

Intermediate Macroeconomics: Economic Growth and the Solow Model

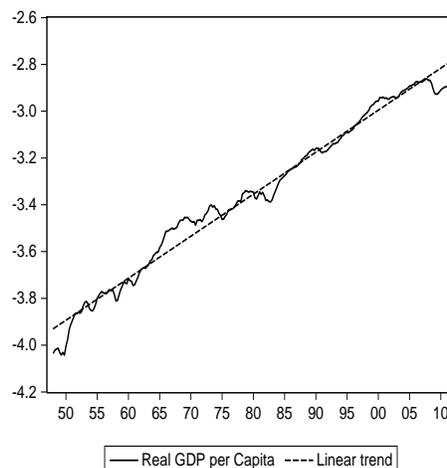
Eric Sims
University of Notre Dame

Fall 2012

1 Introduction

We begin the course with a discussion of economic growth. Technically growth just refers to the period-over-period percentage change in a variable. In the media you hear lots of talk about current “growth” in GDP as a reference to the business cycle. When economists talk about growth, however, we are usually referencing changes in GDP at a lower frequency – i.e. thinking about the sustained increases in GDP over a decade as opposed to what’s happening quarter to quarter.

The sustained increases in GDP over time dominate any discussion of what happens at higher frequencies. Below I plot log real GDP per capita in the US from 1947 to the second quarter of 2012. I also fit a linear time trend and show that as the dashed line.



As we’ve noted before, the trend dominates any gyrations about the trend. The trend line that I fit grows at a rate of 0.45 percent per quarter, or about 1.8 percent at an annualized rate.

Small growth rates can compound up to very big differences in levels over long time periods. If a variable is growing at a constant rate, its level j periods into the future relative to the present is

given by (where g_x is the constant growth rate):

$$X_{t+j} = (1 + g_x)^j X_t$$

Suppose that we take the unit of time to be a year, and that a variable in question is growing at 2 percent. This means that, relative to the present, the value of the variable will be equal to:

$$\frac{X_{t+10}}{X_t} = (1 + 0.02)^{10} = 1.22$$

In other words, if a variable grows at 2 percent per year for 10 years straight, the level of the variable will be 22 percent bigger in 10 years. Suppose instead that the variable grows at 2.5 percent per year. 10 years later we'd have:

$$\frac{X_{t+10}}{X_t} = (1 + 0.025)^{10} = 1.28$$

That extra half of a percentage point of growth nets 6 percentage points more growth over a 10 year period. The differences are even more remarkable if you expand the time horizon – let's go to, say, 30 years, about the gap between generations. Growing at 2 percent per year, we'd have:

$$\frac{X_{t+30}}{X_t} = (1 + 0.02)^{30} = 1.81$$

Growing at 2 percent per year nets us a level that is 80 percent higher after 30 years. Growing at 2.5 percent per year, we'd have:

$$\frac{X_{t+30}}{X_t} = (1 + 0.025)^{30} = 2.10$$

With just a half of a percentage point more of growth per year, over a 30 year horizon the level of X would more than double, increasing by 110 percent. That extra half of a percentage point of growth, which on its own seems quite small, gets us an extra 30 percentage points in the level over 30 years. This is a big number. Current real per capita GDP in the United States is \$50,000, give or take. If that were to grow at 2 percent per year for the next 30 years, per capita real GDP would be about \$90,500. If instead we grew at 2.5 percent per year, 30 years from now real per capita GDP would be about \$105,000. That's about a \$15,000 difference, which is big.

The bottom line here is that growth translates into large differences in levels over long periods of time. This means that it is critically important to understand growth. If we could get the economy to grow even just a little faster on average, this would have large benefits down the road.

2 Stylized Facts

“Stylized facts” are broad generalizations that summarize recurrent features of data. Kaldor (1957) looked at empirical data on economic growth and came up with the following list of “stylized facts.” By “stylized” it should be recognized that, as written, these facts are not literally true, but seem

to hold in an approximate sense over a long period of time.

1. Output per worker grows at a roughly constant rate over time
2. Capital per worker grows at a roughly constant rate over time, the same rate at which output grows (so that the capital-output ratio is roughly constant)
3. The rate of return on capital (closely related to the real interest rate) is roughly constant
4. The return on labor (the real wage) grows at a roughly constant rate, the same rate as output and capital

These are *time series* facts in that they describe the behavior of a single economy over time. There are also *cross-sectional* facts, which look at variation across countries at a given point in time. These are:

1. There are very large difference in per capita GDP across countries
2. There are examples where poor countries “catch up” to rich countries (growth miracles)
3. There are also examples where countries do not catch up (growth disasters)

We are going to construct a model which is going to help us think about economic growth. We will compare the predictions of that model to some of the facts in the data. To the extent to which the model has predictions that align with the facts, we can be confident that the model is a pretty good description of reality. If we think the model is a good description of reality, we can be comfortable in using that model to draw some inference about what kind of policies might be desirable.

The model we are going to build is called the “Solow model,” or sometimes the “neoclassical growth model” after Solow (1957). A downside of the model is that it does not explain where growth comes from; but if there is something like “knowledge” or “productivity” that ones takes as given as growing over time, the model does a very good job at explaining the time series facts. The model has the important implication that the primary determinant of growth is productivity. Saving, which leads to more capital accumulation, cannot sustain growth.

On its surface, the Solow model does less well at the cross-sectional facts. For example, differences in saving rates (and hence different levels of capital accumulation) cannot account for the large disparities in levels of GDP per capita that we observe across countries (for the same reason that saving rates cannot sustain growth either). Also, if some countries are poor only because they don’t have enough capital, the model predicts that these countries should grow faster to catch up to rich countries. Though there are some examples of countries that “catch up” to rich countries, there are also lots of examples where this does not happen, where the large differences in standards of living persist through time. The only way for the Solow model to account for large, persistent differences in standards of livings across countries is for there to be large differences in the levels of productivity, which is sometimes called “static efficiency.” This means that it is really important to better understand the sources of productivity.

3 The Basic Model

Time is discrete, and we denote it by $t = 0, 1, 2, \dots$. This time could denote different frequencies – e.g. $t = 0$ could be 1948, $t = 1$, 1949, and so on (annual frequency); $t = 0$ could be 1948q1, $t = 2$, 1948q2, and so on (quarterly frequency); or $t = 0$ could be 1948m1, $t = 1$ 1948m2, and so on (monthly frequency). Most macroeconomic data from the NIPA accounts are available at best at a quarterly frequency, so, for the most part, I think of dates as being quarters, but it could be months, years, or even weeks or days.

The economy is populated by a large number of households and firms. For simplicity, assume that these households and firms are all identical. Since they are all identical, we can normalize things such that there is one firm and one household (though later I will allow the “size” of the household to grow to account for population growth).

3.1 Firm

The firm produces output using two factors of production: capital and labor. Both of the factors of production are owned by the household and are leased to the firm on a period-by-period basis. It is helpful to fix ideas to think about output as being “fruit” – pineapple, banana, whatever. The reason I like the fruit analogy is that we are going to assume that output is not storable – it is produced in a period (say t), and it can be consumed or re-invested in that period, but you can’t simply hold on to it and eat it tomorrow. Fruit has this property of non-storability and is therefore convenient.

Labor is denominated in units of time – it is how much time people spend working to produce stuff. The household only has so much labor it can supply in a given period – if that labor is not used in a period, it is forever lost. Capital is denominated in units of goods. Capital is different from labor in the following two ways: (i) it must itself be produced (whereas labor is an endowment – you have time available exogenously) and (ii) the supply of capital is not exhausted within a period (using capital today does not preclude you from using it to produce tomorrow). Think about capital as a fruit tree, which itself had to be planted via un-eaten fruit at some point in the past. The fruit tree itself can exist across time and can yield fruit in multiple periods. Hence, the production process involves trees which yield fruit (capital) and people which spend time picking the fruit off the trees (labor).

We assume that these two factors of production (labor hours and capital) are combined using some function to yield output (fruit). Output is a flow concept – it is the amount of new fruit picked in a period. Denote capital at time t as K_t and labor at time t as N_t . New output (fruit) produced at time t is given by:

$$Y_t = AF(K_t, N_t) \tag{1}$$

$F(\cdot)$ is the aggregate production function, and A is a productivity shifter that we will sometimes call “static efficiency.” You can think about A being different both across space (e.g. states or

countries) or across time (e.g. 2011 vs 2010), although for now I'm going to omit a time subscript. For example, one area of the country (say Indiana) may have more fertile soil than another (say Nevada). This means that, for a given amount of capital and labor, the firm could produce more in Indiana than in Nevada, so the A in Indiana would be higher than in Nevada. Alternatively, you could think about this evolving over time. One year may have more rainfall than the other year. Since rain is good for growing fruit, the year with more rainfall would yield more fruit for given amounts of capital and labor, and hence would have a higher A .

