Walrasian Markets (continued)

1 Where we are

- Tuesday, we discussed commodity goods,

  and the difference between markets and marketplaces,

  and gave a very brief introduction to game theory

  including the Prisoner's Dilemma

- Then we jumped into U.S. financial exchanges like the New York Stock Exchange as an example of a market that might look pretty Walrasian –

  a commodity good, price-taking behavior by buyers and sellers, and so on –

  but found an unexpected twist:

  arbitrage opportunities across different exchanges creating an "arms race" to be a little bit faster than other traders,

  leading to socially wasteful spending on very small speed upgrades

- Today, we'll loop back for a few more thoughts on the high-frequency trading "arms race",

  think a little bit more about the "frequent batch auctions" idea proposed by Budish Cramton and Shim to replace continuous trading,

  and then move on to another setting where batch auctions like this are actually used:

  electricity markets

- Any questions before we start?

2 Back to frequent batch auctions

- we discussed trading stocks and other financial assets via **continuous limit order books**
- orders consist of a direction (buy or sell), a number of shares to trade, and a price
- orders are processed as they come in –

  if there's an order in the book to trade at a compatible price, trade happens immediately,

  and if not, the order goes into the book and stays there until it's executed or cancelled
- the point of the paper by Budish, Cramton and Shim is that continuous-time trading creates a race to be a little bit quicker (or at least as quick) as your competitors,

  to take advantage of arbitrage opportunities when the price changes in another market,

  or to avoid being taken advantage of by other traders

- I mentioned that Budish, Cramton and Shim propose an alternative market structure to fix this problem: frequent batch auctions
- rather than trading in continuous time, they propose the exchange store up all the orders that arrive over, say, a tenth of a second,

  and then clear all of those orders together at a single market-clearing price,

  rather than one by one as they come in
- I basically ran of time at the end of lecture, so I didn't get to show how that would work –

  that is, how you would actually process the collection of orders at the end of each "period"
- The set of non-cancelled buy orders gives a demand curve,

  the set of non-cancelled sell orders gives a supply curve, and then there are four possibilities:

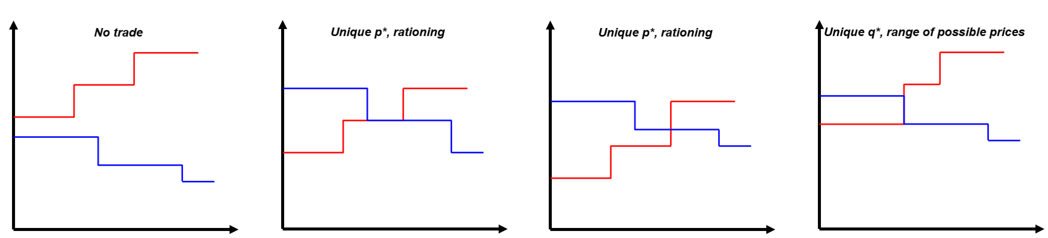

- Any orders that don't trade, stay in the book to start the next "period"

- Like I said, Budish Cramton and Shim argue that this will work better for two reasons

- First, when the price in Chicago changes, fast traders will have time to cancel their standing New York orders that they no longer want to honor,

  so they're exposed to less risk and can therefore demand a lower spread when providing liquidity

- And second, even slow traders are protected, because all the fast traders will compete to trade with them – and that competition will force them to bid more aggressively

- For example, if I'm a slow trader with a standing limit order to buy in New York at 160, and the price in Chicago drops to 150,

  then under continuous-time trading,

  every fast trader would submit an order to sell in New York at 160,

  and whichever order arrived first would sell to me at 160

- But under batch processing, the market clears by price, not by who got there first,

  so if all the fast traders want to sell to me at 160, one of them will instead offer an ask of 159, in order to be chosen over the other fast traders;

  and one of them would then offer an ask of 158, to get chosen instead;

  and so on, until all the fast traders are submitting asks of 150.1, and I hardly get burned at all by the price change

- so, the paper proposes a switch, from continuous-time trading, to very frequent batch auctions

- so far, the exchanges haven't done it

- the paper by Melton[1] talks about one exchange that has experimented with adding a delay to new orders – but not to messages cancelling existing orders – as a potential fix

[1] Hayden Melton (2017), "Market Mechanism Refinement on a Continuous Limit Order Book Venue: A Case Study," *ACM SIGecom Exchanges* 16.1. See also: Oliver Linton and Soheil Mahmoodzadeh (2018), "Implications of High-Frequency Trading for Security Markets," *Annual Review of Economics* 10, for more background info.

