Lecture 4: Cost Minimization (among other things)

1 Where we are

- So far, we've stated the firm's problem under our basic model of production,

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

  and shown some properties of the firm's profit function π and supply correspondence Y^*

- We've shown in general that optimal supply must satisfy the Law of Supply,

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

  for any $y \in Y^*(p)$ and $y' \in Y^*(p')$,

  and that $\pi(\cdot)$ must be convex and homogenous of degree 1

- A big result that you hopefully all watched in the last week was the Envelope Theorem, relating the derivative of an objective function to the derivative of the maximized value:

  if $V(t)$ is defined as $\max_{x \in X} f(x,t)$,

  the Envelope Theorem says that whever V' exists, $V'(t) = \frac{\partial f}{\partial t}(x,t)$ for any $x \in x^*(t)$

- When Y^* is single-valued, the Envelope Theorem gives us Hotelling's Lemma:

  since $\pi(p) = \max_{y \in Y} \sum_j p_j y_j$, the Envelope Theorem says

$$\frac{\partial \pi}{\partial p_i} \quad = \quad \frac{\partial}{\partial p_i}\left(\sum_j p_j y_j\right)\bigg|_{y=y^*(p)} \quad = \quad y_i|_{y=y^*(p)} \quad = \quad y_i^*(p)$$

- For the special case where π is twice differentiable, and $y(p) = Y^*(p)$ is therefore differentiable, we've seen some additional properties having to do with the Jacobian of y

- We also saw that rationalizability is equivalent to $Y^I \subseteq Y^O$
- With finite data, this boils down to the Weak Axiom of Profit Maximization, or $p \cdot y \geq p \cdot y'$ for any two observations (p, y) and (p', y') in the data
- When π is differentiable and we have data from every possible price vector, ratinoalizability is equivalent to π being convex and Hotelling's Lemma holding
- (More specifically, we showed that if π and y are both observed and satisfy adding-up, then the data is rationalizable if and only if π is convex and Hotelling's Lemma holds.

- I believe you saw in section, so I won't spend time on it here,
  that when you only observe π,
  the data is rationalizable if and only if π is convex and homoegenous of degree 1;
  and if you only a differentiable supply function y,
  it's rationalizable if and only if it's homogeneous of degree 0
  and has a Jacobian matrix which is symmetric and positive semidefinite)

- Today, we think a little about actually solving the firm's problem,
  and start heading into a powerful new tool that we'll spend a couple lectures on,
  Monotone Comparative Statics

- But first... any questions?

2 (Skip in lecture – in notes for completeness)

- The result we saw last time:

Proposition. Let P be an open, convex subset of $\mathbb{R}^k_+$. A single-valued supply function $y : P \to \mathbb{R}^k$ and a differentiable profit function $\pi : P \to \mathbb{R}$ that satisfy adding-up are jointly rationalizable if and only if:

1. $y_i(p) = \frac{\partial \pi}{\partial p_i}(p)$ for every i and every $p \in P$ (Hotelling)
2. π is convex

- What if we didn't observe both π and y?
- **Proposition.** A differentiable profit function $\pi : P \to \mathbb{R}$ (P open and convex in $\mathbb{R}^k_+$) is rationalizable if and only if it's homogeneous of degree 1 and convex.
- (To prove this, we'd define $y(p)$ by letting $y_i = \frac{\partial \pi}{\partial p_i}$, so Hotelling holds automatically;

  we already have convexity of π; we can show that if π was homogeneous of degree 1,

  then with y defined this way, the adding-up condition holds,[1]

  and then we can just use the main result and know it's rationalizable.)
- **Proposition.** A differentiable supply function $y : P \to \mathbb{R}^k$ (P open and convex in $\mathbb{R}^k_+$) is rationalizable if and only if it's homogeneous of degree 0 and $D_p y(p)$ is symmetric and positive semidefinite.
- (To prove this, we would let $\pi(p) = p \cdot y(p)$, so adding-up would automatically hold;

  we'd then confirm that Hotelling holds and that π is convex,[2] and invoke the main result.)

