

**Econometrics is a key component to gauging user satisfaction and advertisers' profits.**

BY DENIS NEKIPELOV AND TAMMY WANG

# Inference and Auction Design in Online Advertising

THE TRANSITION OF the largest online advertising platforms to auction-based automated real-time designs has transformed the advertising industry. The advertisers had an opportunity to design flexible goal-specific advertising campaigns, target focused small groups of consumers, and perform fast and efficient experimentations. At the same time, consumers can be exposed to a smaller number of higher quality advertisements. However, in order to achieve the design goals, such as optimal placement of ads that maximizes the sum of utilities of users or revenue generated by the auction platform, the auction needs to be carefully optimized. To optimize the auction, the auctioneer typically chooses a relatively small number of auction parameters that determine the allocation of

items and prices based on submitted bids. Typical auction parameters include reserve prices, which determine the minimum bid and can correspond to any allocation. Another parameter used in sponsored search auctions is the quality score, which is a positive bidder-specific weight used to discount or inflate submitted bids before they are ranked. The setting of auction parameters requires a knowledge of advertisers' preferences and consumers behavior, which can be acquired from data. This makes econometric inference from observed data of high importance for the design and analysis of online advertising auctions.

The structure of online advertising exchanges is becoming significantly complex, and requires multiple parameters to be input by auction designers. These parameters are required to yield consistent advertising allocations, relevant user ad experiences over time and catering ad placements to advertisers' goals. This operational structure of online ad exchanges has been incorporated into the algorithmic implementation of advertising auctions, see Muthukrishnan,<sup>29</sup> Aggarwal and Muthukrishnan,<sup>30</sup> and Muthukrishnan<sup>31</sup> for details of such an implementation. However, the increased heterogeneity and dynamics of the marketplace call for quick "on demand" adjustments of auction parameters in order to pursue auction platform revenue goals and relevant ad experiences for users. Recent evidence from the sponsored search advertising

## » key insights

- **Modern advertising platforms are automated systems powered by auctions that allocate and price advertising slots contemporaneously based on the bids submitted.**
- **Real-time auction-based systems require good predictions of user behavior, need to adapt to the changing marketplace, and require fast and robust methods for evaluating performance.**
- **These challenges can be addressed by applying econometric methods to the data from advertising platforms.**



marketplace indicates that there are significant gains (both in terms of overall welfare and ad exchange revenue) from adopting data-driven designs. For example, Ostrovsky and Schwarz<sup>34</sup> described the experiment on Yahoo!’s search engine where the advertisers’ per-click values were recovered from observed bids. The distribution of recovered values was then used to set search query-specific reserve prices in the advertising auctions. The data shows a significant increase in the search engine’s revenue resulting from this switch.

As with other economic markets, advertising markets balance supply and demand. The webpage content attracts consumers and supplies their “attention.” Advertisers demand user attention since it converts advertised products and services to purchases. As in all economic markets, the equilibrium of demand and supply is driven by price. The dominant pricing model for advertising puts payment control on the side of advertisers and lets them determine their preferred prices. At the same time, consumers receive a “free” service. This feature has made advertiser behavior of direct importance for monetization and, therefore, has been studied in detail in recent years.

We survey the work devoted to the analysis of advertiser behavior, starting from market equilibrium models to more recent work where advertisers must adopt learning strategies to set their bids. We then discuss the “supply” side of the market, which to a large extent remains model-free. Recent empirical work has focused on determining actual advertisers’ value extracted from users. This work turned out to be significantly more delicate since user behavior is highly correlated with user characteristics, and thus, it is very difficult to distinguish the effect of advertising from the inherent propensity of particular users to purchase particular products. Finally, we describe recent changes in user behavior caused by the transformation of browsing and purchasing patterns generated by new devices, and new ways of accessing the Internet. We predict this trend will generate the need to do more in-depth consumer-side research of advertising markets.

The complexity of advertising markets highlights the importance of inference, as inference informs customized market-specific design and optimization of advertising markets. Here, we review existing and emerging approaches to inference for advertising auctions both aimed at the diagnostic analysis of advertising marketplaces and as an input to decisions regarding marketplace changes.

### Platform Design and Advertiser Behavior

*Bidder value inference.* While the general auction platform design can be based on pure economic theory, platform optimization requires the knowledge of the bidders’ underlying preferences and the prediction of user action. This knowledge is derived from the auction outcomes data. The general problem of data-driven auction design can be formulated as the *counterfactual inference*—from observations of agent behavior in the platform, we want to predict how they will behave if the platform design is changed. With the new design, participating agents will re-optimize their strategies, and, thus, this new design leads to the change in bids. Counterfactual inference, therefore, requires the platform to anticipate strategic responses of bidders to the mechanism change.

The counterfactual inference is based on the economic model that characterizes agent behavior on the platform. Based on this model, one can use data to infer the “primitives” of the model (such as agent preferences) that are consistent with the observed behavior. The prediction for the new design will be based on computing bidders’ responses to the new auction mechanism, and the changed bids of the other bidders.

In an advertising auction model, the utility of a bidder (that characterizes the expected gain of the bidder from participating in an auction) is typically set equal to the click probability multiplied the bidder’s value per-click less the cost. The click probability depends on both the identity of the bidder and the placement of the ad. Therefore, it ultimately depends on the bid that determines where the ad is placed. In this model, we characterize each bidder by a single parameter that

expresses this bidder’s value per click. Inference of bidders’ values per click is, therefore, one of the key ingredients of counterfactual inference. Another ingredient is the computation of the *Bayes-Nash* equilibrium of the auction.