2.1 Microwaves and weather

- one other fun related note
- there's another paper – Andriy Shkilko and Konstantin Sokolov (2017), "Every Cloud has a Silver Lining: Fast Trading, Microwave Connectivity and Trading Costs" (working paper)
- they want to compare how a market behaves when there are or aren't some traders with a speed advantage
- (they point out it's not obvious which will be better from a liquidity point of view –

  whether a small number of "fast" traders provide useful liquidity,

  or sabotage the market, relative to having lots of equally-fast traders)
- they use weather as a natural experiment
- in 2011-2012, the microwave technology sending over-air messages faster than fiber-optic was new and very expensive, hence only available to a few traders,

  while the older fiber-optic lines were more widely used
- but the microwave technology that sends over-air messages faster than fiber-optic was disrupted when there was heavy rain between Chicago and New York
- so they could compare market activity when there were a few traders who were faster than anyone else,

  to market activity when there were not a few traders with this speed advantage
- they do indeed find a greater lag between changes in one market and trades in the other market on days with "heavy precipitation";

  and they find that on rainy days, when there aren't a few traders with a speed advantage,

  the market is more liquid and less volatile

3 Bidding strategy

- Now, Budish Cramton and Shim make frequent batch auctions sound pretty tempting
- What we haven't yet thought about:
  how do we expect traders to behave in auctions like these?
- We're not going to answer that right now;
  but armed with the little bit of game theory we saw Tuesday,
  we are going to think about a simpler auction model,
  in order to get some initial useful intuition

- There are lots of different ways we could model an auction setting –
  different ways to think about what information bidders have,
  and what payoffs they face
- We'll spend more time on this later in the semester;
  for now, I'm going to introduce one particular way to think about this
- For now, we're going to get away from thinking about strategic behavior by both sides of the market – buyers and sellers – and think of a setting with fixed supply
- So there are k identical objects for sale, and there are a bunch of potential bidders, who can submit bids (potentially multiple bids if they want)
- If we take all the bids, and use them to construct a demand curve, it will look like this:

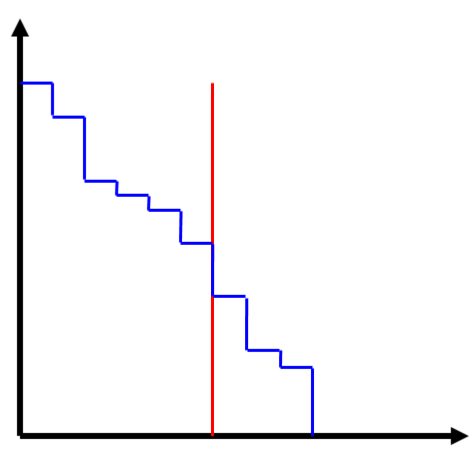

- Note that there is a range of prices that could clear the market:

  any price that's below the k^{th} highest bid, but above the $k+1^{st}$-highest, would work

- It turns out, it's very convenient to assume that the price used is the highest rejected bid,

  i.e., the lowest possible market-clearing price – in this case, the $k+1^{st}$-highest bid

- So that's what we'll suppose: each buyer submits one or more bids, and the "market clears" at the highest rejected bid –

  meaning, the k highest bids "win" the k objects, and the price paid for each one is the $k+1^{st}$ highest bid

- Now suppose that each bidder knows exactly how much getting one of these things is worth to her

- And let's first consider a bidder who only wants one unit:

  so our bidder attaches a value of v to buying a single unit, minus whatever price she pays

- That is, our bidder will get a payoff of 0 if she's not allocated a unit,

  and a payoff of $v-p$ if she gets to buy a unit at price p

- Obviously, our bidder will only bid for one unit;

  the question is, how should she bid?

- Now, without worrying too much about the technical details,

  the essence of game theory is,

  which of my possible moves is best for me, *taking as given* what everyone else is doing?

- That is, assume I can't influence how other people behave;

  all I can do is find the strategy that's best for me, given what they're doing

- And it turns out, we can show that in this case, the bidder should bid exactly her value v.