- (Something cool here: we know the profit function must be homogeneous of degree 1,

  and the optimal supply correspondence must be homogeneous of degree 0
- If we observe both π and y, though, we don't have to check – these conditions are redundant:

  if π and y satisfy both Hotelling's lemma and the adding-up condition,

  they *have to* be homogeneous of the right degree
- But if we only observe y *or* π, we need to check that it's homogeneous of the right degree)

[1] To prove it, note that π is assumed to be homogeneous of degree 1, or $\pi(\lambda p) = \lambda \pi(p)$; differentiating this with respect to λ gives $p_1 \frac{\partial \pi}{\partial p_1}(p) + \ldots + p_k \frac{\partial \pi}{\partial p_k}(p) = \pi(p)$ which is exactly the adding-up condition since $y_i(p)$ was defined as $\frac{\partial \pi}{\partial p_i}(p)$.

[2] Differentiating $\pi(p) = p \cdot y(p)$ gives $\frac{\partial \pi}{\partial p_i} = y_i(p) + \sum_{j=1}^k p_j \frac{\partial y_i}{\partial p_j}$. Since y is homogeneous of degree 0, $0 = \frac{\partial}{\partial \lambda} y_i(\lambda p) = \sum_{j=1}^k p_j \frac{\partial y_i}{\partial p_j}$, so $\frac{\partial \pi}{\partial p_i} = y_i$ and Hotelling holds. Once we have Hotelling, $D^2_p \pi(p) = D_p y(p)$, which is positive semidefinite, so π is convex.

3 Let's think about solving the firm's problem

- Next, we're going to think a little more about how we might actually go about solving a firm's profit maximization problem

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

3.1 Graphically...

- If we're given the production set graphically, then conceptually, this is pretty simple (draw it in two dimensions)
- But that's not usually how we're going to be given a production set

3.2 Another way to define technology: the transformation function

- Another way to represent the production set is with a Transformation Function
- This is a function

$$T \quad : \quad \mathbb{R}^k \quad \rightarrow \quad \mathbb{R}$$

  such that $T(y) \leq 0$ if $y \in Y$ and $T(y) > 0$ if $y \notin Y$

- That is, T is weakly negative whenever y is feasible, and strictly positive when y is infeasible
- (Informally, I like to think of T being how many years in the future a given technology is. If a technology is a positive number of years away in the future, it's not feasible today; if $T(y)$ is zero or negative, the technology has already arrived.)

- The boundary $T(y) = 0$, coinciding with the boundary of Y, is called the **transformation frontier**
- I already made the point – when proving $[D_p y]p = 0$ – that a profit-maximizing firm always chooses production plans on the boundary of Y, meaning on the transformation frontier
- (If T is differentiable, we can define the **Marginal Rate of Transformation**

$$MRT_{i,j}(y) \quad = \quad \frac{\partial T / \partial y_i}{\partial T / \partial y_j}$$

  as the extra amount of good j obtained for each unit reduction of good i (meaning an increase in its use as an input or a decrease of its production)

  or the slope of the boundary at y)

- (If i and j are both being used as inputs,
  this is the rate at which you can substitute one for the other –
  how many extra units of labor will replace a unit of something else
- If j is an output,
  this is how much more of output j you can get by increasing your use of input i
- If they're both outputs, it's how much more of one you can get with the same inputs, by sacrificing some of the other)

- If we write the production set Y as a transformation function $T : \mathbb{R}^k \rightarrow \mathbb{R}$,
  we can phrase profit maximization as a constrained maximization problem

$$\max\{p \cdot y\} \quad \text{subject to} \quad T(y) \leq 0$$

  and we could potentially solve it via Lagrangians if we wanted to

- But in practice, when we want to solve the firm's problem (or prove sharper results),
  we typically simplify the problem by considering firms with a **single output**
  - In IO, multi-product firms face a more complicated problem because their products may be substitutes for each other –
    if Samsung gets more people to buy one model Galaxy,
    fewer people will buy a different model –
    but in our price-taking model, this doesn't matter
  - If you think of a large firm as having separate teams and maybe even separate factories to produce its different products,
    there's little problem thinking of each product as a separate single-product firm

