Why do economists like markets? (Lectures 2-3)

1 Overview

- Today, we'll review why most economists have an initial presumption that markets work
- That is, while most economists recognize that specific markets may fail and give bad results,
  most of us also believe that *if there is no specific reason to expect a market to fail*,
  it will likely work well
- (As one example: I have a colleague who, any time someone discusses a proposed regulation or governmental intervention, reflexively asks, "where's the market failure?"
  that is, where's the explanation for why an unregulated market would not function well,
  that then begs for a solution?)

- We'll focus on two results that underlie this intuition –
  the First Welfare Theorem and the Coase Theorem –
  and also see the limitations of these results

2 First – what is a good outcome?

- let's start with an easier question: what is a *bad* outcome?
- we could say, an outcome is bad if there's another outcome that *everyone* agrees is better
- suppose you and your friend are both registering for classes this semester,

  and you ended up in here, and she ended up in Game Theory,

  but you'd rather take Game Theory, and she'd rather take this class
- There would be a way to reallocate things such that *both of you* were better off
- it should be pretty obvious that that's a sub-optimal outcome

- we have a name for outcomes that aren't "bad" in this way:

  Pareto efficient
- an outcome is Pareto efficient if there's no way to rearrange things to make some people strictly better off, and leave everyone else exactly as well-off as before
- (in settings where money is involved, we can usually instead make everyone strictly better off, since someone who's strictly better off can give one penny each to everyone who's indifferent between the two outcomes;

  but in settings where there's no way to "transfer utility" from one person to another, we need to allow for the possibility that a Pareto improvement leaves some people as well off as before and some better off)

- so, should we be happy with an outcome that's Pareto efficient?
- well, it depends
- in a sense, yes – it means we've at least captured all the "free" surplus, and any further improvements will require some thoughtful tradeoffs
- in a sense, no – something can be Pareto efficient and still pretty dramatically unfair

  (in a huge worldwide economy, one person owning all the wealth is Pareto efficient, since giving any up to someone else would make the rich guy a little bit worse off – but still isn't a result we're likely to strive for)
- however, it's at least a starting point –

  it's obvious that if a market doesn't lead to Pareto-efficient outcomes, it's leaving money on the table
- and in some sense, unless we're willing to make value judgments about how we weigh one person's gains against another person's, Pareto dominance is often about the only statement we're able to make about when one outcome is better than another
- so for now, we'll use Pareto efficiency as a notion of whether a market is giving a good outcome

- (as we'll see later, whether Pareto efficiency says a lot or a little depends on the details.

  In markets with transferable utility – when money is involved, so we can easily shift resources in continuous amounts from one person to another –

  Pareto efficiency often pins down a unique allocation – that is, up to transfers, there's a unique Pareto efficient outcome,

  which means when we've found it, we know that's really the best we can hope to do.

  In markets without transferable utility, there can be many different Pareto efficient allocations, so knowing we've found one doesn't necessarily tell us that much.

  But at least in settings with money, at least up to worries about the distribution of wealth, Pareto efficiency is a pretty good measure of "good outcomes")

3 First Welfare Theorem, baby version

- Before we get to general equilibrium and the first welfare theorem, let's go back to...

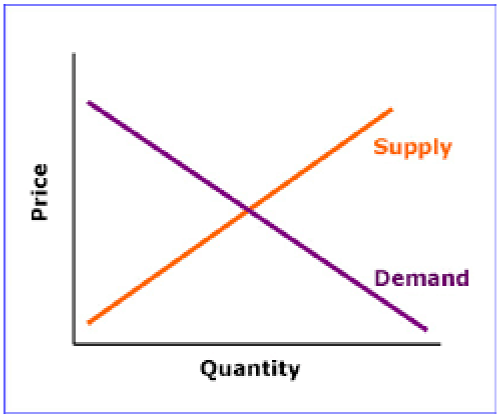

- What does this simple setup say about efficiency?
- That is, if we believe that the equilibrium will be reached – that the market price will be the equilibrium price – will the outcome be Pareto efficient?
- Turns out, yes, but it's easier to see why if we move from the continuous case to a discrete case
- Suppose instead of continuous curves, these are made up of a bunch of sellers with different costs to produce a single unit, and different buyers with different willingness to pay
- Suppose there are ten producers, with costs of 10, 9, 8, 7, 6, 5, 4, 3, 2, and 1

  and ten consumers, with valuations of 10, 9, 8, 7, 6, 5, 4, 3, 2, and 1
- First, what does a Pareto efficient outcome look like?
  - first: if a high-cost seller produces and a lower-cost seller doesn't, that can't be Pareto efficient

