Recurring Themes in Auction Theory and Mechanism Design

Part II: Pre-Auction Choices and Externalities

Dan Quint
U of Wisconsin
Visitor, UTokyo and UTMD

February 14, 2023
University of Tokyo

Overview

- Yesterday:
  - Expected revenue is EV of winner's virtual value
  - We took the *set of bidders*, their *valuations*, and their *information* as given
  - Fixed set of *n* bidders, private values drawn from known distributions F_i

- Today: *pre-auction* decisions
  - Bidder entry, information acquisition, investment
  - Focus on efficiency, rather than revenue
  - We'll use the lens of externalities

Externalities
in auctions

Consider a second-price auction with private values

- $n - 1$ other bidders will bid their valuations
- Let v_{max} be highest valuation among the other bidders, v_s seller's cost/valuation for object
- Consider payoffs of other players

	seller	***v_{max} guy***	***other bidders***	***everyone but me***
I bid $b > v_{max}$ and win	$v_{max} - v_s$	0	0	$v_{max} - v_s$
I bid $b < v_{max}$ and set price	$b - v_s$	$v_{max} - b$	0	$v_{max} - v_s$
I bid $b < b_2$ and don't matter	$b_2 - v_s$	$v_{max} - b_2$	0	$v_{max} - v_s$
I oversleep and don't show	$b_2 - v_s$	$v_{max} - b_2$	0	$v_{max} - v_s$

Consider a second-price auction with private values

- A bidder's decision of whether and how to bid imposes ***no net externality*** on the rest of the game
- So decisions that affect any of these are likely to be made efficiently!

	seller	***v_{max} guy***	***other bidders***	***everyone but me***
I bid $b > v_{max}$ and win	$v_{max} - v_s$	0	0	$v_{max} - v_s$
I bid $b < v_{max}$ and set price	$b - v_s$	$v_{max} - b$	0	$v_{max} - v_s$
I bid $b < b_2$ and don't matter	$b_2 - v_s$	$v_{max} - b_2$	0	$v_{max} - v_s$
I oversleep and don't show	$b_2 - v_s$	$v_{max} - b_2$	0	$v_{max} - v_s$

Consider a second-price auction with private values

- This is *not* true for first-price auctions – my entry or value distribution may change sum of others' payoffs…
- …and we can use sign of this net externality to see how choices are distorted away from efficient

	seller	***v_{max} guy***	***other bidders***	***everyone but me***
I bid $b > v_{max}$ and win	$v_{max} - v_s$	0	0	$v_{max} - v_s$
I bid $b < v_{max}$ and set price	$b - v_s$	$v_{max} - b$	0	$v_{max} - v_s$
I bid $b < b_2$ and don't matter	$b_2 - v_s$	$v_{max} - b_2$	0	$v_{max} - v_s$
I oversleep and don't show	$b_2 - v_s$	$v_{max} - b_2$	0	$v_{max} - v_s$

Example: entry

Auction with endogenous entry

- *n* potential bidders
- Costs a bidder *c* to “enter” and learn valuation
- Potential bidders decide simultaneously whether to enter
- Symmetric mixed strategy equilibrium where entrants earn expected surplus of exactly *c* from entering
- What reserve price induces efficient level of entry?

D Levin and J Smith (1994), Equilibrium in Auctions with Entry, *American Economic Review* 84(3)

We can think about the symmetric mixed-strategy equilibrium

- Let $\pi(m, r)$ be a bidder's expected surplus in *m*-bidder auction with reserve *r*

- If each bidder enters with probability *q*, then

$$c = \sum_{j=0}^{n-1} \binom{n-1}{j} q^j (1-q)^{n-1-j} \pi(j+1, r)$$

- This looks messy – is there an easier way?

