

Adverse Selection and Auction Design for Internet Display Advertising

By NICK ARNOSTI AND MARISSA BECK AND PAUL MILGROM*

We model an online display advertising environment in which “performance” advertisers can measure the value of individual impressions, whereas “brand” advertisers cannot. If advertiser values for ad opportunities are positively correlated, second-price auctions for impressions can be inefficient and expose brand advertisers to adverse selection. Bayesian-optimal auctions have other drawbacks: they are complex, introduce incentives for false-name bidding, and do not resolve adverse selection. We introduce “modified second bid” auctions as the unique auctions that overcome these disadvantages.

When advertiser match values are drawn independently from heavy tailed distributions, a modified second bid auction captures at least 94.8% of the first-best expected value. In that setting and similar ones, the benefits of switching from an ordinary second-price auction to the modified second bid auction may be large, and the cost of defending against shill bidding and adverse selection may be low.

Since the pioneering papers by Edelman, Ostrovsky and Schwarz (2007) and Varian (2007), there has been a growing body of research focusing on auctions for sponsored search advertising on the internet. These automated auctions, which are initiated when a consumer enters a search query, allow advertisers to bid in real time for the opportunity to post ads near unpaid search results.

Markets for internet display advertising, which determine ad placements on all other kinds of web pages, traditionally operated quite differently. Following practices established for offline media, most impressions were sold through contracts that guaranteed an advertiser a large volume of impressions over a pre-specified time period, for a negotiated price. More recently, the high revenues associated with sponsored search auctions prompted the display advertising industry to introduce “real-time bidding,” in which advertisers bid in auctions for individual impressions. Proponents of real-time bidding argue that it alone has the potential to allow each online impression to be routed to the advertiser who values it most highly. Achieving this theoretical ideal, however, requires that each advertiser

* Arnosti: Stanford University, Department of Management Science and Engineering, 475 Via Ortega, Stanford, CA 94305 (e-mail: nicholas.arnosti@gmail.com); Beck: Charles River Associates, 5335 College Ave., Suite 26, Oakland, CA 94618 (e-mail: mbeck@crai.com); Milgrom: Stanford University, 579 Serra Mall, Stanford, CA 94305 (e-mail: milgrom@stanford.edu). Paul Milgrom and Marissa Beck were coinventors of the MSB auction design. They were paid advisors to OpenX Corporation, which has applied for a patent based on this design. Our subsequent study of the properties of this design received support from the National Science Foundation under award 1525730.

knows its value for each individual impression.

What advertisers know about their values depends in part on the types of feedback they collect. Some advertisements are designed to elicit an immediate response from the consumer – most commonly, a click on a link to a website where a sale may occur. By aggregating click and sales data from many impressions, an advertiser might learn how various characteristics of past opportunities have correlated with ad performance. With this knowledge, such a “performance advertiser” could predict the value of future advertising opportunities with reasonable accuracy.

Other advertisements, meanwhile, do not seek an immediate response from the user. Instead, they aim to notify the user of a sales event, movie opening, or other upcoming opportunity, or to raise the user’s awareness of an individual brand or product. Recent work by Lewis and Rao (2015) highlights the difficulty of estimating the return from such an ad campaign; distinguishing the values of individual ads is even more difficult. The limited feedback available to these “brand advertisers” makes it hard for them to reliably estimate values for individual ad opportunities.

For brand advertisers, real-time bidding presents an additional challenge by introducing the possibility of adverse selection. Advertisers’ values for different ad opportunities tend to be positively correlated: all advertisers want to avoid showing their ads to automated web crawlers and most prefer to show more ads to consumers with greater disposable incomes and responsiveness to online advertising. As a result, performance advertisers tend to bid higher on opportunities that are more valuable to brand advertisers and brand advertisers may have difficulty measuring the resulting losses and adjusting their bids accordingly.

This work compares the performance of different auction designs when advertiser valuations are positively correlated and advertisers are differentiated in their abilities to estimate the values of individual impressions. Most of the paper assumes that the mechanism designer chooses an auction with the goal of maximizing the efficiency of the allocation, although in Section VI we briefly address the alternate objective of maximizing revenue.

We introduce our model in Section II. Each performance bidder’s value for an ad opportunity is the product of two factors: the overall quality of the opportunity, which affects the value for all advertisers, and an idiosyncratic match quality, which is unique to the particular bidder. If values are primarily determined by overall quality with little variation in idiosyncratic match qualities, then the benefits from using real-time bidding to assign individual impressions are minimal. If, on the other hand, differences in bidders’ values arise entirely from differences in match quality with little variation in overall quality, then adverse selection is absent and real-time bidding using a second-price auction achieves an efficient allocation. Ideally, one would want a mechanism that performs well in both of these extreme cases as well as in intermediate ones in which common and idiosyncratic components are both significant determinants of advertiser valuations. Such a ro-

bust mechanism might be especially attractive if advertisers disagree about their ability to measure value and about the importance of adverse selection.