    (if the guy with cost 7 is producing and the guy with cost 3 is not, they'd both be better off if the guy with cost 7 bought the good for 5 from the low-cost producer instead of making it himself)
  - similarly: if a low-value buyer gets the good and a higher-value buyer doesn't, it can't be Pareto efficient

    (both would be better off if the high-value buyer bought it from the low-value buyer at a price in between their valuations)
  - so the only candidates for Pareto-efficient outcomes have the k lowest-cost producers selling, and the k highest-value buyers consuming, for some k

- next: how many trades?
- well, if there are fewer than 5 trades, it can't be Pareto-efficient,

  because some buyer with valuation at least 6 isn't buying,

  and some seller with cost at most 5 isn't producing,

  so there's an obvious Pareto improvement
- what about if there are more than 5 trades?

  suppose there were 6 trades, with buyer 10 buying from seller 6, buyer 9 from seller 5, and so on,

  with 10 buying from 6 at price 8, and 5 buying from 1 at price 3
- this means all four of them – buyers 10 and 5, and sellers 6 and 1 – each get surplus of 2
- What if instead, buyer 10 bought from seller 1, and paid $7\frac{3}{4}$ total – with $3\frac{1}{4}$ going to seller 1, $2\frac{1}{4}$ to buyer 5 (who no longer buys), and $2\frac{1}{4}$ to seller 6 (who no longer sells)
- turns out, everyone's better off – so six trades isn't Pareto efficient

• turns out, any Pareto efficient outcome must have five trades – with the five lowest-cost sellers selling, and the five highest-value buyers buying

• (what about prices? not pinned down – those trades, at any prices, would be Pareto-efficient, since changing the prices would make one side better off and the other side worse off)

• what's the market outcome?

• well, the market-clearing price is any price between 5 and 6, and it leads to five trades – exactly the five we need

• so the market clearing price, and the trades that follow, give a Pareto-efficient outcome

4 First Welfare Theorem, real version

- There are two problems with what we've done so far

- One, we've said that *if we get* to equilibrium, that's Pareto-efficient,

  but we haven't said whether we think the market will automatically get to equilibrium

- We'll put aside that question until later

- The other weakness is that we're only thinking about one good

- That is, this is a *partial equilibrium* story –

  we're thinking about the market for steel,

  but assuming that's somehow independent of the market for cars,

  and the market for skyscrapers, and the market for trans-Atlantic shipping

- In fact, markets for different goods are often linked –

  the market for iPhones may impact the market for labor, because Apple might need to hire more people to produce more iPhones,

  and the market for labor may impact the market for iPhones, because if workers earn more money, more of them may want to buy iPhones

- So what we really want to do is think about all markets together

- That's what we'll do next

4.1 overview of model

- A bunch of **goods**, a bunch of **consumers**, and a bunch of **firms**
- Consumers have some initial endowment of goods – you're born with some stuff – and can sell some of what they have, and buy other stuff with the money
- Money is a means of exchange – nobody is born with it, and it has no consumption value, it just facilitates trade
- (If you don't like the idea of consumers being born with endowments of goods – steel, iPhones, etc. – you can think of them having an endowment of *time*, and then can sell some of it to firms, by providing labor)
- Firms are technologies for turning some goods into some other goods – they use some of the goods as inputs, and generate other goods as outputs
- We'll think of both firms and consumers as being **price takers**:

  each good has a "market price" – nobody knows where prices come from,

  but everyone's free to trade as much as they want at the market price
- So again, consumers can sell some of their initial endowment and use that money to buy other stuff
- The key assumptions in the model are:
  - *price-taking behavior* – individuals, and firms, assume that market prices are fixed and outside their own control, they can just choose to trade (or not trade) at the prices they see
  - *commodity goods* – the goods are standardized, so you don't care who you buy from, and there's no uncertainty about the quality of the goods – you know exactly what you're getting and how much you like it
  - *no externalities or network effects* – you only care what you get to consume, so nobody else's choices impact you at all
- We'll be thinking about *general equilibrium*, or *competitive equilibrium*, or *Walrasian equilibrium*, which is when market prices equate supply and demand for each good
  - Given market prices, individuals demand the best consumption bundle they can afford
  - Given market prices, firms choose production to maximize profit
  - And when they do, markets clear
- And the result, which is known as the First Welfare Theorem:

  any such equilibrium is a Pareto-efficient allocation.