4 Single-Output Firms

- So now consider a single-output firm
- In our general production-set model, this is a firm where for every $y \in Y$,
  $y_i \leq 0$ for every i but one
- Or we can think of every y in the production set as taking the form

$$y = (q, -z_1, -z_2, \ldots, -z_m)$$

- (Here q is the quantity of output, and $z_i \geq 0$ the amount of input i used)
- Once we're thinking about a single output,
  if we let q denote the quantity of the firm's output, and z a vector of inputs,
  we can summarize the firm's technology with a function

$$f : \mathbb{R}_+^m \rightarrow \mathbb{R}_+$$

  where $f(z)$ is the most output the firm can product with inputs z
  (Note the signs – now inputs are positive, not negative)
- We call f the **production function**
- If our firm also has free disposal, we can link back to the old notation by observing that

$$Y = \{(q, -z) : q \leq f(z)\}$$

- Similar to the Marginal Rate of Transformation,

  we can define the **Marginal Rate of Technical Substitution**

$$MRTS_{k,\ell} \quad = \quad \frac{\partial f/\partial z_\ell}{\partial f/\partial z_k}$$

  as the ratio at which you can substitute input k for good ℓ while holding production constant

- If we draw "isoquants" in input space (the set of input combinations that give the same output level), the MRTS is the slope

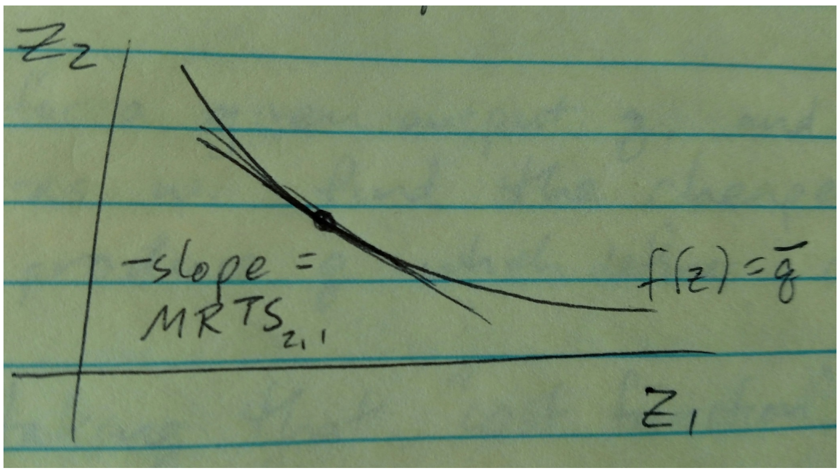

4.1 Solving the Single-Output Firm's Problem

- Now, we let $p \in \mathbb{R}_+$ be the price of the output good,

  and $w \in \mathbb{R}^m_+$ a vector of input prices

- We can write the firm's profit maximization problem as

$$\max_{q,z} \{pq - w \cdot z\} \quad \text{subject to} \quad q \leq f(z)$$

- Of course, if $p > 0$, the firm will maximize profits by setting $q = f(z)$ (not throwing away any of its possible output), so we could write this as

$$\max_z \{pf(z) - w \cdot z\}$$

  and solve directly for z as a function of p and w

- And in fact, you'll do this a bit on the next homework

4.2 Decomposing the Firm's Problem

- However, instead of just solving the firm's problem all at once,

  we can also decompose it into two pieces:

  1. For each output level q, given input prices w, what is the cheapest way to produce q?
  2. Taking that minimum cost to produce q as given,

     what output level q maximizes revenue minus costs?

- That is, rather than directly solving

$$\max_{z \in \mathbb{R}^m_+} \{pf(z) - w \cdot z\}$$

  we can first define

$$c(q) \quad = \quad \min_{z : f(z) \geq q} w \cdot z$$

  and then solve

$$\max_q \{pq - c(q)\}$$

Why is this appealing?