We consider several ad allocation mechanisms:

- 1) Second-price auctions with reserves, resembling current practice.
- 2) An optimal auction, which maximizes the expected value to the advertisers, given the auctioneer’s beliefs.
- 3) A new parameterized class of “modified second bid” (MSB) auctions, which allocate each impression to the highest performance bidder provided that the ratio of its bid to the second highest performance bid exceeds some factor $\alpha \geq 1$.
- 4) The “omniscient” or “first-best” benchmark, which is the auction that maximizes allocative efficiency assuming that the auctioneer, in addition to observing the performance bids, also observes the overall quality of each impression.

We evaluate the performance of these mechanisms under two conditions. The first allows any distributions for the overall quality variable and the match quality variables. The second restricts the match values to follow power law distributions (also known as “Pareto” distributions). This restriction aims to capture the idea that much of an advertiser’s value from real-time bidding comes from relatively few matches of very high quality, such as individuals who recently viewed a product in its online store.¹

Section III demonstrates that even in the limited class of power law environments, the second-price auction with an optimal reserve can capture as little as half of the value delivered by the benchmark omniscient mechanism. The Bayesian-optimal mechanism “OPT” does better by using information in performance advertiser bids to maximize efficiency. However, we show in Section IV that OPT encourages “false-name bidding,” in which the high bidder increases its chance of winning and reduces its expected price by placing a low bid using an additional account. In this auction, placing just one bid may be a dominated strategy. Even if false-name bidding could somehow be prevented, another drawback of OPT is adverse selection against brand advertisers: the impressions awarded to those advertisers may be of disproportionately low value. If the parties have different beliefs about the unobservable distributions of overall and match qualities, the presence of adverse selection could prevent them from agreeing on a contract price.

¹The power law distribution is given by $F(x) = 1 - x^{-a}$ for $x \in [1, \infty)$. The parameter a determines the weight of the power law tail. For $a > 1$, the mean of the distribution is finite (in particular, it is $\frac{a}{a-1}$). A random variable X has this power law distribution if and only if $\ln(X)$ has an exponential distribution with mean $1/a$, which holds if and only if $E[X|X > y] \equiv yE[X]$ for all $y > 1$. The Wikipedia article on the “Pareto Principle,” describes how the 80-20 rule of marketing is modeled by the power law distribution of sales across customers.

These drawbacks of the optimal auction lead us to seek out a class of deterministic, anonymous mechanisms that eliminate adverse selection and the incentives for false-name bidding. We show in Section V that this class is exactly the family of modified second bid (MSB) auctions. This axiomatic approach to identifying the MSB mechanism is unusual: auction theory analyses more often focus on optimal mechanisms for specific environments. Nevertheless, our analysis follows a respected tradition in economic theory and mechanism design, in which authors have often imposed properties like stability, envy-freeness, and strategy-proofness that do not depend on detailed distributional assumptions. One goal of this axiomatic approach is to respect the Wilson doctrine by creating a mechanism whose performance is not too sensitive to detailed assumptions about the economic environment.

Section VI establishes that in addition to eliminating false-name bidding and adverse selection, MSB auctions also perform well relative to both optimal and second-price auctions. Without restrictions on match value distributions, all three have the same worst-case performance, which is 50% of the performance of the omniscient mechanism. More surprisingly, if match values follow a power law, then the worst-case performance of second-price auctions is still only half of the value from an omniscient mechanism, whereas MSB auctions guarantee 94.8%! This is not an artifact of using efficiency as our objective. In the same model, MSB also guarantees 94.8% of the *revenue* generated by any strategy-proof mechanism.

Why do MSB auctions perform so well when match quality has a power law distribution? For an auction to lead to efficient allocations, it must distinguish cases in which some performance advertiser has a high match quality (and so should win the impression) from ones in which no performance advertiser has high match quality (so that the brand advertiser should win). Intuitively, when match quality has a power law distribution, a large fraction of the total value to performance advertisers is generated when one performance advertiser has a match quality that is at least α times higher than the next highest match quality, while this same inequality is unlikely to be satisfied if the impression is a poor match for all performance advertisers. So, with a careful choice of α , MSB can capture a large fraction of the total value from both performance and brand advertisers.

I. Related Work

There is a large and growing literature on auctions for internet advertising. Various papers focus on the role of risk aversion, budget constraints, dynamic contract formation and fulfillment in the presence of uncertain supply, revealing and/or selling cookies, the role of intermediaries, and many other concerns. In this work, we ignore dynamic considerations and budgets to focus on the information structure and the potential difficulty of even the seemingly simple goal of determining a myopically efficient allocation.

The line of work most closely related to our own is that which asks the question

of how to jointly allocate impressions across advance contracts and spot market sales. The primary differences between these papers and our own are the assumptions about the correlation between advertiser valuations. Prior work has focused on extreme cases by assuming either that spot market bids do not provide information about the brand advertiser’s value (Chen, 2010; Balseiro et al., 2014) or that they perfectly reveal this value (Ghosh et al., 2009). In either case, determining a (myopically) efficient allocation is trivial, so these papers focus on other topics. By contrast, we formulate a model with a mix of common and private values, so that spot market bids provide information about the brand advertiser’s value for the impression, but do not perfectly determine it.

