Notes on the Solow growth model
Principles of Macroeconomics

Spring 2025

This is a (hard copy) version of the Solow model notes available at `https://people.carleton.edu/~estruby/solow/`. The electronic version has some embedded links, color, and nicer formatting, but this version has blanks for the exercises you should work out.

1 Introduction: What facts do we want to explain?

Before we jump into the model, we should think about what economic outcomes the model is attempting to make predictions about.

We most often compare countries in per capita" (per person) terms. When economists talk about a countries' level of "economic development" we usually mean their level of GDP per capita.

However, the Solow model really makes predictions in per-worker terms. Countries with more people usually have more workers (but obviously, this could depend on labor force participation and demography).

Economists who have examined this data have found a set of "stylized facts"[1] that are reasonably consistent across countries since around the early 19th century (around the time of the Industrial Revolution). A partial list of these facts we would hope a model of long-run growth is consistent with:

1. There is a wide distribution of GDP/person or GDP/worker across countries.
2. The countries that are richer than others are similar in several dimensions:
   - They tend to have more physical capital per worker.
   - They save and invest at a higher rate than other countries.
   - Their workers tend to have a higher level of education and are healthier (they have more human capital).
   - They seem better at transforming physical inputs into output than poorer countries (they are more productive).
3. People who live in countries that are richer consume more on average.

There are some other facts - which are perhaps more subtle - that we might want to try to understand or be able to explain in light of the Solow model. Two of these are:

4. Rich countries tend to grow at similar rates in the long run.
5. Countries at a similar level of economic development often have similar geographic, institutional, and cultural characteristics, and also similar policies.

Economists have found these facts to be generally true for the post-Industrial-Revolution period using a variety of approaches and datasets. In the rest of this section, we'll illustrate these facts using data from the Penn World Tables. This is the most commonly used dataset for examining long-run growth across countries in economics.

[1] A stylized fact is, more or less, something that is broadly true. The term is common in economics, and specifically originates from debates over macroeconomic growth theory. Nicholas Kaldor argued that theorists "should be free to start off with a stylised view of the facts – i.e. concentrate on broad tendencies, ignoring individual detail." So while these facts are not true for every country at all times, they are broadly what we've come to expect.

1.1 Fact 1: Some countries are richer than others on a per-capita basis

There is an enormous degree of inequality across countries.[2] In 2019, the richest country in the world (on a per capita basis) was Qatar, with a real GDP per capita more than 125 greater than that of the poorest in the world (the Central African Republic). Many of the very richest countries in the world are really more like extremely wealthy city-states (like Macau and Singapore) or have large amounts of oil resources (such as Qatar and or Norway), or are financial centers (Ireland, Switzerland).

Table 1: 15 countries with least GDP per capita (millions of 2017 USD) in 2019.

Country	Real GDP	Population (millions)	Real GDP per capita
Venezuela (Bolivarian Republic of)	7160	29	251
Burundi	9110	12	790
Central African Republic	4642	5	978
D.R. of the Congo	88673	87	1022
Malawi	21635	19	1161
Niger	28226	23	1211
Mozambique	37316	30	1229
Liberia	6212	5	1258
Madagascar	41507	27	1539
Haiti	17503	11	1554
Chad	25756	16	1615
Yemen	51828	29	1777
Guinea-Bissau	3554	2	1850
Sierra Leone	14651	8	1875
Uganda	92619	44	2092

Table 2: 15 countries with most GDP per capita (millions of 2017 USD) in 2019.

Country	Real GDP	Population (millions)	Real GDP per capita
Australia	1364678	25	54147
China, Hong Kong SAR	407576	7	54810
Netherlands	950078	17	55569
Cayman Islands	4004	0	61651
Kuwait	261069	4	62055
United States	20595844	329	62589
United Arab Emirates	645956	10	66113
Brunei Darussalam	31738	0	73249
Norway	396254	5	73669
Switzerland	646920	9	75299
Singapore	477908	6	82336
Luxembourg	55711	1	90479
China, Macao SAR	59874	1	93488
Ireland	501054	5	102622
Qatar	323141	3	114101

Of course, looking at the extremes only tells us so much. Most countries fall somewhere in the middle, although the data has a heavy left tail (there are more very, very poor countries than there are very very rich ones). The histogram below illustrates the distribution. (Notice that the x axis has a log scale!)

[2] All of these comparisons are using real GDP, as measured using 2017 US dollars adjusted for purchasing power parity (roughly, adjusting for differences in price levels across countries).

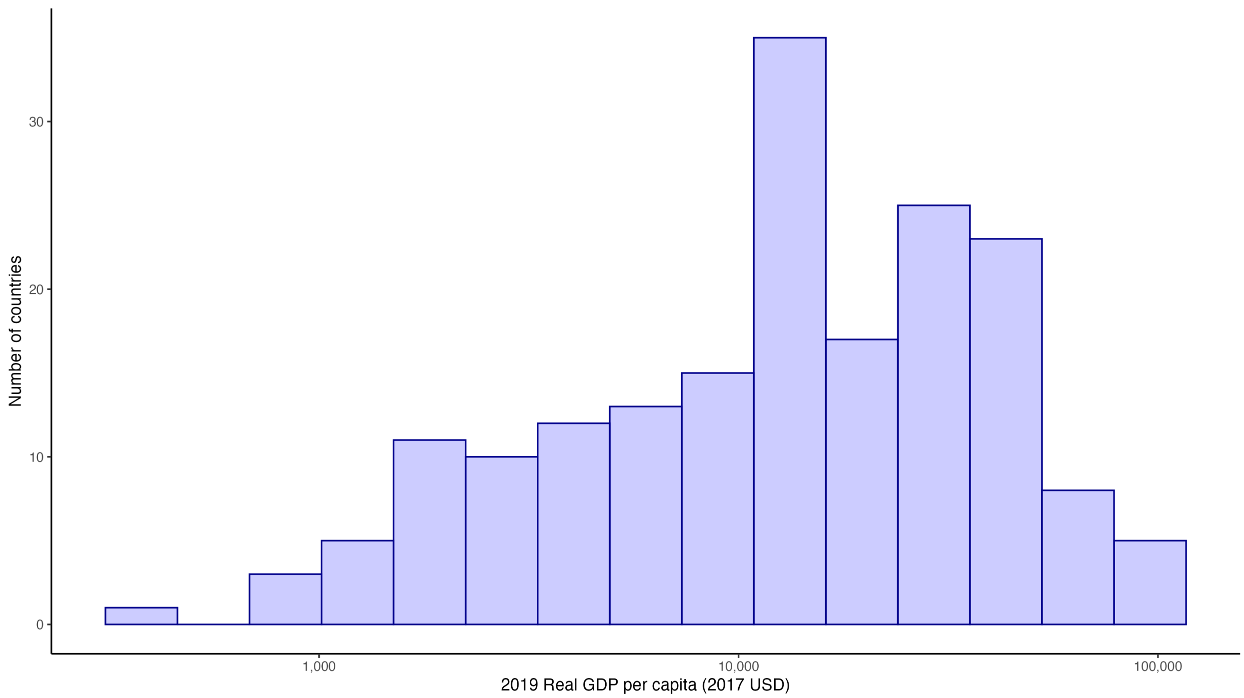

One thing that's important to remember is that looking at the distribution of wealth in a given year doesn't show how the distribution has changed over time. In particular, there has been an enormous amount of growth in GDP per capita over the past several decades in most countries. That can be seen by the change in distribution over time. This can be seen in the histogram below, which shows the same graph as above in blue, but with the distribution of real GDP per capita in 1960 (the dashed outline with red fill). Keeping in mind that this is real GDP, this means that (adjusting for prices), the quantity of output per person has grown dramatically over the past several decades almost everywhere.

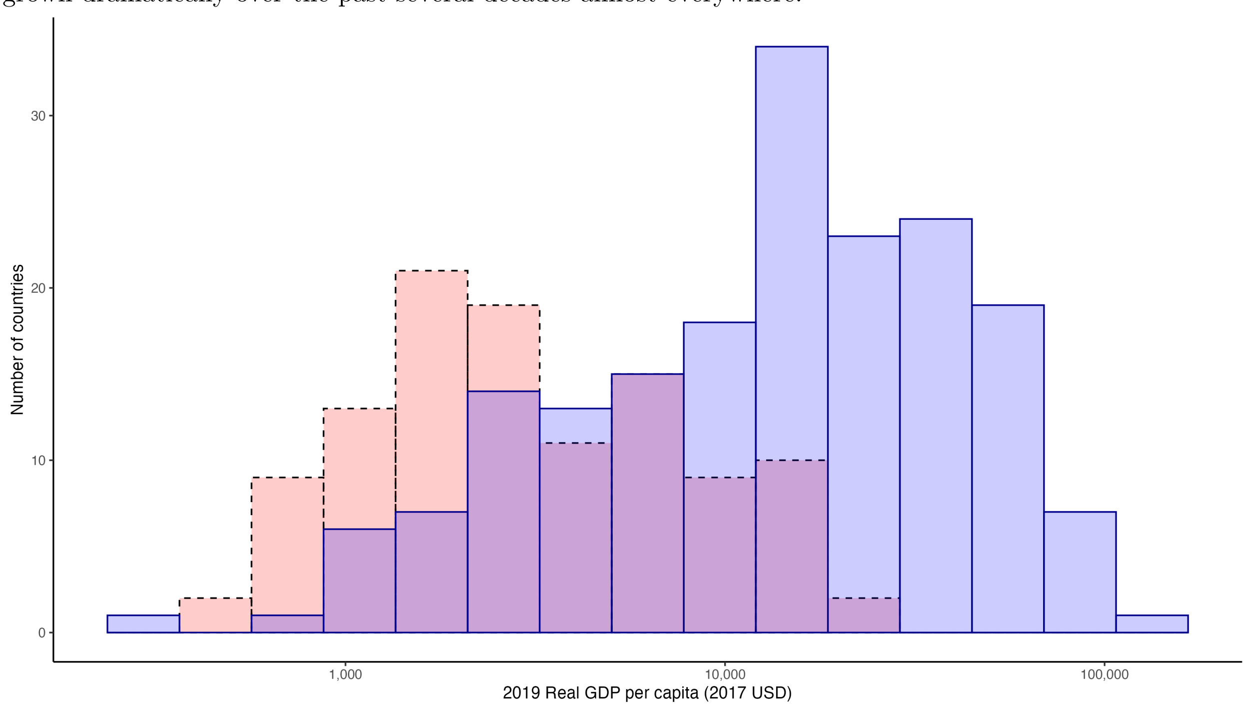

We can also track the evolution of the distribution over decades to get a sense of how it's shifted over time. The distributions below show the distribution at the start of each decade since 1950, with 2019 at the end. The vertical lines indicate quantiles – to the left of the first line is the lowest 25% of the distribution, the middle line indicates the median, and to the right of the last line is the upper 25% of the distribution. The country at the 25th percentile of the distribution in 2019 has a GDP/worker (in real terms) that is a bit higher than the median country in 1950.

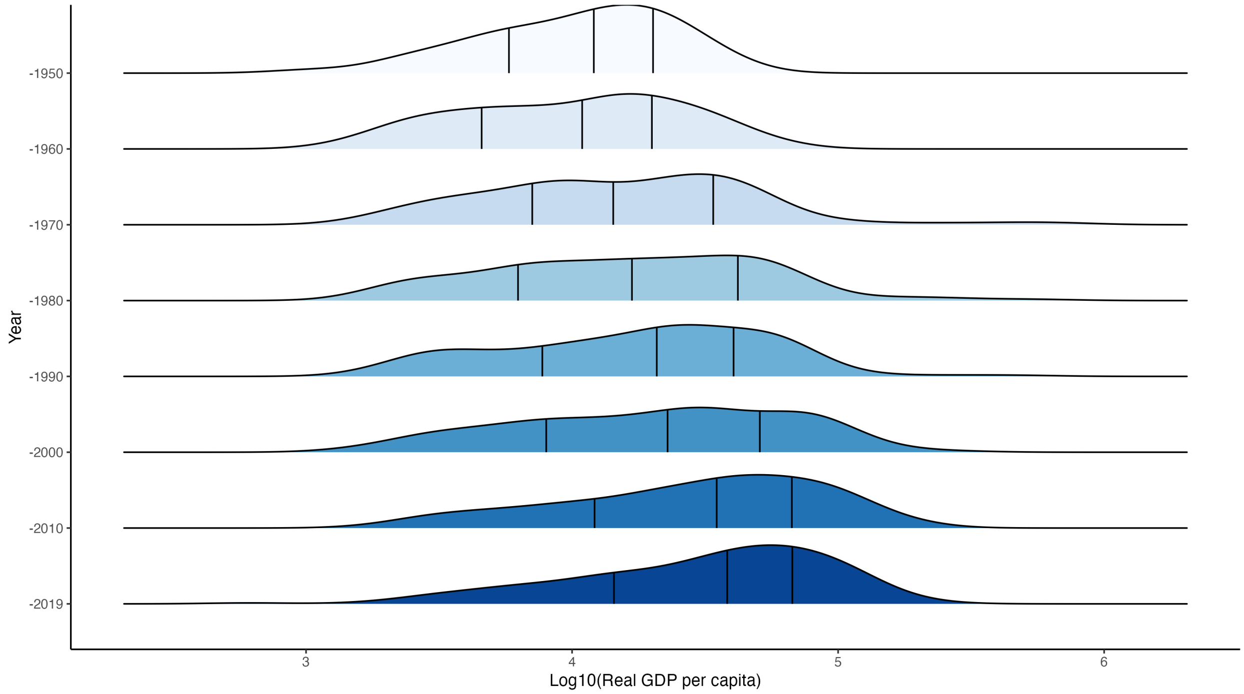

We can also think about the number of people living in countries at different levels of development. Many fewer people live in countries that are extremely poor on average. This can be seen in the next plot, where countries are weighted by their population. The punchline is that there's been a shift since in 1960, where a huge mass of the world population lived in countries where real GDP per capita was roughly 1000, but now much of that has shifted to places where it's closer to 10,000. (And very, very few people live in countries where the GDP per capita is close to 100,000 dollars).[3]

[3]Of course, this doesn't say anything about the distribution of income within countries.

