for any arbitrary choice of i $(1+g)^{\sigma}/\beta$. The annualized rate of return for every asset in this simplified economy is the same because the endowment grows at a constant rate and the preference weights change at a constant rate. So all $r_i = r$.

If you re-arrange this you get

$$1+g = (\beta(1+r))^{1/\sigma}. \tag{6.14}$$

The growth rate of consumption in this Arrow-Debreu exchange economy *has* to equal the growth rate of the endowment because there is no other option for people. That growth rate dictates the rate of return that has to hold such that everyone in this economy acts to ensure that their optimal choices add up to match this constraint.

We can formalize this equilibrium concept, which note is just a way of saying we are formalizing the problem we are writing down and the things we are solving for. Of note, it specifies who knows what and who takes what as given.

Definition 6.4 (Arrow-Debreu Equilibrium) *An Arrow-Debreu Equilibrium is a sequence of prices $\{p_t\}_{t=0}^{\infty}$ and choices $\{c_t^i\}_{t=0}^{\infty}$ for all individuals i such that*

1. *Given $\{p_t\}_{t=0}^{\infty}$ the choices $\{c_t^i\}_{t=0}^{\infty}$ maximize utility for an individual i in (6.1)*

2. *..subject to the constraint (6.2) for individual i*

3. *and that markets clear in each period t so that $\sum_i c_t^i = \sum_i w_t^i$ as in 6.1.*

All this has done is give us a way of thinking about how our individual consumption problem people could co-exist in an economy where things have to add up.

6.2 Sequential markets

That's one way of conceiving of how an exchange economy could work. But it requires us to believe there are a complete set of securities out there allowing everyone to make all possible future trades, right now. You might not like that, or you might easily imagine situations where that would break down (think of where there is a lot of uncertainty). An alternative is to look at an economy where individuals can trade with one another where to buy/sell assets today in exchange for returns tomorrow, and every period they can do that again. But they cannot go out into the future more than one period. Again, we'll assume all the assets are perfectly enforced.

Definition 6.5 (One period asset) *In period j there is an asset that can be purchased for price q_j that pays off one dollar (or one unit of consumption) in period $i+1$.*

They still maximize utility as in 6.1,

$$V = \sum_{j=0}^{J} \beta_j U(c_j) \tag{6.15}$$

but now we're thinking about dynamic budget constraints, not a lifetime budget constraint. In each period they see this dynamic constraint

$$q_j a_{j+1} = w_j + a_j - c_j \tag{6.16}$$

where they have some income/endowment of w_j, they have some financial holding a_j to begin with, and they choose some consumption c_j or they buy a_{j+1} units of the asset for next period at the price q_j (timing and notation here is tricky). The a_j was the payoff or cost of the asset they bought or sold *last* period, and it either added or subtraced from their ability to consume today.

Set up the Lagrangian and solve for their first-order conditions and you have

$$\beta_j U'(c_j) - \lambda_j = 0 \tag{6.17}$$

$$\beta_{j+1} U'(c_{j+1}) - \lambda_{j+1} = 0 \tag{6.18}$$

$$\lambda_{j+1} - \lambda_j q_j = 0 \tag{6.19}$$

as first-order conditions. You can re-arrange these to be that

$$\frac{\beta_j U'(c_j)}{\beta_{j+1} U'(c_{j+1})} = \frac{1}{q_j} \tag{6.20}$$

In this case, the ratio of marginal utilities has to equal the ratio of prices, as usual. The price of consumption in j is one, and the price of consumption in period $j+1$ is q_j. Using the CRRA assumption it would be that

$$c_{j+1} = \left(\frac{\beta_{j+1}}{\beta_j} \frac{1}{q_j} \right)^{1/\sigma} c_j. \tag{6.21}$$

This has to hold for each person in this economy, so there is a first-order condition like this for each. And it has to be that their consumption still adds up each period as in 6.1,

$$c_j^A + c_j^B = w_j^A + w_j^B, \tag{6.22}$$

and therefore it must be that $a_{j+1}^A + a_{j+1}^B = 0$. That is, there can't be any case where both individuals think they can consume more than their endowment in a period.

The only way to violate this would be if $a_0^A + a_0^B > 0$ and there were some assets sitting around to start with.

Combining the first-order conditions and this adding up condition it has to be that

$$w^A_{j+1} + w^B_{j+1} = \left(\frac{\beta_{j+1}}{\beta_j} \frac{1}{q_j} \right)^{1/\sigma} (w^A_j + w^B_j) \tag{6.23}$$

or that

$$q_j = \frac{\beta_{j+1}}{\beta_j} \left(\frac{w^A_j + w^B_j}{w^A_{j+1} + w^B_{j+1}} \right)^{\sigma}. \tag{6.24}$$

Which looks a lot like the AD price equation. The price of an asset that pays off tomorrow depends on the relative preference weight of tomorrow versus today, and it depends inversely on the endowments, modified by how much you are willing to respond to relative prices.

Definition 6.6 (Rate of return) *The rate of return on the one period asset with price q_j that pays off one dollar the next period is defined by*

$$q_j(1 + r_j) = 1 \tag{6.25}$$

Given that definition the whole pricing equation is

$$\frac{1}{1 + r_j} = \frac{\beta_{j+1}}{\beta_j} \left(\frac{w^A_j + w^B_j}{w^A_{j+1} + w^B_{j+1}} \right)^{\sigma} \tag{6.26}$$

Assume again that $\beta_j = \beta^j$, and that $w^A_{j+1} + w^B_{j+1} = (1+g)(w^A_j + w^B_j)$ and we have

$$\frac{1}{1 + r_j} = \frac{\beta}{1} \left(\frac{1}{(1+g)} \right)^{\sigma} \tag{6.27}$$

which means that

$$1 + r_j = \frac{(1+g)^{\sigma}}{\beta}. \tag{6.28}$$

Again, this doesn't depend on the time period j we choose, so $r_j = r$ for all the time periods. Again, the rate of return is identical over time here and that rate of return has to work such that it ensure total consumption grows at the rate of the endowment. You can re-arrange that to

$$1 + g = (\beta(1+r))^{1/\sigma} \tag{6.29}$$

or that the Euler equation must deliver a growth rate of $1 + g$ as the choice for individuals.

Definition 6.7 (Sequential Markets Equilibrium) *A sequential markets equilibrium is a sequence of prices $\{q_t\}_{t=0}^{\infty}$ and choices $\{c^i_t\}_{t=0}^{\infty}$ and $\{a^i_t\}_{t=0}^{\infty}$ for all individuals i such that*

1. *Given $\{p_t\}_{t=0}^{\infty}$ the choices $\{c^i_t\}_{t=0}^{\infty}$ and $\{a^i_t\}_{t=0}^{\infty}$ maximize utility for an individual i in (6.1)*

2. *..subject to the dynamic constraints (6.16) for individual i*

3. *and that markets clear in each period t so that* $\sum_i c_t^i = \sum_i w_t^i$

4. *.. and that* $\sum_i a_t^i = 0$ *(no net assets)*

This is again just a way of consolidating information about how our model is structured, and what people are choosing and what they know. It helps clarify how to arrive at the solution.

Conclusion 6.2 (Equivalence of AD and Sequential Markets) *Given an AD equilibrium with* ${p_t}_{t=0}^{\infty}$ *and* ${c_t^i}_{t=0}^{\infty}$*, then there exists a sequential markets equilibrium with* ${q_t}_{t=0}^{\infty}$ *such the solution for* ${c_t^i}_{t=0}^{\infty}$ *is identical. The reverse holds. Given a sequential markets equilibrium with* ${q_t}_{t=0}^{\infty}$ *and* ${c_t^i}_{t=0}^{\infty}$*, there exists an AD equilibrium with* ${p_t}_{t=0}^{\infty}$ *such that the solution for* ${c_t^i}_{t=0}^{\infty}$ *is identical.*

The choice on how to structure the financial markets for these individuals isn't important, you get the same answer. That's true given these limited parameters. Probably the most important one is that all these AD assets are perfectly enforceable, meaning a contract made at time zero can be enforced in period 1,938. Once that fails, so does this equivalence.

7
Dynamic General Equilibrium

We will approach this decision in this chapter from the perspective of a single decision-maker. You can conceive of this person as a "central planner" or "social planner" who makes decisions about s_I taking into account all the ramifications of those decisions. An alternative is to conceive of this decision-maker as a "representative agent" whose decisions reflect how everyone acts. Neither of those are realistic descriptions of economies in any sense, of course, but this approach will give us a decent way of understanding the choices involved. Then we can examine the conditions under which a decentralized economy of consumers, firms, etc. with limited information will reach the same aggregate decisions.[1] This model is something I usually refer to as the "Ramsey model", but you can find it referred to at times as the "Cass-Ramsey-Koopmans" model, or the "neo-classical growth model".

[1] F. P. Ramsey. A mathematical theory of saving. *The Economic Journal*, 38(152):543–559, 1928; David Cass. Optimum growth in an aggregative model of capital accumulation. *The Review of Economic Studies*, 32(3):233–240, 1965; and Tjalling Koopmans. On the concept of optimal economic growth. In *The Economic Approach to Development Planning*, pages 225–287. Rand McNally, 1965

7.1 The intuition of the DGE model

Before we get into math, let's go through what we're doing. We want to take the Solow model and make the s_I choice endogenous, and from there we want to know if by making it endogenous that we still have a BGP and that the economy is still stable around that BGP. A better way to say this is that we want to understand what conditions on individual consumption choices ensure stability, as that seems to be what we see in the data.

Mathematically and economically this is going to get complicated because the production side stuff from the Solow model will depend on the consumption choice, and the consumption choice will depend on the production side stuff. Individuals are choosing c_t based on a time path of r_{t+1}, and that return depends on the capital/output ratio. The capital/output ratio depends, through s_I, on the choice of c_t.

The consumption problem comes down to, in large part, choosing the initial value c_0, and then letting the Euler equation go from there. Any choice of c_0 will give you *a* solution for c_t across time. That is, if you just plug in *some* value for c_0 you can iterate the Euler equation and capital accumulation equation to generate sequences for c_t and K_t/Y_t and just see what happens. You could do it by hand, but it's also easy to do this on a spreadsheet or computer program.

Not every choice of c_0 will be the optimal solution from the perspective of the individual, though. If we choose c_0 too high to start with, then ultimately the path will become infeasible, in the sense that your spreadsheet or program will want to make K/Y negative at some point. Consumption in this case is too high, and because it is high s_I is too low, and because of that K/Y starts to decline as people consume existing capital to keep up. But as K/Y gets lower the rate of return on capital goes up, which just makes people want to have higher growth in consumption, which exacerbates the problem. This is the Ponzi game situation. If c_0 is too high to start we can keep consuming more but you cannot keep this up.

If you choose c_0 too low you end up with the opposite problem. s_I is now big and K/Y grows rapidly, but that means the rate of return goes down a lot. A low return means you want slow consumption growth, which means you never really consume enough to keep the K/Y ratio in check. As time goes on the value of the assets you hold remains positive even out to infinite time, and remember we don't care about the future that much. This doesn't make sense because you could raise c_0 and have higher consumption at every point in time and still not violate the no-Ponzi condition.

Somewhere in between these situations is *the* solution for c_0 which delivers a time path for consumption that maximizes lifetime utility. This solution ensures that consumption does not go to infinity (infeasible) or to zero (stupid), but rather ends up at growing at a nice constant rate forever. In other words, it ensures we end up at a steady state. That's good. It tells us that having individuals make consumption choices is still consistent with their being a BGP.

But this still doesn't necessarily tell us if that BGP is stable. If we make the "right" choice of c_0 does that always push us back towards the BGP? The answer is going to be yes. The economy is stable around the BGP when the consumption choices of individuals lead to the capital accumulation rate working to push the capital stock back towards steady state. Why does that work here? If K/Y is below steady state, then this would imply that the rate of return is high, and in that case the individuals would want a high growth rate of consumption via the Euler equation. A high growth rate of consumption is only feasible if you start with relatively low consumption today, which means a relatively high accumulation rate for capital.

A high accumulation rate for capital pushes K/Y up, towards the steady state. Mathematically, the exact form this takes is not solvable with pencil and paper, but the core idea is that even with individual choices over consumption (and thus s_I) it remains that $g_{K/Y}$ is negatively related to K/Y.

Let's think again about instability. The Ramsey or neo-classical model tells us a few things that do *not* cause instability, even though intuitively we might have thought they would. One is that people are impatient. We might think that economies would be instable becuase no one is patient enough to accumulate capital. But here this appears to be offset by the rate of return, which at some point is high enough to overcome even the most impatient person's unwillingness to save.

Another possibility for instability is that individuals are not taking into account their effect on the macroeconomy. No one necessarily is thinking they *have* to save or invest in a manner to ensure stability. But so long as the rate of return that individuals face is informative about the real rate of return on capital - the capital market is not totally screwed up - then their individual decisions will lead to stability. We'll be able to establish that it's possible that this kind of decentralized economy could actually end up finding itself on the same path as one in which an all-knowing planner chose the path. Possible, not inevitable. But the point is that decentralized markets do not *inevitably* lead to instability.

This gives us some clues to what could cause instability, though. One is the capital market. If for some reason that market is not telling individuals the real rate of return on capital, then they'll make decisions that might not support stability or even the existence of a BGP. You can think of this as a an important way in which an economy might stray from the BGP and stay off the BGP for a while, and that might be because the de-centralized market has some friction or information problem.

Other failures of stability would be these poor choices of c_0. One is the over-accumulation problem, which at the risk of being simplistic, might be called the Soviet problem. Accumulating capital for the sake of accumulating capital is setting c_0 too low and ending up with a massive K/Y ratio (which, note, means a very low average product of capital) and very little consumption. Eventually people will realize they are not maximizing consumption and will complain and/or revolt. A different version of this failure mode might be the Great Depression and deflationary pressures, where people consume too little in anticipation of deflation, and the over-accumulation of K/Y occurs via a significant *drop* in Y, as opposed to building too much K.

The second failure state is the Ponzi scheme situation, or under-accumulation. Here the immediate term involves rapid growth in consumption which can only be sustained by using up the real capital stock, and as K/Y declines the rate of return goes up and incents people to raise consumption even further. We have not talked about financial markets in any serious way, but this at least sounds like a financial bubble or something akin to it. Perhaps housing bubbles might be a good model for this kind of failure state?

Why would people choose the "wrong" c_0? Why not is the better question. You can come up with a lot of reasons that people are not forward looking or do not have full information. And that may not be because markets have failed at all to pass along information on rates of return. Nevertheless, because we see most rich economies at least in stable situations it seems as if something is working to get c_0 close to the stable situation most of the time.

This gives us a map for what to think about in terms of the model. Don't think about this model as telling us how the world works. Think of this model as telling us that there are conditions under which a BGP can exist and that is is *possible* for the economy to be stable even with individual decisions and all the possible failure states. Instability isn't inevitable.

7.2 *Consumption side*

The consumption problem concerns a single person, or household, or planner, and we didn't talk specifically about how population growth might influence things. We're going to assume that lifetime utility is continuous, as this makes this particular problem easier to solve.

$$V = \int_0^\infty L_0 e^{g_L t} e^{-\theta t} U(c_t) dt \tag{7.1}$$

but now be clear that c_t is consumption *per capita*. The $L_0 e^{g_L t}$ term is capturing that the number of people enjoying this $U(c_t)$ is growing over time. The value of θ is the discount rate, as usual. We'll assume that $L_0 = 1$ for convenience and then

There is a long philosophical argument to have about whether V should depend directly on population growth or not. We're letting V go up with population size, but it need not, it just changes how you interpret θ.

$$V = \int_0^\infty e^{-(\theta - g_L)t} U(c_t) dt \tag{7.2}$$

Assets *per capita* evolve according to

$$da = w + (r - g_L)a - c, \tag{7.3}$$

and we still need a no Ponzi condition like

$$\lim_{t\to\infty} a_t e^{-\int_0^t r_s ds} \geq 0 \tag{7.4}$$

and these people have some initial value of assets per capita, a_0.

We know the answer to the consumption side of this problem, and we'll assume people have the BGP preferences discussed before.

$$g_c = \frac{1}{\sigma}(r - \theta). \tag{7.5}$$

We're going to characterize this consumer side more formally so that we can refer back to it in the future.

Definition 7.1 (Household Optimization problem) *Taking a series of prices $w_t, r_t{}_{t=0}^{\infty}$ and an initial value of a_0 as given, the household chooses a series $c_t, a_t{}_{t=0}^{\infty}$ that:*

- *Maximizes $V = \int_0^\infty e^{-(\theta - g_L)t} U(c_t) dt$*
- *Subject to the dynamic budget constraint $da = w_t + (r_t - g_L)a_t - c_t$*
- *Subject to the no-Ponzi condition $\lim_{t\to\infty} a_t e^{-\int_0^t r_s ds} \geq 0$*
- *And requiring that $c_t > 0$ for all t*

7.3 Production and asset markets

In continuous time, the work from the Solow model translates without too much trouble. We need to do a little work to establish what the rate of return r is that faces individuals. And for that we need both a better understanding of how firms operate and a capital market that matches up firm's demand for capital with consumer's supply of capital.

For the firms, we presume that, like the Solow model, there are firms that operate with a CRS cost function. That set of firms will operate such that the marginal product of capital, which depends on the elasticity, is equal to the price of renting a unit of capital, R. Note that this R is different than the r rate that individuals will earn on their assets, but we'll see how they are connected in a bit. There will be a similar condition setting the marginal product of labor equal to the wage rate,

$$\epsilon_K \frac{Y}{K} = R \tag{7.6}$$

$$\epsilon_L \frac{Y}{L} = w \tag{7.7}$$

where the ϵ terms are the elasticities of production with respect to the input, and with CRS those elasticities sum to one. We'll tweak this later to see that allowing for economic profits doesn't change the stability result. To be formal,

Definition 7.2 (Firm optimization problem) *Taking a series of prices $w_t, R_t{}_{t=0}^{\infty}$ as given, taking the price of output as the numeraire and equal to one, the firm chooses a series $K_t, L_t{}_{t=0}^{\infty}$ such that*

- *They maximize profits at every time t:* $\pi_t = Y_t - w_t L_t - R_t K_t$
- *Subject to the constraint that that* $Y_t = F(K_t, L_t)$ *where* $F()$ *is constant returns to scale and the elasticities of* Y_t *with respect to inputs are* ϵ_K *and* ϵ_L*, respectively*
- *Subject to the constraint that* $K_t > 0$ *and* $L_t > 0$

How does R relate to r, and how do we move the savings of individuals over to firms to use as capital? We need some structure for a financial market.

Definition 7.3 (Zero profit financial market) *Taking a series of per-capita assets* $\{a_t\}_{t=0}^{\infty}$ *and capital* $\{K_t\}_{t=0}^{\infty}$ *as given, the financial market sets a series* $\{r_t\}_{t=0}^{\infty}$ *such that*

- $r_t a_t L_t = (R_t - \delta) K_t$
- *Total assets equal the total capital stock,* $a_t L_t = K_t$

This helps make clear that the assets of households are used, and can only be used, as capital by the firms. It also establishes how the rate of return paid to households, r_t, is related to the rate of return paid by firms, R_t, which has to be that $r = R - \delta$.

Putting this all together. If assets evolve according to

$$da = w + (r - g_L)a - c, \tag{7.8}$$

then with $a = k$, $r = R - \delta$, and the firm's first order conditions we can write

$$\begin{aligned} dk &= w + (R - \delta - g_L)k - c && (7.9) \\ dk &= \epsilon_L y + (\epsilon_K y/k - \delta - g_L)k - c && (7.10) \\ dk &= (\epsilon_L + \epsilon_K)y - (\delta + g_L)k - c && (7.11) \\ dk &= y - (\delta + g_L)k - c. && (7.12) \end{aligned}$$

The evolution of capital (per capita) in this economy looks a lot like how the Solow model works, only here we have $y - c$ rather than $s_I y$.

7.4 *Equilibrium*

That's kind of it. You now have two differential equations relating consumption and capital to one another, and this additional twist that GDP per capita, y, depends directly on capital per capita, k. So there isn't really a third variable, just another version of k.

$$g_c = \frac{1}{\sigma}(\epsilon_K y/k - \delta - \theta) \tag{7.13}$$

and

$$g_k = y/k - (\delta + g_L) - c/k. \tag{7.14}$$

We can ask if there is a steady state at $g_{k/y} = 0$. If $g_{k/y} = 0$, then $g_y = 0$ and $g_c = 0$, given that from the firm side we have $g_y = \epsilon_K/\epsilon_L g_{k/y}$ and $g_A = 0$.

The Euler equation tells us that

$$\left(\frac{k}{y}\right)^* = \frac{\epsilon_K}{\delta + \theta}. \tag{7.15}$$

This looks a lot like the steady state we established two chapters ago. The difference here is that we set $g_A = 0$ and we have that $s_K = \epsilon_K$ here.

Given that, the steady state condition we get is

$$\left(\frac{c}{y}\right)^* = 1 - \epsilon_K \frac{(\delta + g_L)}{\delta + \theta} \tag{7.16}$$

which means that we have

$$s_I^* = \epsilon_K \frac{(\delta + g_L)}{\delta + \theta}. \tag{7.17}$$

The steady state rate of return in this economy is

$$r^* = \theta. \tag{7.18}$$

Yes, there is a point at which both $g_c = g_k = 0$, as usual this is a BGP because once that is true we know that s_I is unchanged as well. Now the question becomes whether this is stable, and that requires us to understand how this operates away from the BGP.

First, think about starting out with some value a_0 (which is also your k_0). This is given, and our individuals need to choose some initial value c_0. Once they have this choice in place, the mechanics of the Euler equation and the capital accumulation equation take over and we can just iterate this process forward to see what happens. Some of those choices of c_0 will be too big and will end up violating the No Ponzi condition in that they will lead k/y into negative territory as individuals try to borrow, which can't happen here because there is no one else to lend you capital. Some of the choices of c_0 will lead to big values of k/y but c will go towards zero (g_c will become negative).

It can be mathematically shown that there is a single value of c_0 such that starting with this c_0 the dynamics of the model put c and k/y on a path such that they hit the steady state. The set of these c_0 values - one for each value of a_0 - constitutes the "stable arm" of this dynamic system. The stable arm is not just the right starting value of c_0 to hit the steady state, but it is the path along which c moves

The dynamic system is thus "saddle path stable". The proof of this is in the supplemental material.

towards the steady state regardless of which value of c_0 you start with. It's the one path from any initial value of a_0 to the steady state.

So yes, it is possible for this system to be stable. If the consumption choice is always along the stable arm, then no matter what the system heads to steady state. The economic question is whether anything would ensure that individuals always *choose* the value of c_0 that makes sense. The answer is that if we assume people are trying to maximize lifetime utility - which we did assume - then the utility-maximizing answer is to pick the stable arm value of c_0. If you start too high, you violate the Ponzi condition, and you cannot keep borrowing forever, it's infeasible. If you start too low, you accumulate many assets - it's feasible - but you could consume a little more and still have a lot of assets, and consuming more would raise lifetime utility.

One can get hung up on "equilibrium concepts" here, but in practice we mean "solves all the equations I've written down".

Stability in this neo-classical model depends on individuals maximizing their utility in their consumption choices. Let's not get hung up on the idea that economies are stable because individuals are incredibly precise about picking exactly the right value of c_0 at all times. They are not, because this is a model, and because people are generally dumb.

So long as people are *close* to the stable arm the dynamics will push *close* to the steady state, and it's always possible to correct by adjusting c later on. If we figure out we are not exactly on the stable arm, we can "jump" our value of c to get on the stable arm, even if we can't "jump" the value of assets or capital. You can think of the economy as being stable in the same way that a distracted driver is stable in their lane on the highway. Yes, they drift sideways some times, but in general they stay in the lane. Whether they are exactly in the middle of the lane all the time isn't as important as that they don't drive sideways into a barrier at sixty miles per hour.

While not necessarily true, a stable arm is going to involve a negative relationship of s_I and k/y under most realistic parameters, which ensures stability of the k/y ratio.

Robert Solow said it better:

> The problem is not just that perfect foresight into the indefinite future is so implausible away from steady states. The deeper problem is that in practice - if there is any practice - miscalculations about the equilibrium path may not reveal themselves for a long time. The mistaken path gives no signal that it will be ultimately infeasible. It is natural to comfort oneself: whenever the error is perceived there will be a jump to a better approximation to the converging arm. But a large jump may be required. In a decentralized economy it will not be clear who knows what, or where the true converging arm is, or, for that matter, exactly where we are now, given that some agents (speculators) will already have perceived the need for a mid-course correction while others have not. This thought makes it hard even to imagine what a long-run path would look like. It strikes me as more or less devastating for the interpretation of quarterly data as the solution of an indefinite time optimization problem.

7.5 *Add growth rate*

The preliminary model excluded any trend growth g_A from consideration because it makes the accounting a little harder to handle. But we can readily adapt now that we know how things work. Incorporating growth in productivity doesn't change the problem that the individual perceives, because they still just care about maximizing consumption per capita.

As per our Solow model, let productivity growth mean that

$$g_y = \epsilon_K g_k + \epsilon_L g_A \tag{7.19}$$

so

$$g_y = \epsilon_K / \epsilon_L g_{k/y} + g_A. \tag{7.20}$$

Even given that, nothing changes about the actual accumulation equation for capital per capita,

$$g_k = y/k - (\delta + g_L) - c/k. \tag{7.21}$$

The consumption Euler equation also doesn't look any different as people just take the path of wages and returns as given.

$$g_c = \frac{1}{\sigma}(\epsilon_K y/k - \delta - \theta). \tag{7.22}$$

But as we know, the steady state now requires that $g_c = g_A$, not zero. Steady state is where $g_{k/y} = 0$, so that means a steady state is where $g_y = g_A$, and hence that $g_k = g_A$. If $g_k = g_A$, then it has to be that

$$g_A = y/k - (\delta + g_L) - c/k \tag{7.23}$$

or that

$$g_A + \delta + g_L = \frac{y}{k}\left(1 - \frac{c}{y}\right). \tag{7.24}$$

Because a steady state is defined by the point where k/y is constant, this only works if c/y is constant, or there is a constant savings rate. Thus a steady state only holds if $g_c = g_y$, or $g_c = g_A$. Using that in the Euler equation gives us

Conclusion 7.1 *Neo-classical growth model steady state* For the decentralized neo-classical model with BGP preferences the steady state outcomes are:

$$\left(\frac{k}{y}\right)^* = \frac{\epsilon_K}{\theta + \delta + \sigma g_A}. \tag{7.25}$$

and that the steady state rate of return is

$$r^* = \theta + \sigma g_A. \tag{7.26}$$

The savings rate in this steady state is

$$s_I^* = \epsilon_K \frac{\delta + g_L + g_A}{\theta + \delta + \sigma g_A}. \tag{7.27}$$

All this is exactly what we thought it should be in 4.5. The decentralized problem with individuals making consumption decisions and firms taking the rental rate as given leads to a BGP with the same properties we thought it would.

The same idea applies to choosing an initial c_0 along a stable arm. There is a stable arm which leads the economy to the steady state, and that stable arm is the utility-maximizing choice, which is why it seems to make sense that economies are stable.

We can actually say quite a bit about how s_{It} will behave *along* the stable arm. See A.20. The general point is that so long as ϵ_K is fixed, it will rise or fall monotonically as we approach steady state. It never jumps around.

After we looked at 4.5 we talked about why s_I^* could or would be less than s_K. In this case, the question is whether s_I^* is less than ϵ_K. But the same idea holds, which is that the savings rate will be lower than capital's share or elasticity if $g_A + g_L < \theta + \sigma g_A$ or $\theta - g_L > (1-\sigma)g_A$.

Why does this condition hold? Look at the PDV of *total* consumption along a BGP (assuming that we're already on the BGP). That's

$$\int_0^\infty c_0 e^{g_c t} L_0 e^{g_L t} e^{-rt} dt. \tag{7.28}$$

On the BGP we know that $r = \theta + \sigma g_A$ and $g_c = g_A$ so this whole thing is

$$\int_0^\infty c_0 L_0 e^{(g_A + g_L - \theta - \sigma g_A)t} dt. \tag{7.29}$$

This PDV is finite only if the exponent is negative, so only if $g_A + g_L - \theta - \sigma g_A < 0$ or $\theta - g_L > (1-\sigma)g_A$.

The point here is that the only *feasible* answer for an economy is to have $s_I^* < \epsilon_K$, or for the economy to "under-save" relative to the importance of capital in production. That happens because people are impatient, but that's what ensures this whole thing holds together.

Conclusion 7.2 (Below the Golden Rule) *In a Ramsey/neo-classical economy with impatient individuals, $\theta - g_L > (1-\sigma)g_A$ is required for the equilibrium to be feasible. Therefore the steady state savings rate is below the Golden Rule savings rate, $s_I^* < s_I^{GR} = \epsilon_K$.*

If the value of θ is just right the optimal savings rate will be constant, as in the Solow model. See A.19. Thus the constant rate of s_I in the data doesn't tell us by itself that we're at steady state or on a BGP.