Milgrom and Weber (1982) provide one of the earliest analyses of auction environments with a mix of common and private values. Due to the symmetry in the assumed environment, standard auctions yield an efficient allocation. Some later work has investigated information asymmetry in settings where auctions have a common value component. Most recently, Abraham et al. (2013) study information asymmetry and adverse selection in the context of internet advertising. They focus on the auctioneer’s revenue and most of their work considers an environment with pure common values, so that achieving an efficient allocation is trivial.

II. Model

We consider the allocation of a single impression. There are $N + 1$ advertisers competing for this impression, with N random and $P(N \geq 2) = 1$. Advertiser i has value X_i for the impression. The value X_i is the product of two terms: a “common value” C and a “match value” M_i . We interpret C as capturing attributes of the user that are valuable to all advertisers, such as the user’s income and responsiveness to online advertising. Meanwhile, the value M_i captures idiosyncratic components that contribute to the quality of the match for advertiser i . We assume that, given N , the M_i are drawn independently from (not necessarily identical) distributions F_i , that C is drawn from a distribution G , and that $M = (M_0, \dots, M_N)$ is independent from C . Furthermore, we assume that both match and common values have finite expectations. We use $X_{(k)}$ and $M_{(k)}$ to denote the k^{th} highest value and match value factor, respectively.

We assume that advertisers $i \in \{1, \dots, N\}$ (the “performance advertisers”) observe their values X_i , but not the components C and M_i . Meanwhile, advertiser 0 (whom we refer to as a “brand advertiser”) **cannot observe** X_0 .² The assumption that there is a single uninformed advertiser and multiple informed advertisers is made for expositional simplicity and could be easily relaxed as part of a richer model.

²If the brand advertiser can measure the value of its overall campaign and nothing more, then it observes $E[X_0]$, but not the distributions G and F_0 separately. This motivates the axiom of adverse-selection freeness discussed in Section V.

Making use of the revelation principle and the assumed risk neutrality of bidders, we consider a *mechanism* to be a mapping from the privately held information $X = (X_1, \dots, X_N)$ to allocation probabilities z and payments p . For $i \in \{0, \dots, N\}$, we let $z_i(X)$ be the probability that advertiser i wins and $p_i(X)$ be advertiser i 's expected payments, given X . For fixed F_i and G , given an allocation rule z , we define the total surplus from impressions awarded to the brand advertiser by $V_B(z) = E[X_0 z_0(X)]$. Similarly, we define the surplus from impressions awarded to performance advertisers by $V_P(z) = E\left[\sum_{i=1}^N X_i z_i(X)\right]$. We consider objectives that are weighted sums of these two terms. In particular, for fixed $\gamma > 0$, we define

$$V(z) = \gamma V_B(z) + V_P(z).$$

When $\gamma = 1$, this corresponds to the total efficiency of the allocation. We show in Section VI that when match qualities are drawn from a power law distribution, if γ is chosen to reflect the relative bargaining power of the publisher and the brand advertiser, then this objective is proportional to the publisher's expected revenue, allowing our results to apply to that objective as well.

Throughout, we use as a benchmark the “first best” or “omniscient” mechanism, which operates as the publisher would if the common value C were observed. In this case, the publisher could deduce match values M from performance bids X and condition the allocation directly on match values. The omniscient allocation rule is simple: award the impression to the performance advertiser with the highest value whenever $M_{(1)} > \gamma E[M_0]$ and to the brand advertiser otherwise. We denote this allocation rule by OMN. This allocation rule is in general not implementable, but it provides an upper bound on the value attainable by *any* mechanism and a way to measure the losses due to the fact that C is unobservable.

III. Second-Price Auction

In this section, we consider one simple mechanism for the allocation problem: a second price auction. Although this mechanism is efficient if all advertisers know their own values, Proposition 1 shows that this may lose up to *half* of the value that could be obtained by OMN (i.e. if C were observed).

Because the brand advertiser does not observe its value, it has no dominant strategy, so our analysis must rely on some other criterion to predict its bid. Naively, it might bid its expected value $E[X_0]$ in each auction. This would be optimal if its value were independent of the values of other advertisers (i.e., if C were constant), but could be a very poor strategy if the correlation between the brand value and the values of other advertisers is high. Instead, we assume that the brand advertiser plays a best response to the dominant strategies of the other bidders. It chooses b to maximize its expected profit, as given by

$$\Pi(b) = E[(X_0 - X_{(1)})\mathbf{1}_{X_{(1)} \leq b}]$$

Note that in order to compute the optimal bid, the brand advertiser must know the distribution of the top performance bid and the correlation between $X_{(1)}$ and X_0 (as determined by the distribution of C). In practice, these could be challenging to learn.

If the publisher has more information about competing bidders, a reasonable alternative is for the publisher to submit a “proxy bid” on behalf of the brand advertiser. This resembles a solution to the allocation problem commonly used in practice: the publisher signs a contract with the brand advertiser and submits each opportunity, along with a corresponding reserve price, to a real-time exchange. If the reserve is met, the impression is awarded to the top bidder. Otherwise, it is allocated to the brand advertiser. We let SP_b denote the allocation rule of the second-price auction when the brand advertiser submits a bid of b (or the publisher does so on the brand advertiser’s behalf). This mechanism is simple and intuitive, and allows performance advertisers to win impressions for which they have very high values. However, for any reserve price, the brand advertiser wins more often when C is low than when C is high, which causes the resulting allocation to be inefficient.

