Economics 1011B
Section 1

Spring 2023

Introduction

- Welcome to Ec1011B!
- A big virtue of this course is that much of the material is self-contained in the lecture slides and the relevant sections of the textbook.
- In general, we will not discuss new material in section. Instead, focus on reviewing key ideas and technical tools that we have developed in the previous week, and that are likely to be useful for preparing for the midterms/final.
- Where useful, we will also try to provide an additional perspective on the material we have covered.

Section Logistics

- Both of us (TFs) will hold two sections each week on Thursday/Friday:
  - Mike's sections: 1:30pm and 3:00pm on Thursdays, Sever 306.
  - Jake's sections: 12:00pm and 1:30pm on Fridays, Sever 206.
- Mike's email is mdroste@fas.harvard.edu. Jake's email is jacob_furst@fas.harvard.edu. Feel free to e-mail us at any time with any questions you have!
- I (Mike) am free any time to schedule 'ad hoc' meetings to discuss the class, problem sets, economics, or anything else on your mind. I respond quickly and have no judgment whatsoever if you're emailing me on the day a problem set is due, so feel free to reach out any time.

Course Logistics

- Course Assistants: We have six wonderful course assistants (CAs) this semester. They will contribute by holding very regular office hours each week (syllabus), plus more for the psets (below).

- Office Hours: The professors, the TFs, and CAs will hold office hours many times a week, as listed on the syllabus.

- More Office Hours: In addition to the regular office hours held by CAs, the professors, and the TFs, the course assistants will hold additional office hours before each problem set and exam to help you prepare.

Course Survey

- If you have a few minutes, I have a short and completely optional course survey I'd love for you to take. I just want to get a sense for what brought you to this class and if any particular topics excite you, so that we can spend more time on them in section!
- The link is here (this is clickable from the pdf): https://forms.gle/esm2abZBB5JdRwHT7
- Based on responses from Thursday's section, I am likely to change or add additional office hours to make them more convenient.

Section Slides

- We will prepare section slides each week (like these ones), and they will be made available on the course Canvas page before section.
- At the end of the prepared slides, we will sometimes include a few 'bonus slides' that attempt to connect the material to the frontier of macroeconomics research, or important historical context that we won't cover in class.
- This material is completely optional and we will not cover it in section, it is just intended for you to read in case you're curious.

Today's Outline

- Course / Section Logistics
- Production Functions
- Solow Model
  - Setup
  - Steady State
  - Dynamics
- Intro to LaTeX and Overleaf
- Bonus Slides
  - Aggregate Production Functions

Solow Model

- This week we'll talk about the Solow model - an important and foundational model that is the starting point for thinking about economic growth.

- Economic growth deals with the important big-picture questions that Ludwig discussed on the first day of class: Why do some nations grow? Can societies grow forever with capital accumulation? Is there an 'optimal' level of savings, or is it possible for society to save "too much"?

- The Solow model is very simple, and in some respects unique relative to the other models we will discuss in this course. There is no explicit optimization - everything is done by positing relationships between aggregate variables rather than describing the behavior of individual consumers or firms.

- In the next two weeks, we will develop a model of consumer behavior that we can synthesize with the Solow model to deliver us a 'complete' model of consumption, savings, output, and growth from first principles.

Production Functions

- Key conceptual ingredient of the Solow model: production functions.
- A production function tells us something about the productive capacity of an economy with labor supply L and capital stock K, and technology (or productivity) A. The quantity $\bar{F}(K, A, L)$ is a generic production function that tells us how much output
- The existence of an aggregate production function is a strong modeling assumption! Not without loss of generality. We will gloss over this completely, but see this week's Bonus Slides for more.
- We will assume that technology enters in a labor-augmenting fashion in our models, working with a production function of the form $F(K, AL)$. As Ludwig notes, this is also a nontrivial assumption - more in a bit.

Production Functions: Nice Properties

- Well-behaved production functions exhibit positive marginal products (first derivatives positive) and diminishing marginal products (second derivatives negative).

  $\implies$ *More inputs increase output, but at a diminishing rate.*

- Very well-behaved production functions obey the Inada conditions, which are additional nice-to-have properties describing the limiting behavior of the first derivatives of F:

$$\lim_{K \to 0} F_K(K, AL) = \infty \quad \text{and} \quad \lim_{AL \to 0} F_L(K, AL) = \infty$$
$$\lim_{K \to \infty} F_K(K, AL) = 0 \quad \text{and} \quad \lim_{AL \to \infty} F_L(K, AL) = 0$$

- If I give you a specific functional form for a production function, we can check and verify whether these general properties hold.