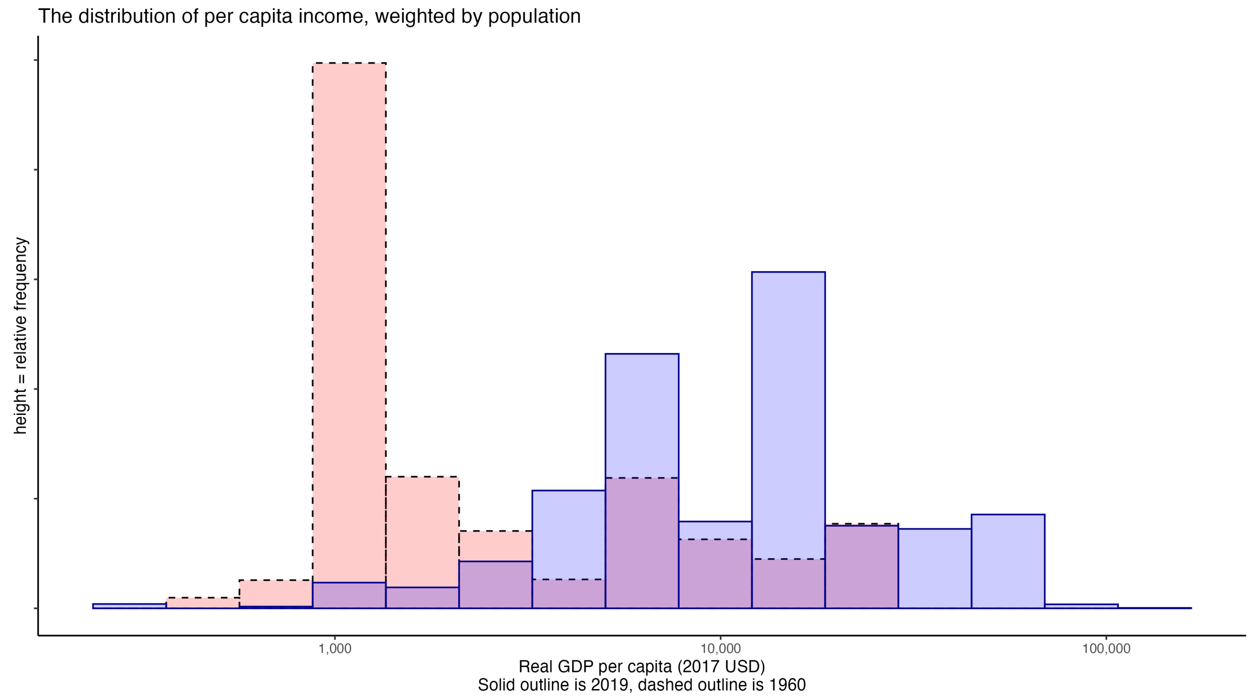

1.2 Fact 2: Countries that are richer (per capita) tend to have more physical capital per worker,save and invest at a higher rate, have a higher level of human capital, and are more productive

> Happy families are all alike; every unhappy family is unhappy in its own way.
>
> Tolstoy, *Anna Karenina*

Economists have noted that countries that have higher GDP per capita also tend to be similar along several dimensions. To an extent, this is unsurprising; if we think about real GDP as the amount of production occurring within a country in a given period of time, then those countries must either be using more of the factors of production or using the existing factors more effectively. Solow was focused on post-Industrial Revolution growth when investment in physical capital (think machinery, factories, and so on) was important, and it turns out those are still correlated with high levels of GDP per capita.

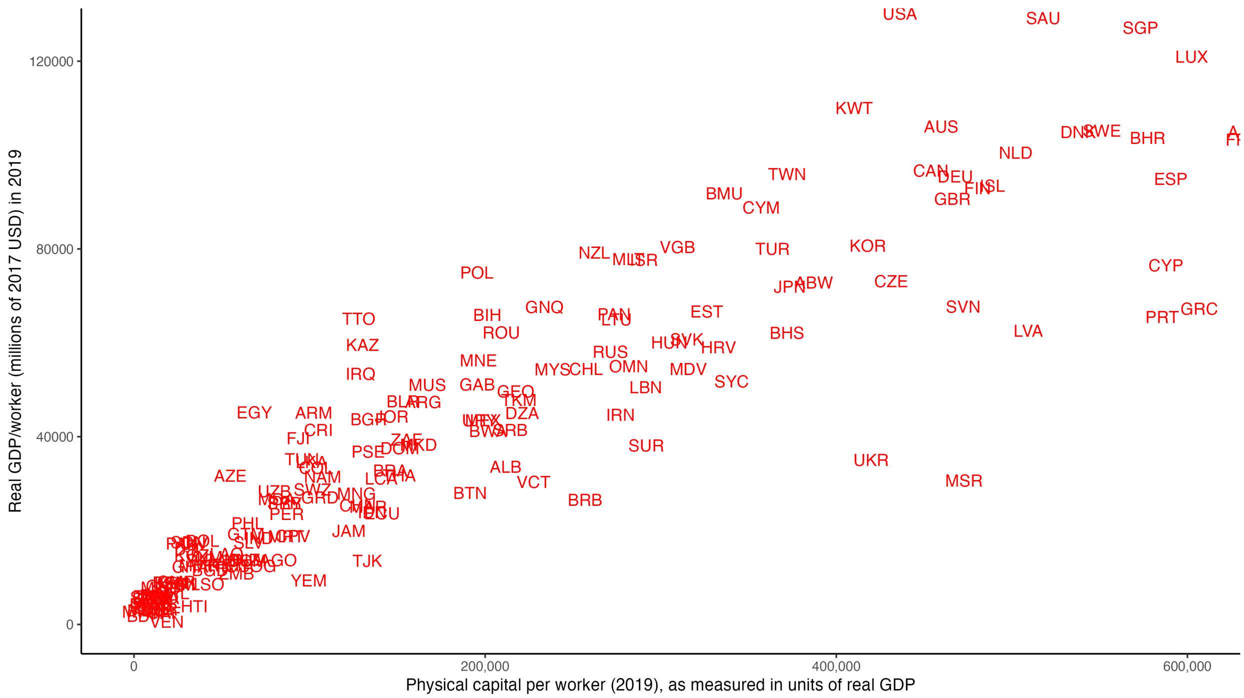

A key feature of physical capital is that it is itself produced, through investment. Investment is financed by saving. Although not all investment is financed by domestic saving (and not all investment is investment in physical capital), the rate of saving (how much output that isn't consumed) is also positively correlated with GDP per capita.

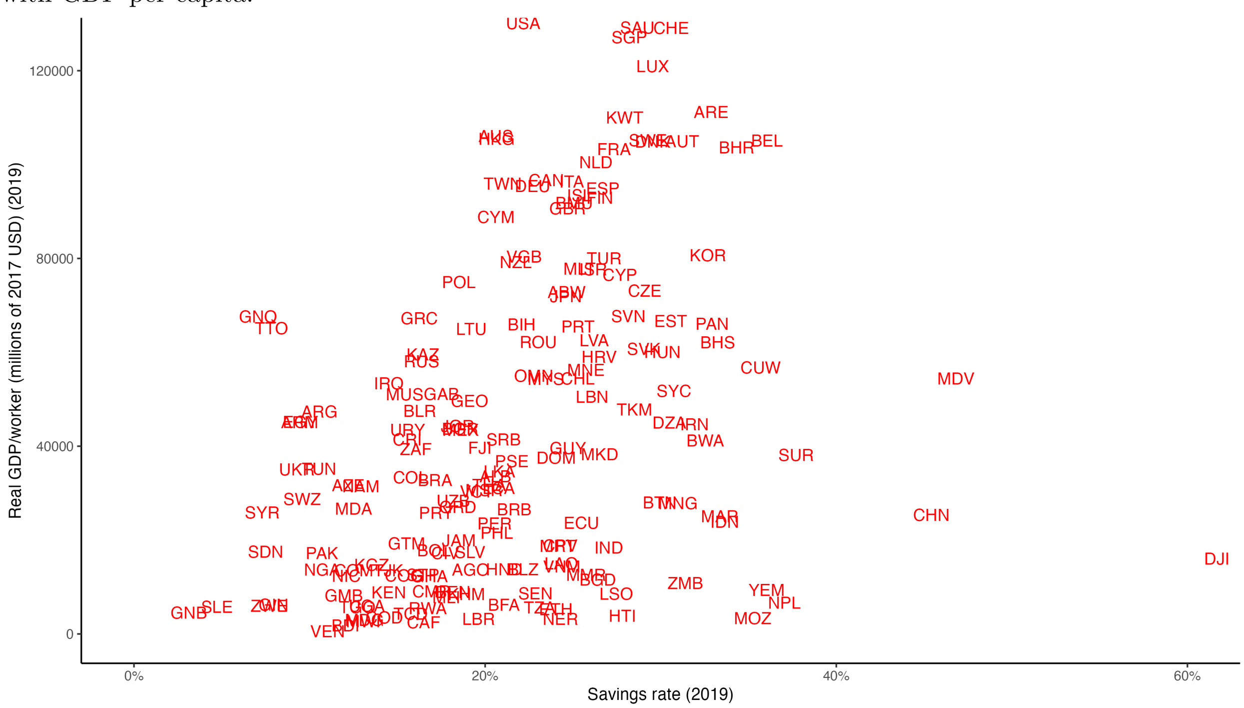

Human capital is the knowledge, skills workers acquire through education, training, and experience, as well as things like their health. The most obvious - although by no means most complete - way of measuring human capital is through measuring the education of workers (and how that education translates into the ability to produce things in a given amount of time, and earn higher wages). The scatterplot below shows that countries that are richer tend to also have higher measured values of human capital.

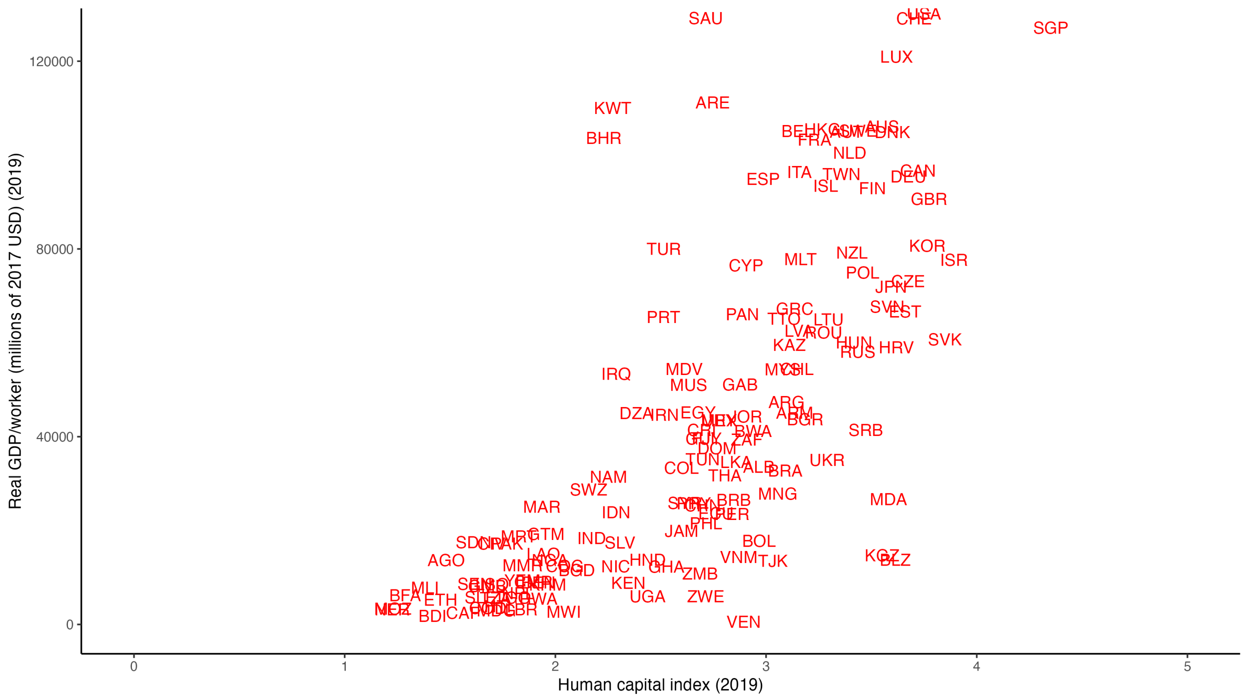

Even after we account for physical and human capital, workers in some countries generate more GDP (on average) than other workers; that is to say, they are more productive. Put another way workers in different countries appear to produce different amounts, even with the same measured amount of physical machinery to work with and education.