The Bayes-Nash equilibrium of an auction is the set of bids and beliefs of bidders regarding the distribution of uncertainty in the auction such that the bids optimize bidders’ expected utilities, and the beliefs of bidders correctly capture the realized distribution of uncertainty. When bidders use Bayes-Nash equilibrium strategies and the equilibrium is unique, the values can be inferred directly from data by inverting bidders’ best response functions. The inferred bidder value makes the observed bid the best response to the observable empirical distribution of opponent bids. Athey and Nekipelov<sup>3</sup> adopted this strategy to infer bidders’ values in sponsored search auctions. Sponsored links that appear beside Internet search results on major search engines are sold using real-time auctions. Advertisers place standing bids where they are associated with search phrases that form part or all of a user’s search query. Each time a user enters a search query, applicable bids are entered to an auction. The ranking of advertisements and the prices paid depend on advertiser bids as well as “quality scores” assigned for each advertisement and user query. These quality scores are traditionally set equal to the predicted probability of a user click on an ad. They vary over time because the statistical algorithms predicting quality produce different predictions based on the features of the user issuing the search query.

In theory, the allocation and payment functions in standalone position auctions are discontinuous with respect to bid, because the price only changes when the bidder goes up in the ranking with an increased bid. However, in ad auctions each bid is typically applied to many auctions. With sufficient uncertainty over auctions, the expected auction outcomes will be continuous. As a result, to accurately characterize the behavior of bidders in this setting, Athey and Nekipelov<sup>3</sup> propose a model that has these properties.

The model proposed by Athey and

Nekipelov<sup>3</sup> deviates from the standard model of a Bayesian auction game where the values of the players are drawn from the distribution. The main reason for this deviation is that in real-world advertising auctions the sets of bidders are fixed and most bidders' bids do not change over long time intervals. Athey and Nekipelov's model rationalizes this observation. In the rest of the survey we adhere to the model of Athey and Nekipelov and assume the values of the bidders are fixed.

Consider the model where there are  $I$  bidders whose ads are eligible to be displayed to a user alongside the search results and to be placed in  $J$  available advertising positions. The user clicks on the ad  $i$  shown in the position  $j$  with probability  $c_{ij}$ , where  $c_{ij} = \alpha_j \cdot s_i$  is the product of the ad-specific click probability  $s_i$  (for  $i = 1, \dots, I$ ) and the advertising slot-specific click probability  $\alpha_j$  (for  $j = 1, \dots, J$ ). The bidder-specific click probability  $s_i$  for the particular bidder  $i$  is a random variable whose realizations vary across search queries of user. The randomness of the bidder-specific click probability reflects different propensities of different users to click on a given ad.<sup>a</sup> The bidder-specific click probability  $s_i$  is referred to as the bidder's *quality score* (or, simply, the score).

This description shows how an ad platform conducts a score-weighted auction. In each user query, the platform runs an auction where each advertisement  $i$  is assigned the score  $s_i$ , and bids are ranked in order of the product  $b_i \cdot s_i$ . The per-click price  $p_i$  that bidder  $i$  in position  $j$  pays is determined as the minimum price such that the bidder remains in her position,

$$p_i = \min \{b_i \cdot s_i b_i \geq s_k b_k\} = s_k b_k / s_i, \quad (1)$$

where bidder  $k$  is in position  $j + 1$ . Note that advertiser  $i$ 's bid does not directly influence the final price she gets charged, except when it causes to change positions. In effect, an advertiser's choice of bid determines which position she attains, where the price per click for each position is exogenous

<sup>a</sup> In practice,  $s_i$  is a prediction from a machine learning algorithm run by the search engine that uses a large set of user features which vary from query to query.

**Advertisers demand user attention since it converts advertised products and services to purchases.**

to the bid and rises with position. An auction with this type of pricing rule is called the (score-weighted) *generalized second price auction*.

Per-click values of bidders  $v_i$  for  $i = 1, \dots, I$  are fixed for each bidder across user queries and are supported on  $[0, \bar{v}]$ . The pay-off of bidder  $i$  in a query where this bidder receives a click is the surplus  $v_i - p_i$ , where  $p_i$  is bidder  $i$ 's price per click defined by Equation (1).

While we assume all  $I$  bidders are eligible to be displayed in each user query, the actual sets of participating bidders in a real search query page will vary. For instance, some bidders may be excluded from certain queries based on the targeting settings (that is, advertisers may prefer to advertise in specific geographic locations). Suppose each bidder knows the entire set of bidders  $I$ , but can not observe the set of competitors in a given query and can not observe neither her own nor her competitors' scores. In this case, bidders form beliefs regarding the distribution of scores of all bidders (since they affect the prices in individual user queries), and beliefs regarding the distribution of realizations of sets of their competitors across user queries. The standard assumption in the sponsored search auction literature (for example, Edelman et al.<sup>16</sup> and Varian<sup>41</sup>) is that bidders have full information regarding each other's bids. This reflects the idea of a high frequency environment where bidders can quickly determine the bids of their opponents, even if the auction platform does not explicitly provide that information. Bidder  $i$  then maximizes the expected pay-off (with per click value) with the bidder's beliefs regarding the distribution of uncertainty of scores and sets of competitors. This pay-off (a.k.a. utility) can be expressed as,

$$U(b_i, b_{-i}; v_i) = E[c_{ij}(v_i - p_i)], \quad (2)$$

where the expectation  $E$  is taken over the distribution of scores  $s_1, \dots, s_I$  and the distribution of sets of bidders participating in the auction. We emphasize that the pay-off depends on the bid of the bidder  $i$  ( $b_i$ ) as well as the bids of all competing bidders ( $b_{-i}$ ) that may determine the rank and the price that bidder  $i$  pays depending on a realization of the set of scores.

Here, we consider the question of bidders' per-click values that can be recovered from data that contains a large number of queries for a given set of potential bidders. For each query, the data available to the advertising platform contains the bids, a set of entrants, and the bidders' scores relating to each user query.