D Levin and J Smith (1994), Equilibrium in Auctions with Entry, *American Economic Review* 84(3)

Think about externality caused by a bidder's decision to enter a SP auction

- If at least one other entrant, 0 net externality
- If no other entrant, then…
  - by entering, he'll win and pay r
  - seller will get surplus of $r - v_s$ instead of 0
- Net externality from a bidder's decision to enter is

$$Pr(\text{no other entrants})\ (r - v_s)$$

  - If $r > v_s$, entry has positive externality so "not enough entry"
  - If $r < v_s$, entry has negative externality so "too much entry"
  - If $r = v_s$, externality is 0 so "efficient entry"
  - And $r = v_s$ is also efficient post-entry
- So reserve of $r = v_s$ maximizes social surplus

D Levin and J Smith (1994), Equilibrium in Auctions with Entry, *American Economic Review* 84(3)

Is there a tradeoff between revenue and efficiency?

- With fixed *n*, *r* solving $r - \frac{1-F(r)}{f(r)} = 0$ maximizes revenue
  - Or $r - \frac{1-F(r)}{f(r)} = v_s$ maximizes seller profit
  - But requires $r > v_s$, which is ex post inefficient
- With endogenous entry, $r = v_s$ maximizes total surplus, and *also maximizes seller profits*
  - Buyers decide to enter before learning valuations
  - Mixed-strategy equilibrium → zero expected surplus
  - Seller captures all surplus, so maximizing surplus also maximizes profits
- Extends to first-price auctions via revenue equivalence (if buyers observe # of entrants before bidding)

D Levin and J Smith (1994), Equilibrium in Auctions with Entry, *American Economic Review* 84(3)

Caveats?

- Setting $r = v_s$ maximizes surplus and seller profit within class of auctions with *unrestricted entry*

- But randomness from mixed strategies is inefficient
  - Post-entry surplus is concave in number of bidders

- Seller can improve by rationing entry to be close to the expected number from the mixed equilibrium
  - Instead of 10 potential bidders all mixing 50-50…
  - …better to have 5 bidders entering for sure

D Levin and J Smith (1994), Equilibrium in Auctions with Entry, *American Economic Review* 84(3)
RP McAfee and J McMillan (1987), Auctions with Entry, *Economics Letters* 23

What if buyers know valuations when deciding whether to enter?

- Symmetric equilibrium with entry threshold
- "Marginal entrant" only wins if he's only entrant, pays r
- Externality is still Pr(no other entrants)($r - v_s$)
  - So $r = v_s$ still maximizes total surplus
- ***But*** seller no longer captures all the surplus
  - At $r = v_s$, increasing r slightly gives "second-order" reduction in total surplus…
  - …but first-order reduction in bidder surplus…
  - …so $r > v_s$ maximizes seller profits

WF Samuelson (1985), Competitive Bidding with Entry Costs, *Economics Letters* 17

Example: value-enhancing investments

Consider a pre-auction investment that affects a bidder's valuation

- Symmetric, IPV setting with fixed *n*
- Before auction, I can make costly investment that will increase my valuation (in FOSD sense)
- Will first- or second-price auction lead to more investment?
  Which is more efficient?

L Arozamena and E Cantillon (2004), Investment Incentives in Procurement Auctions, *Review of Economic Studies* 71(1)

What externalities does my investment cause?

	Second price auction
Effect on other bidders	*negative*
Effect on seller	*positive*
Total net externality	*zero*
Investment level	*efficient*

L Arozamena and E Cantillon (2004), Investment Incentives in Procurement Auctions, *Review of Economic Studies* 71(1)

What externalities does my investment cause?

	Second price auction	***First price auction***
Effect on other bidders	*negative*	
Effect on seller	*positive*	
Total net externality	*zero*	
Investment level	*efficient*	

- (Revenue equivalence does *not* make this question moot
- Even if outcome is symmetric so revenue equivalence "should hold"…
- ..."off-equilibrium-path" outcomes are asymmetric, determine when investment stops being worthwhile)

L Arozamena and E Cantillon (2004), Investment Incentives in Procurement Auctions, *Review of Economic Studies* 71(1)

What externalities does my investment cause?