4.2 more formally: the enivronment

- There are M different goods
  - We'll let p_m denote the price per unit of good m, and $p \in \mathbb{R}_+^M$ the vector of all prices
- There are N consumers
  - Each consumer $i \in \{1, 2, \ldots, N\}$ has some endowment $e_i \in \mathbb{R}_+^M$ of goods,
    and a utility function $u_i : \mathbb{R}_+^M \to \mathbb{R}$ over how much he/she consumes of each good
  - We'll let $x_i \in \mathbb{R}_+^M$ denote how much i consumes of each good, giving utility $u_i(x_i)$
  - (If it feels weird to think of consumers being born with some amount of steel, and iPhones, and all the other consumption goods,
    you can think of one of the goods being "time,"
    and imagine consumers are endowed only with time, and can sell some of that time as labor to get money to buy the other stuff they want to consume)
- There are F firms
  - Each firm $j \in \{1, 2, \ldots, F\}$ is represented by a *production set* $Y_j \subset \mathbb{R}^M$, which consists of all feasible production plans for that firm
  - Negative numbers represent use of inputs, and positive numbers represent outputs
  - For example, if $M = 5$ and the vector $y_j = (-2, -2, 1, 0, 0)$ is in Y_j,
    it means firm j can turn two units each of goods 1 and 2 into one unit of good 3
  - (Technology need not scale – if $y \in Y_j$, it does not imply that $2y \in Y_j$)
  - We'll let $y_j \in Y_j$ denote the production plan that firm j chooses
  - Note that a production plan y gives profits of $p \cdot y$,
    since the firm collects $p_m y_m$ in revenue for each good it produces,
    and pays costs of $p_m(-y_m)$ for each input it has to buy
- Finally, if the firms make money, we need to know what they do with it, so we assume the firms are owned by the consumers
  - Consumer i owns a share θ_{ij} of firm j, with $\sum_{i=1}^N \theta_{ij} = 1$ for each j
  - If firm j produces y_j at prices p, it earns profits $\pi_j = p \cdot y_j$, and pays $\theta_{ij}\pi_j$ to consumer i

- Finally, we make one technical assumption on consumer preferences:

  that each consumer's preferences are **locally non-satiated**.

- This means for any i, any x_i, and any $\epsilon > 0$,

  there's some x_i' such that $\|x_i' - x_i\| \leq \epsilon$ and $u_i(x_i') > u_i(x_i)$

- This just says consumer preferences have no local maxima –

  there's always something you'd prefer to consume a little more (or a little less) of

- (We could instead assume something stronger – there's at least one good in which utility is strictly increasing, or something like that – but this will be good enough for us)

4.3 Where are we headed?

- so that's our environment
- what are we trying to do?
- we want to show any general equilibrium is Pareto efficient
- but first, we need to define general equilibrium, and Pareto efficient

4.4 Pareto efficiency

- Let $x = (x_1, x_2, \ldots, x_N) \in \mathbb{R}^{MN}$ be a consumption plan for each consumer

- x is *feasible* if it's technologically possible to produce that much of each good:

  that is, if there is some production plan $(y_1, y_2, \ldots, y_F) \in Y_1 \times Y_2 \times \ldots \times Y_F$ such that

$$\sum_{i=1}^{N} x_i \quad \leq \quad \sum_{i=1}^{N} e_i + \sum_{j=1}^{F} y_j$$

- A feasible consumption plan x is *Pareto efficient* if it's not Pareto-dominated by any other feasible consumption plan:

  that is, if there does not exist a feasible plan $x' = (x_1', x_2', \ldots, x_M')$ with

$$u_i(x_i') \quad \geq \quad u_i(x_i)$$

  for every i, with strict inequality holding for at least one i.

