Econ 805 – Advanced Micro Theory I
Dan Quint
Fall 2007
Lecture 23 – Nov 27 2007

- Thursday, we start student presentations:
- Kyoung Jin: "Continuous Auctions and Insider Trading," by Albert Kyle
- Michael Rapp: "Auction Based Queue Disciplines," by Kittsteiner and Moldovanu

So, what have we done this semester?

- We began by introducing the workhorse model of an auction as a Bayesian game: the Independent Private Values model. Each bidder knows their own reservation value for the object, but views competing bidders' valuations as random variables; and bidding functions form a symmetric Bayesian Nash equilibrium of the game of imperfect information. Bidders are symmetric, risk-neutral, and their valuations are independent.
- We showed that with independent private values, a number of common auction formats were well-approximated by either first- or second-price auctions with sealed bids.
- We proved the Envelope Theorem, which establishes that the value function of a parameterized optimization problem is the integral of the derivative of the objective function, evaluated everywhere at the optimum. We used this to show that **any** auction mechanism allocating the object to the bidder who values it the most highly, and giving 0 expected payoff to a bidder with the lowest possible valuation, gives the same expected payoff to each type of bidder, and therefore the same expected revenue to the seller.
- Recasting auctions as mechanisms, and without loss of generality as direct-revelation mechanisms, we showed that incentive-compatibility was basically equivalent to the Envelope Theorem, and that any incentive-compatible mechanism was uniquely pinned down by the rule for allocating the object and the expected payoffs realized by the lowest type of each bidder. Setting these "lowest type" payoffs to 0 and plugging in expected payoffs for bidder types from the Envelope Theorem, we found that the seller's expected revenue was equal to the expected value of the winning bidder's "marginal revenue."

- So maximizing expected revenue boils down to maximizing the expected value of the winner's marginal revenue. With symmetric IPV, this can be accomplished with (among others) a second-price auction with a reserve price; with asymmetric IPV, the seller-optimal mechanism does not allocate the object efficiently.

- (There is an additional wrinkle: when the type distribution is not regular, we can't simply give the object to the bidder with the highest marginal revenue, since the chance of winning the object must be increasing in reported type for incentive compatibility to hold. So non-regular distributions require additional smoothing of the allocation rule.)

- We also saw that with correlated private values, the seller could force bidders to accept "side bets" about their opponents' types to punish bidders who underreported their own types, and that sellers could therefore allocate the object efficiently and extract all bidder surplus. Here, though, the seller-optimal mechanism requires very precise knowledge of the joint distribution of bidder types.

- We related the optimal auction problem to that of a monopolist selling a divisible good in multiple markets; and we proved the Bulow and Klemperer result that with symmetric bidders and $\underline{v} \geq v_0$, adding another bidder is more valuable than getting the reserve price right

- We looked at the effect of risk aversion. We saw that in a first-price auction, a risk-averse bidder, holding everything else constant, would bid higher than a risk-neutral bidder, since this transfers some payoffs from the "winning" cases to the "losing" cases, where the marginal utility of wealth is higher. And we used the ranking lemma to show that this result carries through to equilibrium bidding, and that a first-price auction therefore revenue-dominates a second-price auction when bidders are risk-averse. (Maskin and Riley give some characterization of the optimal auction under a more general formulation of risk-averse bidders.)

- We used stochastic dominance to show that risk-averse sellers also prefer first-price to second-price auctions.

- We returned to risk-neutrality and considered asymmetric bidders, in particular, bidders whose ex-ante distributions of possible valuations were ranked by stochastic dominance. We already mentioned the Myerson result that with asymmetric bidders, the seller-optimal auction was discriminatory. We proved the results (from Maskin and Riley) that in first-price auctions, weak bidders bid more aggressively than strong bidders with the same type; and that weak bidders prefer sealed-bid (first-price) auctions, while strong bidders prefer ascending (second-price) auctions, although either one can revenue-dominate.

- We introduced the more general framework of Milgrom and Weber, who allowed bidders' valuations to depend on their own and their competitors' types, and allowed for positive correlation among bidder (and seller) information. With affiliated signals, we proved several results collectively thought of as the linkage principle: that ascending auctions revenue-dominate sealed-bid auctions, that revealing the prices at which losing bidders dropped out increases expected revenue, and that revealing the seller's information also increases expected revenue. (In each case, the winning bidder's payment is linked to more variables which are correlated with his actual type, giving him less incentive to shade his bid.)

