Bid-Limited Targeting

Patrick Hummel
Google Inc.
Mountain View, USA
phummel@google.com

Uri Nadav
Google Inc.
Mountain View, USA
urinadav@google.com

ABSTRACT

This paper analyzes a mechanism for selling items in auctions in which the auctioneer specifies a cap on the ratio between the maximum and minimum bids that bidders may use in the different auctions. Such a mechanism is widely used in online advertising through the caps that companies impose on the minimum and maximum bid multipliers that advertisers may use in targeting. When bidders' values are independent and identically distributed, using this mechanism results in higher revenue than allowing bidders to condition their bids on the targeting information in an arbitrary way and also almost always results in higher revenue than not allowing bidders to target. Choosing the optimal cap on the ratio between the maximum bid and the minimum bid can also be more important than introducing additional competition in the auction. However, if bidders' values are not identically distributed, pure-strategy equilibria may fail to exist.

CCS CONCEPTS

• **Applied computing** → **Online auctions**; **Economics**;

KEYWORDS

Targeting, Bid multipliers, Mechanism design, Online auctions

ACM Reference Format:

Patrick Hummel and Uri Nadav. 2018. Bid-Limited Targeting. In *WWW 2018: The 2018 Web Conference, April 23-27, 2018, Lyon, France*. ACM, New York, NY, USA, 10 pages.
<https://doi.org/10.1145/3178876.3186038>

1 INTRODUCTION

Good targeting is essential for effective online advertising. In a typical online auction, advertisers will have the ability to target their advertisements to different users by making different bids for advertising opportunities based on a user's gender, age, geographical location, and recent interests, auctions that differ by the time of day, day of the week, or month of the year, auctions that take place on different devices such as mobile, tablet, or desktop, etc. This ability to target enhances both user satisfaction and advertiser welfare by making ads more relevant to the user. But despite the

benefits of targeting, there is tremendous concern that targeting technology will ultimately become so effective that revenues from online advertising will decline.

The reason for these concerns is that if targeting enables one to always identify the one product that would be of interest to a user, then only one advertiser should be interested in advertising to that user, and that advertiser will be able to bid without competition. More generally, improved targeting enables advertisers to better identify advertising opportunities that would not be of interest to them, and thereby decreases the quantity of advertising demanded by advertisers as well as the amount of competition in the auctions. Indeed, various academic papers such as [6], [19], and [27] have identified circumstances under which improvements in targeting technology would decrease revenue from auctions.

What possible solutions might there be to the problem of improving targeting while ensuring that these improvements will not reduce revenue from online advertising? While there may be multiple possible solutions to this problem, such as only revealing sufficiently coarse information or only revealing targeting information in auctions with a sufficiently large amount of competition, in this paper we investigate one other possible solution to this problem. Would the auctioneer be able to solve this problem by tying the minimum bid that advertisers may make in auctions to the maximum bid that they make in other auctions? Imposing these constraints might solve this problem because if advertisers wish to be able to make large bids in some auctions, then they will also have to make non-negligible bids in all the auctions, thereby increasing competition in a typical auction.

Such a solution has appeal because the type of technology that is needed to implement such a solution is already established in practice. When companies implement targeting in practice, they do so by allowing advertisers to specify a base bid representing the typical amount that they will bid as well as a number of bid multipliers that specify by how much they would like to multiply their base bid if certain targeting criteria are satisfied [2]. For instance, a lipstick advertiser might make a base bid of \$0.25, and then set a bid multiplier of 2 if the user is known to be female, effectively making a bid of \$0.50 on females. Furthermore, major companies such as Google and Microsoft both impose caps on the minimum and maximum bid multipliers that advertisers may use.¹ Tying

This paper is published under the Creative Commons Attribution-NonCommercial-NoDerivs 4.0 International (CC BY-NC-ND 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW 2018, April 23-27, 2018, Lyon, France

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY-NC-ND 4.0 License.

ACM ISBN 978-1-4503-5639-8/18/04.

<https://doi.org/10.1145/3178876.3186038>

¹See, for example, <http://advertise.bingads.microsoft.com/en-us/blog/27543/upcoming-changes-to-bid-adjustments-in-the-yahoo-bing-network> for the caps that Microsoft imposes on the bid multipliers that advertisers may use in advertising on Bing and <https://support.google.com/adwords/answer/2732132> for the caps that Google imposes on the bid multipliers that can be used in its AdWords auction.

the maximum bid that advertisers may make in some auctions to the minimum bid that they make in other auctions is mathematically equivalent to imposing a cap on the bid multipliers that advertisers may use.