This isn't a failure of some kind. The Golden Rule from 3.5 is an abstract benchmark that ignores time preferences of all kinds. Under the Golden Rule the PDV of consumption goes to infinity, which is fine in the Golden Rule because we're not letting anyone make any choices. But if you have people making choices, and tell them that they have an infinite PDV of consumption to play with, they will very much not save at the Golden Rule rate. They will not save at all, because why bother. An equilibrium with positive savings requires people to be impatient to some degree.

7.6 Planner and market outcomes

The equilibrium (with growth of g_A or not) depends on individuals making utility-maximizing decisions for *themselves*. But how does that equilibrium compare to one where someone with full knowledge of how everything work would pick? Does the decentralization of decision-making here - individuals and firms - lead to the economy missing some obvious gains? I've already alluded to the answer, which is no. The reason is that the capital market is sending both firms and individuals accurate information about the rate of return on capital, and hence they choices consistent with having full information about .... the rate of return on capital.

But if you were an all-knowing planner, what would the problem look like? You'd again want to maximize lifetime utility

There is another point of argument about why a planner would have the same per-capita preferences as individuals. They also might have a different discount rate, etc.

$$V = \int_0^\infty e^{-(\theta - g_L)t} U(c_t) dt. \tag{7.30}$$

Because they are all-knowing, they don't think of themselves as having assets that get translated through a capital market to firms, or anything like that. They know that capital accumulates according to

$$dk = y - (\delta + g_L)k - c \tag{7.31}$$

and to avoid tedious math we're going to set $g_A = 0$ again. They also know that y is a function of k.

Their Hamiltonian looks like this

$$H(t) = e^{-(\theta - g_L)t} u(c) + \lambda(y - (\delta + g_L)k - c). \tag{7.32}$$

The planner is picking c with k as the state variable. Using the standard Hamiltonian conditions we have

$$\frac{\partial H}{\partial c} = e^{-(\theta - g_L)t} u'(c) dc - \lambda = 0 \tag{7.33}$$

$$d\lambda = -\frac{\partial H}{\partial k} = -\lambda(\epsilon_K y/k - \delta - g_L) \tag{7.34}$$

$$dk = \frac{\partial H}{\partial \lambda} = y - (\delta + g_L)k - c. \tag{7.35}$$

Solve these together again like with the consumption problem

$$-\theta e^{-(\theta - g_L)t} u'(c) + e^{-(\theta - g_L)t} u''(c) dc - d\lambda = 0 \tag{7.36}$$

and use with the second condition to get

$$-\theta e^{-(\theta - g_L)t} u'(c) + e^{-(\theta - g_L)t} u''(c) dc = -\lambda(\epsilon_K y/k - \delta - g_L) \tag{7.37}$$

and using the first condition again we get

$$-\theta e^{-(\theta - g_L)t} u'(c) + e^{-(\theta - g_L)t} u''(c) dc = -e^{-(\theta - g_L)t} u'(c)(\epsilon_K y/k - \delta - g_L) \tag{7.38}$$

which solves down to

$$u''(c)dc = \theta u'(c) - u'(c)(\epsilon_K y/k - \delta - g_L) \tag{7.39}$$

and to

$$\frac{dc}{c} = g_c = (\epsilon_K y/k - \delta - g_L - \theta)\left[-\frac{u'(c)}{u''(c)c}\right]. \tag{7.40}$$

This is the same as the consumption Euler equation from the individual problem once we accounted for how r was determined by the market. The planner gets the same answer.

The other condition in the decentralized model was that

$$g_k = y/k - (\delta + g_L) - c/k \tag{7.41}$$

which is just the exact constraint the planner uses. Everything is identical, so the solutions are identical.

Don't over-interpret this result. This is not a statement that decentralized markets invariably or always or inevitably reach the same results as centralized planning. It is not an argument that centralized planning can achieve the same outcomes as markets. It is an observation that in a situation where markets pass on undistorted information *on literally everything that matters* to decentralized decision-makers they will not leave any gains on the table. It works here because this economy is so simple there really is only one piece of information to pass on, the rate of return. The only thing the capital market is doing in the decentralized model is telling indviduals what is happening to R on the prodution side. There's only 1 asset and no other choices that could be made, so there's only one piece of information to pass on.

The other thing we learn from this it is not *necessary* to have a planner to achieve stability. If the market passes along sufficient information the decentralized economy can have stability too. That's interesting to us because a large part of the argument around the origin of macro was whether markets were inherently unstable and that you had to have a central planner to ensure stability.

7.7 *Distribution and stability*

It turns out that you don't even need perfect information from the markets to get stability. What I mean is that even if the signals are distorted to some extent you can still get stability in a decentralized market. You won't reach the same level of utility - gains will be left on the table - but it won't necessarily mean instability.

What we're going to change is that we'll give the capital market some wedge or friction they impose on returns. In our GDP accounting terms we said that

$$R = s_K \frac{Y}{K} \text{ and } r = s_K \frac{Y}{K} - \delta \tag{7.42}$$

where s_K was capital's share of income. That R and r are the returns that the capital market pays to the asset-holders. But they might *charge* a rate of $R' > R$ to firms.

Firms will still spend a fraction ϵ_K of their output on capital services, so it will be that

$$\epsilon_K Y = R'K. \tag{7.43}$$

The capital markets accounts now look like this

$$rK + \Pi = R'K - \delta K \tag{7.44}$$

where π is a profit that the capital market collects and we could see this as

$$s_K Y - \delta K + \pi = \epsilon_K Y - \delta K \tag{7.45}$$

so that

$$\Pi = (\epsilon_K - s_K)Y \tag{7.46}$$

and the profits of the capital market are positive so long as they pay out s_K at a lower rate than ϵ_K. How can they do this? Pick your favorite I/O story about how banks or financial markets are organized that would allow them to collude or compete in some kind of Bertrand/Cournot manner to provide assets to firms. It's not relevant for our problem how they get the profits, just that they do.

What happens to those profits? Well, the world has to add up, so the individual dynamic budget constraint now looks like this

$$da = \pi + w + (r - g_L)a - c \tag{7.47}$$

where $\pi = \Pi/L$ are per-capita profits. From the perspective of the individual these profits just appear (think of them owning stock in the capital market companies). Assume that they don't *know* how this all works; the checks just show up. They don't know that they are getting π profits at the expense of their own r return on capital. That's the point at which information is not flowing to them.

Because π is taken to just appear, like the wage, it changes nothing about the Euler equation, which is still

$$g_c = \frac{1}{\sigma}(r - \theta) \tag{7.48}$$

but that now resolves to

$$g_c = \frac{1}{\sigma}(s_K y/k - \delta - \theta), \tag{7.49}$$

and the rate of return here is *lower* because s_K is lower. All else equal, they'd choose to have lower consumption growth, meaning they save less, which makes sense because the rate of return is depressed. We know the steady state capital/output ratio (assuming productivity growth of g_A) is

$$\left(\frac{k}{y}\right)^* = \frac{s_K}{\theta + \delta + \sigma g_A}, \tag{7.50}$$

and is lower than when the capital market had no profits.

The growth rate of capital is still

$$g_k = y/k - (\delta + g_L) - c/k \tag{7.51}$$

and with the steady state finding this implies that

$$s_I^* = s_K \frac{g_A + \delta + g_L}{\theta + \delta + \sigma g_A} \tag{7.52}$$

which is exactly what we found in 4.5, which didn't rely on any assumptions about efficiency of the capital markets.

What changes with this inefficient capital market? This economy is still stable. Nothing about the dynamics changed in any material way. There is still a stable arm and all that, and in equilibrium individual still choose to be on that stable arm, as that still maximizes utility. It's not that different than just saying from their perspective that the elasticity ϵ_K changed.

What does change is that the planner's solution is now *different*. Either the planner doesn't use a capital market (why bother when you know all) or the planner knows that the capital market makes profits and they know exactly how those profits affect r. The planner has full information. So the planner would still choose

$$s_I^{Planner} = \epsilon_K \frac{g_A + \delta + g_L}{\theta + \delta + \sigma g_A} > s_I^*. \tag{7.53}$$

In this case the planner can sustain a higher k/y ratio and higher consumption in steady state.

Thus frictions in the economy, at least in the simplistic way we've looked at, are not *necessarily* destabilizing. They are inefficient in that they lower lifetime utility. But even with these frictions the capital market still tells people that when k/y goes down r goes up, even if r is always "too low" compared to an efficient market.

7.8 Equilibrium concepts

Very often you will see research which makes a clear statement of what an "equilibrium" consists of. One of the margin notes above made an offhand comment about this just being "what I'm solving for", and that's not wrong as a general point. But it is often useful, in particular as models get more complicated, to have a clear statement of the equilibrium that one is looking at or for. One thing that the equilibrium statement makes clear is how you, as the modeler, view the decision process of the agents (people, firms, etc.) in your model. What optimization problem are they in fact solving, what kind of prices are they aware of, and what are the adding-up conditions or market clearing conditions that hold?

In the case of our de-centralized neo-classical model we might say something like this

Definition 7.4 (De-centralized Equilibrium) *A de-centralized equilibrium consists of a series of prices $\{w_t, r_t\}_0^\infty$, allocations $\{c_t, a_{t+1}\}_0^\infty$ for the household, and allocations $\{K_t, L_t\}_0^\infty$ for the firm such that*

1. *Given prices the allocations solve the household maximization problem in (7.1)*
2. *Given prices the allocations solve the firm maximization problem in (7.2)*
3. *The financial market clears according to (7.3)*
4. *The labor market clears $L_t = L_0 e^{g_L t}$*
5. *The aggregate economy is closed $Y_t = c_t L_t + dK_t + \delta K_t$*

The exact terminology is going to vary depending on who wrote down the model and their choice about how to define things. The important aspects of writing down an equilibrium statement like this are that it contains the information about what is being chosen, what constraints have to bind, and what each agent knows. Here this helps make it obvious that in our neo-classical model the choice is the entire series of consumption and assets from time zero to infinty, made all at once. If you like, the equilibrium statement collects all the moments I might say "we also know XXXX" when working through problems more loosely in class. An equilibrium statement is answering :what are all the things we know about the economy?

This equilibrium is also an Arrow-Debreu Equilibrium, in the sense that we're presuming everything can be decided at period zero. When we move to a value-function approach we will explicitly switch to a sequential equilibrium concept.

Naming these is ambiguous. Some people will refer to this as the "competitive equilibrium", which I avoid because the word competitive denotes a very particular market structure for firms that is not necessary for this de-centralized equilibrium to hold. What this equilibrium is about is the fact that the agents (households and firms) each take the series of prices as given when they solve their maximization problems.

8
Dynamic programming methods

This chapter is technical. It is about a method for solving a dynamic optimization problem. This method, sometimes called dynamic programming, uses something called a Bellman equation to break down the dynamic problem into what amounts to a two-period problem. You can think of this method as flipping how we look at dynamic problems.

Lagrangians and Hamiltonians are asking you (or the optimizer) to decide on the sequence of choices you'll make - c_t for example - over the entire time frame of the problem. In the case where we know with certainty what will happen, that works. And that's true even though the choice of c_t influences the future. There is no uncertainty about *how* c_t affects the future. So conceptually it is plausible to pick the entire sequence of c_t from the initial period.

That's not a very intuitive way of thinking about how people make decisions, but that isn't a problem with Lagrangians and Hamiltonians, it's because we know the world has uncertainty and there is no possibility of picking a full path from today's perspective because things will change in unknown ways.

The Bellman equation approach flips, like I said, the perspective on the problem. Rather than picking a sequence of c_t all at once, it asks the optimizer to set up a decision rule (or algorithm or menu) of what c_t they should pick at time t given the information they have on the state of the world. It's asking you for an answer on what to do whatever the situation. Now, in the case of full certainty, this decision rule will replicate the sequence you'd get with the Lagrangian/Hamiltonian. But one value of the Bellman equation approach is that it will more naturally lead us to a way of thinking about how to optimize in the face of uncertainty, when the *best* thing we can come up with is a decision rule.

The other thing the Bellman equation does is give us a means of numerically solving for the decision rule. You can solve numerically for the Lagrangian/Hamiltonian answer, too, but the Bellman equation is recursive (it refers to itself) and recursion is something that computers do very well. Others can and will teach you all about optimizing the numerical techniques associated with solving these problems. I will not. I'll show you very slow, very inefficient scripts that, to me, explain what is going on without being fast. But the structure of how to solve problems using Bellman equations is powerful and can be used in far more situations.

There isn't anything here that changes how we think about stability or growth. It's purely a way of approaching a problem.

8.1 Setting up the Bellman equation

Nothing here will look weird. We're going to solve the planner's Ramsey problem just as before. We're going to do it in discrete time, however, because this will help make sense of the Bellman equation and because this will be more natural to use in settings with uncertainty. Utility is

$$V = \sum_{t=0}^{\infty} \beta^t U(c_t). \tag{8.1}$$

Depending on your opinions you can imagine that β incorporates an adjustment for population or not. Doesn't matter. For the dynamic budget constraints we have that

$$k_{t+1} = y_t + (1 - \delta - g_L)k_t - c_t. \tag{8.2}$$

The planner knows how capital accumulates, and we know that this doesn't necessarily change the problem's solution unless there is a capital market friction, and even then we could work this out. They know y_t is a function of k_t (and possibly some growth rate g_A). There is also a No Ponzi condition to consider. The other important thing is that the planner has some k_0 initial asset stock.

Let's assume you solved this problem. Then there would be indirect utility function of something like this

$$v(k_0) = \max_{k_{t+1}{}_{t=0}^{\infty}} \left[\sum_{t=0}^{\infty} \beta^t u(y_t + (1 - \delta - g_L)k_t - k_{t+1}) \right] \tag{8.3}$$

which says that the indirect utility, conditional on k_0, is the choice of the sequence of k_{t+1} terms that maximizes lifetime utility subject to the dynamic constraint that relates c_t to k_{t+1}. Note that we still have the notion of a sequence of answers here. $v(k_0)$ is the utility value of being on the stable arm, so to speak.

Now just break apart that summation term this way

$$v(k_0) = \max_{k_{t+1}{}_{t=0}^{\infty}} \left[u(y_0 + (1 - \delta - g_L)k_0 - k_1) + \sum_{t=1}^{\infty} \beta^t u(y_t + (1 - \delta - g_L)k_t - k_{t+1}) \right]. \tag{8.4}$$

This just says the indirect utility depends on the utility in period zero plus a discounted sum of utility from period 1 forward. But look at that thing inside the summation. It's looks like an indirect utility of an optimization problem as well, but a problem that starts at time 1 rather than at time zero. We can write this as

$$v(k_0) = \max_{k_1} \left[u(y_0 + (1 - \delta - g_L)k_0 - k_1) + \beta \max_{k_{t+2}{}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(y_{t+1} + (1 - \delta - g_L)k_{t+1} - k_{t+2}) \right]. \tag{8.5}$$

I did a few notational things. I reset the summation to period zero, which is just a choice on how to number things. The choices still start with k_2 and go forward (see the maximand). Because I reset the time period back to zero, I need to discount it by β, as otherwise when $t = 0$ we'd lose the proper discounting.

And we are doing two maximizations. This seems worse, but it's better. The first optimization is easy, because now we are just picking k_1, one thing. We're trading off utility today - $U(c_0)$ - against the utility we could get by sending k_1 into the future, where that *second* maximiation says "do the best thing possible with k_1". All we've done is push the rest of the optimization into the future.

But notice the form of this future indirect utility looks exactly like indirect utility today in form. The future indirect depends on an initial value k_1 which comes from our first choice, but that doesn't matter. We can write

$$v(k_0) = \max_{k_1} \left[u(y_0 + (1 - \delta - g_L)k_0 - k_1) + \beta v(k_1) \right]. \tag{8.6}$$

This $v(k_0)$ is our Bellman equation. It is recursive in that the function $v(k_0)$ is defined implicitly as it depends on itself in terms of $v(k_1)$. This has reduced our big sequential choice down to choose k_1. In that sense we've lost the time dimension of the problem, taking the future $v(k_1)$ as given. What we've replaced it with is a choice of this function. The Bellman equation is a functional equation, in the sense that it is a function that depends on another function. We can make this generic, as it doesn't depend on just periods 0 and 1,

This isn't that weird. $\ln(xy) = \ln(x) + \ln(y)$ is a functional equation because the function of xy on the left depends on functions of x and y on the right.

$$v(k_t) = \max_{k_{t+1}} \left[u(y_t + (1 - \delta - g_L)k_t - k_{t+1}) + \beta v(k_{t+1}) \right] \tag{8.7}$$

and this holds for any periods t and $t + 1$.

The Bellman equation is a functional equation, but don't think of $v(k)$ as a function, per se. It's better to think of it as a list. It's a list that tells you the value of having a specific amount of k. In that it's a little like the Lagrangian multiplier, but it's not the marginal value of wealth, it's the total value. If you have $k = 1$, then the utility value of that is 3.26 (for example). If you have $k = 2$, the utility value of that is 5.93 (again, example). Thinking of $v(k)$ as a list will help in understanding how we solve this numerically.

Thinking of $v(k)$ as a list is also why Bellman equations are often written with slightly different notation that doesn't use the t subscripts. It's because the Bellman setup is not really about time, per se, but about lists. So normally what's written is this

$$v(k) = \max_{k'} \left[u(y + (1 - \delta - g_L)k - k') + \beta v(k') \right] \tag{8.8}$$

where k is "what I have" and k' is "what I'm choosing".

8.2 *Value function iteration*

$v(k)$ is cool, but it isn't really the "answer" to the optimization problem. The answer to the problem is a *control or policy function* which is the choice you make in response to the argument of the Bellman equation. Given k and given a known list $v(k')$ of the value of having capital in the future, you can do the simple maximization and pick a value of k'. That value of k' implies a value of c because they are linked by the budget constraint $c = y + (1 - \delta - g_L)k - k'$.

The Bellman approach kind of works backwards. Let's assume that we know what the list $v(k)$ looks like and that we therefore know some control function

$$k' = g(k), \tag{8.9}$$

where again I'd encourage you to think of $g(k)$ not as a function, per se, but as a list of what to do if you have k capital to start with. It's the answer to the maximization problem inside the Bellman equation.

This means we should have

$$g(k) = \arg\max_{k'} \left[u(y + (1 - \delta - g_L)k - k') + \beta v(k') \right] \tag{8.10}$$

and that

$$v(k) = u(y + (1 - \delta - g_L)k - g(k)) + \beta v(g(k)). \tag{8.11}$$

Great, we've assumed we have an answer. But we don't know what that answer is. Here's the power of the Bellman equation approach, value function iteration. The value function $v(k)$ is what is know as a *contraction mapping*, and because it is this thing, it means that we can use an iterative process to close in on the actual function $v(k)$.

The mathematical proof that Bellman equations have a solution given certain conditions on u and $\beta < 1$ and the allowable set of choices of k is not something we are going to go through. For our purposes, it's enough to know it can be proven. The intuition is that u has to have diminishing marginal utility, meaning it is bounded, and that there is at least one value of k that you can pick.

Definition 8.1 (Value function iteration) *Iterate the value function as follows:*

1. *Start with any function/list $v(k)_0$, where the subscript 0 refers to the step in the iteration.*

2. *Each subsequent step in the iteration, n, produces a new function/list* $v(k)_{n+1} = \max_{k'}\left[u(y + (1-\delta-g_L)k - k') + \beta v(k')_n\right]$

Then the following is true

1. *The sequence of $v(k)_n$ for $n = 0, ..., \infty$ converges to $v(k)$, the function that solves the Bellman equation*

2. *The distance of $v(k)_n$ to the solution $v(k)$ gets smaller as n increases.*

The "distance" in the definition is between two lists. In practice take the element-wise absolute value of the difference between the two lists, and the distance is the largest of those differences. This is a "sup-norm".

You could apply this iteration by hand, but that would be dumb. This is what computers are for. Give yourself a starting list $v(k')_0$. Do the maximization problem, choosing k', given that starting list, *for every possible value of k*. That will give you an answer for what the value of each k is. That's your new iteration of the value function, and you use that to start the next iteration, $v(k')_1$.

To do this you have to make a few assumptions because computers are great but they cannot actually deal with infinite values of k or run the iteration to infinity:

Assumption 8.1 (Implementing value function iteration) *To apply the iteration method:*

1. *"Discretize the state space": Pick a discrete number of values of k that are allowable in the calculations. This can be a very large number, but needs to be discrete.*

2. *Assign a tolerance: Pick a value of the distance between $v(k)_{n+1} - v(k)_n$ such that once the distance is below this tolerance you can stop iterating.*

You are thus going to get a numerical approximation to the true function $v(k)$. In theory the $v(k)$ list is infinitely long, but you will have 1,000 or 10,000 or 1,000,000 values on it, which will be sufficient to see what is happening.

By the way, once you have iterated and found your solution $v(k)$, you then also have your policy function $g(k)$, because in your last iteration $g(k)$ is the list of choices you made for k' for each value of k that you started with.

This all makes far more sense when you see it on a computer, which is what we'll do in class.

8.3 Properties of the Bellman equation

The Bellman equation gives you a way to numerically solve for $v(k)$ and the policy function $k' = g(k)$, which remember also implies a policy for consumption, c. Remember that's what we actually want to optimize over? There is nothing about the numerical answer, however, that tells us anything necessary about what the solution looks like. It would be weird - and wrong - if the Bellman iteration gave us an answer different than what we got just doing the regular old maximization and finding a stable arm. These better give us the *same* answer.

If the Bellman approach and the Lagrangian/Hamiltonian give us the same answer, why bother? Because the Bellman approach is more flexible for solving problems with uncertainty, and can be applied in other settings where doing the Lagrangian/Hamiltonian isn't plausible. It's a more generic method.

Let's go back to our generic form of the Bellman

$$v(k) = \max_{k'} \left[u(y + (1 - \delta - g_L)k - k') + \beta v(k') \right] \tag{8.12}$$

What kind of properties must the solution, in the form of $v(k)$ and $g(k)$, have?

Start with the maximization problem itself. This says to optimize over k', and the first order condition is

$$-u'(c) + \beta v'(k') = 0 \tag{8.13}$$

or that the marginal utility of consumption today (which depends on our choice of k') should equal the discounted marginal value of k' in the future. That should make some intuitive sense. We should be trading off on the margin between consuming today and consuming in the future, and $v(k')$ is telling us how valuable consumption "in the future" will be. Nothing too strange here.

But this FOC does depend on $v'(k')$, or the marginal value of k', which we don't know. But we do know how $v(k)$ works. It's kind of weird to think about taking the derivative of $v(k)$ on the right-hand side because this is a max problem. But we know that there is this notional policy function $g(k)$ out there that let's us write

$$v(k) = u(y + (1 - \delta - g_L)k - g(k)) + \beta v(g(k)). \tag{8.14}$$

Now, what's the derivative of $v(k)$ with respect to k, where $g(k)$ is accounting for all the impacts of k on the problem?

$$v'(k) = u'(c) \left[\partial y / \partial k + (1 - \delta - g_L) - g'(k) \right] + \beta v'(g(k)) g'(k). \tag{8.15}$$

The change in value when k changes depends on how that impacts the immediate budget through y and the amount of capital available, which is valued based on marginal utility. But it also depends on how k' is affected by having a bigger budget, and that influences today's utility as well as the value in the future.

This, however, simplifies. Re-arrange this to be

$$v'(k) = u'(c) \left[\partial y / \partial k + (1 - \delta - g_L) \right] + \left[\beta v'(g(k)) - u'(c) \right] g'(k). \tag{8.16}$$

The term in the second set of brackets is, given that $k' = g(k)$, just $\beta v'(k') - u'(c)$, which we know from the first-order condition must equal zero. That means that the derivative of the value function with respect to k is just

$$v'(k) = u'(c) \left[\partial y/\partial k + (1 - \delta - g_L)\right]. \tag{8.17}$$

The value function changes with k *only* because that expands the current budget (and that gets valued at marginal utility).

This is an example of the Envelope Theorem. Because $v(k')$ is presumed to be maximized already, adjusting the constraint on the margin cannot change the value. If it could, then $v(k')$ must not have been maximized to begin with.

The last thing we cna do is use this condition for $v'(k)$ and note that it holds for any given value of k. So we can write

$$v'(k') = u'(c') \left[\partial y'/\partial k' + (1 - \delta - g_L)\right]. \tag{8.18}$$

Put that back into the first-order condition from the maximization

$$-u'(c) + \beta v'(k') = 0 \tag{8.19}$$

and we have

$$u'(c) = \beta u'(c') \left[\partial y'/\partial k' + (1 - \delta - g_L)\right] \tag{8.20}$$

which we could write as

$$\frac{u'(c)}{u'(c')} = \beta \left[\partial y'/\partial k' + (1 - \delta - g_L)\right]. \tag{8.21}$$

That should look familiar, as it is just an Euler equation. Presuming that production works in our normal way and that c' means "next period" we have

$$\frac{u'(c_t)}{u'(c_{t+1})} = \beta \left[\epsilon_k y_{t+1}/k_{t+1} + (1 - \delta - g_L)\right]. \tag{8.22}$$

Value function iteration using the Bellman equation provides an answer for what $v(k)$ and $g(k)$ are. Those answers conform to the Euler equation we had from our typical Lagrangian/Hamiltonian problem. Nothing about the Bellman value function iteration gives us a different answer. It is just a tool for finding the answer. But it's a tool that works in lots of settings.

Equilibrium

Like we did with our earlier Ramsey model, we can write down a more formal equilibrium statement. In this case it will be a *recursive equilibrium*, which just refers to the fact that we'll be solving for a value function and policy function as opposed to a set of initial choices (AD). A recursive equilibrium is different than a sequential equilibrium, although as you can see they feel related to one another because they are both about having a "rule" to use in picking what happens next.

Remember that we've done this entirely in terms of a planner so far, or if you like a single person existing on an island who understands all the implications of their actions.

Definition 8.2 (Recursive Planner Equilibrium) *A recursive planner equilibrium consists a value function $v(k)$ and policy function $g(k)$ such that*

1. *$v(k)$ solves the planners Bellman equation $v(k) = \max_{k'} [u(c) + \beta v(k')]$ and $g(k)$ is the associated policy function*
2. *c evolves according to $c = y + (1 - \delta - g_L)k - k'$*

We could consider a decentralized recursive equilibrium where a household that doens't appreciate the implication of their actions is making decisions, and then we need some additional aggregate constraints that govern how the prices that individuals see (wages, returns) are determined by the state of the aggregate economy. Where recursive equilibria statements look different here is that we want to be explicit about which aggregate state variable (capital, labor) is dictating the price, as opposed to just talking about w or r. Why? Because these settings are about solving for the evolution of that state variable, from which everything else follows. That's not some new idea, we did it in the Ramsey model, but in recursive settings we're more explicit about this in part because that is how things get solved numerically and how the math associated with Bellman equations is worked out.

These get more complicated because if we have individual(s) making decisions without knowing the overall effect of their actions, we need additional constraints or terms that ensure that all those decisions are consistent with the aggregate state. To put that in terms that would be familiar from last chapter, we could write down a decentralized equilibrium as

Definition 8.3 (Recursive Decentralized Equilibrium) *A decentralized recursive equilibrium consists a value function $v(a, k)$ and policy functions $a' = g(a, k)$ and $c' = h(a, k)$, and price functions $r(k)$ and $w(k)$,*

1. *$v(a, k)$ solves the planners Bellman equation $v(a, k) = \max_{a'} [u(c) + \beta v(a')]$ and $g(a, k)$ is the associated policy function,*
2. *c evolves according to $a' = w(k) + (1 + r(k) - g_L)a - c$,*
3. *$r(k)$ and $w(k)$ satisfy the firm maximization problem in (7.2),*
4. *The financial market clears according to (7.3),*
5. *Consistenty, $g(k, k) = y + (1 - \delta - g_L) - h(k, k)$*

9
Shocks and fluctuations

The Ramsey model we have is a reasonable platform to describe long-run trends in most economies (developed ones, at least) and why they appear to have a substantial amount of stability across many dimensions. Those economies, though, never appear to be perfectly stable, however, in the sense that their growth rates and allocations like consumption shares are not always at a fixed number. In particular, these economies appear to have notable fluctuations around those trends. These fluctuations are often referred to as "business cycles", although I dislike the term "cycles" because it implies some kind of necessary symmetry (equal time "above" trend and "below" trend) that I don't think is accurate.

Leaving aside my concern we can adapt our Ramsey model to allow for fluctuations. We'll do this in pieces as there are substantial technical issues we need to deal with to get there. The path looks something like this. We have the Solow model, which is a *dynamic* model of capital accumulation. To this we can add *stochastic* shocks to productivity which will create fluctuations around the trend, but which will illustrate for us that because of the dynamic nature of the economy (e.g. capital tomorrow depends on capital today) these shocks will have lingering effects on the economy even if they only exist for one period of time. The stability of the Solow will keep pushing the economy back towards the trend. We'll need a set of tools for describing and handling stochastic processes.

Once we have all that we'll go back to our Ramsey model, which we referred to as a *general equilibrium* model as it allowed the key decision point of the Solow model - the investment rate - to be determined endogenously. We'll put those stochastic shocks into this model as well. That, though, requires two new sets of technical material. First we'll have to understand more about those optimizing consumers and how they deal with uncertainty about what might happen in the future, and second we'll introduce a new way of formulating the Ramsey problem that allows us to solve it (numerically) with the stochastic shocks involved. That beast will be a *dynamic stochastic general equilibrium* (DSGE) model of the economy. It forms the baseline from which almost any modern understanding of fluctuations is built.