In this case, we can define function  $Q_i(\cdot)$  to be equal to the probability of click on the ad of bidder  $i$  in a search query as a function of all bids. We also define function  $TE_i(\cdot)$  equal to the expected payment of bidder  $i$  in a search query as a function of all bids. The utility of bidder  $i$  can be written in terms of these functions as,

$$U(b_i, b_{-i}; v_i) = v_i Q_i(b_i, b_{-i}) - TE_i(b_i, b_{-i}).$$

The best-responding bid maximizes the expected utility given the bids of opponent bidders ( $b_{-i}$ ). Athey and Nekipelov<sup>3</sup> demonstrate that with sufficient smoothness and support size of the distribution of scores, functions  $Q_i(\cdot)$  and  $TE_i(\cdot)$  are strictly increasing and differentiable with respect to bids. As a result, the value per click for each bidder can be recovered using the necessary first order condition for optimality of the bid.

$$v_i = \frac{\partial TE_i(b_i, b_{-i})}{\partial b_i} \bigg/ \frac{\partial Q_i(b_i, b_{-i})}{\partial b_i}. \quad (3)$$

Functions  $Q_i(\cdot)$  and  $TE_i(\cdot)$  can be recovered directly from the data<sup>b</sup> (for example, using splines or any other appropriate method). Thus, the per-click value can also be recovered from a data using Equation (3). Smoothness, and monotonicity of  $Q_i(\cdot)$  and  $TE_i(\cdot)$  can be directly tested in data and has been verified using the auction data by Athey and Nekipelov.<sup>3</sup>

Equation (3) provides a simple practical method for estimating value per click. For each bidder we change bid

<sup>b</sup> Statistical properties of the corresponding estimators, such as the rate of convergence, are determined by either the complexity of the circuit that is used to compute the allocations and payments from the bids or using their functional properties such as smoothness or concavity directly. The methods that can be used to establish these statistical properties parallel those applied to derive sample complexity of mechanisms such as Cole and Roughgarden,<sup>14</sup> Cummings,<sup>15</sup> and Morgenstern and Roughgarden.<sup>28</sup>

## An advertiser's choice of bid determines which position she attains, where the price per click for each position is exogenous to the bid and rises with position.

by a small amount and then compute the change in the outcome of the auctions where the original bid of the bidder was applied. The evaluated change in expenditure will serve as an estimator for the derivative of  $TE_i(\cdot)$  function and the evaluated change in number of clicks will serve as an estimator for the derivative of  $Q_i(\cdot)$  function.

*Evaluation of new auction mechanisms using data.* The approach adopted in the empirical economics literature (see Bajari et al.<sup>4,5</sup>) provides a simple recipe for platform optimization subject to the inferred bidder preferences, assuming the bids constitute the Nash equilibrium (defined previously in terms of best responses and beliefs of bidders) under the new settings. To follow this recipe, first we estimate the components of the first-order condition (Equation 3) using data from historical realizations of auctions. Value per click can then be reverse engineered from the bids for each bidder. Second, we plug in the new proposed auction parameters (or consider a new auction format). For each set of bids we can simulate the auction outcomes under the new auction parameters. These new auction outcomes (prices and allocations for hypothetical bids under the new auction rules) will generate *counterfactual* functions  $\widetilde{TE}_i(b_i, b_{-i})$  and  $\widetilde{Q}_i(b_i, b_{-i})$  for each bid profile. Finally, assuming the system converges to the long-term Nash equilibrium under new auction rules, and the values of the bidders do not change, we can predict the new bids. To do so, we find the set of bids for all participating bidders such that the bid of each bidder maximizes her utility under the new auction rules given the bids of other bidders. Since the maximum utility will be attained at the point where its derivative with respect to the bid is equal to zero, finding the new equilibrium bids is reduced to solving the system of nonlinear equations for  $i = 1, \dots, I$

$$\frac{\partial \widetilde{TE}_i(b_i^*, b_{-i}^*)}{\partial b_i} \bigg/ \frac{\partial \widetilde{Q}_i(b_i^*, b_{-i}^*)}{\partial b_i} = v_i, \quad (4)$$

using values  $v_i$  inferred from the data (using Equation (3)) and solve for the set of bids  $b_1^*, \dots, b_I^*$  that makes all  $I$  equalities hold. Unlike Equation (3) where we use actual bids  $b_i$  to obtain the values, in Equation (4) bids  $b_i^*$  are the

outcome and the values are the input. We note that this strategy only works if functions  $TE_i(\cdot, b_{-i})$ ,  $Q_i(\cdot, b_{-i})$ ,  $\widetilde{TE}_i(\cdot, b_{-i}^*)$  and  $\widetilde{Q}_i(\cdot, b_{-i}^*)$  are strictly monotone and the equilibrium is unique.

Monotonicity of payment and allocation functions along with the uniqueness of the Nash equilibrium allow us to use numerical continuation methods to find the new equilibria. The numerical continuation approach to finding this new equilibrium is based on transforming the set of conditions Equation (4) into system of ordinary differential equations. Athey and Nekipelov<sup>3</sup> show that solving this system of differential equations is equivalent to finding the new equilibrium (that is, it will not diverge or generate new extraneous solutions). The solution of numerical approximation of ordinary differential equations will yield the new equilibrium.

Although formulating the equilibrium computation problem in terms of solving the system of differential equations is mathematically elegant, the process can be complicated since it requires many calls to functions  $\widetilde{TE}_i(\cdot, b_{-i}^*)$  and  $\widetilde{Q}_i(\cdot, b_{-i}^*)$ , which need to be computed from the data for each evaluation. This problem is further accentuated in most advertising markets where the large number of eligible bidders participating in each auction leads to difficulties in computing the new equilibrium, even when sufficient conditions hold for existence and uniqueness of the Nash equilibrium.