We consider a model in which an auctioneer imposes some cap $\alpha > 1$ that reflects the maximum bid multiplier that advertisers may use. Advertisers have prior beliefs about the distribution from which their values are drawn unconditional on the targeting information. The advertisers use this information to select the minimum bid that they will ever make, b . Selecting a minimum bid of b in turn means that the maximum bid that an advertiser may ever make is αb . When the advertisers ultimately bid in the auctions, they are able to condition their bids on the realization of the targeting data, but they are required to make bids between b and αb . Thus bidders who value the advertising opportunity for less than b will be required to make a bid of at least b , and bidders who value the advertising opportunity at more than αb will not be able to bid more than αb . Throughout we restrict attention to second-price auctions.

We characterize the equilibrium strategies that bidders would use in such an auction and show that if the bidders' values are independent draws from some distribution, then there is a symmetric equilibrium in which bidders all choose the same minimum bid b and then, conditional on the realization of the targeting data, follow the strategy of bidding b in the auction if their value is less than b , bidding their value if their value falls between b and αb , and bidding αb if their value is greater than αb .

We then analyze how revenue in this mechanism compares to revenue from other natural mechanisms. Throughout our revenue analysis, we focus on the situation in which each bidder's value is independent of that of the other bidders. This assumption is common in the literature on targeting (*e.g.* [3], [9], [15], [27]), and is reasonable in auctions for display advertising where information about the user such as gender, geographical information, or the user's interests affects the different advertisers differently. Furthermore, empirical evidence from Microsoft's Ad Exchange [8] indicates that there is hardly any correlation between bidders' values within an auction, which [8] indicates implies that "bidder valuations are private, driven by idiosyncratic match quality, rather than a common component." As such, our assumption of independent private values is one that has empirical backing from bidding behavior on a major ad exchange.

In our revenue analysis, we compare the performance of bid-limited targeting with several other benchmarks. The other benchmarks are a no targeting scenario, in which bidders bid in the auction without using targeting data, a full targeting scenario, in which bidders learn the realization of the targeting data before they place their bids and there are no restrictions on the bids they may make, and the full targeting scenario with the optimal reserve price.

When the bidders' values are independent and identically distributed, we find that bid-limited targeting performs favorably compared to these other benchmarks. We show that if advertisers' values are drawn from a continuous distribution,

then there always exists a value of the cap on the bid multiplier such that bid-limited targeting leads to higher revenue than full targeting. We also show that bid-limited targeting will lead to higher revenue than no targeting as long as there are at least three bidders in the auction and the bidders' values are random draws from a distribution satisfying the standard hazard rate condition, and that even when there are only two bidders, there are conditions under which bid-limited targeting leads to greater revenue than no targeting. The mechanism can also lead to higher revenue than setting the optimal reserve price, and choosing the optimal cap on the bid multiplier can increase revenue by more than adding an additional bidder to the auction.²

However, we also show that these conclusions hinge crucially on certain assumptions. When bidders' values are drawn from a bimodal distribution, setting the optimal reserve price may lead to significantly greater revenue than bid-limited targeting. And when the bidders' values are not identically distributed, we show that there typically will not exist a pure-strategy equilibrium to the game. The lack of pure-strategy equilibria with asymmetric bidders may make bid-limited targeting less suitable for practical applications.

There has been extensive work related to how improving targeting or giving bidders more precise information about their values affects revenue, but none of this work considers questions related to tying the minimum bid that bidders may make in some auctions to the maximum bid that bidders may make in other auctions. [1] and [24] have analyzed how revenue in common-value auctions is affected by the level of bidder information about the value of an object. [3], [4], [6], [9], [12], [14], [15], [16], [19], and [27] have analyzed similar questions for private-value auctions. [10] and [13] consider questions related to how information disclosure interacts with prices sellers may charge either for information or for their products. [5] and [17] consider questions related to markets for private data, which may be useful in targeting. [2] considers questions related to devising algorithms to appropriately set bid multipliers in bidding in online auctions, and [21] analyzes the price of anarchy in auctions with multiplicative bid adjustments. But none of these papers considers a model of targeting in which the minimum bid that bidders may make in some auctions explicitly constrains the maximum bid that bidders may make in other auctions.³

2 THE MODEL

We consider a model in which there are n bidders bidding simultaneously in an auction. Each bidder's value v_i is an independent draw from the cumulative distribution function $F_i(\cdot)$ with support equal to some subset of $[0, \infty)$. The bidders compete in a second-price auction. The types of strategies available to bidders depend on whether they are able to target

²This contrasts with results in [7] which illustrate that adding an additional bidder to the auction always increases revenue by more than choosing the optimal reserve price.

³Standard models of sponsored search auctions ([11], [18], and [29]) and budget constraints [22] also do not consider this possibility.

and whether the auctioneer imposes a cap on the maximum bid multiplier that bidders may use.