Some of this is due to differences in the available technology of production (broadly defined), or due to some degree of mis-measurement of inputs. But some of it is also probably due to workers using existing resources more efficiently.[4]

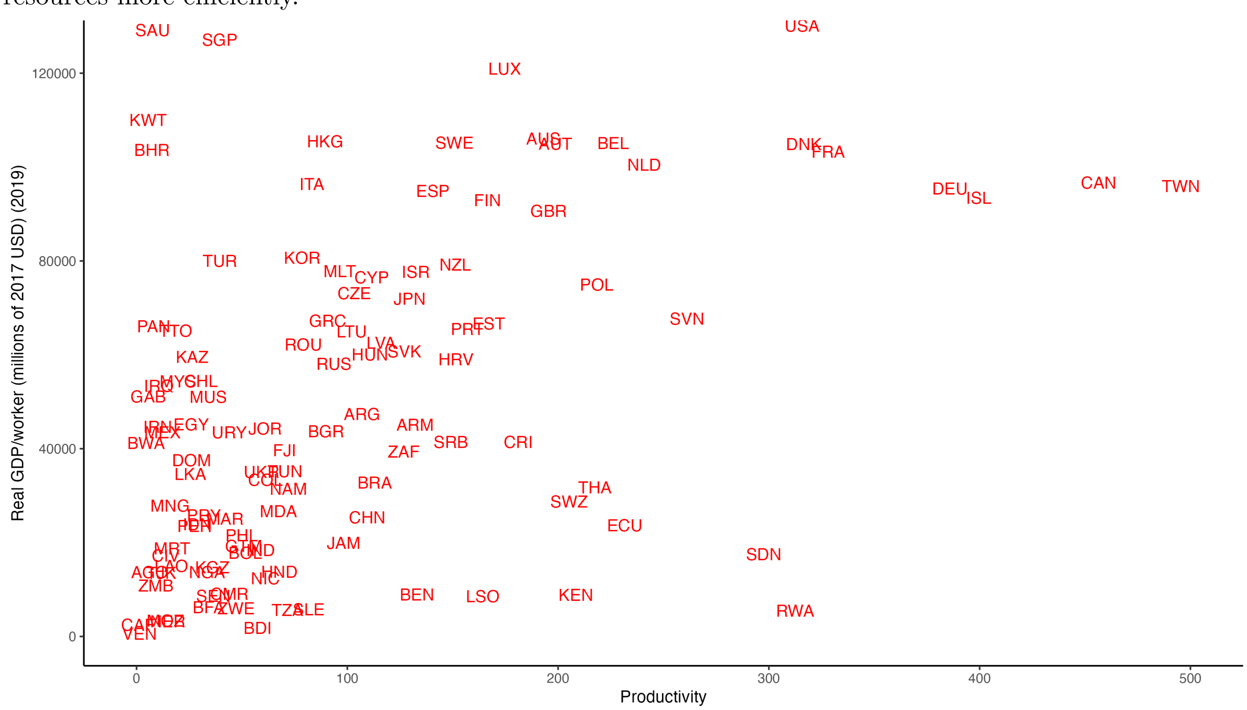

[4]Since everything we can't account for gets thrown in with productivity, it is sometimes called the "Solow Residual." The economist Moses Abramovitz referred to it as "the measure of our ignorance."

In short, countries that are richer (on a per capita or per worker basis) tend to be similar in a number of dimensions: they have more physical capital (financed through higher savings rates), higher levels of human capital, and are more productive. These will also be true in the Solow growth model.

1.3 Fact 3: People in richer countries consume more than people in poor countries

Should we care about whether or not countries are rich? That's a normative question, and to an extent it depends on what we think people value. In general, though people do seem to like consumption, and it turns out that people who live in countries where their workers produce more also seem to be places where the average person consumes more.

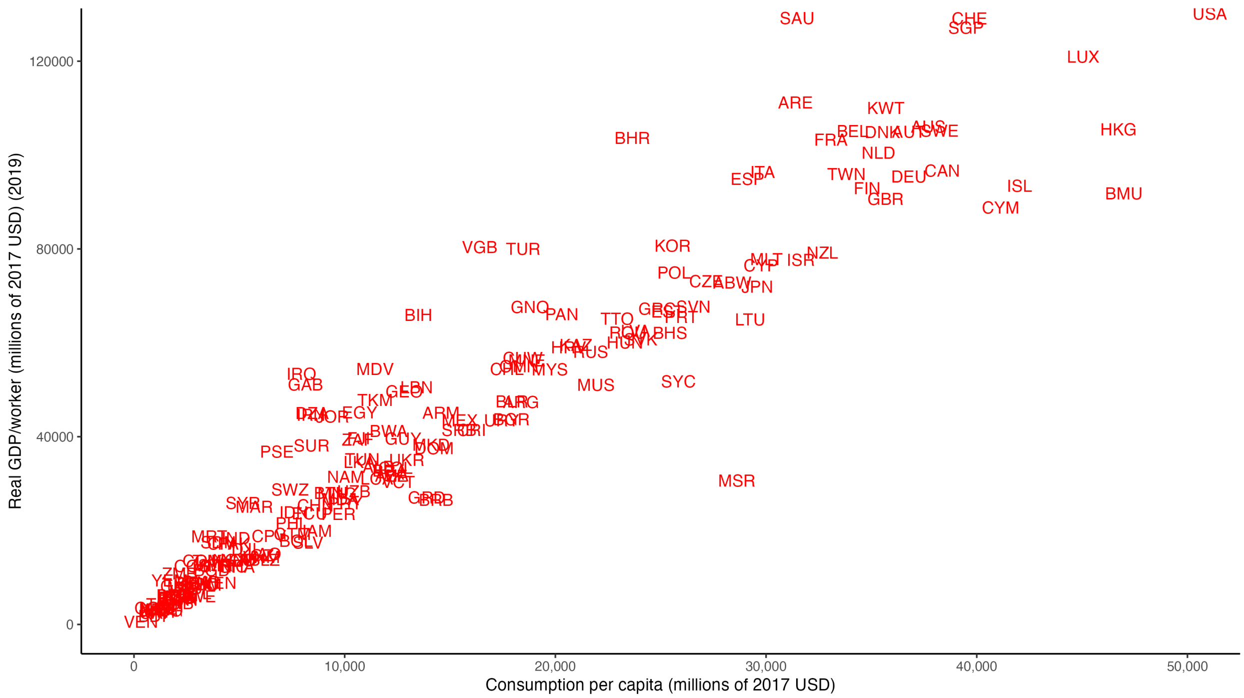

Obviously, this doesn't speak to all aspects of human welfare, or to the distribution of resources within societies. It's a very incomplete picture of whether growth matters for human flourishing. But if consumption is related to welfare, then we should care about growth in GDP/worker so we can understand welfare as well.

1.4 Fact 4: Rich countries tend to grow at a similar rate

Many economies - especially the richest - seem to grow at similar rates. Although this phenomenon –called "convergence in growth rates" — can be made more nuanced, in general we might want a growth model that correctly predicts that a country like the United States and a country like Germany grow at similar rates. We can see this fact by looking at the following scatterplot. The horizontal axis is GDP per worker (on a log scale) in 1960. The vertical axis measures the annual (exponential) growth rate from 1960-2019. Countries that were already (relatively) rich in 1960, in general, grew at relatively similar rates after 1960.

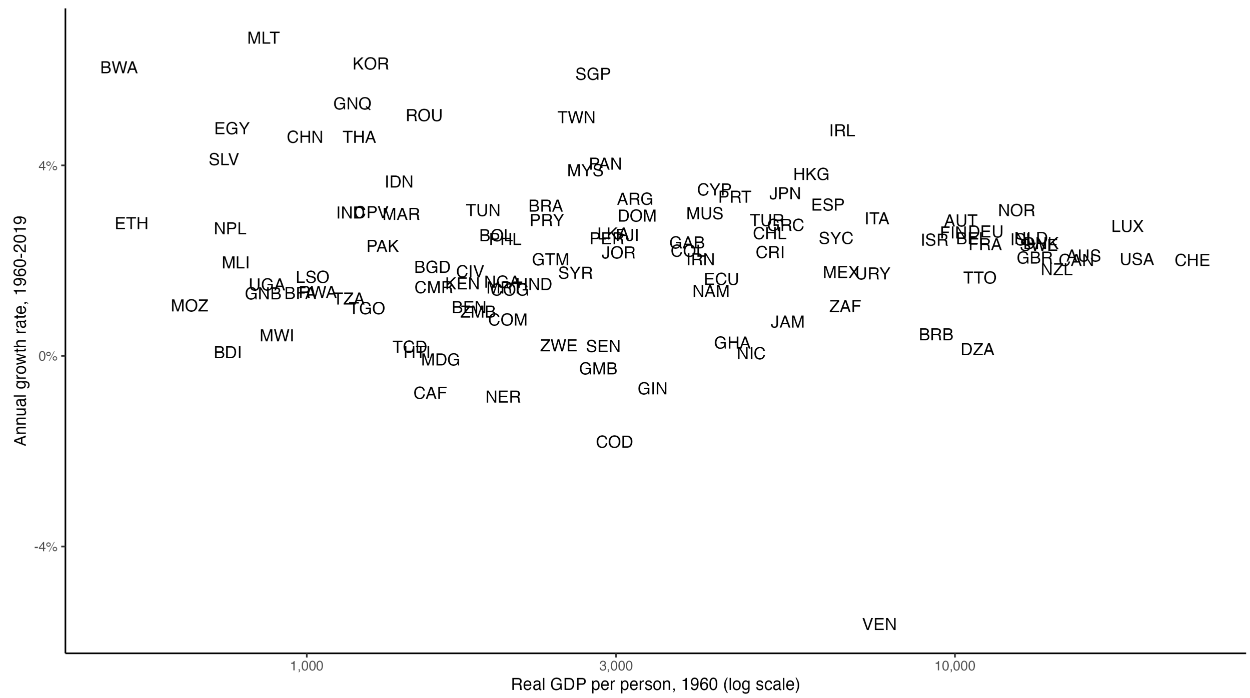

1.5 Fact 5: Countries at a similar level of economic development often have similar geographic, institutional, and cultural characteristics, as well as similar policies.

In an effort to explain facts 1-4, economists have examined a number of "fundamental" causes of growth. For example, it may seem unsatisfying to say that countries that are more productive are richer. Assuming the increased productivity causes countries to be rich, the natural question is "why are countries more productive?" Productivity is a proximate" cause of growth; we're interested in fundamentals.

The Solow model only tells us about proximate causes. At the end of this mini-textbook, we discuss some theories of about differences in things like productivity and savings rates. We can preview those theories by noticing that countries with a higher level of GDP per worker tend to have geographic, institutional, and cultural features in common.

Below, we'll discuss a few examples; there are many others.

1.5.1 Rich countries tend to be further from the equator

As we can see from the below plot, real GDP per capita tends to be higher in countries located at a higher (in absolute value) latitude - that is, countries further from the equator.

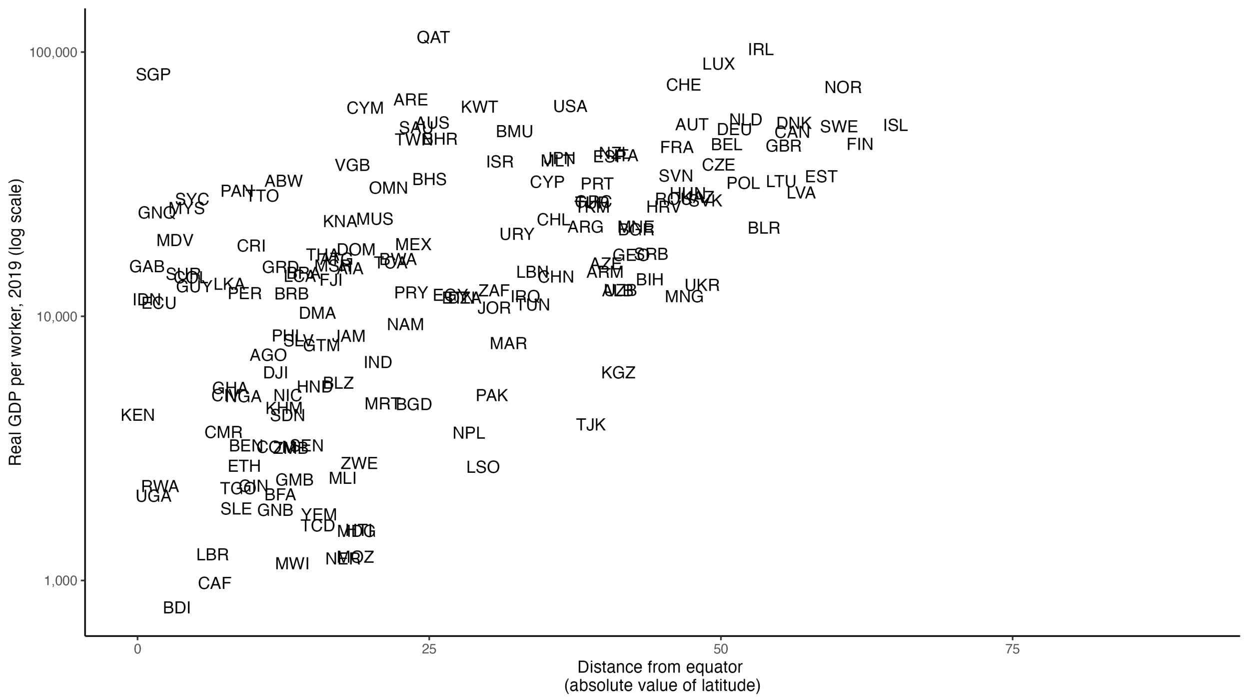

1.5.2 Rich countries tend to have more democratic governments

Many different features of formal and semi-formal institutions have been related to long-run growth. For instance, the Center for Systemic Peace (a research nonprofit) has compiled a historical dataset (Polity IV) on the extent to which governments are characterized as democratic versus autocratic. The scores run from -10 (strongly autocratic) to 10 (strongly democratic). Broadly speaking, rich countries (on a per capita basis) tend to be closer to the strongly-democratic end of the spectrum. The big exception is that there are some extremely rich (on a per capita basis), autocratic oil exporting states like Qatar.