- Well, for one thing, the cost function will accord with lots of earlier economic intuition about marginal costs, average costs, and stuff like that
- For another, the firm's optimal choice of q sometimes stretches the assumptions of our model
  - In order for the firm's problem to have a solution,

    we typically need the production technology to have decreasing returns to scale –

    if not globally, at least locally near the optimum
  - Are we really comfortable with this assumption?
  - If not, there are other ways we can explain firms producing non-infinite amounts,

    but they go outside our model a bit
  - maybe some input level are fixed in the short-term,

    so that even if in the long run there are increasing returns,

    there are decreasing returns in the short-run when the firm can't adjust all inputs;
  - or perhaps the firm isn't really a price-taker in its output good,

    and therefore faces a tradeoff between price and quantity,

    rather than a market price at which it could sell any amount

- These are all interesting questions,

  but they make it hard sometimes to work with the "simple" model and feel good about it

- On the other hand, however the firm chooses q – whether or not it fits our model well – whatever q is and however it's chosen,

  the firm will still want to choose the cost-minimizing way to produce it

- And so whatever properties we find for the cost function,

  those might be more believable than properties of the firm's optimal supply behavior

- And on the other hand,

  if the firm's true cost function is not based on price-taking behavior in input markets,

  as long as the cost function is well-defined,

  we can still think about the firm's choice of output level in the same way

4.3 The Cost Minimization Problem

- So let's formalize the first half of the firm's problem –

  cost minimization for a fixed output level

- Write the problem as

$$\min_{z \geq 0} w \cdot z \quad \text{subject to} \quad f(z) \geq q$$

- Define the **cost function** as its value,

$$c(q, w) \quad = \quad \inf w \cdot z \quad \text{subject to} \quad z \geq 0 \quad \text{and} \quad f(z) \geq q$$

  and the **Conditional Factor Demand Correspondence**

$$Z^*(q, w) \quad = \quad \arg \min_{z : f(z) \geq q} w \cdot z \quad = \quad \{z \ : \ f(z) \geq q \text{ and } w \cdot z = c(q, w)\}$$

  as its solution.

- As we said, once we've solved this, we can consider profit maximization,

$$\max_{q \geq 0} \{pq - c(q, w)\}$$

  as a problem in just one dimension

- Before we do, though, we'll examine the properties of c and Z^*

4.4 Properties of $c(q, w)$ and $Z^*(q, w)$

- Cost Minimization has nearly the same structure as Profit Maximization:

  in both cases, the firm is optimizing a function over some set –

  either the production set Y,

  or the set of input vectors that generate enough output

- And in both cases, it's optimizing a linear function –

  in one case, maximizing $p \cdot y$,

  in the other case, minimizing the cost of the inputs

- So the properties of $c(q, w)$ and $Z^*(q, w)$ will be almost exactly the same as the properties of $\pi(p)$ and $Y^*(p)$,

  and we won't bother to re-do the proofs

- The main difference is that, since now we're minimizing instead of maximizing,

  the cost function will be concave in prices rather than convex,

  and the Jacobian matrix of optimal supply choices will be negative semidefinite rather than positive semidefinite

- **Proposition.**

  1. $c(q, w)$ is homogeneous of degree 1 in w, increasing in q
  2. c is concave in w
  3. Shepard's lemma: if $Z^*(q, w)$ is single-valued (and we write it as $z(q, w)$), then c is differentiable w.r.t. w and
  $$\frac{\partial c}{\partial w_i}(q, w) \quad = \quad z_i(q, w)$$
  4. If $Z^*(q, w)$ is single-valued and differentiable (and we write it as $z(q, w)$), then $D_w z$ is symmetric and negative semidefinite and $[D_w z]w = 0$

- The **Law of Supply** also holds at a fixed level of output:

  if $z \in Z^*(q, w)$ and $z' \in Z^*(q, w')$ then

  $$(z' - z) \cdot (w' - w) \leq 0$$

  so z_i is decreasing in w_i

- One result that doesn't have an analog in what we've done so far:

  Proposition. If f is concave, then c is convex in q (marginal costs are rising).

  (This is trivial with one input, but a result with more than one;

  it's actually problem 1 on the next homework.)