If bidders are not able to target, then when bidders submit their bids, they only know the distributions $F_i(\cdot)$ from which their values are drawn, and not their actual values v_i , so they can only condition their bids on the distributions $F_i(\cdot)$. If bidders are able to make full use of targeting information, then each bidder i learns his value v_i before bidding, and the bidders face no restrictions on how they can bid.

The final possibility we consider is a bid-limited targeting scheme in which bidders are given access to all available targeting information, but the auctioneer imposes a cap α on the maximum bid multiplier that bidders may use. In this case, each bidder i must commit to a minimum bid b_i before bidder i learns his value. This minimum bid b_i represents the lowest possible bid that bidder i will ever make in the auction. Bidders then learn the realizations of their values before bidding in the auction,⁴ but if bidder i commits to a minimum bid of b_i , then the bidder will not be allowed to submit a bid higher than αb_i in the auction. Thus under bid-limited targeting, each bidder i first chooses a minimum bid b_i before learning the realization of his value and then makes a final bid in the auction that is chosen from the interval $[b_i, \alpha b_i]$ after learning the realization of his value.⁵ The bidders choose their minimum bids b_i simultaneously and thus do not know the minimum bids chosen by the other bidders at the time they choose their own minimum bids.

3 EQUILIBRIUM PROPERTIES

When bidders are not able to target, all bidders have an incentive to submit bids equal to their expected values, and revenue equals the second-highest expected value. Similarly, when bidders can make full use of targeting information, all bidders have a dominant strategy of making a bid equal to their value, and revenue is the second-highest realized value.

Also, under bid-limited targeting, once a bidder has chosen a minimum bid b_i , the optimal bidding strategy for the bidder in the second stage of the auction is simple. If the bidder's value v_i satisfies $v_i \in [b_i, \alpha b_i]$, then the bidder has a weakly dominant strategy of bidding his value. If the bidder's value v_i satisfies $v_i < b_i$, then the bidder would like to submit a bid of v_i but is required to make a bid of at least b_i , so the bidder does the next best thing and makes a bid of b_i . Similarly, if the bidder's value v_i satisfies $v_i > \alpha b_i$, then the bidder makes a bid of αb_i . Thus the only interesting decision for a bidder is the choice of the minimum bid b_i , as the rest of the bidder's strategy is completely determined from this.

⁴Under standard bid-limited targeting in online auctions, bidders are able to condition their bids on the realization of the targeting information, effectively enabling them to condition their bids on their values for the advertising opportunity, as we assume in this model.

⁵In practice an advertiser i would choose a base bid b_i and a number of bid multipliers $m(v_i)$ that depend on the advertiser's value for the advertising opportunity. Then if the advertiser wanted to make a bid $b(v_i)$ when the advertiser had a value v_i , the advertiser would set a bid multiplier of $m(v_i) = b(v_i)/b_i$. Bid-limited targeting only imposes the additional restriction that $m(v_i) \in [1, \alpha]$ for all v_i .

In analyzing the equilibrium in the first bidding stage of bid-limited targeting, we first consider the case in which each bidder's value is an independent and identically distributed draw from the same distribution $F(\cdot)$. In this case, we search for a symmetric equilibrium in which all the bidders use the same minimum bid b . If such an equilibrium exists, then it must be the case that no bidder can profitably deviate by instead using a lowest possible bid of $b + \epsilon$ for some $\epsilon > 0$. Changing a bidder's lowest possible bid from b to $b + \epsilon$ would affect the bidder's payoff in two ways:

First, it would be more likely that a bidder with a value $v_i < b$ will win an advertising opportunity that the bidder did not want to win. In particular, if a bidder with value $v_i \leq b$ uses a lowest possible bid of $b + \epsilon$ for some $\epsilon > 0$, then the bidder will always win any advertising opportunities where the other $n - 1$ bidders have values $v \leq b$, but this bidder would have only won these advertising opportunities with probability $\frac{1}{n}$ before deviating.

Second, it would be more likely that a bidder with a value $v_i > \alpha b$ will win an advertising opportunity that the bidder wants to win. A bidder who deviates to a lowest possible bid of $b + \epsilon$ for some $\epsilon > 0$ will always win any advertising opportunities where the bidder has a value of $v_i > \alpha b$. However, this bidder would have only won these advertising opportunities with probability $\frac{1}{k+1}$ if there were k other bidders who also had a value of $v_i > \alpha b$ previously.