We impose that the production function has the following properties:

1. Both factors are necessary to produce anything
2. For a given amount of one factor, more of the other factor results in more output
3. The amount by which an additional factor increases output (holding the other factor fixed) is decreasing in the amount of that factor
4. If you double both factors, you double output

Mathematically, these properties can be represented:

$$\begin{aligned}
 F(K_t, 0) &= F(0, N_t) = 0 \\
 F_K(K_t, N_t) &> 0, \quad F_N(K_t, N_t) > 0 \\
 F_{KK}(K_t, N_t) &< 0, \quad F_{NN}(K_t, N_t) < 0 \\
 F(\gamma K_t, \gamma N_t) &= \gamma F(K_t, N_t), \quad \gamma > 0
 \end{aligned}$$

Mathematically this means that the production function is increasing and concave in both of its arguments and is homogeneous of degree one (equivalently we say that the production function features constant returns to scale). A particular production function that satisfies these requirements is the Cobb-Douglas production function, which we will use throughout the semester:

$$Y_t = AK_t^\alpha N_t^{1-\alpha}, \quad 0 \leq \alpha \leq 1 \tag{2}$$

The firm wants to maximize its profit, which is equal to revenue minus cost. Revenue is just total output which ends up being sold to the household (this is in real terms and we are completely abstracting from money, meaning everything is denominated in units of goods, e.g. fruit). Total cost is the wage bill plus the capital bill. Let w_t be the real wage rate – it is the number of goods the firm must pay each unit of labor. Let R_t be the real “rental rate” – it is the number of goods the firm must give up to lease a unit of capital. The firm is a price-taker, so it takes these as given. Profit is therefore:

$$\Pi_t = AF(K_t, N_t) - w_t N_t - R_t K_t \tag{3}$$

The firm wants to pick capital and labor to maximize profit. The problem is therefore:

$$\max_{K_t, N_t} AF(K_t, N_t) - w_t N_t - R_t K_t$$

The solution is characterized by taking the partial derivatives of the production function with respect to each input and setting them equal to zero:

$$\frac{\partial \Pi_t}{\partial K_t} = 0 \Leftrightarrow AF_K(K_t, N_t) = R_t \tag{4}$$

$$\frac{\partial \Pi_t}{\partial N_t} = 0 \Leftrightarrow AF_N(K_t, N_t) = w_t \tag{5}$$

Because of the concavity assumption, these two conditions imply downward sloping demand curves for each factor input – the bigger the wage, for example, the less labor a firm will want, holding all factors constant. An increase in A will shift the factor demand curves out for both capital and labor, meaning that firms will want more of both inputs at given factor prices.

Because of the constant returns to scale assumption, it turns out that the firm will earn no profits. This is easiest to see by using the Cobb-Douglas form of the production function, so that the optimality conditions are:

$$\begin{aligned} \alpha K_t^{\alpha-1} N_t^{1-\alpha} &= R_t \\ (1 - \alpha) K_t^\alpha N_t^{-\alpha} &= w_t \end{aligned}$$

With these factor demands, we see that $R_t K_t = \alpha Y_t$ and $w_t N_t = (1 - \alpha) Y_t$. Therefore $\Pi_t = Y_t - \alpha Y_t - (1 - \alpha) Y_t = 0$, so there are no profits. Also, with this functional form, α has the interpretation as the share of total income that gets paid out to capital, and $1 - \alpha$ as the share of total income paid out to labor. So α will sometimes be called “capital’s share.”

You may take issue with the notion that the firm earns no profits, because firms in the real world do earn profits. It’s important to draw the distinction between accounting and economic profit. The way I’ve set the model up here there is no distinction between the two, and this is because of a particular way of modeling the ownership structure of capital. In the real world, firms typically own their own capital, and firms are owned by households via shares of common stock. The way I’ve modeled it here households own the capital and lease it to firms. For most specifications, these setups turn out to be isomorphic, but it is often easier to think about the household owning the capital stock. At an intuitive level, the reason is that the household owns the firm, so whether the firm “owns” the capital stock or not is immaterial – the household really owns it either way. Had I set up the model where the firm owns the capital stock, the firm would earn an accounting profit that would be remitted back to the household via dividends. There would be no economic profit, however, because the accounting profit would just be equal to the best outside option, which would be to put the capital in a different firm (remember, there are many identical firms, which we treat

as one firm). Essentially profits would take the role of what amounts to “capital income” for the household.

3.2 Household

The household problem is particularly simple. In fact, we eschew optimization altogether to make life easy. Later on in the course we will make the household problem more interesting.

Households own the capital and have an endowment of labor each period. They earn income from leasing their capital to firms $R_t K_t$ and supplying their labor $w_t N_t$. Total household income is then income from their factor supplies plus any profits remitted from the firm, Π_t (which, as we saw above, is going to be zero given our assumptions).

Households can use their income (which is denominated in units of goods, i.e. fruit) to do two things: consume, C_t , or invest in new capital, I_t . Hence, the household budget constraint is:

$$C_t + I_t \leq w_t N_t + R_t K_t + \Pi_t \tag{6}$$

A quick note. What appears in a budget constraint is a weak inequality sign – consumption plus investment must be less than or equal to total income. Put differently, expenditure cannot exceed total income. Nothing prevents the household from “wasting” some of its income, so that expenditure could, in principle, be less than income. As long as preferences are such that households always “like” more consumption, however, this shouldn’t happen, so most of the time we’ll just assume that the budget constraint binds with equality, and we’ll therefore usually write it with an equal sign.

We make two simple assumptions: first, the household supplies one unit of labor inelastically each period, i.e. $N_t = 1$; and, second, the household consumes a constant fraction of its income each period, equal to $1 - s$, where $0 < s \leq 1$, where s is the saving rate. These are assumptions and don’t necessarily come out of consumer optimization, but they are pretty good approximations to the behavior of the economy over the long haul. Relating this back to some definitions of the aggregate labor market, if we are going to assume that labor is inelastically supplied, we need not make specific assumptions about how this translates into the extensive versus the intensive margin. You could think of the household having many members, and some fraction of those members always work h hours each period, or there could just be one member who works h each period. All that really matters is the total labor input, N , and if we take that to be inelastically supplied (and normalized here to 1), we don’t need to worry about other details.

The current capital stock, i.e. the capital at time t , is predetermined, meaning it cannot be changed within period. This reflects the fact that capital must itself be produced – to get more capital, you first have to produce more and choose not to consume all of it. Capital tomorrow can be influenced by investment today. Investment, from the budget constraint above, is just income (the right hand side) less consumption. Think about it this way. Suppose the household has income of 10 units of fruit and eats 8 fruits. It takes the remaining 2 fruits and plants them in the ground, which will yield 2 additional fruit trees (e.g. capital) available for production in the next period

(we assume there is a one period delay, but could generalize it to multiple periods). Some of the existing capital (e.g. trees) decay each period. We call the fraction of capital that becomes obsolete (non-productive) each period the depreciation rate, and denote it by δ . The capital accumulation equation is given by:

$$K_{t+1} = I_t + (1 - \delta)K_t \quad (7)$$

This just says that the capital available tomorrow (think of today as period t) is investment from today (new contributions to the capital stock) plus the non-depreciated component of the existing capital stock, $(1 - \delta)K_t$.

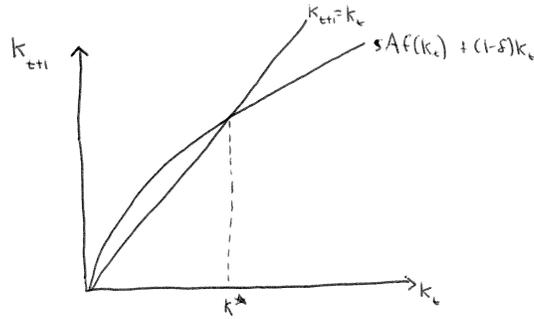
3.3 Aggregation and the Solow Diagram

Now we combine elements of the household and firm problems to look at the behavior of the economy as a whole. Since the firms earn no profits, $\Pi_t = 0$ and $Y_t = w_t N_t + R_t K_t$. Since, by assumption, the household consumes a fixed fraction of income, and from above total income is equal to total output, we have $C_t = (1 - s)Y_t$. Plugging this all into the household budget constraint reveals that $I_t = sY_t$. Hence, we can equivalently call s both the saving rate and the investment rate, since, in equilibrium, investment must be equal to saving. Now using the fact that there is just one household that inelastically supplies one unit of labor to the firm, $N_t = 1$, we get $Y_t = AF(K_t, 1)$. Define $f(K_t) = F(K_t, 1)$. For the Cobb-Douglas production function, this would be $f(K_t) = K_t^\alpha$. Using this, plus the expression for investment, and plugging into the capital accumulation equation yields the central equation of the Solow growth model:

$$K_{t+1} = sAf(K_t) + (1 - \delta)K_t \quad (8)$$

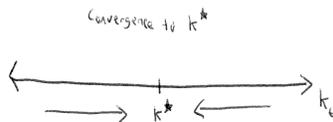
This single difference equation summarizes the model completely (it is called a “difference” equation in the sense that it relates a value of a variable in two adjacent periods of time . . . the continuous time analogue is a differential equation). It is helpful to analyze it graphically. We want to plot K_{t+1} against K_t . Given the assumptions we’ve made, this function will be increasing at a decreasing rate (this is driven by the concavity of $f(K_t)$). When $K_t = 0$ the function starts out at zero with a steep slope. As K_t gets bigger the slope gets flatter. Eventually the slope flattens out to $(1 - \delta)$. This is because $f_K(\cdot) > 0$ but $f_{KK}(\cdot) < 0$, so the slope of $f(K_t)$ has to go to zero, so that the slope of the RHS is just $(1 - \delta)$ (remember the slope of the sum is just the sum of the slopes, since the slope is just a derivative and the derivative is a linear operator).

When plotting this it is helpful (for reasons which will become clearer below) to also plot a “45 degree line” which shows all points in the plane where $K_{t+1} = K_t$. This has slope of 1. Because $sAf(K_t) + (1 - \delta)k_t$ eventually has slope $(1 - \delta) < 1$, we know that the curve must cross the straight line exactly once, provided the curve (which starts out at the origin) starts out with a slope greater than 1. For any of the production functions we use in this class that will be the case.



