Mean Field Equilibria of Dynamic Auctions with Learning

KRISHNAMURTHY IYER and RAMESH JOHARI
Stanford University
and
MUKUND SUNDARARAJAN
Google Inc.

We study learning in a dynamic setting where identical copies of a good are sold over time through a sequence of second price auctions. Each agent in the market has an *unknown* independent private valuation which determines the distribution of the reward she obtains from the good; for example, in sponsored search settings, advertisers may initially be unsure of the value of a click. Though the induced dynamic game is complex, we simplify analysis of the market using an approximation methodology known as *mean field equilibrium* (MFE). The methodology assumes that agents optimize only with respect to long run average estimates of the distribution of other players' bids. We show a remarkable fact: in a mean field equilibrium, the agent has an optimal strategy where she bids truthfully according to a *conjoint valuation*. The conjoint valuation is the sum of her current expected valuation, together with an overbid amount that is exactly the expected marginal benefit to one additional observation about her true private valuation. Under mild conditions on the model, we show that an MFE exists, and that it is a good approximation to a *rational* agent's behavior as the number of agents increases. We conclude by discussing the implications of the auction format and design on the auctioneer's revenue. In particular, we establish a dynamic version of the revenue equivalence theorem, and discuss optimal selection of reserve prices in dynamic auctions.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: —*Economics*

General Terms: Economics, Theory

Additional Key Words and Phrases: Dynamic auctions, Learning, Mean field equilibrium

1. INTRODUCTION

Auctions are observed as a market mechanism in a wide range of economic transactions: sponsored search markets run by Google and Yahoo!, online marketplaces such as eBay and Amazon, crowdsourcing, procurement auctions for public service contracts, licensing auctions (e.g., for mining or oil tracts), etc. Nearly all of these examples are characterized by two important features. On one hand, the auction format is typically relatively straightforward to describe, consisting of repetitions of a simple one-shot auction format. On the other hand, despite the simplicity of the mechanism itself, such markets can give rise to complex dynamic incentives for bidders. As a result, many basic questions become quite challenging: determining optimal bidding strategies for bidders; characterizing dynamic equilibrium behavior among the bidders; and determining optimal choices of market parameters for the auctioneer, such as auction format and reserve prices.

Authors' addresses: `kriyer@stanford.edu`, `ramesh.johari@stanford.edu`, `mukunds@google.com`

As a concrete example, consider online sponsored search auctions [Edelman et al. 2007; Varian 2007]. These repeated auctions operate on a per keyword basis; in a typical scenario, advertisers bidding on a particular keyword estimate their underlying valuation based on the conversion rate from an ad click to revenue (e.g., from sales). As the advertisers win more ad placement through the auction, they *learn* this conversion rate, which informs their bidding decision in the auction [Ghose and Yang 2009; Rey and Kannan 2010; Sculley et al. 2009]. In these settings, the bidders face a trade-off between *exploration*, where they bid higher to obtain more information about their value, and *exploitation*, where they bid optimally given their current information.

The exploration-exploitation trade-off has significant ramifications for market operation and design. *First*, it complicates the design of optimal bidding strategies. In sponsored search markets, for example, identifying optimal learning strategies would lead to better design of bidding agents that incorporate advertisers' uncertainty about their conversion rate. *Second*, as the bidders are playing a complex dynamic game, it can be intractable to characterize equilibrium behavior among many interacting bidders, and in particular to determine how bidders' uncertainty affects the distribution of bids seen over time. *Third*, as a consequence, we lose the ability to guide market operation and design. For instance, auctioneers usually set reserve price in such markets to increase their revenue. As we later demonstrate, setting a reserve without incorporating the learning among the bidders may cause unwarranted restriction of allocation and ultimately yield *lower* revenue.

In our paper [Iyer et al. 2011], we study an abstract dynamic auction model that consists of repetitions of a simple one-shot mechanism. We primarily study a setting where identical copies of a good are sold through a sequence of second price auctions over time; as an example, a copy of the good may denote a click on an advertiser's ad in sponsored search settings. (We also analyze repetitions of other standard one-shot auction formats.) Each agent in the market has an independent private valuation that determines the distribution of the reward she obtains from the good; the private valuation may denote an advertiser's conversion rate in sponsored search. Although agents are initially unaware of their own private valuation, every time an agent wins an auction and obtains a copy of the good, her realized reward from the good incrementally informs her about her valuation. The strategic interactions among the agents along with their beliefs about their valuation influence their bids in the auction. Thus, we naturally obtain a dynamic game among the agents in our model.