What needs to be very clear is that DSGE models are not models of what *causes* business cycles or fluctuations. They are models of how economies *respond* to some exogenous source of fluctuations. The best way to think about DSGE models is like forecasting hurricane paths or major storms. Meteorologists cannot with great reliability tell you precisely when a hurricane will form, or how many might form in the next few months (they try, but they are invariably wrong, often spectacularly). They *can* tell you with a shocking amount of accuracy where a hurricane - once formed - will go. We're engaged in the same process here, building a model of how economies react to unexplained shocks, and I'll remind you that one of the main things we're interested in here is what is necessary in these models to keep delivering the stable return to trend that we see in the data.

We have to keep in mind how important it is to understand why modern developed economies are stable even if we don't have a great handle on what drives fluctuations. The stability is remarkable in and of itself, and I'll remind you that it is not obvious. The origin of macroeconomics was in the Great Depression when a valid question was whether this signalled the collapse and end of the industrial market economy that had developed over the course of 1800 to 1930. While Britain might have claimed around 100 years of experience with such an economy by 1930 (and even that is debatable) the rest of the industrial world had, generously, only fifty years of exposure. Within the lifetime of many people in 1930 were economies that were poor, agrarian, and stagnant. Hence the question of whether modern industrial market economies can persist was, and remains, a live question.

It will be easy to lose track of that doing the technical work behind the DSGE model. The assumptions will seem pointless and unrealistic. The math will seem overly complex. The utter absence of any consideration of political or social issues will be mystifying. Throughout this process I want you to keep in mind that you are in this for the long haul, and this represents only the initial step on the way to making better assumptions or having more realistic discussions. For better or worse the DSGE is the vocabulary used to talk about macroeconomics for the most part. So learn that vocabulary even if your intention is to ultimately invent a better one.

9.1 Fluctuations

Data on fluctuations with respect to A. Time series plot and growth rate to growth rate.

9.2 The stochastic Solow model

Start with the Solow setting so we aren't worried about how individuals adjust their investment behavior yet, and let's focus solely on how shocks create dynamics in the economy via the capital stock. To create a stochastic economy with fluctuations we're going to allow productivity, A_t, to be hit by random shocks each period, and each shock will only last one period. This will make the productivity shocks different than in Figure 3.4, where there was a *permanent* shock to $A(0)$ the pushed the economy to a higher BGP. Here, we'll assume that the productivity shocks have mean zero so that the economy is always fluctuating around the same old BGP.

Assuming that the fluctuations come from stochastic shocks to productivity is why I said that this isn't a model or description of *why* business cycles happen. This is a purely exogenous unexplained set of shocks that we model as if productivity is fluctuating. They are a stand-in for any real economic shock you want to think about, be it financial, monetary, trade-related, policy-related, or even, as it happens, to a major natural disaster like a hurricane or a global pandemic. More advanced models or theories you might look into if you're interested in this will be much more specific about the nature of these shocks, but in many ways they always boil down to "the economy gest more or less productive given the set of capital and labor it has". One way to think about these productivity shocks is that they are capturing capacity use in firms and establishments that might fluctuate for various unmodelled reasons.

In the regular Solow productivity growth was a deterministic process, meaning we knew exactly what would happen at every point in time,

$$\ln A_t = \ln A_0 + g_A t. \tag{9.1}$$

Now, what we want to do is to introduce some stochastic element to this so that we don't know what A_t is going to be at A_0 (or at A_{t-1}). A simple way to do this is to add a stochastic shock to the above process, as in

$$\ln A_t = \ln A_0 + g_A t + \varepsilon_t \tag{9.2}$$

where ε_t is distributed as a Normal with mean zero and a variance of σ_ε^2. Of note, this expression says that productivity is always "around trend" in that A_t always starts out at what we'd expect it to be in the deterministic case, and *then* there is some shock. But that shock does not linger at all into the next period. Just to make this clear, note that

$$\ln A_{t+1} = \ln A_0 + g_A(t+1) + \varepsilon_{t+1} \tag{9.3}$$

does not include any information about A_t or ε_t in it. That will make this the easiest stochastic process to work with, and we'll add things to get more complex later.

If we plug this process back into the expression for GDP per capita we have

$$\ln y_t = \frac{\epsilon_K}{\epsilon_L} \ln K_t/Y_t + \ln A_0 + g_A t + \varepsilon_t, \tag{9.4}$$

which looks a lot like our regular Solow model, except now there is this additional shock to GDP per capita that happens. However, we have to be careful here because through the capital/output ratio the process for GDP per capita does "remember" shocks that happen in the past. For example, if ε_t is a big positive shocks, the GDP per capita will be above trend and K/Y_{t+1} will be a little higher than normal because the economy was a little richer than normal. But then in period $t+1$ the economy doesn't "start" on trend (like productivity does) but rather above trend. The ε_t shock still matters in period $t+1$.

The key lesson here is that even if we understand the stochastic properties of the productivity process well, that doesn't mean we necessarily know the stochastic properties of the GDP per capita process well, given that the dynamics of the capital/output ratio necessarily create a link over time. There are, in rough terms, three ways we can deal with this:

1. Theory. It's possible to work out more carefully the stochastic properties of $\ln y_t$ for this (or other) stochastic properties of A_t. That involves linearizing the Solow around a steady state. This is possible, but goes into details that are not of first-order importance for this course.

2. Simulate and estimate. We could simply plug the stochastic process for A_t and the equations for $\ln y_t$ and the accumulation for K/Y into a computer, and run simulations over and over again where we let a random number generator fill in values for ε_t over many periods. We could then just "look" at the stochastic properties of $\ln y_t$ that emerge from this simulation. That requires us to estimate those stochastic properties given data on $\ln y_t$, which is possible, but is the kind of thing that you will learn in econometrics and, in particular, in macro-econometrics. We will play with the simulations, but leave the estimation to those courses.

3. Assume away the problem. The last option is to assert that this complication from K/Y is a small enough issue that we can ignore it. In practice that means assuming that the stochastic process for A_t produces "small" shocks, and that the economy is always so close to steady state that $K/Y \approx K/Y^*$, and is just a constant. Then $\ln y_t$ inherits the precise stochastic properties of A_t.

For the purposes of this chapter we will take option three, and assume away the problem. That's because for the most part it turns out to be about right, and because the point here is to learn about stochastic processes, not learn about how to estimate time series processes econometrically. The stochastic processes and properties we'll learn about are useful in a variety of contexts, so trying to only solve the Solow (or in the future the full DSGE) model is too limiting.

With that in mind let's assume that the economy is always so close to steady state that $K/Y_t \approx K/Y^*$, meaning the dynamics of the capital/output ratio are insignificant to us. Then our process for GDP per capita is

$$\ln y_t = \kappa + \ln A_t \tag{9.5}$$

where $\kappa = \epsilon_K/\epsilon_L \ln K/Y^*$. This means that

$$\ln y_t = \kappa + \ln A_0 + g_A t + \varepsilon_t = \ln y_0 + g_A t + \varepsilon_t. \tag{9.6}$$

Thus in this case our process for productivity can just be translated directly to the process for GDP per capita.

Now, given this, we can establish some stochastic properties. But before we get into the math of it, just consider the intuition. GDP per capita here follows a trend, and again the value of y_t in any given period is based on that trend, and then is shocked away from that trend. But there is no memory or history to this process. All it does is create noise around the trend. In that sense this noisy process is stable, in the sense that it always is centered on the trend line. The process does not "blow up" or experience fluctuations that accelerate into complete collapse or infinite growth, which is consistent with what we think is true about the world. So the question for our stochastic processes in general is if they "behave" and keep the economy stable - even if there is noise - or if they imply collapse or explosive growth.

In the original Solow model good behavior and stability was ensured by the self-correcting nature of the capital process, even if it wasn't stochastic. If the economy moved away from the steady state for any reason, it naturally moved back towards the steady state. We'll want some similar thing to be true about our stochastic processes. If they experience a big shock (negative or positive) is there something about the process that ensures it reverts back to "normal". In this initial case that's due to the memory-less nature of the shock; it doesn't matter what happened yesterday at all. For other more complicated processes we'll have to evaluate what creates the same condition.

9.3 Stability and Stationarity

Let's establish a better definition of what we mean by stability of a stochastic process. To do that let's first define a few properties of a stochastic process:

Definition 9.1 (Properties of stochastic processes) *Define the following properties of a stochastic process x_t*

1. *Expected value at time t:* $E[x_t] = \mu_t$
2. *Variance at time t:* $V[x_t] = E[(x_t - \mu_t)^2]$
3. *Autocovariance of length k:* $E[(x_t - \mu_t)((x_{t+k} - \mu_{t+k})]$

Our notion of stability is that this stochastic process should "behave" around the trend, which informally means we think the expected value of that process should not change (or not change a lot) and, perhaps most relevant, that the variance of that process does not get bigger and bigger and t goes up. We want the noise in the system to stay contained. If the variance explodes as time goes on then the process isn't stable in the sense we want.

Statistically we can define the following property:

This definition is for weak stationarity. *Strong* stationarity is about the joint distribution of values of x_t realizations doesn't depend on t itself.

Definition 9.2 (Stationarity) *A stochastic process for x_t is (weakly) stationary if it satisfies the following properties*

1. $E[x_t] = \mu$*, where μ is a constant and independent of t*
2. $V[x_t] = E[(x_t - c)^2] = \gamma(0)$ *where $\gamma(0)$ is finite and independent of t*
3. $E[(x_t - c)((x_{t+k} - c)] = \gamma(k)$ *where $\gamma(k)$ is finite and independent of t*

In short, a stationary process is a time series of random variables (x_t) that maintain the same stochastic properties over time. The $\gamma(k)$ things are purely notational laziness. They are just a standard shorthand way of writing variances and covariances, rather than writing out things like $E[(x_t - \mu)((x_{t+k} - \mu)$ all the time.

We can also define an autocorrelation as $\gamma(k)/\gamma(0)$, just a scaled autocovariance.

A (weakly) stationary process is one where the expected value, variance, and auto-covariances are all the same regardless of the time period t. Even for those higher-level auto-covariances, the point is that their size is determined only by the value of k (how far apart in time they are) and not by the specific t (the actual time).

Conclusion 9.1 (Stability and stationarity) *Define stability of a process x_t as the variance of that process being finite and independent of t, then*

1. *A process that is stationary is also stable*
2. *A process that is not stable is not stationary*
3. *A process that is stable can be either stationary or non-stationary*

What this implies is that if we study stationary processes, we will also be studying stable ones, and invariably that will be what we study. Note that this is technically more restrictive that we want, but we'll see below that we can study a stable time series that isn't technically non-stationary, and it won't really change our conclusions about anything.

No, we're not proving this conclusion. It can be proved by a more competent econometrician.

Let's look at a specific kind of process,

Definition 9.3 (White noise processes) *A stochastic process x_t is called "white noise" if it has the following properties*

1. $E[x_t] = 0$
2. $E[x_t^2] = \gamma(0)$ *is finite*
3. $E[x_t x_{t+k}] = \gamma(k) = 0$ *for all* $k \neq 0$.

which means that a white noise process is stationary (and therefore stable).

This is the simplest kind of stochastic process, and in some sense lots of other processes are built off of manipulations of these white noise processes. In the example before, the noise ε_t is a white noise process, and hence is stationary. The variance of ε_t is always σ_ε^2 at any given point in time and all the auto-covariances are zero, meaning that ε_t doesn't have any statistical relationship with any past value or future value of the shock - it's "memoryless".

What about our actual process for GDP per capita, though? The expected value of that in time t is

$$E[\ln y_t] = \ln y_0 + g_A t \tag{9.7}$$

which very much *does* depend on t, so we already know it isn't stationary. Moving on, the variance of this process is

$$V[\ln y_t] = \sigma_\varepsilon^2 \tag{9.8}$$

which is constant, so that part is okay. If we did the upper level auto-covariances we'd see they are all still zero. So this isn't stationary, but it's close.

The problem is obviously this whole trend thing with $g_A t$. But that is simple to deal with by looking at how the series for GDP per capita behaves *relative to the trend*. If we de-trend this process,

$$\ln y_t - \ln y_0 - g_A t = \varepsilon_t \tag{9.9}$$

then the whole left-hand side - the de-trended GDP per capita - is just a white noise process and thus is stationary.

De-trending here is obvious because we know the process at work. De-trending a data series where we don't know the process is subject to lots of assumptions and arguments about those assumptions.

Definition 9.4 (Trend stationarity) *A stochastic process for x_t is trend stationary if it can be written as*

$$x_t = f(t) + z_t \tag{9.10}$$

where $f(t)$ is a deterministic function of t and z_t is some stationary process.

Our simple model is trend stationary. It's essentially just noise around our typical BGP. This perhaps makes it obvious that we are not explaining anything about what creates these fluctuations, we're just assuming they exist so that our model looks a little more like the data.

9.4 *Moving average processes*

To take a step closer to something meaningful, we might want to consider shocks that linger in the economy. We already have something that does this - the capital/output ratio - but here we're going to build in that the actual shocks themselves retain some effect period by period.

Consider the alternative process

$$\ln y_t = \ln y_0 + g_A t + \varepsilon_t + b_1 \varepsilon_{t-1}. \tag{9.11}$$

This allows for the prior period shock, ε_{t-1}, to still have an effect in time t. The coefficient b_1 tells us how powerful this effect is. Does this change our evaluation of the stationarity? First, let's simplify by de-trending and calling that detrended variable x_t,

$$x_t = \ln y_t - \ln y_0 - g_A t = \varepsilon_t + b_1 \varepsilon_{t-1}. \tag{9.12}$$

The expected value of x_t is

$$E[x_t] = 0 \tag{9.13}$$

and the variance is

$$V[x_t] = (1 + b_1^2)\sigma_\varepsilon^2 \tag{9.14}$$

which is finite and not dependent on t itself. We can also evaluate the auto-covariance between two observations adjacent to one another in time

$$E[x_t x_{t-1}] = E[\varepsilon_t \varepsilon_{t-1}] + b_1 E[\varepsilon_t \varepsilon_{t-2}] + b_1 E[\varepsilon_{t-1}^2] + b_1 E[\varepsilon_{t-1}\varepsilon_{t-2}] = b_1 \sigma_\varepsilon^2. \quad (9.15)$$

This is finite and doesn't depend on t.

The process is no longer white noise, but it is still stationary. The addition of the lingering shock creates some auto-correlation between x_t and x_{t-1}, which seems kind of obvious, but this doesn't make things get out of control. These shocks aren't feeding on one another, so to speak. We can generalize,

Definition 9.5 (Moving average process) *A moving average process of order q, MA(q), is defined as*

$$x_t = \mu + \epsilon_t + b_1\epsilon_{t-1} + b_2\epsilon_{t-2} + ... + b_q\epsilon_{t-q} \quad (9.16)$$

and is stationary for any values of b_j.

We can thus model our shocks to productivity and/or GDP per capita with any kind of arbitrary set of shocks that continue to have effects in later periods and still have a stationary series. This is stable because of the finite number of lags involved. Eventually any shock that has occurred disappears completely from influencing x_t, even if the values of b_q are huge. That's a big hint. Stability and stationarity depend on the shocks "disappearing" from influence eventually.

Let's think about an infinite number of backward-looking shocks,

Definition 9.6 (Infinite moving average process) *An infinite moving average process, MA(∞), is defined as*

$$x_t = \mu + \sum_{j=0}^{\infty} b_j \epsilon_{t-j} \quad (9.17)$$

and is stationary so long as $\sum_{j=0}^{\infty} |b_j| < \infty$.

Note that the condition here is just that the combined effect of the shocks has to stay finite, which in an MA(q) is true by default. The lesson we have is that our process can remain stable - stationary - even if shocks have continued effects period after period. These kind of moving-average shocks don't create any kind of feedback loop that could generate unstable behavior. Even if one shock is big, the next one is no more likely to be big. The fact that the b_j terms can't get too big just means that shocks can't have a big effect forever.

9.5 *Auto-regressive process*

A different kind of stochastic process does allow for the possibility of feedback, and thus does open up the possibility of instability. This kind of auto-regressive process is very commonly used in economics, so we'll need to see what kind of properties it needs to stay stable.

Let's keep working with this idea of a de-trended process for GDP per capita, $x_t = \ln A_t - \ln A_0 - g_A t$. But now we're going to let that de-trended GDP per capita depend on itself,

$$\ln x_t = \rho \ln x_{t-1} + \varepsilon_t. \tag{9.18}$$

The value of ρ is the auto-regressive coefficient as it determines how much of an impact the lagged value has on the present. There is also the truly stochastic element ε_t that is creating noise in the process that pushes the values away from trend. This process is called an AR(1), as it depends only on 1 lag of the x_t variable.

The deterministic Solow model is already auto-regressive in the sense that the $t-1$ value of GDP per capita influences the t value via the K/Y ratio. It's stable because of the dynamics of K/Y.

The question here is what kind of conditions on ρ make this stochastic process act in a stable manner, meaning it acts like our regular Solow model with an exogneous shock. Let's assess this new process for whether it is stationary. To do this we're going to iterate this process out so that we can look at x_t in terms of x_0 and the set of shocks between 0 and t, which it should be intuitive will all matter. What this shows us is that this x_t is just a summation of all the various shocks back to time ... negative infinity.

$$x_t = \varepsilon_t + \rho\varepsilon_{t-1} + \rho^2\varepsilon_{t-2} + ... + \rho^\infty\varepsilon_{-\infty} = \sum_{j=0}^{\infty} \rho^j \varepsilon_{t-j}. \tag{9.19}$$

Does it make sense to think about a process running back to negative infinity? Not really, but in principle we can think of things working this way.

There are a few ways of assessing this auto-regressive process. First, note that as written above this is now just a moving average, and we know how to assess the stationarity of a moving average. This process will be stationary if

$$\sum_{j=0}^{\infty} |\rho^j| < \infty. \tag{9.20}$$

Note that this is just the summation of the values of ρ, ignoring the ε terms. If these terms stay finite then we know the series doesn't get out of control. Now, in this case the coefficients of the moving average have a clear structure, as opposed to just being a series of b_j terms.

This is an infinite sum, and we can manipulate the absolute values, so we get

$$\sum_{j=0}^{\infty} |\rho^j| = \sum_{j=0}^{\infty} |\rho|^j \tag{9.21}$$

and if $|\rho| < 1$ then we have

$$\sum_{j=0}^{\infty} |\rho|^j = \frac{1}{1 - |\rho|} < \infty. \tag{9.22}$$

This process is stationary if the values of ρ are not too big. It's non-stationary if $|\rho| \geq 1$ because eventually the process "blows up" to infinity.

In the end, the auto-regressive process is stationary if the auto-regressive parameter ρ is below one in absolute value, meaning the shocks are dampening over time.

More general AR processes that depend on an arbitrary number of lags are discussed in A.18.

Definition 9.7 (Stationarity of AR(1) processes) *An AR(1) process for* $x_t = \rho x_{t-1} + \varepsilon_t$ *is stationary if* $|\rho| < 1$ *and* ε_t *is white noise.*

We can work with an AR(1) process that has $\rho = 1$. In this case the process for GDP per capita would look like this

$$\ln x_t = \ln x_{t-1} + \varepsilon_t. \tag{9.23}$$

The value of GDP per capita (relative to trend) is just last period plus some shock, and there is nothing pushing this back towards that presumed trend. This process has a "memory" in that all the past shocks are always relevant (as opposed to the dwindling impact of past shocks in a stationary AR(1) process). This type of process has a name:

Definition 9.8 (Random walk) *A random walk is a stochastic process of the form*

$$x_t = x_{t-1} + \varepsilon_t \tag{9.24}$$

where ε_t *is white noise. This process is not stationary.*

You can work out that this is not stationary by looking at the variance, which is $V[x_t] = \sigma_\varepsilon^2 t$, and obviously expands continuously with time.

However, if one takes the first difference of a random walk, we have

$$x_t - x_{t-1} = \varepsilon_t \tag{9.25}$$

which means the *change* in x_t is stationary.

We often refer to random walk process as having a "unit root", which refers to how we evaluate AR processes for stability, but in practice you can think of a unit root as meaning $\rho = 1$.

Definition 9.9 (Difference stationarity) *A stochastic process for* x_t *is difference stationary if it can be written as*

$$x_t - x_{t-1} = z_t \tag{9.26}$$

where z_t *is some stationary process.*

For the case of GDP per capita, if we thought that the level of GDP per capita was related as in the random walk, and thus was not stationary, we might still think that the growth rate of GDP per capita - the difference in log GDP per capita - was stationary and stable. This would mean that there is no tendency of the growth rate to explode, even though we can't necessarily call the path of GDP per capita stable.

Note that it isn't immediately obvious from the data on the level of GDP per capita that we should use either an AR(1) with $\rho < 1$ or a random walk as the right model. We presume that there is a trend in GDP per capita - the BGP - but it could just be a realization of a random walk process. Econometrically, one should be able to test for which model is correct, but in practice this turns out to be difficult because a process with $\rho \approx 0.9$ is stationary but looks a lot like one with $\rho = 1$ even over a long time frame.

9.6 *Impulse response functions*

A standard object of interest is something called an "impulse response function" for things like GDP per capita, it's growth rate, the capital/output ratio, and other items of interest in the economy. You've already created impulse response functions, as they are the graphs of what happens in response to a single shock to the economy in "period zero" in a typical homework problem.

Even though all our stochastic processes (AR, MA, or ARMA) get shocked every period, the general way to "see" how shocks propogate is to do this kind of homework problem for the model. That is, assume there is a single shock at time one, ε_1, and then all future $\varepsilon_t = 0$.

Definition 9.10 (Impluse response function for AR(1)) *For a stationary AR(1) process $x_t = \mu + \rho x_{t-1} + \varepsilon_t$ the impulse response function is the series of derivatives of x_{t+j} for $j \geq t$ with respect to the shock in period t, ε_t.*

$$\frac{\partial x_{t+j}}{\partial \varepsilon_t} = \rho^j. \tag{9.27}$$

Perhaps not a surprise, but the effect of a shock to ε_t on x_{t+j} is just the autocorrelation, ρ^j, between those periods.

10
Uncertainty

Adding shocks or fluctuations to the standard Solow model via productivity leads to a model that can at least replicate something that looks like the data, but as before there isn't any allowance for the idea that people make decisions about s_I and/or g_c. The idea that people would not respond or react to these fluctuations seems wrong. Knowing that there is a chance of a bad outcome or good outcome could, at least in principle, influence what people choose to do. Here we're going to introduce that uncertainty into the individual decision process and see that there are reasons to believe that this will affect their behavior in predictable ways.

There are two principles we can think about in this context. The first is that people will choose g_c based on their expectations of what will happen, and that as shocks occur they'll make adjustments. There will be a plan, and then there will be reality. In terms of the kind of solution we find for individual behavior, this means there isn't a full lifetime path. We can't solve for the whole stable arm, so to speak, because we don't know what will happen at any given moment. The nature of the consumption choice people make with uncertainty is that it takes the form of a rule or algorithm. If good thing happens, do this. If bad thing happens, do that. That rule will help maximize expected utility given the expectation of shocks, but there will be continuous adjustment as shocks are revealed, and consumption and utility will not end up equal to expectations.

The second principle is that the presence of uncertainty will induce additional savings relative to a certain future. It's not just that individuals will choose $E[g_c]$ equal to the growth rate of consumption with certainty, it's that they'll raise $E[g_c]$ - have lower present consumption - because they will want to insure themselves against future bad shocks. Marginal utility falls as c goes up, so a bad shock hurts more than a good shock feels good. Because people don't like the bad shocks, they'll consume less today (when they have certainty) to give themselves more consumption in the future (when they face uncertainty). This is called precautionary savings and it means that the answer to the full forward-looking value function Ramsey problem is not just the solution to the regular Ramsey problem with some stochastic shocks attached (like it is with the Solow).

10.1 Intuition

Go back to the situation where we just think about picking c_1 and c_2, as in Section 4.3. If people made a tiny change to consumption today, dc_1, then they'd lose $U'(c_1)dc_1$ in utility. They'd gain what? $-dc_1(1+r)$ in consumption tomorrow, and what's the utility value of that? Well, now with uncertainty it depends on what consumption in period 2 looks like. Instead of knowing that we have $U'(c_2)$, the best we can say is that marginal utility is $E[U'(c_2)]$ because maybe there was a bad shock (and marginal utility is big) and maybe there was a good shock (and marginal utility is low).

If we knew the expected value of marginal utility we'd know that in equilibrium you'd choose

$$U'(c_1)dc_1 - \frac{E[U'(c_2)]}{1+\theta}dc_1(1+r) = 0.$$

That's fine and that gives us an Euler equation like before of

$$\frac{E[U'(c_2)]}{U'(c_1)} = \frac{1+\theta}{1+r}. \tag{10.1}$$

Like before, we can approximate $U'(c_2)$, but now we're going to do a second-order approximation. We could have done this before, but the second order terms all drop out because there isn't any uncertainty. Here, they'll still matter.

$$U'\left((1+g_c)c_1\right) \approx U'(c_1) + U''(c_1)dc_1 + \frac{1}{2}U'''(c_1)(dc_1)^2.$$

In expectation that is

$$E[U'(c_2)] \approx U'(c_1) + U''(c_1)c_1E[g_c] + \frac{1}{2}U'''(c_1)c_1^2E[g_c^2]. \tag{10.2}$$

and in the Euler equation that gives us

$$1 + \frac{U''(c_1)c_1}{U'(c_1)}E[g_c] + \frac{1}{2}\frac{U'''(c_1)c_1^2}{U'(c_1)}E[g_c^2] \approx 1+\theta-r \tag{10.3}$$

There's some tedious algebra but we can re-write this as

$$E[g_c] \approx \frac{-U'(c_1)}{U''(c_1)c_1}(r-\theta) + \frac{1}{2}\frac{-U'''(c_1)c_1}{U''(c_1)}E[g_c^2]. \tag{10.4}$$

Feels like we have not accomplished much. But we've got the intuition we need. On the left is now just the expected value of consumption growth. That's an expected value because we don't know what's going to happen in period 2.

On the right, we've got a few familiar things. The first term is the typical growth rate of consumption, under certainty, which depends on the IES and the gap between r and θ. Nothing new here. But we're *adding* a second term to this. That term involves two things. The first is the $E[g_c^2]$, which is kind of like the variance of consumption growth (but not the actual variance). This is measuring the possible "noise" in consumption growth that occurs because of uncertainty. If the possible shocks to c_2 are big, then $E[g_c^2]$ will be big as well.

The second terms is this fraction which deals with the nature of the third(!?) derivative of the utility function.

Definition 10.1 (Coefficient of relative prudence) *The coefficient of relative prudence (CRP) is defined as*

$$CRP \equiv \frac{-U'''(c_1)c_1}{U''(c_1)}$$

and measures how sensitive marginal utility is to changes in consumption. The CRP measures whether $U'(c)$ is convex (CRP is positive) or concave (CRP is negative) with respect to c.

Our typical CRRA preferences with $U(c) = c^{1-\sigma}/(1-\sigma)$ have a CRP of $1+\sigma$. That's positive, so people with CRRA preferences act prudently and save extra when uncertainty is present.

We can make this statement about how people consume based on the CRP.

Conclusion 10.1 (Prudence and precautionary savings) *If individuals have a $CRP > 0$ then their expected consumption growth $E[g_c]$ with uncertainty is greater than g_c with certainty. That implies that individuals save more when uncertainty is present.*

A useful way of seeing what is going on here is by plotting out the marginal utility curve, as in Figure 10.1. The marginal utility curve, $U'(c)$, slopes down, which is just the assumption of diminishing marginal utility. Mathematically that negative slope occurs because $U''(c) < 0$. In the Figure I've drawn the marginal utility cuve as convex to the origin, which is due to the fact that $U'''(c) > 0$, or the second derivative *of marginal utility* is positive. A positive CRP is saying that this is what the marginal utility curve looks like.

In practice, that means marginal utility falls as c goes up, but that the decline slows down as c gets higher, creating this convex shape. Once consumption is high enough marginal utility falls, but not by a lot. That's going to create the incentive to save more. To what's going on, consider the uncertainty plotted in the figure. The person

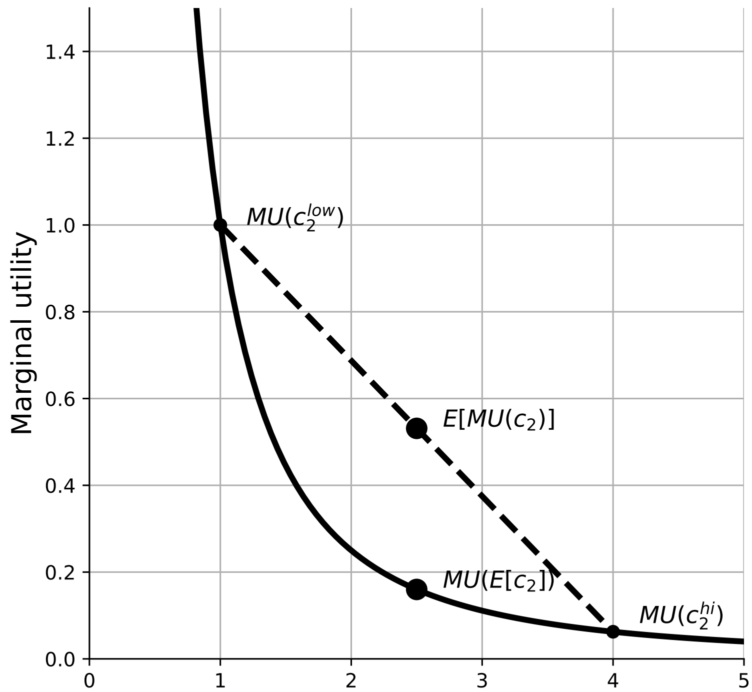

Figure 10.1: Marginal utility and precautionary savings

faces either c_2^{low} or c_2^{hi} in the second period of life (say with 50/50 probability). The expected value of their consumption is around 2.5 here. The *marginal utility of expected consumption*, $U'(E[c_2])$, is on the marginal utility curve.

