# Algorithmic Game Theory and Econometrics

VASILIS SYRGKANIS
Microsoft Research, NYC

The traditional econometrics approach for inferring properties of strategic interactions that are not fully observable in the data, heavily relies on the assumption that the observed strategic behavior has settled at an equilibrium. This assumption is not robust in complex economic environments such as online markets where players are typically unaware of all the parameters of the game in which they are participating, but rather only learn their utility after taking an action. Behavioral models from online learning theory have recently emerged as an attractive alternative to the equilibrium assumption and have been extensively analyzed from a theoretical standpoint in the algorithmic game theory literature over the past decade. In this letter we survey two recent works, [Nekipelov et al. 2015, Hoy et al. 2015], in which we take a learning agent approach to econometrics, i.e. infer properties of the game, such as private valuations or efficiency of observed allocation, by only assuming that the observed repeated behavior is the outcome of a no-regret learning algorithm, rather than a static equilibrium. In both works we apply our methods to datasets from Microsoft's sponsored search auction system.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics; K.4.4 [**Computers and Society**]: Electronic Commerce; F.1.2 [**Modes of Computation**]: Online computation

General Terms: Theory, Economics, Econometrics, Experimental

Additional Key Words and Phrases: Econometrics, no-regret learning, sponsored search, value inference, welfare guarantees

## 1. INTRODUCTION

One of the main goals of the econometric analysis of strategic interactions is the inference of the private parameters of participants based solely on their observed actions. For instance, by observing a sequence of bids of a set of bidders participating repeatedly in an auction for a single item, one aims to infer the private value each player has for the item. Another quantity of interest in such environments is the efficiency of the outcome of the strategic interaction, i.e. was the item sold to the player with the highest or approximately highest valuation.

Any such task requires an assumption on how the players make decisions in a repeated game setting. One of the main assumptions that has been overwhelmingly used in traditional econometrics is that the actions that we observe in the data are the product of a Nash equilibrium behavior of the participants, i.e. a state of mutual best-responses [Athey and Nekipelov 2010; Bajari et al. 2013]. Such an assumption is rather strong, especially in complex environments such as online sponsored search auctions, where the players do not even know who they are competing against and do not even know all the parameters of the auction rule.

In such settings, players typically only observe periodic aggregate feedback of what their utility would have been for any possible action they could have taken in the last period. Therefore, models of strategic behavior should be better suited

Author's address: vasy@microsoft.com

to such feedback structures and allow for bounded rationality of the players. One such model of behavior that has been proposed in the game theory literature [Foster and Vohra 1997; Freund and Schapire 1999] and which has been extensively analyzed in the algorithmic game theory literature in the past decade is that of no-regret learning [Blum et al. 2008; Roughgarden 2009; Syrgkanis and Tardos 2013]. No-regret learning simply assumes that players use a learning algorithm which, over-time, guarantees them that their utility is at least as good as the best fixed action in hindsight. Thus unlike the Nash equilibrium assumption, no-regret learning allows for dynamic player behavior and requires only an approximate best-response property and only on average over a time period. Moreover, there exist many learning algorithms that achieve this property and which work even in the aforementioned utility feedback model.

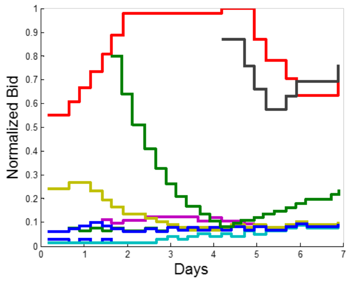


Fig. 1. Normalized bid for listings of a single advertiser over the course of a week.

Dynamic behavioral models, such as learning agent models, seem of practical importance, since in many real sponsored search datasets we observe bidders changing their bids very frequently. For instance, in Figure 1 we depict the bids of a subset of the listings of a single advertiser in Microsoft's sponsored search auction system over the period of a week.

## 2. VALUE INFERENCE FOR LEARNING AGENTS

In [Nekipelov et al. 2015] we address the problem of inferring player valuations from a sequence of bid observations in a repeated sponsored search auction environment. We propose an approach that solely assumes that the sequence of bids is the outcome of a vanishing regret learning algorithm.

In the setting that we analyzed, advertisers submit a bid for being allocated a position in the sponsored section of a search page. Advertisers are allocated positions based on some quality score and their bid. When an advertiser is clicked we assume that she receives some value $v$ which is private and known only to her. This per-click valuation is the parameter that we want to infer. We also assume that the utility of an advertiser is quasi-linear in money, i.e., her utility is her value minus her payment.

Assuming that the sequence of bids of an advertiser is an $\epsilon$-regret sequence implies that the utility that she derived over the entire period that we observe must be at least as high as what any fixed bid would have achieved less some $\epsilon$. This condition gives a set of inequalities that the value of a player must satisfy, one inequality per fixed bid. The intersection of these inequalities, is the set of values that are *rationalizable* under the assumption of $\epsilon$-regret. Varying $\epsilon$, we get a set of pairs $(v, \epsilon)$, such that value $v$ is rationalizable under the $\epsilon$-regret assumption. We refer to this set as the *rationalizable set*. We show that the rationalizable set is convex and characterize its statistical learning properties. We show that the statistical learning rate of the rationalizable set is remarkably comparable with the statistical learning rates of methods that make the stronger equilibrium behavior assumption [Athey and Nekipelov 2010].