The curve, $K_{t+1} = sAf(K_t) + (1 - \delta)K_t$, crosses the line, $K_{t+1} = K_t$, at a point I mark as K^* . As noted above, as long as the curve starts out with slope greater than one, and finishes with slope less than one, these can cross exactly once at a non-zero value of K . This is a “special” point, and we’ll call it the “steady state” capital stock. We call it the steady state because if $K_t = K^*$, then $K_{t+1} = K^*$, and, moving forward in time one period, $K_{t+2} = K^*$. In other words, if and when the economy gets to K^* , it will be expected to stay there forever. I say “expected” because it’s possible that A could change at some point in the future (more on those possibilities below).

It turns out that, for any initial K_t , the economy will be expected to approach K^* . So not only is K^* a point of interest because it’s a point at which the economy will be expected to sit, it’s also interesting because the internal dynamics of the model are working to take the economy there. You can see this by noting that, for a given K_t , you can “read off” K_{t+1} from the curve, while the 45 degree line reflects K_t onto the vertical axis. At current capital below K^* , it is easy to see that $K_{t+1} > K_t$ (the curve is above the line). At current capital above K^* , we see that $K_{t+1} < K_t$. This means that if we start out below the steady state, capital will be expected to grow. If we start out above the steady state, capital will be expected to decline. Effectively, no matter where we start we’ll be headed toward the steady state from the natural dynamics of the model.



We can algebraically solve for the steady state capital stock assuming the Cobb-Douglas production function: $F(K_t, N_t) = K_t^\alpha N_t^{1-\alpha}$, which implies $f(K_t) = K_t^\alpha$. The capital accumulation equation is:

$$K_{t+1} = sAK_t^\alpha + (1 - \delta)K_t$$

We can solve for K^* by setting $K_t = K_{t+1} = K^*$ and simplifying:

$$\begin{aligned} K^* &= sAK^{*\alpha} + (1 - \delta)K^* \\ \Rightarrow K^* &= \left(\frac{sA}{\delta}\right)^{\frac{1}{1-\alpha}} \end{aligned} \tag{9}$$

Given the steady state capital stock, we can compute steady state output (which corresponds to steady state output per worker, since we've normalized the labor input to one) and steady state consumption:

$$Y^* = A \left(\frac{sA}{\delta}\right)^{\frac{\alpha}{1-\alpha}} \tag{10}$$

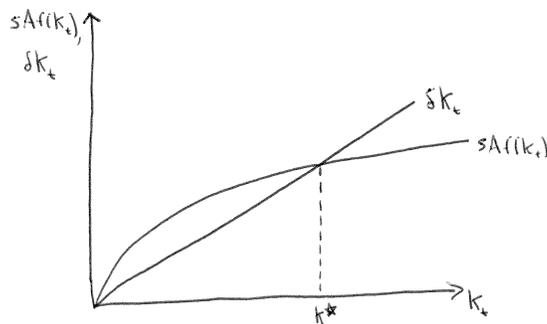
$$C^* = (1 - s)A \left(\frac{sA}{\delta}\right)^{\frac{\alpha}{1-\alpha}} \tag{11}$$

3.3.1 Alternative Graphical Depiction

Many textbooks draw the “Solow diagram” a bit differently. In particular, define $\Delta K_{t+1} = K_{t+1} - K_t$. Subtract K_t from both sides of (8) to get:

$$\Delta K_{t+1} = sAf(K_t) - \delta K_t \tag{12}$$

The alternative figure plots the two components of the right hand side against K_t , with δK_t plotted without the negative sign. $sAf(K_t)$ is sometimes called “saving” or “investment” and δK_t is called “break even investment.” If saving is greater than break even investment, the capital stock will be expected to grow. If saving is less than break even investment, the capital stock will be expected to decline.



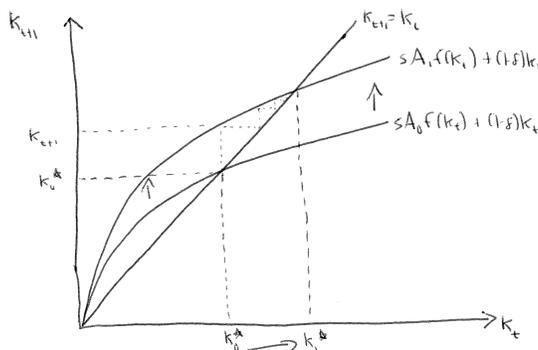
The figure has a very similar interpretation. To the left of K^* , the point of intersection and exactly the same K^* in the other picture, saving exceeds break even investment, and so the capital

stock grows. To the right of K^* , break even investment exceeds saving, and so the capital stock declines.

3.4 Comparative Statics

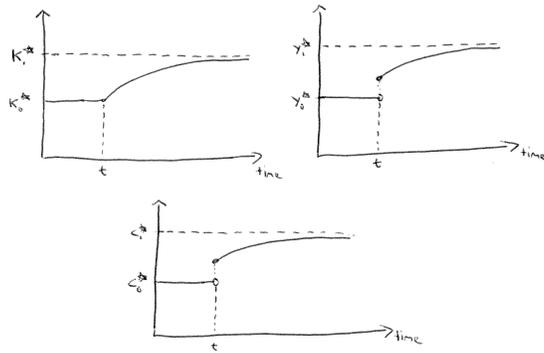
In this section we want to consider two exercises. What happens to the steady state and along the transition path when (i) A permanently increases or decreases and (ii) s permanently increases or decreases.

Suppose an economy initially sits in steady state. Suppose that at time t there is an immediate, surprise, and permanent increase in A from A_0 to A_1 . This will shift the curve in the Solow diagram up, so that it intersects the 45 degree line further to the right. The economy does not immediately go to the new steady state – recall that K_t is predetermined, and new capital must be produced. We can read the new K_{t+1} off of the vertical axis at the initial $K_t = K^*$ (the economy starts in a steady state) and the new curve. Then you can “iterate that forward” by moving the picture forward in time (essentially by just changing the time subscripts) to see how the economy transitions.

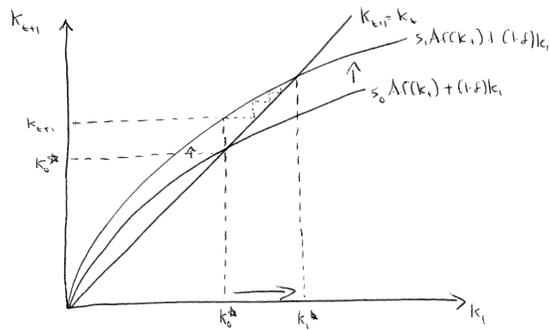


Eventually we will end up in a new steady state with a higher level of capita, $K_1^* > K_0^*$. Since the capital stock is, in a sense, what determines everything else, we know that we must also end up with higher output and consumption.

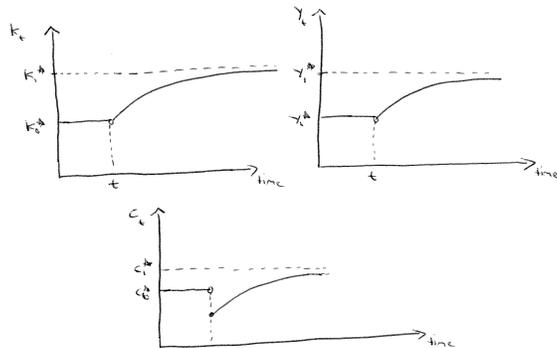
But how do we get there? What happens immediately when A increases? The capital stock does not move, but since output is the product of A with a function of K , output must jump up immediately. Since consumption is a fixed fraction of output, consumption must also jump up immediately. Starting in period $t+1$, the capital stock is growing, approaching its new steady state. If the capital stock is growing, output must continue to grow after its initial jump. Consumption must also do the same. Below I plot “impulse responses” of K_t , Y_t , and $C_t = (1-s)Y_t$. Though K_t does not react immediately, Y_t and hence C_t do, since A_t is immediately higher. After the initial jump in Y_t and C_t , all three of these series smoothly approach the new higher steady state:



Next let's consider an economy initially sitting in steady state, and suppose at time t that there is an immediate, surprise, and permanent increase in s , the saving rate. Similarly to the increase in A , this will shift the curve in Solow diagram up:



Below are the impulse responses of capital, output, and consumption:



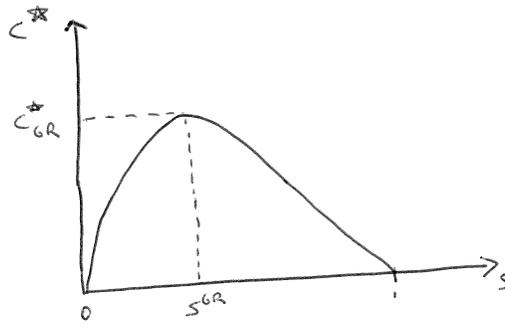
Now, even though the main "Solow diagram" looks similar to the case of an increase in A , the

dynamics that play out in output and consumption are not the same. Following the increase in the saving rate, there is an immediate decline in consumption. This is because output, which is a function of capital and A , cannot move within period since the capital stock is predetermined and A is not moving by assumption. Effectively, an increase in the saving rate means that households are consuming a smaller part of a fixed pie. The initial drop in consumption, coupled with no change in output, means that investment goes up (the economy is accumulating more capital). Hence, in the period after the change in the saving rate, output will start to grow, because the economy will be accumulating more capital. This means that consumption will start to grow in period $t + 1$, relative to its initial decline in period t . Effectively the pie starts growing in $t + 1$, and hence consumption begins to go up. But eventually, this growth goes away. We'll approach a new steady state K with an associated higher steady state Y . C may end up higher or lower (though I've drawn the figure where it ends up higher). There are countervailing effects going on – households are consuming a smaller fraction of the pie, but the pie is getting bigger. Which effect dominates is unclear (more on this below).