Putting all this together, we see that if all bidders have values $v \leq b$ (which happens with probability $F(b)^n$), then a bidder increases his probability of winning an advertising opportunity by $\frac{n-1}{n}$ by changing his minimum bid from b to $b + \epsilon$, and in expectation this costs the bidder $b - E[v|v < b]$ in these cases. But if a bidder has value $v_i \geq \alpha b$ and exactly k other bidders have values $v \geq \alpha b$ (which happens with probability $\binom{n-1}{k} F(\alpha b)^{n-k-1} (1 - F(\alpha b))^{k+1}$), then the bidder increases his probability of winning an advertising opportunity by $\frac{k}{k+1}$ by changing his minimum bid from b to $b + \epsilon$, and in expectation the bidder's utility increases by $E[v|v > \alpha b] - \alpha b$ in these circumstances. From this it follows that a necessary condition for there to exist a symmetric equilibrium in which each bidder uses a minimum bid of b is that the following inequality must be satisfied:

$$\frac{n-1}{n} F(b)^n [b - E[v|v < b]] \geq \sum_{k=1}^{n-1} \frac{k}{k+1} \binom{n-1}{k} F(\alpha b)^{n-k-1} (1 - F(\alpha b))^{k+1} [E[v|v > \alpha b] - \alpha b] \quad (1)$$

Similarly, if there is an equilibrium in which all bidders use a minimum bid of b , it must be that no bidder can profitably deviate by instead using a minimum bid of $b - \epsilon$. Similar analysis to that used to derive equation (1) then shows that another necessary condition for there to exist a symmetric equilibrium in which each bidder uses a minimum bid of b is

the following:

$$\frac{1}{n}F(b)^n [b - E[v|v < b]] \leq$$

$$\sum_{k=1}^{n-1} \frac{1}{k+1} \binom{n-1}{k} F(\alpha b)^{n-k-1} (1 - F(\alpha b))^{k+1} [E[v|v > \alpha b] - \alpha b]$$

Now note that the left-hand side of equation (2) is exactly $\frac{1}{n-1}$ as large as the left-hand side of equation (1), while the right-hand side of equation (2) is at least as large as the right-hand side of equation (1) multiplied by $\frac{1}{n-1}$. From this it follows that if equation (1) is satisfied with equality, then the inequality in equation (2) will also be satisfied. Thus a sufficient condition for it to be an equilibrium for all bidders to use a minimum bid of b is the following:

$$\frac{n-1}{n}F(b)^n [b - E[v|v < b]] =$$

$$\sum_{k=1}^{n-1} \frac{k}{k+1} \binom{n-1}{k} F(\alpha b)^{n-k-1} (1 - F(\alpha b))^{k+1} [E[v|v > \alpha b] - \alpha b]$$

Our first theorem illustrates that there is a symmetric equilibrium in which bidders make use of a minimum bid b that satisfies equation (3). Moreover, this is the lowest possible value of b that can arise in a symmetric equilibrium.

THEOREM 3.1. *If the bidders' values are independent and identically distributed, then there is a symmetric equilibrium in which bidders use a minimum bid b that satisfies equation (3). This is the unique symmetric equilibrium if there are two bidders. Moreover, any other symmetric equilibrium will necessarily entail using a minimum bid b that is strictly greater than the unique b satisfying equation (3).*

PROOF. We prove that there exists a symmetric equilibrium in which bidders use a minimum bid b that satisfies equation (3) here and defer the proof of the remainder of the theorem to the full version of the paper [20]. To see this, first note that if b is the lower bound of the support of $F(\cdot)$, then the left-hand side of equation (3) is zero, while the right-hand side is positive. And if b is $\frac{1}{\alpha}$ times the upper bound of the support of $F(\cdot)$ (which may be infinite), then the right-hand side of equation (3) is zero, while the left-hand side is strictly positive. Since both the left-hand side and the right-hand side of equation (3) vary continuously with the parameters, it follows from the intermediate value theorem that there exists a value of b that is a solution to this equation. Thus there exists a symmetric equilibrium in which bidders use such a minimum bid. $\square$

While there are potentially multiple symmetric equilibria, the equilibrium in Theorem 3.1 in which bidders all use a minimum bid b satisfying equation (3) is the most natural equilibrium to focus on because this is the equilibrium in which bidders use the smallest possible minimum bid b , and thus enables the bidders to win the advertising opportunities as cheaply as possible.⁶ Since this is the best equilibrium for

⁶We formally prove that this is the best equilibrium for the advertisers in the full version of the paper [20].

the advertisers, we focus on this equilibrium throughout the remainder of the paper. Next we illustrate how the players' equilibrium strategies vary with the parameters:

LEMMA 3.2. *The equilibrium value of b satisfying equation (3) is decreasing in α .*

We defer the proofs of most results to the full version of the paper [20]. To understand this result, note that when α is relatively larger, it is relatively less likely that multiple bidders will have values greater than αb , while the probability that all n bidders have values less than b remains unchanged. Thus increasing α decreases the benefits from using a slightly larger minimum bid, while having no effect on the costs from using such a minimum bid. From this it follows that if α is relatively larger, then bidders will have an incentive to use a relatively lower value of b in equilibrium.⁷ A similar analysis can illuminate how the equilibrium minimum bid varies with the number of players.