Production Functions: Returns to Scale

- One final property of production functions that deserves its own slides: returns to scale. It is often convenient to characterize a given production function by whether it exhibits constant, increasing, or decreasing returns to scale.

- This answers the question: if we multiplied our production inputs by some factor λ, would we get more than, less than, or exactly λ times the output we had before?

- Formally, F has constant/increasing/decreasing returns to scale if, for any λ and any (K, AL), one of these conditions holds:

$$F(\lambda K, \lambda AL) = \lambda F(K, AL) \qquad \text{(constant returns to scale)}$$
$$F(\lambda K, \lambda AL) > \lambda F(K, AL) \qquad \text{(increasing returns to scale)}$$
$$F(\lambda K, \lambda AL) < \lambda F(K, AL) \qquad \text{(decreasing returns to scale)}$$

- Intuition: with CRS, doubling the inputs yields double the output.

Production Functions: Cobb-Douglas

- In this course, we will typically work with a Cobb-Douglas production function (with labor-augmenting technology), which takes the form:

$$F(K, AL) = K^{\alpha}(AL)^{\beta}$$

- F has strictly positive first derivatives, strictly negative second derivatives with respect to each (non-negative) argument (concave), and satisfies the Inada conditions.

- F exhibits constant returns to scale if $\alpha + \beta = 1$, increasing returns to scale if $\alpha + \beta > 1$, and decreasing returns to scale if $\alpha + \beta < 1$.

- Useful exercise: Show that the Cobb-Douglas production function satisfies all these properties.

- For the Solow model, we will assume that F exhibits *CRS* by assuming the exponents sum to 1. Sometimes we will express the Cobb-Douglas exponents as α and $1 - \alpha$.

Production Functions: Constant Returns to Scale

- There is one very useful property of constant returns to scale production functions that we will exploit in the Solow model.

- Recall that a constant returns to scale production function satisfies:

$$F(\lambda K, \lambda AL) = \lambda F(K, AL) \qquad \text{for any } \lambda \text{ and any } (K, AL)$$

- Choose $\lambda = 1/L$. Then we can rewrite the production function as:

$$F(K, AL)/L = F(K/L, A)$$

- In words: output $F(K, AL)$ is a function of two things, K and AL, but when F is CRS, output per worker $F(K, AL)/L$ can be written as a function of K/L and a constant A - easier to work with.

- Keep this insight on the back burner for now.

Solow Model: Production

- We've finished talking about production functions for now - let's set up the Solow model using that language we just learned.

- We start by assuming that each period t, Y_t goods are produced according to a CRS Cobb-Douglas production function involving capital K and 'effective labor' AL:

$$Y_t = F(K_t, AL_t) = K_t^{\alpha}(AL_t)^{1-\alpha} \tag{1}$$

where K_t is capital, L_t is labor, and A is technology (constant and exogenous).

Solow Model: Investment and Production

- Next, we assume that existing capital stock depreciates at a constant rate δ each period. The *law of motion for capital* is:

$$K_{t+1} = (1 - \delta)K_t + I_t \tag{2}$$

  where I_t is new investment into the capital stock. This equation says that capital next period is equal to undepreciated capital this period plus investment.

- What determines investment? We will assume that an exogenous fraction s of output is saved and invested into the capital stock:

$$I_t = sY_t \tag{3}$$

Solow Model: Investment and Production

- Lastly, we will assume that the population (or labor supply) grows at an exogenous rate n:

$$L_{t+1} = (1+n)L_t \tag{4}$$

- Another common variant of the Solow model allows technology to grow at an exogenous rate (call it g) - just for simplicity, we will exclude this from our classroom Solow model. Return to discussing that simple extension at the end of class.

Solow Model: Intensive Form

- As Prof. Straub discussed today, one distinctive trick in the Solow model is to redefine variables. This is called 'detrending' the model or casting it in 'intensive form'.

- Define output per worker $y_t \equiv Y_t / L_t$, and capital per worker $k_t \equiv K_t / L_t$.

- Why are we doing this? Recall from the previous slide that we've assumed L grows forever at rate n – it never settles down to a finite number.

- This will imply (through the other equations - how?) that capital K, output Y, and investment I grow forever, too.

- When we divide K (and other objects) by L, it turns out that this object ('capital per worker', or the capital-labor ratio) will converge to a unique, finite value - precisely because we've divided by the stuff that keeps on growing. We can solve for this value (in a few slides).

Solow Model: Taking Stock (In Words)

- Our economy starts with some initial stock level of technology, labor supply, and capital: A, L_0, K_0.