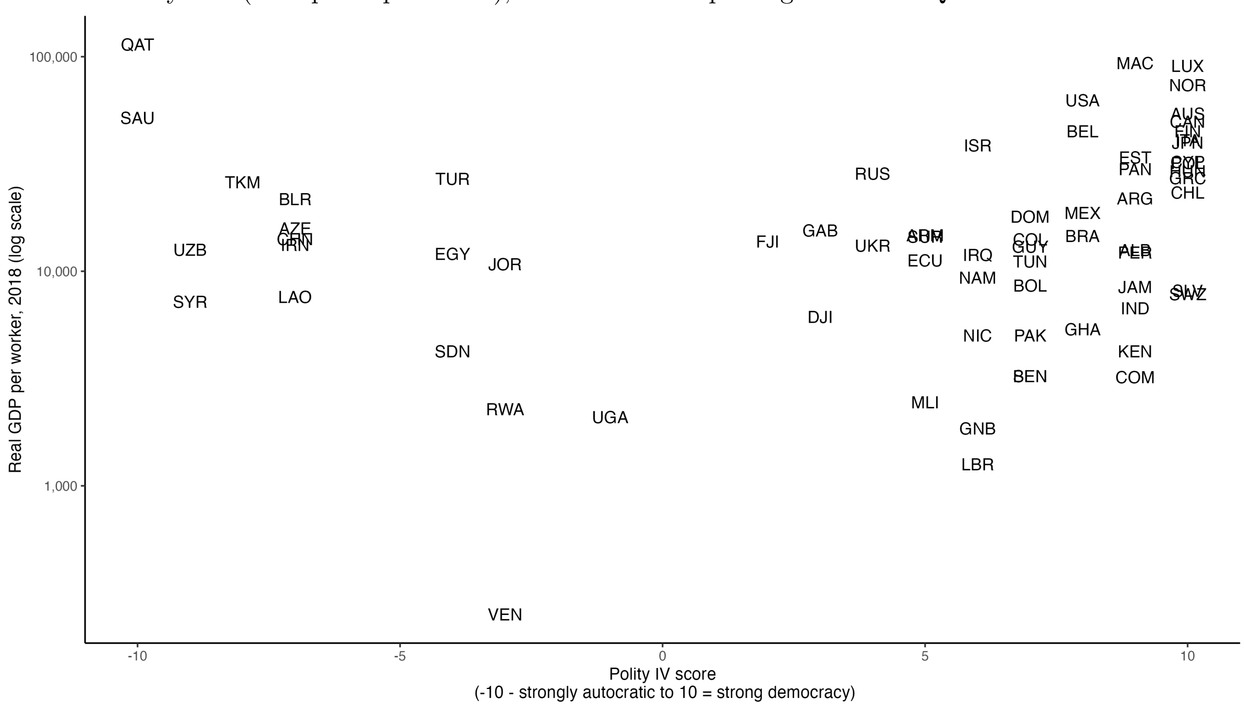

1.5.3 People in rich countries often express similar cultural attitudes

The World Values Survey, conducted since 1981, asks individuals questions about both what attitudes they cultivate in children, the importance of institutions like family or government, and a number of other questions that are broadly related to "culture" - norms, attitudes, etc. Economists have explored whether peoples' stated values are related to their stage of economic development. For example, it appears to be the case that people who live in countries with higher GDP per capita also tend to say that they think higher education is equally important for men and women.

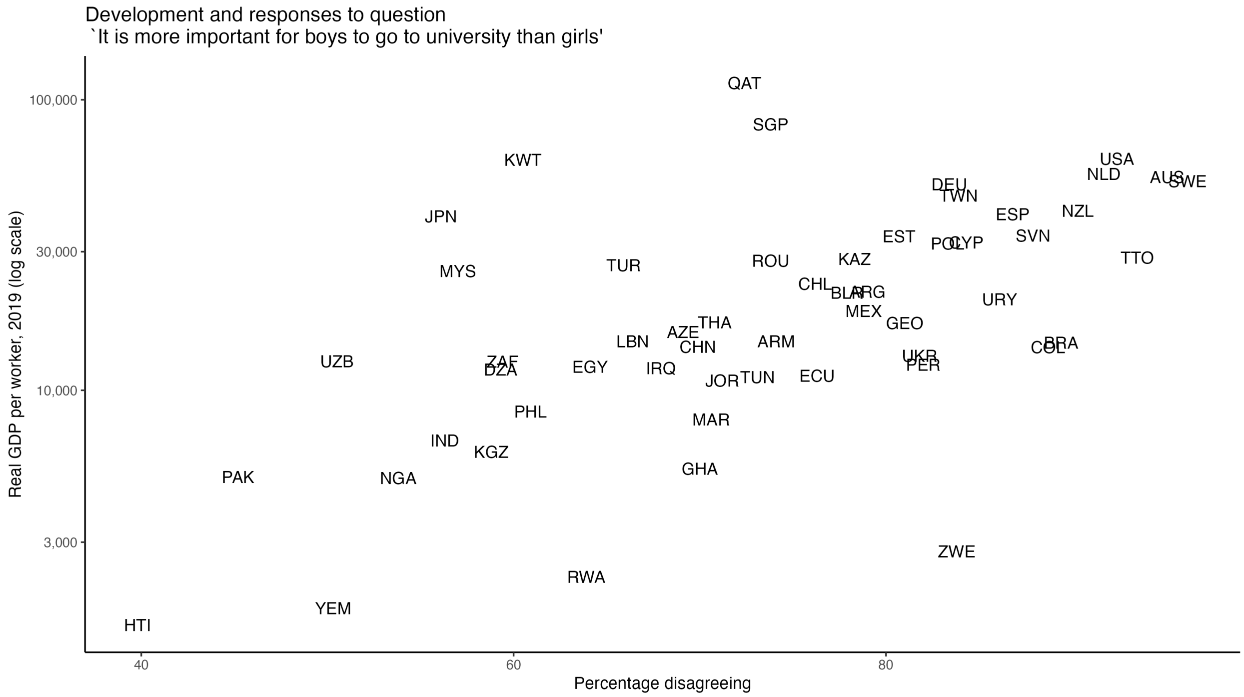

Summing up the facts, and a word of caution

1. Some countries are richer than others on a per-capita basis (there is a wide distribution of growth outcomes).

2. The countries that are richer than others are similar in several dimensions:

   - They tend to have more physical capital per worker.
   - They save and invest at a higher rate than other countries.
   - Their workers tend to have a higher level of education and are healthier (they have more human capital).
   - They seem better at transforming physical inputs into output than poorer countries (they are more productive).

3. People who live in countries that are richer consume more on average.

4. Rich countries tend to grow at similar rates in the long run.

5. Countries at a similar level of economic development often have similar geographic, institutional, and cultural characteristics, and pursue similar policies.

We'll next move to actually developing an economic model that we can relate to these facts.

It's important to remember that these plots we've shown are essentially correlations; they show when two variables appear to move in similar directions on average. However, these kinds of plots can't tell us about causal relationships - whether, for instance, countries are rich because they value gender equality or if it gender equality happens to be associated with some other social or economic that facilitates economic growth. Similarly, it could just be that democracies happen to be located further from the equator and the geographic and institutional correlations we observe are related to each other, without one obviously being the true cause of economic growth.

We would like to understand not just the what, but the why. The model we develop going forward makes certain claims about causality. We would want to have a more sophisticated set of tools than graphs to test whether those causal relationships are, in fact, true. This is the role of econometrics (which is a course you'll take if you continue in economics!)

More briefly, the Nobel Laureate Robert Lucas wrote in 1988:

> Is there some action a government of India could take that would lead the Indian economy to grow like Indonesia's or Egypt's? If so, what, exactly? If not, what is it about the 'nature of India' that makes it so? The consequences for human welfare involved in questions like these are simply staggering: Once one starts to think about them, it is hard to think about anything else. This is what we need a theory of economic development for: to provide some kind of framework for organizing facts like these, for judging which represent opportunities and which necessities.

2 Main Assumptions

This is a particular version of the Solow model, one where we are making some special assumptions (more special than the ones Solow made). Some of them are technical assumptions meant to make the model easier to analyze. Throughout, I will try to be clear about what those special assumption are, and you should think about how things might change if those assumptions changed. As the opening of Solow's article put it:

> "All theory depends on assumptions which are not quite true. That is what makes it theory. The art of successful theorizing is to make the inevitable simplifying assumptions in such a way that the final results are not very sensitive."

2.1 Consumption and investment

To set things up, we're going to make a few simplifying assumptions about the economic environment.

Assumption 1 (Closed economy without government)**.** *We assume that there is no trade, so net exports are zero. We also assume there is no government spending and taxes are zero.*

This makes things very simple, because all GDP is either consumption or investment:

$$Y_t = C_t + I_t$$

Important thing to remember: This is a statement about accounting, not about causal relationships! It will always be true, but it won't be the case that we can make statements like "an increase in consumption increases GDP" in this model – that's a more complicated question.

Along those lines, we will make the further assumption about how investment is related to the level of GDP.

Assumption 2 (Constant rate of saving)**.** *Assume a constant fraction of output, $0 < \gamma \leq 1$, is saved. (γ is the lower case Greek letter "gamma"). That is, γY_t is the amount of investment that occurs in a given period of time. We call γ the saving rate. γ is a parameter of the model – one that takes on a constant value by assumption.*

It's hopefully easy to see that if $\gamma Y_t = I_t$, then as a matter of accounting,

$$C_t = (1 - \gamma) Y_t$$

2.2 The aggregate production function

Now that we've set up some accounting facts, we need to figure out how much production there actually is. We'll imagine a very stylized version of the economy where there are several different different "factors of production" – inputs – which are combined to create a single output. That output is real GDP.We summarize the relationship between the quantity of inputs and the quantity of output using a mathematical relationship, which we'll call the **aggregate production function**

$$Y_t = f(A_t, L_t, K_t, H_t)$$

Where at a particular time t (say, a given year):

- Y_t is real GDP in that year
- L_t is the size of the **labor force -** the number of workers in the country
- K_t is the stock of **physical capital** – the amount of equipment, physical structures, etc that are used to produce goods and services.
- H_t is the stock of **human capital** – the knowledge, skills, and ability that individual workers have. They acquire human capital via education, training, experience, and other investments.
- A_t is **productivity**. Productivity is the ability to combine factors of production into output. As A_t increases, a given amount of labor, physical, and human capital will produce more final output. We can think about this as improvements in technology, but also improvements in *efficiency* with which we use the factors of production.
- $f(A_t, L_t, K_t, H_t)$, the **production function**. This summarizes the mathematical relationship between the amount of labor, physical capital, human capital, productivity, and the quantity of real GDP produced.[5]
- If we don't have one of these factors (e.g., if $K_t = 0$), then there is no production ($Y_t = 0$).

We will conventionally use upper-case letters for aggregate variables. Lower case letters will be for "per worker" variables. That is, Y_t is GDP in a country in year t, and $y_t \equiv Y_t/L_t$ is GDP per worker in year t.

We're going to talk about some general assumptions we need for this production function, and then we'll make an assumption about the functional form (which will be a stronger assumption than we need for the model, but it's convenient).

2.3 Assumptions about the aggregate production function

Assumption 3 (Constant returns to scale)**.** *We assume that the production function is* ***constant returns to scale (CRS) in the "reproducible factors of production" (labor, physical capital, and human capital).*** *This means that if we increase each of these factors by the same proportion, output also increases proportionately. Mathematically, if z is a constant:*

$$zY_t = f(A_t, zL_t, zH_t, zK_t)$$

[5]To keep all of these straight, it may be helpful to think about production of a particular good, and pretend that good is the only good produced in the country. The economist Paul Romer (who won the 2018 Economics Nobel Prize for his contributions to growth theory) uses the analogy of producing a chair. The chair is the final good (real GDP). The worker who produces the chair is part of the labor force. The tools she uses to produce the chair are physical capital. Her training as a carpenter is her human capital; nobody else can "use" her human capital without hiring her to use it. Productivity, on the other hand, is like the Pythagorean theorem or other facts about geometry that anyone can use to make a chair once they're discovered. Physical and human capital, clearly, take resources to produce; we have to spend resources to create tools or to train carpenters. However, in general, we do not need to expend resources to re-invent geometry once we've discovered it. And once we know how many workers we have, their individual skill as carpenters, the tools they work with, and how productive they are using their physical and human capital, we can know how many chairs they produce.

The reasoning behind CRS is what's often referred to as a "replication argument." Suppose you have a factory with a certain amount of workers, and it produces a given amount of output. If you build an identical factory across the street, and put another identical set of workers in the new factory, you should be able to produce twice as much output as before.

We make one more *critical* assumption about the production function

Assumption 4 (Diminishing marginal product of capital)**.** *The amount of output per worker always increases when we add additional capital per worker, but at a diminishing rate. Each additional unit of capital per worker adds less and less output, holding everything else fixed.*

This mean if you have a very small amount of capital, adding an additional machine makes a big difference in the amount you produce, all else equal. If you have lots of machines, and the same number of workers, one additional machine won't help as much.

1. Explain the difference between assuming constant returns to scale and assuming diminishing marginal product of capital.

The most common example of a constant returns to scale production function is something like the following

$$Y_t = A_t K_t^{\alpha} \left(h_t \times L_t\right)^{1-\alpha}$$

This function may look a little strange. It's an example of what's called a "Cobb-Douglas" function.[6]

The term $h_t \equiv H_t/L_t$ is the average amount of human capital workers have. Multiplying that by the number of workers gives us the total quantity of human capital inputs. The way human capital works in this particular production function is by making it "as if" we have more workers than actually physically exist.

2. Verify that the Cobb-Douglas production function

$$Y_t = A_t K_t^{\alpha} \left(\frac{H_t}{L_t} L_t\right)^{1-\alpha}$$

exhibits constant returns to scale in K_t, L_t, H_t. (That is, if we multiply each of the factors by some number $z > 0$, that we get $z \cdot Y_t$)

[6]Cobb-Douglas functions are functions that are the products of exponential terms. If you're feeling rusty with exponents, there are some reminders in the appendix). A constant returns to scale Cobb-Douglas function is one where all the exponents sum to 1. You could also have decreasing returns or increasing returns to scale with this function, by changing the exponential terms. You often see Cobb Douglas production functions, but sometimes you see them pop up in other contexts in micro- and macroeconomics. They're very convenient to work with and have some nice properties.