The most important insight from this exercise is the following: a permanent increase in the saving rate *cannot* lead to a permanent change in growth. We started in a steady state with zero growth, and ended up in a new (higher) steady state, also with zero growth. We do get output growing for a while as it approaches the new steady state, but this effect is temporary. Saving rates can affect levels of output, capital, and consumption over the long run, but not growth rates. Sustained growth must come from something else. You may be thinking that this “negative” result is special because I designed a model in which the variables do not grow at all in the long run anyway. As we'll see below, this result does not depend on that assumption.

Now I want to return to a point made above. It is unclear whether consumption ultimately goes up, down, or stays the same following an increase in the saving rate. Since people get utility from consumption (not from output per se), a reasonable question to ask is what is the saving rate that maximizes steady state consumption? Note we are focusing on consumption in the long run. The saving rate that would maximize current consumption is zero – consume the whole pie. The problem is that this would lead to low consumption in the future.

The saving rate that maximizes steady state consumption is what we call the “golden rule” saving rate. Intuitively, we can see that C^* is going to be a function of s – if $s = 0$, then $C^* = 0$, because we will converge to a steady state with no capital. If $s = 1$, we will converge to a steady state with a lot of capital but none of it will be consumed, so $C^* = 0$. From this we can surmise that steady state consumption ought to be increasing in the saving rate for low values of s and decreasing in the saving rate for high values of s . For example, moving from $s = 0$ to $s = 0.01$ has to raise C^* , because we go from zero steady state capital to something positive. Likewise, moving from $s = 1$ to $s = 0.99$ also has to raise C^* , because we'd be moving from 0 to positive consumption (consuming part of the pie instead of nothing of a larger pie). As such, we can intuit that C^* as a function of s looks something like the figure below, with the s associated with the maximum C^* being the “Golden Rule” saving rate.



As an exercise you'll be asked to find an expression for the Golden rule saving rate for the Cobb-Douglas production function. For empirically plausible production functions, we have a good idea that it is probably something like 0.3-0.4. Most modern economies (e.g. the US) have saving rates that are below this, which means that raising the saving rate would raise steady state consumption. Does this mean that we should enact policies to encourage more saving? Not necessarily. Keep in mind that the Golden rule is about consumption in the “long run” – as we saw above, an increase in the saving rate must be associated with an immediate decline in consumption. If the saving rate is below the Golden Rule, we know that the immediate decline in consumption will be followed by an ultimate increase in consumption. But we cannot say whether that is a good or a bad thing without knowing something about how households value present versus future consumption (e.g. how they “discount” the future). So we cannot definitively make a claim that we “should” try to encourage higher saving rates.

In contrast, we could make a judgment about a saving rate “above” the Golden Rule. If the saving rate were that high, a reduction in the saving rate would lead to an immediate increase in consumption and a long run increase in consumption. Regardless of how individuals discount the future, they would be better off by reducing the saving rate. We refer to a situation in which the saving rate exceeds the Golden Rule as “dynamic inefficiency” – it is inefficient in the sense that people could be better off both today and in the long run just by saving less. As noted above, for most modern economies, this does not seem to be much of an issue.

4 Accounting for Growth

A key result of the previous sections is that the economy converges to a steady state in which the capital stock (and hence also output) does not grow. This is flatly at odds with the data, and seems a strange result if our objective is to better understand the sources of growth. In this section we make two modifications that make the model better fall in line with the data, but which do not fundamentally alter any of the lessons from the simpler model. These are (i) accounting for population growth and (ii) accounting for productivity growth.

4.1 Population Growth

In the previous section we assumed that there was only one household and that it supplied its labor inelastically. We are going to continue to make the assumption of inelastic labor supply (over long horizons, average hours worked per person are roughly constant), but we will allow the “number” or “size” of households to grow to match facts about population growth. It does not really matter how one thinks about it – you can think of many households (that are all nevertheless the same), and the number of households growing over time; or one household that keeps increasing in size.

Since we continue to assume that labor is supplied inelastically, there is no material difference between population (L) and aggregate labor input (N) – if population grows and labor is inelastic, then labor must grow at the same rate. So we will assume that that aggregate labor input grows at rate g_n each period:

$$N_t = (1 + g_n)N_{t-1} \quad (13)$$

It is easy to manipulate this to see that g_n is the period-over-period growth rate in N . If you take logs and use the approximation that $\ln(1 + g_n) \approx g_n$, you will see that the growth rate is approximately equal to the log first difference. If you solve this expression “backwards” to the beginning of time (say $t = 0$), you get:

$$N_t = (1 + g_n)^t N_0 \quad (14)$$

Where N_0 is the population at the “beginning of time,” which we might as well normalize to one.

The rest of the Solow model is the same. This means that we can still reduce the model to the difference equation:

$$K_{t+1} = sAF(K_t, N_t) + (1 - \delta)K_t \quad (15)$$

For mathematical reasons, to “solve” difference equations like this, it is helpful to write the model in stationary terms. Define lowercase variables as “per-worker” or “per-capita” terms – since we are not modeling labor supply decisions, per-worker and per-capita will be the same (up to a constant). That is, let $y_t = \frac{Y_t}{N_t}$, $k_t = \frac{K_t}{N_t}$, and $c_t = \frac{C_t}{N_t}$. To transform the central Solow model equation, divide both sides by N_t and simplify:

$$\frac{K_{t+1}}{N_t} = \frac{sAF(K_t, N_t)}{N_t} + (1 - \delta)\frac{K_t}{N_t}$$

Now we need to do a couple of intermediate steps. First, because we assumed that $F(\cdot)$ has constant returns to scale, we know that $\frac{1}{N_t}F(K_t, N_t) = F\left(\frac{K_t}{N_t}, \frac{N_t}{N_t}\right)$. As before, define $f(k_t) = F(k_t, 1)$. Then we can write the difference equation as:

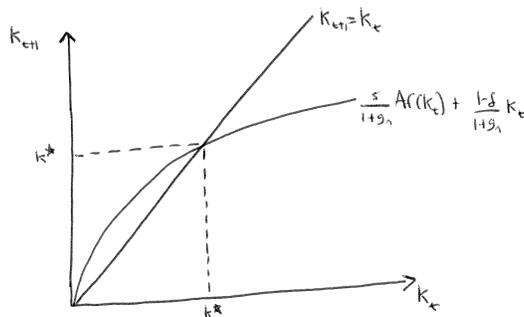
$$\frac{K_{t+1}}{N_t} = sAf(k_t) + (1 - \delta)k_t$$

In the Cobb-Douglas case, for example, we just have that $f(k_t) = k_t^\alpha$. Now, we are not finished

because we don't have the correct normalization on the left hand side. We need to multiply and divide by N_{t+1} to get this in terms of the per-capita variables we have defined:

$$\begin{aligned} \frac{N_{t+1} K_{t+1}}{N_t N_{t+1}} &= sAf(k_t) + (1 - \delta)k_t \\ (1 + g_n)k_{t+1} &= sAf(k_t) + (1 - \delta)k_t \\ k_{t+1} &= \frac{s}{1 + g_n}Af(k_t) + \frac{1 - \delta}{1 + g_n}k_t \end{aligned} \tag{16}$$

(16) is the modified main equation of the Solow model when we allow for population growth. Graphically we can proceed just as we did before, plotting k_{t+1} against k_t . The first part of the right hand side is a constant $\frac{s}{1+g_n}$ times a concave function, while the second part is linear in k_t . This means that the slope of the right hand side starts out steep and eventually goes to $\frac{1-\delta}{1+g_n}$. This is less than one, which means that the right hand side must cross the 45-degree line at some positive, finite value, call it k^* . This point where the series cross is called the steady state capital stock per worker (not the steady state capital stock).



As in the case with no population growth, we can do comparative statics which end up being qualitatively similar. In particular, an increase in either A or s leads the economy to transition to a new steady state in which the capital stock per worker is higher. To get to that new steady state the per-capita variables have to grow for a while, but eventually they hit a new steady state. So, the main conclusion of the previous section still holds – an increase in the saving rate leads to faster growth for a while, but not forever.

As in the case with no growth, sometimes you see the main Solow equation written in “difference” form. Subtracting k_t from both sides and defining $\Delta k_{t+1} = k_{t+1} - k_t$, we'd have:

$$\Delta k_{t+1} = sAf(k_t) - (\delta + g_n)k_t \tag{17}$$

This looks just like (12), but the “effective” depreciation rate is $\delta + g_n$. This makes some sense – the physical capital stock depreciates at rate δ , while the number of works grows at rate g_n . To

keep the capital stock per worker constant, new investment must be at least $(\delta + g_n)k_t$ – the δ covers the physical depreciation, while the g_n covers the increase in the number of workers.