But they'll never actually consume $E[c_2]$. They'll get either the high or low outcome. The marginal utility of the low outcome is very high because of the positive third derivative, while the marginal utility of the high outcome is low, but not that much lower than the certainty case. The *expected value of marginal utility*, $E[U'(c_2)]$, is the probability weighted average of the two marginal utilities, and it is distinctly *higher* than $U'(E[c_2])$. Because of the shape of the marginal utility curve this person expects their marginal utility to be very high.

In the Euler equation, a very high marginal utility in the second period says to move more consumption towards the second period. In other words, this person needs to lower c_1 and save more in order to raise the expected value of consumption in the second period to offset the fact that their marginal utility in period 2 is quite high because of the uncertainty. You can think of what this person is doing is sliding the consumption here to the right so that the distinction in marginal utilities is not quite as big. Precautionary savings is essentially self-insuring against future risk.

10.2 *Expected savings*

Let the uncertainty hitting this economy be such that log productivity is what is fluctuating around a trend, so that we have a trend stationary process for $\ln A_t$. In other words, we're going to let

$$\ln A_t = \ln A_0 + g_A t + x_t \tag{10.5}$$

where $x_t = \rho x_{t-1} + \varepsilon_t$ with $\rho < 1$ and ε_t be white noise, so that x_t is a stationary process. That means that log GDP per capita is

$$\ln y_t = \epsilon_K/\epsilon_L \ln K_t/Y_t + \ln A_0 + g_A t + x_t \tag{10.6}$$

is also a trend stationary process.

Consumption in this economy would be, at any given point in time, $\ln c_t = \ln(1 - s_{It}) + \ln y_t$, so consumption is also going to be trend stationary. What we're after here is some sense of what that implies for the level of savings and the implied expected value of the rate of return, using what we know about consumption from above. That Euler expression depends on $E[g_c]$ and $V[g_c]$, so let's evaluate those.

First, note that $g_c \approx \Delta \ln c_t$, where I've used the Δ to be clear we're thinking of a discrete time change, and not a derivative. So we're after $E[\Delta \ln c_t]$ and $V[\Delta \ln c_t]$. We're only going to evaluate these "close" to the steady state or around a balanced growth path. That is, what are the long-run effects of this uncertainty on the rate of return and savings? For short-run effects we'd have to be more careful with things.

Around a BGP s_I will be stable, so $E[\Delta \ln c_t] \approx E[\Delta \ln y_t]$ and similar for the variance. What's the difference in log GDP per capita?

$$\Delta \ln y_t \approx g_A + x_{t+1} - x_t \tag{10.7}$$

because we are close to the BGP and so the capital/output ratio isn't changing a lot. The expected value of this is

$$E[\Delta \ln y_t] \approx g_A + E[x_{t+1}] - E[x_t]. \tag{10.8}$$

What's the expected value of the AR(1) process we're working with? We know from prior work that the AR(1) is just a moving average over infinite shocks in the past

$$x_t = \sum_{j=0}^{\infty} \rho^j \varepsilon_{t-j} \tag{10.9}$$

so

$$E[x_t] = 0 \tag{10.10}$$

for any period t (including $t+1$) and

$$V[x_t] = \sigma_\varepsilon^2 \sum_{j=0}^{\infty} \rho^2 j = \frac{\sigma_\varepsilon^2}{1-\rho^2} \tag{10.11}$$

while we are at it.

Therefore,

$$E[g_c] \approx E[\Delta \ln c_t] \approx E[\Delta \ln y_t] \approx g_A \tag{10.12}$$

or the expected growth rate of consumption is just g_A, which should make some sense. We're just bouncing around a trend line. At the same time, the variance of this thing is not so obvious. we need to do a little work here because what we want is

$$V[\Delta \ln y_t] \approx V[x_{t+1}] + V[x_t] - 2Cov(x_{t+1}, x_t) \tag{10.13}$$

where the minus sign comes from the fact that x_t is subtracted. Filling in some things we know

$$V[\Delta \ln y_t] \approx 2\frac{\sigma_\varepsilon^2}{1-\rho^2} - 2Cov(\rho x_t + \varepsilon_{t+1}, x_t) \tag{10.14}$$

and applying rules for covariances out we get

$$V[\Delta \ln y_t] \approx 2\frac{\sigma_\varepsilon^2}{1-\rho^2} - 2\rho\frac{\sigma_\varepsilon^2}{1-\rho^2} = \frac{2\sigma_\varepsilon^2}{1+\rho}. \tag{10.15}$$

Which means in the end we have that

$$V[g_c] \approx V[\Delta \ln c_t] \approx V[\Delta \ln y_t] \approx \frac{2\sigma_\varepsilon^2}{1+\rho}. \tag{10.16}$$

Now, go back to the approximation of the Euler equation and let's talk about what this implies will be the rate of return in expectation along the balanced growth path.

$$E[g_c] \approx \frac{1}{\sigma}(r-\theta) + \frac{1}{2}(1+\sigma)E[g_c^2]. \tag{10.17}$$

One thing we have to do is deal with this $E[g_c^2]$ thing, because that isn't quite a variance. But the definition of a variance is

$$V[g_c] = E[g_c^2] - E[g_c]^2 \tag{10.18}$$

so we can write

$$E[g_c] \approx \frac{1}{\sigma}(r-\theta) + \frac{1}{2}(1+\sigma)\left(V[g_c] + E[g_c]^2\right). \tag{10.19}$$

Plugging in what we know

$$g_A \approx \frac{1}{\sigma}(r-\theta) + \frac{1}{2}(1+\sigma)\frac{2\sigma_\varepsilon^2}{1+\rho} + \frac{1}{2}(1+\sigma)g_A^2. \tag{10.20}$$

Solve this for r and you get

$$r \approx \sigma g_A + \theta - (1+\sigma)\frac{\sigma_\varepsilon^2}{1+\rho} - \frac{1}{2}(1+\sigma)g_A^2. \tag{10.21}$$

That's a lot. But note that it contains all the intuition from the non-stochastic case. If σ_ε is zero and there are no fluctuations, then this retains the idea that $r \approx \theta + \sigma g_A$, except that this extra little term is also now hiding in there with respect to g_A^2. That arises because we did a second-order approximation to the Euler equation, and even if there is no noise, we've added something here to the mix. To a second order, people care about the curvature of their utility function, and if g_A is big, then this means the growth of consumption will be very large in steady state. Large growth makes people with big σ values unhappy to some extent, so they save a lot (have low r) so that their consumption will be very *big* in steady state and then the utility differences that g_c creates across periods are not as bothersome. It's the same logic as uncertainty, just applied to rapid growth. In this sense you can justify or think about why savings rates would be high even in places with rapid growth (where otherwise you might imagine they would not bother).

But onto the uncertainty part. Here, the bigger is the variance of the shocks, σ_ε^2, the less people like it, and the value of σ here (I know) dictates how much they hate it. The uncertainty bothers them because they'd prefer to have smooth consumption. So what they do to respond is save more, meaning k/y is higher, meaning r will be lower along the BGP. If the shocks are big, then saving means higher average consumption levels (remember, we're below the Golden Rule) and therefore the disruption of the shocks is not as bad because utility has "flattened out". You save to make yourself feel rich to self-insure against the shocks.

Note that if ρ goes up then the implied size of this effect is *lower*. That is, if the shocks are persistent, then there is a smaller effect on r. Why? Because high ρ means smooth consumption. If the shocks aren't big but are persistent, then you get what you want, consumption is similar period by period because each shock lasts and continues to push up (or down) your consumption for many periods. All σ is doing is saying you want smooth consumption, it doesn't care about the level.

Ultimately, the point is that what dictates the BGP level of r, and hence all the BGP values of k/y and s_I, which recall are just

$$\left(\frac{k}{y}\right)^* = \frac{\epsilon_K}{r^* + \delta} \tag{10.22}$$

which comes from the definition of the rate of return. Once we've pinned down r^* we know what the level of the k/y ratio is that supports this. You could plug in to see that

$$\left(\frac{k}{y}\right)^* = \frac{\epsilon_K}{\sigma g_A + \theta - (1+\sigma)\frac{\sigma_\epsilon^2}{1+\rho} - \frac{1}{2}(1+\sigma)g_A^2 + \delta} \tag{10.23}$$

if you don't like yourself. All the adjustments to r for uncertainty mean k/y responds in the long run as people try to smooth consumption. That also shows up in the savings rate, which has to be the right rate to hit that k/y ratio. Remember, it has to be that in the end

$$\left(\frac{k}{y}\right)^* = \frac{s_I^*}{\delta + g_L + g_A} \tag{10.24}$$

just due to the Solow mechanics of the capital stock. This only ever holds in steady state (even in the Solow model), and it is what relates the savings rate to the production side restrictions on k/y including population growth. But since we know what k/y ratio people want from their need to smooth, we know what s_I they have to choose to support that, or

$$s_I^* = (\delta + g_L + g_A)\frac{\epsilon_K}{\sigma g_A + \theta - (1+\sigma)\frac{\sigma_\epsilon^2}{1+\rho} - \frac{1}{2}(1+\sigma)g_A^2 + \delta}. \tag{10.25}$$

If we can pin down the rate of return using the Euler equation, we can pin down the steady state of everything else. Note how r^* moves inversely to both k/y and s_I, and that was true with or without uncertainty. All uncertainty is doing is making this something like $E[r^*]$ along the BGP because we'll always be bouncing around it.

Here's where you can see that higher variance in the shocks raises the s_I along the BGP. In addition, we've learned that because of second-order effects, if g_A is very big that will also raise the s_I along the BGP.

10.3 Value function

We want to incorporate this kind of intuition into our larger dynamic model, where individuals make decisions about consumption over time, subject to a series of shocks. We'll have some work to do in the next section to describe those shocks in a better way, but for the moment let's just think about there being two possible shocks to productivity, good and bad, so A_t^{High} or A_t^{Low}. The probability of each is p and $1-p$, and these shocks are i.i.d. in that there is no persistence in them and the past values don't matter. This is a crude "white noise" kind of thing.

This is where the Bellman equation approach becomes valuable. Let's write down *two* value functions, one for each state of productivity at time t.

$$\begin{aligned} v(k, Low) &= \max_{k'} \left[u(y^{Low} + (1-\delta-g_L)k - k') + \beta E[v(k')] \right] \\ v(k, High) &= \max_{k'} \left[u(y^{High} + (1-\delta-g_L)k - k') + \beta E[v(k')] \right]. \end{aligned}$$

Note how the state shows up. If we're in the Low state, that is affected y via productivity, so it's y^{Low}, and similar for the good state. The maximization problem is different depnending on whether you have a good or bad shock, and it's different because it changes your set of resources you can work with. But note that we know the state we're in here.

The other part of this is that the value function depends on the expected value of k' in the future, $E[v(k')]$. That's like our intuitive approach above where we cared about $E[u'(c_2)]$, and like above it's what creates some issues in working with this. We can characterize that expectation, though

$$E[v(k')] = pv(k', High) + (1-p)v(k', Low) \tag{10.26}$$

which means that both of our value functions $v(k, Low)$ and $v(k, High)$ depend on *both* possible future value functions with k'. That also should make sense. When I make a decision today about k', I have to take into account the possibility that tomorrow I'll wake up and either be in the good or bad state. The logic from before says that relative to a situation with certainty we'll increase k' because we want to self-insure against the bad state.

We could also put this in terms of a single "state" variable S, which is either A^{High} or A^{Low}, and then we could write a single value function of

$$v(k, S) = \max_{k'} \left[u(y(S) + (1-\delta-g_L)k - k') + \beta E[v(k', S)] \right]. \tag{10.27}$$

Either way, what we're doing is saying that the value of having k today *conditional on being in state S* is equal to our normal maximization problem, which takes S as given, and trades this off against the *expected* value of k' in the future.

Before we operationalize this, we'll put some more formality on how these shocks work.

10.4 Stochastic shocks

To extend this concept out over an indefinite period of time, all the way up to infinite periods, we have to do a little work in defining how the stochastic shocks (to productivity, but could be other things) work. We already know about stochastic processes like AR and

MA processes. What we need to do here is introduce a means of talking about these in terms of a discrete set of shocks as opposed to a continuous one. In the stochastic material before, ϵ_t was a shock that was distributed normally (or log normal) meaning it could take on any value between minus infinity and infinty. It's not plausible to create an infinite number of value functions for all possible outcomes. Instead we're going to discretize the uncertainty, and in the simplest form to two states (e.g. good or bad).

We'll start with some additional definitions.

Definition 10.2 (Markov property) *A stochastic process x_t satisfies the "Markov Property", or is called a "Markov Process", if for all t*

$$P(x_{t+1}|x_t, x_{t-1}, ..., x_0) = P(x_{t+1}|x_t). \tag{10.28}$$

The Markov property says that a random variable next period only depends on todays realization of the random variable, and no earlier realizations. If you think of shocks in the future (x_{t+1}), present (x_t), and past (x_{t-1} and earlier) then what we're saying is that the future shock depends only on the present, and not on the past.

An AR(1) process like $x_t = \rho x_{t-1} + \varepsilon_t$ satisfies the Markov property. The only realization that matters is from $t-1$. In contrast, an AR(2) process does not satisfy the Markov property.

The Markov property is quite general, and we'll narrow this down for our purposes.

Definition 10.3 (Markov chain) *A stochastic process x_t is a "Markov chain" if*

- *It satisfies the Markov property*
- *$Pr(x_{t+1})|x_t)$ does not depend on t (e.g. it is stationary)*
- *x_t is drawn from a finite set of values S*

Basically, a Markov chain is a discrete state space stationary stochastic process that satisfies the Markov property. Because a Markov chain has a finite number of values it can take, we can define a transition matrix that is dimension $S \times S$ which tells us what the probabilty of moving from state 1 to state 2 (or 3 or 4) is, and the probability of moving from state 4 to state 1 (or 2 or 3 or whatever). To be clear, S is the set of values that x_t can take, and the transition matrix is the probability of any given state happening, given the information on the current state.

Definition 10.4 (Transition matrix) *The transition matrix Q of a Markov chain is an $S \times S$ matrix with the following properties.*

- *Element q_{ij} is defined as $q_{ij} = P(x_{t+1} = S_j|x_t = S_i)$.*
- *For each row i (e.g. value that x_t can take) the row sums to one, $\sum_j q_{ij} = 1$.*

- *There is no restrictino on the sum of a column*

An example of a Markov chain would be a state space $S = \{-1, 1\}$ so that x_t can only ever be -1 or 1. The transition matrix, which we'll denote by Q, could be formed like this:

$$Q = \begin{bmatrix} 3/4 & 1/4 \\ 1/5 & 4/5 \end{bmatrix} \tag{10.29}$$

The row indicates the current state, so row 1 refers to the state where $x_t = -1$ and row 2 to $x_t = 1$. The first column indicates the probability of x_{t+1} being in the -1 state and second column to being in the 1 state. As you can see, rows have to sum to one, as given today's outcome, x_{t+1} has to take on some value. But the column need not sum to one, and their relative size indicates which state is more or less likely in the future.

Markov chains are handy because they have nice simplifying properties, because the Markov property removes the dependence on multiple past realizations. In short, you can just apply the transition matrix an arbitrary number of times to find out the probabilities an arbitrary number of periods in the future.

Conclusion 10.2 (Arbitrary period probability) *Given x_t is in state i, the probability of state j being the realization of x_{t+k} is Q^k_{ij}, the i-jth entry in the transition matrix to the kth power.*

The transition matrix Q gives you the *conditional* probability of j occurring given x_t is in the i state. What about the unconditional probability of j occurring? A different way to think of this is what the long-run distribution of values of x_t will be.

Conclusion 10.3 (Unconditional distribution) *The vector P is the $S \times 1$ vector of unconditional probabilities of each state occurring (or alternately the expected share of periods spent in each state). Given a Markov chain with transition matrix Q, that vector is defined by*

$$P = Q'P \tag{10.30}$$

And now that we know all this about Markov chains, let's connect this back to our more flexible stochastic process, the AR(1). It can be shown that

Conclusion 10.4 (Markov representation of AR(1)) *The AR(1) process $x_t = \rho x_{t-1} + \varepsilon_t$ with variance σ^2_ε for the shock can be approximated with a Markov chain where*

$$S = \left[\frac{-\sigma^2_\varepsilon}{1-\rho^2}, \frac{\sigma^2_\varepsilon}{1-\rho^2}\right] \tag{10.31}$$

and the transition matrix is

$$Q = \begin{bmatrix} \frac{1+\rho}{2} & \frac{1-\rho}{2} \\ \frac{1-\rho}{2} & \frac{1+\rho}{2} \end{bmatrix} \tag{10.32}$$

Note that an AR(1) approximation has a transition matrix that is symmetric in the sense that the columns sum to one, meaning that you have an equal chance of either state. That should fit the intuition on AR(1) processes. The shock was normal around zero. There was an equal probability of a good (positive) or bad (negative) draw. That's retained here. The persistence of the shock (the ρ) value shows up in the transition matrix. If ρ is close to one, then once you are in a state you tend to stay there. And if ρ is close to zero it's just completely random. Note that ρ also influences the two states. The higher is ρ the bigger the shocks are, which is capturing the fact that if ρ is big you are carrying around the persistent effect of past shocks with you.

Conclusion 10.5 (Properties of Markov AR(1)) *The AR(1) process $x_t = \rho x_{t-1} + \varepsilon_t$ with variance σ_ε^2 for the shock can be approximated with a Markov chain with a long-run expected value of*

$$E[x_t] = P_1 \frac{-\sigma_\varepsilon^2}{1-\rho^2} + P_2 \frac{\sigma_\varepsilon^2}{1-\rho^2} = (P_2 - P_1)\frac{\sigma_\varepsilon^2}{1-\rho^2} = \mu_x, \tag{10.33}$$

and a variance of

$$V[x_t] = E[(x_t - \mu_x)^2] = \left[\frac{\sigma_\varepsilon^2}{1-\rho^2}\right]^2 \left(1 - (P_2 - P_1)^2\right) \tag{10.34}$$

Think of these properties as being the expected value and variance along a balanced growth path over many iterations of the process, as opposed to the expected value or variance of *next period's* shock. The variance depends on how noisy the white noise is, naturally,

10.5 Value function iteration with shocks

It turns out that our simplistic two-state shock is not that simplistic, and we can approximate a lot of things with is. So let's go back to this characterization

$$\begin{aligned} v(k, Low) &= \max_{k'} \left[u(y^{Low} + (1-\delta-g_L)k - k') + \beta E[v(k')|Low] \right] \\ v(k, High) &= \max_{k'} \left[u(y^{High} + (1-\delta-g_L)k - k') + \beta E[v(k')|High] \right]. \end{aligned}$$

where now we know that *Low* and *High* could refer to the outcomes of an AR(1) process that we approximated with a two-state Markov chain with some state space S and transition matrix Q.

We have to be a little careful here. Note that the expectation terms are now altered. In $v(k, Low)$ we have $E[v(k')|Low]$, and that's because we're taking the expectation of the value of having k' *conditional of having the Low state today*. We need this because the probability weights inside that expectation may not be identical to those in $E[v(k')|High]$. The Markov transition matrix is what tells us how to form those weights.

$$
\begin{aligned}
E[v(k')|Low] &= q_{Low,Low}v(k', Low) + q_{Low,High}v(k', High) \\
E[v(k')|High] &= q_{High,Low}v(k', Low) + q_{High,High}v(k', High)
\end{aligned}
$$

Like our original problem, we are kind of stuck here, and like our original problem we can get an answer using value function iteration. We're just going to iterate through two Bellman equations.

- Start with an initial guess of the value functions, like $v_0(k', Low) = 0 \quad \forall k'$ and $v_0(k', High) = 0 \quad \forall k'$. The subscript 0 here refers to the iteration of our guess about the value function, and not to a time period.
- Solve the two problems, $v_1(k, Low)$ and $v_1(k, High)$ given the guesses for v_0. Thus we're solving for what the planner would do if they had k, and either value of productivity, given the v_0 value functions.
- Iterate again, and solve the problems for $v_2(k, Low)$ and $v_2(k, High)$ given the functions/lists v_1 from the prior step.
- Continue until both value functions have converged (not to each other, each to their individual long-run answer) $v_j(k, Low) \approx v_{j+1}(k, Low)$ and $v_j(k, High) \approx v_{j+1}(k, High)$.

Before we needed the computer to figure out the stable arm, because the Bellman equation (or the Lagrangian or Hamiltonian) didn't have a nice solution. But given the computer, we could solve for the whole stable arm. Here, we're going to use the computer to solve for the value function, and that in turn will implicitly give us the policy functions, $g(k, Low)$ and $g(k, High)$. But that's it. We can't solve for a "stable arm" because every period there will be a shock. The best we could do is use the computer to simulate the economy, using a random number generator to give us a series of shocks. With uncertainty the key is to be able to understand how the decision process works.

11

Linearizing and solving the Ramsey

These notes are just about linearization and are not new material you need to know for a midterm or the comp. They are very rough and are only intended to help understand the script we'll use in class.

We've had several long-winded ways of evaluating and solving the Ramsey model. The forward looking problem (AD equilibrium and trying to pick c_0) as well as the recursive problem (Bellman equation and value function iteration). Those are just ways of "seeing" the problem and finding a way to a numeric solution given that there isn't an obvious equation one can write down for c_0 (or any c_t) given the initial conditions.

There is yet another way to solve the Ramsey model which involves cheating. Not really cheating, but solving it in an approximate way near the steady state. That means we'd ignore some of the non-linearity that happens "far" from steady state, but if we're willing to believe that the shocks or changes that hit the economy are not too big, this would be fine. More to the point, this retains all the important intuition about how various parameters change the outcomes and transitions speeds in response to shocks.

Doing this requires a lot of mathematical machinery. You have to take a first order Taylor series expansion of the Ramsey model around steady state to get a linear system of differential equations. Then you have to use standard techniques for solving linear systems of differential equations to find a final solution, and that involves things like eigenvalues and eigenvectors. Tedious, but possible. In the end what you get is an expression like the following for the exact time path of the capital/output ratio.

$$(k/y)_t = (k/y)^* + e^{-\lambda t}((k/y)_0 - (k/y)^*) \tag{11.1}$$

or the capital/output ratio can be written as a weighted sum of the initial and the steady-state. The weight evolves with time, and given $\lambda > 0$, the weights evolve such that when $t = 0$, $(k/y)_t = (k/y)_0$ and as $t \to \infty$ we have $(k/y)_t = (k/y)^*$. Thus we get the expected result that the capital/output ratio slowly transitions to the steady state from any jinitial value we start with. The value of λ is a collection of parameters that arise out of the Taylor series expansion and solving the differential system, but for the moment just accept that this format works.

From this you can also get things in log terms and then growth rates. What I mean is that

$$\ln((k/y)_t/(k/y)^*) \approx \frac{(k/y)_t - (k/y)^*}{(k/y)^*} \approx e^{-\lambda t}\frac{(k/y)_0 - (k/y)^*}{(k/y)^*} \tag{11.2}$$

and therefore

$$\ln(k/y)_t = \ln(k/y)^* + e^{-\lambda t}(\ln(k/y)_0 - \ln(k/y)^*). \tag{11.3}$$

Given this form, the growth rate of the capital/output ratio is

$$g_{KY} \approx d\ln(k/y)_t = \lambda e^{-\lambda t})\ln(k/y)^* - \lambda e^{-\lambda t}\ln(k/y)_0 = \lambda e^{-\lambda t}\left(\ln(k/y)^* - \ln(k/y)_0\right) \tag{11.4}$$

so that the growth rate is positive when the economy is below steady state, $\ln(k/y)^* - \ln(k/y)_0 > 0$, and the growth rate slows down as t gets bigger and the economy gets closer to steady state. If you want to think of the $g_{K/Y}$ diagram from the Solow model, then what you're thinking about is graphing g_{KY} against $(k/y)_0$, or what is the growth rate for any given state of k/y. That curve is than evaluated at $t = 0$,

$$g_{KY} \approx \lambda\left(\ln(k/y)^* - \ln(k/y)_0\right) \tag{11.5}$$

which is now just a line with a negative slop. This is linear in logs, which means it isn't necessarily linear in k/y itself. because we did the Taylor series expansion. Given that λ moves with parameters (like σ)) we can evaluate how this curve compares across economies. Regardless, note that this system is stable so long as λ is positive, because then the as k/y_0 goes up, g_{KY} goes down.

Given that result, everything else can be inferred from the dynamic budget consraint and the Euler equation.

$$\ln y_t = \frac{\epsilon_K}{\epsilon_L}\ln(k/y)_t + \ln A_0 + g_A t, \tag{11.6}$$

and we know exactly the path for this now, while for consumption we can work in several ways to get that. One way is to note that we have this relationship

$$g_{KY} = \epsilon_L\left(\frac{s_I}{k/y} - (\delta + g_A + g_L)\right) \tag{11.7}$$

which defines g_{KY} in terms of the savings rate, and so we can rearrange and see that we could write

$$s_{It} = \left(g_{KY}/\epsilon_L + (\delta + g_A + g_L)\right)(k/y)_t \tag{11.8}$$

and we know how to calculate both g_{KY} and $(k/y)_t$ in terms of known entities now. The ambiguity in what happens to the savings rate comes because we know (if we are below steady state) that $(k/y)_t$ is growing over time (pushing up on the savings rate) but g_{KY} is going to be falling over time as we approach steady state (pushing down on the savings rate), so it's a race between those two for what happens to savings. Given s_{It} and $\ln y_t$ we could easily solve for $\ln c_t$.

$$\ln c_t = \ln(1 - s_{It}) + \ln y_t, \tag{11.9}$$

and that includes for c_0. So we know how this all works.

We know that

$$R_t = \epsilon_K (y/k)_t \tag{11.10}$$

(or alternatively that it is $R = s_K y/k$), so that we can solve immediately for R and $r_t = R_t - \delta$. Note that once we have this we can solve for

$$g_{ct} = \frac{1}{\sigma}(\epsilon_K (y/k)_t - \delta - \theta) \tag{11.11}$$

and we know the growth rate of consumption at each point in time. We could have gotten that as well from knowing $\ln c_t$, but this is an immediate way of seeing it.

What this amounts to is that we can characterize the *response* to any shock if we can characterize what it means for $(k/y)^*$ and $\ln(k/y)_0$ at the point the shock occurs. The other thing we need to know is that the shock does immediately, which includes what it does to $\ln(k/y)_0$. But it also might include that prior to the shock we were on an old BGP (a different $(k/y)^*$) or had a different level of $\ln A_0$.

For all intents and purposes, if you work with this setting, taking the value of λ as some given, then you can solve the entire Ramsey model for all periods, including c_0.

11.0.1 Comparing σ

If we're comparing economies with different σ values, then one of the problems we'd have was ensuring that σ is only influencing the transition path, and not the steady state itself. The way to ensure this or to assume that this is the case is to assert that both economies have the same r^* regardless of the value of σ. Since

$$r^* = \theta + \sigma g_A \tag{11.12}$$

this amounts to assuming that θ differs between them by just the right amount to offset the difference in σ. If you let g_A be different this would have additional effects on things (and they'd be hard to compare). So in essence a high σ economy that is unwilling to substitute has to be a little less patient in order to end up with the same r^* as the low σ economy. That's fine, as it ensures they have the same $(k/y)^*$ and that we can plot or compare them to one another in an easy way.

The difference in σ shows up in the λ parameter on convergence speed. In the derivation below you can see exactly how it influences this, and it's complicated, but in practice what happens is that when σ is high, λ is low (convergence is slow) while when σ is low, λ is high (convergence is fast).

11.0.2 Uncertainty

This works with stochastic shocks as well. All the above equations would give us the time paths of the *expected* values of things like y_t and c_t and k/y. We could solve the above equations and then add in a stochastic shock to y_t if we wanted to generate an example of a time path for GDP per capita (or other variables).

The stochastic shocks did add a term in the Euler equation, so that now we have

$$E[g_{ct}] = \frac{1}{\sigma}(\epsilon_K (y/k)_t - \delta - \theta + Z) \tag{11.13}$$

where Z is just some parameters that govern how noisy the shocks are (remember how the AR(1) showed up). By adding this in here, we know that in steady state it would have to be that in steady state where $E[g_{ct}] = g_A$

$$\epsilon_K (y/k)^* = \delta + \theta + \sigma g_A - Z \tag{11.14}$$

or

$$(k/y)^* = \frac{\epsilon_K}{\delta + \theta + \sigma g_A - Z} \tag{11.15}$$

is *higher* than in the certainty case because of Z being here. This also means that savings must be higher to support that higher k/y ratio. For our purposes all Z does is push up the savings rate, push up the BGP of GDP per capita, and push up the BGP for consumption.