LEMMA 3.3. *For fixed α , the equilibrium value of b satisfying equation (3) is increasing in n .*

The fact that the equilibrium value of b is increasing in n is also intuitive. When there are a larger number of bidders in an auction, if the value of b stays the same, it will be relatively less likely that all n bidders will have values less than b , and relatively more likely that multiple bidders will have values greater than αb . Thus it will be relatively less likely that a bidder with value $v \leq b$ will be forced to win an advertising opportunity that this bidder did not want to win, and it will be relatively more likely that a bidder with value $v \geq \alpha b$ will fail to win an advertising opportunity that he really wanted to win. Since a larger number of bidders means the costs from using a slightly larger minimum bid are relatively lower and the benefits from using a slightly larger minimum bid are relatively higher, the equilibrium minimum bid b is increasing in the number of bidders n .

With these results in mind, we now turn to the question of how bid-limited targeting compares to full targeting and to not allowing the bidders to make use of targeting information when bidders' values are independent and identically distributed.

4 COMPARISON WITH NO TARGETING AND FULL TARGETING

We begin our comparison of bid-limited targeting to other mechanisms by studying how bid-limited targeting compares to not allowing the bidders to target at all. First we note how general changes in the number of bidders affects the relative attractiveness of bid-limited targeting compared to not allowing the bidders to target at all.

THEOREM 4.1. *Suppose that bid-limited targeting leads to higher revenue than no targeting for some number of bidders*

⁷Theorem 3.1 and Lemma 3.2 can also be proven under the weaker assumption that bidders' values are independent and identically distributed draws from some unknown distribution.

n. Then bid-limited targeting also leads to higher revenue than no targeting for any larger number of bidders (for either the same α or the optimally chosen α).

PROOF. We have seen in Lemma 3.3 that for fixed α , the equilibrium value of b in bid-limited targeting is increasing n , so the range of bids that bidders may submit is strictly greater where there are more bidders. In addition, the distribution of the second-highest value when there is a larger number of bidders strictly first order stochastically dominates this distribution when there are fewer bidders. Thus the auctioneer's revenue from bid-limited targeting is strictly increasing in n if α remains unchanged (and is therefore also increasing in n for the optimally chosen α).

However, the revenue from no targeting is independent of the number of bidders since bidders always submit bids equal to $E[v]$ under no targeting. From this it follows that if bid-limited targeting leads to higher revenue than no targeting for some fixed number of bidders n , then bid-limited targeting also leads to higher revenue than no targeting for any larger number of bidders. $\square$

This result parallels findings in [19] which illustrate that if an auctioneer prefers full targeting to no targeting for some small number of bidders, then the auctioneer also prefers full targeting to no targeting for larger numbers of bidders. We now seek to obtain detailed results on when bid-limited targeting is preferred to no targeting as a function of the number of bidders. First we consider what happens when there are exactly two bidders:

THEOREM 4.2. *When there are two bidders, the auctioneer obtains greater revenue from bid-limited targeting with α slightly greater than 1 than by not allowing the bidders to target at all if and only if $F(E[v]) < \frac{1}{2}$.*

PROOF. In the limit as $\alpha \rightarrow 1$, the lowest possible value of b that the bidders use approaches the value of b satisfying $F(b)^2[b - E[v|v < b]] = (1 - F(b))^2[E[v|v > b] - b]$. This expression can instead be rewritten as $[F(b)^2 + (1 - F(b))^2]b = F(b)^2E[v|v < b] + (1 - F(b))^2E[v|v > b]$ or $b = \frac{F(b)^2}{F(b)^2 + (1 - F(b))^2}E[v|v < b] + \frac{(1 - F(b))^2}{F(b)^2 + (1 - F(b))^2}E[v|v > b]$.

Now we know that $F(b)E[v|v < b] + (1 - F(b))E[v|v > b] = E[v]$. Thus if $w < F(b)$, then $wE[v|v < b] + (1 - w)E[v|v > b] > E[v]$, and if $w > F(b)$, then $wE[v|v < b] + (1 - w)E[v|v > b] < E[v]$. We also know that if $b = E[v]$, then $\frac{F(b)^2}{F(b)^2 + (1 - F(b))^2} < F(b)$ holds if and only if $F(E[v]) < \frac{1}{2}$. Thus if $F(E[v]) < \frac{1}{2}$, then $\frac{F(b)^2}{F(b)^2 + (1 - F(b))^2}E[v|v < b] + \frac{(1 - F(b))^2}{F(b)^2 + (1 - F(b))^2}E[v|v > b] > E[v]$ when $b = E[v]$, so $b > E[v]$ must hold in equilibrium. Similarly, if $F(E[v]) > \frac{1}{2}$, then $\frac{F(b)^2}{F(b)^2 + (1 - F(b))^2}E[v|v < b] + \frac{(1 - F(b))^2}{F(b)^2 + (1 - F(b))^2}E[v|v > b] < E[v]$ when $b = E[v]$, so $b < E[v]$ must hold in equilibrium. Finally, if $F(E[v]) = \frac{1}{2}$, then $\frac{F(b)^2}{F(b)^2 + (1 - F(b))^2}E[v|v < b] + \frac{(1 - F(b))^2}{F(b)^2 + (1 - F(b))^2}E[v|v > b] = E[v]$ when $b = E[v]$, so $b = E[v]$ in equilibrium.

