Recurring Themes in Auction Theory and Mechanism Design

Part III: Robustness in Auctions

Dan Quint
U of Wisconsin
Visitor, UTokyo and UTMD

February 20, 2023
University of Tokyo

Overview

- Last week:
  - Seller’s problem, and relating revenue to virtual value of winner
  - Pre-auction decisions, and understanding them through externalities
- Today:
  - “Robustness” in auctions
- Tomorrow:
  - Empirics (from a theorist’s point of view)

Let's start with an example

What's the optimal auction when buyer valuations are correlated?

- Two bidders, valuations are…
- Ask the buyers their valuations, and then…
  - Sell to the higher-value bidder for their value
  - Anyone who claims to have $10 valuation, must also bet $30 to win $15 that the other bidder will report $10 as well!
- Efficient, and seller gets all surplus, so clearly optimal

	$v_2 = 100$	$v_2 = 10$
$v_1 = 100$	$1/3$	$1/6$
$v_1 = 10$	$1/6$	$1/3$

	$v_2 = 100$	$v_2 = 10$
$v_1 = 100$	Sell to either for $100	Sell to 1 for $100, also charge 2 $30
$v_1 = 10$	Sell to 2 for $100, also charge 1 $30	Give one $15, give the other $5 and the object

R. Myerson (1981), Optimal Auction Design, *Mathematics of Operations Research* 6(1)
J Crémer and R McLean (1988), Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions, *Econometrica* 56(6)

If this is optimal, why don't we ever see sellers doing this?

- Weird and complicated
- Must be precisely tailored
  - Seller needs to know exact distribution of valuations…
  - …and buyers' beliefs about each others' valuations…
  - …and even buyers' beliefs about each others' beliefs
- Fails if buyers are risk-averse
- Has a "bad" equilibrium too
- Vulnerable to collusion
- Seems… "fragile"?

	$v_2 = 100$	$v_2 = 10$
$v_1 = 100$	$1/3$	$1/6$
$v_1 = 10$	$1/6$	$1/3$

	$v_2 = 100$	$v_2 = 10$
$v_1 = 100$	Sell to either for \$100	Sell to 1 for \$100, also charge 2 \$30
$v_1 = 10$	Sell to 2 for \$100, also charge 1 \$30	Give one \$15, give the other \$5 and the object

R. Myerson (1981), Optimal Auction Design, *Mathematics of Operations Research* 6(1)
J Crémer and R McLean (1988), Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions, *Econometrica* 56(6)

Here's another alternative

- Just post a price of $99
  - Doesn't matter what bidders believe about each others' valuations or beliefs
  - Doesn't matter if bidders are risk-averse
  - No "bad" equilibrium where bidders get object for free
  - Not vulnerable to collusion
  - Expected revenue is $66…
  - …versus $70 from the optimal mechanism
- This seems more… "robust"

	$v_2 = 100$	$v_2 = 10$
$v_1 = 100$	$^1/_3$	$^1/_6$
$v_1 = 10$	$^1/_6$	$^1/_3$

	$v_2 = 100$	$v_2 = 10$
$v_1 = 100$	Sell to either for $100	Sell to 1 for $100, also charge 2 $30
$v_1 = 10$	Sell to 2 for $100, also charge 1 $30	Give one $15, give the other $5 and the object

Robustness in auctions

- Informally, I think of “robust” as…
  - An auction that would still work pretty well if your model was a little wrong?
  - An auction that performs acceptably in a wide range of settings?
  - An auction whose performance doesn’t depend critically on your modeling assumptions being true?
- Still lots of different things this could mean
- Today: notions of robustness

What if you're not sure what the bidders know?

Common knowledge

- To fully describe a model, we need to specify…
  - Distribution of primitives (like valuations)
  - Each buyer's beliefs about those distributions
  - Each buyer's beliefs about each other buyer's beliefs about them
  - …and so on, to infinity

- To do this in a reasonable way, we usually just assume the environment is common knowledge
  - Everyone shares a common prior on all the details of the environment, and knows that, and knows everyone else knows it…
  - Which pins down beliefs, but also beliefs about beliefs and so on

- But that's kind of a strong assumption

“The Wilson Doctrine”

“Game theory has a great advantage in explicitly analyzing the consequences of trading rules that presumably are really common knowledge;

it is deficient to the extent it assumes other features to be common knowledge, such as one player’s probability assessment about another’s preferences or information.