The CRS assumption is a standard assumption. The next assumption is special: We don't need to make it, the standard Solow model doesn't make it, but it will make solving the model easier.

Assumption 5 (Special assumptions about the production function)**.** *We will assume that the level of productivity, the number of workers, and the amount of human capital is fixed over time.*

Why make these assumptions? They allow us to go from the aggregate production function to one that is in per worker terms easily.

3. Starting from

$$Y_t = A_t K_t^{\alpha} \left(h_t \times L_t\right)^{1-\alpha}$$

   Derive the per worker production function

$$y_t = f(A, k_t, h) = A k_t^{\alpha} h^{1-\alpha}$$

4. Draw the per-worker Cobb-Douglas production function (with capital per worker on the horizontal axis, and output per worker on the vertical axis) you derived. On the same set of axes, draw an investment function Then, draw a function for investment per worker $i_t = \gamma y_t$. For a particular k, label the amount of output, investment, and consumption.

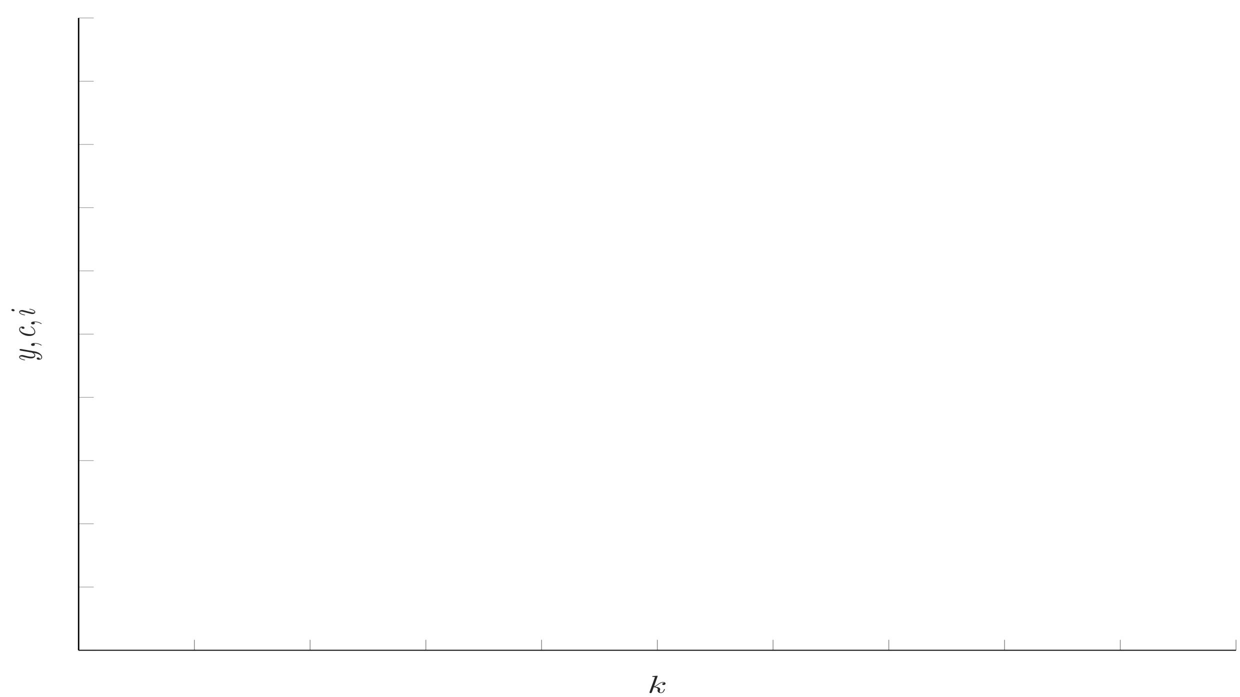

2.4 The evolution of physical capital

Given the assumptions we've made, the only factor of production that changes over time is the amount of physical capital (the capital stock) per worker. The next set of assumptions is about describing how the capital stock changes over time.

To think about this, we'll focus on two special features about physical capital. First, it is created by the act of investing. Investing is the decision to use resources today to create more physical capital that will be available for use in the future. The other feature is is that physical capital depreciates (wears out) over time. For instance, a car will break down over time as its parts wear out, and those parts will have to be replaced.

In other words, the amount of capital available at a given time could be generally described as

$$\begin{aligned} \text{Capital this year } &= \text{Capital last year} \\ &- \text{Capital that “wears out” (depreciates) after last year} \\ &+ \text{New capital built (investment) last year} \end{aligned}$$

We need to make an assumption about how much capital wears out in each period.

Assumption 6 (Constant rate of depreciation)**.** *Assume that the *rate* of depreciation is fixed over time at a number* $0 < \delta \leq 1$*. (*δ *is the lower case Greek letter “delta”.) That is,* δK_t *is the amount of physical capital that wears out in a period of time.*

5. Graphically depict how the level of depreciation of capital per worker changes as the amount of capital per worker changes. That is, draw a “depreciation function” with capital per worker on the horizontal axis, and depreciation per worker on the vertical axis that is consistent with a constant rate of depreciation.

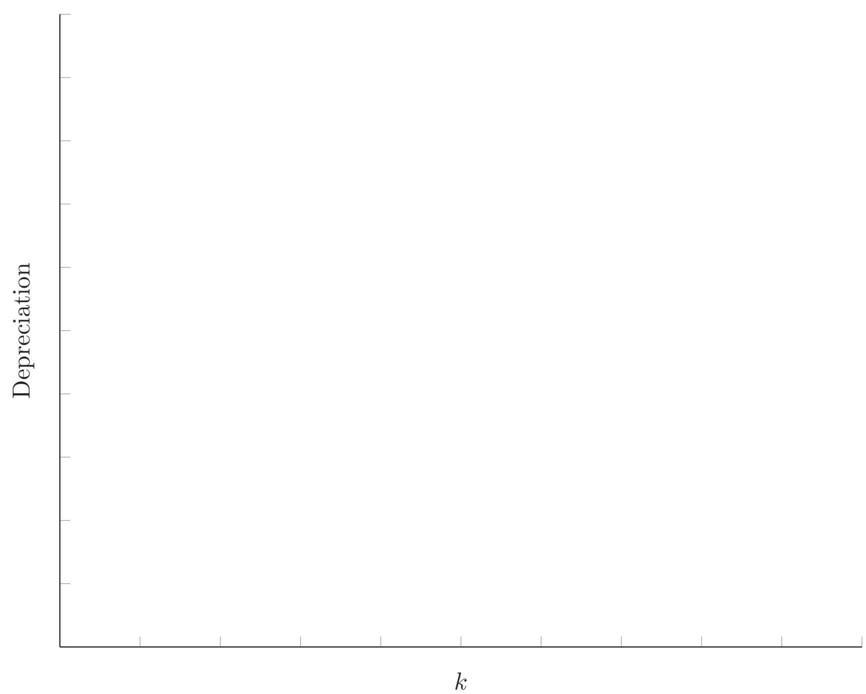

Mathematically, we can represent the evolution of the capital stock in the aggregate as

$$K_t = K_{t-1} - \delta K_{t-1} + I_{t-1}$$

where I_{t-1} is the aggregate amount of investment in year $t-1$.By dividing by the number of workers, we can come up with the per-worker capital evolution equation

$$k_t = k_{t-1} - \delta k_{t-1} + i_{t-1}$$

(Hopefully, it's easy to intuit how this would change if we assumed population was growing over time: the amount of machines per worker would depend on how many machines were wearing out, how many new machines we were building, and how quickly new workers were arriving).

2.5 Summarizing the main assumptions

In this section we've outlined the crucial mathematical assumptions behind our version of the Solow model:

1. Assumption 1: Closed economy without government, so, (in per worker terms)

$$y_t = c_t + i_t$$

2. Assumption 2: Constant rate of saving and investment, so that (in per worker terms)

$$i_t = \gamma y_t$$

3. Assumption 3: The aggregate production function is CRS and Cobb-Douglas:

$$Y_t = f(A_t, L_t, K_t, H_t) = A_t K_t^{\alpha} (h_t L_t)^{1-\alpha}$$

4. Assumption 4: Diminishing marginal product of capital (this follows from the functional form we assumed)

5. Assumption 5: For simplicity, assume productivity, labor, and human capital do not change over time:

$$Y_t = f(A, L, K_t, H)$$

$\Rightarrow$ Combining the previous two assumptions,

$$y_t = f(A, k_t, h)$$

6. Assumption 6: Constant rate of depreciation

$$0 < \delta \leq 1$$

In the next section, we'll start combining these assumptions to figure out what they imply for the predictions of the model.

3 Solving for steady state

Now, we're ready to explore the predictions of the model.

To be precise, the model has several variables - A, h, δ, γ that we take as "exogenous" (determined outside the model).

Given those variables, we want to be able to explore the predictions of the model for "endogenous" variables (determined inside the model) – y, c, k, i. Roughly speaking, we're asking ourselves two sets of questions (1) If we knew A, h, δ, γ what level of y, c, k, i would we expect the economy to arrive at eventually? How are countries that have different values of the parameters different from each other in terms of endogenous variables? (2) If A, h, δ, γ change, how would y, c, k, i change over time?

How do we approach these questions? Let's take a step back.

The assumptions the model made meant that we could summarize (mathematically) the economy (in per worker form) in four equations:

1. $y_t = f(A, k_t, h)$
2. $k_t = k_{t-1} - \delta k_{t-1} + i_{t-1}$

3. $i_t = \gamma y_t$
4. $y_t = c_t + i_t$

The first equation says that output depends on parameters A, h, and the amount of capital per worker k_t. The third and fourth equations say that once we know output, and given the parameter γ we also know the amount of investment and consumption (per worker) in the economy. So we need to know what capital per worker k_t is today, and then we can figure everything else out about the economy today.

Equation 2 says that capital today depends on capital yesterday, and investment yesterday. But we know investment yesterday depends on output yesterday (just think about equation 3, but for $t-1$), and output yesterday depends on capital per worker yesterday (think about equation 1, rolled back to $t-1$). To sum up **knowing capital per worker in the previous period tells us capital per worker today. Knowing capital per worker today tells us everything else about output per worker today, investment per worker today, and consumption today**

3.1 Digging into how capital changes over time

Let's look at how capital per worker evolves over time again:

$$k_t = k_{t-1} - \delta k_{t-1} + i_{t-1}$$

Another way of writing this is

$$k_t - k_{t-1} = i_{t-1} - \delta k_{t-1}$$

That is, the change in the number of machines per worker is related to investment, the rate of depreciation, and the size of the capital stock in the previous period.

6. Using the capital evolution equation, explain (mathematically and intuitively) when the amount of capital per worker will be increasing, when it will be decreasing, and when it will be constant. (And indicate it graphically)

3.2 Steady State

It turns out that the assumptions we've made imply that the economy will eventually converge to a steady state where $k_t = k_{t-1}$ and then nothing will change after that. As the previous exercise reveals, this is when $i_{t-1} = \delta k_{t-1}$. Intuitively, this is when the amount of new capital being built is exactly enough to replace the capital that's wearing out. It's easiest to see how this happens graphically, by drawing the functions for δk and i on the same graph:

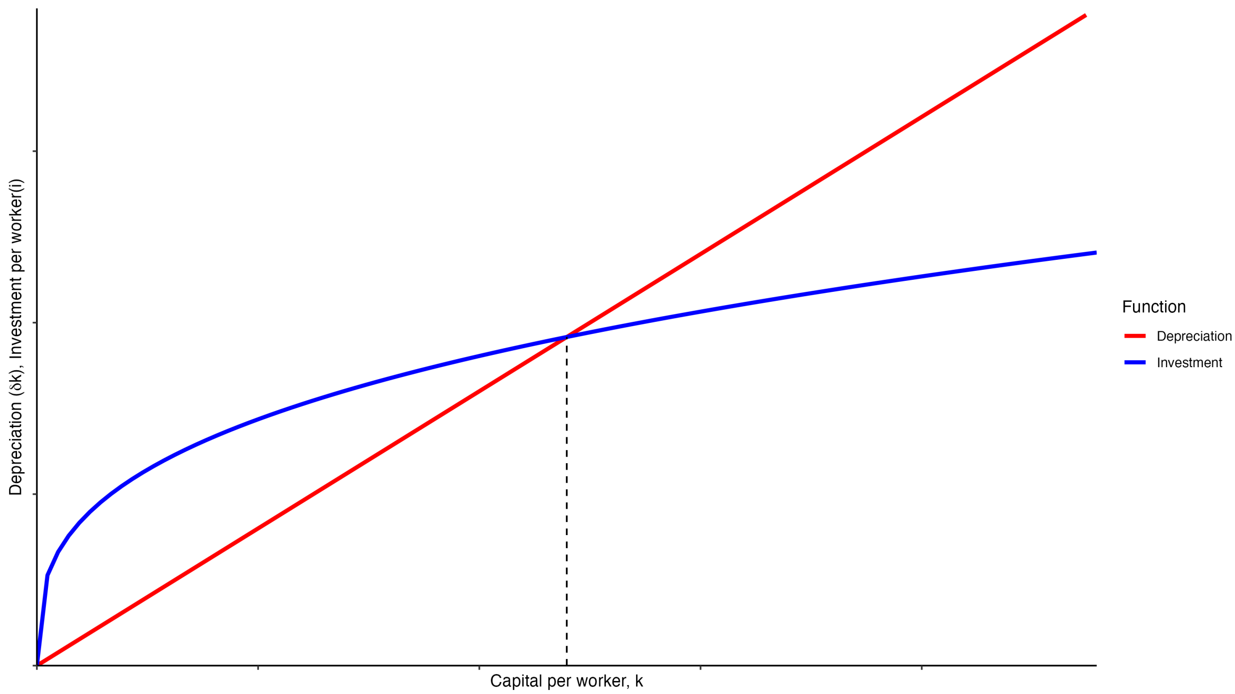

The intersection of the depreciation and investment functions is where they are equal (the dashed line indicates this level). At that point, capital per worker will not change over time.