But when the bidders are not allowed to target, the auctioneer's revenue is $E[v]$. And when the auctioneer employs bid-limited targeting with a value of α infinitesimally greater than 1, the auctioneer's revenue is arbitrarily close to the equilibrium minimum bid b , since all bidders submit bids arbitrarily close to b . From this and the results in the previous paragraph, it follows that the auctioneer's revenue is greater under bid-limited targeting with α arbitrarily close to 1 than under no targeting if and only if $F(E[v]) < \frac{1}{2}$. $\square$

The fact that bid-limited targeting can potentially lead to larger revenue than no targeting when there are two bidders stands in contrast to results comparing full targeting to no targeting when there are two bidders. [6], [19], and [27] have noted that when there are two bidders and there are no restrictions on the strategies that bidders may use, revealing information about the bidders' values will always hurt the auctioneer's revenue. By contrast, the results in Theorem 4.2 indicate that under bid-limited targeting, the auctioneer may be able to increase her revenue by allowing bidders to target if the auctioneer imposes significant restrictions on the bid multipliers that bidders may use.⁸

It is worth noting that the revenue differences in Theorem 4.2 can be substantial. When α is close to 1, bidders in bid-limited targeting will all make bids in some narrow interval, but the bids in this interval can be substantially lower (or higher) than the fixed bid that bidders would make under no targeting. For example, if there are two bidders and $F(E[v]) < \frac{1}{2}$, then if the bidders use a minimum bid of $b = E[v]$, the probability a bidder with value $v < b$ will be forced to win an advertising opportunity it does not want to win (because both bidders bid b) will be less than the probability a bidder with value $v > b$ will be unable to win an advertising opportunity that the bidder wants to win (because both bidders bid αb). Thus a bidder will have an incentive to raise its minimum bid above $b = E[v]$ because this will do more to increase the bidder's probability of winning advertising opportunities the bidder wants to win than it will to hurt the bidder's probability of winning advertising opportunities the bidder does not want to win.

Using bid-limited targeting with α close to 1 also results in greater bidder surplus than no targeting. If bidders are not allowed to target, then both bidders make bids equal to their expected values and expected bidder surplus is zero. By contrast, under bid-limited targeting, bidders are more likely to win the auction if they have larger values, so bidders are able to achieve a positive expected surplus. This also implies that allowing endogenous entry would only strengthen the circumstances under which bid-limited targeting is preferred to no targeting with two potential entrants.

While not allowing the bidders to target frequently leads to higher revenue than bid-limited targeting when there are two bidders, when there are at least three bidders, much

⁸Nonetheless, the fact that revenue from bid-limited targeting can be lower than revenue from no targeting when $F(E[v]) > \frac{1}{2}$ implies that the seller could still have an incentive to reduce the information provided to bidders if there is little competition in the auction.

milder conditions guarantee that bid-limited targeting will lead to higher revenue than no targeting. To illustrate this, it is helpful to first present the following lemma:

LEMMA 4.3. $F(E[v]) \leq 1 - \frac{1}{e}$ for any distribution $F(\cdot)$ with a non-decreasing hazard rate.

We now use this lemma to prove the following result:

THEOREM 4.4. *Suppose there are at least three bidders and the bidders' values are drawn from a distribution with non-decreasing hazard rate. Then there exists some parameter α under bid-limited targeting such that bid-limited targeting will lead to greater revenue than no targeting.*

PROOF. From Theorem 4.1, it suffices to show that when there are exactly three bidders and the bidders' values are drawn from a distribution with non-decreasing hazard rate, then the auctioneer's expected revenue under bid-limited targeting in the limit as $\alpha \rightarrow 1$ is greater than auctioneer's expected revenue from not allowing the bidders to target.