How inefficient might the second-price auction (with optimally chosen reserve)³ be? Two extremes are to set the reserve at zero (so that the brand advertiser never wins) or to set it arbitrarily high (so that the brand advertiser always wins). Thus, the second-price auction can always deliver a value of $\max(\gamma E[X_0], E[X_{(1)}])$. Clearly, no mechanism can deliver more than $\gamma E[X_0] + E[X_{(1)}]$, so for any distribution of match and common values, the second-price auction (with appropriate reserve) delivers at least half of the available value. The next result states that this bound is tight, even if performance match values are restricted to be IID.

PROPOSITION 1:

1) For any $\gamma > 0$, $N \geq 2$, F_i , G , and $E[M_0]$, there exists $b \in \{0, \infty\}$ such that

$$V(SP_b) \geq \frac{1}{2}V(OMN).$$

2) For any $\gamma, \epsilon > 0$, there exists F, G, N and $E[M_0]$ such that if performance match values are drawn IID from F ,

$$\sup_{b \geq 0} V(SP_b) < \left(\frac{1}{2} + \epsilon \right) V(OMN).$$

To show that the bound of $1/2$ is tight, we use a sequence of examples in which both match and performance values follow power law distributions: for

³Note that in our model, the brand advertiser and publisher have the same information, and thus when $\gamma = 1$ (so that the publisher wishes to maximize total surplus) the second-price auction and “contract with proxy bidding” are mathematically equivalent. In particular, if the publisher uses a reserve of b , the resulting total surplus is $E[X_0 \mathbf{1}_{X_{(1)} \leq b}] + E[X_{(1)} \mathbf{1}_{X_{(1)} > b}] = \Pi(b) + E[X_{(1)}]$. In other words, the brand advertiser’s optimal bid in a second-price auction is precisely the bid that maximizes allocative efficiency.

$x \geq 1$, $P(M_i > x) = x^{-a}$ and $P(C > x) = x^{-c}$, where $a, c > 1$ are parameters determining the weights of the power law tails.

For the performance of the second-price auction to approach the lower bound of $1/2$, three things must be true. First, the omniscient mechanism needs to deliver nearly $\gamma E[X_0] + E[X_{(1)}]$. Second, the highest bid in the auction must convey virtually no information about the optimal assignment, so that $\sup_b V(\text{SP}_b) \approx \max(\gamma E[X_0], E[X_{(1)}])$. Third, it must be the case that $\gamma E[X_0] = E[X_{(1)}]$. The first condition holds as $a \downarrow 1$, for in that case nearly all of the performance value can be captured from a vanishingly small fraction of the impressions, while delivering the remainder to the brand advertiser. For fixed a , the second condition holds as $c \downarrow 1$ because then nearly all of the variation in the highest bid comes from variations in C (which are irrelevant when determining the optimal assignment). Given the distribution of performance match values, the third condition is satisfied by a suitable choice of $E[M_0]$. All three conditions can thus hold simultaneously, proving that the bound is tight.

Although the circumstances above are quite specific and the worst-case bound of 50% is never exactly met, this family of examples demonstrates that even an optimally selected second-price auction may be unsatisfying if the publisher faces substantial uncertainty about both match and common value factors. Motivated by this fact, we next consider the properties of the Bayesian-optimal mechanism.

IV. Bayesian-Optimal Auction

One approach to mitigating the potential inefficiency of the allocation from a second-price auction would be to solve for the Bayesian-optimal auction in our setting. However, we argue in this section that even when this rule is monotonic and, therefore, implementable by a strategy-proof mechanism, it is an impractical solution to the allocation problem. Given bids $X = x$, the Bayesian-optimal auction (“OPT”) would award the impression to the performance advertiser with the highest value whenever $E[M_{(1)}|X = x] > \gamma E[M_0]$ and to the brand advertiser otherwise. This allocation rule depends in a detailed way on the distributions F_i and G , and will in general be much more complicated than a second-price auction with a reserve. Furthermore, it suffers from two other significant drawbacks:

- 1) Performance advertisers may have a clear incentive to submit shill bids.
- 2) The quality of impressions awarded to the brand advertiser may be disproportionately low.

The first concern arises because in many web-based environments, bidders reveal only virtual identities. A single bidder could create several accounts and submit multiple bids for the same impression. If the publisher uses OPT to select the winner, such behavior can be profitable for the bidder. The reason is that the optimal allocation generally depends on the low bids as well as the high ones: the impression should be awarded to the top bidder only if its bid exceeds the

expected value of the impression to the brand advertiser, conditional on the complete bid profile. By making one or more low bids, the top bidder can reduce the conditional expectation of C , thereby winning impressions that would otherwise go to the brand advertiser. This manipulation results in a loss of efficiency.