7. Explain why we would end up at the level of capital per worker indicated by the dashed line if we started out at any other (positive) level of capital per worker.

Notice that we can also add the production function to graphically see the level of GDP per worker in steady state as well.

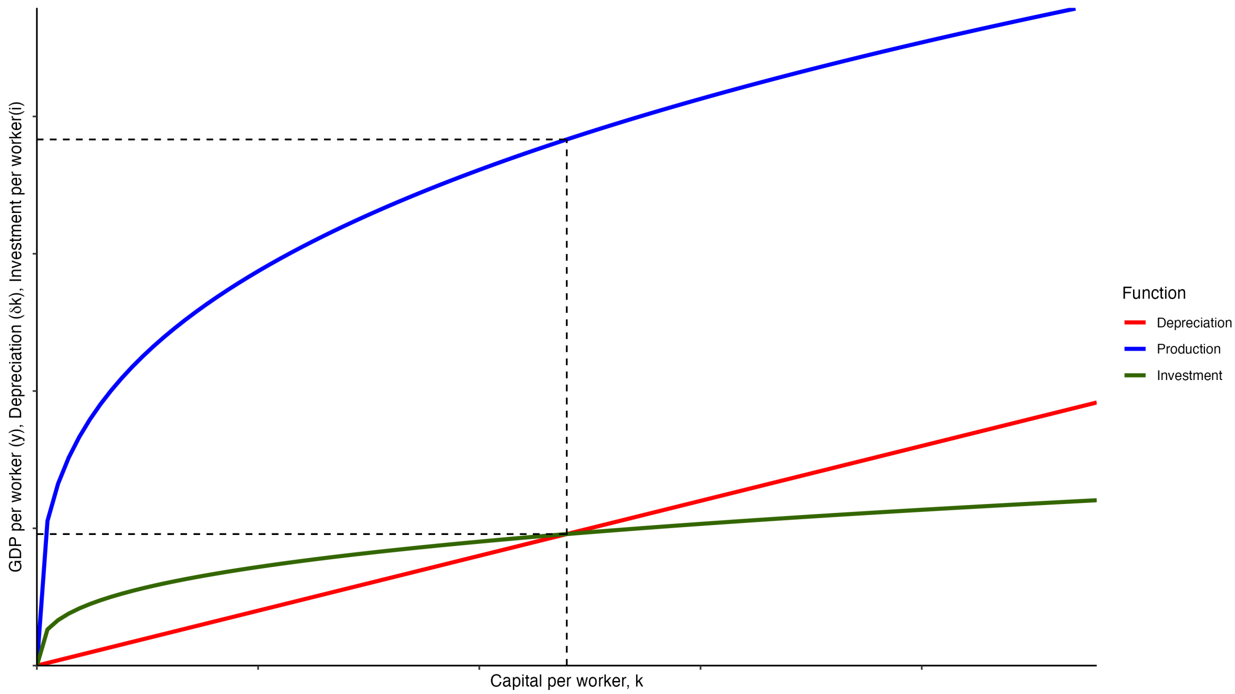

3.3 Preview

Now that we know how to find a steady state, we can compare two countries that are identical in all ways except a single parameter. For instance, let's imagine we are using the Solow model to predict how two countries that are identical in all ways except their savings rate will be different.

8. On a single graph, illustrate production, investment, and depreciation curves for two countries that have the same productivity, human capital per worker and rate of depreciation, but different savings rates. Clearly label the steady state levels of capital per worker and output per worker.

   Based on your graphs, which country has higher output per worker? Which country has higher *consumption* per worker?

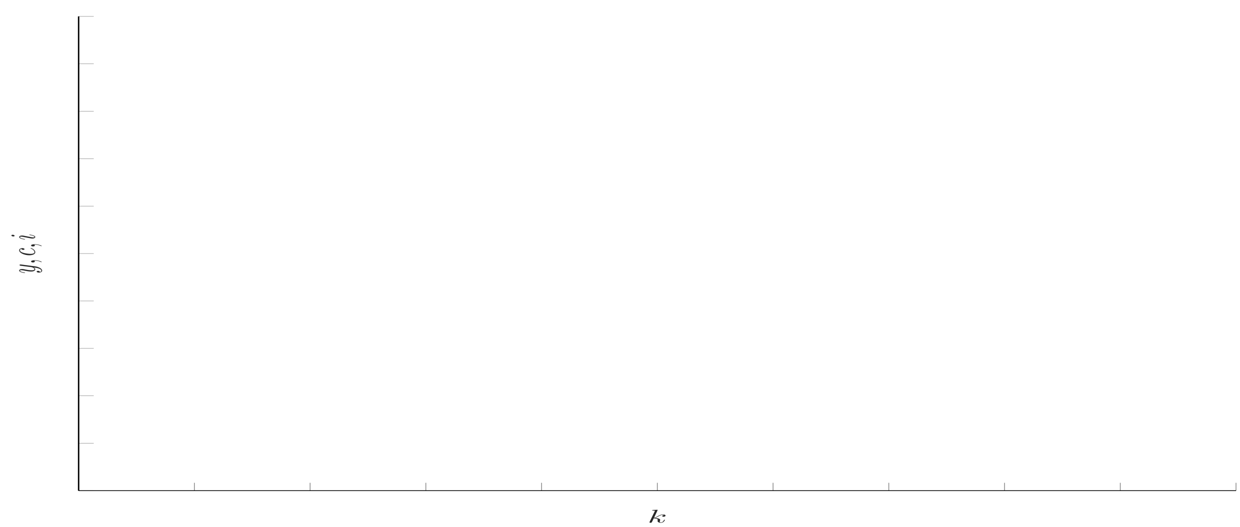

4 Comparing two countries

So a few things to remember when doing this kind of exercise: 1. All the curves should always come out of the origin – without capital, we don't produce. 2. Changes in capital per worker will be movements *along* the curve, and changes in other variables rotate the curves upward or downward. In this case, a higher savings rate implies that at any given level of k, i is higher (because $i = \gamma y = \gamma f(A, k_t, h)$). γ doesn't appear in any other equation, so none of the other curves move. Intuitively, the depreciation curve shouldn't change because it's only describing how many machines wear out through production; the production function doesn't change because it just says how much we produce given the inputs available, not what we do with them.

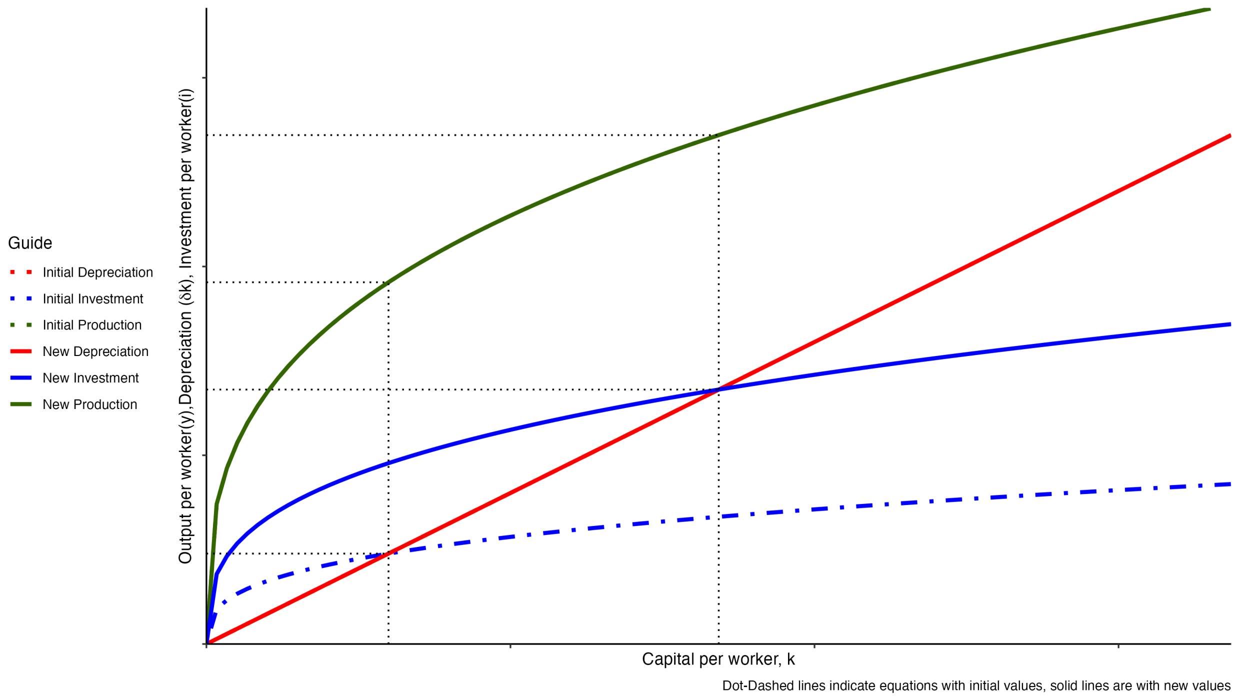

The higher-saving country will have higher capital per worker in the steady state (because they are saving enough to offset a higher level of depreciation), and hence higher output per worker. However, they will not necessarily have higher consumption per worker. Saving more means consuming less *at any given level of output*. Output is increasing, but so is saving. So, it depends on which change is relatively larger. In the above graph, it is in fact the case that consumption per worker goes down despite output per worker increasing quite a bit.Increasing saving means the workers in this country face a tradeoff between lower consumption today in favor of higher output tomorrow. Depending on how much they're saving, they may or may not have higher *consumption* in the future.

To summarize: According to the Solow model, a country that saves more, all else equal, will be able to sustain a higher level of capital per worker and hence have a higher level of output per worker. They may or may not have a higher level of consumption per worker.

4.1 Numerical comparisons

We already introduced the Cobb-Douglas production function. If we are willing to assume that the Cobb-Douglas function is a good description of production, we can use it to make numerical predictions about the levels of capital per worker (and hence output, investment, and consumption) in different countries by picking different numerical values for the parameters.

In the online version of this text, there's an interactive application which lets you select parameters for two countries, calculate steady state values, and plots the comparative graphs. This may be helpful for checking your work as you practice using the model. It's also available at the following web address: `https://shinyapps.carleton.edu/estruby/solow_comparative/`

5 Dynamics: What happens over time when parameters change

Before, we thought about using the Solow model two compare two countries that were identical, except one country had a different savings rate than another. But we could just have thought about it in terms of a single country that decided to save a different amount than it had in the past. We can also use the model to think about what happens over time in that country.

5.1 Changes in savings rates

Let's go back to the previous example. Imagine the low savings-rate country decides to increase its saving. Graphically, we could depict this change as a shift in the investment curve. We know that the economy's new steady state is at the intersection of the higher investment curve and the (same) depreciation curve. We want to explain intuitively how the economy gets from the old steady state to the new.

The verbal explanation would be:

> At the "old" level of capital per worker, the economy jumps from the old to new investment curve. Because investment is now greater than depreciation at this level of capital per worker, the economy begins to accumulate more capital per worker. This means that the economy produces more. The process of growth of the capital stock and greater output continues over time, with the additional output gained diminishing as capital per worker increases. Eventually, the country reaches the new steady state where depreciation is equal to the new, higher level of investment.

If we were to look at data for this country, we would see that their level of investment (per worker) initially increases. We would then observe (in per worker terms) capital, output, investment, and depreciation increasing over time. The initial increase in investment (and output) is large, but gradually slows (because of diminishing marginal returns), until growth stops once the economy reaches the new steady state.[7]

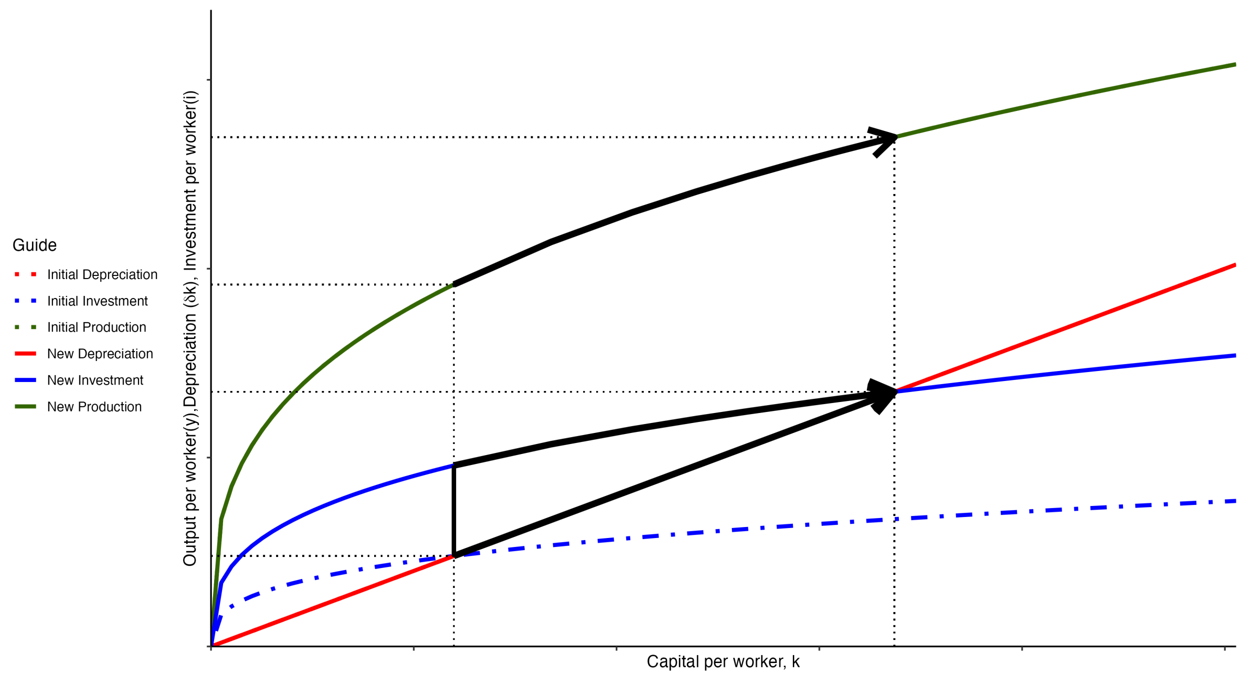

[7]This gives us some insight about how consumption will evolve. Initially, it drops because savings have increased. It begins to increase from its new low level over time because output per worker is gradually increasing. It may or may not reach (or surpass) its level at the old steady state depending on how much savings change and the rate at which returns are diminishing)

5.2 Changes in productivity

Preview

Suppose that, due to an improvement in efficiency, a particular country becomes more productive. How would we depict this graphically and describe the change over time?