When there are three bidders, in the limit as $\alpha \rightarrow 1$, we know from the equilibrium condition in equation (3) that the equilibrium value of b approaches the solution to the equation $\frac{2}{3}F(b)^3[b - E[v|v < b]] = (1 - F(b))^2(\frac{2}{3} + \frac{1}{3}F(b))[E[v|v > b] - b]$, meaning $b[2F(b)^3 + (1 - F(b))^2(2 + F(b))] = 2F(b)^3E[v|v < b] + (1 - F(b))^2(2 + F(b))E[v|v > b]$, and

$$b = \frac{2F(b)^3}{2F(b)^3 + (1 - F(b))^2(2 + F(b))}E[v|v < b] + \frac{(1 - F(b))^2(2 + F(b))}{2F(b)^3 + (1 - F(b))^2(2 + F(b))}E[v|v > b]. \quad (4)$$

Now we know that $F(b)E[v|v < b] + (1 - F(b))E[v|v > b] = E[v]$. Thus if $w < F(b)$, then $wE[v|v < b] + (1 - w)E[v|v > b] > E[v]$. But

$$\frac{2F(b)^3}{2F(b)^3 + (1 - F(b))^2(2 + F(b))} < F(b) \quad (5)$$

holds if and only if $2F(b)^2 < 2F(b)^3 + (1 - F(b))^2(2 + F(b))$. Computationally it follows that this inequality holds if and only if $F(b) < \frac{2}{3}$. But from this and the fact in Lemma 4.3 that $F(E[v]) \leq 1 - \frac{1}{e} = 0.632$ for any distribution $F(\cdot)$ with non-decreasing hazard rate, it follows that if $b = E[v]$, then the inequality in equation (5) holds, and the right-hand side of equation (4) is greater than $E[v]$. Thus in order to satisfy the condition in (4), it must be the case that $b > E[v]$.

But when the bidders are not allowed to target, the auctioneer's expected revenue is $E[v]$. And when the auctioneer employs bid-limited targeting with a value of α infinitesimally greater than 1, the auctioneer's revenue is arbitrarily close to the equilibrium minimum bid b . From this and the results in the previous paragraph, it follows that if there are at least three bidders whose values are drawn from a distribution with non-decreasing hazard rate, then the auctioneer obtains greater revenue by employing bid-limited targeting with α slightly greater than 1 than by not allowing the bidders to target at all. $\square$

This result is again in contrast to existing results comparing full targeting to no targeting when there are three bidders. [19] has noted that with three bidders, the auctioneer's revenue from no targeting is higher than the auctioneer's revenue from full targeting as long as the bidders' values are drawn from a distribution with density $f(v)$ that is decreasing in v . In fact, [19] notes that no targeting can be preferred to full targeting even when the auctioneer sets the optimal reserve price under full targeting.

By contrast, Theorem 4.4 guarantees that as long as bidders' values are drawn from a distribution with non-decreasing hazard rate, bid-limited targeting will be preferred to not allowing the bidders to target. Between this result and Theorem 4.2, we see that by imposing caps on the maximum bid multipliers that bidders can use, an auctioneer can allow the bidders to make use of targeting information in situations where this would ordinarily be unprofitable.

It is worth noting, however, that the non-decreasing hazard rate assumption in Theorem 4.4 is crucial to obtain this result. If the bidders' values are drawn from a distribution with sufficiently fat tails, it need not be the case that bid-limited targeting is preferred to no targeting, even when there are a large number of bidders. In fact, our next result shows that for *any* number of bidders, it is possible to find some distribution with sufficiently fat tails such that no targeting is preferred to bid-limited targeting.

THEOREM 4.5. *For any number of bidders, n , there exists some distribution such that if the bidders' values are independent random draws from that distribution, then the auctioneer obtains greater revenue by not allowing the bidders to target than by employing any bid-limited targeting scheme.*

To understand the intuition behind this result, suppose that bidders' values are drawn from a skewed distribution such that a bidder is very likely to have a small value but in the unlikely event that the bidder has a high value, the bidder's value is extremely high. In this case, if the bidders are not able to target, then all bidders will make a moderate bid equal to their expected values, and the auctioneer will obtain a moderate amount of revenue. But if the bidders do learn the realization of their values, it is very likely that the second-highest of the bidders' values will be small. Moreover, under bid-limited targeting, the bidders will have an incentive to use small minimum bids because it is much more likely that they will be faced with a situation in which they are forced to win an advertising opportunity that they did not want to win than it is that they will be unable to win an advertising opportunity that they wanted to win. Thus even when there are a large number of bidders, if the bidders' values are drawn from a distribution with sufficiently fat tails, the auctioneer may prefer not to allow the bidders to target than to use bid-limited targeting.

It is worth noting that the result in Theorem 4.5 is not due to bid-limited targeting being an inferior way to allow bidders to use targeting information. When bidders' values are drawn from a distribution with fat tails, allowing bidders to make full use of targeting information will also lead to

lower revenue than no targeting. In fact, allowing bidders to make full use of targeting information can never lead to higher revenue than bid-limiting targeting with an appropriately chosen cap on the bid multiplier:

THEOREM 4.6. *For any number of players $n \geq 2$ and any distribution $F(\cdot)$ with compact support and strictly positive density throughout its support, there exists some parameter α under bid-limited targeting such that bid-limited targeting leads to greater revenue than full targeting.*

The reason for this result is the auctioneer’s revenue under full targeting is the same as the auctioneer’s revenue under bid-limited targeting in the limit as $\alpha \rightarrow \infty$. But for large values of α , the auctioneer can always improve her expected revenue by slightly reducing the maximum bid multiplier the bidders can use. From this it follows that the auctioneer will always do better by using bid-limited targeting with an appropriately chosen bid multiplier than by allowing the bidders to make full use of targeting information.

