Recap/Review/Wrap-up

1 Plan

- Tonight: HW5 due
- Online course evaluations are open till tomorrow
- Bunch of office hours before the final:
  - Friday (tomorrow), 10-11 a.m.
  - Monday 10-noon
  - Tuesday 1:30-3:30
  - Wednesday 10-noon
- Final exam next Thursday, 7:45 am, Soc Sci 6102

- Today: recap the semester and tie it all together a bit

- Any questions before we get started?

2 Recap

- For those of you who have read Roth's book, or at least flipped through it,

  you know it's organized around some of the different ways that markets can fail –

  some of which we've seen this semester –

  and solutions he's seen to instances of each type of failure

- For this class, I decided to categorize markets, more or less, by how they clear,

  and we looked at six types of markets:

  1. Walrasian-style markets where trade occurs at a market price,
  2. markets where price is fixed (or zero) and supply is rationed by queuing or other methods,
  3. matching markets where trade is decentralized and requires search,
  4. matching markets cleared by a central matching mechanism,
  5. markets which use auctions,
  6. and markets which use, or are well-approximated as, all-pay auctions or contests

- Rather than just running through everything we did this semester with no clear structure,

  I thought it would be useful to go back over the semester's topics by specifically flagging:

  1. the **markets** that we looked at
  2. the **model or models** we used
  3. the **results** we found
  4. and the mathematical **tools and techniques** we applied

  so I'll divide the board up like this:

model	results	markets	tools

- The rest of today will be filling in that six-by-four grid
- So, onward...

3 Walrasian markets

3.1 Model

- the model we used here was a model of competitive production and consumption
- consumers are characterized by their preferences,
  firms by their production technologies,
  and everyone is a price-taker –
  maximizes utility or profits given market prices, which they can't influence
- implicitly, this was assuming that there are no information problems –
  the goods are commodities, and you don't care who you trade with;
  that firms have no market power;
  and that there are no externalities,
  as your utility depends only on your own consumption

3.2 Results

- the main result was the First Welfare Theorem –
  that given these assumptions, any equilibrium allocation is Pareto-efficient
- we also saw the Coase Theorem, which wasn't based on a particular model –
  the idea that voluntary negotiations will always lead to efficient outcomes,
  based on the assumptions that property rights are well-defined and tradeable and,
  crucially, that there are no transaction costs
- these two results were meant as a baseline –
  to explain why economists typically expect market outcomes to be good,
  while still allowing that some things would upend this result:
  externalities or public goods,
  market power,
  information asymmetries,
  or poorly-defined or incomplete property rights
  (or, in the case of Coase, transaction costs)

3.3 Markets

- We looked first at how U.S. financial markets work,

  with continuous-time limit order books used to clear markets for stocks and other financial instruments

- We saw that the design of these markets created an arbitrage opportunity – free money –

  for whoever could transmit information fastest between exchanges in Chicago and New York

- This created an "arms race" to be fastest –

  each firm having an individual incentive to invest in speed,

  even though everyone investing in speed lowered everyone's profits

- The paper we looked at proposed a change in how the market cleared –

  switching from continuous-time processing as orders come in one by one,

  to discrete-time but very frequent auctions that processed a bunch of orders at once

- They showed this would eliminate the arms race problem,

  while still allowing people to trade in a fraction of a second

- since the proposed fix to the arms race problem was frequent batch auctions,

  we looked next at a market that uses frequent auctions to clear a market

- specifically, we looked at hour-ahead auctions in the market for electricity supply in Texas

- power plants have already committed to supply a certain amount of power each hour,

  and are now bidding to supply more, or withdraw some promised supply,

  to balance the real-time demand an hour later

- we hadn't yet formalized how to think about bidding behavior in auctions,

  but we somewhat foreshadowed our later auction model,

  and considered a firm's incentives bidding in such an auction

- we showed that, if a firm thought of the bids of their competitors as random variables,

  that a firm had an incentive to bid their true marginal cost for the first unit of electricity they were bidding on

- (this is the same as our observation much later that in a second-price private value auction,

  your best move is to bid your valuation)