As seen in the graph, there exists a steady state capital stock per worker. Just as in the case with no population growth, we can see that the economy will naturally tend towards that point on its own, given any initial starting value for the capital stock. For the Cobb-Douglas production function, we can algebraically solve for the steady state capital stock per worker as:

$$k^* = \left(\frac{sA}{\delta + g_n} \right)^{\frac{1}{1-\alpha}} \quad (18)$$

This is the same expression as we had before, just amended for population growth. We can also derive expressions for steady state output per worker and steady state consumption per worker. Again, these look almost the same as before, just with the new term related to population growth:

$$y^* = \left(\frac{sA}{\delta + g_n} \right)^{\frac{\alpha}{1-\alpha}} \quad (19)$$

$$c^* = (1 - s) \left(\frac{sA}{\delta + g_n} \right)^{\frac{\alpha}{1-\alpha}} \quad (20)$$

These expressions show us that, in a model with exogenous population growth, the economy will converge to a steady state in which per-capita variables (capital, consumption, and output) do not grow. In this steady state the levels of capital, consumption, and output must be growing if labor input is growing. To see this clearly, suppose that the economy has converged to a steady state by period t . This means that the capital stock per worker cannot be expected to grow between t and $t + 1$:

$$k_{t+1} = k_t \Rightarrow \frac{K_{t+1}}{N_{t+1}} = \frac{K_t}{N_t} \Rightarrow \frac{K_{t+1}}{K_t} = \frac{N_{t+1}}{N_t}$$

$\frac{K_{t+1}}{K_t}$ is just the gross growth rate of the capital stock, $1 + g_k$. Once we've hit the steady state, this expression says that the growth rate of the capital stock must be equal to the growth rate of population. That is, $g_k = g_n$. So, in the steady state in which capital per worker is not growing, it must be the case that capital is growing at the same rate as population.

What is happening to output in the steady state? We can do the same exercise as above:

$$y_{t+1} = y_t \Rightarrow \frac{Y_{t+1}}{N_{t+1}} = \frac{Y_t}{N_t} \Rightarrow \frac{Y_{t+1}}{Y_t} = \frac{N_{t+1}}{N_t}$$

So we see that output growth must also equal population growth, $g_y = g_n$. Since consumption is just a fixed fraction of output, consumption must also grow at the same rate: $g_c = g_n$. So our conclusion is that adding population growth to the model does not change any of its basic implications; all it does is to get the aggregate variables of the model to grow at the same rate as population. The model still does not generate growth in per capita variables.

4.2 Productivity Growth

Our previous twist was to add population growth to the model. That got us growth in levels of capital, output, and consumption. But there is still no growth in per capita variables.

Let's introduce a new variable, Z_t . This represents the level of “labor-augmenting technology.” Basically, it is a measure of productivity. It multiplies the labor input, N_t , to yield “efficiency units of labor,” $Z_t N_t$. We will assume that Z_t grows over time, in much the same way as N_t does. With Z_t growing, the intuitive idea is that workers are becoming more efficient – one worker in 1950 is the equivalent of, say, two workers in 1990, because workers in 1990 are more efficient.

The process for Z_t is:

$$Z_t = (1 + g_z)^t Z_{t-1} \quad (21)$$

Iterating this back to the “beginning” of time, and normalizing $Z_0 = 1$, we have:

$$Z_t = (1 + g_z)^t \quad (22)$$

The production technology is:

$$Y_t = AF(K_t, Z_t N_t) \quad (23)$$

What this is saying is that efficiency units of labor (or effective units of labor) is what enters the production function, not actual units of labor per-se. Note that Z_t and A play essentially the same role – an increase in either allows the economy to produce more output given capital and labor; they can thus both be interpreted as measures of productivity. The main difference between the two is that we are assuming/allowing Z_t to grow over time, while A has no trend growth (though one can entertain it moving around). Basically, one can think of Z_t as governing growth rates and A as governing the level of productivity (for example, if two economies have similar trend growth rates but different levels, this would be picked up by A – more on this below). For this reason A is sometimes referred to as “static efficiency,” implying that Z_t is a measure of “dynamic efficiency.”

The main Solow model equation is still the same, subject to this new addition:

$$K_{t+1} = sAF(K_t, Z_t N_t) + (1 - \delta)K_t \quad (24)$$

Last time we defined lower-case variables as per-capita variables. Now let's define lower case variables with a “hat” atop them as “per-efficiency units of labor” variables. That is, $\hat{k}_t = \frac{K_t}{Z_t N_t}$, $\hat{y}_t = \frac{Y_t}{Z_t N_t}$, $\hat{c}_t = \frac{C_t}{Z_t N_t}$.

Similarly to what we did before, divide both sides of this difference equation by $Z_t N_t$:

$$\frac{K_{t+1}}{Z_t N_t} = \frac{sAF(K_t, Z_t N_t)}{Z_t N_t} + (1 - \delta) \frac{K_t}{Z_t N_t}$$

Again, similarly to what we did last time, because of the constant returns to scale assumption $\frac{sAF(K_t, Z_t N_t)}{Z_t N_t} = sAF\left(\frac{K_t}{Z_t N_t}, \frac{Z_t N_t}{Z_t N_t}\right)$. As before, define $f(\hat{k}_t) = F(\hat{k}_t, 1)$. Using the newly defined

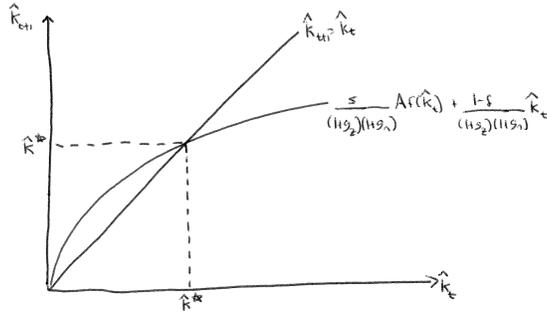
per-efficiency unit variables, we have:

$$\frac{K_{t+1}}{Z_t N_t} = sAf(\hat{k}_t) + (1 - \delta)\hat{k}_t$$

As before, we again need to manipulate the left hand side. Multiply and divide it by $Z_{t+1}N_{t+1}$ and simplify:

$$\begin{aligned} \frac{Z_{t+1}N_{t+1}}{Z_{t+1}N_{t+1}} \frac{K_{t+1}}{Z_t N_t} &= sAf(\hat{k}_t) + (1 - \delta)\hat{k}_t \\ \frac{Z_{t+1}N_{t+1}}{Z_t N_t} \frac{K_{t+1}}{Z_{t+1}N_{t+1}} &= sAf(\hat{k}_t) + (1 - \delta)\hat{k}_t \\ (1 + g_z)(1 + g_n)\hat{k}_{t+1} &= sAf(\hat{k}_t) + (1 - \delta)\hat{k}_t \\ \hat{k}_{t+1} &= \frac{s}{(1 + g_z)(1 + g_n)} Af(\hat{k}_t) + \frac{1 - \delta}{(1 + g_z)(1 + g_n)} \hat{k}_t \end{aligned} \quad (25)$$

We can graph this difference equation just as we have before, here putting \hat{k}_{t+1} on the vertical axis and \hat{k}_t on the horizontal axis. As before, \hat{k}_{t+1} is an increasing and concave function of \hat{k}_t , with the slope starting out large and flattening out to $\frac{1-\delta}{(1+g_z)(1+g_n)} < 1$. This means that the curve must cross the 45 degree line at one point, which we will call the steady state capital stock per effective worker.



Qualitatively this picture looks the same as the earlier cases; the math is just a little more complicated. Algebraically we can solve for the steady state values of capital, output, and consumption per effective worker just as we did before. We can do the same exercises we did before and reach the same conclusions. Increases in s or A lead to temporary bouts of higher than normal growth, but eventually we level off to a steady state in which per effective worker variables do not grow.

The main difference about this setup relative to the previous two is that there will be steady state growth in per capita variables. To see this, suppose that the economy has converged to a steady state in which the capital stock per efficiency unit of labor is not growing; that is, $\hat{k}_{t+1} = \hat{k}_t$. Manipulating this mathematically, we see:

$$\begin{aligned} \widehat{k}_{t+1} &= \widehat{k}_t \\ \frac{K_{t+1}}{Z_{t+1}N_{t+1}} &= \frac{K_t}{Z_{t+1}N_{t+1}} \\ \frac{K_{t+1}}{N_{t+1}} &= \frac{Z_{t+1}}{Z_t} \frac{K_t}{N_t} \\ \frac{k_{t+1}}{k_t} &= (1 + g_z) \\ \frac{K_{t+1}}{K_t} &= (1 + g_z)(1 + g_n) \end{aligned}$$

In other words, in the steady state capital per worker, k_t , grows at rate g_z , the rate of growth of labor augmenting technology. We can also see that output per worker, y_t , and consumption per worker, c_t , also grow at rate g_z . In contrast, the growth rate of the level of capital is equal to $(1 + g_z)(1 + g_n) - 1$, which is approximately $g_z + g_n$. The levels of output and consumption will also have this growth rate in the steady state.

Thus, this version of the Solow model predicts that per-capita variables will grow at constant rates in the “long run” (e.g. in the steady state). That is consistent with the data, but at some level it’s a disappointing result – these variables grow because we assumed that one of the variables we fed into the model grows (in particular Z_t). That does not seem like a great result – essentially we get growth by assuming it. The important result about the Solow model is actually a negative result – long run growth does not come from saving and capital accumulation. It comes from productivity growth. This negative result has important implications for policy that we will explore below.