	Second price auction	***First price auction***
Effect on other bidders	*negative*	*less negative than second-price*
Effect on seller	*positive*	*more positive than second-price*
Total net externality	*zero*	*positive*
Investment level	*efficient*	*less than efficient*

- Investment makes me “strong bidder” in asymmetric auction
- Asymmetric FP auction can be higher- or lower-revenue…
- …but under many conditions make it higher-revenue…

L Arozamena and E Cantillon (2004), Investment Incentives in Procurement Auctions, *Review of Economic Studies* 71(1)

So under many (but not all) conditions…

- Under first-price auction, value-enhancing investments induce positive externality…
- …so first-price auction induces less than efficient amount of investment
- In symmetric setting where all bidders can invest and equilibrium is symmetric…
  - first-price auction has lower than efficient investment…
  - second-price auction has efficient investment…
  - …and by revenue equivalence, same level would be efficient for both, so first-price has lower investment

L Arozamena and E Cantillon (2004), Investment Incentives in Procurement Auctions, *Review of Economic Studies* 71(1)

Arozamena and Cantillon explain it differently

- “We find that after the investment, the investor’s opponents will collectively bid more aggressively.
- …In the language of industrial organization, investment has a negative strategic effect in the FPA. This erodes its benefits.
- …Under the same condition… the FPA will induce less investment than the SPA.
- …The fact that the SPA generates the socially efficient investment incentives provides us with a clear normative interpretation of this underinvestment result.”

L Arozamena and E Cantillon (2004), Investment Incentives in Procurement Auctions, *Review of Economic Studies* 71(1)

Example: information acquisition

Suppose bidders must invest to learn their valuation more precisely

- Symmetric, IPV setting with fixed *n*
- Before auction, bidders simultaneously choose how precise a signal to get about their own valuation
- Will first-price or second-price auction lead to more information acquisition? Which is more efficient?

D Bergemann and J Valimaki (2002), Information Acquisition and Efficient Mechanism Design, *Econometrica* 70(3)
D Hausch and L Li (1993), Private Value Auctions with Endogenous Investment: Revenue Equivalence and Non-Equivalence, working paper
N Persico (2000), Information Acquisition in Auctions, *Econometrica* 68(1)

With risk-neutral bidders, easy interpretation of “more information”

- To risk-neutral bidder, what matters is expected value of ex post valuation, conditional on information he has pre-auction

- More precise signal about unobserved truth corresponds to a mean-preserving spread of this expected value
  - Bidder with no info has point beliefs at expected value
  - Bidder with perfect info has distribution *F* on posterior expected value

- Think of “choosing more precise information” as “switching to a more disperse distribution of valuations”
  - Recall that in any mechanism, $U_i(v_i) = U_i(a_i) + \int_{a_i}^{v_i} E_{v_{-i}} p(s, v_{-i}) ds$
  - $U_i{}'(v_i) = E_{v_{-i}} p(v_i, v_{-i})$ is increasing in v_i
  - Expected surplus $U_i(v_i)$ is convex in v_i, so more info is always valuable!
  - (*if* no strategic response from other bidders)

Second-price auction

- Bidder’s bid imposes no net externality…
- …so information acquisition imposes no net externality…
- …so information acquisition should be efficient
- Doesn’t matter whether bidders see much information their rivals acquire
- What about first-price auction?

First-price auction

- More complicated – and depends on whether information acquisition is observable

- What externality does "overt information acquisition" impose?
  - Acquiring better information makes you "well informed" bidder in asymmetric first-price auction
  - Not much known about asymmetric FP auctions where one bidder is "higher-variance" than others
  - More information makes you "stronger when you're strong," but also "weaker when you're weak"

- A useful special case might be large *N*:
  - Winner will be high-value, so top of bidder's value distribution matters
  - More information makes a bidder stronger, so results from before apply
  - (Under certain conditions, FP auction leads to less info acquisition)

First-price auction

- What about covert information?
- If rivals don't see how much information you acquire, there's no strategic response
- So once you know the interim expected value of your valuation, you face same optimization problem regardless of how much information it's based on…
- …so optimal bid, and expected payoff at that point, are same
- But what is effect of your valuation on other players' surplus?