I foresee the progress of game theory as depending on successive reductions in the base of common knowledge required to conduct useful analyses of practical problems.

Only by repeated weakening of common knowledge assumptions will the theory approximate reality.”

R Wilson (1987), Game-Theoretical Analyses of Trading Processes, in TF Bewley, Ed., *Advances in Economic Theory, Fifth World Congress*, Cambridge U Press

How would we make an auction less reliant on common knowledge?

- One way: don't rely on buyers' *beliefs* at all
- In private value settings: direct mechanism where truth-telling is a *dominant strategy*, not just a best response
  - In IPV settings, generally no loss – seller can get same revenue with dominant strategies, so "no need" to worry about beliefs
- With common values, dominant strategies don't exist
  - Analogous concept is ex post implementation
  - "Truth-telling is dominant if you expect others to tell the truth, regardless of what you believe about the distribution of their valuations"
- In certain environments, any implementable outcome can be ex post implemented

D Bergemann and S Morris (2005), Robust Mechanism Design, *Econometrica* 73 (6)

What if seller doesn’t know what buyers know about valuations?

- Suppose *buyers* have common knowledge of environment…
- …but *seller* only knows distribution of valuations, not what information buyers have about them
- For example: first price auction, two bidders, pure common value $v \sim U[0,1]$, no idea what info bidders have about it
  - Could be: neither bidder knows anything, so revenue = 1/2
  - Could be: both know it exactly, revenue again = 1/2
  - Could be: one knows it exactly, one doesn’t, then revenue = 1/3
  - Could be any other information structure
- Question: what expected revenue can seller “robustly” predict?
  - Lower bound on expected revenue over all information structures?

D Bergemann, B Brooks and S Morris (2017), First Price Auctions with General Information Structures: Implications for Bidding and Revenue, *Econometrica* 73 (6)

What if seller doesn't know what buyers know about valuations?

- Stick with example (n = 2, v ~ U[0,1], unknown info structure)
- Suppose revenue (winning bid) is deterministic function β of v, weakly increasing, and either bidder is equally likely to win
- Revenue $R = \int_0^1 \beta(v)dv$, bidder surplus $\frac{1}{2}\int_0^1 (v - \beta(v))dv$
- For this to be an equilibrium, β can't be too low
  - If β was uniformly close to 0…
  - …each bidder would be getting surplus ≈ ¼…
  - …but even without information, could deviate to a slightly higher bid and get surplus close to ½
- Ruling out profitable "upward deviations" puts a lower bound on β, therefore lower bound on revenue

D Bergemann, B Brooks and S Morris (2017), First Price Auctions with General Information Structures: Implications for Bidding and Revenue, *Econometrica* 73 (6)

What if seller doesn't know what buyers know about valuations?

- One possible deviation: for some w, "whenever my equilibrium bid is below $\beta(w)$, bid $\beta(w)$ instead"
- Lots of math to show that if this is not a profitable deviation,

$$\beta(w) \geq \frac{1}{2w} \int_0^w \big(v + \beta(v)\big) dv$$

- Define a mapping Λ from functions to functions by

$$\Lambda(\beta)(w) = \frac{1}{2w} \int_0^w \big(v + \beta(v)\big) dv$$

- Λ turns out to be monotone, so

$$\beta(v) \geq \square\ \beta)(v) \geq \Lambda^2(\beta)(v) \geq \Lambda^3(\beta)(v) \geq \cdots$$

- And Λ is a contraction, so its fixed point gives the lower bound on β

D Bergemann, B Brooks and S Morris (2017), First Price Auctions with General Information Structures: Implications for Bidding and Revenue, *Econometrica* 73 (6)

What if seller doesn't know what buyers know about valuations?