If one wants to make a point prediction on the value of a player, then a selection rule is needed, to select among the points in the rationalizable set. We analyze the point that corresponds to the smallest multiplicative regret (the sequence has multiplicative regret $\lambda$ if the current utility of the bidder is at least $(1-\lambda)$ times the utility of any fixed bid). We apply this point-prediction approach to a dataset from Microsoft's sponsored search system. Figure 2 depicts the results of our analysis when applied to all the listings of a single account. For the inferred values, we depict the distribution of how much a player shades his value on average and the distribution of the smallest rationalizable error across listings. We find that on average for many accounts, advertisers bid around 60% of their inferred valuation and that the smallest rationalizable error, though small, is bounded away from zero for almost 70% of the listings (i.e. doesn't satisfy the exact best response property).

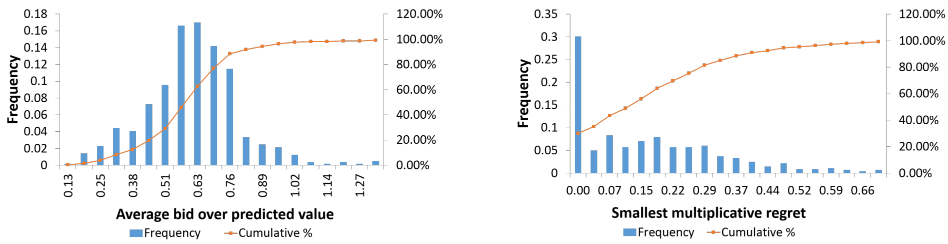


Fig. 2. Distribution of bid shade ratio and smallest multiplicative regret across listings of a single advertiser.

## 3. DATA-DRIVEN ROBUST EFFICIENCY GUARANTEES

In [Hoy et al. 2015], we give an econometric approach for directly inferring a lower bound on the efficiency of the resulting allocation in a repeated auction setting, without even inferring first the valuations of the players. Our approach is an empirical analogue of the smoothness approach on quantifying the worst-case inefficiency in games [Roughgarden 2009; Syrgkanis and Tardos 2013; Hartline et al. 2014] and therefore inherits several robustness properties of smoothness. For instance, the lower bound on the efficiency that is derived via our method holds regardless of whether the data that we observe are the product of a Bayes-Nash equilibrium where the player valuations are stochastic, or whether they are the product of a learning process employed by an advertiser with a fixed valuation. Moreover, our method enjoys fast statistical learning rates when only a sub-sample of the strategic interactions is observed.

The smoothness approach of [Syrgkanis and Tardos 2013] and its refinement for single-parameter mechanism design environments, via the revenue and value covering formulation of [Hartline et al. 2014] is based on the following argument: at any outcome of the game that satisfies a best-response or approximate best-response property either the player is getting high utility and hence high allocation probability, or the payment that he needs to make in order to achieve a high allocation probability given the competition must be high. The latter quantity is typically referred to as the threshold payment.[1] Subsequently, if this threshold payment is closely related to the revenue that the auction receives then we can attribute this term to the current welfare of some other bidder. Combining these two arguments gives a lower bound on the efficiency of the allocation.

The crucial observation in [Hoy et al. 2015] is that both the threshold payment quantity and the revenue are observed in the data. Thereby we do not need to theoretically prove a relation between the two quantities. One simply needs to analyze the relation of the two quantities from the data. This can potentially lead to better efficiency guarantees than the theoretically provable ones. More importantly, our approach can be used to infer efficiency lower bounds even in auctions where no worst-case theoretical relation is known between the two quantities and therefore no worst-case efficiency lower bound can be inferred simply from the rules of the auction without observing the data. This is the case with the actual complex sponsored search auction that is being used in Microsoft's sponsored search system. For instance, the application of our approach to real datasets for a selection of high-revenue keywords yielded significant efficiency guarantees, ranging from 30% to 70% of the efficiency of the optimal allocation.

## REFERENCES

Athey, S. and Nekipelov, D. 2010. A structural model of sponsored search advertising auctions. In *Sixth ad auctions workshop*.

Bajari, P., Hong, H., and Nekipelov, D. 2013. Game theory and econometrics: a survey of some recent results. *Advances in Economics and Econometrics 3*.

Blum, A., Hajiaghayi, M., Ligett, K., and Roth, A. 2008. Regret minimization and the price of total anarchy. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*. STOC '08. ACM, New York, NY, USA, 373–382.

Foster, D. P. and Vohra, R. V. 1997. Calibrated learning and correlated equilibrium. *Games and Economic Behavior 21,* 12, 40 – 55.

Freund, Y. and Schapire, R. E. 1999. Adaptive game playing using multiplicative weights. *Games and Economic Behavior 29,* 1, 79–103.

Hartline, J., Hoy, D., and Taggart, S. 2014. Price of Anarchy for Auction Revenue. In *ACM Conference on Economics and Computation*. ACM Press, New York, New York, USA, 693–710.

Hoy, D., Nekipelov, D., and Syrgkanis, V. 2015. Robust data-driven efficiency guarantees in auctions. *1st Workshop on Algorithmic Game Theory and Data Science and CoRR abs/1505.00437*.

Nekipelov, D., Syrgkanis, V., and Tardos, E. 2015. Econometrics for learning agents. In *Proceedings of the 16th ACM Conference on Economics and Computation*. EC'15.

Roughgarden, T. 2009. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st annual ACM symposium on Theory of computing*. STOC '09. ACM, New York, NY, USA, 513–522.

Syrgkanis, V. and Tardos, E. 2013. Composable and efficient mechanisms. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*. STOC '13. ACM, New York, NY, USA, 211–220.

[1]The actual quantity that goes into the formulation in order to produce tight efficiency results is slightly more involved and the reader is referred to the paper for a full exposition.