First-price auction

- What about covert information?

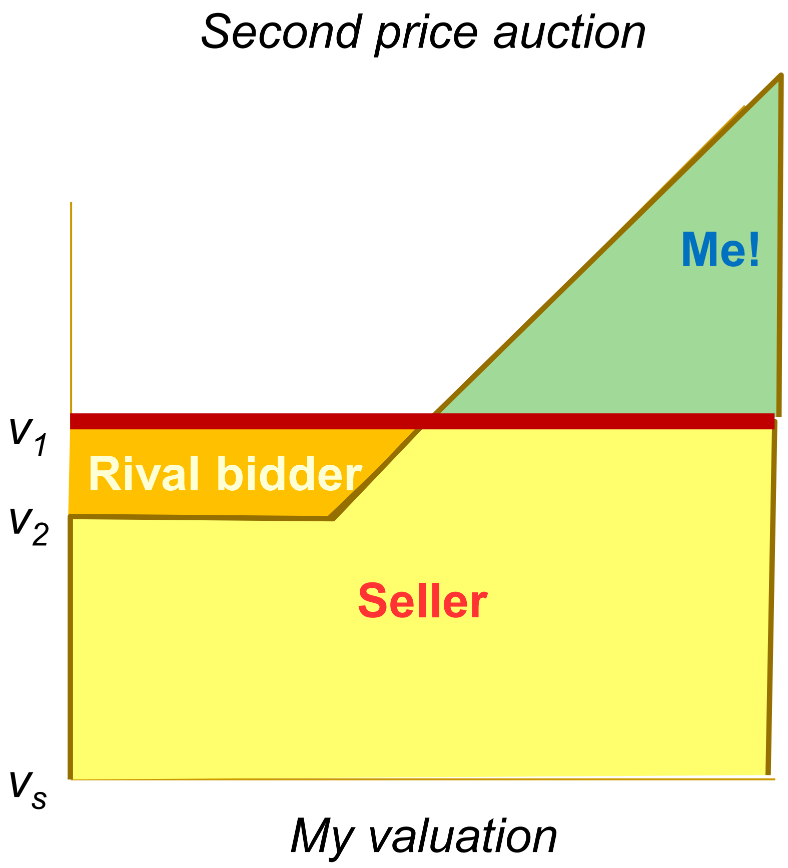

First-price auction

- What about covert information?