- In the example (N = 2, $v \sim U[0,1]$), they find $\beta(v) \geq v/3$, so $R \geq 1/6$
- And they find an information structure that achieves that revenue
- For any joint distribution of valuations: tight lower bound on revenue, upper bound on bidder surplus for first-price auctions
- What's neat
  - In optimal mechanism with correlated values (earlier today), key constraint was that high valuation buyers not want to imitate low valuation buyers
  - Here, key constraint is that bids must be high enough so that bidders don't want to deviate to higher bids – this implies a lower bound on revenue
  - Out of all possible deviations, ruling out particular upward deviations is key
  - And making these constraints hold with equality yields the "worst case" information structure that achieves minimal revenue
- What's "robust" here isn't the auction, but the revenue bound

D Bergemann, B Brooks and S Morris (2017), First Price Auctions with General Information Structures: Implications for Bidding and Revenue, *Econometrica* 73 (6)

What if you're worried about resale?

Resale

- In asymmetric settings, optimal mechanism sometimes sells to bidder who doesn’t have highest valuation
- What if winner could resell the prize to a losing bidder?
- Would this change “strong” bidder’s behavior in original auction, reducing revenue?

An example

- Two bidders, with v_1 ~ U[0,10] and v_2 = 2
  - Virtual values

$$VV_1 = v_1 - \frac{(10 - v_1)/10}{1/10} = 2v_1 - 10$$

and $VV_2 = 2$

- Optimal auction sells to 1 if v_1 > 6, otherwise to 2
  - Seller should offer to buyer 1 for 6, sell to buyer 2 if he declines

- But…
  - Why wouldn't buyer 1 just wait, let buyer 2 "win," and try to buy it from him afterwards?

C Zheng (2002), Optimal Auction with Resale, *Econometrica* 70

An example

- Interestingly, this may not be a problem
  - Suppose after "winning," buyer 2 can run optimal auction
  - Since $VV_1 = 2v_1 - 10$ and $v_2 = 2$, buyer 2's optimal auction again sells for 6 when $v_1 > 6$
  - So 2's expected surplus from getting the good is
$$\frac{4}{10}(6) + \frac{6}{10}(2) = 3.6$$
  - So original seller can just sell to buyer 2 for 3.6
  - This is same revenue as optimal mechanism w/o resale!

- But, this doesn't always work when $n > 2$…
  or if "resale market" takes different form

C Zheng (2002), Optimal Auction with Resale, *Econometrica* 70

So let's ask a different question

- What if you have no idea how the resale market works, or what information players will have at that point…

- …but want a mechanism that performs well regardless?

- Can we solve

$$\max_{mechanisms} \min_{resale\ protocols} R(m)$$

- Turns out: yes!

G Carroll and I Segal (2019), Robustly Optimal Auctions with Unknown Resale Opportunities, *Review of Economic Studies* 86

Model

- n bidders, with independent private values θ_i
- Initial mechanism leads to allocation $p(\theta)$ and payments
- Resale market is modeled in reduced-form way
  - Some post-auction resale might happen
  - Depends on initial allocation p, and on bidders' types θ
  - Let $v_i(p,\theta)$ be final payoff achieved by bidder i, excluding payment made to original auctioneer
  - (for now, assume all private information revealed before resale, so resale doesn't depend on other details of play in auction)
- Resale markets replaces payoffs $p_i\theta_i$ with payoffs $v_i(p,\theta)$

G Carroll and I Segal (2019), Robustly Optimal Auctions with Unknown Resale Opportunities, *Review of Economic Studies* 86

Optimal auctions with unknown resale

- Restrictions on $v_i(p, \theta)$
  - Object can only be resold if it was sold, can't generate more surplus than highest valuation, so $\sum_i v_i(p,\theta) \leq \left(\max_i \theta_i\right)\left(\sum_i p_i(\theta)\right)$
  - And original buyer could keep object, so $v_i(p,\theta) \geq p_i \theta_i$
- If we knew how resale market operated, we'd know $v_i(p,\theta)$
  - could design optimal mechanism with that as each bidder's "valuation"
  - A little complicated, since v_i depends on other bidders' θ_j,
  - and can be nonzero even when $p_i = 0$,
  - but we could handle it
- But we want to handle "any" resale protocol, so "any" v_i
- Trick: "guess" the *worst* one and optimize for that!

