Lecture 1: Firms and Production[1]

1 Overview – what are we doing here?

1.1 everyone optimizes; so what?

- Economic theory: "Everyone optimizes – so what?"
- That is, what does it mean to assume that everyone is optimizing,

  and what are the implications of that assumption?
- You'll be confronting that same question in a number of different contexts this year:
  - Consumer and producer theory: what does it mean for an individual consumer, or an individual firm, to optimize given the prices they see?
  - Game theory and information economics: what does it mean to optimize when other peoples' behavior affects you and they're optimizing too?
  - General equilibrium: if everyone's optimizing, how does the market clear?
- More specifically, for right now:

  if producers optimize given their technological abilities,

  or consumers optimize given their preferences and budget constraint,

  what does that imply about how they will behave?

[1] These notes borrow from notes prepared by Jonathan Levin, Paul Milgrom, and Ilya Segal for the first-year PhD micro sequence at Stanford University. Typos and other errors are mine, though.

1.2 economics $\approx$ theory $\cap$ data

- So that's the first question: what does the assumption of optimization –
  either profit maximization or utility maximization –
  imply about behavior?
- But there's a flip side of that question as well, which is,
  what restrictions does this assumption place on observables you might see,
  and to what extent would observables allow you to infer the "hidden" part of the model,
  either individual preferences or producer technologies?
- And this hints at a point I want to make right at the outset:
  theory isn't just for theorists

- Most of you won't go on to be theorists,
  and that's fine – the world needs theorists, but maybe not so many
- But I, at least, think a lot of economics
  is really about the intersection between theory and data

- That is, there are a few people doing very abstract high theory that never touches data;
  and a few people doing RCT-type experiments with no need for theory,
  or data mining with atheoretic machine learning algorithms;
  but for the most part, most of economics involves both theory and data
- Empirical observation can inform theory –
  hinting at what situations are interesting to apply theory to understand,
  what assumptions might be reasonable, and so on
- And theory can tell you what predictions to test empirically,
  and give you a structure within which to do empirical estimation

- The theorists here are at UW are applied theorists –
  we're interested in using models to better understand parts of the real world;
  we rely on observation, at least casual observation, to inform our models,
  and hope that more systematic observation will prove consistent with them
- Most of the macro people here use theory models as well,
  just with different questions in mind
- And most of the empirical people here use structural models,
  which are basically theory models with parts that need to be estimated from data

- So when we think about theoretical models of profit-maximizing firms or utility-maximizing consumers,
  we won't just be looking for theoretical predictions of the model
  in order to know how people and firms "should" behave,
  but we'll be thinking about when data might be consistent with those models,
  or be inconsistent with those models and thus tell us those models are wrong;
  and when data might let us learn the unknown parts of those models –
  infer firms' technological capabilities or individuals' preferences,
  from observation of their choices

- So for those of you interested in micro theory, great,
  hopefully you'll love this class;
  and for those of you not interested in micro theory, that's fine –
  hopefully this class will be full of stuff that is useful to you when you do empirical work

- I used to teach consumer theory before producer theory;
  starting two years ago, I teach producer theory first
- While people naturally identify with being a consumer first,
  the mathematical tools build more naturally when we start with the firm's problem,
  so that's the way I'll be teaching it

- and from there, onward!

2 Basic model of firms and production

- I'll jump right into the basic model we'll use for firms and production,
  and then loop back to discuss some of the key assumptions implicit in that model,
  and whether they make sense

- The basic "unit" we'll think about in producer theory is the firm
- The basic model doesn't consider where firms come from,
  or how they're organized, or why they exist;
  we just take them as given
- And the basic model of production makes three key assumptions about firms,
  which I'll list first and then flesh out:

  1. firms have **fixed, exogenous technological capabilities**
  2. firms are **price-takers**, i.e., they see market prices as outside their control
  3. firms act to **maximize profits**, given the technological capabilities they have and the prices they see

2.1 firms have technologies

- When I say that firms have fixed, exogenous technological capabilities,

  what I mean is that each firm is characterized by its ability to turn certain combinations of goods into certain other combinations of goods