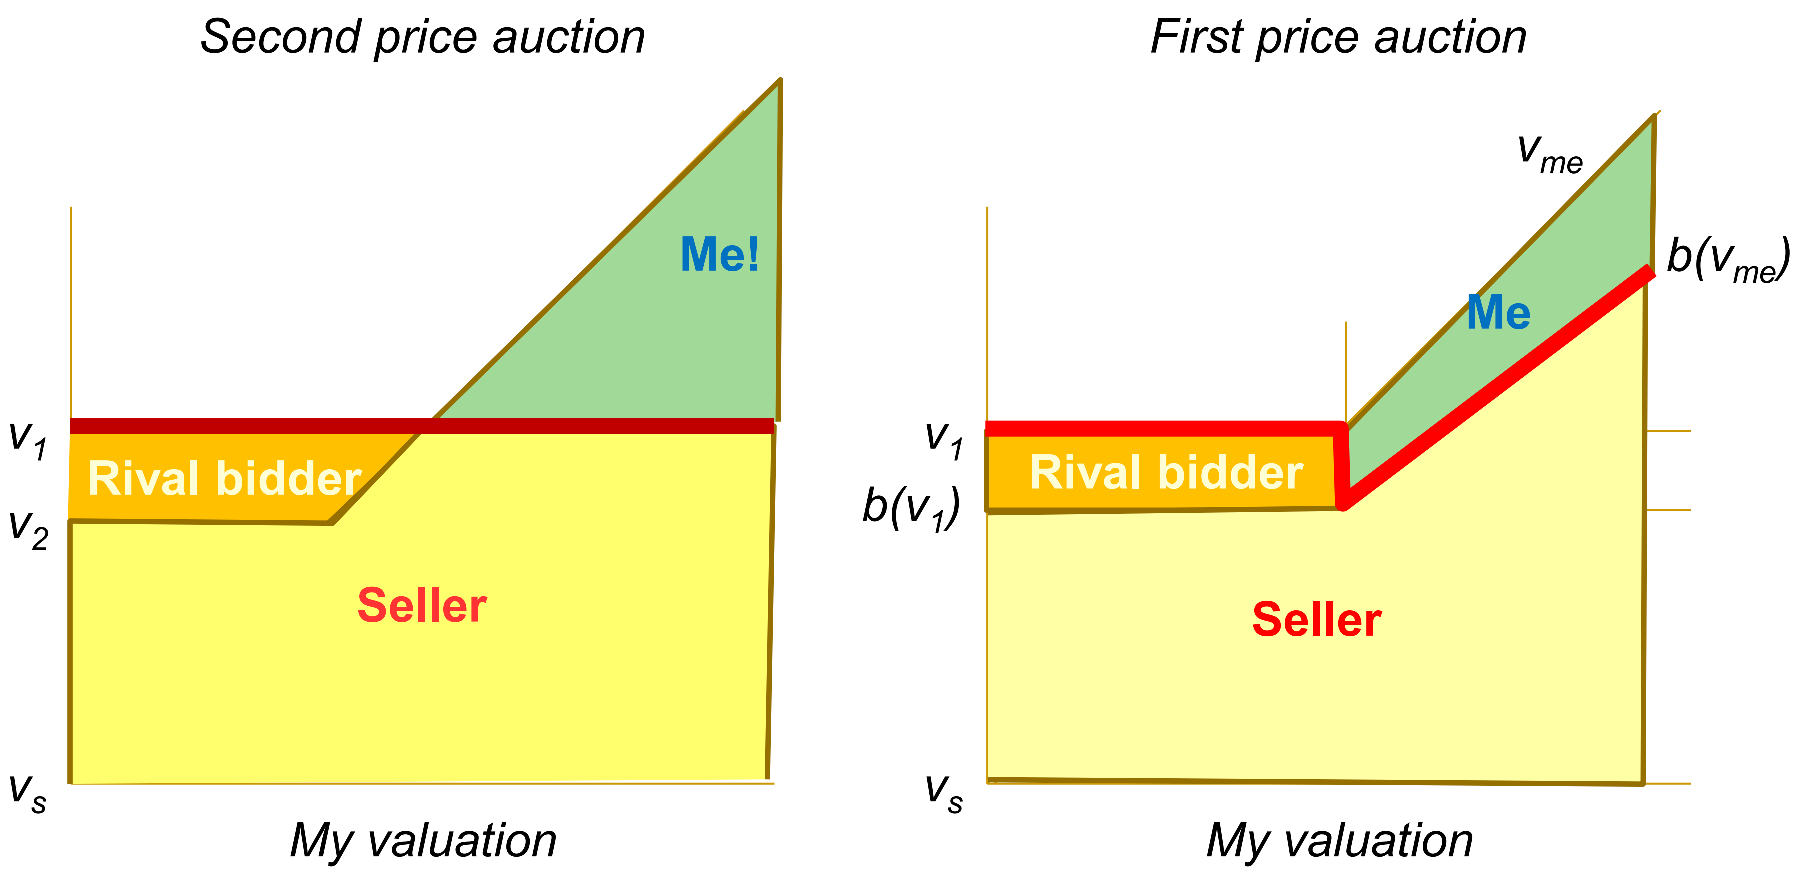

First-price auction

- What about covert information?
  - Covertly increasing my valuation can impose a positive *or* negative externality!
  - Acquiring *a lot* more information probably imposes a positive externality…
  - …but we care about the incentive on the margin
  - If *n* is high, $b(v)-v$ will be small, so “dip” will be small…
  - …but strongest rival will be near the top of value distribution
  - Tricky to sign externality this way!