G Carroll and I Segal (2019), Robustly Optimal Auctions with Unknown Resale Opportunities, *Review of Economic Studies* 86

“Worst case” resale market

- Suppose: after the auction, buyer with highest valuation learns all private information and has all the bargaining power
- Why is this worst from seller’s perspective?
  - “Problem” is high-value buyer skipping auction to wait for resale market
  - This resale protocol makes that most appealing
- In our two-buyer example from before (v_1 ~ U[0,10], v_2 = 2)
  - If high-value buyer has all the bargaining power post-auction…
  - …buyer 1 will skip initial auction, buy for 2 after the auction if $v_1 > 2$
  - Buyer 2 won’t get any extra surplus from resale, so can’t pay more than 2 to original seller
- Next: what’s the optimal auction if this is the resale protocol?

G Carroll and I Segal (2019), Robustly Optimal Auctions with Unknown Resale Opportunities, *Review of Economic Studies* 86

“Worst case” resale market

- Suppose: after the auction, buyer with highest valuation learns all private information and has all the bargaining power
- Let $i^*(\theta)$ be identity of buyer with highest value, then

$$v_i(p,\theta) = \begin{cases} p_i\theta_i + \sum_{j\neq i} p_j(\theta_i - \theta_j) & if \quad i = i^* \\ p_i\theta_i & otherwise \end{cases}$$

- And from here we end up with

$$VV_i = \begin{cases} \theta_i - \frac{1-F_i(\theta_i)}{f_i(\theta_i)} & if \quad i = i^* \\ \theta_i - \frac{1-F_i(\theta_i)}{f_i(\theta_i)} - \frac{1-F_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} & otherwise \end{cases}$$

- With this resale protocol, buyer with highest valuation *always* also has highest virtual valuation!
- So it’s never optimal to “mis-allocate”

Optimal auction with "worst case" resale

- Suppose: after the auction, buyer with highest valuation learns all private information and has all the bargaining power
- With this resale protocol, seller should never sell to buyer who doesn't have the highest valuation
  - So resale market "no longer matters," since there will never be resale
  - Also means this must be "worst-case" resale protocol we should plan for
- So, just solve original mechanism design problem with added constraint: if you sell, must sell to buyer with highest valuation!
- Only remaining question: when should you sell and when should you not?

Optimal auction with “worst case” resale

- Naively: sell to highest-value bidder whenever his VV > 0
- But…

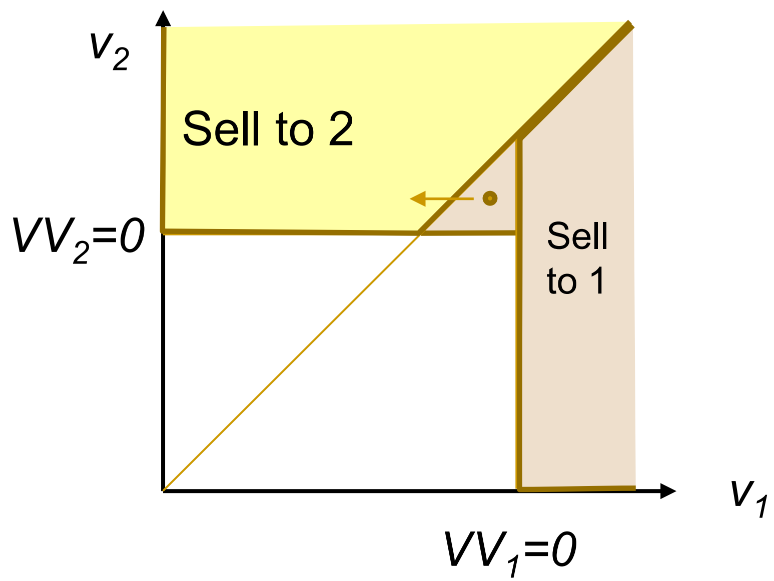

Optimal auction with "worst case" resale

- Naively: sell to highest-value bidder whenever his VV > 0
- Optimal mechanism turns out to be…
  - give each bidder a personal value threshold r_i
  - allocate good to bidder with highest value if *any* bidder's value exceeds their own threshold
  - (Mechanism was proposed much earlier by Ausubel and Cramton, who suggested it would be optimal if "resale was perfect" but didn't formalize what that meant)
  - Still need to calculate the thresholds, paper gives an algorithm
- This is auction that gives highest "worst-case" revenue, where worst case is over all possible resale markets
- So this auction is "robust to all possible forms of resale"
  - Performs same for any resale protocol, since never mis-allocates so resale never happens

What if you're worried about collusion?