- Why?
- Well, let's let y be the k^{th} highest bid among our bidder's opponents
- (Our bidder doesn't know the value of y – she sees it as a random variable – but it will turn out to have some value)
- If she bids more than y, she'll be among the k highest bids, and win an object; and the $k+1^{st}$ highest bid, which sets the price, will be y
- So if she bids more than y, she'll win an object and pay y
- On the other hand, if she bids less than y, then there are at least k bids higher than hers, and she doesn't win
- So... if she knew what y was, she would have a choice of getting a payoff of...
  - $v - y$ by bidding more than y, or
  - 0 by bidding less than y
- So when $v > y$, she's happy with any bid above y;

  and when $v < y$, she's happy with any bid below y
- You know's above y when $v > y$, and below y when $v < y$?

  v is!
- So if she just submits a bid of v, she gets what she wants:

  she wins an object whenever she's happy to pay the price y that she has to pay for it,

  and she doesn't win an object when the price would be too high for her

- Another way to think about it:

  since she doesn't know what y is, suppose she thinks of it as having some probability distribution F with density f
- What's her expected payoff if she bids b?
- Well, her payoff $v - y$ when $y < b$, and 0 when $y > b$, so its expected value is

$$u(v, b) \quad = \quad \int_0^b (v - y) f(y) dy + \int_b^\infty 0 f(y) dy \quad = \quad \int_0^b (v - y) f(y) dy$$

- And if we differentiate with respect to b, this is

$$\frac{\partial u}{\partial b} \quad = \quad (v - b) f(b)$$

- So basically, if $b < v$, I can increase my expected payoff by raising my bid;
  and if $b > v$, I can increase my expected payoff by lowering my bid;
  and obviously, I'm perfectly happy when $b = v$

- So, in an auction with fixed supply,
  where the price is equal to the highest rejected bid,
  if I know exactly what the prize is worth to me,
  my best response is to simply bid my value

- Cool!

- (What's also cool is that if everyone does this,
  then the outcome is efficient!
  The objects go the bidders with the highest bids,
  and if everyone is bidding their valuation,
  that means the bidders with the highest valuations win the objects,
  which is Pareto efficient.)

3.1 bidding for multiple units

- But what if our bidder wants more than one unit?
- For example, what if she values the first unit she receives at v_1,
  and the second unit at $v_2 \leq v_1$?
  - (If $v_2 > v_1$, things get very complicated, so let's ignore that case for now.)
- Should she submit two bids, $b_1 = v_1$ and $b_2 = v_2$?
- It turns out, no!
- Let's think about the intuition first, and then show this with math first

- When you only want one unit, your bid can determine whether you trade or not,
  but the only time it affects the *price* is when you lose,
  and therefore don't care what the price is anyway;
  so your bid never simultaneously wins and determines the price
- For this reason, there's no reason to distort your bid away from your true valuation
- When you're bidding for two units, your second bid does two different things:
  it determines whether you get a second unit (if you were already going to get a first),
  which makes you want to bid close to your value;
  but if your second bid loses, it could also determine the market-clearing price,
  which is the price you pay for your first unit!
- If you think of your two bids as different "people,"
  then bidding aggressively for a second unit imposes a negative externality on the profits you get from your first unit!
- This gives you an incentive to "shade" your second bid below your second valuation

- Now let's see the math
- Let y still be the k^{th} highest opposing bid, and now let z represent the $k-1^{st}$-highest, and let F_y and F_z be their distributions, with f_y and f_z their densities
- Now by definition, $z \geq y$; and we'll assume our bidder submits two bids, with $b_1 \geq b_2$
- So, let's think about the different scenarios for what could happen:

  1. If $b_1 < y$, then both of her bids are lower than the k^{th} highest opposing bid, and she wins nothing
  2. If $b_1 > y > b_2$, then her higher bid wins, and y becomes the highest rejected bid, meaning she wins one object and pays y for it
  3. If $z > b_2 > y$, then her higher bid is among the k highest, but her lower bid is the $k+1^{st}$, which means she wins one object and pays b_2 for it
  4. Finally, if $b_2 > z$, both of her bids are above the $k-1^{st}$-highest opposing bid – meaning both her bids win, making z the highest rejected bid, so she wins two objects and pays z for each

- If we write out her expected payoff as integrals like before, we get

$$u(v_1, v_2, b_1, b_2) = 0 + \int_{b_2}^{b_1} (v_1 - y) f(y) dy + (F(b_2) - G(b_2))(v_1 - b_2) + \int_0^{b_2} (v_1 + v_2 - 2z) g(z) dz$$

- And if we take derivatives, we find...

$$\frac{\partial u}{\partial b_1} = (v_1 - b_1) f(b_1)$$

  so just like before, $b_1 = v_1$ is optimal

- And as for b_2,

$$\begin{aligned} \frac{\partial u}{\partial b_2} &= -(v_1 - b_2) f(b_2) + (f(b_2) - g(b_2))(v_1 - b_2) - (F(b_2) - G(b_2)) + (v_1 + v_2 - 2b_2) g(b_2) \\ &= -(F(b_2) - G(b_2)) + (v_2 - b_2) g(b_2) \end{aligned}$$