As noted in the stylized facts, in the data real wages grow over time, whereas the return to capital (which is closely related to the real interest rate, as we will see later) does not. Put differently, economic growth seems to benefit labor not capital, which sometimes runs counter to popular perceptions. Does the Solow model match these predictions? Assume that the production function is Cobb-Douglas:

$$Y_t = AK_t^\alpha (Z_t N_t)^{1-\alpha}$$

The profit-maximization problem of the firm is similar to before, with the modification that now it is efficiency units of labor that enter the production function. The representative firm can choose labor and capital to maximize profits, taking the real wage and the real rental rate as given. Note that the firm takes Z_t as given – it cannot choose efficiency units of labor, just physical units of labor.

$$\max_{N_t, K_t} \Pi_t = AK_t^\alpha (Z_t N_t)^{1-\alpha} - w_t N_t - R_t K_t$$

The optimality conditions are:

$$w_t = (1 - \alpha)AK_t^\alpha Z_t^{1-\alpha} N_t^{-\alpha} \quad (26)$$

$$R_t = \alpha AK_t^{\alpha-1} Z_t^{1-\alpha} N_t^{1-\alpha} \quad (27)$$

We can re-write these conditions in terms of capital per efficiency units as:

$$w_t = (1 - \alpha)A \left(\frac{K_t}{Z_t N_t} \right)^\alpha Z_t = (1 - \alpha)A \widehat{k}_t^\alpha Z_t \quad (28)$$

$$R_t = \alpha A \left(\frac{K_t}{Z_t N_t} \right)^{\alpha-1} = \alpha A \widehat{k}_t^{\alpha-1} \quad (29)$$

Divide w_t and R_t by w_{t-1} and R_{t-1} , respectively:

$$\frac{w_t}{w_{t-1}} = \left(\frac{\widehat{k}_t}{\widehat{k}_{t-1}} \right)^\alpha \frac{Z_t}{Z_{t-1}} \quad (30)$$

$$\frac{R_t}{R_{t-1}} = \left(\frac{\widehat{k}_t}{\widehat{k}_{t-1}} \right)^{\alpha-1} \quad (31)$$

The left hand sides are just the gross growth rates of the real wage and the rental rate. If we evaluate these in steady state, we see that the wage will grow at the same rate as Z , with $g^w = g^z$. The rental rate will not grow. In other words, consistent with what we observe in the data, in the long run of the Solow model wages grow at the same rate as productivity (and hence also output, consumption, and investment), while the return to capital does not grow.

Before moving on, let us summarize the the predictions of the full blown model:

- Output per worker, capital per worker, and consumption per worker all grow at the same rate g_z in the steady state
- The real wage grows at the same rate, g_z in the steady state
- The return on capital does not grow in the steady state

These predictions are consistent with the time series stylized facts (1)-(4) that we mentioned earlier. The fact that the model matches the data means that we can be reasonably comfortable in using the model to draw policy conclusions, even in spite of the fact that we don't have a theory of where long run productivity growth comes from.

5 Quantitative Experiment

To see some of this in action, I am going to conduct a couple of quantitative experiments. By quantitative I mean that I am going to take the model as presented above, assign some numbers

to parameter values, and simulate out time paths of the variables under different scenarios.

For these exercises I am going to use the Cobb-Douglas production function. The parameters which need values are s , α , δ , A , g_n , and g_z . Let's start with α . With the Cobb-Douglas production function, $1 - \alpha$ will be equal to total payments to labor. In US data this amounts to somewhere between 60-70 percent of total income. So let's set $\alpha = 0.33$. Let's take the unit of time to be a year. In US data, real GDP grows at about 3 percent per year, with about 1 percent of this coming from population growth and about 2 percent coming from growth in real GDP per capita. So let's set $g_n = 0.01$ and $g_z = 0.02$. A ends up being uninteresting unless we are making cross-country comparisons (more on this below), so we can just set it to $A = 1$. Finally, in the US the share of investment in output is something around 15 percent. So let's set $s = 0.15$. Finally, $\delta = 0.1$ is consistent with the amount of depreciation that we observe at an annual frequency.

Written in per efficiency units of labor variables, the model is summarized by the following equations:

$$\begin{aligned}\widehat{k}_{t+1} &= \frac{1}{(1 + g_n)(1 + g_z)} \left(sA\widehat{k}_t^\alpha + (1 - \delta)\widehat{k}_t \right) \\ \widehat{y}_t &= A\widehat{k}_t^\alpha \\ \widehat{c}_t &= (1 - s)\widehat{y}_t \\ \widehat{i}_t &= s\widehat{y}_t\end{aligned}$$

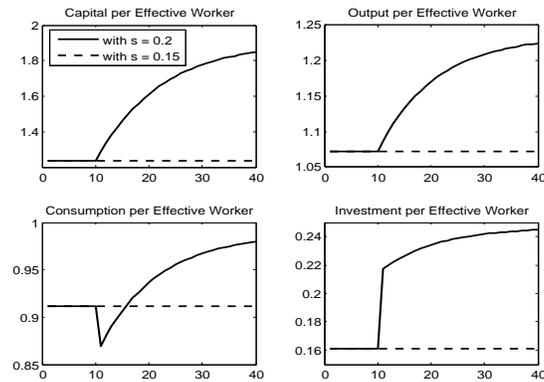
In terms of the levels of the variables, the model is summarized by:

$$\begin{aligned}K_{t+1} &= sAK_t^\alpha(Z_tN_t)^{1-\alpha} + (1 - \delta)K_t \\ Y_t &= AK_t^\alpha(Z_tN_t)^{1-\alpha} \\ C_t &= (1 - s)Y_t \\ I_t &= sY_t \\ Z_t &= (1 + g_z)^t \\ N_t &= (1 + g_n)^t\end{aligned}$$

In terms of the levels, I am normalizing $Z_0 = N_0 = 1$.

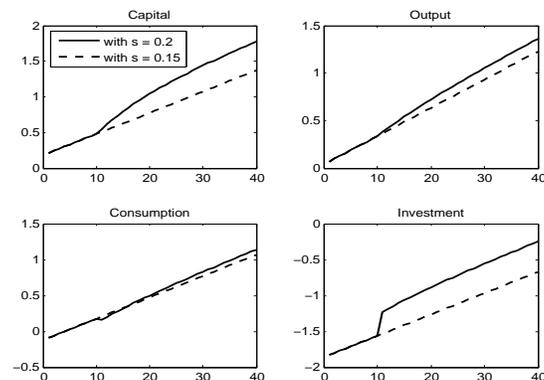
Given the parameterization of the model, we can solve for the steady state values of the per efficiency units of labor variables as: $\widehat{k}^* = 1.24$, $\widehat{y}^* = 1.07$, $\widehat{c}^* = 0.91$, and $\widehat{i}^* = 0.16$.

The experiment I consider is the following. At time 0 the economy sits in a steady state, meaning $\widehat{k}_0 = \widehat{k}^*$. The economy sits in the steady state for 10 periods. This means that the per efficiency unit variables remain constant at their steady state values from period 0 to 9, while the large variables all grow at rate $(1 + g_z)(1 + g_n) - 1$. In period 10, the saving rate increases from $s = 0.15$ to $s = 0.20$, and is expected to remain there forever. Below is a plot of how the per-efficiency unit variables evolve over time:



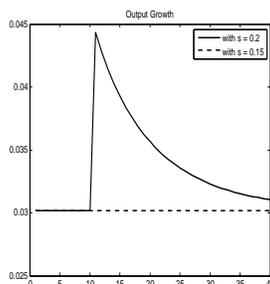
In the figure the dashed lines show what the time paths of the variables would have looked like had the saving rate remained at $s = 0.15$ – since the economy began in a steady state, it would have stayed in the steady state, so these lines are flat. The solid lines show what happens when the saving rate increases to $s = 0.2$. We see that consumption immediately jumps down, while investment immediately jumps up. Consumption remains below its pre-shock steady state for about 6 periods, after which time it is higher than where it began. Put differently, since consumption eventually increases following the increase in s , we are below the Golden Rule. The capital stock, and hence also output, does not react immediately within period, but grows for an extended period of time, smoothly approaching a new, higher steady state. Note that the units of time here are years. Hence, convergence to the new steady state takes a while – after 30 years, it’s still not there.

Now in reality, we don’t really care about the per efficiency units of labor variables – that’s just a construct to help us analyze the model. What we really care about are the levels of variables. Below I plot the log levels of the four variables. I plot these in the log because it makes the picture easier to read – a variable growing at a constant rate looks linear in the log, but exponential in the level. The dashed line shows the path that would have obtained had we remained at a saving rate of $s = 0.15$, whereas the solid line shows the response when s increases to 0.20 in period 10.



Because we are plotting the levels here, all of the variables would have kept growing without

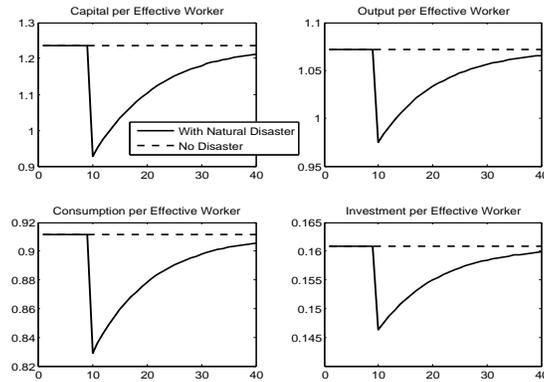
the change in the saving rate (the dashed line). The increase in the saving rate makes them grow faster (for a while), which means that the variables end up on a permanently higher trajectory. In particular, output ends up 15 percent higher after a while than it otherwise would have been had we kept the lower saving rate. Consumption ends up about 8 percent higher eventually.