A significant simplification in equilibrium settings can be achieved when the object of interest is not the entire vector of bids, but a specific lower dimensional auction outcome, such as revenue or welfare. When a simple outcome is defined, the entire vector of counterfactual equilibrium bids is no longer of direct interest, which effectively reduces the scale of the computational problem. Next, we consider this approach in application to social welfare in an auction.

*The price of anarchy bounds.* While equilibrium-based analysis can be widely applied, it is largely based on the idea that bidders use equilibrium strategies and the corresponding set of bids constitutes a Nash equilibrium. A common assumption used in the economics literature is the equilibrium is unique in a given market. In practice, however, it is impossible to

guarantee that sufficient conditions for Nash equilibrium uniqueness are satisfied in a given auction mechanism. For instance, the non-monotonicity of either the payment or allocation function leads to several possible values for each click that are consistent with observed bids. Then the simple equilibrium-based analysis is no longer valid.

The economics literature has developed an approach to the analysis of data from outcomes of games with multiple equilibria and games in non-equilibrium settings (for example, see Aradillas-Lopez et al.<sup>1</sup> and Beresteanu et al.<sup>6</sup>). The approach is to directly consider the optimality conditions for the players in these new settings. In auctions such optimality conditions lead to sets of possible values of bidders referred to as *identified sets*. This approach, however, has proven to be very difficult with reported computational time significantly exceeding those in the equilibrium framework even for simple games (see Ciliberto and Tamer<sup>13</sup>).

Koutsoupias and Papadimitriou<sup>23</sup> have developed a theoretical framework for analyzing simple auction outcomes such as revenue and welfare. The framework produces bounds for these outcomes comparing them to the welfare-maximizing allocation over all feasible realizations of bidders' values. This approach is referred to as *the price of anarchy* approach, which is defined as the ratio of the worst-case objective function value of the outcome of a game (such as Nash equilibrium outcome) and that of an optimal outcome. It turns out the price of anarchy can be established under minimal assumptions on the underlying preferences for many commonly used auction mechanisms (for instance, see Roughgarden<sup>36</sup>).

There are two difficulties with practical application of the price of anarchy bounds for a given game. First, these bounds can be quite conservative. Given they are derived taking the worst case equilibrium and the value profile, these bounds can be large. However, the values of players and the equilibria that generate the worst-case outcomes may not occur in reality. Even when the data from the mechanism is available, the price of anarchy bounds do not take it into account. Second, the price of anarchy is mechanism-specific and has so far been derived on the case-by-case basis.

Hoy et al.<sup>20</sup> bridge the gap between the robust but coarse theoretical price of anarchy bounds and precise but difficult to compute identified sets. The idea is to integrate data directly into the price of anarchy style analysis that can be applied to large classes of existing auction mechanisms. This new concept is called the *empirical price of anarchy*.

Here, we discuss the idea behind the empirical price of anarchy for auction settings where bidders adhere to playing full information equilibrium best responses in the auction mechanism  $A$ , but where those best responses are not unique (for example, due to non-monotonicity of allocations and prices in the auction). We assume, like in the model of Athey and Nekipelov,<sup>3</sup> the auction parameters are random while the profile of bidders' values is fixed and we use  $\mathbf{V} = (v_1, \dots, v_r)$  to denote that profile. For a given profile of values auction  $A$  can have multiple Nash equilibria that constitute set  $\Sigma(A, \mathbf{V})$ . Combined with the distribution of auction parameters across user search queries, an equilibrium (which is an element of  $\Sigma(A, \mathbf{V})$ ) generates the distribution of observable auction outcomes (allocations and payments of bidders)  $D(\cdot)$  across those user queries.

If the set of all equilibria  $\Sigma(A, \mathbf{V})$  is not a singleton, then there will be many equilibrium distributions of outcomes  $D(\cdot)$  that correspond to value profile  $\mathbf{V}$ . Conversely, for each distribution of auction outcomes  $D(\cdot)$  there may be multiple value profiles of the bidders that could have generated that distribution of outcomes. Our next step is based on exploring this information to make statements regarding the counterfactual welfare (the sum of bidder utilities and the revenue of the auctioneer) in auctions.

We define the Empirical Bayesian Price of Anarchy (EPOA) as the worst-case ratio of welfare of the optimal allocation to the welfare of an equilibrium in a given auction  $A$ , taken over all value profiles of the bidders  $\mathbf{V}$  and Nash equilibria that could have generated the observed distribution of auction outcomes  $D(\cdot)$ . The notion of EPOA allows us to provide bounds on the (unobserved)

welfare of the optimal allocation (OPT) and the welfare of the auction  $A$  that is actually implemented,

$$\frac{\text{WELFARE}(\text{OPT})}{\text{WELFARE}(A)} \leq \text{EPoA}(A; D). \quad (5)$$

Thus, the EPoA is the characteristic of an auction that allows us to measure the efficiency of the current auction mechanism without estimating (the set of) values of the bidders.

The empirical price of anarchy (and, subsequently, the bound for optimal welfare) is defined for the true distribution of auction outcomes  $D(\cdot)$ . Then, the idea is to replace the true distribution of auction outcomes with empirical distribution  $\widehat{D}(\cdot)$  of auction outcomes observed in the data. Given that the empirical distribution is a consistent estimator of the true distribution<sup>c</sup> the bound for the welfare constructed using EPoA ( $A; \widehat{D}$ ) will bound the actual auction welfare with probability approaching one (as we get more samples from the distribution of auction outcomes  $D(\cdot)$ ). We call EPoA ( $A; \widehat{D}$ ) the *estimator* for EPoA.

We note that even the estimator for EPoA is defined as a potentially complex constrained optimization problem. It turns out that it is possible to avoid solving this problem by invoking the *revenue covering* approach. The revenue covering approach is based on establishing the ratio of the actual auction revenue and maximum payment that can be extracted from participating bidders in equilibrium. This ratio can be used to establish a simple bound for EPoA. We now describe this approach in more detail in application to the sponsored search auction.