An example

- Two buyers, independent private values $v_i \sim U[1,2]$
- Seller has $v_0 = 0$
- Optimal auction is a second-price auction with reserve price of 1, expected revenue is 4/3
- But what if buyers collude?
  - Before auction, they meet and run “knockout auction” to determine which of them has higher value, he’ll bid unopposed
  - Then revenue falls to 1
- Can we find an auction that isn’t vulnerable to collusion, even if we don’t know what form collusion will take?

YK Che and J Kim (2006), Robustly Collusion-Proof Implementation, *Econometrica* 74 (4)

A “collusion-proof” mechanism

- Two buyers, independent private values $v_i \sim U[1,2]$
- Alternative mechanism:
  - Each buyer submits a bid
  - Higher bidder gets object
  - Lower bidder pays 4/3, receives winner’s bid
- Equilibrium is for each to bid $b(v_i) = \frac{5}{6} + \frac{1}{3} v_i$
- Seller receives 4/3 no matter what – doesn’t care if buyers collude, or how!
- (Has flavor of “selling the firm to the agent”)

YK Che and J Kim (2006), Robustly Collusion-Proof Implementation, *Econometrica* 74 (4)

Analogous mechanisms can be designed for a lot of settings

- Assuming bidders are risk-neutral…
  - If values are independent and all are colluding together, any implementable outcome can be made "robustly collusion-proof"
  - If values are correlated, need an additional condition
  - If only a subset are colluding, seller needs to know identities of at least two of them
- But…
  - This mechanism is back to being weird and complicated
  - Needs to be tailored precisely to environment
  - Buyers need to commit to participate before they have a chance to pool their information – if they already knew they both had low valuations, neither would agree to participate

YK Che and J Kim (2006), Robustly Collusion-Proof Implementation, *Econometrica* 74 (4)

What if the buyers worry the seller will cheat?

Trusting the seller

- We've been assuming seller has commitment power
- In a direct revelation mechanism, buyers reveal their valuations
  - Have to trust seller won't change the rules…
  - …or lie about other buyers' valuations, and therefore what outcome and payment they get
- What if seller isn't completely trustworthy?

M Akbarpour and S Li (2020), Credible Auctions: A Trilemma, *Econometrica* 88 (2)

"Credible" mechanisms

- Suppose I bid $100 in a **second-price** sealed-bid auction, second-highest bid was $60
  - Seller has incentive to lie about second-highest bid so I'll pay more
  - If he tells me someone bid was $95, I have no way to know
- Suppose I bid $85 in a **first-price** sealed-bid auction, second-highest bid was $70
  - If he asks me for more than $85, I know he's changing the rules
  - If he approaches the losing bidder and offers to sell to him at $90, that bidder knows he's changing the rules
  - Any way seller could increase revenue, *some* buyer would know he cheated
- An auction is *credible* if any way a seller could profit by breaking the rules, some single bidder would detect

M Akbarpour and S Li (2020), Credible Auctions: A Trilemma, *Econometrica* 88 (2)

