CSCE 475/875 Multiagent Systems

Handout 22: Revenue Equivalence in Auctions

October 24, 2017

(Based on Shoham and Leyton-Brown 2011)

Revenue Equivalence

Of the large (in fact, infinite) space of auctions, which one should an auctioneer choose? To a certain degree, the choice does not matter, a result formalized by the following important theorem.

Theorem 11.1.4 (Revenue equivalence theorem) *Assume that each of n risk-neutral agents has an independent private valuation for a single good at auction, drawn from a common cumulative distribution $F(v)$ that is strictly increasing and atomless on $[v\_, \bar{v}]$. Then any efficient auction mechanism in which any agent with valuation $v\_$ has an expected utility of zero yields the same expected revenue, and hence results in any bidder with valuation v_i making the same expected payment.*

The theorem states that any allocation mechanism/auction in which

1. the bidder with the highest type/signal/valuation always wins
2. the bidder with the lowest possible type/signal/valuation expects zero surplus
3. all bidders are risk neutral, and
4. all bidders are drawn from a strictly increasing and atomless distribution will lead to the same expected revenue for the seller (and player *i* of type *v* can expect the same surplus across auction types). (*based on Wikipedia*)

Thus, when bidders are risk neutral and have independent private valuations, **English, Japanese, Dutch, and all sealed bid auction protocols are revenue equivalent**.

Risk Attitudes

One of the key assumptions of the revenue equivalence theorem is that agents are risk neutral. It turns out that many auctions *cease to be revenue-equivalent when agents' risk attitudes change.*

(***Note***: Risk averse agents prefer the sure thing; risk-neutral agents are indifferent; risk-seeking agents prefer to gamble.)

To illustrate how revenue equivalence breaks down when agents are not risk neutral, consider an auction environment involving n bidders with IPV valuations drawn uniformly from $[0, 1]$. Bidder i, having valuation v_i, must decide whether it would prefer to engage in a first-price auction or a second-price auction. Regardless of which auction it chooses (presuming that the bidder, along with the other bidders, follows the chosen auction's equilibrium strategy), i knows that it will gain positive utility only if it has the highest utility.

In the case of the first-price auction, i will always gain $\frac{1}{n} v_i$ when it has the highest valuation.

In the case of having the highest valuation in a second-price auction, i's *expected* gain will be $\frac{1}{n} v_i$, but because he or she will pay the second-highest actual bid, the amount of i's gain will *vary* based on the other bidders' valuations.

Thus, in choosing between the first-price and second-price auctions and conditioning on the belief that it will have the highest valuation, i is presented with the choice between ***a sure payment and a risky payment with the same expected value***.

If i is ***risk averse***, it will value the sure payment more highly than the risky payment, and hence will bid more aggressively in the first-price auction, causing it to yield the auctioneer a higher revenue than the second-price auction. (***Note*** that it is *i's behavior in the first-price auction that will change*: the second-price auction has the same dominant strategy regardless of i's risk attitude.)

If i is ***risk seeking*** it will bid *less* aggressively in the first-price auction, and the auctioneer will derive greater profit from holding a second-price auction.

(***Note***: Implications for an auctioneer?)

The strategic equivalence of Dutch and first-price auctions continues to hold under different risk attitudes; likewise, the (weaker) equivalence of Japanese, English, and second-price auctions continues to hold as long as bidders have IPV valuations. These conclusions are summarized in Table 11.1.

Risk-neutral, IPV		=		=		=		=	
Risk-averse, IPV	Japanese	=	English	=	2nd	<	1st	=	Dutch
Risk-seeking, IPV		=		=		>		=	

Table 11.1: Relationships between revenues of various single-good auction protocols. ('>' = more money for auctioneer)

Addendum: Why $\frac{1}{n} v_i$?

Proposition 11.1.2 *In a first-price auction with two risk-neutral bidders whose valuations are drawn independently and uniformly at random from the interval* $[0, 1]$, $(\frac{1}{2} v_1, \frac{1}{2} v_2)$ *is a Bayes–Nash equilibrium strategy profile.* (***Note***: That is: bidder 1's best response to bidder 2's strategy is $\frac{1}{2} v_1$.)

Theorem 11.1.3 *In a first-price sealed-bid auction with n risk-neutral agents whose valuations are independently drawn from a uniform distribution on the same bounded interval of the real numbers, the* ***unique symmetric equilibrium*** *is given by the strategy profile* $(\frac{n-1}{n} v_1, \ldots, \frac{n-1}{n} v_n)$. (In other words, the unique equilibrium of the auction occurs when each player bids $\frac{n-1}{n}$ of its valuation.)

Thus the gain is the utility (or valuation of the good) minus the amount paid for the good: $v_i - \frac{n-1}{n} v_i = \frac{1}{n} v_i$.