The second concern is that the quality of impressions won by the brand advertiser under OPT may differ systematically from a random sample. The value delivered to the brand advertiser under such a contract depends on the correlation between brand and performance values, which is determined by the distributions of M and C . Since these distributions are not directly observed, the publisher and brand advertiser might disagree about them and so be unable to negotiate a suitable contract price.

We illustrate both of these drawbacks by considering the case when match values follow a power law distribution.

PROPOSITION 2: *Suppose that match values are IID draws from a power law distribution. Then*

- $(N, X_{(N)})$ is a sufficient statistic for X_0 , and $E[X_0 | N, X_{(N)}]$ is increasing in $X_{(N)}$.
- If C follows a power law distribution, then under OPT, $E[z_0(X)|C, N]$ is decreasing in C .

The first bullet point of the proposition is useful for illustrating the impact of false-name bidding. It implies that OPT sells the impression to the highest-bidding performance advertiser if $X_{(1)} > \gamma E[X_0|N, X_{(N)}]$, at a price equal to $\max(\gamma E[X_0|N, X_{(N)}], X_{(2)})$. If, for example, C is distributed according to a positive density on $[0, \infty)$, then an advertiser can ensure that $\gamma E[X_{(0)}|N, X_{(N)}] < \epsilon$ by submitting an arbitrarily small false-name bid in addition to its truthful bid. This bid carries essentially no risk, removes the brand advertiser from the competition, and effectively converts the mechanism into a second-price auction among the performance advertisers.⁴

The second point of Proposition 2 establishes the existence of adverse selection: the higher the impression's quality, the less likely it is to be awarded to the brand advertiser. The intuition for this point can be seen by regarding the conditional expectation $E[X_0|N, X_{(N)}]$ as a function of *three* variables: $(N, C, M_{(N)})$, using $X_{(N)} = CM_{(N)}$. Intuitively, the expectation blends an unbiased estimate implied by the data with the prior expected value, so it responds less than proportionately to changes in C , holding N and M constant. As a result, for any realizations of N and $M_{(N)}$, the ratio $X_{(1)}/E[X_0|N, X_{(N)}]$ is increasing in C . This results in adverse selection against the brand advertiser, who wins precisely when this ratio is less than γ . We show in Section VI that these drawbacks of OPT may

⁴This conclusion does not rely on the power law distribution. So long as match values are distributed on $[1, \infty)$ and C is distributed on $(0, \infty)$, $E[C|X_{(1)}, \dots, X_{(N)}] < X_{(N)}$, so a single low bid can reduce the estimate of X_0 sufficiently to exclude the brand advertiser from the allocation.

not be compensated by especially strong performance: for some settings, OPT delivers only half of the value that could be achieved if C were directly observed. Motivated by the limitations of the optimal mechanism, in the following section we seek a mechanism that discourages shill bidding and eliminates adverse selection.

V. An Axiomatic Approach

The previous section demonstrated that the optimal mechanism can in general be manipulated by shill bids and may award brand advertisers impressions of disproportionately low value. In this section, we study mechanisms that avoid these concerns.

We consider an extended game in which performance advertisers have the option to create and bid from multiple accounts. In such a setting, it may be infeasible to discriminate against individual bidders, so we seek a mechanism that treats bidders symmetrically. Furthermore, we require a mechanism for which it is a dominant strategy for each advertiser to play as a single bidder and to report its value truthfully. This restriction is similar to that imposed by Yokoo, Sakurai and Matsubara (2004). Compared to their work, we introduce the additional requirement that bidders cannot reduce other bidders' surplus by submitting additional low bids. This is intended to capture the possibility that advertisers with low values might be tempted to raise the prices paid by their competitors, particularly if they can do so by submitting bids that are certain not to win.

DEFINITION 1: *Given a mechanism (z, p) , let z_{-0} and p_{-0} denote the allocations and payments of performance advertisers. The mechanism is **anonymous** (among performance advertisers) if, for any $n \geq 2$, any permutation σ on $\{1, \dots, n\}$, and any $x \in \mathbb{R}_+^n$, the following hold:*

$$\sigma(z_{-0}(x)) = z_{-0}(\sigma(x)) \quad \text{and} \quad \sigma(p_{-0}(x)) = p_{-0}(\sigma(x)).$$

DEFINITION 2: *The mechanism (z, p) is **strategy-proof** if, for all $n \geq 2$, all $x \in \mathbb{R}_+^n$, all $i \in \{1, \dots, n\}$, and all \hat{x}_i ,*

$$x_i z_i(x) - p_i(x) \geq x_i z_i(\hat{x}_i, x_{-i}) - p_i(\hat{x}_i, x_{-i}).$$

*The mechanism (z, p) is **winner false-name proof** if no bidder can benefit by submitting multiple bids, meaning that for all $n \geq 2$, $x \in \mathbb{R}_+^n$, all $m \geq 1$, and all $y \in \mathbb{R}_+^m$:*

$$x_i z_i(x) - p_i(x) \geq x_i \left(z_i(x, y) + \sum_{j=1}^m z_{n+j}(x, y) \right) - \left(p_i(x, y) + \sum_{j=1}^m p_{n+j}(x, y) \right).$$

*A mechanism (z, p) is **loser false-name proof** if no bidder i can harm its*

competitors by submitting lower additional bids; that is, if for all $n \geq 2$, $x \in \mathbb{R}_+^n$, $m \geq 1$, and $y \in \mathbb{R}_+^m$ satisfying $\max y \leq x_i$, the following holds for $j \neq i$:

$$x_j z_j(x) - p_j(x) \geq x_j z_j(x, y) - p_j(x, y).$$

The mechanism (z, p) is **fully strategy-proof** if it is strategy-proof, winner false-name proof and loser false-name proof.