Consider the sponsored search auction model with uncertainty as we described in detail. We can define the average price per click for bidder  $i$  with bid  $b_i$  as  $\text{ppc}_i(b_i) = \text{TE}_i(b_i)/Q_i(b_i)$ . The typical function that provides the expected number of clicks as a function of the bid  $Q_i(b_i)$  in the sponsored search auction is continuous and monotone. As a result, we can construct its inverse  $b = Q_i^{-1}(z)$  that specifies the bid that is needed to get the expected number of clicks  $z$ . Then, the average price per click

can be redefined in terms of expected clicks as  $\text{ppc}_i(z) = \text{ppc}_i(Q_i^{-1}(z))$ , which is the average price per click that bidder  $i$  getting  $z$  clicks pays. Function  $\tau_i(z) = \text{ppc}_i^{-1}(z)$  is called the *threshold* because it corresponds to the minimum price per click that bidder  $i$  needs to pay to get  $z$  clicks in expectation. The *cumulative threshold* for agent  $i$  who gets  $Q_i$  expected clicks is  $T_i(Q_i) = \int_0^{Q_i} \tau_i(z) dz$ .  $T_i(Q_i)$  can be interpreted as the total payment that bidder  $i$  would have paid, had she purchased each additional click at the average price of previously purchased clicks when she purchases a total of  $Q_i$  clicks.

**DEFINITION 1.** Strategy profile  $\sigma$  of auction  $A$  (defining the mapping from bidders' values into their bids) is  $\mu$ -revenue covered if for any feasible allocation  $\{Q_i\}_{i=1}^I$

$$\mu \text{Revenue}(A, \sigma) \geq \sum_i T_i(Q_i). \quad (6)$$

Auction  $A$  is  $\mu$ -revenue covered if for any strategy profile  $\sigma$ , it is  $\mu$ -revenue covered.

The inequality Equation (6) can be defined in expectation over the realizations of possible equilibrium profiles and all thresholds corresponding to the realizations of bid profiles. Then, the expected revenue from the auction can be measured directly by the auction platform (as a sum of payments of bidders per auction over the observed auction realizations). The thresholds can be computed via simulation from the observed auction realizations given that the allocation and pricing rule is controlled by the auction platform, and thus empirical equivalents of  $\text{TE}_i(\cdot)$  and  $Q_i(\cdot)$  can be computed for each bid  $b_i$ . As a result, the platform can compute the revenue covering parameter for given auction mechanism  $A$ .

The next result, developed in Syrgkanis and Tardos,<sup>37</sup> takes revenue covering parameter and provides the EPoA for the auction mechanism  $A$ .

**THEOREM 1.** *The welfare in any  $\mu$ -revenue covered strategy profile  $\sigma$  of auction  $A$  is at least a  $\frac{\mu}{1-e^{-\mu}}$ -approximation to the optimal welfare. In other words,  $\text{EPoA}(A; D) \leq \frac{\mu}{1-e^{-\mu}}$ .*

The estimator for the EPoA is then obtained by replacing the true revenue covering parameter  $\mu$  with its estimator

recovered from the data. As a result, this approach allows one to establish the bounds for auction welfare bypassing complex computations required in the approaches previously used in the economics literature by combining statistical inference and results from the algorithmic game theory. The Syrgkanis and Tardos<sup>37</sup> approach may potentially be applied to other bounds, for example, comparing the welfare of a given auction mechanism to the welfare of another auction mechanism (instead of the welfare of the optimal allocation). We believe that such an analysis is an important direction for future research.

*Advertising markets with learning agents.* Real-world advertising platforms are complex systems with thousands of bidders who compete in the same auction with bidders dynamically changing their bids, entering and exiting auctions. In this case, information requirements for bidders to derive their Bayes-Nash equilibrium profiles are truly impractical since they are required to form beliefs over the actions of all of their thousands of opponents, as well as the dynamic adjustment of auction parameters by the advertising platform.

In practice, most bidders in these large advertising platforms use algorithmic tools that allow them to automatically and dynamically update their bids for multiple ads and advertising campaigns. The algorithmic solutions implemented in these tools take the advertisers goals (in terms of yield of auction outcomes) as inputs, and adjust bids using dynamic feedback from the auction outcomes. Such implementations can be associated with algorithmic learning, where the bidding strategy is treated as the goal of online statistical learning procedure.

Recent work by Blum et al.,<sup>8,9</sup> Blum and Y. Mansour,<sup>10,11</sup> Caragiannis et al.,<sup>12</sup> Hartline et al.,<sup>19</sup> Kleinberg et al.,<sup>22</sup> Roughgarden,<sup>36</sup> and Syrgkanis and Tardos<sup>37</sup> shows that some of the worst case outcome properties of full information pure Nash equilibria extend to outcomes when all players use no-regret or low-regret learning strategies, assuming the game itself is stable. The assumption that players use low-regret learning to adjust their strategies is attractive for a number of reasons. First, low-regret learning outcome

<sup>c</sup> Formally, the infinity norm  $\sup_t |\widehat{D}(t) - D(t)|$  converges to zero in probability under mild assumptions regarding  $D(\cdot)$ .

generalizes the Bayes-Nash equilibrium assumption. Rather than assuming that at any time players' actions form a Bayes-Nash equilibrium, it is only necessary to assume that each player has no-regret for any fixed action over a longer time period. Thus, this assumption reduces the requirement to gather player aggregate behavior over multiple runs of the game rather than the constraint of optimality in each game. If behavior of all agents is stable over the time period, this exactly fits the Nash equilibrium assumption; however, this also allows evolving play. Second, there are many well-known learning strategies that guarantee players achieve no regret, including the weighted majority algorithm<sup>2</sup> also known as Hedge,<sup>17,26</sup> and regret matching<sup>18</sup> just to name a few simple ones.