- Each period, economy produces output Y_t using a CRS Cobb-Douglas production function $F(K_t, AL_t)$.

- A constant fraction δ of capital K depreciates each period (crumbles to dust). Capital next period is equal to undepreciated capital from this period plus investment, which is a fixed fraction s of output this period.

- The population or labor supply L grow at a constant rate n each period.

Solow Model: Taking Stock (In Math)

- Solow model is characterized by the following equations (given K_0, L_0):

$$
\begin{aligned}
Y_t &= K_t^{\alpha}(AL_t)^{1-\alpha} && \text{(Output is CRS Cobb-Douglas)} \\
K_{t+1} &= (1-\delta)K_t + I_t && \text{(Law of motion for capital)} \\
I_t &= s \cdot Y_t && \text{(Investment is some fraction of output)} \\
L_{t+1} &= (1+n)L_t && \text{(Law of motion for population)} \\
k_t &\equiv \tfrac{K_t}{L_t} && \text{(Definition of capital per worker)} \\
y_t &\equiv \tfrac{Y_t}{L_t} && \text{(Definition of output per worker)}
\end{aligned}
$$

- Endogenous variables: K_t, L_t, Y_t, I_t (and thus by definition also k_t, y_t).

- Key parameters: α, δ, s, n, A.

Solow Model: Steady States

- The Solow model is dynamic - it describes how quantities change over time. A steady state is an important concept in dynamic models: it describes values that endogenous variables converge to in the long-run.

- There exists a unique steady state for k_t and y_t in Solow. Solving for these will be very useful for us to characterize how quantities like Y/L or Y change over time.

- The steady state k^* is defined as the value of k_t such that $k_{t+1} = k_t = k^*$. Ludwig's slides provide one way to solve for k^*: first solving for k_{t+1} as a function of k_t and parameters, then imposing $k_t = k_{t+1} = k^*$.

- I'll try to be more explicit about these steps.

Solow Model: Solving for Steady State

- First step: get an equation for k_{t+1} in terms of y_t, k_t and parameters.

- Start from definition of k_{t+1} then use model equations to rewrite:

$$
\begin{aligned}
k_{t+1} &= \tfrac{K_{t+1}}{L_{t+1}} && \text{(definition of } k\text{)} \\
&= \tfrac{(1-\delta)K_t + I_t}{L_{t+1}} && \text{(law of motion for capital)} \\
&= \tfrac{(1-\delta)K_t + s \cdot Y_t}{L_{t+1}} && \text{(investment fixed fraction of output)} \\
&= \tfrac{(1-\delta)K_t + s \cdot Y_t}{(1+n)L_t} && \text{(law of motion for labor supply)} \\
&= \tfrac{1}{1+n}\Big(s \cdot y_t + (1-\delta)k_t\Big) && \text{(definitions of } y,\ k\text{)}
\end{aligned}
$$

Solow Model: Solving for Steady State

- Second step: rewrite y_t in terms of k_t:

$$\begin{aligned} y_t &= \tfrac{Y_t}{L_t} && \text{(definition of } y\text{)} \\ &= \tfrac{F(K_t, AL_t)}{L_t} && \text{(output is CRS Cobb-Douglas)} \\ &= F(k_t, A) && \text{(CRS: divide arguments by } L\text{)} \\ &= k_t^{\alpha} A^{1-\alpha} && \text{(plug in production function)} \end{aligned}$$

- Plugging this equation into the previous slide, get key difference equation:

$$k_{t+1} = \tfrac{1}{1+n}\left(s \cdot k_t^{\alpha} A^{1-\alpha} + (1-\delta) k_t\right) \tag{5}$$

Solow Model: Steady State

- How do we solve for the steady state value of k, denoted k^*? The steady state k^* is defined such that $k_{t+1} = k_t = k^*$. So plug in $k^* = k_t = k_{t+1}$ into equation 5 and then rearrange using algebra to solve for k^*.