11.1 Solving linear differential equation systems

In principle, you can stop reading now. Everything that follows is *how* to get to the equation in (11.1). If you take that equation on faith, and pick a value of λ, you can solve the Ramsey and evaluate the consequences of any shock. You can also use it to infer what could have explained any data series you are looking at. All the nuance of the Ramsey gets shoved into the λ parameter. If you want to know how to get to that, everything else in these notes tells you that.

Take the simplest version of this first. Let's say you have a linear differential equation of $dx = ax + b$, so that the change (with respect to time) of x depends on x itself, and a and b are constants. We're looking for a few things. What is the steady state? Is that steady state stable? Can we characterize the value of x at all points of time? Here the answers are pretty simple.

First, the steady state is where $dx = 0$, so this is when $x^* = -b/a$. Stability requires a little more work, as now we need to know whether x will approach the steady state or not. To see this, we'll make the first adjustment to notation which is to look at everything relative to the steady state, so let

$$z = x - x^* \tag{11.16}$$

be the variable denoting how far from steady state we are. Then it follows that

$$dz = az \tag{11.17}$$

because $dz = dx$ and $az = ax - ax^* = ax - a(-b/a) = a + b$. Anyway, now that we have a differential equation without the bothersome constant, we can immediately see that it implies

$$dz/z = a \tag{11.18}$$

or that the growth rate of z is constant. The *gap* between where we start and where we end up will grow at a constant rate. We want that rate to be *negative* so that z approaches zero. But we already know that this system has a solution of

$$z_t = z_0 e^{at} \tag{11.19}$$

because this is just exponential growth. The initial value $z_0 = x_0 - x^*$ or how far we start from steady state. Knowing what we know, we can write this as

$$x_t = x^* + e^{at}(x_t - x^*) \tag{11.20}$$

so that the value of x_t approaches x^{ast} only if $a < 0$ and the term e^{at} goes to zero over time. We know exactly what x_t will be at each point in time, and we know that this depends in part on how far from steady state it starts out.

One last thing to establish, which will seem trivial, is that if we scale our modified system by some constant v, like this

$$z_t/v = z_0/ve^{at} \tag{11.21}$$

then nothing changes about the solution for steady state or about the assessment of stability. If $a < 0$, then as t goes to infinity $z_t/v \rightarrow 0$, meaning it must be that $z_t \rightarrow 0$, or this system goes to steady state.

11.2 Systems of differential equations

What about when there is a system of differential equations? In particular, a linear system of differential equations like

$$\begin{pmatrix} dx_1 \\ dx_2 \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} + \begin{pmatrix} b_1 \\ b_2 \end{pmatrix} \tag{11.22}$$

where the change in both x_1 and x_2 depends on the level of x_1 and x_2, where the coefficients in the matrix tell us how much they depend on them. Writing this in vector form we would have

$$dx = Ax + b \tag{11.23}$$

where A is the matrix of coefficients. Again, the steady state is going to be

$$x^* = -A^{-1}b \tag{11.24}$$

and notice the similarity to before (inverse of a times b). Again, we can write things in terms of deviations from steady state

$$dz = Az \tag{11.25}$$

where dz is the vector $(dx_1, dx_2)'$ and z is the vector $(x_1 - x_1^*, x_2 - x_2^*$.

Just like the simple one-dimensional version of this, there is a standard answer, and it looks similar. The answer is that

$$z_t = e^{At} z_0 \tag{11.26}$$

or

$$x_t = x^* + e^{At}(x_0 - x^*) \tag{11.27}$$

where everything is interpreted as vectors/matrices. Now, our problem is *stable* if we have something like $A < 0$, which isn't quite right because A is a matrix. So we need to understand how to figure out if A is negative-ish with respect to the elements of x such that as t gets bigger this e^{At} term goes to zero.

The problem, of course, is that evaluating e^{At} is not easy, and moreover we don't just want a theoretical condition that ensures it converges to steady state, but in an ideal world a way to solve for the exact value of x_t at every given point of time, meaning we want the exact form of e^{At}.

We can get somewhere by noting that we can break down a square matrix A into the following

$$A = VDV^{-1} \tag{11.28}$$

where V is a matrix of eigenvectors (columns) and D is a diagonal matrix containing the eigenvalues of A. All standard linear algebra stuff. The more important quality of this relationship for us is that another linear algebra result is that

$$e^{At} = Ve^{Dt}V^{-1} \tag{11.29}$$

or that the exponential of A depends on the exponential of the diagonal part of it. Use this and go back to our simple relationship of $z_t = e^{At}z_0$, and we have

$$V^{-1}z_t = e^{Dt}V^{-1}z_0. \tag{11.30}$$

Notice that what we've done is "scale" z_t and z_0 by the same matrix V, just like in the simple one-dimensional case. And just like we did in that case, we can see that if this scaled form converges to steady state then so does z_t.

So we're going to change notation once again and write the following

$$w_t = e^{Dt}w_0 \tag{11.31}$$

where $w_t = V^{-1}z_t$ is just the scaled version of z_t. Now we have a system that we can evaluate with relative ease, in the sense that e^{Dt} has a easy form to work with, as it is diagonal. For a simple 2-variable system this says that

$$w_{1,t} = e^{d_1 t}w_{1,0} \tag{11.32}$$

$$w_{2,t} = e^{d_2 t}w_{2,0} \tag{11.33}$$

and now it's just two separate dynamic relationships that we can understand in a straightforward way.

If both d_1 and d_2 are negative, then both w_1 and w_2 go to steady state, and if that's true then it must be that (both elements of) z_t goes to steady state. If both d_1 and d_2 are positive, then things spin out of control, and therefore so does z_t.

If only $d_1 < 0$ (WLOG) and $d_2 > 0$, then it's *possible* for this system to converge to steady state, and in that sense we say it is "saddle-path stable", meaning there is one route towards the steady state. The only way this works is if the explosive compoenent *starts* at steady state, meaning $w_{2,0} = 0$. Because $d_1 < 0$, the value of $w_{1,0}$ can be anything and things will converge.

This saddle path case is the most relevant for the Ramsey problem, as we'll see. It's worth establishing some results on what that looks like though using this generic form of the problem. We know that

$$\begin{pmatrix} w_{1,0} \\ w_{2,0} \end{pmatrix} = \frac{1}{v_{11}v_{22} - v_{12}v_{21}} \begin{pmatrix} v_{22} & -v_{12} \\ -v_{21} & v_{11} \end{pmatrix} \begin{pmatrix} z_{1,0} \\ z_{2,0} \end{pmatrix} = \frac{1}{v_{11}v_{22} - v_{12}v_{21}} \begin{pmatrix} v_{22}z_{1,0} - v_{12}z_{2,0} \\ -v_{21}z_{1,0} + v_{11}z_{2,0} \end{pmatrix} \tag{11.34}$$

where all the v terms are elements of the inverse of V. If we assert that $w_{2,0} = 0$, then it has to be that

$$-v_{21}z_{1,0} + v_{11}z_{2,0} = 0 \tag{11.35}$$

or that

$$z_{2,0} = \frac{v_{21}}{v_{11}}z_{1,0} \tag{11.36}$$

or that there is a strict proportional relationship between the initial state of both variables. Remember, these z variables are deviations from steady state, and this is telling us that in the saddle path case the system converges to steady state only if the distance to steady state of z_2 is proportional to z_1.

Can we count on that happening? It depends on what kind of system we're trying to describe. In this economic case, we have that one of the initial variables (capital) will be given, so that's not possible to change. But the other initial value (consumption) will be a choice, and we can set it to the right amount to ensure this holds.

Moreover, assuming we set things this way, we get that

$$\begin{pmatrix} w_{1,0} \\ w_{2,0} \end{pmatrix} = \frac{1}{v_{11}v_{22} - v_{12}v_{21}} \begin{pmatrix} v_{22} & -v_{12} \\ -v_{21} & v_{11} \end{pmatrix} \begin{pmatrix} z_{1,0} \\ z_{2,0} \end{pmatrix} = \begin{pmatrix} z_{1,0}/v_{11} \\ 0 \end{pmatrix} \tag{11.37}$$

Meaning we know how to characterize the entire initial state of w that *ensures* we converge to steady state. That means our system $w_t = e^{Dt}w_0$ is now

$$w_{1,t} = e^{d_1 t}z_{1,0}/v_{11} \tag{11.38}$$

$$w_{2,t} = 0. \tag{11.39}$$

Okay, but we want this in terms of z_t, not w_t. Reversing course, we know that $w_t = V^{-1}z_t$ so $Vw_t = z_t$

$$\begin{pmatrix} z_{1t} \\ z_{2t} \end{pmatrix} = \begin{pmatrix} v_{11} & v_{12} \\ v_{21} & v_{22} \end{pmatrix} \begin{pmatrix} e^{d_1 t}z_{1,0}/v_{11} \\ 0 \end{pmatrix} = \begin{pmatrix} e^{d_1 t}z_{1,0} \\ e^{d_1 t}\frac{v_{21}}{v_{11}}z_{1,0} \end{pmatrix} \tag{11.40}$$

and we know the entire time path of z_t for both variables. Both converge to steady state here given that $d_1 < 0$, and *remember* that this only works because we picked the initial value $z_{2,0}$ to be precisely the right value to ensure that our second dimension was always in steady state.

11.3 Eigenvalues and Eigenvectors

We've gotten pretty far, but note that we need the eigenvalues and eigenvectors for the matrix A to make this work. For a given matrix

A the eigenvalues are found by solving the characteristic equation of A which is

$$\det(A - \lambda I) = 0 \tag{11.41}$$

for all values of λ that work. For a generic matrix with entries a, b in the first row and c, d in the second, this is

$$(a - \lambda)(d - \lambda) - bc = 0 \tag{11.42}$$

which expands into the polynomial

$$\lambda^2 - (a + d)\lambda + (ad - bc) = 0 \tag{11.43}$$

and now you have to take the roots of this thing. That's do-able, and the answers are

$$\lambda_1 = (a + d)/2 - \left((a + d)^2/4 - (ad - bc)\right)^{1/2} \tag{11.44}$$

$$\lambda_2 = (a + d)/2 + \left((a + d)^2/4 - (ad - bc)\right)^{1/2}. \tag{11.45}$$

Great. Note that the difference in these two is simply the sign in the middle, and so whether either of these will be negative or positive depends on the relative size of the terms. But in principle λ_1 is the one we are thinking of as being negative.

We also need the eigenvectors for each one of these values so we can collect them into V. Without going through all the steps here, for a simple 2x2 we can establish that the V matrix will be either

$$V = \begin{pmatrix} b & b \\ \lambda_1 - a & \lambda_2 - a \end{pmatrix} \tag{11.46}$$

if $b \neq 0$ or

$$V = \begin{pmatrix} \lambda_1 - d & \lambda_2 - d \\ c & c \end{pmatrix} \tag{11.47}$$

if $c \neq 0$, and we'd choose the form (or either arbitrarily) depending on the values of b and c in the original A matrix.

Note that with our prior section we established something about the eigenvalues and eigenvectors that depends on A, so that we could then say

$$\begin{pmatrix} z_{1t} \\ z_{2t} \end{pmatrix} = \begin{pmatrix} e^{\lambda_1 t} z_{1,0} \\ e^{\lambda_1 t} \frac{\lambda_1 - a}{b} z_{1,0} \end{pmatrix} = \begin{pmatrix} e^{\lambda_1 t} z_{1,0} \\ e^{\lambda_1 t} \frac{c}{\lambda_1 - d} z_{1,0} \end{pmatrix} \tag{11.48}$$

as our solutions. It matters that we've arranged things so that the first eigenvalue is the negative one, given what we established in the prior section.

11.4 Linearizing the Ramsey

So the game is to put the Ramsey into a form that matches something like the $dz = Az$ form of the prior section and then if we can do that we know that it's the eigenvalues of the A matrix that are going to do a lot of work for us.

To get there we're going to cast the whole problem in a specific way that allows us to track things as a system of two variables, K/AL and C/AL. The reason for K/AL is that this is what we've looked at throughout, as this is just $K/Y = (K/AL)^{1-\epsilon_K}$, and the reason for C/AL is that this is a symmetric idea and once we have C/AL we can infer the values of things like C/Y to get the savings rate.

For some sanity, define

$$\hat{k} = K/AL = (K/Y)^{1/(1-\epsilon_K)} \quad (11.49)$$
$$\hat{c} = C/AL = (1-s_I)\hat{y} \quad (11.50)$$
$$\hat{y} = Y/AL = \hat{k}^{\epsilon_K} = (K/Y)^{\epsilon_K/(1-\epsilon_K)} \quad (11.51)$$

where A is the level of productivity and L is the size of the population. These are convenient forms to solve with as c or y are not steady in steady state, but grow at a constant rate. We need something that remains static in steady state. The last relationship just establishes what we already knew about how GDP and capital were related through ϵ_K, but puts that in the form of Y/AL, again to simplify things. The extra relationships show you how these things are related to stuff we already talked about like the K/Y ratio and s_I. We're just making some convenient scalings of variables to make this easier.

This implies that

$$g_{\hat{k}} = \hat{k}^{\epsilon_K - 1} - \hat{c}/\hat{k} - (\delta + g_A + g_L)$$
$$g_{\hat{c}} = \frac{1}{\sigma}\left(\epsilon_K \hat{k}^{\epsilon_K - 1} - \delta - \theta - \sigma g_A\right)$$

where $g_{\hat{c}}$ comes from the fact that we are looking at $g_{\hat{c}} = g_c - g_A$, and $g_{\hat{k}}$ is just the dynamic budget constraint for $g_{\hat{k}} = g_k - g_A$. Note that I've plugged in for $\hat{y}$ here already

The differential systems are in the form of changes, dx, not growth rates, so this system is

$$d\hat{k} = \hat{k}^{\epsilon_K} - \hat{c} - (\delta + g_A + g_L)\hat{k} \quad (11.52)$$
$$d\hat{c} = \frac{1}{\sigma}\left(\epsilon_K \hat{k}^{\epsilon_K - 1} - \delta - \theta - \sigma g_A\right)\hat{c} \quad (11.53)$$

This isn't quite what we need yet. The vector $dx = (d\hat{k}, d\hat{c})'$ is our vector of changes in variables, and so we should have this in a linear system with $x = (\hat{k}, \hat{c})'$, but as you can see this isn't quite how things work because there are non-linear relationships between dx and x.

We can get this into a linear form we can analyze by doing a Taylor series expansion around the steady state.

For a system of equations, a Taylor expansion of a non-linear system works just like a regular one-variable Taylor expansion. In our case, if we have $dx = g(x)$ as our non-linear system, then

$$dx \approx g(x^*) + g'(x^*)(x - x^*) \tag{11.54}$$

where $g(x^*) = 0$ because that is the definition of a steady state, and $g'(x^*)$ is the Jacobian (the set of derivatives of each element of x with respect to all other elements) of g evaluated at steady state itself. Even though $g(x)$ is non-linear, with $g'(x^*)$ evaluated at steady state with fixed values of x^*, it is now just a matrix of fixed coefficients. In other words, $dx \approx g'(x^*)(x - x^*)$ is a linear system of differential equations. In fact, because $dx = d(x - x^*)$ because steady states are, well, steady, it's the case that our Taylor expansion is already in this form

$$dz \approx g'(x^*)z \tag{11.55}$$

and here the A matrix of coefficients is just $A = g'(x^*)$.

All that means we need to find this Jacobian thing for our existing Ramsey model. What we are after is this general Jacobian,

$$g'(x) = \begin{pmatrix} \partial d\hat{k}/\partial\hat{k} & \partial d\hat{k}/\partial\hat{c} \\ \partial d\hat{c}/\partial\hat{k} & \partial d\hat{c}/\partial\hat{c} \end{pmatrix} \tag{11.56}$$

and then we want to fix this Jacobian as a set of fixed coefficients by evaluating each of those partials at the steady state.

For our system of equations we get

$$g'(x) = \begin{pmatrix} \epsilon_K \hat{k}^{\epsilon_K - 1} - (\delta + g_A + g_L) & -1 \\ \frac{\hat{c}}{\sigma}\epsilon_K(\epsilon_K - 1)\hat{k}^{\epsilon_K - 2} & \frac{1}{\sigma}\left(\epsilon_K \hat{k}^{\epsilon_K - 1} - \delta - \theta - \sigma g_A\right) \end{pmatrix} \tag{11.57}$$

but evaluated at the steady state we know that

$$g'(x^*) = \begin{pmatrix} \epsilon_K (\hat{k}^*)^{\epsilon_K - 1} - (\delta + g_A + g_L) & -1 \\ \frac{\hat{c}^*}{\sigma}\epsilon_K(\epsilon_K - 1)(\hat{k}^*)^{\epsilon_K - 2} & 0 \end{pmatrix} \tag{11.58}$$

and simplifying further we have

$$g'(x^*) = \begin{pmatrix} (\epsilon_K - s_I^*)(y/k)^* & -1 \\ -\frac{\epsilon_K \epsilon_L}{\sigma}(1 - s_I^*)[(y/k)^*]^2 & 0 \end{pmatrix} \tag{11.59}$$

What this all means is that we have a dynamic system that looks like this

$$dz \approx \begin{pmatrix} (\epsilon_K - s_I^*)(y/k)^* & -1 \\ -\frac{\epsilon_K \epsilon_L}{\sigma}(1 - s_I^*)[(y/k)^*]^2 & 0 \end{pmatrix} z \tag{11.60}$$

where $z = (\hat{k}, \hat{c})'$. Putting this in more familiar terms from the prior section, we have a system that is of the form $dz = Az$, and we know what A looks like now. One thing to note here is that if two economies have the same steady state k/y and s_I (possible), then the only thing that differentiates them is the value of σ, and higher σ is going to mean a smaller element in that lower left, which will result ultimately in a slower convergence speed and lower λ_1.

Thus for the Ramsey model we need to know the eigenvalues and the eigenvectors. Plugging into what we've got from the prior section.

$$\begin{aligned}\lambda_1 &= (\epsilon_K - s_I^*)(y/k)^*/2 - \left(((\epsilon_K - s_I^*)(y/k)^*)^2/4 + \frac{\epsilon_K\epsilon_L}{\sigma}(1-s_I^*)[(y/k)^*]^2\right)^{1/2} \\ \lambda_2 &= (\epsilon_K - s_I^*)(y/k)^*/2 + \left(((\epsilon_K - s_I^*)(y/k)^*)^2/4 + \frac{\epsilon_K\epsilon_L}{\sigma}(1-s_I^*)[(y/k)^*]^2\right)^{1/2}.\end{aligned}$$

and have fun with that. We also need some eigenvector information, and that's

$$\begin{aligned}v_{11} &= (\epsilon_K - s_I^*)(y/k)^*/2 - \left(((\epsilon_K - s_I^*)(y/k)^*)^2/4 + \frac{\epsilon_K\epsilon_L}{\sigma}(1-s_I^*)[(y/k)^*]^2\right)^{1/2} \\ v_{21} &= -\frac{\epsilon_K\epsilon_L}{\sigma}(1-s_I^*)[(y/k)^*]^2\end{aligned}$$

so again, have fun with that.

Just going back to the generic solution to the differential. Remember that we had to set one of the $w_{2,0}$ terms to zero to ensure the system didn't explode. That's equivalent here to setting c_0 to the right thing to ensure we are on a stable arm towards the steady state. We have to pick just the right consumption value so that things work. That value in our case is

$$\hat{c}_0 - \hat{c}^* = \frac{-\frac{\epsilon_K\epsilon_L}{\sigma}(1-s_I^*)[(y/k)^*]^2}{(\epsilon_K - s_I^*)(y/k)^*/2 - \left(((\epsilon_K - s_I^*)(y/k)^*)^2/4 + \frac{\epsilon_K\epsilon_L}{\sigma}(1-s_I^*)[(y/k)^*]^2\right)^{1/2}}(\hat{k}_0 - \hat{k}^*)$$

A
Supplemental material

A.1 Logs and growth rates

Throughout the notes we'll be talking about growth, and so the mathematics and notation of dealing with things that grow is crucial. Let's start with a simple function that the value of some variable z depends on time, $z(t)$. Take the total derivative of this function

$$dz(t) = z'(t)dt \tag{A.1}$$

or the absolute change in $z(t)$ depends on the derivative $z'(t)$ and the actual change in time. If you divide both sides of this by $z(t)$, you get

$$\frac{dz(t)/dt}{z(t)} = \frac{z'(t)}{z(t)}. \tag{A.2}$$

On the left is the growth rate of z; the change in z for a given change in time divided by z itself. We'll be concerned a lot with growth rates, so I'm going to introduce some new notation here to avoid writing out this whole fraction over and over again.

$$g_z \equiv \frac{dz(t)/dt}{z(t)}. \tag{A.3}$$

Any time you see g_z anywhere, it means the growth rate of z (or whatever variable is in the subscript). As we'll often implicitly assume that the change in time is $dt = 1$, I will also often refer to $dz(t)/z(t)$ as g_z. We'll sometimes use the following notation as well to save on notation in certain situations,

$$G_z \equiv 1 + g_z, \tag{A.4}$$

and G_z is sometimes referred to as the "growth factor" of z.

That definition works fine, but we can establish a nice relationship between the (natural) log of $z(t)$ and the growth rate that will save us a lot of time in the work ahead. Start with $\ln z(t)$, and again take the total derivative,

$$d \ln z(t) = \frac{z'(t)}{z(t)} dt. \tag{A.5}$$

This expression means that

$$\frac{d \ln z(t)}{dt} = \frac{z'(t)}{z(t)} = g_z. \tag{A.6}$$

The derivative of *log* $z(t)$ for a given change in time is equal to the growth rate of z. This is incredibly useful for us. It effectively says that the derivative of the log of something with respect to time is equal to its growth rate. As we'll often look at figures which plot the log of a variable against t, this gives us a simple visual way to evaluate the growth rate. The slope of $\ln z(t)$ with respect to time is equal to the growth rate.

Things get even more stark if we have exponential growth, meaning there is a constant growth rate. Let

$$z(t) = z(0)e^{bt} \tag{A.7}$$

where $z(0)$ is some initial value of z, and the value of z at any time t is $z(t)$. b is a parameter and is equal to the growth rate, as we will see. Take logs and you get

$$\ln z(t) = \ln z(0) + bt \tag{A.8}$$

and now this is log-linear. If we take the derivative of this with respect to time we get $d \ln z(t)/dt = b$, or the growth rate is $g_z = b$. Moreover, note that this last equation is simply the equation for a straight line. Which can use this visually as well. If we see that the relationship of (log) $z(t)$ with time is linear, we know it has a constant growth rate. The intercept of this line tells us about the size of $z(0)$.

You can confirm this by finding $z'(t) = bz(0)e^{bt}$, and dividing by $z(t)$, so that $z'(t)/z(t) = bz(0)e^{bt}/z(0)e^{bt} = b$.

Finally, you should be familiar with the properties of growth rates of products, ratios, and powers, which all can be confirmed by taking logs and then differentiating with respect to time.

$$\begin{aligned} X = ZW \quad &\Rightarrow \quad g_X = g_Z + g_W \\ X = \frac{Z}{W} \quad &\Rightarrow \quad g_X = g_Z - g_W \\ X = Z^\alpha \quad &\Rightarrow \quad g_X = \alpha g_Z \end{aligned} \tag{A.9}$$

A.2 Taylor series expansions

At times it's useful to use Taylor series expansions to approximation functions. Formally, given some function $f(x)$ evaluated around the point a is

$$f(x) \approx f(a) + \frac{f'(a)}{1!}(x-a) + \frac{f''(a)}{2!}(x-a)^2 + ... = \sum_{n=0}^{\infty} \frac{f^{(n)}(a)}{n!}(x-a)^n. \tag{A.10}$$

We would normally work with a first-order (just the $f'(a)$ term) Taylor approximation, eliminating the other higher-order terms. The reason we do this is because as the weight in the denominator (the factorial) is growing rapidly and decreasing the relevance of these higher-order terms.

An example is approximating the log function using the Taylor series, which is where we get our rule of thumb that differences in logs are approximately equal to growth rates. Let $f(x) = \ln(x)$. Evaluate this at the point a and a first-order approximation is

$$\ln(x) \approx \ln(a) + \frac{1}{a}(x-a). \tag{A.11}$$

Let $x = (1+g)a$, where g is a small growth rate like $g = 0.02$. Then the Taylor series says

$$\ln(x) \approx \ln(a) + \frac{1}{a}((1+g)a - a) = \ln(a) + \frac{ga}{a} = \ln(a) + g. \tag{A.12}$$

This means that

$$\ln(x) - \ln(a) = \ln((1+g)a) - \ln(a) = \ln(1+g) \approx g. \tag{A.13}$$

A.3 From nominal to real GDP

In theory we want to evaluate the change in real GDP, meaning the aggregation of changes in real quantities consumed, each weighted by their expenditure share, as shown in equation (**??**). How does the data on real GDP produced by a national statistical agency match up to this?

To explain, let's start with a simple two-good situation, and build the intuition for the issues that come up with measuring real GDP. In our little two-good world, the real amount of c_1 and c_2 produced and consumed is a result of an interaction of our preferences for the two goods and the production possibilities frontier that tells us the physical trade-off in producing good 1 in terms of good 2. But we cannot observe the utility function or the PPF, we only observe the actual market outcomes. That is, we see the nominal price that each sells for p_1 and p_2, and we know the quantity produced, c_1 and c_2 (or we see expenditure on good 1 and good 2, and divide by price to recover the quantity).

Moreover, we observe these market outcomes both in the present, and in the past. And we'd like to compare the bundle of goods produced in those two periods and figure out whether we are better or worse off - whether economic growth occurred. If $c_{2,Present} > c_{2,Past}$ and $c_{1,Present} > c_{1,Past}$, then it is straightforward to conclude that the present is better off than the past (assuming typical utility over the two goods).

Figure A.1 gives a simple example of this. That figure includes a line that has slope $-p_1/p_2$, which gives the relative price of the two goods. As drawn, that relative price is different in past and present, but note that this doesn't change the conclusion that the present is better off than in the past in this case. A second thing to note is that pure inflation in prices is not an issue in the Figure. We have the real quantities graphed, and the slope of the line is unaffected by the absolute price level, because it captures the relative price of the two goods. What we cannot do just by looking at the Figure, however, is put a firm number on the increase in real consumption.

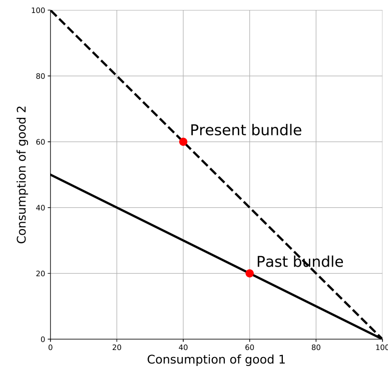

Figure A.1: A change in consumption where the Present bundle is an obvious improvement over the Past bundle.

The bigger issue comes when we have a situation as in Figure A.2, where the present consumes more of good 2, but less of good 1. Now, it is not obvious whether the economy is better off or not. If we valued everything at the relative prices of the past, then the present is obviously better (i.e. it would appear that the constraint shifted up). But if we valued everything at the relative prices of the present, then it would look like the *past* is obviously better (i.e. it would again appear that the constraint shifted up). Notice that this is not an issue with pure inflation in prices, as again all that we're focusing on here is the relative price of goods within a given period.

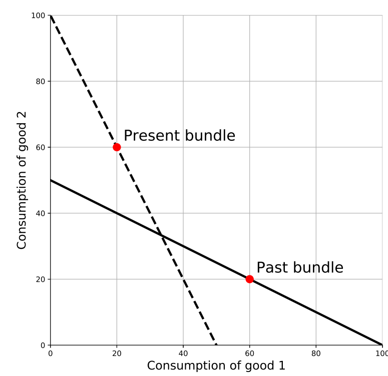

Figure A.2: A change in consumption where it is not obvious if the Present or Past bundle is better.

How do we compare past and present, then, and figure out whether the real value of the consumption bundle has increased or decreased? The answer is that we're going to compare them from both perspectives (e.g. compare them using the relative prices of the past, and then compare then using the relative prices of the present) and do something like average those perspectives. "Holding prices constant" in the data means averaging the prices of past and present, and if we do that, we can map the measured change in real GDP to the theoretical definition.

We want to do this for far more than two goods, however, so several pieces of notation are necessary before starting. Let P_{jt} be the nominal price of good j at time t, and let c_{jt} be the real quantity of good j produced/consumed at time t. Let $e_{jt} = p_{jt}c_{jt}$ be the nominal expenditure on good j at time t. Finally, let total nominal expenditure on the goods from time t, *valued at the prices from time* s, be denoted E_{st}. Putting this all together, we can establish the following equalities,

$$E_{st} = \sum_{j \in J} p_{js}c_{jt} = \sum_{j \in J} \frac{p_{js}}{p_{jt}} e_{jt}.$$

The value E_{tt} is therefore nominal expenditure on goods purchased at time t, valued at the prices of time t, or nominal GDP.