In the context of bidders using algorithmic strategies, we cannot directly estimate their preferences from observed actions. While in Nash equilibrium, bid is the expression of bidder's value per click, but with algorithmic strategy, the bid is solely the outcome of the algorithm and thus the link between the bid and the value can be lost if the algorithm only approximately optimizes the bidder's objective function. In this case there may be a range of possible values that are compatible with the observed bids.

Nekipelov et al.<sup>33</sup> use the concept of  $\epsilon$ -regret learning to study online learning in sponsored search auctions. In online learning settings the regret of a given learning algorithm measures the difference between cumulative loss achieved by predictions given by the algorithm and loss of the best predictive function (a.k.a. "expert") from a given reference set. This concept can be applied to learning in repeated sponsored search auctions. The loss function will be associated with the negative utility in the sponsored search auction described earlier (that is, the loss minimization will correspond to the utility maximization). Nekipelov et al.<sup>33</sup> suggest set the strategies that adhere to a constant bid over the sequence of auctions. This idea is motivated by the empirical observation of Microsoft's Bing advertising platform, where a large fraction of bidders do not change bids over the lifetime of their advertisements.

## Real-world advertising platforms are complex systems with thousands of bidders dynamically changing their bids, entering and exiting auctions.

Then, the average regret of bidder  $i$  over  $T$  periods against the expert that sets the constant bid  $b'$  can be evaluated as,

$$\text{Regret}_i(b') = \frac{1}{T} \sum_{t=1}^T (U_i(b', b_{-i,t}; v_i) - U_i(b_{i,t}, b_{-i,t}; v_i)),$$

where index  $t$  correspond to the time period. The sequence of bids  $\{b_{i,t}\}_{t=1}^T$  achieves  $\epsilon_i$ -average regret if for any expert  $b'$  (from the bid space)  $\text{Regret}_i(b') \leq \epsilon_i$ . Based on this definition Nekipelov et al.<sup>33</sup> introduce the notion of a *rationalizable set*. We note that if a given sequence of bids  $\{b_{i,t}\}_{t=1}^T$  has  $\epsilon_i$ -average regret, then the value per click  $v_i$  of bidder  $i$  and the average regret  $\epsilon_i$  satisfy the set of inequalities

$$\frac{1}{T} \sum_{t=1}^T (U_i(b', b_{-i,t}; v_i) - U_i(b_{i,t}, b_{-i,t}; v_i)) \leq \epsilon_i$$

for each  $b'$ . Since bid sequences are observed in data, and the components of bidder utility (expected clicks and expected payment) can be simulated, these inequalities impose restrictions on pairs  $(v_i, \epsilon_i)$ . The set of these pairs is the rationalizable set is denoted as  $\mathcal{NR}$ . The range of bidder values contained in this set characterizes all possible values per click that rationalize the data from this bidder. Expected regret of a player reflects the properties of the learning algorithm used.

Under continuity and monotonicity of expected payment  $\text{TE}_i(\cdot)$  and expected click  $Q_i(\cdot)$  functions, Nekipelov et al.<sup>33</sup> establish basic geometric properties of the rationalizable set, such as its convexity and closeness. They also present an elegant geometric characterization that suggests an efficient algorithm for computing the rationalizable set. Since closed convex bounded sets are fully characterized by their boundaries, one can use the notion of the *support function* to represent the boundary of set  $\mathcal{NR}$ . The support function of a closed convex set  $\mathcal{NR}$  is  $h(\mathcal{NR}, \mathbf{u}) = \sup_{x \in \mathcal{NR}} \langle x, \mathbf{u} \rangle$ . Geometrically, the support function is the distance and hyperplane that it is tangent to the set  $\mathcal{NR}$  with normal vector  $\mathbf{u}$ . Since  $\mathcal{NR}$  is a two-dimensional set, then  $\mathbf{u}$  is a two-dimensional vector of unit length. An important result from convex analysis shows that support functions fully characterize closed convex sets.

Nekipelov et al.<sup>33</sup> derive the support



function of the rationalizable set by defining the following functions that can be computed directly from the auction data via simulation

$$\Delta P(b') = \frac{1}{T} \sum_{t=1}^T (\text{TE}_i(b', b_{-i,t}) - \text{TE}_i(b_{i,t}, b_{-i,t})),$$

$$\Delta C(b') = \frac{1}{T} \sum_{t=1}^T (Q_i(b', b_{-i,t}) - Q_i(b_{i,t}, b_{-i,t})).$$

These functions characterize how an average outcome in  $T$  auctions changes when bidder  $i$  switches to a fixed alternate bid  $b'$  from an actually applied bid sequence. The characterization of the rationalizable set is given in the following theorem.

**THEOREM 2.** *Under monotonicity and continuity of  $\Delta P(\cdot)$  and  $\Delta C(\cdot)$  the support function of  $\mathcal{NR}$  with  $\mathbf{u} = (u_1, u_2)^T$  and  $\|\mathbf{u}\| = 1$  is  $h(\mathcal{NR}, \mathbf{u}) = |u_2| \Delta C \left( \Delta P^{-1} \left( \frac{u_1}{|u_2|} \right) \right)$ .*

This theorem establishes that valuations and algorithm parameters for  $\epsilon$ -regret algorithms can be recovered from the data (by computing functions  $\Delta C(\cdot)$  and  $\Delta P(\cdot)$ ). If the bids constitute the Nash equilibrium, we can pinpoint the bidders' value per click. As explained earlier, the initial stage of learning may not be a Nash equilibrium and there will be an entire range of values compatible with the data. At the same time, the characterization of this range of values and the entire rationalizable set for learning bidders are reduced to evaluation of two one-dimensional functions from the data. We can use efficient numerical approximation for such an evaluation. The corresponding error in the estimated rationalizable set will be determined by the estimation error of functions  $\Delta C(\cdot)$  and  $\Delta P(\cdot)$ .