- Relative to what we found before, there's an extra term
- Now it's not optimal to set $b_2 = v_2$, because b_2 does two things...
  it determines whether you win a second unit, but also might set the price for your first unit!
- (If we plug in $b_2 = v_2$, then $\partial u / \partial b_2$ is negative – we want to reduce b_2 below v_2

- For your first bid, there's no reason to shade,

  because if your highest bid loses, it doesn't matter if it sets the price since you don't win any objects anyway

- But on your second bid, it could set the price you pay for your first,

  giving you a reason to shade your bid below your valuation

- If you're bidding for a third unit, the incentive to "shade" is even stronger,

  because now if your third bid sets the price, it's setting the price for *two* units you're already winning

- So, the general insight?

- In a uniform price auction with fixed supply,

  there's an incentive to bid your actual valuation for the first unit;

  but if you're bidding for additional units,

  there's an incentive to bid below your actual valuation for each additional one

- (The "size" of this incentive depends on your likelihood of setting the price;

  so in a larger, more competitive market with lots of opponents, the incentive is small,

  but in a smaller market with only a few bidders, it's larger)

3.2 What about auctions to sell?

- Everything we just did was focusing on bidders competing to buy something in fixed supply
- We could also flip it all around, and think about potential sellers competing to sell something in fixed demand
- Consider a market with fixed, inelastic demand for k units of a good
- A bunch of potential sellers, with costs c_1 of supplying a single unit,
  cost c_2 of supplying a second unit, and so on
- Payoffs are now $p - c_1$ from supplying one unit at price p,
  $2p - c_1 - c_2$ from supplying two units, and so on
- And we could once again organize an auction with a price set at the best rejected bid –
  this time, the lowest k bids win the right to supply the market,
  and the $k + 1^{st}$-lowest sets the price they all get paid
- And the insight would be exactly the same:
  each firm would have an incentive to bid c_1 for the first unit they'd supply,
  but to shade their bids for subsequent units to be *above* marginal cost
- And that brings us to...

4 Texas electricity auctions

- Electricity is a nice example of a commodity good –

  you want your phone charger to work,

  but you probably don't really care who you're buying electrons from

- We're going to take a quick look at auctions used to balance supply and demand in an electricity market

- Now, the "buyers" in this setting aren't the end consumers,

  they're local utilities (the equivalent of Madison Gas and Electric)

  who have contracts to provide electrical service to households

  (plus some large industrial customers)

- And the sellers are firms that own and operate power plants –

  in this case, mostly natural gas and coal-powered plants,

  although there are also a couple of nuclear, some wind, some oil, some hydroelectric, and some biomass plants

- We'll be looking specifically at electricity markets in Texas,

  because there's a recently-graduated PhD student from here who wrote his dissertation on Texas electricity auctions, so I learned a bit about them from him

- A bit of background[2]
- After some massive blackouts in the 1960s, the U.S. electrical system was organized into eight regional authorities – here's a picture
- Texas basically is its own region, because, Texas
- Within Texas (the Electric Reliability Council of Texas, or ERCOT, region), the state is divided up into Zones (another picture)
- The electrical grid moves electricity around the state;

  when the grid is uncongested, electricity in all of Texas clears at a single price,

  but there is sometimes congestion between zones, in which case the market has to clear separately within each zone
- (The most common problem is congestion between the West and the North,

  because a lot of wind generation happens in the West,

  and when the wind is really blowing, the grid can't always move all that power.
- The arrows show the inter-zone directions where congestion is sometimes an issue,

  although in 90% of the periods looked at in the paper, there was no congestion.

  The paper uses data from 2009.)

- In the Texas market in 2009, there were 228 total power plants, owned by 37 firms

  (This is good – sounds fairly competitive...

  but the four biggest firms supply over half of the state's electricity,

  and the eight biggest supply 76%)
- Natural gas plants accounted for 48% of the market, coal another 33%, and nuclear 12%; another 5% was wind power, and then negligible amounts of other fuels
- The generating firms supplied power to the grid; the power was demanded by large industrial customers and electricity retailers, basically local utilities who were providing it to consumers

[2] This is drawing on Cristian Hernandez (2018), "Resource Misallocation in an Oligopoly Environment: Evidence from Texas' Electricity Market" (working paper); Ali Hortacsu and Steven Puller (2008), "Understanding Strategic Bidding in Multi-Unit Auctions: A Case Study of the Texas Electricity Spot Market," *RAND Journal of Economics* 39.1; and Collin Cain and Jonathan Lesser (2007), "A Common Sense Guide to Wholesale Electric Markets," *Bates White*.