9. Draw, and describe intuitively, how a change in productivity affects capital per worker, output per worker, investment per worker, and consumption per worker in a country over time.

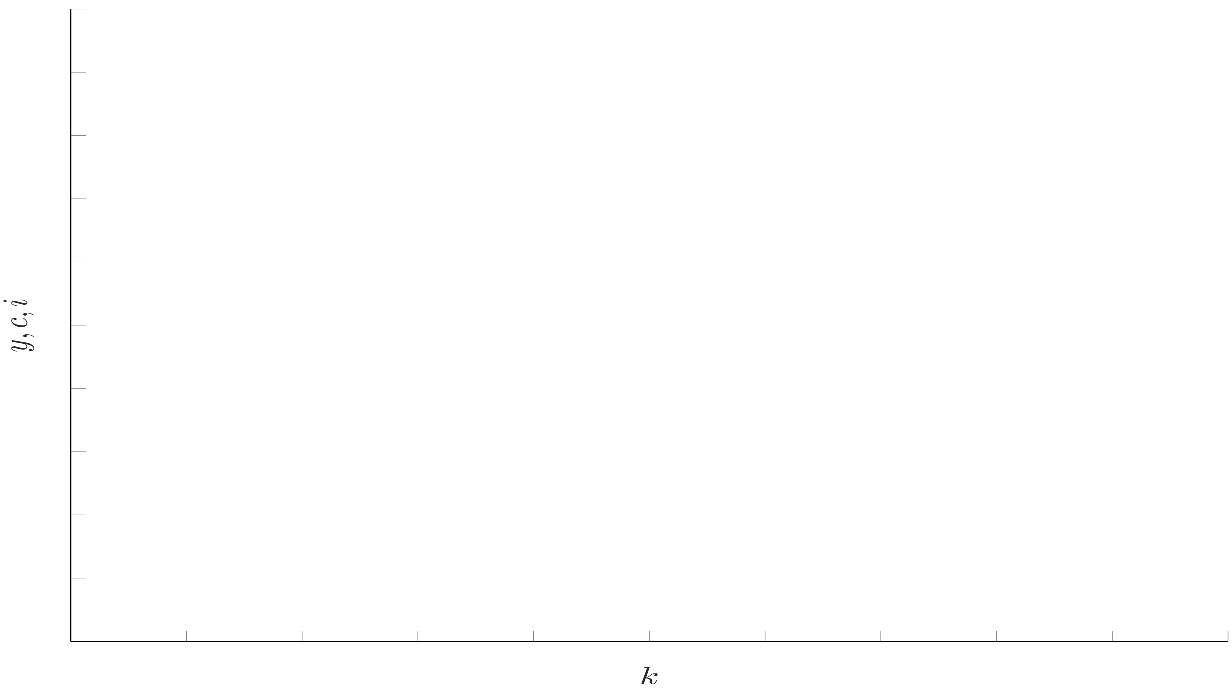

Change in productivity: Solution

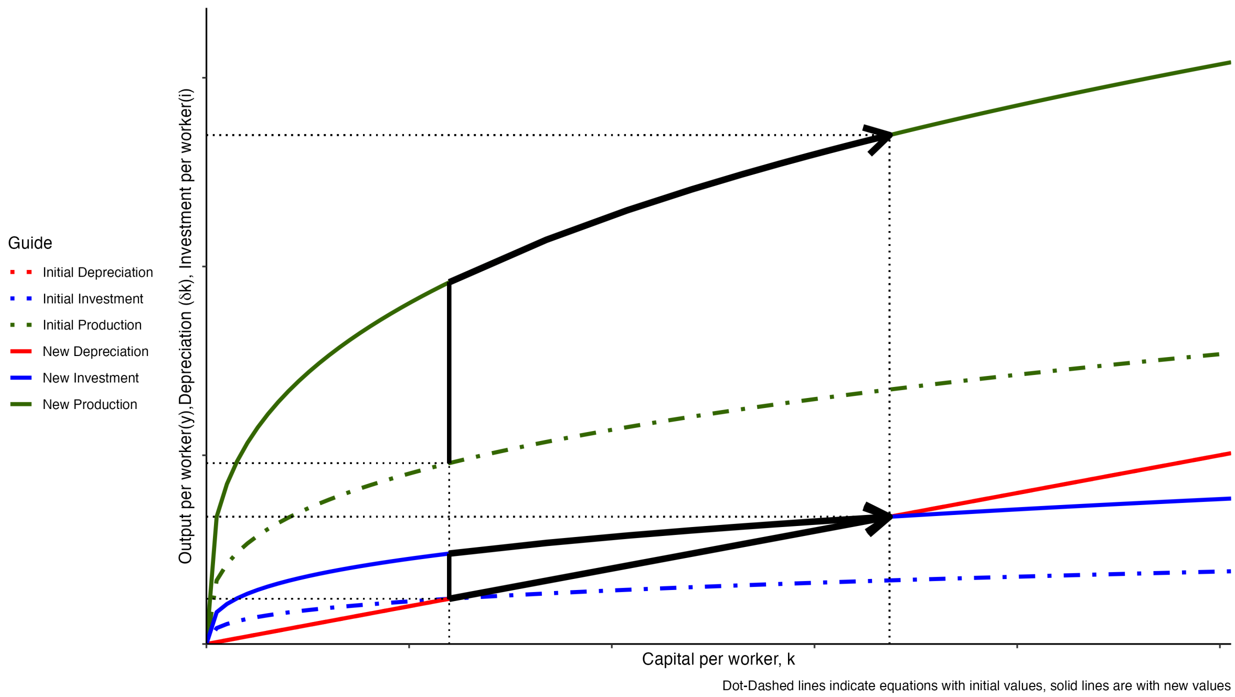

The increase in productivity means that at every level of capital per worker, workers produce more output. This also means that at a given rate of saving, more is invested. So there is an initial increase in output and investment due to the change in productivity. Then, because investment is greater than depreciation, capital per worker begins to grow over time, which leads to further increases in output, investment and depreciation. This continues until the economy reaches its new steady state with a higher level of capital per worker. Note that consumption is a constant fraction of output per worker, and the savings rate hasn't changed, so consumption increases over time as well.

Unlike changes in savings rates, there is no tradeoff between current and future consumption when productivity increases.

This observation is quite important. We assumed that productivity was constant. This means that in steady state, GDP per worker doesn't change; all of the growth in the model is simply transitioning from one steady state to another. This is at odds with the fact that most rich countries seem to grow at a similar, positive rate. However, the "full-blown" Solow model (presented in the Appendix) actually assumes that productivity grows at a constant rate over time. In the long run, this growth rate of productivity is the same as the growth rate of GDP per worker. We could understand this growth rate as the rate at which technology or efficiency is increasing at the "frontier," and perhaps is related to the productivity of research scientists. Trying to explain the growth rate of productivity *endogenously* – within the model – is what earned Paul Romer the 2018 Economics Nobel Prize.

In the online version of this text, there is an application that lets you select "start" and "end" parameters, calculate steady state values, and plots the comparative graphs, including the path that the economy follows to get there. This may be helpful for checking your work as you practive using the model. It's avaialble at the following web address: `https://shinyapps.carleton.edu/estruby/solow_paths/`

6 Summarizing the implications of the Solow model

By looking at steady states, we can examine why (in the long run) some countries are richer than others (Fact 1), and how that relates to their inputs and productivity (Fact 2) and consumption (Fact 3).

10. Relate the predictions of the Solow model to the first three facts about the data shown in section 1.

In our stripped-down version of the Solow model, the only reason countries grow is because of "catch up" or transitional growth. Still, we can imagine that if productivity grows every year, countries that are at their "per worker" steady states initially will also be growing over time.

11. Relate the predictions of the Solow model, and the idea of catch up growth, to explain differences in growth rates observed in the data, and the "convergence" of growth rates in rich countries (fact 4 in Section 1)

You'll notice that we haven't really explained fact 5: the fact that rich countries tend to be similar in some dimensions (geography, culture, and institutions) that the Solow model is silent on. Next, we'll discuss three theories about "fundamental" reasons that some countries are richer than others.

7 Fundamental causes of growth

> Men make their own history, but they do not make it as they please; they do not make it under self-selected circumstances, but under circumstances existing already, given and transmitted from the past.
>
> Karl Max, *The Eighteenth Brumaire of Louis Bonaparte*

The Solow model relates growth to features of economies like their degree of saving or human capital. But why do some countries save more than others? Why do certain countries choose more or less efficient methods of organizing production?

This last question is probably the most critical. We call productivity as a "residual" - what's left over after physical and human capital are accounted for. The fact that productivity is so highly correlated with GDP per worker (and person!) is evidence that that the proximate causes of growth - inputs and their direct determinants like savings rates - are not enough to answer the question "why are some countries richer than others?" We need to look for fundamental causes - things that explain productivity and investment in physical and human capital.

The economist Daron Acemoglu[8] has sorted theories of growth into three big categories: geography, culture, and institutions. In addition to these, we might add another possible candidate – "policy." We'll briefly explore each, although we've already seen examples of those big theories. (Fact 5 in the Data section). We'll briefly discuss some representative arguments in favor of each.

7.1 Geography

Geography is a catch-all term for the "exogenous" features of the physical environment. For example, climate and soil quality a country can affect agricultural growth. Ease of transportation on waterways – like rivers or oceans – can facilitate economic exchange, and coastal countries tend to be richer than landlocked ones. Swampier regions tend to have more diseases like malaria that affect health directly and other aspects of human capital indirectly. Natural resources – like deposits of energy resources, primary metals, etc – have been quite important in shaping some countries' economic and political development over time.

Although natural resources are only part of the overall geographic hypothesis, you might be wondering just how important they are. One, somewhat crude, way of looking at this would be to take real GDP and subtract the value added from the mining and resource extraction industry. (A similar adjustment was used by Robert Hall and Charles Jones in a widely cited paper from 1999). Below, I've plotted a scatter plot with real GDP/worker in 2017 on the horizontal axis and non-mining real GDP/worker on the vertical axis. Countries close to the 45 degree line have very little value added from mining. Countries further below the line have more of their GDP coming directly from the mining industry. You notice the big outliers are mainly oil-rich countries, which we saw are outliers in terms of GDP/worker (and in terms of having relatively autocratic institutions for rich countries, Norway excepted).

[8] Winner, alongside coauthors Simon Johnson and James A. Robinson, of the 2024 Economics Nobel prize!

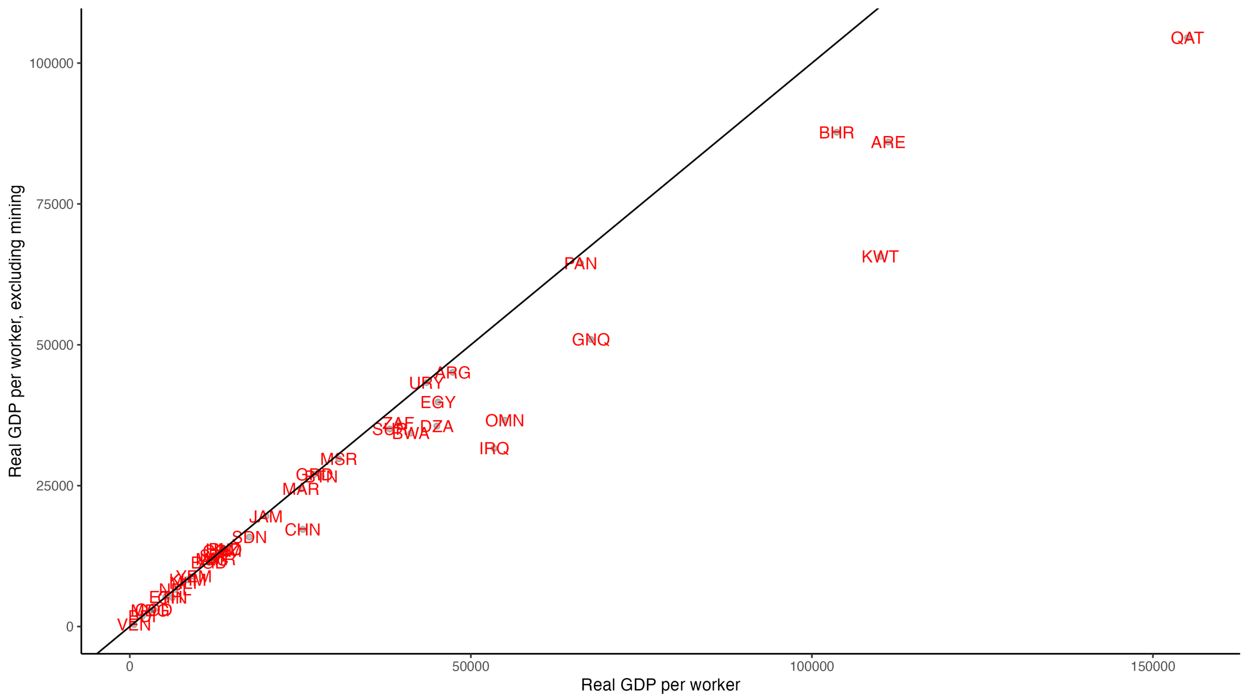

7.2 Culture

Cultural attitudes are persistent features of peoples' beliefs, preferences, and attitudes. These preferences affect peoples' economic choices at an individual level - for example, what sorts of jobs people choose (or whether they choose to work at all), and what sorts of transactions they engage in.