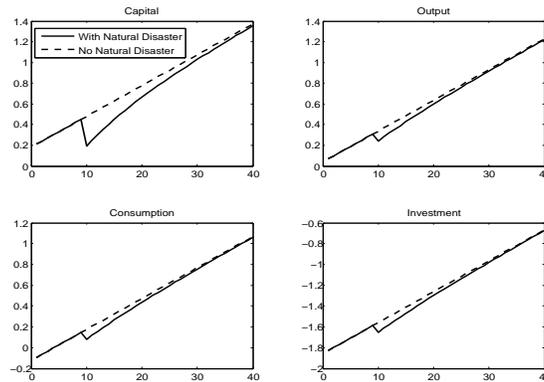
The figure above plots the growth rate of output, both pre and post the change in the saving rate, as well for the trajectory the growth rate would have taken had the saving rate remained at $s = 0.15$ (dashed line). In the steady state, the growth rate is about 0.03, or 3 percent ($(1 + g_z)(1 + g_n) - 1 \approx g_z + g_n = 0.03$). In the period after the saving rate increases, output growth jumps up, from 3 percent to close to 4.5 percent. It then starts to come down, but it remains elevated for a long time, though it eventually returns to its pre-shock level. As we argued qualitatively, here we see that quantitatively an increase in the saving rate leads to a permanently higher level of output; to get there the economy has to grow faster, but only for a while. In the long run the growth rate is independent of the saving rate.

Next I consider a different quantitative experiment, using the same model and basic parameterization. I assume that the economy is in a steady state from periods 0 to 9 (with saving rate $s = 0.15$). In period 10 it is hit with a “natural disaster.” The natural disaster destroys 25 percent of its capital stock in that period. A real life example might be something like a hurricane (e.g. Hurricane Katrina wiped out a lot of refineries, factories, etc., on the Gulf Coast). You could also think about this as a war – WWII destroyed a lot of the physical capital stock in Germany, Britain, and Japan.

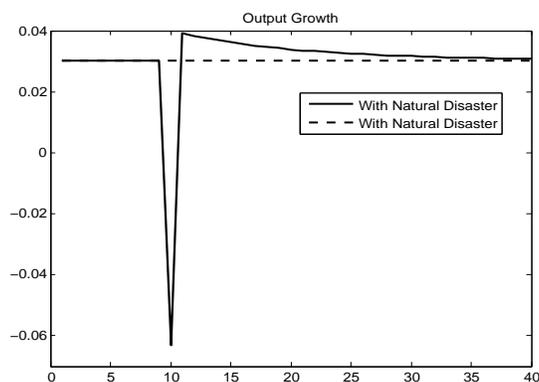
What I do is to trace out the implications of this natural disaster for the per efficiency units variables. Effectively what happens is we go from sitting in a steady state in period 9 to being far below the steady state in period 10. The reduction in the capital stock does not change the steady state capital stock – it just knocks the existing capital stock well below that steady state. When you start out below the steady state the variables have to grow faster than normal to converge



We see exactly that simple intuitive pattern in these responses. All of the variables jump down immediately when the capital stock gets destroyed – output is a function of current \hat{k} , so it must go down; and since investment and consumption are just fixed fractions of output, they also jump down. But after that initial downward jump, they all are going to recover to approach the steady state in which they began. This necessitates that these variables grow faster than normal for an extended period as they approach the original steady state from “below.”



The figure above plots the log levels of the variables, just as I did in the case of a change in the saving rate. The growth is interrupted in period 10, the period of the disaster, but then we can see the levels of the variables growing faster to “catch up” the trend line that would have obtained had the disaster not occurred.



The figure above plots the response of the growth rate of output to the disaster shock. As in the earlier case, the steady state growth rate of output is about 3 percent per year. This growth rate goes sharply negative (to about -7 percent) in the year of the shock, but then it immediately jump up higher than where it was before (to about 4 percent). It then remains elevated for a number of years after the disaster. It is this faster growth that eventually takes the levels back to their pre-shock paths.

6 Why Are Some Countries Rich, and Others Poor?

We've seen that the Solow model does pretty well at the time series stylized facts. How well does it do at matching the cross-sectional facts? In particular, what can the Solow model say about why there are such large differences in standards of living across countries? What can be done to lift poorer countries out of poverty?

As we have seen, the Solow model predicts that countries should converge to steady states in which per efficiency units of variables do not change and in which per capita and levels versions of those variables grow at constant rates. An immediate implication of this is that, if all the parameters of the model are the same, the Solow model predicts that there should be *convergence*: every economy should eventually look the same in the levels. The only reason we'd ever observe any difference in levels is if, by happenstance, some countries were initially endowed with more or less capital (and hence closer to or further away from their steady states). But eventually they'd all end up looking the same.

This convergence property is pretty clearly rejected in the data, at least in an unconditional sense. There are some instances in which the convergence prediction is more or less borne out and some where it fails. In the last section we studied how an exogenous reduction in the capital stock would predict faster growth. We actually do see that for Germany and Japan in the wake of World War II. From 1948 to 1972, output per capita grew more than 8 percent per year in Japan and almost 6 percent per year in Germany, compared with 2 percent in the US. This makes sense in light of the Solow model – Japan and Germany experienced a reduction in their capital stocks, which

put them far below their steady states. The US did not. Japan and Germany had to grow faster to catch back up to their steady states. So there is some evidence in the data that, if countries are initially endowed with “too little” capital relative to the steady state, they will grow faster to converge to the steady state.¹

An interesting, and ultimately revealing, fact is that, though Japan and Germany grew faster than the US from 1948 to 1972, they did not from 1972 on. In fact, relative to per capita GDP in the US, Japan and Germany are now both about where they were 40 years ago (per capita GDP in both countries is about 70-80 percent of what it is in the US, with this ratio fairly stable since then). In other words, if you look at the data in the world, there appears to be convergence, but it is *conditional* convergence – countries seem to converge to their *own* steady states. Countries which for obvious reason (war) start out with little capital grow faster than average for a while, and eventually approach their own steady state. But as we see in the case of Germany and Japan, these steady states evidently differ across countries, and any cursory look at the data shows that the differences are very large.

A good data source are the Penn World Tables, which have annual data on per capita real GDP relative to the United States for almost 200 countries. We’ve already seen that Japan and Germany have had per capita GDP of about 75 percent of the US, and that ratio has been relatively stable. This suggests that their steady state levels of output are about 75 percent of the United States. Table 1 shows per capita GDP in 1970 relative to the US and again in 2010 for a (psuedo-random) selection of countries.

There are some very interesting patterns that pop out of this table. First, there are a lot of countries for which the ratio is about the same in 2010 as it was in 1970 – this suggests that these countries and the US had both more or less reached their steady states by 1970 and have grown at common rates since, though the ratios show very wide income disparities, suggesting that those steady states are not the same. These countries are primarily in Europe (Germany, Denmark, France, Spain) and South America (Bolivia, Brazil, Ecuador). There are some growth “disasters” where countries got significantly worse relative to the US – these countries tend to be in Africa (Ghana, Liberia, Zimbabwe), with Barbados another interesting outlier in this regard. Finally, there are some growth “miracles” whereby countries have grown enormously relative to the US – these countries tend to be in Asia (Hong Kong, South Korea, Singapore). China is another good example, but we don’t really have reliable data on them, particularly going back to 1970. An interesting and pressing question is whether this fast growth in Asia will continue. One interpretation is that, when these countries began opening up to the western world, they had too little capital. Their institutions and culture are conducive to strong economies, and so capital has been flowing in. If accumulation of capital is the source of the growth, then we would expect this faster than world-average growth to disappear – these countries are just converging to their steady states from a position of low capital, similarly to Japan and Germany after WWII. An alternative interpretation is that these economies are fundamentally different from the US and the rest of the

¹The numbers in this paragraph are adopted from the case study in Chapter 7 of Greg Mankiw’s *Macroeconomics*, 7th edition.

Table 1: GDP per Capita Relative to the United States

Country	Relative GDP in 1970	Relative GDP in 2010
Algeria	13.6	15.5
Barbados	135.7	63.8
Bolivia	13.3	9.5
Brazil	18.9	20.9
Cambodia	4.8	5.3
Denmark	81.8	83.2
Ecuador	15.6	15.8
France	77.5	75.6
Ghana	9.2	4.9
Hong Kong	32.2	90.0
Jamaica	40.7	20.8
South Korea	13.0	61.8
Liberia	7.5	1.1
Portugal	36.6	48.5
Singapore	31.8	128.0
Spain	57.1	66.1
Sudan	5.2	5.5
Taiwan	18.3	69.4
Zimbabwe	1.6	0.8

Notes: The numbers in this table are the ratio of real per capita GDP in a country to the United States times 100.

world, and may continue to experience faster productivity growth due to large g_z . Only time will tell, but some existing research shows that much of the recent growth in the Asian countries has come from capital accumulation. This would suggest that these countries will be slowing down in the near future.

The other main thing that pops out of the table is the following: not only are there apparent differences in steady state output per capita, these differences can be quite large. For example, US GDP per capita is about 100 times bigger than Zimbabwe and Liberia, and about 20 times bigger than most other African and South American countries. What parameter(s) in the Solow model could account for such large differences in standards of living?