- (On the homework, you'll show that a convex production set implies a concave production function,

  and a concave production function implies a convex cost function.

  Once you know this result, it's the key piece of the Weitzman proof of separating hyperplanes that I posted as "bonus material" on the Canvas site.)

4.5 what about rationalizability? (might skip in lecture)

- Suppose we have lots of observations of a firm's input choices,

  targeting the same level of output q,

  at dififerent input price vectors w

- The result on rationalizability looks just like the case of profit maximization:

- **Proposition.** Let $W \subset \mathbb{R}^m_+$ be an open, convex subset of the input price space.

  For fixed q, a conditional factor demand function $z : W \to \mathbb{R}^m_+$ and differentiable cost function $c : W \to \mathbb{R}_+$ are jointly rationalizable by some production function if and only if

  1. $c(q, w) = w \cdot z(q, w)$ (adding-up)
  2. $\frac{\partial c}{\partial w_i}(q, w) = z_i(q, w)$ (Shepard's Lemma)
  3. c is concave in w

4.6 solving it...

- Once we have the cost function, the firm's profit maximization problem is just

$$\max_{q \geq 0} \{pq - c(q, w)\}$$

  which has first-order condition

$$p = \frac{\partial c}{\partial q}$$

  at an interior solution

- so if production occurs, price must equal marginal cost

- As you'll show on the homework, if the production set is convex (or the production function is concave), then the cost function is convex in q, so the problem

$$pq - c(q)$$

  is concave, and so the first-order condition is sufficient: setting price equal to marginal cost, the firm is maximizing profits

- If c is not convex, though, the FOC isn't sufficient

  So if there are fixed costs, or when there are increasing returns to scale over some production ranges, then it's a little more complicated

- We can rewrite firm profits as

$$pq - c(q,w) \quad = \quad q\left(p - \frac{c(q,w)}{q}\right)$$

  so we're interested in the firm's average cost function $\frac{c}{q}$

- **Proposition.** If the production function has increasing returns to scale, $\frac{c}{q}$ is decreasing in q; if decreasing returns, $\frac{c}{q}$ is increasing.[3]

- A common assumption people like to make is that marginal costs are u-shaped –

  decreasing at first, then increasing

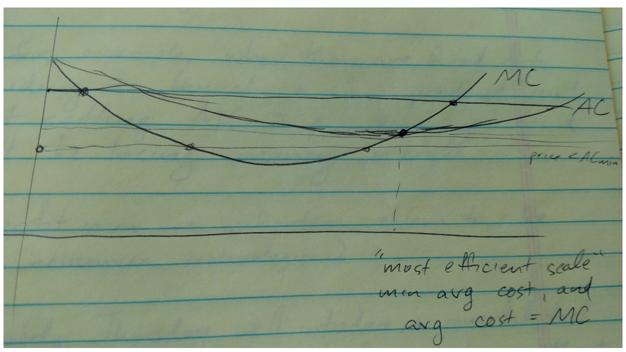

- When that's the case, average costs are also u-shaped,

  with marginal costs coinciding with average costs at the bottom of the average cost curve

- Recall we can write firm profits as $q(p - \frac{c(q)}{q}) = q(p - AC)$;

  if we let AC_{min} denote the minimum value of average costs, then...

  - if $p < AC_{min}$, the firm has no way to make positive profits, so $q^* = 0$
  - If $p > AC_{min}$, price crosses the marginal cost curve twice; q^* is the higher of the two solutions to $c'(q) = p$

[3]**Proof.** If f has IRS, it means $f(\lambda z) \geq \lambda f(z)$ for $\lambda > 1$. So fix input prices w, and let $z^* \in Z(q,w)$ be a cost-minimizing way to produce q; by definition, $c(q,w) = w \cdot z^*$. With increasing returns, $f(\lambda z^*) \geq \lambda f(z^*) = \lambda q$, so if you need to produce λq, λz^* gives you at least enough output; this means $c(\lambda q, w) \leq w \cdot (\lambda z^*) = \lambda c(q,w)$, so $\frac{c(\lambda q)}{\lambda q} \leq \frac{\lambda w \cdot z^*}{\lambda q} = \frac{w \cdot z^*}{q} = \frac{c(q)}{q}$. With decreasing returns, the proof is identical, just with all the statements we just made holding for $\lambda < 1$ rather than $\lambda > 1$.