### Platform Design and User Behavior

*Inference of user actions.* As discussed, the clickthrough rate measuring the probability that a given user will click on a given ad is a crucial input in pricing and allocation rules for advertising auctions. Advertising platforms use sophisticated machine learning tools to make such predictions on a user-by-user basis. However, despite a significant amount of effort and sophisticated statistical approaches, measuring the casual impact of advertising on the actions of users has been shown to be exceedingly difficult.<sup>25</sup> In addition, since a larger emphasis has been placed

**Despite a significant amount of effort and sophisticated statistical approaches, measuring the causal impact of advertising on the actions of users has been shown to be exceedingly difficult.**

on the demand side of the advertising market (that yields revenue to advertising platforms), the work on the analysis of user behavior is significantly more sparse. Here, we briefly discuss some recent empirical results on the analysis of user behavior. This aspect of the market leaves a plethora of opportunities for future research.

The basic idea that explains the emergence of this problem is that user actions (such as the clicking on an ad) are correlated with observable and unobservable user characteristics. At the same time, ad delivery ensures ads are, by their very design, delivered to customers or during time periods that have different purchasing probabilities as compared to the baseline population. As a result, the straightforward approach based on predicting user action by using user characteristics can lead to overestimating the impact of advertising as shown in Blake et al.<sup>7</sup> and Lewis et al.<sup>35</sup>

To validate observational estimates of user behavior (such as the click-through rates for different advertising slots), advertising platforms rely on fully randomized experiments. These experiments are used by advertising platforms to infer important quantities of interest. For example, search engines exogenously re-order ad ranking and limit the number of ads shown in order to produce training data for many algorithms that govern ad delivery systems. These experiments are conducted on a small percentage of traffic so as to limit revenue risk, but nonetheless, still capture millions of searches. The experiments provide ground truth estimates of user behavior such as the causal effect of an ad as well as the impact of "organic" results on ad effectiveness. However, given their size, these estimates often need to be aggregated over many advertisers to obtain reasonable precision.

*New directions for further work.* Recently, advertising platforms have begun shifting toward advertiser goal-based pricing. In other words, if an advertiser's ultimate goal is sales to customers, then an advertiser may not be interested in the clicks per se. However, the noise in advertiser-specific estimates makes "reasonable size" experiments inadequate for supporting market design solutions that would have advertiser only pay for marginal clicks (i.e., the clicks that were actually

caused by the ad). This is because search engines, quite naturally, do not want to rest a pricing solution on the estimates with a large amount of noise with millions, if not billions, of dollars at risk. The current practice is for advertisers to figure their true return to advertising for themselves. This approach is inefficient because the platform possesses far more information to conduct the inference, and thus is in a much better position to conduct price-to-value conversion. New generations of platform design will break through the bottleneck of the inefficiency leading to a decreased loss of social welfare and a potential increase in the platform revenue.

The detailed logging of digital advertising, combined with scale achieved on major platforms, opens the potential for new observational methods that could fulfill the promise of routinely measuring advertising effectiveness in an unbiased, precise manner. Inference of user intentions and prediction of user actions are exceedingly difficult outside of search engines where advertisers explicitly specify queries as proxy of intention. In social, entertainment-, or task-oriented contexts, rendering potential advertisers in front of a user via dynamic mechanism is the most efficient, but not an easily implementable approach.<sup>38</sup>

While advertisers and advertising platforms clearly move to performance-based pricing, the user behavior is becoming more complex. Current trends show that device usage will continue to change how users consume information, enjoy leisure time, and communicate with each other (see Fulgoni and Lipsman<sup>39</sup> and Xu et al.<sup>40</sup>). In this context, the advertising ecosystem is evolving. Brands have changed ad formats to engage and interact with users in-video, in-game, and in other dynamic content. The hyper-local nature of mobile applications provides a new type of a signal about the user. This more refined user information, in turn, changes the composition of advertisers. The typical persona an advertiser has shifted from large business entities like big-name brands or e-commerce platforms on Google and Facebook, to deep vertically integrated professionals. For example, on Yelp, Zillow, or Grubhub, advertisers are small business owners or individual professionals

like real estate agents or plumbers. These business users do not have the sophisticated business knowledge or technology skills to aid them in ad campaign management. The user actions relevant for these new breeds of advertisers shifted away from clicks to in-app text, direct call, or email communication. Ideally, new vertical marketplaces will provide users with better experiences and advertisers with more accurate user intent and direct access to relevant users. But, it requires advertising platforms to adapt and evolve. Advanced machine learning technology is essential for these new platforms to accurately account for user actions with a large volume of data points across devices. Another key necessity is the design of new auction mechanisms that encompass the need to provide the advertisers with simple and easy ways to manage work flows. The technological change that has been occurring over the past decade creates the need for new platforms that are both grounded in the game theory but also account for complex behavioral responses of users and advertisers.

### Acknowledgments

The authors are indebted to Éva Tardos who suggested the theme of this article and helped them shape it. □