- but firms faced an incentive to bid above costs on subsequent units of supply,
  or below costs to buy back subsequent units,
  so beyond the first unit, firms were not bidding their true costs
- we saw a couple of papers that found evidence of exactly this type of strategic bidding,
  and estimated its effect on allocative efficiency:
  costs were 10-30% higher than the lowest-cost production possible

3.4 Tools

- For the model, we were just using individual optimization
- To make sense of the particular markets we looked at,
  we also introduced the tools of static game theory and Nash equilibrium
- We did this, in part, to define the Prisoner's Dilemma,
  so we could use that as an analogy for the high-speed trading arms race –
  every trading firm has an individual incentive to invest in greater speed,
  but when they all invest, everyone is worse off than if they all hadn't
- without formally defining a model of bidding in auctions,
  we also introduced the idea of thinking about your opponents' bids as random variables
- we used this to show that in a multi-unit auction (like an electricity auction)
  where the price everyone pays is set by the highest rejected bid,
  firms have an incentive to bid their true valuation for the first unit they bid on,
  but to shade their bids on subsequent units;
  we used this to predict the pattern we saw in the Texas electricity auction data

4 "Non-market" markets

- Next, we looked at some markets that don't clear on the basis of price –
  typically, goods where the price is held artificially low,
  and the resulting excess demand leads to some form of rationing

4.1 Markets

- rationing of car tires and certain foods during WW2, gasonline in the 1970s
- markets where you wait in line – Duke basketball tickets, Soviet bread lines, event tickets like Obama's speech on campus

4.2 Model

- the phenomenon we modeled was waiting in line
- we used a few different versions of a model
- the "richest one" had heterogeneity in both consumers and products,
  with consumer x getting utility
  $$V + \beta x\theta - ct$$
  from attending the show and sitting in seat θ
  if she had to wait in line for time t to get in

4.3 Tools

- here, we introduced some cool tools that we'd use again later on
- we saw that when the payoff function has increasing differences,
  any equilibrium must have assortative matching –
  the more eager fans getting the better seats
- and once we knew the equilibrium allocation,
  we introduced the envelope theorem,
  and showed how we could calculate players' equilibrium payoffs,
  and from that, the equilibrium strategies

4.4 Results

- this let us calculate the equilibrium strategies –
  exactly when each type of bidder would get in line –
  which let us understand how a queuing market will perform

- among the results we found:
  if all consumers are the same, with or without heterogeneous products,
  we get full rent dissipation – equilibrium payoffs are 0

- with heterogeneous consumers, some consumers get positive surplus,
  and it's the same surplus as if the products were priced correctly to clear the market –
  but the "money" is wasted rather than kept by the venue

- we found that decreasing the cost of waiting wouldn't make things any better –
  if waiting is less costly, people just wait longer in equilibrium –
  and that increasing the baseline value of the product V also wouldn't make things better

- we found that a purely random allocation could outperform queuing,
  but doesn't have to – it depends on the details of the environment

- We also thought about when the assumptions of our model would be violated –
  and whether there were conditions under which queuing might make sense

4.5 More Markets

- the queuing model led us to look at concert tickets –

  even when there's no longer physical waiting in line to buy tickets,

  if tickets are priced below market value, there are profits to be made reselling them,

  which leads to rent-seeking behavior by ticket resellers

- we talked about efforts to reduce this,

  and why auctions for event tickets haven't really caught on

- on rationing more generally,

  we looked at taxi rides, which are rationed at peak times, because price and supply are both regulated,

  and contrasted with Uber, which uses surge pricing to try to match supply and demand in realtime,

  and looked at how that worked

- surge pricing led us to think about price gouging after natural disasters,

  and the pros and cons of whether it should be allowed

- and related to rent dissipation,

  we talked about whether allowing riders to tip drivers would actually make Uber drivers better off,

  or whether free entry would lead drivers to "compete away" those gains

5 Matching Markets and Search

- next, we defined matching markets –

  markets with two clear "sides" of the market, where the point of trade is to partner up with someone on the other side

- we started with matching markets where trade was decentralized,

  and so participants had to search for each other

5.1 Model

- we started with a simple two-sided model of men and women,

  where all men had the same preferences over women,

  and all women had the same preferences over men,

  and "your payoff was your partner"

- we assumed there was a fixed cost c of searching in each period,

  and if you searched, you randomly matched to someone on the other side of the market who was also searching