- so a consumption plan is Pareto efficient if – given the productive capabilities of the economy – there's no way to make some consumer strictly better off, without making another consumer worse off

4.5 Walrasian Equilibrium

Walrasian equilibrium is defined, basically, as two things happening:

1. Consumers and firms behave optimally given the prices they see

2. Market prices are such that the demand for each good equals supply, i.e., all markets clear

Formally, a Walrasian equilibrium is a vector of prices p^*, a consumption plan for each consumer $x^* = (x_1^*, \ldots, x_N^*)$, and a production plan for each firm $y^* = (y_1^*, \ldots, y_F^*)$ such that:

- Given prices p^*, each firm is maximizing profits: for each $j \in \{1, 2, \ldots, F\}$,

$$y_j^* \quad \in \quad \arg\max_{y_j \in Y_j} p^* \cdot y_j$$

- Given prices p^* and their wealth (from both their endowment of goods and their income from the firms they own), each consumer is maximizing utility: for each i,

$$x_i^* \quad \in \quad \arg\max_{x_i} \left\{ u_i(x_i) \quad : \quad p^* \cdot x_i \quad \leq \quad p^* \cdot e_i + \sum_j \theta_{ij}(p^* \cdot y_j^*) \right\}$$

- Each market clears: element by element,

$$\sum_i x_i^* \quad = \quad \sum_i e_i + \sum_j y_j^*$$

4.6 The Big Result

Theorem (The First Fundamental Theorem of Welfare Economics)**.** *If* (p^*, x^*, y^*) *is a Walrasian equilibrium, then* x^* *is Pareto efficient.*

4.7 OK, let's prove it

- First, let's prove a smaller result:

  Claim 1. *Fix prices p, suppose u_i is locally non-satiated, and let x_i^* be a solution to*

  $$\max u_i(x_i) \quad s.t. \quad p \cdot x_i \leq w_i$$

  1. *if $u_i(x_i') > u_i(x_i^*)$, then $p \cdot x_i' > p \cdot x_i^*$*
  2. *if $u_i(x_i') \geq u_i(x_i^*)$, then $p \cdot x_i' \geq p \cdot x_i^*$*

- The first part is a simple proof by contradiction:

  if $p \cdot x_i' \leq p \cdot x_i^*$, then $p \cdot x_i' \leq w$, and so x_i' is one of the bundles consumer i could afford;

  and if $u_i(x_i') > u_i(x_i^*)$, this means the consumer would have made a mistake in choosing x_i

- The second part is more subtle, and is where we use our assumption of local non-satiation

- Suppose that $u_i(x_i') \geq u_i(x_i^*)$, but $p \cdot x_i' < p \cdot x_i^*$

- Let $\delta = p \cdot x_i^* - p \cdot x_i' > 0$, and note that $p \cdot x_i' = p \cdot x_i^* - \delta \leq w - \delta$ –

  so δ is how much money the consumer would have left over if he had chosen x_i' instead of x_i^*

- Now, we will use local non-satiation to show there's some other consumption plan x_i'',

  which costs within δ of x_i', but is strictly better

- Let $\epsilon = \delta / \|p\|$, and recall that LNS preferences means there's a plan x_i'' such that $\|x_i'' - x_i'\| \leq \epsilon$ and $u_i(x_i'') > u_i(x_i')$

- So we know $u_i(x_i'') > u_i(x_i') \geq u_i(x_i)$, so the consumer would have preferred x_i'' to x_i if he could afford it; what's left is to show he could afford it

- Now,

  $$p \cdot x_i'' \quad = \quad p \cdot x_i' + p \cdot (x_i'' - x_i') \quad \leq \quad w - \delta + p \cdot (x_i'' - x_i')$$

- Now if you go back to trigonometry, for any two vectors a and b, $a \cdot b = \|a\|\|b\| \cos \theta$,

  where θ is the angle between them, and therefore $a \cdot b \leq \|a\|\|b\|$; so

  $$p \cdot (x_i'' - x_i') \quad \leq \quad \|p\|\|x_i'' - x_i'\| \quad \leq \quad \|p\|\epsilon \quad = \quad \|p\|\frac{\delta}{\|p\|} \quad = \quad \delta$$

  and so $p \cdot x_i'' \leq w - \delta + \delta = w$

- So x_i'' is strictly preferred to x_i, but still costs no more than w;

  so if this were all true, the consumer would have had to have made a mistake in choosing x_i