So how do we compare the output of products at time a to the output of products at time b, where without any loss of generality we'll assume that b comes after a? As we did with the two-good case, we want to value products from both periods in time a prices, and then value products from both periods in time b prices.

The ratio of real output from period b to real output from period a, using period a prices, is

$$\left(\frac{Y_b}{Y_a}\right)_a = \frac{\sum_{jinJ} p_{ja}c_{jb}}{\sum_{jinJ} p_{ja}c_{ja}} = \frac{E_{ab}}{E_{aa}}. \tag{A.14}$$

The alternative is to use period b prices, or

$$\left(\frac{Y_b}{Y_a}\right)_b = \frac{\sum_{jinJ} p_{bj}c_{bj}}{\sum_{jinJ} p_{bj}c_{aj}} = \frac{E_{bb}}{E_{ba}} \tag{A.15}$$

These both satisfy our requirement that we compare the quantities from periods a and b using a common set of relative prices. But they might differ because the relative prices of goods in a could be different than in b. As neither one has a greater claim on the truth, we want to combine the information in both.

A simple average will not work, because these are ratios, but a geometric average will work.

We want a geometric average so that symmetric ratios (e.g. 2 and 1/2) result in an average ratio of 1, while an arithmetic average would give us 1.25.

$$\left(\frac{Y_b}{Y_a}\right)_F = \left(\left(\frac{Y_b}{Y_a}\right)_a \left(\frac{Y_b}{Y_a}\right)_b\right)^{1/2} = \left(\frac{E_{ab}}{E_{aa}}\frac{E_{bb}}{E_{ba}}\right)^{1/2} \tag{A.16}$$

where I used the subscript F for this ratio because it is related to something called a "Fisher Ideal Price Index", which is described below.

Note that the best we can do is create this ratio of real output between b and a. There is no way to get an absolute number for real output, as it is meant to measure something like "living standards", which have no units. We might decide that a given year (e.g. 2009 or 2013) is our "base" year, and assign it a real index of 100, and use the ratios we calculate with other years to create index values relative to that. For example, if $(Y_{2010}/Y_{2009})_F = 1.034$, and 2009 was our base year, then we might report real GDP in 2010 as $Y_{2010} = 103.4$.

This is what we're after, but the common way to arrive at these ratios involves deflating nominal GDP in given years by prices indices. What we can show is that this deflation process delivers exactly the same ratio we just calculated. The price deflator used is again a geometric average of two different underlying price deflators (Laspeyres and Paasche) that differ in which base period they use.

The Laspeyres price index computes the ratio of prices in b relative to a, using the quantities from a (the prior period) as weights,

$$P^L_{ab} = \frac{\sum_{j\in J} P_{bj}c_{aj}}{\sum_{j\in J} P_{aj}c_{aj}}. \tag{A.17}$$

In contrast, the Paasche price index computes the ratio of prices in b relative to a, using the quantities from b (the later period).

$$P^S_{ab} = \frac{\sum_{j \in J} P_{bj} c_{bj}}{\sum_{j \in J} P_{aj} c_{bj}} \tag{A.18}$$

where I used P^S to denote this because P^P would be confusing. Again, there is nothing about the Laspeyres or Paasche that makes one better than the other. So again, we geometrically average them to get something called a Fisher Ideal Index,

$$P^F_{ab} = \left(P^L_{ab} \times P^S_{ab}\right)^{1/2} = \left(\frac{E_{ba}}{E_{aa}} \frac{E_{bb}}{E_{ab}}\right)^{1/2} \tag{A.19}$$

You can probably already see some symmetry of the Fisher price index with the way we valued real GDP above.

Note that the Fisher price index for a given year in terms of its own prices is just equal to one, or $P^F_{aa} = 1$.

Now go back and rethink calculating the ratio of real GDP across periods b and a. With some tedious algebra, you can show that

$$\left(\frac{Y_b}{Y_a}\right)_F = \left(\frac{E_{ab}}{E_{aa}} \frac{E_{bb}}{E_{ba}}\right)^{1/2} = \frac{E_{bb}/P^F_{ba}}{E_{aa}}. \tag{A.20}$$

This says that you can recover the real ratio by first deflating nominal GDP in period b (E_{bb}) by the appropriate Fisher price index that uses period a as a base (P^F_{ba}). This would give you something like the implicit nominal spending you'd have done for period b products at period a prices. Second, you divide this deflated nominal GDP by nominal GDP in period a (E_{aa}), and this gives you the ratio of real GDP across the two periods.

National accounts typically will give you the pieces of information to do this deflation method. But note that what we are really trying to do is compare quantities across periods using a common set of prices.

A last topic here concerns making comparisons over long periods of time. We want to look at economic growth over a range of years, not just two. We could imagine using a single base year, and comparing all other years to that one. But then we'd have ratios like $(Y_{1964}/Y_{2009})_F$, and we'd be measuring real GDP in 1964 using 2009 prices. The disconnect of those prices may be substantial, and we haven't even touched yet on the issues that would come up with new products arriving, or old products disappearing.

Instead what happens is that we "chain" together estimates of real GDP ratios for adjacent years. For example, let's say we calculate $(Y_{2008}/Y_{2009})_F = 0.98$, $(Y_{2007}/Y_{2008})_F = 0.97$, $(Y_{2006}/Y_{2007})_F = 0.99$. If we set real GDP in 2009 to 100, then we can get real GDP in 2008 by using the first ratio, meaning real GDP in 2008 is 98. Real GDP in 2007 would then be $98 \times 0.97 = 95.06$, and real GDP in 2006 would be

$95.06 \times 0.99 = 94.11$. Our chained series for real GDP would be 94.11 in 2006, 95.06 in 2007, 98 in 2008, and 100 in 2009.

Chained real GDP calculations are the closest thing we've got in the data to the theoretical idea of growth we're working with. In the theory, we're talking about $d \ln Y$, or the change in real GDP holding prices constant, and chained GDP does something very similar. It uses (relatively) small changes in time, constructs a price index that is an average of the prices across the two time periods in question, and shows us the change in real GDP at that average of the prices.

A.4 National Income Product Account Basics

The principles of how to add up economic activity seem simple, but the actual practice of collecting the data and aggregating it are not trivial. The National Income Product Accounts (NIPA) are the rules and standards used to track economic activity in a consistent manner over time. The NIPA track the *flows* of economic activity, not the *stock* of assets or financial values. These flows are the economic transactions that take place in a given time period. As each transaction has two sides, the NIPA uses T-accounts to match "sources" of funds on the right (e.g. expenditures) with "uses" of the left (e.g. incomes). These two sides have to balance.

The Integrated Macroeconomic Accounts (IMA) combine the NIPA accounts with flow of funds data from the Federal Reserve to reconcile the flows in the NIPA with changes in stocks in the flow of funds.

Account 1. Domestic Income and Product Account

Line			Line		
1	Compensation of employees, paid	8,618.5	15	Personal consumption expenditures (3–3)	11,050.6
2	Wages and salaries	6,938.9	16	Goods	3,739.1
3	Domestic (3–12)	6,924.0	17	Durable goods	1,191.9
4	Rest of the world (5–15)	14.9	18	Nondurable goods	2,547.2
5	Supplements to wages and salaries (3–14)	1,679.6	19	Services	7,311.5
6	Taxes on production and imports (4–15)	1,132.1	20	Gross private domestic investment	2,511.7
7	*Less:* Subsidies (4–8)	58.0	21	Fixed investment (6–2)	2,449.9
8	Net operating surplus	4,131.7	22	Nonresidential	2,007.7
9	Private enterprises (2–19)	4,151.0	23	Structures	448.0
10	Current surplus of government enterprises (4–28)	–19.3	24	Equipment	937.9
11	Consumption of fixed capital (6–14)	2,534.2	25	Intellectual property products	621.7
			26	Residential	442.2
12	**Gross domestic income**	**16,358.5**	27	Change in private inventories (6–4)	61.8
			28	Net exports of goods and services	–565.7
13	Statistical discrepancy (6–20)	–203.3	29	Exports (5–1)	2,198.2
			30	Imports (5–13)	2,763.8
			31	Government consumption expenditures and gross investment (4–1 plus 6–3)	3,158.6
			32	Federal	1,292.5
			33	National defense	817.8
			34	Nondefense	474.7
			35	State and local	1,866.1
14	**Gross domestic product**	**16,155.3**	36	**Gross domestic product**	**16,155.3**

Figure A.3: Domestic income and product account. Account 1 from National Income Product Accounts

Figure A.4 shows the first T-account from NIPA, the domestic income and product account (the example is from 2012, and all numbers are in billions). The left-hand side shows gross domestic income, which should be identical to gross domestic product, although the sources don't quite agree, so there is a discrepancy of about 200 billion. Regardless, what this side of the table shows is the equivalent of the terminology in the notes about wages and operating surplus. Here, there is more detail, and some additional small terms. Compensation of employees (W) is broken down to wages and salaries and supplements (e.g. health insurance payments). Taxes on production and imports include things like tariffs. Net operating surplus is the combination of returns to capital and economic profits ($RK + \Pi$). Consumption of fixed capital is the depreciation of capital, δK. This is on the income side, which remember is really the "use" side of the T-account. Depreciation is one "use" of the payments made by the source side. The right-hand side is that source side, and it shows the standard macroeconomic breakdown of GDP into consumption (C), investment (I), government (G) and net exports (NX).

Account 6. Domestic Capital Account

Line			Line		
1	Gross domestic investment	3,126.1	10	Net saving	327.0
2	Private fixed investment (1–21)	2,449.9	11	Personal saving (3–8)	946.7
3	Government fixed investment (1–31)	614.4	12	Undistributed corporate profits with IVA and CCAdj (2–17)	691.1
4	Change in private inventories (1–27)	61.8	13	Net government saving (4–9)	−1,310.8
5	Capital account transactions (net)	−6.5	14	*Plus:* Consumption of fixed capital (1–11)	2,534.2
6	Transfer payments for catastrophic losses (7–3)	−7.7	15	Private	2,038.0
7	Other capital account transactions (7–4)	1.1	16	Government	496.2
8	Net lending or net borrowing (–), NIPAs (7–5)	−461.7	17	General government	436.1
			18	Government enterprises	60.1
			19	*Equals:* Gross saving	2,861.2
			20	Statistical discrepancy (1–13)	−203.3
9	**Gross domestic investment, capital account transactions (net), and net lending, NIPAs**	**2,657.9**	21	**Gross saving and statistical discrepancy**	**2,657.9**

Figure A.4: Domestic capital account. Account 6 from National Income Product Accounts

The NIPA contain other T-accounts. An example is given in Figure A.4, which shows the domestic capital account. The "use" on the left is the mainly gross domestic investment (I), which you can see here is attributable in large part to private fixed investment (e.g. homes and business equipment). There is a meaningful adjustment for net borrowing of -461 billion, which represents foreign purchases of investment goods in the United States. The "source" side of the table is, to me, a bit of misnomer here. I tend to think of the right-hand side as a description of where the investment spending on the left side went. Most went to "consumption of fixed capital", or to cover depreciation. Net saving is the change in capital stock due to the gross investment done. Simplifying a little, we've got (Net Savings = Gross Domestic Investment - Consumption of Fixed Capital), which in the notation established in the notes, would be $\Delta K_{t+1} = I_t - \delta K_t$.

In the notes and the models we develop, we tend to simplify a lot, excluding things from these tables that are small or immaterial to the larger message. Hence the notes do not account for taxes on production and imports, or the statistical discrepancy, or capital account transactions. But the intention of the models is to be consistent with these T-accounts, in that ultimately we want to match the data obtained from them.

A.5 Input/output accounting and Leontief inverses

The connection of GDP (Y) to gross output (Q) depends on the input/output relationships of all the different units of production. Input/output refers to the fact that some of the output of one unit of production (i) is used as an input by another unit of production (j). As noted, those intermediate transactions are not part of GDP. But they do determine how much gross output is nessecary to produce a given amount of GDP.

A quick example can help illustrate this. Consider an oil drilling company that produces barrels of oil. The barrels are used by a refining company that uses them to produce gasoline. But it is also the case that the oil drilling company needs to buy gasoline in order to run their equipment. In addition, there are consumers who have final demand for gasoline. Let's say that 100 barrels of oil can produce 1,000 gallons of gasoline. And let's say that consumers demand 1,000 gallons of gasoline. But it is also the case that it takes 50 gallons of gasoline to produce 100 barrels of oil. So how much oil production is there? To get 1,000 final gallons, we need 100 barrels produced. But that requires an additional 50 gallons of gas, which requires an additional 5 barrels of oil. But that 5 barrels of oil requires 2.5 gallons of gas, which requires and additional 0.25 barrels of oil. And so on and so on.

We can construct this interaction in a little system of equations.

$$
\begin{aligned}
y_B &= 0.1 \times y_G + 0 \\
y_G &= 0.5 \times y_B + 1,000
\end{aligned}
$$

where y_B are barrels produced, which are equal to one-tenth of the total gallons of gas produced, as that is the demand from the refinery. The zero in the first line represents the fact that there is zero final demand for oil (no one uses oil directly, it is only an intermediate good here). The second line shows that the total gallons of gas produced are equal to one-half the number of barrels produced (that is the demand from oil producers) plus the 1,000 gallons demanded by final consumers. This is just a two-equation, two-unknown situation. It can be solved to find $y_B = 105.26$ is the gross output of the oil drilling company, and $y_G = 1,052.63$ is the gross output of the refinery.

Note that we only get hard numbers because I asserted that final demand for gas was 1,000 gallons.If final demand was some arbitrary c_G, then we could give a generic solution of $y_B = 0.1 \times c_G/0.95$, and $y_G = (0.5 \times 0.1/0.95 + 1) \times c_G$.

This same logic can be extended to the case of J arbitrary units of production. Ultimately we'd have a J equation system, with J unknown gross outputs of $p_j y_j$, taking the final demands $p_j c_j$ of each unit as given. This is tedious, but it is a straightforward linear algebra problem, and can be set up in matrix form. To do this, take the expression for gross output of a unit in equation (**??**) and "stack" these, rather than adding them up. The only difference with the little example of oil and gas is that here we're going to use values (i.e. with relative prices included) rather than raw quantities.

$$\begin{aligned} p_1 y_1 &= p_1 y_{11} + p_1 y_{21} + ... + p_1 y_{J1} + p_1 c_1 \\ p_2 y_2 &= p_2 y_{12} + p_2 y_{22} + ... + p_2 y_{J2} + p_2 c_2 \\ &... \\ p_J y_J &= p_J y_{1J} + p_J y_{2J} + ... + p_J y_{JJ} + p_J c_J \end{aligned}$$

where recall that $p_i y_{ji}$ is the intermediate demand by unit j for output from unit i. Now we can simplify this by building it into matrix form. Define the matrices as follows

Depending on the level of analysis you are working with - establishment, firm, sector - you might be able to assert that $y_{ii} = 0$. But we don't have to assume this. It is okay if we think of units of production purchasing their own output to use as an intermediate good. Think of a refinery buying its own gas to run its own trucks and equipment.

$$\mathbf{q} = \begin{bmatrix} p_1 y_1 \\ p_2 y_2 \\ ... \\ p_J y_J \end{bmatrix} \mathbf{c} = \begin{bmatrix} p_1 c_1 \\ p_2 c_2 \\ ... \\ p_J c_J \end{bmatrix} \mathbf{A} = \begin{bmatrix} \frac{p_1 y_{11}}{p_1 y_1} & ... & \frac{p_1 y_{J1}}{p_J y_J} \\ \frac{p_2 y_{12}}{p_1 y_1} & ... & \frac{p_2 y_{J2}}{p_J y_J} \\ ... & \ddots & ... \\ \frac{p_J y_{1J}}{p_1 y_1} & ... & \frac{p_J y_{JJ}}{p_J y_J} \end{bmatrix} \tag{A.21}$$

The vector q captures the gross output of each unit, the vector c the final demand for each units output. The matrix A is called a "technical coefficients" matrix. It tells us how much of each intermediate good is necessary to produce one (real) dollar of additional output from a given unit of production. It takes $p_2 y_{12}/p_1 y_1$ in purchases of intermediate good 2 to produce one real dollar of output of good 1, for example. A thus tells us the "recipe" for producing gross output.

The raw data for these matrices come from something called a "use table", which are part of input/output accounts produced by statistical agencies. See A.6 for a description of what that use table looks like.

With these matrices, we can simplify our system of equations down to

$$\mathbf{q} = \mathbf{A}\mathbf{q} + \mathbf{c}, \tag{A.22}$$

and solve this using normal matrix operations,

$$(\mathbf{I} - \mathbf{A})\mathbf{q} = \mathbf{c} \tag{A.23}$$

where $\mathbf{I}$ is an identity matrix, and so

$$\mathbf{q} = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{c}. \tag{A.24}$$

This is the matrix equivalent of our earlier example with the barrels and gas. The $(\mathbf{I} - \mathbf{A})^{-1}$ matrix is a "Leontief inverse".[1] It isn't obvious

[1] Wassily W. Leontief. *Input-Output Economics*. Oxford University Press, 1966

here, but this Leontief inverse summarizes all of the direct and indirect effects of final demand for c on gross output q. You could sum up the elements of $\mathbf{q}$ to get gross output Q, or sum up the elements of $\mathbf{c}$ to get GDP, Y.

To see what is going on, go back to the simple example involving barrels and gasoline. Write the matrix form of this as

$$\begin{bmatrix} y_B \\ y_G \end{bmatrix} = \begin{bmatrix} 0 & 0.1 \\ 0.5 & 0 \end{bmatrix} \begin{bmatrix} y_B \\ y_G \end{bmatrix} + \begin{bmatrix} c_B \\ c_G \end{bmatrix} \tag{A.25}$$

where I've allowed both oil, c_B, and gas, c_G, to have some final demand. Our general setting says that we should form the Leontief inverse first, which in this case is

$$(\mathbf{I} - \mathbf{A})^{-1} = \begin{bmatrix} 1.052 & 0.105 \\ 0.526 & 1.052 \end{bmatrix}. \tag{A.26}$$

The entries in this table tell us the total effect of final demand, whatever that may be. For example, the upper right entry (0.105) tells us that every additional gallon of gas demanded induces 0.105 barrels of oil to be produced. We know that technically, only 0.1 barrels are necessary, but 0.105 are produced because oil production requires some gasoline itself. The bottom left entry says that every additional gallon of oil demanded (if anyone wanted oil as a final good) would lead to 0.526 gallons of gas being produced, 0.5 because it takes that much gas to produce a barrel, and an additional 0.026 because that gas requires some oil itself.

Given that Leontief inverse, we know that

$$\begin{bmatrix} y_B \\ y_G \end{bmatrix} = \begin{bmatrix} 1.052 & 0.105 \\ 0.526 & 1.052 \end{bmatrix} \begin{bmatrix} c_B \\ c_G \end{bmatrix}. \tag{A.27}$$

Note that just knowing the matrix $\mathbf{A}$, and thus the Leontief inverse, doesn't tell us how big GDP or gross output will be. It only allows us to solve for the relationship of gross output and GDP. Determining actual GDP or gross output would still depend on things like the supply of factors of production (e.g. labor, capital) and the productivity of individual units of production. Nevertheless, this structure is still useful because it gives us the basis for understanding how to account for the interactions of different units of production, which will be necessary for several different results.

A.6 Use tables

As part of input/output accounting for an economy, something called a "use" table is produced. This contains information on the intermediates used by a given unit of production, along with additional information on value-added and gross output. A sample of such a table is shown in Figure A.6, from a BEA manual on I/O accounting.[2] This particular use table is shown for major industries of the U.S., but it is not necessary that a use table is at the industry level. One could, hypothetically, produce a use table for every single production unit (e.g. establishment or firm) in the economy, it just would get very large.

[2] Karen J. Horowitz and Mark A. Planting. *Concepts and Methods of the U.S. Input-Output Accounts*. U.S. Department of Commerce Bureau of Economic Analysis, 2009

Table 1.2 Use table: Commodities used by industries and final uses

		Industries																Final uses (GDP)							
		Agriculture, forestry, fishing, and hunting	Mining	Utilities	Construction	Manufacturing	Wholesale trade	Retail trade	Transportation and warehousing	Information	Finance, insurance, real estate, rental, and leasing	Professional and business services	Educational services, health care, and social	Arts, entertainment, recreation,	Other services, except government	Government	Total Intermediate	Personal consumption expenditures	Private fixed investment	Change in private inventories	Exports of goods and services	Imports of goods and services	Government consumption expenditures and gross	Total final uses (GDP)	TOTAL COMMODITY
Commodities	Agriculture, forestry, fishing, and hunting																								
Commodities	Mining																								
Commodities	Utilities																								
Commodities	Construction																								
Commodities	Manufacturing																								
Commodities	Wholesale trade																								
Commodities	Retail trade																								
Commodities	Transportation and warehousing																								
Commodities	Information																								
Commodities	Finance, insurance, real estate, rental, and leasing																								
Commodities	Professional and business services																								
Commodities	Educational services, health care, and social assistance																								
Commodities	Arts, entertainment, recreation, accommodation, and food services																								
Commodities	Other services, except government																								
Commodities	Government																								
Commodities	Other																								
Commodities	Scrap, used and secondhand goods																								
	Total Intermediate																								
Value added	Compensation of employees																								
Value added	Taxes on production and imports, less subsidies																								
Value added	Gross operating surplus																								
Value added	Total value added																								
Total industry output																									

Total industry output

Total commodity output

Figure A.5: Sample use table. Source: Horowitz and Planting (2009).

The rows of the table are labeled "commodities", although you'll notice they are identical to the "industries" listed across the top. Technically, a cell in the table reports the total amount expended by the industry (column) on a given commodity (row), but at the industry level the commodities are typically assumed to be identical with a given industry. Referring back to the accounting in A.5, if rows are denoted by i and columns by j, then $p_i y_{ji}$ is what in each cell, or the purchase of commodity i by industry j for use as an intermediate good.

If you read across a row, you can find total intermediate usage of commodity i, or $\sum_{j \in J} p_i y_{ji}$. To the right of the total intermediate usage are categories of final use, or purchases of commodity i by final users for things like consumption, investment, government use, or the like. The total of those final uses sum to the 2nd-to-last column, "Total final use (GDP)". In the notation from A.5, this total final use is $p_i c_i$. Finally, the final column of the table is labelled "Total Commodity", which in our terminology should be equivalent to $p_i y_i = p_i c_i + \sum_{j \in J} p_i y_{ji}$.

If you look down a column, you instead see the purchases by a given industry of the commodities of other industries. The row "Total intermediate" is thus $\sum_{j \in J} p_j y_{ij}$, or the sum spent by industry i on intermediates from other industries (commodities) j. Below that are three components of value-added: compensation of employees (W_i, taxes on production and imports (this is usually small relative to value-added), and gross operating surplus ($R_i K_i + \Pi_i$). The three components by necessity sum to total value-added, or $p_i c_i = W_i + R_i K_i + \Pi_i$. Finally, the last row of the table is "Total industry output", which again should be gross output, $p_i y_i = W_i + (R_i K_i + \Pi_i) + \sum_{j \in J} p_j y_{ij}$.

From our perspective, the gross output reported in the column "Total commodity output" should equal the gross output reported in the row "Total industry output" for each of the industries. In practice, these are not equal in the data. First, you can see that there can be commodities (e.g. scrap and second-hand goods) that are not counted as an industry. Second, the things produced by a given industry do not necessarily slot neatly into commodities. For example, think of a firm that does estate planning. It may be classified under "Finance, insurance, etc.." as an industry, because the BEA classifies firms by their primary product. But the commodities it produces could consist of both financial advice (in the "Finance, insurance, etc.." commodity group) and legal advice (in the "Professional and business services" commodity group). Hence the table need not be entirely symmetric. In practice, we'll pick one "side" to use, which here will be to look at industries, and thus pull our information on value-added and gross output from the final *rows* of the table. The last important note is that the use table provides the raw input to the "technical coefficients" matrix $\mathbf{A}$ from A.5. The table reports values like $p_j y_{ij}$, and we just need to divide these each by gross output of the given industry, $p_j y_j$, to arrive at the technical coefficients.

We could also use the "make" table from the input/output accounts to get the technical coefficients. The value of the use table is that it has the value-added and labor compensation numbers as well.

A.7 I/O weighted labor cost shares

The statements about elasticities ϵ_L and ϵ_K laid out in assumption 2.2 come from work by David Baqaee and Emmanuel Fahri.[3] They establish that under the simple assumptions about production found

[3] David Baqaee and Emmanuel Farhi. A Short Note on Aggregating Productivity. NBER Working Papers 25688, National Bureau of Economic Research, Inc, March 2019

in assumption 2.1 that the values of ϵ_L and ϵ_K are appropriately measured by what I'll refer to as an input/output weighted labor share and capital share, respectively. Given that these add to one, if we can identify just the labor share, we will know both. I'm simply going to show the actual calculation of this labor share. Those authors provide the proof that this is in fact the right way to measure ϵ_L.

To find the input/output weighted labor share, we need to construct something similar to the technical coefficients matrix from A.5. The difference is that we need those coefficients with respect to costs, not with respect to gross output.

Let the total *costs* of any given industry i be given by

$$costs_i = W_i + R_i K_i + \sum_{j \in J} p_j y_{ij} \tag{A.28}$$

which consists of total compensation to labor, W_i, total compensation to capital, $R_i K_i$, and then the sum of total payments to all intermediate good providers.

To incorporate labor and capital into this analysis fully, we're going to treat them as industries themselves. This isn't a statement about their actual operation, but rather an accounting trick so that we can simplify the analysis. The labor "industry" uses no intermediate goods, and it doesn't hire labor or capital itself. It simply provides labor to other industries.

We're going to build a matrix of cost shares for each industry, including our fake industries for labor and capital. This matrix, $\mathbf{\Omega}$, is going to capture the fraction of total costs accounted for by each input used by an industry.

$$\mathbf{\Omega} = \begin{bmatrix} \frac{p_1 y_{11}}{costs_1} & \dots & \frac{p_1 y_{J1}}{costs_J} & 0 & 0 \\ \frac{p_2 y_{12}}{costs_1} & \dots & \frac{p_2 y_{J2}}{costs_2} & 0 & 0 \\ \dots & \ddots & \dots & \dots & \dots \\ \frac{p_J y_{1J}}{costs_1} & \dots & \frac{p_J y_{JJ}}{costs_J} & 0 & 0 \\ \frac{W_1}{costs_1} & \dots & \frac{W_J}{costs_J} & 0 & 0 \\ \frac{R_1 K_1}{costs_1} & \dots & \frac{R_1 K_1}{costs_1} & 0 & 0 \end{bmatrix} \tag{A.29}$$

Each column represents the cost breakdown of a single industry. Take a look at the second row, first column. This says that the fraction of industry 1's total costs that are accounted for by inputs from industry 2 is $p_2 y_{12}/costs_1$. As you go down the rows, these are all the different intermediate inputs, up to $p_J y_{1J}/costs_1$. The final two rows of the matrix give you the fraction of costs accounted for by wages and capital costs for an industry. Thus summing up a column delivers the same summation as in equation (A.28). The final two columns show the cost breakdown for our two fake industries for labor and capital. All the entries in these columns are zero because they don't incur any costs. If we had more of these factors of production (e.g. different types of capital or workers) we'd add columns for each of these factors, and have more zeros.

If you look back at A.5, you'll see that this matrix is similar to the technical coefficients matrix **A** that was used in the Leontief inverse. The difference is that the coefficients here are divided by total costs, whereas in that prior matrix they were divided by gross output (which is equivalent to revenues). To the extent that revenues are not equal to costs, which would be the case if there were economic profits, the two matrices entries will differ. What we have here in $\mathbf{\Omega}$ tells us how much the industry j will spend on inputs from i if we we decided to raise the total costs of industry j by one (real) dollar.

There is a similar feedback loop involved with costs as we saw with gross output in A.5. That is, if there is some final demand for good 1, say, $p_1 c_1$, then this requires us to spend money on inputs to produce good 1, costs. But if producing good 1 requires us to purchase some input from industry 9, say, industry 9 may in turn require some inputs of good 1. And this could go around and around, with a logic similar to what we saw in the oil and gas example.

The last thing we need is a vector of final good use by industry, but in this case scaled by total GDP, which is just $Y = \sum_{j \in J} p_j c_j$. Let's define

$$b = \begin{bmatrix} p_1 c_1 / Y \\ p_2 c_2 / Y \\ \dots \\ p_J c_J / Y \\ 0 \\ 0 \end{bmatrix} \tag{A.30}$$

where the final two rows are the final good use shares for labor and capital, respectively, and those are equal to zero because those two factors are not demanded for final use, only as part of producing other goods.

Baqaee and Farhi (2019) build a matrix algebra statement to account for the feedback of costs in one industry on the others. This is similar in form to what was in A.5 in accounting for gross output and GDP's relationship.

$$\mathbf{\Lambda} = \mathbf{\Omega}\mathbf{\Lambda} + \mathbf{b}, \tag{A.31}$$

where **b** is the vector of final use shares of GDP, $\mathbf{\Omega}$ gives us the coefficients governing the feedback of one industry on another's costs, and **L** is a new vector. It captures what I called the input/output weighted cost shares. In particular, the entry of $\mathbf{\Lambda}$ associated with the labor "industry", call it $\mathbf{\Lambda}_L$, will give us the cost share equal to ϵ_L.