The standard game-theoretic tool used to analyze such dynamic games is the equilibrium concept known as perfect Bayesian equilibrium (PBE). However, there are two central problems with this approach. First is that such equilibria are *intractable*: the state space complexity is enormous (since bidders must maintain beliefs over all that is unknown to them), and grows exponentially with the number of bidders and with time. Second, and partly in consequence, is that such equilibria are *implausible*: in equilibrium, a bidder's optimal bidding strategy is intricately predicated on what she believes *other* bidders' strategies are, and a PBE requires each bidder to accurately forecast and estimate *exactly* how her competitors will respond to any bid she makes today.

The complexity of PBE motivates us to consider an approximation methodology that we refer to as *mean field equilibrium* (MFE). (See [Adlakha et al. 2010; Huang et al. 2007; Jovanovic and Rosenthal 1988; Lasry and Lions 2007; Tembine et al. 2009].) MFE is inspired by a *large market* approximation: with a large number of bidders, tracking and forecasting the exact behavior of individual bidders is impractical and implausible. In an MFE, individuals take a simpler view of the world. They postulate that fluctuations in the empirical distribution of other agents' bids have "averaged out", and thus optimize holding the bid distribution of other agents fixed. MFE requires a consistency check: the postulated bid distribution must arise from the optimal strategies agents compute. The benefit of analyzing a large market using MFE is that for the agents to optimize their behavior, it is sufficient for them to just maintain beliefs about their *own* private valuation. This reduces the dimension of the system state that each agent needs to track, simplifying the analysis tremendously.

Furthermore, we believe that MFE corresponds more closely to an equilibrium concept that might be applicable in practice, particularly in settings with a large number of bidders. For example, in sponsored search auction markets, bidders generally do not have access to complete information about the bid history for auctions they participated in. Rather, bidders are usually provided with various tools by the auctioneer to aid in strategizing how to bid; e.g., Google provides the advertisers with a bid simulator that simulates how often an ad would get displayed and clicked upon on making a particular bid [Friedman 2009]. The bid simulator bases its predictions on aggregated historical data, that gives a "bid landscape" of competitors' bids on the same category or keyword (i.e., the distribution of bids). Bid landscapes inherently assume stationarity in the market, at least for a limited time horizon of interest; thus bidders are reacting to average information about their competitors. It is reasonable to expect that for many bidders, therefore, their own decision of how to bid will *not* explicitly forecast opponents' reactions, and instead will assume that these reactions have averaged out in any forecasting about future auction outcomes. This type of example illustrates how the rationality assumptions in MFE might naturally arise in practice.

Our main contributions address the challenges raised above.

(1) *Characterizing optimal strategies for bidders: Conjoint valuations.* We show that in the large market model, the optimal strategy of an agent takes a remarkably simple form: given her current belief about her valuation, the optimal strategy is to bid according to a *conjoint valuation.* The conjoint valuation is the sum of her current expected valuation, together with an overbid. This overbid denotes an agent's value for learning about her true private valuation, and we show that it is exactly the expected marginal benefit to one additional observation about her valuation. Thus the conjoint valuation presents a structurally simple and *plausible* strategy that captures how an agent in the large market balances the trade-off between exploration and exploitation.

(2) *Consistency and validity of the mean field model: Existence of MFE and an approximation theorem.* We show that the mean field model is consistent by proving the existence of an MFE. This involves showing that the stationary distribution of a market where each agent follows the mean field strategy turns

out to be the market distribution that each agent had assumed to solve their decision problem. We extend this result under mild conditions to a dynamic auction setting consisting of repetitions of a fixed *standard auction*; for example, standard auctions include second price, first price, and all-pay auctions. Thus, we obtain, in fairly general settings, the existence of informationally simple equilibria in mean field models, where agents make bids taking into consideration only their own belief about their valuation and the bid distribution in the market. This result provides evidence of the *tractability* of the mean field model.