5 The Monopolist's Problem

- Now, so far we've assumed firms are price takers – both for their inputs, and for their outputs
- We can also think of a firm that makes a unique product, and is able to set the price of that product, but faces a downward-sloping demand curve – a tradeoff between price and the quantity they'll sell
- Assume the firm is still a price taker in inputs, so the cost minimization problem is unchanged, and take $c(q)$ as given

- There are two ways we can formulate the firm's problem: choosing price, or choosing quantity
- First, imagine the firm sets price p, knowing that leads to demand $Q(p)$
- The firm, then, solves

$$\max_p \{pQ(p) - c(Q(p))\}$$

- For the second formulation, the firm chooses quantity, knowing this determines the price it will have to sell at:

  we let $P(q)$ be the **inverse demand function**

- Then the firm solves

$$\max_q \{P(q)q - c(q)\}$$

- This formulation is easier to work with; if both P and c are differentiable, we get the FOC

$$P'(q)q + P(q) \quad = \quad c'(q)$$

- Since P' is negative, this tells us that at the optimum, $P > c'$, or price is greater than marginal costs
- If c is convex, this implies the monopolist sells less than a price-taking firm with the same technology
- If we write the FOC as

$$P'(q)q + \big[P(q) - c'(q)\big] \quad = \quad 0$$

  the first term is the change in *infra-marginal* profit – the reduction in the profits I get from the other customers I'm selling to anyway if I lower my price to serve another customer – and the second term is the profit I make on this new marginal customer

6 Motivating Monotone Comparative Statics

- Consider a firm maximizing profits, given a production set Y
- The Law of Supply tells us that if p_1 goes up, y_1 can't go down,
  but on its own, it doesn't tell us anything about the response of the other y_i

- But there may be other things we'd like to know,
  and that we think might hold pretty generally
- If the price of steel goes up, we know the firm will use less steel;
  should we expect the firm to use more aluminum?
  should we expect the firm to use fewer steel-piercing drill bits?

- When we get to consumer theory,
  we'll give formal definitions of complements and substitutes –
  when the price of peanut butter goes up,
  perhaps a consumer will demand less jelly, and more ham?
- Is there an analog in producer theory –
  can we think of inputs as being complements, and substitutes, in a way we can formalize?

- On one extreme, if we know the firm's production set exactly,
  we can of course look directly at how Y^* changes with p,
  and give precise answers for what will happen
- On the other extreme, if we want results that will hold for *any* Y,
  the Law of Supply tells us something, but it leaves open a pretty wide range of possibilities
- But is there an intermediate case –
  are there conditions we can put on a firm's technology,
  which are fairly general and likely to hold much of the time,
  but could give us more detailed results?

- In the late 1980s, Paul Milgrom and John Roberts looked closely at US manufacturing,[4]

  and noticed that many firms were breaking away from the traditional, highly-specialized production line (which had been around since Henry Ford)

  and had made a number of changes to their production process:

  introducing more flexible tools and programmable, multi-task production equipment;

  producing output in smaller batches;

  having shorter production cycles, smaller inventories, and lower levels of back orders

  (summarized as "a general emphasis on speeding up all aspects of the firm's operation");

  along with broader product lines, a widespread emphasis on higher quality,

  and new organizational strategies and workforce management policies

- Milgrom and Roberts noted that these changes tended to happen together –

  many firms had adopted all or most of these changes,

  while very few firms had adopted a small subset of them –

  so they concluded the changes were likely complementary to each other –

  a firm that had adopted some, would be more likely to find it worthwhile to adopt others

- They argue that adoption was triggered by technological advances which effectively lowered the costs of certain inputs, but not necessarily all of them –

  focusing on reductions in the cost of data collecting and processing, computer-aided design, and flexible manufacturing,

  due to advances in computers, computer networks, and robots and other programmable production equipment

- They also noted that many of these changes were

  "not a matter of small adjustments made independently at each of several margins,

  but rather have involved substantial and closely coordinated changes in a whole range of the firm's activities.