- We'll assume there are a finite number of distinct goods

- Each good is homogeneous – once you specify which good it is, you don't care who made it or where it's coming from, steel is steel and electricity is electricity

- We typically think of the goods as being divisible – you could in principle use any fractional amount of each one – although this isn't essential

- The goods could each potentially either be inputs for the firms or outputs produced by the firm – we don't need to label some goods as inputs and some as outputs, they're all just goods

- (for intuition, it may help to think of inputs like labor, capital, land, and raw materials,

  and outputs that are finished products or services, like Teslas or haircuts;

  but the formal model doesn't distinguish between input goods and output goods –

  there are just a bunch of goods, and a firm chooses to use some of them and produce others)

- And we'll say there are k of them

- A **production plan** for a firm tells you how much of each good the firm will use or produce

- A production plan can be represented as a vector $y \in \mathbb{R}^k$, where each element of the vector tells you how much of good k the firm will produce or use up –

  positive numbers indicate outputs, and negative numbers represent inputs

- So, for example, if $k = 3$ and the goods are steel, labor, and Toyota Corollas,

  a firm might have a production plan of $(-200, -10000, 100)$,

  indicating it plans to turn 200 tons of steel and 10,000 hours of labor into 100 Corollas

- (Another firm, or even the same firm, might also have a production plan of $(100, -5000, -100)$

  it could melt down 100 Corollas for raw steel, but it would have to put in more labor,

  and would get back less steel than went into the cars in the first place.)

- To define a firm's technological capabilities,

  we need to categorize which production plans are actually feasible and which ones aren't

- I'd love to start a business producing iPhones out of beach sand and recycled plastic,

  but I probably don't have the ability to do that

- We define a firm's **production set** as the set of production plans which are feasible,

  that is, the ones which are possible given the firm's technology

- Formally, $Y \subset \mathbb{R}^k$ is the subset of $\mathbb{R}^k$ containing all the production plans that are feasible for the firm

- So a point $y \in Y$ represents something the firm could do – turn some amount of steel and labor into some number of Corollas –

  while points $y \notin Y$ are things the firm wouldn't be able to do

- And the set Y tells us everything we need to know about a firm –

  it fully defines the firm's productive abilities

- For a little visual intuition, here's what a very simple production set might look like,

  in a world with just two goods and a firm that uses one as input to produce the other one:

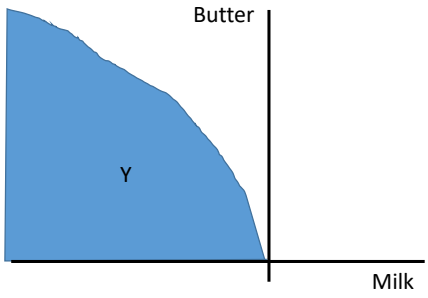

2.2 firms are price-takers

- Next, we make the assumption that firms are **price-takers**:
  firms don't get to set prices; instead, there's a market price for each good,
  and the firm can buy or sell as much of each good as it wants at those prices
- So there are per-unit market prices $p_1, p_2, \ldots, p_k$ for each good,
  which the firm can't control, but just takes as given
- (We'll always assume prices are **weakly positive**,
  and often that they're strictly positive)

- Recall that when we write down a production plan y,
  outputs produced are positive numbers, and inputs used are negative numbers
- So for a good that the firm is producing,
  the revenue the firm receives from selling that good is $p_i y_i$
  and for an input the firm is using,
  $p_i y_i$ is negative, and its magnitude is the cost the firm pays for that input
- So the firm's profit is

$$\begin{aligned}
\text{profits} &= \text{revenues} - \text{costs} \\
&= \textstyle\sum_{y_i>0} p_i y_i - \sum_{y_i<0} p_i(-y_i) \\
&= \textstyle\sum_{i=1}^{k} p_i y_i \\
&= (p_1, \ldots, p_k) \cdot (y_1, \ldots, y_k) \\
&= p \cdot y
\end{aligned}$$