- So we found that first-price auctions yield higher revenue under risk-aversion, while ascending auctions give higher revenue with correlated signals. We considered a hybrid of the two auctions types, first proposed by Klemperer: an initial ascending auction, followed by a sealed-bid auction among the last few remaining bidders, called an Anglo-Dutch auction. We showed an example (by Levin and Ye) with both correlation and risk aversion, where an Anglo-Dutch auction revenue-dominated a pure ascending auction; but when risk aversion was sufficiently strong, it still underperformed the pure first-price auction.

- Next, we considered another special case: pure common values. We looked at the drainage tract model, credited to Wilson, of a single informed bidder and a bunch of uninformed bidders, and showed that the unusual theoretical predictions of the model seemed to be supported by empirical evidence.

- The drainage tract auction exhibited a funny symmetry: the distribution of the highest uninformed bidder matched the distribution of the informed bidder. We discussed a couple other asymmetric cases where this occurs: any two-bidder common-value auction with independent types, and asymmetric private-value auctions with post-auction resale.

- We looked at the problem of equilibrium multiplicity in second-price common-value auctions. And we looked at auctions with "almost common values," where a small private value advantage had an explosive effect on the equilibrium allocation and revenue.

- We looked at the question of information aggregation – whether the winning bid in a common-value auction approached the true value as the number of bidders gets large. We found that with a fixed number of prizes, information aggregation only occurs with extremely powerful signals; but that when the number of prizes grows along with the number of bidders, information does aggregate under reasonable conditions. And we mentioned a result that this conclusion is dependent on certainty (or near certainty) in the number of bidders, even as this number grows; nonvanishing uncertainty about the fraction of bidders who will win leads to failures in information aggregation.

- Next, we enlarged the scope of the problem to include more than just the auction itself.

- First, we considered endogenous information acquisition. We looked at Persico, who showed that first-price auctions are more risk-sensitive than second-price auctions, since better information about the value of the object also contains more information about your opponent's bids, which is relevant in first-price but not in second-price auctions. Persico therefore found that when information acquisition is unobservable and costly, first-price auctions lead to greater information acquisition than second-price (or ascending) auctions; which can potentially reverse the Milgrom-Weber revenue ranking result.

- We looked at a paper by Hernando-Veciana, who considered additively separable private plus common value components, and assumed that choices of how much to spend on gathering better information are observable. He found that when incremental information is about the common value piece, sealed-bid auctions lead to greater information acquisition than ascending auctions, since there is less ability to free-ride off of the information of a better-informed opponent. When incremental information is about the private value piece, the reverse is true; ascending auctions lead to greater information acquisition. He points out that both these results suggest that ascending auctions are more efficient.

- We looked at a paper by Larson, who looked at two-bidder second-prices with almost common values, and proved some cool results that were hard to interpret (and impossible to recap without getting into a lot of notation).

- And looking again at pure common-value auctions, we found that "new" information is more valuable when it is gathered publicly, while redundant information is more valuable when gathered in secret.

- Next, we looked at endogenizing the number of bidders in an auction, focusing on symmetric mixed-strategy equilibria in a first-stage entry game. With costly entry, entrants expect zero expected profit (net of entry costs), so the seller's problem coincides with choosing the socially optimal auction. We looked at Levin and Smith, who find that with common values, the seller gains by charging an entry fee to discourage entry, and gains by setting a positive reserve price when entry fees are impossible. With independent private values, the seller is better off charging no entry fee or reserve price. With affiliated private values, no entry fee or reserve price is optimal prior to an ascending auction, but a positive entry fee is optimal prior to a first-price auction. And a result from McAfee and McMillan: the seller can gain further by capping the number of bidders, rather than allowing unlimited entry.

- Next, we considered the problem of collusion among bidders. We looked at the optimal collusive mechanisms for weak cartels (who can coordinate bidding strategies) and strong cartels (who can also make side payments), and actions that sellers can take to discourage collusion: setting a higher reserve price; using sealed-bid auctions or, in ascending auctions, not allowing bidders to name their own bids; and making enforcement more difficult for the cartel by revealing less information after the auction and choosing a non-random tiebreaking rule.

- We looked at double-auctions – settings where both buyer and seller have private information and submit bids. In a one-buyer/one-seller setting, if both bids affect price, there is a huge multiplicity of equilibria; and these equilibria are generally inefficient (relative to the optimal mechanism, which is already inefficient relative to first-best by Myerson-Satterthwaite). On the other hand, if either the buyer's or seller's bid is used as the transaction price when trade occurs, there is generally a unique equilibrium, which is (second-best) efficient.

- We also looked at "big" double-auctions – common-value double-auctions where the number of buyers and sellers get large. Reny and Perry consider the problem with discrete strategy spaces; for a fixed grid and fixed number of buyers and sellers, even existence of a well-behaved equilibrium in these settings can be problematic, as best-responses need not be monotonic in types. However, Reny and Perry prove that as the grid of allowed bids gets sufficiently fine, and as the number of buyers and sellers get large, an equilibrium exists in the double-auction which approximates a rational-expectations equilibrium: up to an arbitrarily small ϵ, bidders bid their expected valuation conditional on being pivotal, information aggregates, and the outcome is efficient.