  Even though these changes are implemented over time...

  the full benefits are achieved only by an ultimately radical restructuring."

[4] Paul Milgrom and John Roberts (1990), "The Economics of Modern Manufacturing: Technology, Strategy and Organization," *American Economic Review* 80.3. Note that Theorems 7, 8 and 9 are not correct as published. If you want to go deep into the weeds on this, see Bushnell and Shepard (1995), "The Economics of Modern Manufacturing: Comment"; Donald Topkis (1995), "The Economics of Modern Manufacturing: Comment"; and Milgrom and Roberts (1995), "The Economics of Modern Manufacturing: Reply," all in the *American Economic Review* 85.4. Or just use Milgrom and Roberts (1990) for context, motivation, and to understand the general setup.

- Milgrom and Roberts thus wanted a mathematical framework for considering these changes which allowed for non-convexities in firms' choice sets –

  allowing for discrete jumps in technology, not just continuous incremental changes –

  and they therefore needed mathematical tools which don't depend on differentiability

- The approach they introduced has come to be known as Monotone Comparative Statics –

  Comparative Statics being the question of how the solution to an optimization problem (like a firm's choice of production) changes in response to an exogenous parameter change (like a falling input price),

  and Monotone emphasizing that the approach focuses on predicting the *sign* of a change, but not necessarily its magnitude

- Today, I'll introduce a "baby version" of the Monotone Comparative Statics approach,

  to give some intuition for how it works

- Next week, we'll extend that to the full-on version,

  wrap up producer theory, and move on to consumer theory

7 Monotone Comparative Statics

7.1 General Setup – one dimension

- We'll start with the simple case – a one-dimensional choice variable – and then move on to the more interesting case of a multi-dimensional problem

- Like when we introduced the Envelope Theorem, we'll start out with a generic parameterized optimization problem
$$\max_{x \in X} g(x,t)$$
where x is a choice variable and t is a parameter

- You can think of t as a price,

  and x as a choice made by the firm, such as how much output to produce

- Like with the Envelope Theorem, the choice set X can be continuous or discrete –

  so x could be how much output to produce, or how many factories to open

- We'll let
$$x^*(t) \quad = \quad \arg\max_{x \in X} g(x,t)$$
be the solution to this problem

- Our question is, when can we say that x^* is increasing, or decreasing, in t

7.2 When is a set above another set?

- One complication is that the problems we're considering won't always have a unique solution
- That is, we want to say the solution is increasing, or decreasing, in a parameter...

  but the solution might be a *set* of solutions
- So how do we say that one set of solutions is "greater than" another set?

- We'll define a partial order on sets in the following way
- Let $A, B \subset \mathbb{R}$ be sets
- We'll say A is greater than B in the *strong set order*, or

$$A \geq_{SSO} B$$

  if for any $a \in A$ and $b \in B$, $\max\{a, b\} \in A$ and $\min\{a, b\} \in B$
- So there can be elements of B bigger than elements in A –

  but only if those points are in both sets
- This means that all the points that are in B but not A

  are below all the points that are in both sets

  which are below all the points that are in A and not B
- If A and B have just one element each, then $A \geq_{SSO} B$ just means $a \geq b$ –

  so this really can be thought of as a generalization of "greater than",

  to allow for sets that have nontrivial overlap

- That's how we'll rank sets – and if we have a set defined as a function of a parameter, $X(t)$, we'll say that X is increasing in t if it's increasing via the strong set order – that is, if

$$X(t') \quad \geq_{SSO} \quad X(t)$$

  whenever $t' > t$

7.3 Increasing Differences

- We'll say our objective function g has **increasing differences**

  if the change in g when you move from a lower value of x to a higher value

  is increasing in t

- That is, g has increasing differences if $x' > x$ implies

$$g(x', t) - g(x, t)$$

  is weakly increasing in t; or

$$g(x', t') - g(x, t') \quad \geq \quad g(x', t) - g(x, t)$$

  whenever $x' > x$ and $t' > t$

- If g is differentiable w.r.t. either variable, this is the same as

$$\frac{\partial g}{\partial x} \quad \text{increasing in} \quad t \qquad \text{or} \qquad \frac{\partial g}{\partial t} \quad \text{increasing in} \quad x$$