2.3 firms are profit maximizers

- And finally, we make the assumption that firms are **profit maximizers** –
  that the firm's objective is to make as much money as possible,
  so the firm will choose whichever feasible production plan leads to the highest profits
- Or to put it in math terms, the firm solves

$$\max_{y \in Y} p \cdot y$$

- And that's our basic model:
  a firm is defined by a set of feasible production plans,
  faces market prices for each good,
  and chooses whichever production plan maximizes profits

3 Is this a good model?

- George Box (statistician): "all models are wrong, but some are useful"
- Too early to tell if it's useful – you haven't yet seen what the model teaches us
- Obviously the assumptions are strong, but how big a problem are they?

3.1 Let's start by talking about some cool cars

- This is the Tesla Roadster.

- All-electric car, based on the Lotus Elise, introduced in 2008.
- Cost $110,000 at the time, more than three times the price of a baseline BMW 3-series.
  I believe they initially had a range around 100 miles on a charge,
  but subsequent improvements have boosted that to about 215 miles.
- I suspect they're insanely fun to drive –
  they could do 0 to 60 in under 4 seconds,
  and weighed under 3000 pounds
- You've probably never seen one, because Tesla only sold about 2500 of them.
- But that doesn't mean the product was a failure.
- Elon Musk started Tesla; when they rolled out the Roadster, he made it clear that the company's end goal was not "to make toys for rich people."
- The goal of the Roadster was to show that an electric vehicle could actually be desireable – could be awesome.
- And also to get better at building electric vehicles – to do a bunch of R&D that would pay off later.
- And also to make some money, which could then be spend developing the next product.

- This is the Tesla Model S.

- Introduced in 2012.
- They cost $75,000 and up depending on the exact model,
  and Tesla has sold about 260,000 of them.
- I was in Silicon Valley for Fall of 2014, and these cars were everywhere.
- They are amazing.
  They go zero to 60 in 3.2 seconds.
  The range is about 265 miles.
  They have a perfect five-star safety rating.
  My wife test-drove one and loved it.
- The goal of the Model S was to show that a mass-market electric car could be great –
  not just great for an electric car, but a great car.
- And to get even better at building electric cars.
- And to make more money, to put into development on the next product.

- This is the Tesla Model 3.

- Unveiled in March 2016, the base model was intended to cost $35,000.
- Tesla accepted deposits for 325,000 of them in the first week after announcing them.
- Tesla hoped, from the beginning, to build 500,000 per year. That would require the entire world's output of lithium-ion batteries, so Tesla built its own enormous plant to build the batteries.

  (Seriously. They put a $5 billion battery plant in the Nevada desert, to double the world's annual supply of lithium ion batteries, because otherwise they couldn't build as many Model 3s as they want to.)[2]
- They started producing them in summer 2017, and after some production delay and hiccups, have now delivered about 1.3 million of them,

  and are now producing around 200,000 per quarter
- This is the car Elon Musk wanted to build from the beginning.

  With tax credits and gas savings, it can compete with Accords and Camrys
- (Tesla's now moved on to talking about a $25,000 car in the next few years,

  but the Model 3 was the end goal when the company launched –

  to iterate through two earlier models, and then get to the Model 3 as a practical car for normal people)

[2]Sources for Tesla stuff:
http://waitbutwhy.com/2015/06/how-tesla-will-change-your-life.html#part3
https://en.wikipedia.org/wiki/Tesla_Roadster
https://en.wikipedia.org/wiki/Tesla_Model_S
https://en.wikipedia.org/wiki/Tesla_Model_3