Using basic matrix operations, we have

$$\mathbf{\Lambda} = (\mathbf{I} - \mathbf{\Omega})^{-1}\mathbf{b}. \tag{A.32}$$

The inverse is the cost equivalent of the Leontief inverse. It tells us how much total costs for an industry i go up given an increase in the final output of another industry j.

Now, given all this machinery, can we calculate the value of $\mathbf{\Lambda}_L$ and use it to see what happens to ϵ_L over time? The answer is yes. To do this, we need information on the coefficients in Ω. We can get the raw costs for industry i from purchasing an input from industry j, $p_j y_{ij}$ from a "use" table, see A.6. However, we have a problem with computing the total costs. This problem arises because what the use table reports as "gross operating surplus" is a combination of both capital costs ($R_i K_i$) and economic profits (Π_i). To proceed I'll make some assumptions about what percentage of gross operating surplus is accounted for by $R_i K_i$, and compute Λ from that.

We could try to leverage a different data source on industry-level costs, but this raises issues in mapping industries used in the use table to industries used in the other data sources. They are not always compatible.

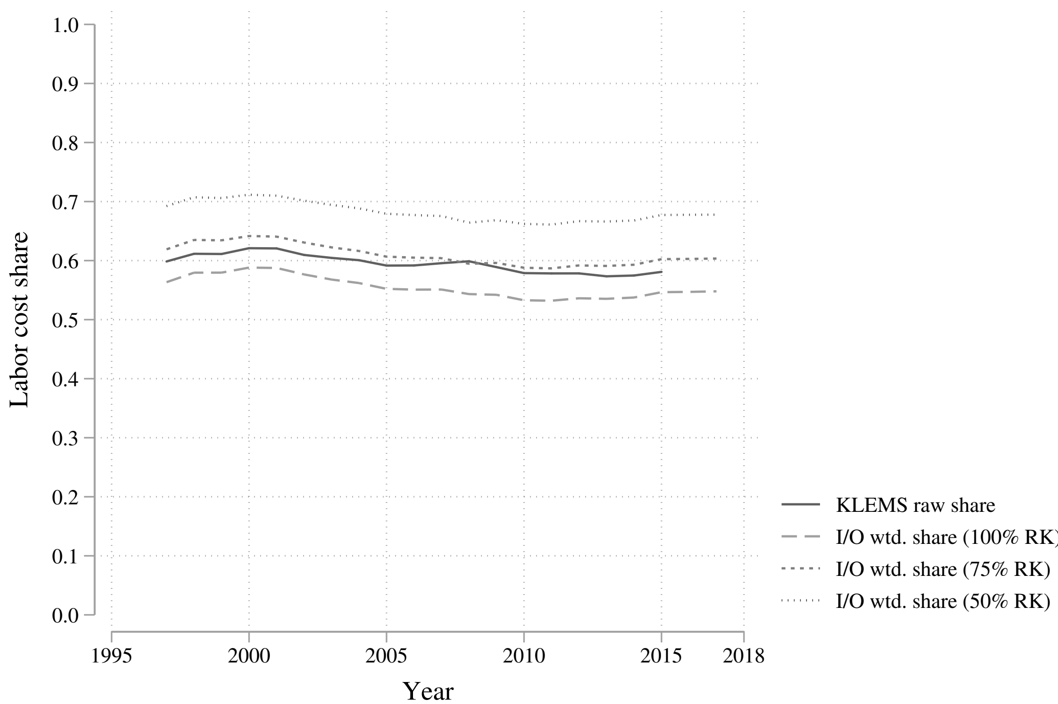

Figure A.6: Input/output weighted labor cost shares over time. Source: authors calculations.

All the data we need is from the use tables at the BEA. For each year, I compute Λ three times: (1) assuming $R_i K_i$ is 50% of gross operating surplus, (2) assuming assuming $R_i K_i$ is 75% of gross operating surplus, and (3) assuming $R_i K_i$ is 100% of gross operating surplus (meaning zero economic profits).

See https://apps.bea.gov/iTable/index_industry_io.cfm and access the Historical Benchmark Input-Output Data. I'm using the Use-SUT-Framework-1997-2017-SUM.xlsx sheet from that source. I use only the 1997-2017 data because it reports labor compensation. There are 71 industries in the use table.

Figure A.7 plots the results of these calculations. The three dashed lines show the values of $\mathbf{\Lambda}_L$ retrieved from $\mathbf{\Lambda}$ under the different assumptions about $R_i K_i$. As can be seen, when $R_i K_i$ is assumed to be 100% of gross operating surplus, and capital payments are largest, we get the lowest values for $\mathbf{\Lambda}_L$. As we assume that capital payments make up smaller and smaller shares of costs, the implied labor cost share rises. Nevertheless, note that under any of the three assumptions, the values of $\mathbf{\Lambda}_L$ are roughly stable over time. For comparison, the dark line plots the naive calculation of the labor share of costs taken from a different source (KLEMS) that was shown in Figure 1.5. All series show the same lack of any distinct trend, although they are not of course perfectly constant.

One concern is that the value of $\mathbf{\Lambda}_L$ shifted substantially over time because the $R_i K_i$ share of gross operating surplus changed over time. There is evidence that payments to capital as a share of GDP fell over time, while the share of GDP going to economic profits rose.[4] At the most extreme, the $R_i K_i$ share of gross operating surplus fell from around 75% in 1995 to around 50% in 2015. This would imply that the true value of $\mathbf{\Lambda}_L$ rose somewhat over time, from about 0.62 to about 0.67. Whether than constitutes a significant trend in the cost share of labor is unclear. Each of the three series has a range of about 0.05 even if we assume capital payments did not change as a share of gross operating surplus.

[4] Simcha Barkai. Declining labor and capital shares. *The Journal of Finance*, 75 (5):2421–2463, 2020

Assumption 2.3 should thus be read as a very rough approximation. The theoretical results that are derived rely on the constancy of $\mathbf{\Lambda}_L$, and hence of ϵ_L. Thus they should be seen as approximations as well, and not rigid laws of nature.

A.8 Final expenditure versus value-added

When we think about separate sectors or industries in the economy, we have to be careful to be consistent in how preferences and production are described.[5] Consider the following table. In this simple setup, there was $900 in spending done in say, a month. What it shown in how that spending was broken down in two different dimensions.

[5] Berthold Herrendorf, Richard Rogerson, and Akos Valentinyi. Two perspectives on preferences and structural transformation. *American Economic Review*, 2013

	Value-Added:			
Expenditure:	Farm	Factory	Shop	Exp Total
Food	40	20	100	160
Phone	0	150	150	300
Vacation	30	10	400	440
VA Total	70	180	650	900

The rows represent *expenditure* categories, which may conform closely to how consumers would characterize their purchases. For this month, they purchased some food, a phone, and a vacation. From the perspective of the consumer, they spent $160 on food, $300 on a phone, and $440 on the vacation. On the other hand, the columns represent *value-added* categories, which you might consider more as the "firm" view of spending. There is a farm, a factory, and a shop(s). From the farms perspective, it sold $180 of goods, the factory sold $180, and the shop sold $650.

Entries in the middle of the table show us how the spending from either perspective broke down across the other. Take a look from the expenditure side first. The consumer may have purchased each of those three items at a shops (or perhaps multiple shops). The *shop* knows that of the $160 spent on food, $20 goes to the factory that packaged the food, and $40 to the farm that grew it in the first place. From the shops perspective, it added $100 in value to those other products. From the consumers perspective, though, it probably only knows that it spent $160 on food. The input/output relationships that went into providing that food in the shop are likely not relevant to the consumer.

Similarly, the expenditure categories that the consumer is basing their utility off of are not relevant to the producers. The fact that the vacation costs $440 and the phone only $300 is very important to the consumer, but those costs don't really matter to the shop who sold them, who cares only about the $400 they make on the vacation, and the $150 on the phone.

When we try to model an economy with multiple industries or sectors, we have to be concious of this disconnect. That is, if we talk about industries that produce value-added (e.g. agriculture, manufacturing, and services) then we also have to talk about preferences over those *same* value-added industries. That is, we'd have to describe a utility function where people care about how much *farm* production they want, how much *factory* production they want, and how much *shop* production they want. On the other hand, if we wanted to have preferences defined over food, phones, and vacations, then we'd have to specify production functions which tell us how much *food* is produced, how many *phones* are produced, and how many *vacations* are produced.

It could be difficult to be consistent. For example, farms know how much they produced, but do they know how much of it went to supermarkets to be purchased as "food" versus how much was purchased to be used as meals on vacation? Probably not.

The problem comes because data sources have different assumptions. Most of the production-side data we have classifies firms into industries based on their value-added. So the BEA might report the value-added output of farms, factories, and shops separately. But at the same time, from the consumption side the BEA will report spending on food, phones, and vacations for individuals. The two data sources are telling you about different things. Translating one into the terms of the other requires an input/output matrix.

A.9 Cost minimization, markups, and cost shares

You can gain a lot of intuition by thinking about a simple, single firm problem. Let this firm have profits of π_i given by

$$\pi_i = p_i y_i - p_{mi} M_i - R_i K_i - w_i L_i,$$

where p_i is the price firm i can charge for their output, and we're not going to make any assumption about how this is set. y_i is gross output. M_i is the amount of intermediate good used (there could be more than one, but this will do the job for us here), and it costs p_m per unit. R_i is the rental cost of capital, and K_i is capital used. w_i is the wage the firm faces, and L_i is the labor they hire. Note that all the input costs are denoted as being specific to i, which is just meant to say that we're not assuming anything about how those wages and rental costs compare with other firms. But the firm is taking those as given.

Do a little re-arranging of this, and we have

$$1 = \frac{\pi_i}{p_i y_i} + \frac{p_{mi} M_i}{p_i y_i} + \frac{R_i K_i}{p_i y_i} + \frac{w_i L_i}{p_i y_i} \tag{A.33}$$

which just says that the shares of revenue accounted for by profits and input costs have to add up to one. To simplify notation, let s_π be the share accounted for profits, and similar notation for each input term, so that

$$1 = s_\pi + s_M + s_K + s_L \tag{A.34}$$

where I suppressed the i notation to keep things looking clean.

Further, the firm has some production function that dictates how output is related to inputs

$$y_i = F(M_i, K_i, L_i).$$

Look at the cost minimization problem for this firm. That is, given some target output $\overline{y}_i$, how will the firm choose to hire intermediates, capital, and labor to produce this? That is a constrained optimization problem, and we can form a Lagrangian to solve it

$$\mathcal{L} = p_{mi} M_i + R_i K_i + w_i L_i + \lambda_i(\overline{y}_i - F(M_i, K_i, L_i)).$$

The multiplier λ_i tells us how much the value function - costs - changes in response to a change in the constraint - output. In other words, λ_i tells us the marginal cost of production to firm i.

The first order conditions with respect to the three inputs are:

$$\begin{aligned} p_{mi} &= \lambda_i F_M \\ w_i &= \lambda_i F_L \\ R_i &= \lambda_i F_K \end{aligned}$$

It is worth spending a moment thinking about the intuition behind these conditions. Think of a firm who decides they want to produce one more unit of output. They could spend w_i/F_L on extra labor to do this, because $1/F_L$ is the amount of labor they'd need, and labor costs w_i. Or they could spend R_i/F_K on capital, or p_{mi}/F_M on intermediates. Obviously, they'd pick the lowest of those three. But if one of them (capital, say) was lower than another (labor, say), then they should not just rent more capital, they should actively fire some workers, and use that additional money to rent more capital, as it is a cheaper way to produce. As they are trying to cost minimize, they'll set the amounts of capital and labor and intermediates such that those three ratios equal one another, $w_i/F_L = R_i/F_K = p_{mi}/M_i$. That ratio is the cost of producing a unit of output, or the marginal cost, or $\lambda_i = w_i/F_L$ (and equivalently for the other inputs). Those are our first-order conditions.

Go back to equation (A.33) and plug these in for the input prices, and we have

$$1 = \frac{\pi_i}{p_i y_i} + \frac{\lambda_i}{p_i}\frac{F_M M_i}{y_i} + \frac{\lambda_i}{p_i}\frac{F_K K_i}{y_i} + \frac{\lambda_i}{p_i}\frac{F_L L_i}{y_i}. \tag{A.35}$$

To go forward let's keep introducing notation. First,

$$\mu_i \equiv \frac{p_i}{\lambda_i}$$

is the *markup* of the firms price over its marginal cost. This markup will indicate that the firm is earning economic profits. It could come from market power that the firm has because it is a monopoly, or in an oligopoly, or engaged in Bertrand competition, etc. It doesn't matter for us here *why* the firm has market power. All we are doing is establishing some relationships taking the ability of a firm to charge a markup as given. Second,

$$\epsilon_x \equiv \frac{F_x x_i}{y_i}$$

is the elasticity of production with respect to input x.

Take the definitions of markups and elasticites back and plug them into equation (A.35) and you get

$$1 = s_\pi + \frac{\epsilon_M}{\mu_i} + \frac{\epsilon_K}{\mu_i} + \frac{\epsilon_L}{\mu_i}$$

which can be re-arranged to

$$\epsilon_M + \epsilon_K + \epsilon_L = \mu_i(1 - s_\pi).$$

This gives us a very useful relationship for thinking about how the properties of production functions (the left-hand side) are related to economic profits (the right-hand side).

The left side is the sum of elasticities, which tells us about the returns to scale of production. If these sum to one, then production is constant returns, while if they sum to more than one it is increasing returns, and if they sum to less than one it has decreasing returns. This sum pins down the right-hand side, which says that the markup of prices over marginal costs is positively associated with the share of revenues that go to economic profits (which isn't a surprise).

You can confirm the relationship of the elasticities and returns to scale by asking what percent output grows if you scale up all inputs by some percentage X%.

Consider a few possible situations:

(a) Production is constant returns ($\epsilon_M + \epsilon_K + \epsilon_l = 1$). If $s_\pi > 0$, then it must be because the firm has a markup $\mu_i > 1$. In other words, with constant returns the only way a firm can earn profits is by charging a price over marginal cost. That seems trivial, but is not obvious given other production situations.

(b) Production is increasing returns ($\epsilon_M + \epsilon_K + \epsilon_l > 1$). In this case, the firm could have a markup $\mu_i > 1$ but still have $s_\pi = 0$. With increasing returns, the firm has to charge more than marginal cost in order to pay their inputs, and yet has nothing left over for profits.

The relationship of scale to markups and the profit share involves no assumption about how firms set the markup, or what kind of market they are in. It is a property that arises solely from cost minimization, and so it holds no matter what you assume about how firms set prices.

Let's go back to the firm problem cost minimization problem. Define their total costs as

$$costs_i \equiv p_{mi}M_i + R_iK_i + w_iL_i.$$

Now, call the fraction of costs accounted for by labor ϕ_L, so that

$$\begin{aligned}\phi_L &= \frac{w_i L_i}{costs_i} \\ &= \frac{\lambda_i F_L L_i}{\lambda_i F_M M_i + \lambda_i F_K K_i + \lambda_i F_L L_i} \\ &= \frac{F_L L_i / y_i}{F_M M_i / y_i + F_K K_i / y_i + F_L L_i / y_i} \\ &= \frac{\epsilon_L}{\epsilon_M + \epsilon_K + \epsilon_L}.\end{aligned}$$

The share of costs accounted for by labor is equal to its elasticity relative to the total elasticities (which recall capture the returns to scale). If there are constant returns to scale, then it follows that $\phi_L = \epsilon_L$, or the cost share of labor is equal to the elasticity of production with respect to labor. All the findings here for labor apply to the other inputs as well.

We can utilize observable data on costs to infer unobservable properties of the production function. Note that we cannot use the cost data by itself to know if there are constant returns to scale or not. But if we are willing to assume constant returns and cost minimization, then we can use ϕ_L data to find ϵ_L. This holds for a given firm, but implicitly is the same idea being used to leverage aggregate data on labor costs to infer the elasticity of aggregate output with respect to aggregate labor.

See A.7.

One other simple relationship that comes out of the various definitions developed so far is

$$\phi_x = \mu s_x. \tag{A.36}$$

for each input x. This says that the share of costs accounted for by input x is equal to the markup times the shrae of revenues accounted for by input x. This should be straightforward to understand. The share of input x in revenues is going to be lower than the share in costs, because of the presence of profits. The markup tells us how important those profits are. Note that if $\mu = 1$ and firms charge prices equal to marginal costs, then $\phi_x = s_x$. This highlights an empirical issue that I trace back to a paper by Hall (1990).[6] We may be able to observe the revenue share s_x for an input, but not the cost share ϕ_x (perhaps because we don't have full information on costs). Under only very restrictive conditions (competition that ensures $\mu = 1$) would s_x be useful in inferring ϕ_x and hence ϵ_x.

[6] Robert E. Hall. Invariance properties of solow's productivity residual. In Peter Diamond, editor, *Growth, Productivity, Employment*. MIT Press, Cambridge, MA, July 1990

The last thing to go over with here is the relationship of gross output and value-added for the firm. This is something that matters for any model trying to aggregate across different units of production, but I tend to rely on Basu and Fernald (2002) for an explanation of their relationship.[7] Value-added is gross output minus intermediates,

[7] Susanto Basu and John Fernald. Aggregate productivity and aggregate technology. *European Economic Review*, 46:963–991, 2002

so

$$va_i = p_i y_i - p_{mi} M_i = p_i y_i (1 - s_M) \tag{A.37}$$

and it can be decomposed into.

$$va_i = \pi_i + R_i K_i + w_i L_i \tag{A.38}$$

We can talk about shares of value-added associated with inputs like labor or capital. These would be

$$s_L^V = \frac{w_i L_i}{va_i} = \frac{s_L}{1 - s_M} \tag{A.39}$$

for labor, and a similar expression for capital. We could also work with costs so that

$$\phi_L^V = \frac{w_i L_i}{w_i L_i + R_i K_i} = \frac{\phi_L}{1 - \phi_M}. \tag{A.40}$$

The importance of these relationships is that at an aggregate level we are interested in GDP, which is a sum of value-added. Hence we might have aggregate data that is analagous to value-added, and aggregate data on wages and/or capital payments that could give us measures of ϕ_L^V. But those ϕ_L^V shares cannot necessarily be translated into ϵ_L, because that is equal to ϕ_L. So at the aggregate level we need some way to adjust for the use of intermediates, meaning we need information on ϕ_M. The input/output adjustments made in A.7 are intuitively doing this for us.

A.10 Multiple capital types

A.11 Types of productivity and the Uzawa Theorem

When building the Solow model we assumed that g_A entered into equation (2.9) multiplied through by ϵ_L. Let me now show you why that was an acceptable assumption to make, and that despite this the theory allows for multiple types of productivity growth.

To start, note that our original production setting only allowed for capital and labor to affect output through their raw numbers (K and L). But what may matter to firms is the *effective* labor or capital that they use. Take labor, for example. The actual number of employees (L) may be less relevant to a firm than how much brainpower or skill their employees can bring to bear.

Let's define different kinds of productivity growth that we could allow for.

There would be as many types of productivity growth as we had factors of production, plus disembodied productivity. If we allowed for resources, for example, then there would be "resource-augmenting" productivity growth as well.

Definition A.1 (Types of productivity growth) *Productivity growth can be classified as*

- *"Labor-augmenting" (denoted A_L), so that $A_L L$ is the total amount of effective labor used in production, and g_{A_L} is the growth rate of labor-augmenting productivity.*
- *"Capital-augmenting" (denoted A_K), so that $A_K K$ is the total amount of effective capital used in production, and g_{A_K} is the growth rate of capital-augmenting productivity.*
- *"Disembodied" (denoted A_Y), which works by augmenting the overall production of GDP, and g_{A_Y} is the growth rate of disembodied productivity.*

Traditionally, g_{A_Y} is referred to as "Hicks-neutral" productivity growth, g_{A_L} as "Harrod-neutral" productivity growth, and g_{A_K} as "Solow-neutral" productivity growth.

A_L is meant to capture the effectiveness of an individual worker, whether that comes about because of education or native skill or some combination of the two. It encompasses human capital, which we can analyze more systematically later in the notes. Similarly, for capital A_K captures how effective each unit of capital is, which might depend on its vintage (e.g. older capital may be less useful than newer capital) or other factors. Disembodied productivity affects output regardless of the level of inputs used. An example could be a more efficient inventory management system. Using the exact same workers and exact same capital, a firm might be able to increase their output because they can fill orders faster.

We can go back and write the growth rate of output given these definitions as

$$g_Y = \epsilon_K \left(g_K + g_{A_K}\right) + \epsilon_L \left(g_L + g_{A_L}\right) + g_{A_Y}. \tag{A.41}$$

The basic relationship is the same as before. Note that because capital-augmenting productivity growth (g_{A_K}) affects the effective amount of capital used, it enters multiplied by ϵ_K. Similarly labor-augmenting productivity growth, g_{A_L}, enters multiplied by ϵ_L. Disembodied productivity growth enters with an implicit elasticity of one.

Define g_A as follows:

$$g_A \equiv \frac{\epsilon_K}{\epsilon_L} g_{A_K} + g_{A_L} + \frac{1}{\epsilon_L} g_{A_Y} \tag{A.42}$$

and it is straightforward to show that equation (A.41) can be written as

$$g_Y = \epsilon_K g_K + \epsilon_L g_L + \epsilon_L g_A, \tag{A.43}$$

which is identical to what we started with in equation (2.9) from the main Solow model. Mathematically all we are doing is multiplying and dividing by ϵ_L. Given that the elasticities are constant, given assumption 2.2, we can write *all* types of productivity growth in terms of labor-augmenting productivity growth (e.g. multiplied by ϵ_L). All this is really doing is saving us some notation. Defining g_A

this was does not say that labor-augmenting productivity growth is more important than other types.

While in our case all the types of technological change are interchangable, that comes about because we've used the data to assert that the elasticities are constant at all times (on or off a balanced growth path). The Uzawa Theorem[8] is a more generic statement that if a model has a balanced growth path with a constant g_y and constant elasticities, then it must be possible to write the production function the way we did in equation (2.9), meaning that all productivity growth can be expressed as labor-augmenting. Uzawa's original statment of this Theorem is not clear, and more recently an accessible proof of it is available.[9]

[8] H. Uzawa. Neutral Inventions and the Stability of Growth Equilibrium. *The Review of Economic Studies*, 28(2):117–124, 02 1961

[9] Charles I. Jones and Dean Scrimgeour. A new proof of Uzawa's steady-state growth theorem. *The Review of Economics and Statistics*, 90(1):180–182, 2008

I think there is a common misunderstanding of the Uzawa Theorem that it says productivity growth can *only* be labor-augmenting, meaning that it must be that $g_{A_K} = g_{A_Y} = 0$. But the Uzawa Theorem is about the BGP, and doesn't say anything about how productivity growth needs to be expressed when an economy is not on a BGP. And to the extent that the elasticities like ϵ_L are unchanged over time, any combination of productivity growth is plausible. There are recent papers exploring whether the Uzawa Theorem needs to hold for balanced growth to exist.[10]

[10] Gene M. Grossman, Elhanan Helpman, Ezra Oberfield, and Thomas Sampson. Balanced growth despite Uzawa. *American Economic Review*, 107 (4):1293–1312, April 2017

A.12 Harrod-Domar Model

A.13 Traditional Solow diagram

A.14 Evidence on Convergence

A.15 An exact solution for the Solow model

From equation (3.4) we know how to relate $g_{K/Y}$ and K/Y. To save notation in this section, let $Z = K/Y$, and so $g_z = g_{K/Y}$. Rewrite the equation as

$$g_z = \epsilon_L(s_I/z - \delta - g_L - g_A)$$

and note that $g_z = (dz/dt)/z$, so we can write this as

$$\frac{dz}{dt} = \epsilon_L s_I - \epsilon_L(\delta + g_A + g_L)z,$$

which is a simple, linear, differential equation. This can be solve for the exact time path of $z(t)$ using standard techniques.

First, create the integrating factor $\mu(t) = e^{\epsilon_L(\delta+g_A+g_L)t}$, and multiply both sides of the expression for dz/dt by this

$$e^{\epsilon_L(\delta+g_A+g_L)t}\frac{dz}{dt} + e^{\epsilon_L(\delta+g_A+g_L)t}\epsilon_L(\delta + g_A + g_L)z = e^{\epsilon_L(\delta+g_A+g_L)t}\epsilon_L s_I.$$

Now, note that the entire left-hand side of this expression is simply an application of the product rule for derivatives, so we can write

$$\frac{\partial\left[ze^{\epsilon_L(\delta+g_A+g_L)t}\right]}{\partial t} = e^{\epsilon_L(\delta+g_A+g_L)t}\epsilon_L s_I.$$

Integrate both sides with respect to dt,

$$\int \frac{\partial\left[ze^{\epsilon_L(\delta+g_A+g_L)t}\right]}{\partial t}dt = \int e^{\epsilon_L(\delta+g_A+g_L)t}\epsilon_L s_I dt$$

and you get

$$ze^{\epsilon_L(\delta+g_A+g_L)t} = \frac{s_I}{\delta+g_A+g_L}e^{\epsilon_L(\delta+g_A+g_L)t} + C$$

where C is some constant of integration. Rearrange this to be

$$z(t) = \frac{s_I}{\delta+g_A+g_L} + Ce^{-\epsilon_L(\delta+g_A+g_L)t}.$$

What is the constant C? We can figure this out if we look at what happens at $t = 0$. In that case

$$z(0) = \frac{s_I}{\delta+g_A+g_L} + C$$

or

$$C = z(0) - \frac{s_I}{\delta+g_A+g_L}.$$

Plug that back into the equation for $z(t)$ and we get

$$z(t) = \frac{s_I}{\delta+g_A+g_L}\left(1 - e^{-\epsilon_L(\delta+g_A+g_L)t}\right) + z(0)e^{-\epsilon_L(\delta+g_A+g_L)t}.$$

And finally, at last, let's remember that $z(t) = K(t)/Y(t)$, so that we have

$$\frac{K(t)}{Y(t)} = \frac{s_I}{\delta+g_A+g_L}\left(1 - e^{-\epsilon_L(\delta+g_A+g_L)t}\right) + \frac{K(0)}{Y(0)}e^{-\epsilon_L(\delta+g_A+g_L)t}.$$

What does this tell us? Note that the capital/output ratio at any given moment in time is a mixture of two quantities. The weights $1 - e^{-\epsilon_L(\delta+g_A+g_L)t}$ and $e^{-\epsilon_L(\delta+g_A+g_L)t}$ add up to one. As t grows the first weight goes towards one, and the second weight goes towards zero. What are these two weights acting on? The first fraction is $s_I/(\delta+g_A+g_L)$, which we know is $(K/Y)^{BGP}$. The second fraction is $K(0)/Y(0)$, or the initial capital/output ratio. So the value of $K(t)/Y(t)$ is some weighted average of the balanced growth path capital/output ratio and the inital capital/output ratio. That means $K(t)/Y(t)$ is always between these two, which makes sense because the capital/output ratio is transitioning from one to the other.

How fast does that transition go? It depends on how sensitive the weights are to time. Therefore the speed on transition depends on the parameters $\epsilon_L(\delta + g_A + g_L)$. Note that the investment share s_I does *not* show up in the transition speed. It depends on the elasticity with respect to labor, which also tells us the elasticity with respect to capital. The more elastic production is with respect to capital (and thus the less elastic with labor), the *slower* is the transition. When output is sensitive to capital, then as capital accumulates, output goes up a lot, which means more capital is accumulated, and the K/Y ratio does not change much. When the elasticity with respect to capital is low, when capital accumulates this doesn't change Y much, so the K/Y ratio changes a lot.

The depreciation rate and growth rates of productivity and labor influence the transition speed as well. The higher these are, the faster capital is depreciating and the faster output is growing outside of the effects of capital, so the K/Y ratio adjusts quickly. The logic follows in reverse without too much trouble.

This isn't a pleasant equation to look at or work with, given that the actual path of K/Y through time is non-linear. Nevertheless, it gives us some notion of what drives transitory growth. Note that this conforms to the idea that the bigger the gap between the intial capital/output ratio and the BGP capital/output ratio, the larger the change in the capital/output ratio.

If we try to embed this is the expression for output per capita from equation (3.7), things do not get any better. That tell us

$$\ln y(t) = \frac{\epsilon_K}{\epsilon_L} \ln\left(\frac{s_I}{\delta + g_A + g_L}\left(1 - e^{-\epsilon_L(\delta+g_A+g_L)t}\right) + \frac{K(0)}{Y(0)} e^{-\epsilon_L(\delta+g_A+g_L)t}\right) + \ln A(0) + g_A t.$$

Technically, this is the exact path of output per capita over time in the Solow model. If an economy is on a BGP and $K(0)/Y(0) = (K/Y)^{BGP}$, then you can confirm that this delivers the very simple answer that

$$\ln y(t)^{BGP} = \frac{\epsilon_K}{\epsilon_L} \ln\left(\frac{s_I}{g_A + g_L + \delta}\right) + \ln A(0) + g_A t$$

we found before.