We also formalize the idea that a mechanism should award brand advertisers a “representative” sample of impressions, for *any* distribution of match and common values. We think of this as a robustness criterion: if the mechanism has this property, then the publisher and brand advertiser do not need to agree on details of the environment. Both agree that the value awarded to the brand advertiser is simply $E[X_0]$ times the number of impressions awarded, regardless of their beliefs about C .

DEFINITION 3: A mechanism (z, p) is **adverse-selection free** if, for every joint distribution of (C, M) such that M is independent from C , $z_0(CM)$ is also independent of C .

We now turn our attention to a simple class of mechanisms that have all of the above properties: the family of modified second bid (MSB) auctions. These auctions offer the impression to the top bidder, at a price equal to α times the second highest bid for some $\alpha \geq 1$. If the top bidder is unwilling to pay this price, the impression may be awarded to the brand advertiser.

DEFINITION 4: A mechanism (z, p) is a **modified second bid auction** if there exists $\alpha \geq 1$ such that for $i = 1, \dots, N$,

- 1) either $z_i(x) = \mathbf{1}_{x_i > \alpha \max x_{-i}}$, or $\alpha > 1$ and $z_i(x) = \mathbf{1}_{x_i \geq \alpha \max x_{-i}}$,
- 2) $p_i(x) = z_i(x) \cdot \alpha \max x_{-i}$,
- 3) $z_0(x) \in \{0, 1\}$, with $z_0(x) \leq 1 - \sum_{i=1}^N z_i(x)$.

The sole distinction between the two cases in the first condition is whether performance bidder i wins or loses when it bids *exactly* its threshold price $\alpha \max x_{-i}$. The third condition states that the brand advertiser wins only when no performance advertiser does, but allows for the possibility that the impression remains unassigned.

It turns out that modified second bid auctions are the *only* deterministic⁵ mechanisms that satisfy the properties given above.

THEOREM 1: A deterministic mechanism (z, p) is anonymous, fully strategy-proof, and adverse-selection free if and only if it is a modified second bid auction.

⁵A mechanism (z, p) is deterministic if $z_i(x) \in \{0, 1\}$ for all i and all x .

Theorem 1 characterizes MSB auctions as the unique auctions which are anonymous, deterministic, adverse-selection free, and fully strategy-proof in a setting where bidders may create multiple identities. We note here that other auctions may have any three of these properties. Relaxing anonymity allows the value of α to depend on the bidder. Relaxing determinism permits mechanisms that randomly set aside opportunities for the brand advertiser, and award the rest via a second-price auction among performance advertisers. A second-price auction with reserve (corresponding to the brand advertiser's bid) is deterministic, anonymous, and fully strategy-proof, but not adverse-selection free. Finally, consider the auction in which the highest bidder wins whenever its bid is at least twice the lowest competing bid and, in that case, pays the higher of the second-highest bid and twice the lowest bid. This mechanism is deterministic, anonymous, and adverse-selection free, but not fully strategy-proof.

In the remainder of the paper, we use the notation MSB_α to denote the modified second bid allocation rule with parameter α in which $z_i(x) = \mathbf{1}_{x_i > \alpha \max_{j \neq i} x_j}$ and $z_0(x) = 1 - \sum_{i=1}^N z_i(x)$.

VI. Performance Analysis

Although we have characterized MSB auctions as the only mechanisms that are deterministic, anonymous, false-name proof, and adverse-selection free, it is not clear that these properties are necessary. Perhaps shill bidding can be identified and prevented by other means. Additionally, the publisher may be able to convince the brand advertiser to accept below-average impressions in return for a lower price. If so, it is reasonable to wonder whether these axioms are costly, and whether we could achieve notably better performance by abandoning them.

In this section, we demonstrate that in important cases, MSB auctions perform nearly as well as their optimal counterparts. We consider a setting in which performance match values are IID (i.e., $F_i = F$ for $i \in \{1, \dots, N\}$). Our first result states that for arbitrary F , MSB auctions guarantee half of the value delivered by OMN, and the guarantee provided by OPT is no better. Our second result is that if F corresponds to a power law distribution, then MSB auctions capture at least 94.8% of the value delivered by OMN. In this case, any gains from adopting OPT are necessarily modest.