We next tackle the issue of whether an MFE, which rests on a large market assumption, accurately captures a rational agent's behavior in a finite market. We prove that indeed an MFE is asymptotically a good approximation to agent behavior in a finite market. Formally, we show that if in a finite market, every agent except one follows the MFE strategy, then the remaining agent's loss on playing the MFE strategy converges to zero as the number of agents in the market increases. This result justifies formally the use of an MFE to analyze agent behavior in a finite market as the number of agents increases. We emphasize, however, that MFE may be a useful approximation even when the number of agents is not large, simply because it more accurately captures the information available to bidders when they optimize (e.g., in sponsored search auctions, bidders are responding to bid landscape information).

(3) *Market design: Auction format and reserve prices.* Finally, to illustrate the power of our approach, we leverage the analytical and computational simplicity of MFE to address "second best" market design: how should an auctioneer choose the auction format to maximize revenue, with the constraint of relatively "simple" repeated auction mechanisms?

In static settings, the revenue equivalence theorem states that an auctioneer's expected revenue in any standard auction remains the same. In a dynamic setting, the main difference is that now changing the auction format not only affects an agent's payment in each auction, but also affects her incentive to learn more about her valuation. Nevertheless, we prove a *dynamic revenue equivalence* theorem that extends the static version to dynamic settings. We show this by relating, under some conditions, an MFE of a market with repetitions of a given standard auction format to an MFE of a market with repeated second price auction. This result shows that changing the one-shot auction format will not increase the seller's expected revenue.

We then consider the possibility of increasing revenue by choosing a reserve price. In static auctions, setting a reserve has the effect of extracting greater revenues from high valuation bidders, at the expense of shutting out bidders with lower valuations [Myerson 1981]. In dynamic auctions with learning, however, a reserve has an added effect: it reduces bidders' incentives to learn their valuation. Ignoring this added effect while setting the reserve may cause the auctioneer to incur a high penalty. We develop benchmarks to evaluate this penalty, by comparing the MFE where an auctioneer anticipates bidders' learning behavior, against one where the auctioneer is oblivious to bidders' learning. The computational tractability of MFE allows us to evaluate these benchmarks:

we numerically observe that depending on the uncertainty bidders hold about their valuations, the incremental benefit to setting a reserve for an auctioneer could be as high as 15-30% more than the incremental benefit if the learning is ignored.

REFERENCES

Adlakha, S., Johari, R., and Weintraub, G. Y. 2010. Equilibria of dynamic games with many players: Existence, approximation, and market structure. *arXiv abs/1011.5537*.

Edelman, B., Ostrovsky, M., and Schwarz, M. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review 97*, 242–259(18).

Friedman, D. 2009. Bid like a pro with the bid simulator. `http://adwords.blogspot.com/2009/08/bid-like-pro-with-bid-simulator.html`.

Ghose, A. and Yang, S. 2009. An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Science 55,* 10, 1605–1622.

Huang, M., Caines, P. E., and Malhamé, R. P. 2007. Large-population cost-coupled LQG problems with nonuniform agents: Individual-mass behavior and decentralized ϵ-Nash equilibria. *IEEE Transactions on Automatic Control 52,* 9, 1560–1571.

Iyer, K., Johari, R., and Sundararajan, M. 2011. Mean field equilibria of dynamic auctions with learning. *SSRN eLibrary* 1799085.

Jovanovic, B. and Rosenthal, R. W. 1988. Anonymous sequential games. *Journal of Mathematical Economics 17*, 77–87.

Lasry, J. M. and Lions, P. L. 2007. Mean field games. *Japanese Journal of Mathematics 2*, 229–260.

Myerson, R. B. 1981. Optimal auction design. *Mathematics of Operations Research 6,* 1, 58–73.

Rey, B. and Kannan, A. 2010. Conversion rate based bid adjustment for sponsored search. In *Proceedings of WWW*. 1173–1174.

Sculley, D., Malkin, R. G., Basu, S., and Bayardo, R. J. 2009. Predicting bounce rates in sponsored search advertisements. In *Proceedings of ACM SIGKDD*. 1325–1334.

Tembine, H., Boudec, J.-Y. L., El-Azouzi, R., and Altman, E. 2009. Mean field asymptotics of Markov decision evolutionary games and teams. In *Proceedings of GameNets '09*. 140–150.

Varian, H. R. 2007. Position auctions. *International Journal of Industrial Organization 25,* 6, 1163–1178.