- The steady state value of k, in terms of model parameters, is:

$$k^* = \left(\frac{s}{\delta + n}\right)^{\frac{1}{1-\alpha}} A$$

- For a CRS Cobb-Douglas prod fn., $y^* = F(k^*, A)$ (make sure you see why!), so it's easy to figure out y^*:

$$y^* = F(k^*, A) = \left(\frac{s}{\delta + n}\right)^{\frac{\alpha}{1-\alpha}} A$$

Solow Model: Key Graph

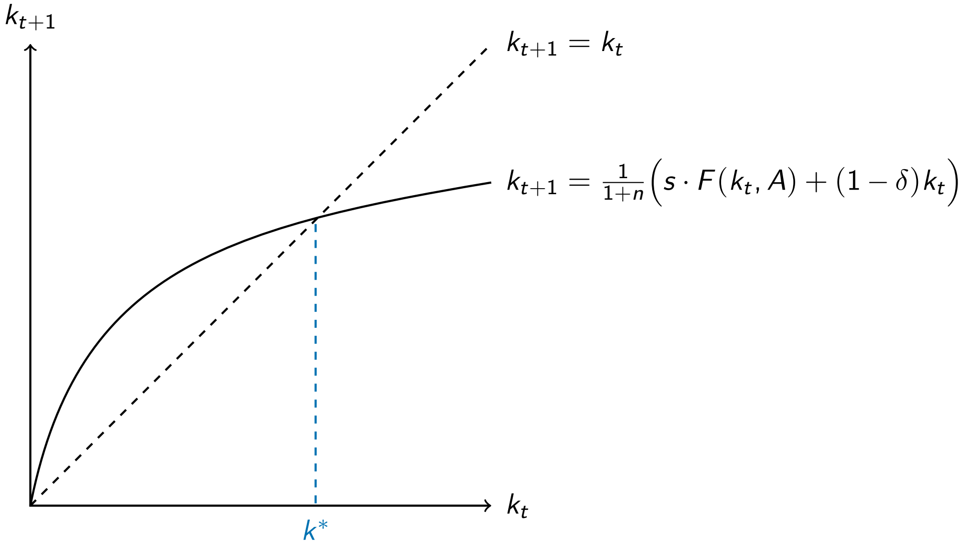

Solow Model: Some lessons

- What does the Solow model say determines economic growth in the short-run and in the long-run?
- In the long-run, income/output per capita, y, is constant: it settles down to a steady state. But in the short-run, output per worker y and capital per worker k can increase (or decrease), depending on where those quantities are relative to their steady states.
- Why? In the short-run, growth is possible due to **capital accumulation**: the marginal product of capital $F_K = \partial F / \partial K$ is high when k is low (Q: why?).
- But eventually, diminishing marginal product of capital 'chokes off' the potential for growth due to capital accumulation, since the constant rate of depreciation of capital ends up offsetting the gains from the marginal product of capital.
- The 'Golden Rule' (see Kurlat 4.3) suggests that there is an optimal savings rate (not 0% or 100%!) that maximizes long-run consumption. What's that about? Two slides.

Solow Model: Comparative Statics

- We can ask a number of questions immediately with the Solow model.
- For each of these cases, you can always get the sign of the right answer graphically (i.e. how the concave line shifts), or use calculus to compute, e.g., $\partial k^* / \partial p$ for any model parameter p, like the savings rate s, the population growth rate n, etc.
- Let's think them through graphically. What happens to k^* as n goes up?
- What about n or δ?
- What about the savings rate s? That deserves its own slide!

Solow Model: Golden Rule Savings Rate

- Recall s is exogenous in this model - just a fixed number. Nonetheless, we can ask: if we could choose s, what maximizes consumption (per worker), which is whatever is left over after saving?

$$c_t = (1 - s)y_t \tag{6}$$

- You can set this up as a maximization problem and solve for s. For any general f, the optimal choice of s leads to a level of capital per worker we'll denote k_{GR} that satisfies $F_K(k_{GR}, A) = n + \delta$, where F_K is the derivative of F with respect to K (the marginal product of capital).

- For the Cobb-Douglas production function, this implies the consumption-maximizing savings rate is $s_{max} = \alpha$ - the parameter from the production function!

- This implies there is such a thing as too much savings in this model: if $s > \alpha$, c^* will be smaller than if $s = \alpha$. Why?

LaTeX

- Now, let's change gears and talk about LaTeX!
- We are requiring that problem sets my typeset in TeX for this course. If you haven't used TeX before, don't worry - we'll provide lots of resources for you to help learn.
- I would highly recommend learning TeX through the website Overleaf, which functions more or less like Google Docs for LaTeX.
- Fundamentally, LaTeX is a scripting language that allows you to create publication-quality documents and slides. The vast majority of academic papers and textbooks you've seen in your life have been formatted in TeX.
- Unfortunately, there's no way we can teach you LaTeX in a few minutes - the only way to learn is to deal with the frustration of having your problem set not compile (believe me, I remember my first semester using LaTeX very well!)

LaTeX

- An important element of learning LaTeX is to start from example 'code' (i.e. the .tex file we use to create your problem sets) and learn how to modify it.