Different characteristics we might want in an auction

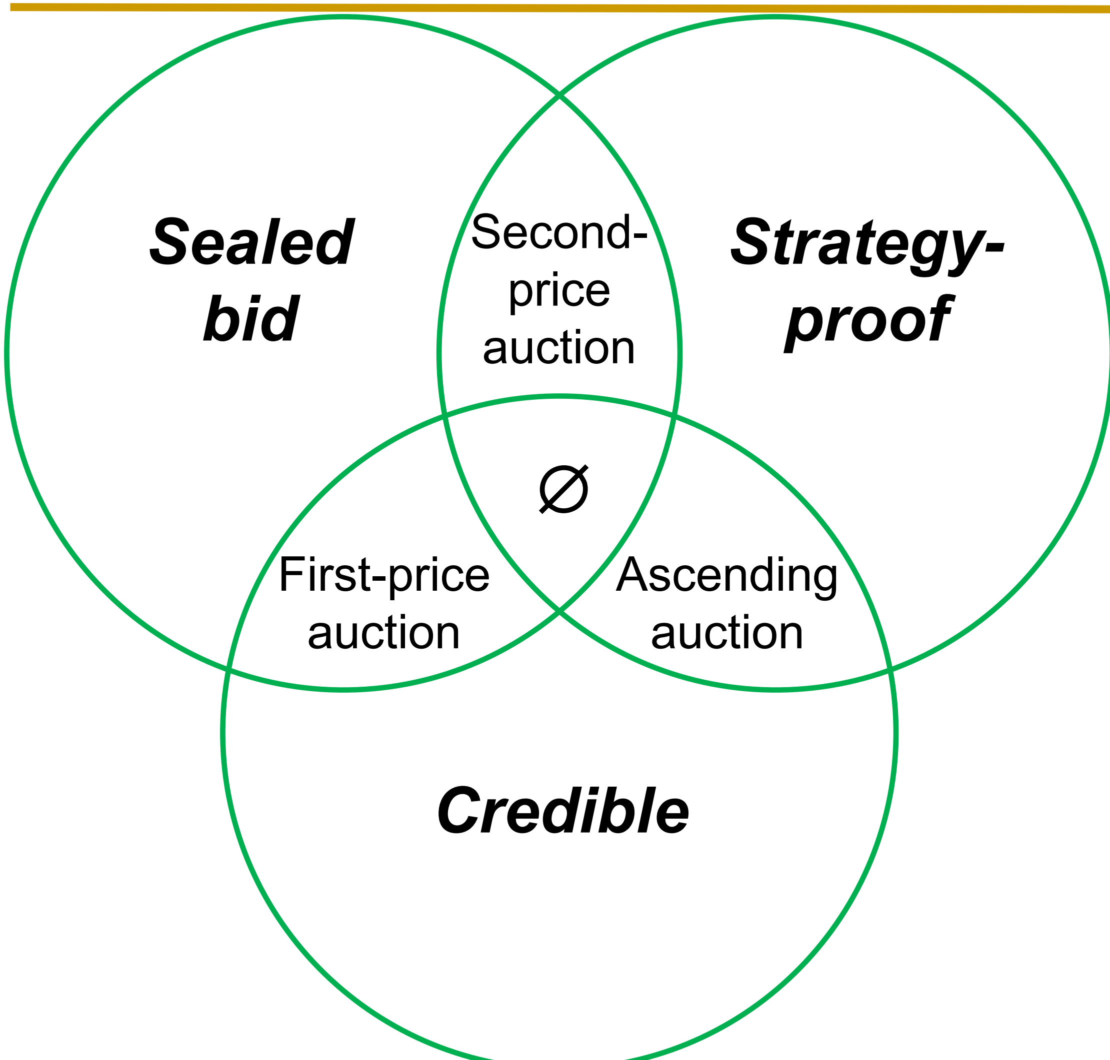

- Credible mechanisms more "robust to buyers not trusting seller"…
- …but you can't have everything!

M Akbarpour and S Li (2020), Credible Auctions: A Trilemma, *Econometrica* 88 (2)

A couple other ideas

What if the seller doesn't know the distribution of buyer valuations?

- Running optimal auction requires "knowing *F*" – what if you don't?

- If $n > 1$, you can just run second-price auction with no reserve
  - Jason Hartline papers on revenue guarantees of "prior-free" mechanisms in different settings

- With symmetric IPV, learn optimal mechanism reasonably fast
  - Example last week with one buyer and one "sample draw"
  - To get $(1-\varepsilon)$ times optimal revenue, need on the order of ε^{-3} observations

T Roughgarden (2014), Approximately Optimal Mechanism Design: Motivation, Examples, and Lessons Learned, *ACM SIGEcom Exchanges* 10 (2)

What if you're worried about "bad equilibria"?