5.2 Results

- the main result from the simple model was that as search costs got lower,

  you got closer and closer to assortative matching –

  high types matching with high types, low types with low types

- when costs are very low, people can afford to be very patient,

  so the 13s will keep searching until they find a 13,

  so the 12s can't match with 13s, but will search till they find a 12,

  and so on

- when costs are very high, people stick with the first partner they find,

  because searching again is too expensive

- and in between, you get an in-between result –

  higher types pool together and match within the higher pool,

  and lower types pool together and match within the lower pool,

  but 12s and 13s sometimes end up together,

  as to 10s and 11s

5.3 Markets

- we looked at an empirical paper on online dating markets
- if we think about online dating versus offline,

  we might imagine that online dating is characterized by lower search costs –

  more opportunities to meet more people cheaply
- and the result we found – based on Korean data –

  was that couples that met online had more assortative matching

  than couples that met in other ways

5.4 back to models and results

- we also saw that if we introduce money –

  if we allow transfers of surplus within a couple –

  the results can change
- with money, what matters is not each individual partner's payoff,

  but the combined payoff, since they can allocate it however they want
- in that case, whether the best match with the best

  depends on whether an increase in one partner's quality is worth more to a high- or a low-quality partner
- that is, whether the combined surplus has increasing or decreasing differences in the two partners' types
- when surplus has increasing differences, then as search costs get low,

  we get positive assortative matching – the better firms hiring the better workers –

  but when surplus has decreasing differences,

  we instead get the better firms hiring the worse workers,

  because firm quality and worker quality are substitutes for each other

- we also defined stability as a static, equilibrium-like concept,

  and saw that what's stable matched the predictions above –

  positive assortative matching when there's no money,

  or when there is money and surplus has increasing differences;

  and negative assortative matching when there's money and surplus has decreasing differences

5.5 Tools

- the big idea we saw here was the ability to calculate value functions recursively –

  writing my expected payoff as a function of my payoff today and my continuation payoff if I keep searching,

  and using the fact that my search problem is stationary –

  it looks the same tomorrow as it did today –

  we could then work out what my expected payoff would have to be

6 Centralized Matching Markets

- from there, we moved onto matching markets with a centralized mechanism, beginning with the National Resident Matching Program

6.1 Markets

- NRMP

6.2 Model

- We modified our model of matching markets to allow for idiosyncratic preferences – each doctor has his or her own strict preferences over hospitals, and each hospital has its own strict ranking of doctors
- We suppose each doctor, and each hospital, submits a rank-order list of possible matches, and we look for matchings that are stable – nobody pairs with a partner when they'd rather be alone, and no hospital-doctor pair who aren't together would both rather be together

6.3 Results

- One big result was simply that a stable matching always exists,
  and that the Deferred Acceptance Algorithm gives a way to find one
- In fact, the DAA always finds the best stable matching,
  from the point of view of everyone on the proposing side;
  but it's also the worst stable matching,
  from the point of view of everyone on the receiving side
- We saw that under the DAA when doctors propose,
  doctors do best by reporting their true preferences,
  but some hospitals might gain from misreporting;
  when hospitals propose, if each hospital only wants a single doctor,
  hospitals do best by reporting truthfully;
  but if hospitals want multiple doctors, as we saw on the midterm,
  they can sometimes gain by misreporting
- we also saw that couples can create a serious problem –
  with couples, there's no guarantee a stable match even exists,
  and the DAA doesn't give a way to find one;
  and we talked about the redesign in the 1990s,
  which was meant to handle couples, and other complications,
  and address the concern that the old match favored hospitals over doctors

6.4 Back to markets...

- After the NRMP, we looked at some other centralized matching problems
- We considered school choice – programs assigning students to public schools in major cities
- a key question was whether schools have actual preferences over the students they enrolled,

  or just accepted students based on priorities they were told they were supposed to respect
- we saw that if schools are assumed to have preferences over students (New York),

  the problem looks just like doctors and hospitals,

  and the Deferred Acceptance Algorithm is a natural way to assign students
- if schools do not have preferences (Boston),

  this creates a tradeoff between efficiency and elimination of justified envy,

  and depending on how seriously we need to respect students' priority orders,

  either the DAA or a different algorithm, based on Top Trading Cycles, might be a better choice
- but we also saw a possible argument in favor of a different mechanism,

  the Boston Mechanism, which market design folks had mostly dismissed early on because it's not strategy-proof
- (ex ante vs ex post efficiency, strategic play expressing cardinal preferences)