3.2 Where am I going with this?

- I'm talking about Tesla to highlight the limitations of the basic production model I just showed you
- Tesla is an illustration that...
  - **technology is dynamic** – it changes over time
    Elon Musk wanted to build the Model 3 in 2008
    It was impossible in 2008, but it's possible now
  - **technology is strategic**
    Elon Musk decided in 2008 that he wanted to build the Model 3,
    so he made decisions that led to it becoming feasible –
    Tesla chose to do things specifically to change their technological capabilities over time
  - **so technology is endogenous** – the result of choices that somebody made,
    and therefore responsive to economic incentives
- The model I just showed you doesn't allow for any of that –
  it assumes technology is static, and exogenous –
  we just take it as given, without thinking about where it comes from or how to influence it
- I could tell a similar story about individuals' tastes – that they're dynamic, and endogenous
- (If preferences weren't endogenous, there'd be very little role for marketing!)
- But when we talk about consumer preferences,
  we'll also be assuming they're static and exogenous –
  we don't know where they come from, and there's no way to influence them

- We also assumed **firms are price-takers**

  That they have no control over the price of goods,

  they can just buy and sell as much as they want at market prices

- The Tesla story also illustrates the fact that this isn't always true.

  Tesla concluded that there wouldn't be enough lithium-ion batteries available at *any* price, which is why they wanted to produce their own.

- When we're talking about small firms and standardized, commodity-type products,

  assuming firms can't affect prices might seem pretty reasonable

- On the other hand, it seems strange to assume Apple has no control over the price they can sell iPhones for, but can sell as many as it wants at that market price

- And firms might also have some market power buying inputs locally –

  for those of you new to Madison, you'll learn about Epic Systems, our local tech firm with 10,000 employees –

  Epic might have some impact on wages around Madison,

  and might face a tradeoff between how much they pay and how many people they hire

- Later on, you'll think about the monopolist's problem, and other models of competition –

  when firms face an explicit tradeoff between price and quantity –

  but for now, we're assuming firms are price-takers

- Finally, even the assumption of profit maximization isn't as innocuous as it sounds

- For decent-sized firms, the owners of the firm aren't the ones making the day-to-day decisions,

  and the day-to-day managers might have other goals besides maximizing profits

- This is called an agency problem, or a principal-agent problem, and there's a large literature on solving it –

  how to create incentives for managers ot act in the interest of owners or shareholders

- (And of course, even without an agency problem,

  some firms may have goals other than profits –

  maintaining a reputation that will be valuable in the long run,

  supporting local causes, etc.)

- for now, we're ignoring all of this

3.3 So if these models are so limited, why learn them?

- Several reasons
- First, to get better at working with models in general
  - part of the point of this class is not just understanding these specific models of consumer or firm behavior,
    but to get used to working with models at all –
    to get practice mapping complex realities to tractible models,
    manipulating those models mathematically,
    using those models to make predictions,
    and understanding the link between those models and hypothetical data
  - So it makes sense to start with the simplest models
  - Even if these models aren't exactly right, they're a good starting point,
    and we can potentially build up from there
  - Also, part of "using" models is thinking critically about them,
    so for our purposes, it's kind of OK these models aren't perfect
- Second, simple models are often used as pieces of richer, more complex models
  - If you want to know how the world works when technology is dynamic and endogenous,
    You need to first understand how it works when technology is fixed,
    So you understand the incentives to develop a new technology in the first place
  - For example, there's a recent AER paper trying to understand the role of reallocation – shifting production from less-productive to more-productive firms – in growth[3]
    The model in the paper is complicated – households face intertemporal tradeoffs between consuming now or saving, firms decide how hard to work on research and development, prices and interest rate are endogenous, there are a lot of interconnected moving parts
    But it's built up on top of simple models of household preferences and firm production
  - Or if you want to model international trade,
    maybe it's enough to imagine technology is static in each country in the short term,
    but different across countries,
    which gives a reason to trade
  - Or if you want to understand whether people save enough for retirement,
    maybe you don't need to model where their tastes for coffee come from,
    and it's enough to have a simple model of consumer behavior within your bigger model of labor supply and investment decisions

[3] D. Acemoglu, U. Akcigit, H. Alp, N. Bloom, and W. Kerr (2018), "Innovation, Reallocation, and Growth," *American Economic Review* 108.11

- And finally, if you want to be able to communicate with economists – to talk with them, to read their research, etc. – you need to know where they're coming from,

  and these simple models are in the background of how most economists think about the world most of the time

- We know they leave things out –

  we know they're not an accurate portrayal of every situation –

  but we always have them in mind as the starting point

- So, even though the real world is incredibly rich and complex and dynamic and intractable,

  we'll be assuming away a lot of the richness and complexity,

  in order to focus on some very simplified, tractable models

  and understand, in the simplest possible cases, the basic implications of a couple of very basic premises:

- **Firms have fixed technological capabilities,**

  and they maximize profits given the prices they see

- **and people have preferences,**

  and they optimize based on their budgets

- (Einstein: "All our science, measured against reality, is primitive and childlike – and yet it is the most precious thing we have.")

