

Crowdsourced Bayesian Auctions

Jing Chen

We investigate the problem of optimal mechanism design, where an auctioneer wants to sell a set of goods to buyers, in order to maximize revenue. In a Bayesian setting the buyers' valuations for the goods are drawn from a prior distribution D , which is often assumed to be known by the seller. In this work, we focus on cases where the seller has no knowledge at all, and "the buyers know each other better than the seller knows them". In our model, D is not necessarily common knowledge. Instead, each player individually knows a posterior distribution associated with D . Since the seller relies on the buyers' knowledge to help him set a price, we call these types of auctions crowdsourced Bayesian auctions.

For this crowdsourced Bayesian model and many environments of interest, we show that, for arbitrary valuation distributions D (in particular, correlated ones), it is possible to design mechanisms matching to a significant extent the performance of the optimal dominant-strategy-truthful mechanisms where the seller knows D .

To obtain our results, we use two techniques: (1) proper scoring rules to elicit information from the bidders; and (2) a reverse version of the classical Bulow-Klemperer inequality. The first technique lets us build mechanisms with a unique equilibrium that guarantee good revenue, even when the players' second and higher-order beliefs about each other are wrong. The second allows us to upper bound the revenue of an optimal mechanism with n players by a $(1 - 1/n)$ fraction of the revenue of the optimal mechanism with $n-1$ players. We believe that both techniques are new to Bayesian optimal auctions and of independent interest for future work.

Joint work with Pablo Azar and Silvio Micali.