- We will provide the underlying source TeX code for all problem sets that we post in this class. You are welcome to typeset your answers directly on the p-set (e.g. underneath each question), and use the equations present on the problem set to help you learn the syntax for writing math.

- Remember also that we encourage collaborating with friends on problem sets to a certain extent. Collaboration is welcome because you can learn a lot from your classmates. However, we will require that each person in the course submit their own typset problem set.

- You should reserve a few hours after completing each p-set (e.g. on paper and pencil, or a tablet, or a whiteboard) just to typeset your solutions. This is really annoying at first - I promise you will get faster!

Bonus Slides

Bonus Slides: Aggregate Production Functions

- This week's bonus slides concern the aggregate production function F that we have seen in the Solow model (and will see again throughout the course).

- These bonus slides are entirely optional and intended to give you some food for thought. None of this will be covered on exams.

- Here are two sets of questions related to this week's material that I think are fun to consider:
  1. To what extent does aggregate US data (on income, capital, and labor) appear to be approximated by an aggregate production function?
  2. Is the notion of an 'aggregate production function' economically meaningful? If firms each have their own production technologies that are well-behaved, does this imply that there is an aggregate production function? Is it meaningful to aggregate disparate notions of capital into a single variable?

Bonus Slides: Aggregate Production Functions

- First, let's take as given that the U.S. data can be approximated by, for instance, a Cobb-Douglas production function, $Y_t = K_t^{\alpha} L_t^{\beta}$.

- We observe data on output Y_t, labor L_t, and capital K_t. How do we estimate the coefficients in the production function?

- Easy to do with linear regression! A Cobb-Douglas function is log-linear; that is, if we take logs of both sides, we see that $ln(Y)$ is an linear function of $ln(L)$ and $ln(K)$, specifically, $ln(Y) = \alpha ln(K) + \beta ln(L)$, naturally motivating a linear regression specification that we can use to estimate α and β.

- ... But we're assuming a functional form for aggregate output! Whether this is reasonable depends on our objective (e.g. do we want to interpret the estimated coefficients as the output elasticity of labor/capital?) Pol Antras (in our department) wrote a nice paper asking whether or not the aggregate production is well-approximated appears well-approximated by a Cobb-Douglas function summarizing many of the key issues involved (link)

Bonus Slides: Aggregate Production Functions

- These questions turn out to have been the focus of an intense academic debate that occurred in the 1950's and 1960's, what is now sometimes called the Cambridge capital controversy.

- One one side were the 'marginalists' in Cambridge, Massachusetts - eminent macroeconomists that include Franco Modigliani, Paul Samuelson, and Bob Solow. They were occupied with analyzing the behavior of an economy described by the models we have studied in this course.

- On the other side, the 'post-Keynesians' based primarily in at Cambridge University in the U.K - notably, Joan Robinson[1] and Pierro Sraffa.

[1] Joan Robinson was a fascinating economist, a pioneer in industrial organization and the study of imperfectly competitive markets - also a controversial figure worth reading about!

Bonus Slides: Aggregate Production Functions

- Joan Robinson kicks off the debate in 1953: claims that the aggregate production function is a 'powerful tool for miseducation' in macroeconomics.

- The heart of her critique is sometimes called the aggregation problem. To what extent is it meaningful to aggregate the capital stock into a single unit, or to assume that the production processes of many firms can be summarized by a single aggregate production function?

- Here is one way to think about it: Suppose that there are N firms in the economy which each produce their own output using capital and labor. Each firm i produces output according to a well-behaved micro production function $Y_i = F(K_i, L_i)$. Under what conditions does this 'aggregate up' to a single production function relating aggregate output to aggregate capital and labor with desirable properties?

- Robinson demonstrated with simple examples that sometimes, individual firms can face well-behaved production functions, and the economy's 'aggregate' production function can be very poorly-behaved (i.e. constant or increasing marginal products).

Bonus Slides: Aggregate Production Functions

- So the Cambridge UK folks won - the aggregate production function is an assumption, and the assumptions we impose are nontrivial to justify in aggregate. Paul Samuelson wrote a paper summarizing the debate, largely conceding to the points made by Robinson and others (link).

- Despite all of this, many economists still work with models that assume the existence of aggregate production functions. Why? 'You need a model to beat a model' - nobody has come up with anything better to replace it with, at least as a general-purpose tool.

- The aggregate production function is still relevant today. The late great Harvard economist Emmanuel Farhi wrote an interesting paper providing conditions under which well-behaved aggregate production functions arise from the problems facing individual firms.