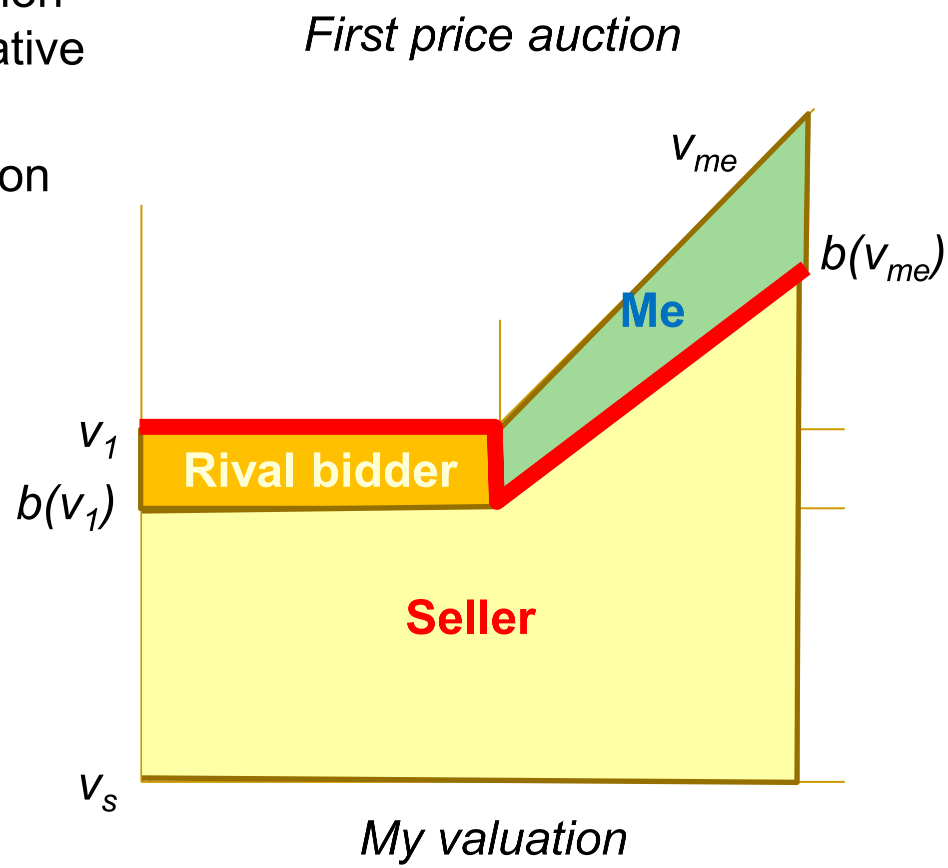

First-price auction

- What about covert information?

- Persico (2000) is the classic
  - Focuses on two-bidder case
  - Different model, with valuations correlated and interdependent

- He finds FPA "more risk-sensitive" than SPA
  - payoff falls off more quickly when you bid sub-optimally
  - so better information is more valuable in FPA on the margin
  - partly because more precise information about your own valuation tells you more about opponent's valuation as well, and therefore you know more about his likely bid

N Persico (2000), Information Acquisition in Auctions, *Econometrica* 68(1)

Beyond single-item auctions

Auctions for multiple goods

- Multiple items, buyers may have different private value for each (and for combinations)

- Vickrey Clarke Groves mechanism generalizes the second-price auction
  - Bidders report their valuations
  - Allocation is set to maximize total surplus
  - Bidder *j* pays difference in other bidders' surplus between efficient allocation with *j* and without *j*

- Famously, ex post efficient and strategy-proof
  - Reporting true preferences is a dominant strategy…
  - …and VCG selects efficient allocation given reported prefs

Auctions for multiple goods

- VCG is designed to eliminate externalities
  - Payment rule gives each bidder payoff equal to their contribution to total surplus
  - So a bidder's report doesn't change combined surplus of other players (other bidders plus seller)

- No externalities → efficient investment
  - Rogerson (1992): "…Groves mechanisms provide not only a first-best solution to the simple collective choice problem (as has been established in the existing literature) but also a solution to the collective choice problem when ex ante investments must be made."