A.16 Continuous-time infinite consumption problem

Assuming that time is continuous is useful at times for analysis. Again we're assuming the individual has full knowledge of the path of interest rates and income that they face, and can borrow/lend at that interest rate at will.

The full dynamic problem is now to mazimize

$$\max_c \int_0^\infty e^{-\theta t} U(c)\, dt \tag{A.44}$$

subject to the constraint that

$$\dot{a} = ra + w - c \tag{A.45}$$

where a is now the instantaneous level of assets, w is the instantaneous wage rate (or any exogenous income process), c is the instantaneous level of consumption, and r is the constant rate of interest. In the utility function, θ represents the discount rate, which exponentially lowers the utility of consumption as time passes.

The level of a is the state variable, meaning that it does not jump around, while c is the control variable, meaning that it can. For all variables, I've dropped the time subscript for convenience.

We want to eliminate the possibility of unlimited borrowing and infinite consumption, so we have the present value budget constraint of

$$a_0 + \int_0^\infty e^{-rt} w dt = \int_0^\infty e^{-rt} c dt. \tag{A.46}$$

You could set up a Lagrangian, but then we'd need an infinite number of multipliers. We'll set up a Hamiltonian, which involves a single multiplier, μ, that is called a *co-state* variable, evolves over time on its own, and implicitly captures the time path of the infinite multipliers. It is the "shadow value" of having assets, meaning it gives you the value of the change in assets at any given moment in time converted into utility terms.

$$H = \max_c \left\{ e^{-\theta t} U(c) + \mu (ra + w - c) \right\} \tag{A.47}$$

Now, to solve this problem you need to apply several conditions to the Hamiltonian. This method is called optimal control theory, and we will use this theory without explicitly discussing the origins of it. First, maximize H with respect to c, as written

$$U'(c) e^{-\theta t} - \mu = 0 \tag{A.48}$$

and this gives you something that looks like the FOC from a simpler consumption problem. It's telling us that we have to balance out the marginal gain in utility from consumption against the marginal cost, which is captured by $-\mu$, because a small increase in consumption means a small drop in assets.

Recover the constraint by taking $\partial H / \partial \mu = \dot{a}$ or

$$\dot{a} = ra + w - c \tag{A.49}$$

which is just ensuring that we meet the constraint on how assets accumulate. The rate of change of the shadow value of assets, $\dot{\mu} = -\partial H / \partial a$ or

$$\dot{\mu} = -\mu r \tag{A.50}$$

and this describes how the shadow value of assets evolves over time. Finally, we have a transversality condition, which keeps the problem from "blowing up" and having infinite consumption later in life. This condition is that

$$\lim_{t\to\infty} \mu a = 0 \tag{A.51}$$

You solve (A.48), (A.49), and (A.50) together in order to find the solution. First, take (A.48) and find the derivative with respect to time

$$\dot{c}U''(c)\,e^{-\theta t} - \theta U'(c)\,e^{-\theta t} - \dot{\mu} = 0$$

which gives us another expression for $\dot{\mu}$. Plug (A.50) into the above equation to get

$$\dot{c}U''(c)\,e^{-\theta t} - \theta U'(c)\,e^{-\theta t} = -\mu r$$

Now notice from (A.48) that $\mu = U'(c)\,e^{-\theta t}$ and plug that in to get

$$\dot{c}U''(c)\,e^{-\theta t} - \theta U'(c)\,e^{-\theta t} = -U'(c)\,e^{-\theta t} r$$

Start going crazy with the algebra and you can get the following statement

$$\frac{\dot{c}}{c} = (r-\theta)\,\frac{U'(c)}{cU''(c)} \tag{A.52}$$

which describes the growth of consumption over time. Assume for the moment that the $\frac{U'(c)}{cU''(c)}$ term is constant. Then whether consumption is growing or falling depends on the relative size of r and θ, or exactly what we saw in a simpler model. If r is larger, then consumption is rising as people save their incomes, and if θ is larger then consumption is falling as people discount the future a lot.

The second term on the right hand side of (A.52) should be familiar. It's just the intertemporal elasticity of substitution. What it says is that your consumption growth will be slower if your willingness to substitute between is lower.

In the CRRA case, we know exactly what the intertemporal elasticity is, $1/\sigma$. So that means that if preferences are CRRA, the optimal consumption growth is

$$\frac{\dot{c}}{c} = (r-\theta)\,\frac{1}{\sigma}$$

Again, solving explicitly for the level of consumption is possible. Equation (A.52) is a first order differential equation with a simple form and has the solution that

$$c_t = c_0 e^{\frac{1}{\sigma}(r-\theta)t}$$

which gives us a nice way to describe consumption in any period. Now we need the budget constraint, which was

$$a_0 + \int_0^\infty e^{-rt} w dt = \int_0^\infty e^{-rt} c dt$$

and we can plug in our formula for c_t, play with some algebra and get

$$a_0 + \int_0^\infty e^{-rt} w dt = c_0 \int_0^\infty e^{\frac{1}{\sigma}(r(1-\sigma)-\theta)t} dt$$

The integral on the right hand side can be evaluated to be a positive, finite number if

$$(1-\sigma) r < \theta$$

which is just a condition that the return on savings is not so high that it would make sense to accumulate infinite assets. Now, evaluating the integral and rearranging we get that

$$c_0 = \frac{1}{\sigma} (\theta - (1-\sigma) r) \left[a_0 + \int_0^\infty e^{-rt} w dt \right].$$

This form of the consumption function yields essentially identical conclusions to a discrete time model. At times the continuous time model proves more convenient to use in further analysis, and at others the discrete model works. It doesn't really matter which form of the problem we use, the logic is the same.

A.17 Constant Relative Risk Aversion

Conclusion 4.1 established that the only way for the growth rate of consumpton to be stable was if the inter-temporal elasticity of substitution (IES) was constant. This led to conclusion 4.2 which said that this implied utility had a specific form, known as "Constant Relative Risk Aversion".

To see how this follows, consider the definition of the IES and set it equal to the constant $1/\sigma$,

$$\frac{-U'(c)}{U''(c)c} = \frac{1}{\sigma}.$$

From the definition of the IES in definition 4.1, we can write this as

$$\frac{-d \ln c}{d \ln U'(c)} = \frac{1}{\sigma}.$$

Re-arrange this to

$$-\sigma d \ln c = d \ln U'(c)$$

and then integrate both sides to find

$$-\sigma \ln c = \ln U'(c).$$

Exponentiate both sides to get

$$c^{-\sigma} = U'(c),$$

which says that the marginal utility of consumption is a power function of consumption itself. In more practical terms, the elasticity of the marginal utility with respect to c is constant, or it falls by σ percent whenever c rises by 1%, no matter how big c gets.

Finally, integrate both sides of this again and you have

$$\frac{c^{1-\sigma}}{1-\sigma} = U(c) \tag{A.53}$$

which is the definition of the BGP preferences given in conclusion 4.2. This is the only form of a utility function that delivers a constant IES no matter how large c gets.

The name given to these preferences, "Constant Relative Risk Aversion", comes from the fact that they also imply risk aversion is constant with respect to c. The definition of the coefficient of risk aversion is

$$\frac{-U''(c)c}{U'(c)} \tag{A.54}$$

and it measures the curvature of the utility function, or how much marginal utility changes when c changes. As you can see, it is simply the reciprocal of the IES (which meaures how much c changes when marginal utility changes). Hence if the IES is constant and equal to $1/\sigma$, then the coefficient of risk aversion must be equal to σ.

A.18 General time series processes

More general time series processes extend the notion of auto-regressive or moving-average models, or potentially combine them. The intuitions are similar but require more mathematical baggage to establish similar points.

Definition A.2 (Autoregressive process) *An auto-regressive process for x_t of order k, referred to as an AR(k) process, is written as*

$$x_t = \mu + \rho_1 x_{t-1} + \rho_2 x_{t-2} + ... + \rho_k x_{t-k} + \varepsilon_t$$

where $\varepsilon_t \tilde{N}(0, \sigma_\varepsilon^2)$ and ε_t is indepdent over time.

We still want to know if this process is stable, in the sense of having a finite variance and not "blowing up" over time. Intuitively we want the combined effects coming through the ρ_j terms to be "less than one", except it is not so simple as to just add them up or anything like that. We will introduce some notation and operators to do this.

Definition A.3 (Lag operator) *The operator L lags the variable x_t by one period*

$$Lx_t = x_{t-1}, \tag{A.55}$$

and the operator L^k lags the variable x_t by k periods

$$L^k x_t = x_{t-k}. \quad (A.56)$$

Note that in an AR(k) process we are still just operating on the same variable, x_t, at different lags, so the lag operators allow us to write a simplified expression.

Definition A.4 (Lag polynomial) *A lag polynomial, $a(L)$, is a series of lag operators*

$$a(L) = a_0 + a_1 L + a_2 L^2 + ... + a_p L^p \quad (A.57)$$

where a_j are coefficients, and when applied to series x_t is

$$a(L)x_t = a_0 x_t + a_1 x_{t-1} + a_2 x_{t-2} + ... + a_p x_{t-p}. \quad (A.58)$$

The important aspect of a lag polynomial is the set of coefficients, the a terms. You can think of every lag polynomial as having infinite lag terms, L^∞. What distinguishes lag polynomials is the set of weights, so a lag polynomial that only is concerned with one lag has $a_0 \neq 0$, $a_1 \neq 0$, and then all $a_j = 0$ for all $j \geq 2$.

Definition A.5 (Lag polynomial autoregressive process) *An auto-regressive process for x_t of order k, referred to as an AR(k) process, can be written as*

$$a(L)x_t = \mu + \varepsilon_t$$

where $a(L)$ is the lag polynomial $a(L) = 1 - \rho_1 L - \rho_2 L^2 - .... - \rho_k L^k$.

All we've done is come up with a short-hand for writing down AR processes. But now that we have this shorthand we can establish some properties.

Conclusion A.1 (Stability of AR process) *An auto-regressive process x_t of the form $a(L)x_t = \mu + \varepsilon_t$ is stable if all roots of the lag polynomial $a(L)$ are greater than one in absolute value.*

The conclusion is the generalized version of the idea that in an AR(1) we need $|\rho| < 1$ for stability. But note that this is slightly different in saying that the *roots* of the polynomial have to be *bigger* than one. Let's establish what those roots are, and then we'll be able to see that this is really saying the same thing.

You'll often see this condition stated as "all roots lie outside the unit circle" which just confuses the issue.

Definition A.6 (Roots of lag polynomials) *Given a lag polynomial $a(L)$ this can be factored as*

$$a(L) = (1 - \alpha_1 L)(1 - \alpha_2 L)....(1 - \alpha_k L) \quad (A.59)$$

and the roots of this polynomial are $1/\alpha_1$, $1/\alpha_2$, ..., $1/\alpha_k$.

In general the values of α_j are going to be polynomial combinations of the various ρ_j terms from the original lag polynomial itself. The conclusion in A.1 is saying that each of those $1/\alpha_j$ terms has to be bigger than one if the process is going to be stable, which means we're really putting some restriction on the value of the ρ_j terms. Let's look at how this lag structure works in the sense of an AR(1) to confirm it works in the way we think.

An AR(1) can be written as $a(L)x_t = \mu + \varepsilon_t$ where

$$a(L) = 1 - \rho_1 L. \quad \text{(A.60)}$$

We can factor this, trivially, as $a(L) = 1 - \alpha_1 L$ where $\alpha_1 = \rho_1$. Stability requires that $1/\alpha_1 > 1$, or that $1/\rho_1 > 1$ or that $\rho < 1$. While this is a lot of additional work for an AR(1), it extends easily to higher-order AR processes.

Some examples of AR(2) processes. Let

$$x_t = 1.5x_{t-1} + x_{t-2} + \varepsilon_t, \quad \text{(A.61)}$$

so the lag polynomial is $a(L) = 1 - 1.5L - 1L^2$. You can find the roots either with the quadratic formula or by factoring this as $a(L) = (1 - 2L)(1 + .5L)$. The roots are 1/2 and -2. Since one of the roots is less than one in absolute value, this process is *not* stable. Which kind of makes intuitive sense here, as that x_{t-1} is emphasizing any shock, and the x_{t-2} term is not doing anything to "dampen" the shock. This will have a variance that increases with time.

On the other hand, this process

$$x_t = -.3x_{t-1} + 0.1x_{t-2} + \varepsilon_t \quad \text{(A.62)}$$

has a lag polynomial of $a(L) = 1 + .3L - .1L^2$ and factors to $a(L) = (1 + .5L)(1 - .2L)$, and the roots are -2 and 5. This process is stable because the roots are both bigger than one in absolute value. Again, there is intuitive to some extent. The coefficients on the lag terms are small so that the effect of a shock "dies out" over time.

The lag polynomial structure can provide the way to "iterate" the AR(k) process into a moving-average, as we did with the AR(1). But with multiple lags doing this by hand is almost impossible. For a given AR(k) process, though

$$a(L)x_t = \varepsilon_t \quad \text{(A.63)}$$

then in principle we could invert the $a(L)$ term and write x_t as a function of ε_t like this

$$x_t = a(L)^{-1}\varepsilon_t. \quad \text{(A.64)}$$

But inverting this whole polynomial thing is not obvious, and we again need a little help with notation.

Definition A.7 (Inversion of lag polynomial term) *Define the inversion $(1-\alpha_j L)^{-1}$ as the term that solves*

$$(1-\alpha_j L)^{-1}(1-\alpha_j L) = 1. \tag{A.65}$$

We can apply this across all the factored terms of a lag polynomial,

Definition A.8 (Inversion of lag polynomial and AR(k) processes) *For the lag polynomial $a(L)$, it can be factored and inverted as*

$$a(L)^{-1} = (1-\alpha_1 L)^{-1}(1-\alpha_2 L)^{-1}....(1-\alpha_k L)^{-1} \tag{A.66}$$

so that an $AR(k)$ can be written as

$$x_t = a(L)^{-1}\varepsilon_t = (1-\alpha_1 L)^{-1}(1-\alpha_2 L)^{-1}....(1-\alpha_k L)^{-1}\varepsilon_t. \tag{A.67}$$

What isn't obvious from this is that this represents a way of writing x_t as an infinite sum of the ε_t terms, or that we've iterated out the entire process and turned it into a moving average of the errors.

Let's work through how that works. There is a property of this (unproved) that

Definition A.9 (Inversion of stable lag polynomial term) *If $|\alpha_j| < 1$ (which recall is required for stability) then*

$$(1-\alpha_j L)^{-1} = \sum_{i=0}^{\infty} \alpha_j^i L^i. \tag{A.68}$$

The inversion of a lag polynomial term can, if that term has a stable coefficient, be expressed as a sum of lag terms. Note that this is turning something that is a function of one lag (the inverted polynomial term) into something with infinite lags, so that it is iterating out a process. Also note that this is nothing more than a generic version of the property that $\sum_{i=0}^{\infty} a^i = 1/(1-a)$ if $|a| < 1$.

For our AR(1) process, $a(L)x_t = \mu + \varepsilon_t$, all this is saying we should be able to write

$$x_t = a(L)^{-1}\mu + a(L)^{-1}\varepsilon_t \tag{A.69}$$

which given $a(L) = (1-\rho_1 L)$ means we should have

$$x_t = (1-\rho_1 L)^{-1}\mu + (1-\rho_1 L)^{-1}\varepsilon_t \tag{A.70}$$

and if $|\rho_1| < 1$, this is stable, then we can write

$$x_t = \mu \sum_{i=0}^{\infty} \rho_1^i + \sum_{i=0}^{\infty} \rho_1^i L^i \varepsilon_t. \tag{A.71}$$

The first term simplifies because $L^i\mu = \mu$ for any lag. The second term is just saying we should take lags of the error term weighted by ρ_1, so we have

$$x_t = \frac{1}{1-\rho_1}\mu + \sum_{i=0}^{\infty} \rho_1^i \varepsilon_{t-i}. \tag{A.72}$$

This is just what we established in the main material, which is that the AR(1) process can be expressed as an infinite moving average of the error terms.

For an AR(2) this is where all the work pays off. Take the example of

$$x_t = -.3x_{t-1} + 0.1x_{t-2} + \varepsilon_t \tag{A.73}$$

which has a lag polynomial of $a(L) = (1+.5L)(1-.2L)$. If we invert the two terms

$$(1+.5L)^{-1} = \sum_{i=0}^{\infty} -.5^i L^i = 1 - .5L + .25L^2 + ... \tag{A.74}$$

and

$$(1-.2L)^{-1} = \sum_{i=0}^{\infty} .2^i L^i = 1 + .2L + .04L^2 + ... \tag{A.75}$$

then

$$a(L)^{-1} = (1-.5L+.25L^2+....)(1+.2L+.04L^2+....) = 1-.3L+.28L^2+... \tag{A.76}$$

so that

$$x_t = \varepsilon_t - .3\varepsilon_{t-1} + .28\varepsilon_{t-2} + ... \tag{A.77}$$

as the infinite moving-average representation of x_t. Not that you'd want to work out the infinite number of terms in the moving average.

We already established the definition of a general moving average process in Definition 9.5. There is a lag operator way of writing these as well.

Definition A.10 (Lag polynomial moving average process) *A moving average process for x_t of order q, referred to as an MA(q) process, can be written as*

$$x_t = \mu + b(L)\varepsilon_t \tag{A.78}$$

where $b(L)$ is the lag polynomial $b(L) = 1 + b_1 L + b_2 L^2 + ... + b_q L^q$.

By itself this isn't very exciting. The same properties apply to $b(L)$ as before. We can factor it into a series of $b(L) = (1-\beta_1 L)(1-\beta_2 L)....(1-\beta_q L)$.

Definition A.11 (Invertability of moving average) *A moving average process for $x_t = \mu + b(L)\varepsilon_t$ is invertible if the absolute value of all the roots of $b(L)$ are greater than one, and an invertible moving average process can be written as*

$$\varepsilon_t = b(L)^{-1}x_t - b(L)^{-1}\mu. \tag{A.79}$$

So in one sense invertibility for a moving-average is like stability for an AR process. The reason invertibility matters here is that we could, given an invertible process, recover the value of ε_t (or other lags of it) given the data on the x_t terms. These are normally unobserved shocks, but what this is saying is that if the lag polynomial is invertible (in essence, is stable) then we can extract those unobserved shocks from observed information.

Definition A.12 (ARMA processes) *A x_t process is ARMA(k,q) if it has the form*

$$x_t = \mu + \rho_1 x_{t-1} + \rho_2 x_{t-2} + ... + \rho_k x_{t-k} + \varepsilon_t + b_1 \varepsilon_{t-1} + b_2 \varepsilon_{t-2} + ... + b_q \varepsilon_{t-q} \quad \text{(A.80)}$$

which can be expressed as

$$a(L)x_t = \mu + b(L)\varepsilon_t \quad \text{(A.81)}$$

where $a(L)$ is the lag polynomial for the AR terms and $b(L)$ is the lag polynomial for the MA terms.

The properties of an ARMA are inherited from the two parts.

Definition A.13 (Stability of ARMA) *An ARMA(k,q) model is stable if the roots of the $a(L)$ polynomial are all larger than one in absolute value, and then it can be written as*

$$x_t = a(L)^{-1}\mu + a(L)^{-1}b(L)\varepsilon_t \quad \text{(A.82)}$$

which is an infinite moving average.

Definition A.14 (Invertibility of ARMA) *An ARMA(k,q) model is invertible if the roots of the $b(L)$ polynomial are all larger than one in absolute value, and then it can be written as*

$$a(L)b(L)^{-1}x_t = b(L)^{-1}\mu + \varepsilon_t \quad \text{(A.83)}$$

which is an AR(∞) model of the ε_t process.

A.19 Ramsey to Solow

In the Ramsey model consumption growth is

$$g_c = \frac{1}{\sigma}(\epsilon_K y/k - \delta - \theta). \quad \text{(A.84)}$$

while from the production side we have that

$$g_y = \epsilon_K(s_I y/k - \delta) + \epsilon_L g_A - \epsilon_K g_L. \quad \text{(A.85)}$$

If the savings rate s_I is going to remain constant, it has to be that $g_c = g_y$, so the game here is to figure out what has to hold such that $g_c = g_y$. This isn't a case of setting $g_c = g_y$ in these equations, because we need this to hold at all y/k. It's identifying which parameter values in the Ramsey model just happen to lead to $g_c = g_y$.

This gets easier if you write

$$g_c = \epsilon_K (1/\sigma) y/k - (1/\sigma)\delta - (1/\sigma)\theta \tag{A.86}$$

and then if

$$(1/\sigma)\theta = -(1/\sigma)\delta + \epsilon_K \delta - \epsilon_L g_A + \epsilon_K g_L \tag{A.87}$$

that

$$g_c = \epsilon_K (1/\sigma) y/k - \epsilon_K \delta + \epsilon_L g_A - \epsilon_K g_L \tag{A.88}$$

ensures that $g_c = g_y$ at all times.

Therefore if

$$\theta = \sigma \epsilon_K (\delta + g_A + g_L) - \sigma g_A - \delta \tag{A.89}$$

the individuals in the Ramsey model will pick a situation where they keep s_I constant at all times no matter the level of k/y, and that savings rate will be

$$s_I^* = \frac{1}{\sigma}. \tag{A.90}$$

In that sense we could justify the Solow model as being the outcome of an individual utility-maximizing problem where they had exactly this specific choice of θ and where it was the case that strictly $\sigma > 1$.

If you take this seriously as a statement about economies, then a 25% s_I implies that $\sigma \approx 4$. What kind of discount rate makes this plausible? If $\epsilon_K \approx .3$, $\delta = 0.05$, $g_A = 0.02$, and $g_L = 0.01$, then you'd need

$$\theta \approx 4(0.3)(0.05 + 0.02 + 0.01) - 4(0.02) - 0.05 \approx -0.034 \tag{A.91}$$

or we'd need people to *prefer* the future, which cannot happen because then utility is infinite and that's incompatible with saving only 25% (you'd want to save everything). To make θ a standard number we'd need to adjust something else, like say ϵ_K. If we make that 0.5 we recover a plausible value of $\theta \approx 0.03$. In this case, the steady state interest rate would be $r^* \approx 0.03 + 4(0.02) = 0.11$ which seems high, but perhaps you could argue yourself into thinking that was plausible if we accounted for risk or something like that.

A different option is to presume that g_A is more like 0.01, rather than 0.02, and you can set $\epsilon_K \approx 0.33$. Then you get an implied θ of around 0.0024, which is at least positive. The steady state interest rate is $r^* \approx 0.0024 + 4(0.01) = 0.0424$, and that seems like a reasonable number.

A.20 Behavior of savings rates in the Ramsey model

What we want to determnine is how the consumption/output ratio, c/y (and hence s_I) changes over time before it gets to steady state. We're dealing here with the (more normal) case that k/y is *below* whatever steady state it has, and therefore that $g_{ky} > 0$ and $d(k/y) > 0$. We know that if we're below steady state then

$$k/y < \frac{\epsilon_K}{(\delta + \theta + \sigma g_A)} \tag{A.92}$$

or

$$\epsilon_K y/k > (\delta + \theta + \sigma g_A). \tag{A.93}$$

The marginal product of capital in this case is higher than the "drag" from depreciation and impatience/preferences.

The growth rate of $c/y = g_c - g_y$

$$g_{cy} = \frac{1}{\sigma}\epsilon_K y/k - \frac{1}{\sigma}(\delta + \theta + \sigma g_A) - \epsilon_K s_I y/k + \epsilon_K(\delta + g_A + g_L) \tag{A.94}$$

and given that

$$s_I^* = \epsilon_K \frac{\delta + g_L + g_A}{\theta + \delta + \sigma g_A}. \tag{A.95}$$

we can write this as

$$g_{cy} = \left[s_I^* - \frac{1}{\sigma}\right](\delta + \theta + \sigma g_A) - \left[s_I - \frac{1}{\sigma}\right]\epsilon_K y/k \tag{A.96}$$

You can see the sign of this depends on two comparisons of the savings rate to $1/\sigma$, the current savings rate and the steady state. Because we are below steady state it's the case that the weight on the second term is bigger than on the first.

We can work through three different conditions based on the steady state.

1. $s_I^* = 1/\sigma$. If it so happens that this is true, then the whole first term is zero, and the sign of g_{cy} depends on $1/\sigma$ relative to the current s_I. If $s_I < 1/\sigma$ it is below the steady state value. But then $g_{cy} > 0$ and so $ds_I < 0$ or s_I would get *farther* from steady state, and g_{cy} would become more positive and it would keep getting farther. Reverse the logic for the $s_I > 1/\sigma$ case and you see it isn't stable either. The only possible solution here that ends up at a steady state is if $s_{It} = 1/\sigma$ for all t. That is, if $s_I^* = 1/\sigma$ then s_I is constant as in the Solow model.

2. $s_I^* < 1/\sigma$. In this case the first term is negative. If $s_I > 1/\sigma$ (above steady state), then the second term is negative as well, and $g_{cy} < 0$ or $ds_I > 0$ and s_I would get *farther* away from steady state. It must be that $s_{It} < 1/\sigma$ at all times.

3. $s_I^* > 1/\sigma$. In this case the first term is positive. If $s_I < 1/\sigma$ (below steady state), then the second term is positive as well, and $g_{cy} > 0$ or $ds_I < 0$ and s_I would get *farther* away from steady state. It must be that $s_{It} > 1/\sigma$ at all times.

We can say something stronger than this. If $s_I > s_I^*$ then to approach steady state we need $g_{cy} > 0$ so that $ds_I < 0$. And if $s_I < s_I^*$ then we need $g_{cy} < 0$ and $ds_I > 0$. Knowing that:

1. $s_I^* > 1/\sigma$. We have that $s_{It} > 1/\sigma$ as well. The only answer consistent with reaching steady state is if $s_I^* > s_{It} > 1/\sigma$ so that $g_{cy} < 0$ and $ds_I > 0$. If the IES is "small" then you want to smooth consumption. You set the savings rate low to start and let the consumption rate fall over time to offset the fact that accumulation and productivity growth will increase the per-capita consumption for you later on.

2. $s_I^* < 1/\sigma$. We have that $s_{It} < 1/\sigma$ as well. The only answer consistent with reaching steady state is if $1/\sigma > s_{It} > s_I^*$ so that $g_{cy} > 0$ and $ds_I < 0$. If the IES is "big" you are more willing to save early and let your consumption grow not just due to accumulation and productivity growth, but also because you consume a bigger share of it.

3. $s_I^* = 1/\sigma$. In this case the only consistent answer is for $s_I = 1/\sigma$ at all times.

The behavior of the savings rate depends a lot on the IES, as that dictates how interested people are in letting consumption grow quickly. But remember that s_I^* depends on both σ and θ, so which situation we end up in depends a lot on whether people are patient or not.

Bibliography

David Baqaee and Emmanuel Farhi. A Short Note on Aggregating Productivity. NBER Working Papers 25688, National Bureau of Economic Research, Inc, March 2019.

Simcha Barkai. Declining labor and capital shares. *The Journal of Finance*, 75(5):2421–2463, 2020.

Susanto Basu and John Fernald. Aggregate productivity and aggregate technology. *European Economic Review*, 46:963–991, 2002.

David Cass. Optimum growth in an aggregative model of capital accumulation. *The Review of Economic Studies*, 32(3):233–240, 1965.

Gene M. Grossman, Elhanan Helpman, Ezra Oberfield, and Thomas Sampson. Balanced growth despite Uzawa. *American Economic Review*, 107(4):1293–1312, April 2017.

Robert E. Hall. Invariance properties of solow's productivity residual. In Peter Diamond, editor, *Growth, Productivity, Employment*. MIT Press, Cambridge, MA, July 1990.

Berthold Herrendorf, Richard Rogerson, and Akos Valentinyi. Two perspectives on preferences and structural transformation. *American Economic Review*, 2013.

Karen J. Horowitz and Mark A. Planting. *Concepts and Methods of the U.S. Input-Output Accounts*. U.S. Department of Commerce Bureau of Economic Analysis, 2009.

Charles I. Jones and Dean Scrimgeour. A new proof of Uzawa's steady-state growth theorem. *The Review of Economics and Statistics*, 90(1):180–182, 2008.

Òscar Jordà, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M Taylor. The Rate of Return on Everything, 1870–2015*. *The Quarterly Journal of Economics*, 04 2019.

Nicholas Kaldor. A model of economic growth. *The Economic Journal*, 67(268):591–624, 1957.

Tjalling Koopmans. On the concept of optimal economic growth. In *The Economic Approach to Development Planning*, pages 225–287. Rand McNally, 1965.

Wassily W. Leontief. *Input-Output Economics*. Oxford University Press, 1966.

Edmund Phelps. The golden rule of accumulation: A fable for growthmen. *The American Economic Review*, 51(4):638–643, 1961.

F. P. Ramsey. A mathematical theory of saving. *The Economic Journal*, 38(152):543–559, 1928.

Robert M. Solow. A contribution to the theory of economic growth. *The Quarterly Journal of Economics*, 70(1):pp. 65–94, 1956.

H. Uzawa. Neutral Inventions and the Stability of Growth Equilibrium. *The Review of Economic Studies*, 28(2):117–124, 02 1961.