- Finally, in the last few lectures, we looked at a number of auctions for multiple objects.

- First, we considered sequential auctions. We saw nice results (some proved, some conjectured, in an early working paper by Milgrom and Weber) that with affiliated, interdependent values, sequential one-unit auctions should revenue-dominate a one-shot k-unit auction, and prices should drift upwards. Unfortunately, the latter was directly contradicted by at least one empirical observation.

- Next, we looked at uniform-price, highest-rejected-bid, sealed-bid auctions for multiple identical objects, and found that bidders shade their bids (or reduce their demand) on units beyond the first, leading to inefficient allocations.

- We looked at simultaneous ascending auctions, which can be used for either identical or heterogeneous items. We found that when bidders view the goods as substitutes, all sorts of good things happen: straightforward bidding is always feasible, and leads to a result approximating the competitive equilibrium. However, when some bidders

view items as complements, a competitive equilibrium may not even exist, and bidders face an exposure problem.

- We examined the Vickrey auction, a direct-revelation mechanism where truthful revelation is a dominant strategy. The Vickrey auction guarantees an efficient allocation, but is subject to a number of problems: truthful revelation disappears under budget constraints; computational intensity (bidders must evaluate and bid on 2^L packages); the existence and persistence of low-revenue equilibria; and, when items are not substitutes, sensitivity to a number of strategic manipulations (such as joint deviations by losing bidders and shill bidding).

- We considered first-price package auctions with full information, and examined "profit-target" equilibria, which coincide with the bidder-optimal points in the core. However, we said very little about first-price package auctions with private information.

- We considered the Ausubel-Milgrom ascending proxy auction, which they sell as the "best of both worlds": under full information, profit-target equilibria coincide with the bidder-optimal points in the core (like in Bernheim and Whinston); and when goods are substitutes, truthful bidding is an equilibrium even with private information, leading to an efficient outcome as in the Vickrey auction. However, we haven't been able to say much on the theoretical side about private information combined with complementarities.

- (Finally, for variety, we wrapped up with some analysis of wars of attrition and a generalized all-pay setting.)

So, where do we go from here?

- First of all, as we saw, the theory side of one-shot single-item auctions is pretty well developed, and not really a lively arena for new research. However, I think there is still some room for work on the strategic questions before the auction occurs – entry, information, collusion, the interaction of auctions with mergers, and so on. There's a clear view (especially in Klemperer, and in Milgrom as well) that the success of a real-world auction often hinges more on these variables – sufficient entry and genuine competition – than on fine-tuning the rules of the auction itself.

- Auctions do present some interesting empirical problems. The main "objective" in empirical estimation is typically to recover the distribution of bidder valuations from bid data. However, without further information, it is often impossible to empirically distinguish private from common values, or (as we'll see in two weeks) bidder collusion from competitive bidding; and allowing for correlation among bidder types adds a further complication.

- There's also some interesting experimental and behavioral work on auctions. I listed as a possible presentation topic a paper by Levin and Charness, which demonstrates that people (at least, undergrads) consistently fail to fully account for the winner's curse, even when all other strategic elements are stripped away.

- As we saw, the theory is underdeveloped on sequential auctions, and auctions for multiple items. I think the sequential auctions question is a very interesting one, since (for example) there are often dozens of auctions for identical objects on eBay going on at the same time. Understanding the dynamics here – how people choose which auction to bid in, how sellers decide when to start and end auctions, and so on – could be very interesting. (Next week, we'll see a paper on bidding in the final minute of an eBay auction. Given the prevalence of last-minute bids, it might be possible to interpret early bids as signals of intent to bid seriously in an auction, in an attempt to solve a coordination problem among bidders.)

- And of course, as we saw, multi-unit auctions, especially with complementarities, are not very well understood theoretically, and there's definitely room for developments here. Extending the usual models to this setting appears to be analytically intractable, so part of the problem is even coming up with simpler, tractable models that say something. In addition, with combinatorial bidding, strategic play by bidders is extremely complex; my own intuition is that it is therefore very sensitive to things we don't usually consider, like the user interface and how bidders are prompted/guided toward bids. Should be lots of room for experimental/behavioral stuff here, as well as for new theory.

- And finally, as I was trying to get at with wars of attrition, some of the tools we've been using – envelope theorem and so on – end up having applications in other areas as well.

So, I hope you all found the course interesting, and learned something. I'm looking forward to seeing your presentations the next few meetings.