W Rogerson (1992), Contractual Solutions to the Hold-Up Problem, *REStud* 59(4)

So that's the good news… ***but***…

- VCG is computationally "hard"
  - Requires finding efficient allocation
  - → computational demand is exponential in number of objects
- "True" VCG isn't feasible in "large" settings
  - (Example: 2017 FCC "incentive auction" to repurpose TV broadcast rights for 5G mobile
  - 705 "sellers," 62 "buyers," 2912 licenses, and millions of pairwise feasibility constraints due to interference between stations)

Approximation-based VCG

- One option in large settings: use faster (polynomial-time) algorithm to find *approximately* optimal allocation
- Example of what such an algorithm might look like:
  - Let *m* be number of objects
  - Pick a “small” number *c*
  - Calculate the most efficient allocation with only *c* winners
  - Ignore bids for more than sqrt(*m*/*c*) objects, and run a greedy algorithm on the remaining bids
  - Take the better of these two allocations
- How do approximation-based VCG mechanisms perform when buyers face investment opportunities?

D Lehmann, R Müller and T Sandholm, The Winner Determination Problem, in P Cramton, Y Shoham and R Steinberg (2006), *Combinatorial Auctions*, MIT Press

How do approximation-based mechanisms perform?

- Suppose we use "fast" algorithm to find approximately optimal allocation, apply VCG payment rule
- Turns out: any "reasonable" VCG-based mechanism like this is not strategy-proof
  - "Reasonable": if only one buyer wants an object, they get it
  - "VCG-based" rule is only strategy-proof if it chooses exactly efficient allocation out of a restricted set of possible ones
  - Rules out "reasonable" VCG approximations besides exact VCG
- So this strategy won't yield mechanisms that are actually strategy-proof
- What about incentives for investment?

N Nisan and A Ronen (2007), Computationally Feasible VCG Mechanisms, *JAIR* 29

How do approximation-based mechanisms perform?

- For ex post efficient mechanisms:
  efficient investment incentives $\leftrightarrow$ strategy-proof
- What about mechanisms that are not exactly efficient *or* exactly strategy-proof?
- Turns out, “almost” ex post efficient + “almost” strategy-proof implies “almost” efficient investment incentives
  - If mechanism always yields surplus within η of optimal,
  - and each bidder’s gain from misreporting is bounded above by ε,
  - then maximum gain from investing an amount other than the social optimum is bounded above by $(\varepsilon + \eta)k$,
  - where k is number of relevant outcomes per player

JW Hatfield, F Kojima, and SD Kominers (2019), Strategy-Proofness, Investment Efficiency, and Marginal Returns: An Equivalence, working paper

That's the good news

- If a mechanism is close to efficient and close to strategy-proof, gain from investing other than socially optimal amount is also "small"
  - Though with a multiplier based on number of alternatives
- But, even if gains from non-socially-optimal investment are small, impact on surplus could still be large
- Alternative approach
  - instead of asking how close to optimal efficient strategies are…
  - …ask how far from efficient outcome is if players follow exactly optimal strategies

M Akbarpour, SD Kominers, KM Li, S Li, and P Milgrom, Algorithmic Mechanism Design with Investment, working paper

Approximation-based mechanisms can still be made exactly strategy-proof

- Any algorithm that chooses approximately efficient allocation…
- …is a mapping from reported preferences to allocations
- As long as mapping is monotone, the right payment rule makes it strategy-proof
- If a given algorithm for choosing allocation performs "pretty well" for fixed preferences…
- …does it still perform "pretty well" when buyers have an opportunity to invest?

M Akbarpour, SD Kominers, KM Li, S Li, and P Milgrom, Algorithmic Mechanism Design with Investment, working paper

Is a "pretty efficient" mechanism still "pretty efficient" with investment?