- The market is cleared in multiple stages:

  there are some long-term contracts, but more immediately, there's a day-ahead market and a "real-time" market

- In the day-ahead market, producers commit their power plants,

  deciding which plants will be online when and the hourly output level for each plant

- Demand for electricity (in the short term) is generally perceived to be perfectly inelastic – there's no way to raise the price of electricity in real-time to get a bunch of people to turn off their air conditioners – so the real-time market is used to balance the difference between the day-ahead plan and the actual real-time demand

- For the real-time market, firms submit bids an hour ahead, for each fifteen-minute interval of the subsequent hour:

  for example, at 3 p.m., I would bid for four blocks, 4-4:15 pm, 4:15-4:30, 4:30-4:45, and 4:45-5

  and I could bid as both a buyer and a seller –

  I can ask a price to increase my output level above what I had scheduled ahead of time,

  and I can also offer a price to reduce my output below my commitment

- Firms' marginal costs consist primarily of fuel costs –
  to make electricity come out, I have to burn fuel, and more fuel means more electricity
- There's also some variable operation and maintenance costs to account for –
  running a power plant is more expensive than shutting it down
- And firms get charged for emissions under two federal and one state air pollution programs:
  they pay for their emissions of nitrogen oxides under the Clean Air Interstate Rule,
  for their sulfur dioxide emissions under the Acid Rain Program,
  and for their total emissions under a Texas Commission on Environmental Quality program
- What's nice is that every part of marginal costs is easily measurable,
  so empirical papers can very cleanly calculate firms' marginal costs
  and compare to their bidding behavior
- (In fact, in one paper, the author didn't even have access to producers' initial commitment – what they had said they would generate ahead of time – but could infer it,
  because firms bid above their marginal cost when they're selling additional power
  and below marginal cost when they're "buying" (offering to reduce production),
  so where the supply curve that the firm bids crosses the firm's marginal cost curve,
  that's what its initial position was)

- So, from an earlier paper (Hortacsu and Puller), here's an example of one firm's bids in the "real-time" market for a particular day (in red),

  along with the marginal costs the authors calculated (in orange)

- As expected, the firm is bidding below its costs as a buyer –

  when marginal costs are around $25 per megawatt-hour,

  the firm is offering to pay $12 per megawatt-hour to "buy back" some of its promised production,

  and will buy back even more at $5 per megawatt-hour;

  on the other hand, it's willing to add a bit more production right around marginal costs,

  and then add some additional production if the price is above $35 per megawatt-hour,

  and some more if it's above $45 per megawatt-hour

- So, each producing firm submits a bid schedule like this for each 15-minute interval,

  and the ERCOT system operator aggregates them into a market supply curve,

  which he crosses with the inelastic (vertical) real-time demand

  to determine the market-clearing price and tell each producer how much to actually generate

- Part of the point of the Hernandez paper is that the problem is not static –

  that is, the firm isn't choosing independently how to operate in each fifteen-minute interval,

  because there are ramping constraints –

  you can't just turn on a plant that's shut off and immediately be operating at full capacity

- And his point is that if you ignored these constraints, you would mis-measure marginal costs, and get everything wrong

- Here's a picture of what he concludes is the true industry marginal cost curve at a particular point in time

  (6 p.m. on August 8, 2009)

- And here's how he would have mismeasured it if he had ignored these ramping constraints

- (Basically, a bunch of relatively cheap plants happened to be shut off,

  so a lot of imagined cheap capacity wasn't really there in the short term.)

- The point of his exercise was to measure how inefficient was the allocation of production due to market power –

  that is, because firms aren't bidding their actual marginal costs, it's not always the lowest-cost firms who are asked to generate the power,

  and he wants to measure how big this distortion is

- What he finds is that it's not negligible, but not huge –

  he finds that marginal costs are between 9% and 30% higher than they would be without any strategic behavior by bidders

- (He prefers the estimate near the low end of that range, but there are reasons that's a lower bound)

- As motivation, he was pushing back against other papers saying that output is badly misallocated in many industries, where the most productive plants aren't always the ones being used, and output could be increased over 40% in the U.S. by reallocating production, maybe twice as much in developing economies

- His point was that in this particular market, the inefficiency doesn't seem to be that large when measured carefully

- I brought it up because I think it's a neat market –

  there are a lot of moving parts, a lot of constraints,

  a fairly complex market,

  but a pretty straightforward real-time Walrasian-style auction at the heart of it,

  and it generally seems to function pretty well

5 What's Next

- That's it for Walrasian-style markets

- Next week, we're going to start in on something very different