- So far, we've just required truth-telling to be *an* equilibrium
  - "Partial implementation"

- "Full implementation" – *every* equilibrium of game should implement seller's desired outcome

- One approach
  - Start with mechanism where that's an equilibrium…
  - …then add "side bets" to disrupt other equilibria
  - Leads to messy, complicated mechanisms

E Maskin (1999), Nash Equilibrium and Welfare Optimality, *Review of Economic Studies* 66 (1)
M Ollar and A Penta (2017), Full Implementation and Belief Restrictions, *AER* 107 (8)

FUN SIDE NOTE: I asked ChatGPT: “what are the best references for "full implementation" in mechanism design in economics?”

1. "Mechanism Design and Implementation" by Roth and Sotomayor
   Doesn’t exist AFAIK, but R+S did write a great book on matching markets!

2. "A Course in Game Theory" by Osborne and Rubinstein
   Textbook on game theory

3. "Mechanism Design: A Linear Programming Approach" by Roughgarden
   Actually written by Rakesh Vohra

4. "Auctions, Theory, and Practice" by Krishna
   Book on auctions

5. "Algorithmic Game Theory" by Nisan, Roughgarden, Tardos, and Vazirani
   Book on algorithmic game theory

6. "Designing Economic Mechanisms" by McAfee and McMillan
   Actually by Leonid Hurwicz and Stanley Reiter

Takeaways

Main point of today

- Auctions/mechanisms that are theoretically optimal in a setting often seem “fragile”
- We’d like mechanisms that don’t fail spectacularly if a modeling assumption is violated
- Lots of different conceptions of what “robust” means, often at odds with each other
  - Simple mechanisms seem more robust…
  - …but eliminating “bad equilibria” or the threat of collusion requires weird, complicated mechanisms
- No silver bullet, just a thing to think about

Another fun example of an auction being "fragile"

German spectrum auction, 1999

T-Mobil					Mannesmann				
20	20	20	20	20	20	20	20	20	20
18.18	18.18	18.18	18.18	18.18					

US spectrum auction, 1996-7

License 378 (Rochester MN)

US West

Two licenses in Iowa

McLeod	**McLeod**
$313,378	**$62,378**
McLeod	**McLeod**

P Klemperer (2002), What Really Matters in Auction Design, *Journal of Economic Perspectives* 16 (1)

Back to big picture

- Last three lectures: theory
  - Understanding seller's problem through virtual value…
  - …bidders' investments through externalities…
  - …and "robustness" in various ways
- To design the right mechanism or predict how outcomes will change under various counterfactuals, need to know primitives of the environment
- Tomorrow: how we learn them from observable data

Thank you!

References

R. Myerson (1981), Optimal Auction Design, *Mathematics of Operations Research* 6(1)

J Crémer and R McLean (1988), Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions, *Econometrica* 56(6)

R Wilson (1987), Game-Theoretical Analyses of Trading Processes, in TF Bewley, Ed., *Advances in Economic Theory, Fifth World Congress*, Cambridge U Press

D Bergemann and S Morris (2005), Robust Mechanism Design, *Econometrica* 73 (6)

D Bergemann, B Brooks and S Morris (2017), First Price Auctions with General Information Structures: Implications for Bidding and Revenue, *Econometrica* 73 (6)

C Zheng (2002), Optimal Auction with Resale, *Econometrica* 70

G Carroll and I Segal (2019), Robustly Optimal Auctions with Unknown Resale Opportunities, *Review of Economic Studies* 86

YK Che and J Kim (2006), Robustly Collusion-Proof Implementation, *Econometrica* 74 (4)

M Akbarpour and S Li (2020), Credible Auctions: A Trilemma, *Econometrica* 88 (2)

T Roughgarden (2014), Approximately Optimal Mechanism Design: Motivation, Examples, and Lessons Learned, *ACM SIGEcom Exchanges* 10 (2)

E Maskin (1999), Nash Equilibrium and Welfare Optimality, *Review of Economic Studies* 66 (1)

M Ollar and A Penta (2017), Full Implementation and Belief Restrictions, *American Economic Review* 107 (8)

P Klemperer (2002), What Really Matters in Auction Design, *J of Economic Perspectives* 16 (1)