## CSCE 475/875 Multiagent Systems Handout 21: "You Can't Handle the Truth!"

October 24, 2017 (Based on Shoham and Leyton-Brown 2011)

#### **Independent Private Value (IPV)**

To analyze properties of the various auction protocols, let's first consider agents' valuations of goods: **their utilities for different allocations of the goods.** Auction theory distinguishes between a number of different settings here. One of the best-known and most extensively studied is the *independent private value (IPV*) setting.

In this setting all agents' valuations are drawn independently from the same (commonly known) distribution, and an agent's type (or "signal") consists only of its own valuation, giving itself no information about the valuations of the others.

An example where the IPV setting is appropriate is in auctions consisting of bidders with personal tastes who aim to buy a piece of art purely for their own enjoyment.

(*Note*: Also think about the *common-value* assumption. Resale value!)

### Second-price, Japanese, and English auctions

**Theorem 11.1.1** *In a second-price auction where bidders have independent private values, truth telling is a dominant strategy.* 

**Proof.** Assume that all bidders other than *i* bid in some arbitrary way, and consider *i*'s best response.

First, consider the case where i's valuation is larger than the highest of the other bidders' bids. In this case i would win and would pay the next-highest bid amount. Could i be better off by bidding dishonestly in this case?

- If *i* bid higher, *i* would still win and would still pay the same amount.
- If *i* bid lower, *i* would either still win and still pay the same amount or lose and pay zero.

Since *i* gets nonnegative utility for receiving the good at a price less than or equal to its valuation, *i* cannot gain, and would sometimes lose by bidding dishonestly in this case.

Now consider the other case, where *i*'s valuation is less than at least one other bidder's bid. In this case *i* would lose and pay zero.

- If *i* bid less, it would still lose and pay zero.
- If *i* bid more, either it would still lose and pay zero or it would win and pay more than its valuation, achieving negative utility.

Thus again, *i* cannot gain, and would sometimes lose by bidding dishonestly in this case.

In the IPV case, we can identify strong relationships between the second-price auction and Japanese and English auctions.

Consider first the comparison between second-price and Japanese auctions. In both cases the bidder must select a number (in the sealed-bid case the number is the one written down, and in

the Japanese case it is the price at which the agent will drop out); the bidder with highest amount wins, and pays the amount selected by the second-highest bidder.

The difference between the auctions is that *information about other agents' bid amounts is disclosed in the Japanese auction*. In the sealed-bid auction an agent's bid amount must be selected without knowing anything about the amounts selected by others, whereas in the Japanese auction the amount can be updated based on the prices at which lower bidders are observed to drop out. In general, this difference can be important—interdependent values, common values, etc.; however, it makes no difference in the IPV case.

# Thus, Japanese auctions are also dominant-strategy truthful when agents have independent private values.

#### **First-price and Dutch auctions**

Let us now consider first-price auctions.

The first observation we can make is that the Dutch auction and the first-price auction, while quite different in appearance, are actually the same auction (**in the technical jargon, they are** *strategically equivalent*).

In both auctions each agent must select an amount without knowing about the other agents' selections; the agent with the highest amount wins the auction, and must purchase the good for that amount. (*Note*: Strategic equivalence is a very strong property: it says the auctions are exactly the same no matter what risk attitudes the agents have, and no matter what valuation model describes their utility functions.)

This being the case, it is interesting to ask why both auction types are held in practice. One answer is that they make a trade-off between **time complexity** and **communication complexity**.

- First-price auctions require each bidder to send a message to the auctioneer, which could be unwieldy with a large number of bidders.
- First-price auctions can be done asynchronously; while Dutch auctions require everybody to be at the same place synchronously.
- Dutch auctions require only a single bit of information to be communicated to the auctioneer, but requires the auctioneer to broadcast prices.

Unlike the case of second-price auctions, here we do not have dominant strategies. In a firstprice auction, an agent's risk attitude also matters. For example, a risk-averse agent would be willing to sacrifice some expected utility (by increasing its bid over what a risk-neutral agent would bid), in order to increase its probability of winning the auction.