- we looked at the problem of letting MBA students into popular courses,

  and the idea of using something like market-clearing prices, where students were given budgets of fake money, to solve the problem
- (here, many of the challenges were in the implementation:

  how to allow students to communicate preferences that includes substitutes and complements,

  and how to solve a really hard computational problem in a reasonable amount of time)

- and finally, we talked about the "market" for donor kidneys –

  how "kidney exchange" programs have allowed kidney patients with a willing but incompatible donor to receive a kidney,

  but how there is still a major shortfall in the U.S., with about 100,000 kidney patients awaiting a transplant from a deceased donor,

  and a typical wait time of over 3 years

- and described the market in the one place that allows for paid donation, Iran,

  and the success this has had in shortening wait times

6.5 Tools

- the main tools we used were the definition of stability,

  and the Deferred Acceptance (or Gale Shapley) algorithm

- in school choice, we also defined an algorithm based on Top Trading Cycles,

  and the Boston Mechanism

- in course match, we saw the idea of finding equilibrium prices for each course,

  based on students reporting their preferences and getting hypothetical budgets of fake money

7 Auctions

- Finally, the last few weeks, we've been talking about auctions

7.1 Model

- we focused on the benchmark model in auction theory, Independent Private Values:

  a fixed set of risk-neutral bidders,

  with valuations which are independent random variables

  drawn from potentially different probability distributions F_i,

  and payoffs of $v_i - p$ from winning (or $-p$ from losing and paying p)

- many of our stronger results assumed symmetry – all the F_i are the same –

  or that the distributions were regular, i.e., $x - \frac{1-F_i(x)}{f_i(x)}$ is increasing

7.2 Tools

- We adapted our game-theory notion of equilibrium to incorporate uncertainty about other bidders' valuations,

  with bidders maximizing their expected payoff,

  in expectation over the other bidders' valuations (and therefore bids)

- We saw the Revelation Principle, which allowed us to think of any possible auction format as simply an allocation rule p and a payment rule x,

  defining allocations and payments for each combination of bidder valuations

- We used the Envelope Theorem to calculate bidder payoffs for a given allocation rule,

  which led us to restating expected revenue as

$$E(Rev) \quad = \quad E_v \sum_i p_i(v) VV_i(v_i) - \sum_i U_i(a_i)$$

  where VV_i is the virtual value of bidder i, $v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$,

  and $U_i(a_i)$ is the surplus earned by bidder i at her lowest possible valuation

7.3 Results

- Writing the seller's problem this way showed that the seller's revenue depends only on the equilibrium allocation and the surplus of low types;

  this established Revenue Equivalence, that in the symmetric setting, all standard auctions give the same expected revenue

- It let us to easily state the revenue-maximizing auction when distributions are regular:

  award the object to the bidder with the highest virtual value, as long as it's positive,

  and use the payment rule that makes that incentive-compatible,

  and doesn't give any surplus to the lowest type of each bidder

- For deterministic mechanisms – such as giving the object to the bidder with the highest virtual value –

  this payment rule takes a nice form:

  the winning bidder's payment is the lowest valuation she could have reported, and still received the object

- This tells us that in the symmetric case,

  the optimal auction is simply an English auction with a reserve price

  (or, by revenue equivalence, any standard auction with a reserve price)

- In the asymmetric case, this is no longer true –

  the bidder with the highest VV may not be the bidder with the highest valuation –

  but at least we know exactly what the optimal auction is

- We also proved the Bulow-Klemperer result: in the symmetric case,

  adding a bidder increases revenue more than setting the reserve price optimally

7.4 Model

- after proving those results, we revisited our model,

  thinking about how its assumptions might be violated in different settings

- we thought about common values, and the winner's curse:

  if the prize is worth the same to whoever wins it, but nobody knows exactly what it's worth,

  then if I win, I find out my guess was the most optimistic,

  which means it's not worth as much as I was thinking

7.5 Markets

- as an example of a common-values setting,

  we thought about auctions for offshore oil and gas drilling rights,

  where the value of winning depends mostly on the volume of hydrocarbons under the sea bed