For this exercise we can abstract from either of the two sources of trend growth we considered. For countries where the ratio to the US has stayed pretty constant, it is clear that trend growth is more or less the same, for the purposes of comparison we can avoid looking at that. For the base model which abstracts from trend growth, steady state output per worker is:

$$Y^* = A \left(\frac{sA}{\delta} \right)^{\frac{\alpha}{1-\alpha}}$$

Take natural logs of this:

$$\ln Y^* = \frac{\alpha}{1-\alpha} \ln s + \frac{1}{1-\alpha} \ln A - \frac{\alpha}{1-\alpha} \ln \delta$$

Let's suppose that a country has GDP per capita that is 20 percent of the US. This would mean that $\frac{y^*}{y_{US}^*} = 0.2$. Taking logs, we'd have $\ln y^* - \ln y_{US}^* = -1.6$. Let's assume that the countries have identical α and δ , but allow s and A to be different. Differencing the above expression, we'd have:

$$\ln Y^* - \ln Y_{US}^* = \frac{\alpha}{1-\alpha} (\ln s - \ln s_{US}) + \frac{1}{1-\alpha} (\ln A - \ln A_{US})$$

We can re-write the above expression to be:

$$\frac{1-\alpha}{\alpha} (\ln Y^* - \ln Y_{US}^*) = (\ln s - \ln s_{US}) + \frac{1}{\alpha} (\ln A - \ln A_{US})$$

A reasonable value for α in the data is $1/3$, which would mean that $\frac{1-\alpha}{\alpha} = 2$. If we're looking at a country with one-fifth (20 percent) of US per capita GDP, this means that the left hand side of this expression would be $-3.2 = 2 \times 1.6$. So:

$$-3.2 = (\ln s - \ln s_{US}) + 3(\ln A - \ln A_{US})$$

Now let's suppose that the A s are the same in both countries, and focus just on differences in saving rates. Suppose that the US saving rate is 0.15. Then we have $\ln 0.15 = -1.9$. Re-arranging and solving for the required saving rate in the poor country, we get:

$$-5.1 = \ln s$$

$$s = \exp(-5.1) = 0.006$$

In other words, if the only thing different between two countries is the saving rate, and the US has a saving rate of 0.15, you'd have to have a saving rate of 0.006 in the poor country to account for that country having GDP per capita that is 20 percent of the US. In other words, you'd have to have a 25 fold difference in the saving rates ($0.15/0.006 = 25$) to account for a 5 fold difference in levels of real GDP per capita (one country having a level that is 20 percent of the other is a factor of five difference). We simply don't observe disparities in saving rates that are anywhere near that large to be able to account for the large differences in output per capita.

This exercise suggests that the only thing that can account for large disparities in levels of output per worker is differences in A . Economists sometimes refer to this variable A as "static efficiency." It's essentially a measure of productivity – the bigger is A , the more a country can produce for given levels of capital and labor. What is a little disappointing is that the model says nothing of what A really is or where it comes from – it's an exogenous variable in the model. In a sense A is a residual – it explains the part of output that we can't explain with observable inputs. In the business cycle literature A is often referred to as "total factor productivity," TFP.

Because of the central importance of A in accounting for differences in standards of living across the world, economists have devoted a lot of attention to potential sources of what A is. One thing that immediately comes to mind is something like "knowledge." There is some truth to it, but simple introspection reveals that A has to encompass a lot more than that – knowledge is pretty easy to transfer, and people in most of the rest of the world at least have access to much of the same knowledge that people in the US do, especially with the internet and improved global communications. One thing economists have emphasized and which seems to have some empirical plausibility is climate. As a general rule (there are, of course, exceptions), the closer countries are to the equator the poorer they are. Climates near the equator are generally hot and muggy. Aside from being uncomfortable and hence making it difficult for people to focus, disease also tends to thrive in these climates, which is also going to hurt productivity. Geography is something else that matters – for example, a country with lots of natural waterways makes transport of goods and services easier, and transport facilitates trade, which leads to gains from specialization. Think about the Nile river in ancient Egypt, or the Mississippi in the US. Another example is the terrain and how easy it is to traverse – think about a very mountainous country like Afghanistan. Another critical factor to which economists have pointed is "institutions" broadly defined. Specifically, countries with stable governments, the rule of law, property protection, and stable social norms (i.e. tribes don't go looking to behead each other) tend to be richer. Another factor, closely related to institutions, is infrastructure. Countries with better roads, running water, electricity, etc. are richer. Since these are goods typically provided by governments, infrastructure sometimes falls under the institutions category.

In terms of trying to lift countries out of poverty, there is only so much one can do. Two of the key factors mentioned above – climate and geography – cannot be changed. Though a bit of an exaggeration, there is some truth to a statement like “The best way to help poor people in Africa is to get them out of Africa.” Terrain and climate there are just not as conducive to economic activity as in other locations, so these countries can probably never be as well off as some others. The biggest thing that can be changed is institutions. If you can get countries to adopt stable democracies without corruption, which have well defined legal rules which protect property and don’t impose exorbitant taxes on entrepreneurship, then you can help lift these countries out of poverty. Unfortunately, and unsurprisingly, you tend to not observe such stable institutions in these countries. Policies to promote free trade (in people, goods, and ideas) will lead to more knowledge diffusion and should also help.

In addition to pointing out what kinds of things might help improve standards of living in poor countries, the Solow model also has implications for what things are not likely to work. In particular, very poor countries are *not* poor because they don’t have enough capital. We can infer this for two reasons. First of all, we don’t see convergence in many countries – see the table above. Second, if countries were just poor because they didn’t have enough capital, there would be large arbitrage (e.g. profit) opportunities to move capital to those countries (if capital is low relative to its steady state, then the marginal product of capital, and hence the return on capital, should be high). We also don’t see that. Rather, very poor countries are poor because they have low A . This means that “aid,” broadly defined, is not likely to lift these countries out of poverty. For example, you could imagine shipping 1000 computers to the Congo. While this seems like a benevolent idea, it’s not likely to do much – there is poor electricity and internet access, the people there may not have the education to know how to use the computers, and the rule of law is such that there is a high probability that the machines would be stolen. You can even extrapolate some of these lessons to poverty reduction in an advanced country like the US. Are people poor because they don’t have enough capital? Or is it because they have poor education and poor social structures? As in the Solow model, there is a lot of evidence pointing to the latter. To lift people out of poverty we need to improve institutions.

7 Beyond the Solow Model

The Solow Model is a widely used model in economics that makes a number of important insights. We’ve discussed two at length. First, sustained growth in GDP per capita does not come from higher saving rates. Rather, growth must come from improvements in productivity. Second, income disparities around the world cannot be explained by differences in saving rates (or more broadly by factor accumulation, which is directly related to the saving rate). This fact has important implications for lifting poor countries out of poverty.

A significant drawback of the Solow model is that it does not explain where growth comes from – rather, it takes growth in productivity as exogenous. A recent literature has tried to “endogenize” growth in the model. We will not study it here, but models of this sort feature increasing returns,

research and development, and/or other similar features. Even though we are not modeling where growth in Z comes from (or where levels of A come from), we can use some common sense and intuition to think about policies that would promote growth. Here is a small listing:

1. Patent protection. Patents play competing roles in economics, as they essentially serve as government granted monopolies, which economists typically don't like. The reason we issue patents is to encourage innovation in the first place – innovators would have little incentive to come up with new inventions if they weren't going to be able to extract some higher than average returns should those projects work out. So strong patent protection should encourage innovation, which should help facilitate growth.
2. Free trade. In microeconomics you often talk about the gains from specialization. More trade allows for greater specialization and hence more productivity.
3. Education. As mentioned earlier, education or “knowledge” is something that can and does impact economic growth and standards of living. Sometimes the accumulation of knowledge is considered as a third factor of production – human capital. Knowledge is like capital in the sense that you have to “produce” it (you have to go to school to acquire skills), and that it does not completely depreciate within period. There are different ways of encouraging education, some of which we see governments doing, and some which are more effective than others.
4. Subsidize research and development. Patent protection is a form of this kind of subsidy. Other ways to accomplish this same goal are somewhat more direct – the government can explicitly fund research in the sciences. This is what the National Science Foundation and the National Institute for Health do, for example. Also, the government running educational institutes like colleges and universities both improves education, but also subsidizes research.
5. Infrastructure. Public capital takes the form of things like roads, bridges, running water, etc. These are all things which make private sector participants more productive. Solid physical infrastructure, much like rule of law and stable political institutions, will promote growth.

The Solow model has the important implication that sustained growth must come from productivity, with factor accumulation resulting from higher saving rates not being able to do the trick. As noted above, that does not mean that raising saving rates would not be a good idea. As we have seen, higher saving rates will raise the *level* of output in the long run, and may raise the level of consumption if we are below the Golden Rule (which all the empirical evidence suggests we are). Differences in saving rates cannot account for 20 fold differences in standards of living that we observe in a cross-section of countries, but they can make fairly sizable differences in welfare within a country – in the quantitative section, we saw that increasing the saving rate from from 15 percent to 20 percent could increase steady state output per worker by about 15 percent, which is large. I think that it is fair to say that most economists would agree that the US saves too little, and believe that policies which promote saving would be beneficial.

The following is a partial listing of the kinds of policies that might result in more saving:

1. Get budget deficits under control. It's not a feature of the model as we've laid it out already, but governments that run deficits reduce the economy's effective saving rate. Consider the accounting identity for a closed economy with a government: $Y_t = C_t + I_t + G_t$. Let T denote total tax payments, paid for by households. We can re-write this identity as: $Y_t = C_t + T_t + I_t + G_t - T_t$, where the equality does not change because we just added and subtracted something from the right hand side. Private saving is: $Y_t - T_t - C_t$, disposable income (income less taxes) minus consumption. Public saving is $T_t - G_t$, basically tax revenues less expenses. Investment is equal to: $I_t = S_t^{PR} + S_t^G$, where $S_t^G = T_t - G_t$, or public saving. Budget deficits reduce total saving in an economy, and will work to prevent saving and factor accumulation.
2. Use the tax system to encourage saving. Most personal income taxes in this country are related to income. At the margin taxes on income reduce the incentive to earn income. Switching to a consumption based tax would reduce that adverse incentive, as well as encouraging saving by dis-incentivizing consumption. The drawback of such policies is that they would likely disproportionately hurt the poor (i.e., on its surface a consumption tax is regressive), though there are potential ways around that. Other policies, like preferential tax treatment of income from savings – like capital gains, interest, and dividends – will also work to stimulate private consumption.