Religion is one facet of cultures that has often been theorized to affect economic outcomes. Max Weber's assertion that the "Protestant Work Ethic" could explain disparate economic outcomes across European countries is one religious theory of economic growth.

There are more complicated theories linking religious beliefs to growth. For instance, countries with more developed financial systems might have higher savings rates because people can earn a return on saving. But religion can influence the development of financial systems; for instance, there were religious restrictions in Catholic countries that forbade lending at interest for many centuries (that variably applied to Catholics and non-Catholics), which may have affected the development of financial systems (and growth) in those countries. Today, many Islamic countries have different degrees to which lending at interest is illegal, and so-called "Islamic banking" systems exist to facilitate saving and investment using alternative arrangements to those in non-Islamic systems. The differences in religion may lead to differences in financial arrangements that affect saving, investment, and growth.

Religion, however, is not the only aspect of culture that might matter. For example, the economists Yuriy Gorodnichenko and Gérard Roland used data on cultural attitudes collected by sociologists across countries to characterize countries' (average) attitudes along a "collectivist-individualist" axis. They argue that more individualist societies encourage innovation and discovery by socially rewarding individuals who stand out. Collectivist societies encourage conformity, which discourages innovation but may make it easier to solve "collective action" problems because individuals think about society's interests when making choices. On balance, Gorodnichenko and Roland find that more individualist countries tend to have higher GDP per worker.

7.3 Institutions

"Institutions" is a broad category of features of social organization that affect (1) economic arrangements (property rights, barriers to entering markets) (2) political arrangements (constraints on politicians and elites). For example, property rights (the extent to which people control their property, and the ease with which that property can be legally or extra-legally taken from them) affect the incentives people have to acquire and use property. Institutions like the choice of legal system, constraints on executive autority, and the extent to which those constraints are effectively enforced can affect property rights.

Proponents of institutional theories often point to historical examples - such as the divergent outcomes of North and South Korea, East and West Germany, or communities on either side of the US-Mexico border - to support the case for institutions shaping economic growth. One such historical example economists have studied is the divergent effects of European colonization. Daron Acemoglu, Simon Johnson, and James Robinson argued in an influential paper that colonizing powers set up two kinds of colonies: "Neo-Europes," which were intended for permanent settlement by Europeans, and "extractive" colonies. The latter had institutions that facilitated taking wealth out of the country, and had few restrictions on executive power. The institutions they set up in the two kinds of colonies were persistent even after colonization ended. For example, the institutional setup in the United States and Kenya were quite different, despite both being British colonies. Acemoglu et. al. argue that the differences in colonizing institutions have had a powerful causal effect on economic outcomes as a result.

7.4 What about policy?

Policy choices are related to institutional choices, but are more mutable. The typical story for institutions emphasizes long-standing, hard-to-change arrangements. On the other hand, some economists and policymakers have emphasized the importance of particular policy reforms that could help countries develop – for example, removing restrictions on particular markets. One set of such policies colloquially became known as the "Washington Consensus" which included smaller budget deficits, controlling inflation, tax reforms, and liberalization of financial markets and trade.[9] It's easy to understand the perceived attractiveness of the policy idea. For example given a widespread consensus in the economics profession that, for example, allowing trade with other countries is good, removing taxes on traded goods to encourage trade may seem like a logical way to boost growth.

The evidence on how effective these reforms were is mixed. The economist Dani Rodrik declared, in 2006, that

> Proponents and critics alike agree that the policies spawned by the Washington Consensus have not produced the desired results... it is fair to say that nobody really believes in the Washington Consensus anymore. The debate now is not over whether the Washington Consensus is dead or alive, but over what will replace it

On the other hand, a 2019 working paper by the economist William Easterly reviewed recent evidence and said

> [...] there is a strong correlation between improvements in policy outcomes and changes in growth outcomes.

The question of policy is also a question about correlation versus causality (as Easterly's quote alludes to!). Countries may have good policies because they have good institutions, for example.

7.5 Summing up

Economists have searched for "fundamental" causes of macroeconomic growth. Three broad categories of theories emphasize geographic, cultural, and institutional fundamentals that generate the outcomes in long-run growth and the distribution of wealth across nations. Economists and policymakers also have debated whether, conditional on a country's history and geographic features, whether policy reforms can result in higher growth.

Although there are partisans for each point of view (and the arguments between their proponents can be caustic), it's probably worthwhile to think of each as having some ability to explain outcomes in different countries. Certainly, "institutional" theories are quite powerful for explaining certain cases (like the split

[9] John Williamson, the coiner of the phrase "Washington Consensus," wrote a number of articles about what the term meant and how it was (mis)used by others.

between North and South Korea's outcomes). But most economists who have tested these theories have found that each explanation has explanatory power even after controlling for the others.

Part of the difficulty in identifying a single fundamental cause is that each seems related. The line between culture and institutions can be somewhat fuzzy, for instance. The Nobel-prize winning economic historian Douglass North, in fact, defined institutions as "the rules of the game in a society or, more formally, the humanly devised constraints that shape human interaction." We know that things we might call "culture" can influence (and are influenced in turn) by institutional environments. And Acemoglu, Robinson and Johnson's theory of why certain colonies were set up the way they are is linked to the disease environment - a geographic feature of the country in question.

A Math appendix

A.1 Exponential functions

If we're working with the Cobb-Douglas production function specifically, we end up using exponential functional forms, and so it can be handy to remember some basic rules of exponents.

Exponential functions are a compact way of writing products. For example:

$$10 \times 10 \times 10 = 10^3$$

We can also have fractional exponents. You may recall that the square root function is

$$\sqrt{(}x) = x^{\frac{1}{2}}$$

More generally

$$\sqrt[n]{x} = x^{1/n}$$

.

Some important rules that you could derive from the definition:

- $x^{-n} = \frac{1}{x^n}$
- $x^a \times x^b = x^{a+b}$
- Combining the two previous rules, $\frac{x^c}{x^d} = x^{c-d}$
- $x^0 = 1$ in general (to be safe, say this is true for $x \neq 0$.)

Remember that $x^a + x^b \neq x^{a+b}$. We add the powers across multiplication, not addition.

An important thing to remember is that something growing at a constant growth *rate* grows exponentially over time. That is, if the stating value is x_0, t time periods have passed, and the variable grows at a rate g every time period:

$$x_t = x_0 \times (1+g) \times (1+g) \times \ldots \times (1+g) = x_0(1+g)^t$$

A.2 Logarithmic Functions

Logarithmic functions are related to exponential functions, so we'll often use them together. The relationship is that: if $x^b = z$, then

$$\log_x(z) = log_x(x^b) = b \times log_x(x) = b \times 1 = b$$

Other useful properties of logs for our purposes are the following (illustrated using a log of base a; think about $a = 10$ if that's helpful, but this is true for any log.)

$$a^{\log_a x} = x$$

$$\log_a(x^y) = y \times log_a(x)$$

$$\log_a(xy) = \log_a x + \log_a y$$

Be careful: $\log_a(x+y) \neq \log_a x + \log_a y$

We will also often use logarithms to make plots of data easier to view. Taking our exponential growth example, notice that if we were to plot x_t against time, it would increase exponentially: the slope would be increasing over time. But if we plot $\log_a(x_t)$, the slope is constant and equal to the growth rate. Moreover, plotting on a logarithmic scale (or a ratio scale) means that equal spaces on a vertical axis are proportional changes in the variable rather than constant changes.

We'll also often use $\log_{10}$, which is handy for plotting. Going from $log_{10}(x) = z$ to $\log_{10}(y) = z + 1$ implies that $y = 10x$.

B List of variables

This is just a short list of variables that are used in the text.

Remember that capital letters represent aggregates, lower case letters represent per-worker quantities, and a t subscript means at time t(e.g., x_t is the quantity of x in year t).

- Y_t - aggregate GDP
- L_t - Number of workers
- y_t - GDP per worker, Y_t/L_t
- K_t - aggregate stock of physical capital
- k_t - physical capital per worker, K_t/L_t
- H_t - Human capital
- h_t - human capital per worker, H_t/L_t
- A_t - Productivity (Note that we don't think about "productivity per worker.")
- I_t - Aggregate investment in physical capital
- i_t - Investment in physical capital per worker, I_t/L_t
- C_t - Aggregate consumption
- c_t - Consumption per worker, C_t/L_t.
- $Y_t = f(A_t, L_t, K_t, H_t)$ is the production function
- γ (lower-case Greek letter "gamma") - The fraction of output that is invested in a period, $I_t = \gamma Y_t$, $0 < \gamma \leq 1$
- δ (lower-case Greek letter "delta") - The fraction of the capital stock that wears out every period, $0 < \delta \leq 1$. The total capital that wears out in a period is δK_t.
- α (lower case Greek letter "alpha") - The exponent on physical capital in the Cobb-Douglas production function. $0 < \alpha < 1$

C An extension to population and productivity growth

The canonical Solow model includes exponential growth in population and productivity growth. That is, $A_{t+1}/A_t = 1 + g$, where g is a constant, and $L_{t+1}/L_t = 1 + n$ where n is a constant. This changes the way we solve the model a little bit. We will assume that technology is "labor augmenting" – that is, having higher productivity is like having extra workers. We then define A_tL_t as the number of "effective workers" in the economy. Steady state in the model is therefore characterized by a constant level of capital (and hence output) per unit of effective labor, rather than per-person.

It's easiest to see this by assuming Cobb-Douglas production and assuming constant human capital per worker. Assume the aggregate production function is

$$Y_t = K^\alpha (hA_tL_t)^{1-\alpha}$$

This is written a little differently than before, because we're explicitly writing productivity as something that augments the number of workers in the economy. Notice that $h \times A_tL_t$ is the total amount of human capital per effective worker.

To get this in terms of output per effective worker, divide both sides by A_tL_t:

$$\frac{Y_t}{A_tL_t} = \frac{K^\alpha (hA_tL_t)^{1-\alpha}}{(A_tL_t)^\alpha (A_tL_t)^{1-\alpha}}$$

Then using lower-case variable names with tildes to denote things in per-effective-unit-of-labor units:

$$\widetilde{y_t} = \widetilde{k_t^\alpha} h$$

We also need to re-write the law of motion for capital in terms of units of effective labor.

Here, we're going to use a fact: Since $\widetilde{k_t} = K_t/(A_t L_t)$, $\ln(\widetilde{k_t}) = \ln(K_t) - \ln(A_t) - \ln(L_t)$

If we take a derivative of this expression with respect to t (notice that each of these variables is a function of t), we get

$$\frac{\Delta k_t}{k_t} \approx \frac{\Delta K_t}{K_{t-1}} - \frac{\Delta A_t}{A_{t-1}} - \frac{\Delta L_t}{L_{t-1}}$$

(This becomes exact as $t \to 0$.)From the law of motion for (aggregate) capital, $\Delta K_t/K_t = \frac{I_{t-1}}{K_{t-1}} - \delta$. So then we can write

$$\widetilde{k}_t - \widetilde{k}_{t-1} \approx \frac{I_{t-1}}{K_{t-1}} \cdot k_{t-1} - (\delta + n + g)\widetilde{k}_{t-1}$$

Multiplying the (I/K) term by $(1/AL)/(1/AL)$ we get $\widetilde{i}_{t-1}/\tilde{k}_{t-1}$. Simplfying we get

$$\widetilde{k}_t = \widetilde{k_{t-1}} + \gamma\widetilde{y_{t-1}} - (n + g + \delta)\widetilde{k_{t-1}}$$

This looks very similar to the aggregate version of this expression, except now we have the $n + g$ term. The intuition is that the amount of units of capital per effective worker may decline for three reasons. One is that it may wear out (δ). A second is that we have more actual human beings (n). That means a given amount of machines would have to be spread around more people. The third is each existing worker is more productive (g). This means that a given amount of machines will be spread around more effective workers. A steady state will be when the amount of investment is enough to offset the wearing out of machines and the arrival of new (effective) workers.Graphically, the story is the same as before except we now find steady state at the intersection of $\gamma\widetilde{k}^\alpha h$ and $(n + g + \delta)\widetilde{k}$. We no longer consider shifts in A_t because it grows at a constant exogenous rate, but we can talk about what happens if that rate (or the rate of population growth) changes.

Finally, notice that if per effective worker variables are constant in steady state, that means that aggregates are growing at the rate $n + g$. That way, $\widetilde{x} = X/(AL)$ remains constant for any $\widetilde{x}$. This means that the Solow model predicts that aggregate GDP grows at the rate of $n + g$ in the long run. On a per capita basis (for similar reasons), the Solow model predicts that per worker variables grow at a rate g. Hence, the Solow model says that in steady state, the rate of productivity *growth* is what determines growth in GDP per capita. Hence, we still need productivity to change over time to explain sustained growth in GDP per capita over the long run.