- this example was interesting because in many of these auctions,

  one bidder – the owner of a neighboring tract – had better information about the value of a new tract for sale,

  and we saw that this had a big impact on bidding behavior, and post-auction profit

- we also looked at a couple of auction settings which went way beyond our basic model

- we looked at internet ad auctions –

  generalized second-price auctions used by Google to sell ads returned with search results

- and we looked at spectrum auctions –

  auctions used by governments for licenses to use particular frequency ranges for mobile communication,

  including some early-2000s 3G auctions in Europe,

  and the recent "incentive auction" in the US,

  which was both buying up underutilized bandwidth from TV stations

  and selling it to mobile providers for 5G service

- and finally, I briefly mentioned a paper on used car auctions,

  finding that the identity of the auctioneer could significantly effect outcomes –

  some auctioneers are better than others at selling cars

- these last couple applications were interesting,

  partly because the markets themselves are interesting,

  but partly because they go beyond easily tractible models –

  there's no theoretical model cleanly establishing the "right" way to do a two-sided auction like the FCC incentive auction,

  and no model that would account for the role of the auctioneer in an English auction,

  showing that while theory can give some powerful results,

  it can't always answer everything

8 Contests

- finally, Tuesday, we talked about all-pay auctions and research contests

8.1 Model

- to model settings like lobbyists trying to influence a politician,
  we used a model of an all-pay auction,
  but with complete information – everyone's valuation is known

8.2 Tools

- we introduced mixed strategies, and mixed-strategy equilibrium,
  to figure out the equilibrium of the complete-info all-pay auction

8.3 Results

- we saw the politician (seller) sometimes increases his revenue by limiting the competitors,
  and is sometimes best off excluding the lobbyist willing to pay the most

8.4 Markets

- aside from corrupt politicians,
  we talked about patent races, and competition for innovation prizes
  like the Ansari X and Google Lunar X Prizes
- we talked about procurement auctions,
  and the multi-phase competitions run by the Department of Defense
  under the Small Business Innovation Research program

8.5 Model

- For innovation contests like this, I introduced a different model,
  where firms choose both how much to spend developing a high-quality product,
  and also what price to offer it at;
  and I mentioned a result that from the procurer's point of view,
  it's often optimal to invite exactly two competitors

8.6 More markets

- We looked at some other settings that look like contests
- grant proposals, which are in some sense socially wasteful,
  and could perhaps be improved by making them more random
- tennis, where giving some prize money to the loser can actually increase effort in some cases
  (when prizes are two-dimensional and the players are different enough in ability levels);
- and some silly auction formats like highest-unique-bid and penny auctions,
  which are nominally set up as auctions but are more like lotteries

9 Three Big Takeaway Lessons

9.1 There's always a market

- if we define markets narrowly – as people buying and selling a good in an unrestricted way – then there are some goods for which there isn't a market
- but even for things that aren't traded freely,
  something still determines who winds up with them
- if we think of markets as defined more broadly –
  as whatever process determines the final allocation of a scarce resource –
  then there's *always* a market
- if the market doesn't clear based on price, it clears based on something else
- if people can't buy and sell freely, this may lead to a black market,
  or allocation based on waiting time, or lottery, or personal connections,
  but we can still think about the performance of a market based on the allocation it leads to

- before Al Roth and colleagues redesigned the New York school assignment system,
  students were somehow ending up in public high schools
- transplantable organ allocation may not work flawlessly,
  but a bunch of people still somehow end up with kidney transplants
- my friend who was TA for the World Food Economy class at Stanford told me the professor liked to say,
  "All markets clear; in this market, starvation is a form of market clearing"

- six days ago, the Atlantic magazine had an article[1] on recent work being done on machine learning software to better place refugees to the U.S. in cities where they are "most likely to be welcomed and find success"
- even in settings we don't typically think of as markets,
  if there's a scarce resource, it has to get allocated somehow;
  and if that allocation matters for efficiency, it's interesting to think about as a market

[1] Krishnadev Calamur, "How Technology Could Revolutionize Refugee Resettlement," *The Atlantic*, Apr 26 2019. https://www.theatlantic.com/international/archive/2019/04/how-technology-could-revolutionize-refugee-resettlement/587383/