- So with all that out of the way,

  let's get back to our model of production,

  and start thinking more about it

4 Production

4.1 Assumptions one might consider making about Y

- So far, we've defined the production set Y as a subset of $\mathbb{R}^k$
  which simply catalogs which production plans are feasible for a firm;
  we haven't yet said anything about what assumptions we might want to make about Y

- We'll always assume Y is **nonempty** and **closed**,
  otherwise the firm's profit maximization problem won't typically have a solution

- Some other properties of Y which we may want to assume:

- Y satisfies **free disposal** if $y \in Y$ and $y' \leq y$ implies $y' \in Y$
- (In $\mathbb{R}^k$, $y' \leq y$ means it's weakly lower in every component –
  that is, $y = (y_1, y_2, \ldots, y_k)$ and $y' = (y'_1, \ldots, y'_k)$, with $y'_i \leq y_i$ for every i)
- This means the firm can freely throw things away –
  it can burn extra inputs, or produce and throw away some outputs
- The production set we drew had free disposal,
  since for each point $y \in Y$, everything below and to the left of it is also in Y
- (For a profit-maximizing firm, all that will matter is the boundary of the production set,
  but it's often convenient to think of Y as thick)

- Y has the **shutdown property** if $0 \in Y$ (the firm is free to do nothing)
- (In a dynamic model, you might assume a firm has a shutdown option in the long run but not the short run)

- Y has (weakly) **increasing returns to scale** if $y \in Y$ implies $\alpha y \in Y$ for every $\alpha > 1$
  (the firm can scale up production – if I can produce y, I can produce $2y$,
  or if I double all my inputs, I at least double all my outputs –
  but I can't necessarily scale it down)
- Y has (weakly) **decreasing returns to scale** if $y \in Y$ implies $\alpha y \in Y$ for every $\alpha \in (0, 1)$
  (the firm can scale down production, but can't necessarily scale it up)
- Y has **constant returns to scale** if $y \in Y$ implies $\alpha y \in Y$ for every $\alpha > 0$
  (the firm can scale production up or down at any ratio – or, CRS = IRS + DRS)

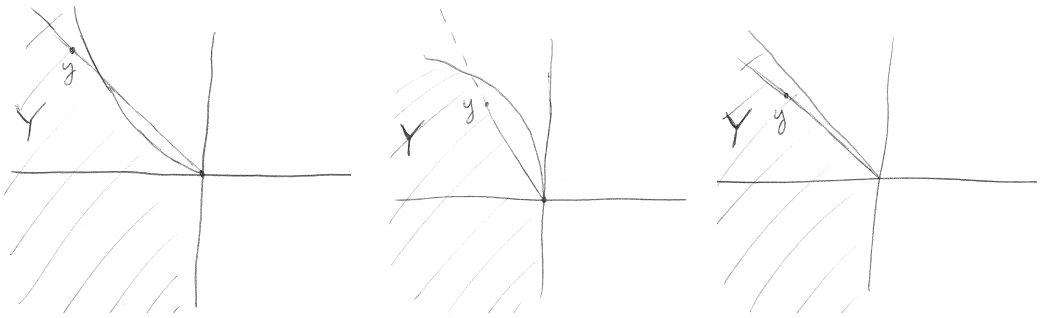

Increasing returns	Decreasing returns	Constant returns
$y \in Y \to \alpha y \in Y$, $y > 1$	$y \in Y \to \alpha y \in Y$, $y < 1$	$y \in Y \to \alpha y \in Y$, $y > 0$