- Focus on mechanism's *surplus guarantee*

$$\inf_{\textit{Instances of environment}} \frac{\textit{Surplus achieved by algorithm}}{\textit{First-best surplus}}$$

- A mechanism is β-efficient if for every possible instance of the environment, total surplus $\geq \beta$ x first-best surplus

- Question: if a mechanism is β-efficient for fixed preferences, how efficient is it with investment?

M Akbarpour, SD Kominers, KM Li, S Li, and P Milgrom, Algorithmic Mechanism Design with Investment, working paper

We already know…

- …if my report doesn't impose an externality on other players…
- …then my valuation doesn't impose an externality…
- …and I'll make efficient investment decisions
- But in richer environment, I have lots of ways to change my report and potentially cause an externality
- Big advance: figuring out **which externalities matter**

M Akbarpour, SD Kominers, KM Li, S Li, and P Milgrom, Algorithmic Mechanism Design with Investment, working paper

Which externalities matter?

- Focus on buyer *j*, who has a vector v_j of preferences over a finite set of outcomes *O*
- Suppose given reported preferences $v = (v_j, v_{-j})$, the algorithm gives buyer *j* outcome *o*
- A change in *j*'s preferences from v_j to v_j' *confirms outcome o* if it increases *j*'s valuation for outcome *o* more than for any other outcome *o'*
  - Change reinforces efficiency of giving outcome *o* to buyer *j*
- Paper shows if a mechanism is β-efficient without investment…
  - *In general*: could have arbitrarily low surplus guarantee with investment
  - *But*, if the allocation rule is such that *confirming changes do not impose negative externalities*, then it remains β-efficient with investment
  - To get an approximately efficient mechanism to still perform well with investment, design it to not have any negative externalities from confirming preference changes

Wrapping up

Takeaway from today?

- With single-good auctions…
  - Second-price auction eliminates externalities, first-price does not
  - Second-price auction leads to efficient entry (when $r = v_s$), efficient investment, efficient information acquisition
  - Signing externalities gives an elegant way to sign distortion from first-price auction

- With multiple-good auctions…
  - VCG eliminates externalities → efficient investment
  - When VCG is infeasible, approximation-based mechanisms that mimic it don't create large perverse incentives…
  - …and can be designed to give good performance when investment incentives are taken into account

Big picture

- So far…
  - Seller's problem with fixed set of bidders, info, valuations
  - Bidders' pre-auction decisions, and effect on efficiency
- Next Monday:
  - Different ways to think about "robustness" in auctions
  - "Robustness" $\approx$ "auctions that still do OK even when some of your modeling assumptions are wrong"…
  - …but can mean many different things
- Next Tuesday:
  - How a theorist thinks about empirical research in auctions
  - Including some of my own work on making it more "robust"

Thank you!

References

D Levin and J Smith (1994), Equilibrium in Auctions with Entry, *Amer. Econ. Review* 84(3)

RP McAfee and J McMillan (1987), Auctions with Entry, *Economics Letters* 23

WF Samuelson (1985), Competitive Bidding with Entry Costs, *Economics Letters* 17

L Arozamena and E Cantillon (2004), Investment Incentives in Procurement Auctions, *Review of Economic Studies* 71(1)

D Bergemann and J Valimaki (2002), Information Acquisition and Efficient Mechanism Design, *Econometrica* 70(3)

N Persico (2000), Information Acquisition in Auctions, *Econometrica* 68(1)

D Hausch and L Li (1993), Private Value Auctions with Endogenous Investment: Revenue Equivalence and Non-Equivalence, working paper

W Rogerson (1992), Contractual Solutions to the Hold-Up Problem, *Review of Economic Studies* 59(4)

P Cramton, Y Shoham and R Steinberg (2006), *Combinatorial Auctions*, MIT Press

N Nisan and A Ronen (2007), Computationally Feasible VCG Mechanisms, *Journal of Artificial Intelligence Research* 29

JW Hatfield, F Kojima, and SD Kominers (2019), Strategy-Proofness, Investment Efficiency, and Marginal Returns: An Equivalence, working paper

M Akbarpour, SD Kominers, KM Li, S Li, and P Milgrom, Algorithmic Mechanism Design with Investment, working paper