And two related posts from Al Roth's blog: https://marketdesigner.blogspot.com/2019/04/refugee-resettlement-in-us-by-hias.html and https://marketdesigner.blogspot.com/2018/10/resettling-refugees-thoughtfully-by.html

9.2 The details matter

- the First Welfare Theorem and Coase Theorem might suggest the details don't matter –

  that market outcomes are guaranteed to be efficient,

  and so the details of how the market is set up are basically irrelevant,

  maybe up to a bit of redistribution

- we've seen that's clearly not the case

- right from the start, we saw that one way to trade stocks – continuous limit order books –

  led to hundreds of millions of dollars of largely wasteful investment in speed,

  while an alternative – frequent batch auctions – would not create that same incentive

- Roth points out that centralized matching markets that use algorithms that select a stable match have tended to catch on and last a long time,

  while algorithms that don't select a stable match tend to be abandoned because people try to circumvent them

- By adjusting for ad "quality," Google found a better auction format for online search ads than Yahoo,

  and Google won and Yahoo lost –

  in 2008, Yahoo started outsourcing its ad placements to Google

- The UK and Germany raised over 600 euros per capita auctioning spectrum rights for 3G,

  while Switzerland raised only 20

- the point is, markets can perform well or perform poorly,

  depending on the specifics of the setting, and on exactly how the market is organized

- the idea that "the market will decide" –

  that "the market" is some neutral, impartial, inevitable problem-solver –

  is appealing, but too simplistic –

- the market could decide different things, depending on how it is designed

9.3 Theory is a useful guide... but the map is not the territory

- Part of the point of this class was to think about how to translate a real setting to a model – and what we could learn by doing that
- In some markets, theory makes very strong predictions, or theoretical results help to explain what we see in the real world
- Once you learn that the NRMP mechanism always selects a stable match, it makes more sense that the medical match caught on even though it was voluntary, and that it's lasted several decades (even if it needed some updating)
- Once you learn about revenue equivalence, it's unsurprising that multiple different auction formats are in widespread use
- And theory can help a lot if you're looking to design a new market – the design of school choice programs, MBA course selection, and spectrum auctions has been heavily influenced by theorists (although some have worked better than others)

- However, theory also has limits
- Theory models are limited by the assumptions they depend on –

  as we saw, real-world settings always have features that go beyond any simple model
- And theory results must be viewed through the lens of the assumptions needed for them to hold
- The Coase Theorem may be true –

  a particular set of assumptions may indeed lead to the efficiency result –

  but it may be irrelevant, if those assumptions rarely hold
- One of the appealing things about the market design approach

  is the willingness to use theory when appropriate, but along with other tools –

  simulation, experimentation – to find a workable design,

  and to not let the perfect be the enemy of the good
- And also to stay true to the real goal –

  if you're really interested in ex ante efficiency,

  don't settle for ex post efficiency because it's an easy problem to solve theoretically;

  embrace that the problem is hard, and can't be solved perfectly,

  and make whatever progress you can toward solving that hard problem

10 To sum up...

- several years ago, when I was interviewing for a job at another school,
  a dean scanned my CV full of papers about auctions,
  and asked why economists love auctions so much
- I said that part of the reason is we see them as a microcosm for markets in general,
  but that unlike with markets in general, it's clear what the rules are

- in this class, we've looked at lots of different markets –
  some, like the NRMP and Google ad auctions, where the rules are quite clear;
  others, like stock trading, queuing, or trying to bribe a congressman,
  where there aren't really any rules,
  and instead we're looking for a model that's a decent descriptive fit for what's going on
- I had three real goals for this semester:

  1. to teach you about a bunch of interesting markets;
  2. to show you how economists like to model some of these markets,
     and what insights we can get from some of those models;
  3. and to start you thinking about how to bridge the gap between market and model –
     to start to learn about choosing a way to model a particular setting,
     and how to think critically about models and their limitations

- we'll all find out in a week how I did
- (based on the early returns – from homework assignments and the midterm –
  I feel pretty good about it)

- I know I made the class hard,
  both technically, but also in not just being a mechanical application of techniques;
  but I feel you've all risen to that,
  and it's been an absolute pleasure teaching you this semester
- thank you all for being here!