PROPOSITION 3:

1) For any $\gamma > 0$, N , F_i , G , and $E[M_0]$, there exists $\alpha \in \{1, \infty\}$ such that

$$V(\text{MSB}_\alpha) \geq \frac{1}{2} V(\text{OMN}).$$

2) For any $\gamma, \epsilon > 0$, there exists N , F , G , and $E[M_0]$ such that if performance match values are drawn IID from F ,

$$V(\text{OPT}) < \left(\frac{1}{2} + \epsilon \right) V(\text{OMN}).$$

There are several possible interpretations of Proposition 3. In theoretical computer science, constant-factor approximations are often celebrated, and Proposition 3 states that no mechanism provides a better guarantee than that of MSB. However, many mechanisms (including the second-price auction) offer the same guarantee, so this argument does not justify the adoption of MSB auctions. Furthermore, the possibility of *doubling* match value is economically very meaningful. Thus, one possible conclusion from Proposition 3 is that worrying about the choice of mechanism is “barking up the wrong tree”: the real gains may come from better information about the common value component.

We offer a third interpretation: perhaps worst-case analysis is overly pessimistic and the distributions used in the proof of Proposition 3 are unrealistic in some way. In particular, it has been observed that in online advertising, a large fraction of the total value comes from a small number of very valuable impressions. That is, the distribution of advertiser values can be said to be heavy-tailed. Motivated by this thought, in Theorem 2, we consider the case where match values are drawn independently from a power law distribution. Under this assumption, the somewhat pessimistic conclusion of Proposition 3 reverses sharply.

THEOREM 2: *Suppose that N is deterministic and that the performance match values are IID draws from a power law distribution. Then for any $\gamma > 0$, there exists $\alpha \geq 1$ such that the following hold simultaneously:*

- $V_B(\text{MSB}_\alpha) = V_B(\text{OMN})$
- $V_P(\text{MSB}_\alpha) \geq 0.885 \cdot V_P(\text{OMN})$
- $V(\text{MSB}_\alpha) \geq 0.948 \cdot V(\text{OMN})$

Although the assumptions in Theorem 2 are strong, so are the conclusions. Suppose that the publisher contracts with the brand advertiser and commits to using an allocation rule z such that $V_B(z) \geq v_B$, for some $v_B \in (0, E[X_0])$. Subject to this constraint, the publisher aims to maximize the value of allocations to performance advertisers.⁶ Theorem 2 states that it is possible to choose α such that under MSB_α , the contract is fulfilled and performance advertisers get at least 88.5% of the value that could be delivered to performance advertisers *if the publisher directly observed C* . Furthermore, the third component of the theorem implies that if the value of the contract with the brand advertiser is chosen optimally, then an MSB auction delivers at least 94.8% of the value of the first-best solution. These results hold for any weight of the power law tail, expected brand value $E[X_0]$, distribution of C , and number of performance advertisers N . In particular, this implies that when match values follow a power law, MSB auctions are “nearly optimal” even in cases where C has a degenerate distribution (so that adverse selection is of no concern). Finally, our result that a second-price

⁶Although in Theorem 2 the publisher’s goal is to maximize $\gamma V_B(z) + V_P(z)$, there is a well-known equivalence between this approach and that of maximizing $V_P(z)$ subject to a lower-bound on $V_B(z)$.

auction may attain only 50% of the performance of OMN is directly comparable, because the examples discussed in Section III are ones in which match values follow a power law distribution. Taken together, these facts suggest that when match values are heavy tailed,

- 1) the benefit of moving from a second-price auction to an MSB auction may be significant,
- 2) protecting against adverse selection and shill bidding may have little cost, and
- 3) the potential gains from accurate information about C may be minimal.

We now turn to the question of how to choose α . Theorem 2 states that there exists a “good” choice of α , but offers no guidance on how to find it. This turns out to be straightforward. So long as the publisher knows the brand advertiser’s average value and the joint distribution of the top two performance bids (which can be learned, for example, using data from previous auctions), it is possible to compute, for any α , the value of the resulting allocation for both brand and performance advertisers. This allows the publisher to select α optimally. Alternatively, if the publisher has already promised the brand advertiser a certain fraction of impressions, α can simply be set to ensure that this guarantee is met. Importantly, these calculations *do not rely on assuming that performance match values are independent, identically distributed, or follow a particular distribution*, nor do they rely on *any* assumptions about the (unobserved) common value C : the only assumption is that C is independent from the vector of match values.⁷ By contrast, the optimal reserve in a second-price auction depends on the joint distribution of brand and performance advertiser values.

We now consider the revenue generated by MSB auctions. Although our primary focus in this paper is allocative efficiency, the publisher naturally cares about how this surplus is divided. In order to address this topic, we must posit a division of the surplus generated by impressions allocated to the brand advertiser – a point on which we have so far remained agnostic. We suppose that the publisher and brand advertiser split this surplus proportionately; that is, the revenue to the publisher is δV_B , for some $\delta \in (0, 1)$. It turns out that when match values follow a power law distribution with parameter a , for any strategy-proof mechanism z , publisher revenues from performance advertisers are at most $(1 - a^{-1})V_P(z)$ (see the Appendix for details). Thus, in this setting, maximizing revenue is effectively equivalent to maximizing $\delta V_B + (1 - a^{-1})V_P$. It follows that there exists an MSB auction that delivers at least 94.8% of the maximum revenue that could be achieved by *any* strategy-proof mechanism, even if C were observed. We state this below.