### References

1. Aradillas-Lopez, A. Tamer, E. The identification power of equilibrium in simple games. *J. Bus. Econ. Stat.* 26, 3 (2008), 261–283.
2. Arora, S., Hazan, E., Kale, S. The multiplicative weights update method: A meta-algorithm and applications. *Theory of Computing* 8, (2012), 121–164.
3. Athey, S., Nekipelov, D. *A Structural Model of Sponsored Search Advertising Auctions*. AdAuctions'10, 2010.
4. Bajari, P., Hong, H., Nekipelov, D. Game theory and econometrics: a survey of some recent results. *Advances in Economics and Econometrics* 3, (2013), 3–52.
5. Bajari, P., Hong, H., Krainer, M., Nekipelov, D. Estimating static models of strategic interactions. *J. Bus. Econ. Stat.* (2011).
6. Beresteanu, A., Molchanov, I., Molinari, F. Sharp identification regions in models with convex moment predictions. *Econometrica* 79, (2011), 1785–1821.
7. Blake, T., Nosko, C., Tadelis, S. Customer heterogeneity and paid search effectiveness: a large scale field experiment. *Econometrica* 83, 1 (2015), 155–174.
8. Blum, A., Even-Dar, E., Ligett, K. Routing without regret: on convergence to Nash equilibria of regret-minimizing algorithms in routing games. *PODC '06*, 45–52.
9. Blum, A., Hajiaghayi, M., Ligett, K., Roth, A. Regret minimization and the price of total anarchy. *ACM STOC*, (2008), 373–382.
10. Blum, A., Mansour, Y. From external to internal regret. *J. Mach. Learn. Res.* 8, (2007), 1307–1324.
11. Blum, A., Mansour, Y. Learning, regret minimization, and equilibria. In *Algorithmic Game Theory*. N. Nisan et al. eds. Cambridge University Press, 2007, 79–101.
12. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M., Lucier, B., Paes Leme, R., Tardos, E. Bounding the inefficiency of outcomes in generalized second price auctions. *J. Econ. Theo.* 156, (2015), 343–388.
13. Ciliberto, F., Tamer, E. Market structure and multiple

- equilibria in airline markets. *Econometrica* 77, 6 (2009), 1791–1828.
14. Cole, R., Roughgarden, T. The sample complexity of revenue maximization. *ACM STOC*, (2014), 243–252.
15. Cummings, R., Ligett, K., Nissim, K., Roth, A., Wu, Z. Adaptive learning with robust generalization guarantees. In *29th Annual Conference on Learning Theory*, (2016), 772–814.
16. Edelman, B., Ostrovsky, M., Schwarz, M. Internet advertising and the generalized second-price auction. *American Economic Review* 97, 1 (2007), 242–259.
17. Freund, Y., Hsu, D. A new Hedging algorithm and its application to inferring latent random variables. *ArXiv:0806.4802*
18. Hart, S., Mas-Colell, A. A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68, (2000), 1127–1150.
19. Hartline, J., Syrgkanis, V., Tardos, É. *No-Regret Learning in Repeated Bayesian Games*. *Advances in Neural Information Processing Systems* 28 (NIPS 2015).
20. Hoy, D., Syrgkanis, V., Nekipelov, D. Robust data-driven efficiency guarantees in auctions. *arXiv:1505.00437*, 2015
21. Kesselheim, T., Kleinberg, R., Tardos, E. Smooth online mechanisms: a game-theoretic problem in renewable energy markets. *EC'15*, 203–220, 2015.
22. Kleinberg, R., Piliouras, G., Tardos, E. Load balancing without regret in the bulletin board model. *Distributed Computing* 24, 1 (2011), 21–29.
23. Koutsoupias, E., Papadimitriou, C. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science* (1999), 404–413.
24. Lahaie, S., McAfee, P. Efficient Ranking in Sponsored Search. *Internet and Network Economics* (2011), 254–265.
25. Lewis, R. Rao, J. The unfavorable economics of measuring the returns to advertising. SSRN 2367103, 2014.
26. Littlestone, N., Warmuth, M. The weighted majority algorithm. *Information and Computation* 108, 2 (1994), 212–260.
27. Lykouris, T., Syrgkanis, V., Tardos, É. Learning and efficiency in games with dynamic population. *SODA'16*, 120–129, 2016.
28. Morgenstern, J., Roughgarden, T. On the pseudo-dimension of nearly optimal auctions. *NIPS'15*, 136–144, 2015.
29. Muthukrishnan, S. Ad exchanges: research issues. *Internet and Network Economics*, 1–12, Springer, 2009.
30. Aggarwal, G., Muthukrishnan, S. Theory of sponsored search auctions. *foundations of computer science*, 2008, FOCS'08. IEEE 49th Annual IEEE Symposium on. IEEE, 2008.
31. Muthukrishnan, S. Internet ad auctions: insights and directions. *Automata, Languages and Programming*, 14–23, Springer, 2008.
32. Myerson, R. Optimal auction design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
33. Nekipelov, D., Syrgkanis, V., Tardos, E. Econometrics for learning agents. *EC '15*, 1–18, 2015.
34. Ostrovsky, M., Schwarz, M. Reserve prices in internet advertising auctions: a field experiment. *EC'11*, 59–60, 2011.
35. Lewis, R., Rao, J., Reiley, D. *Measuring the Effects of Advertising: The Digital Frontier, Economic Analysis of the Digital Economy*. University of Chicago Press, 2015, 191–218.
36. Roughgarden, T. Intrinsic robustness of the price of anarchy. *Journal of the ACM* 62, 5 (2015), article No. 32.
37. Syrgkanis, V., Tardos, É. Composable and efficient mechanisms *STOC '13*, 211–220, 2013.
38. Ballings, M., Van den Poel, D. CRM in social media: predicting increases in facebook usage frequency. *Eur. J. Oper. Res.* 244, 1 (2015), 248–260.
39. Fulgoni, G., Lipsman, A. Numbers, please: digital game changers: how social media will help usher in the era of mobile and multi-platform campaign-effectiveness measurement. *J. Advert. Res.* 54, 1 (2014), 11–16.
40. Xu, J., Forman, C., Kim, J.B., Van Ittersum, K. News media channels: complements or substitutes? evidence from mobile phone usage. *J. Mark.* 78, 4 (2014), 97–112.
41. Varian, H. Online ad auctions. *Am. Econ. Rev.* 99, 2 (2009), 430–434.

**Denis Nekipelov** (denis@virginia.edu) is an associate professor in the departments of Economics and Computer Science at the University of Virginia, Charlottesville.

**Tammy Wang** (twang@rivierapartners.com) is VP of Data Science and Analytics at Riviera Partners, San Francisco, CA.