  and if g is twice differentiable, this is the same as

$$\frac{\partial^2 g}{\partial x \partial t} \quad \geq \quad 0$$

  so this gives us an easy condition to check if we know the objective function

- (Intuitively, this is saying there is complementarity between higher x and higher t –

  an increase in x is worth more when t is bigger,

  or an increase in t is worth more when x is bigger)

- Our main result for today will be that if g has increasing differences,

  then $x^*(t)$ is increasing in t

7.4 Result – Topkis in One Dimension

Theorem ("Baby Topkis")**.** *Let*

$$x^*(t) \quad = \quad \arg\max_{x \in X} g(x,t)$$

where $X \subseteq \mathbb{R}$. If g has increasing differences, then $x^*(t)$ is increasing in t via the Strong Set Order.

- Notice how general this is
- We're not assuming g is differentiable,

  we're not assuming it's concave,

  we're not assuming it has a unique solution;

  we're really not assuming *anything* beyond increasing differences

  - Alternative approach: if x^* is single-valued, and differentiable,

    and g is twice differentiable and strictly concave in x, and some other assumptions hold,

    we can recover the derivative of x^* from the derivatives of g,

    but that's an awful lot of assumptions
  - Here we're not assuming any of that stuff!
- Just increasing differences and that's it!
- In the special case where $x^*(t)$ is single-valued,

  the result is just that $x^*(t)$ is weakly increasing

Proof of Topkis in one dimension

- Let $t' > t$; want to show $x^*(t') \geq_{SSO} x^*(t)$
- So if we pick $x \in x^*(t)$ and $x' \in x^*(t')$,
  we need to show that $\max\{x, x'\} \in x^*(t')$ and $\min\{x, x'\} \in x^*(t)$
- If $x' \geq x$, then this already holds and we're done; so assume $x' < x$
- So we need to show that $x \in x^*(t')$ and $x' \in x^*(t)$

- Since $x \in x^*(t)$ and $x' \in X$, we know that

$$g(x,t) - g(x',t) \quad \geq \quad 0$$

- Increasing differences, with $x > x'$, means $g(x,\cdot) - g(x',\cdot)$ is increasing in the parameter, so

$$g(x,t') - g(x',t') \quad \geq \quad g(x,t) - g(x',t) \quad \geq \quad 0$$

- If $x' \in x^*(t')$, and $g(x,t') \geq g(x',t')$, then x is also in $x^*(t')$
- This means $g(x,t') = g(x',t')$, and therefore

$$g(x,t') - g(x',t') \quad \geq \quad g(x,t) - g(x',t) \quad \geq \quad 0 \qquad \longrightarrow \qquad g(x,t) \quad = \quad g(x',t)$$

- So since $x \in x^*(t)$, $x' \in x^*(t)$ as well, and we're done

7.5 How do we use the result?

- Consider a single-output firm with cost function c
- The firm's problem is

$$\max_{q \geq 0} \{pq - c(q)\}$$

  so $g(q, p) = pq - c(q)$, where p is the parameter and q the choice variable

- Note that $\frac{\partial g}{\partial p} = q$ is increasing in q, so the objective function has increasing differences
- So now we're guaranteed that $q^*(p)$ is increasing in p (via the strong set order)
- If the firm has a unique optimal production level q, then q must be weakly increasing in p
- And we didn't have to make *any* assumptions about c –

  it doesn't have to be convex, or differentiable, or anything!

- (To be fair, we already knew this –

  q^* increasing in output price is a consequence of the Law of Supply –

  but it's still nice to see that this works!)

- That's where we ran out of time –

  Thursday, we'll use the single-dimensional result to build some more intuition,

  then build up to the more interesting multi-dimensional result