⁷In the power law case, the optimal choice is $\alpha = \max(1, \gamma E[X_0]/E[X_{(2)}])$.

COROLLARY 1: *Suppose that N is deterministic, that performance match values are IID draws from a power law distribution, and that the publisher and brand advertiser split surplus in any fixed proportion. Then there exists α such that MSB_α delivers at least 94.8% of the revenue of any strategy-proof auction.*

VII. Conclusion

In this paper, we consider the performance of certain auctions when bidder values are correlated and some bidders are uncertain of their values. We introduce modified second bid auctions as the only anonymous, deterministic mechanisms that, in our model, are fully-strategy-proof and free of adverse selection. We demonstrate that when match quality is distributed according to a power law distribution, MSB auctions may significantly improve upon more traditional second-price auctions and capture nearly all of the value obtained by an omniscient benchmark.

Our work makes two contributions: one methodological and one practical. On the methodological side, we believe that the “not-quite-optimal mechanism design” approach used in this paper is a generalizable and appealing way to derive and analyze new mechanisms. Our approach begins the same way as others in a long tradition in economics, that is, by identifying a candidate class of mechanisms as the only ones with certain desirable properties. These axioms, however, are not optimality properties. So, mechanism in hand, our approach evaluates the costs of our axioms by providing performance guarantees for efficiency and revenue relative to the unrestricted optimal mechanism.

On the practical side, we demonstrate that when allocating impressions jointly to advance contracts and spot market bidders, traditional reserve-based approaches may leave significant match value “on the table.” There may be notable gains to using a mechanism that adaptively sets a reserve based on submitted bids, as the MSB mechanism does. Furthermore, our results suggest that most of these gains can sometimes be realized by a mechanism that is simple to understand and implement.

Like all economic models, the one in this paper includes many simplifications. Although the MSB auction is quite generally fully strategy-proof, two of our other theoretical findings rely on statistical independence assumptions. First, the proof that MSB auctions are adverse-selection free uses the assumption that the brand advertiser’s value is statistically independent of a particular set of ratios – those between the values of different performance advertisers. Without that independence assumption, MSB auctions would not be adverse-selection free, although they would still be approximately so if the independence condition is approximately satisfied. Second, the good efficiency and revenue performance of the MSB compared to an omniscient mechanism relies on the assumed independence among performance advertiser match qualities. If performance advertisers selling similar products target the same users, that could lead to their match qualities to be positively correlated. Then, the signature that the MSB mechanism utilizes

to identify good matches – a single performance bid that is sufficiently larger than the others – might miss some good matches, degrading the mechanism’s performance.

The problem of achieving good matching when bidders’ values have both private and common quality components is not unique to the setting discussed in this paper. One closely related example is the problem of matching advertisers to opportunities on social media sites such as Facebook. Suppose that one advertiser launches a campaign targeting university students, and offers to pay \$1 for each click. Meanwhile, a second advertiser targets students in STEM fields, and offers \$1.25 for each click. If the two ads have the same estimated click-through rates, then all STEM students will be won by the second advertiser, leaving the first with an unrepresentative set of students. If, in fact, the first advertiser has a value of \$1.50 for students in STEM and \$0.50 for all other students, this outcome is inefficient. Rather than requiring every advertiser to estimate its value for each narrow subset of users, it may be possible to mitigate adverse selection and improve allocative efficiency by using mechanisms (such as MSB auctions) that score bids according to the breadth of their target audience. The question of how best to achieve this is one interesting direction for future work.

REFERENCES

- Abraham, Ittai, Susan Athey, Moshe Babaioff, and Michael Grubb.** 2013. “Peaches, Lemons, and Cookies: Designing Auction Markets with Dispersed Information.” *ACM Conference on Electronic Commerce (EC)*.
- Balseiro, Santiago R, Jon Feldman, Vahab Mirrokni, and S Muthukrishnan.** 2014. “Yield Optimization of Display Advertising with Ad Exchange.” *Management Science*, 60(12).
- Chen, Ying-Ju.** 2010. “Optimal Dynamic Auctions for Display Advertising.” *Conference on Information Systems and Technology (CIST)*.
- Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz.** 2007. “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords.” *American Economic Review*, 97(1): 242–259.
- Ghosh, Arpita, Preston McAfee, Kishore Papineni, and Sergei Vassilvitskii.** 2009. “Bidding for Representative Allocations for Display Advertising.” *Internet and Network Economics (WINE)*.
- Lewis, Randall A, and Justin M Rao.** 2015. “The Unfavorable Economics of Measuring the Returns to Advertising.” *The Quarterly Journal of Economics*, 130(4): 1941–1973.
- Milgrom, Paul R., and Robert J. Weber.** 1982. “A Theory of Auctions and Competitive Bidding.” *Econometrica*, 50(5).

- Varian, Hal R.** 2007. "Position Auctions." *International Journal of Industrial Organization*, 25(6): 1163 – 1178.
- Yokoo, Makoto, Yuko Sakurai, and Shigeo Matsubara.** 2004. "The effect of false-name bids in combinatorial auctions: new fraud in internet auctions." *Games and Economic Behavior*, 46: 174–188.