- We can assume Y is **convex**, by the usual definition of a convex set –
  that if y and y' are both in Y, then so is any weighted-average $ty + (1-t)y'$ $(t \in (0, 1))$
  (or if two points are in Y, so is the line connecting them)
- Note that convexity plus the shutdown property together imply decreasing returns to scale

- We can also assume Y is **strictly convex**,
  which means that if $y, y' \in Y$, $y' \neq y$, and $t \in (0, 1)$, then $ty + (1-t)y' \in \text{int}(Y)$

- We also typically assume firms can't make something out of nothing – $Y \cap \mathbb{R}^k_+ = \{0\}$
  (or, if the firm produces anything, it must use some inputs)
- Sometimes people assume production is irreversible – that $Y \cap -Y = \{0\}$
  (or if $y \in Y$ and $y \neq 0$, then $-y \notin Y$)

5 Profit Maximization

5.1 The Firm's Problem

- Now that we've defined firms' technological abilities,

  we can think about the firm's problem, which is simply

$$\max_{y \in Y} p \cdot y, \qquad \text{or} \qquad \max p \cdot y \quad \text{subject to} \quad y \in Y$$

- ($p \cdot y$ is just the dot product, or $\sum_{i=1}^{k} p_i y_i$.

  Since outputs are positive numbers and inputs are negative, assuming prices are positive,

  $(p_i y_i)^+$ is the revenue from selling output y_i,

  and $(p_i y_i)^-$ is the cost of buying input y_i,

  so $p \cdot y$ gives revenues minus costs)

- We'll define the profit function

$$\pi(p) \quad = \quad \sup_{y \in Y} p \cdot y$$

  and the optimal supply correspondence

$$y(p) \quad = \quad Y^*(p) \quad = \quad \arg\max_{y \in Y} p \cdot y \quad = \quad \{y \in Y : p \cdot y = \pi(p)\}$$

- (Note that I said supremum, not maximum, for π,

  because we don't yet know that the firm's problem has a solution;

  it's at least possible that it doesn't, and therefore that $Y^*(p)$ is empty)

5.2 A cool result: the Law of Supply

- Let p and p' be two price vectors,

  and let $y \in Y^*(p)$ and $y' \in Y^*(p')$ be optimal production plans at each

- Of course, this means that both y and y' are in the production set Y

- Now, if y is optimal at p, it must be at least as good as y',

  since we know that's another thing the firm could have done; so

$$p \cdot y \quad \geq \quad p \cdot y'$$

- Similarly, if y' is optimal at p', then

$$p' \cdot y' \quad \geq \quad p' \cdot y$$

- If we add these,

$$p \cdot y + p' \cdot y' \quad \geq \quad p \cdot y' + p' \cdot y$$

  or if we rearrange,

$$p' \cdot (y' - y) \quad \geq \quad p \cdot (y' - y)$$

  or

$$(p' - p) \cdot (y' - y) \quad \geq \quad 0$$

- This is called the **Law of Supply**

- Now, what does this mean?
- Well, if we think of a change in price $\Delta p = p' - p$,
  and the resulting change in the firm's choice of production $\Delta y = y' - y$,
  it means that $\Delta p \cdot \Delta y \geq 0$,
  which means the two changes go more or less in the "same direction" in $\mathbb{R}^k$
- If prices shift in one "direction,"
  optimal production always shifts "broadly in the same direction"
- One extreme example: suppose p' and p differ in only one price
- For example, suppose the change from p to p' is just an increase in the first price, p_1,
  with all other prices staying the same,
  so $p' - p = (\delta, 0, 0, \ldots, 0)$ for some $\delta > 0$
- Well then, $(p' - p) \cdot (y' - y) = \delta(y_1' - y_1) \geq 0$ means $y_1' \geq y_1$
- So if the price of good 1 goes up,
  the firm must produce weakly more of it if it's an output,
  or use weakly less of it if it's an input
- This sounds kind of obvious,
  but we'll see in a few weeks,
  it's not always true when we get to consumer theory!