4.1 Numerical Results

The above theorems all indicate that bid-limited targeting will typically perform at least as well as either no targeting or full targeting, but they do not give a sense of the size of the improvements. We now present some numerical results to address this. We consider both the case in which bidders’ values are independent and identically distributed draws from the uniform distribution and the case in which bidders’ values are drawn from a lognormal distribution with mean and standard deviation equal to $\frac{1}{2}$. The lognormal distribution is of particular interest since [23], [26], and [28] have noted that bidders’ values can be well modeled as being drawn from a lognormal distribution in their studies of sponsored search auctions at Yahoo! and Baidu.⁹

n	α	No targeting	Full targeting	Reserve price
2	1.00	0.0%	50.0%	20.0%
3	1.00	13.5%	13.5%	6.8%
4	1.22	23.3%	2.8%	0.7%
5	1.88	33.9%	0.4%	-0.4%
6	2.45	42.9%	0.1%	-0.3%

Table 1: Optimal α and percentage improvement in revenue from using the optimal bid-limited targeting mechanism over no targeting, full targeting, and full targeting with the optimal reserve price as a function of the number of bidders (n) when each bidder’s value is an independent and identically distributed draw from the uniform distribution.

Tables 1 and 2 indicate that bid-limited targeting leads to substantially higher revenue than no targeting as long as there

⁹[26] also explicitly focuses on the case in which the lognormal distribution has mean and standard deviation of $\frac{1}{2}$ in the context of work on sponsored search auctions at Yahoo!.

n	α	No targeting	Full targeting	Reserve price
2	1.00	-12.4%	57.5%	39.0%
3	1.00	0.5%	21.7%	17.9%
4	1.00	10.8%	7.8%	6.7%
5	1.95	20.8%	1.4%	1.0%
6	4.18	33.5%	0.2%	0.0%
7	7.30	45.6%	0.0%	-0.0%

Table 2: Optimal α and percentage improvement in revenue from using the optimal bid-limited targeting mechanism over no targeting, full targeting, and full targeting with the optimal reserve price as a function of the number of bidders (n) when each bidder’s value is an independent and identically distributed draw from the lognormal distribution with mean and standard deviation equal to $\frac{1}{2}$.

are at least three to four bidders. Bid-limited targeting also leads to substantially higher revenue than full targeting when there are a small number of bidders, but the improvement from using bid-limited targeting becomes small for larger numbers of bidders.

One other curious finding in Tables 1 and 2 is that bid-limited targeting can sometimes lead to higher revenue than the Myerson optimal auction [25] where full targeting is accompanied by the optimal reserve price. The reason this finding does not contradict Myerson’s results is that [25] restricts attention to mechanisms that satisfy *ex post* individual rationality in the sense that all bidders who participate in the auction must be able to guarantee themselves a non-negative expected payoff by participating. By contrast, bid-limited targeting will only satisfy *ex ante* individual rationality since bidders are required to commit to a minimum bid before they learn their values, and a bidder who ends up having a really low value will still have to make this minimum bid and thereby incur a negative expected payoff.¹⁰ For this reason, bid-limited targeting can, and sometimes does, result in higher revenue than using the optimal reserve price.

But there are still settings in which bid-limited targeting leads to considerably lower revenue than setting the optimal reserve price. In Table 3 we analyze a setting in which bidders’ values are drawn from a bimodal distribution in which there is a high probability that a bidder will have a low value and a low probability that a bidder will have a much higher value. While bid-limited targeting still leads to greater revenue than full targeting in this setting, bid-limited targeting also typically leads to much lower revenue than setting the optimal reserve price. Thus when bidders’ values are drawn from such a bimodal distribution, bid-limited targeting is ineffective compared to setting a reserve price.

¹⁰This weaker condition of *ex ante* individual rationality is realistic in the online advertising setting since advertisers bid in millions of auctions a day and do not get feedback on individual impressions but instead only seek to do as well as possible in aggregate.

n	α	No targeting	Full targeting	Reserve price
2	1.00	-36.0%	35.2%	-2.8%
3	1.00	-27.9%	5.4%	-25.3%
4	13.36	-13.1%	1.5%	-30.8%
5	11.70	2.1%	0.8%	-33.5%
6	13.56	16.3%	0.1%	-35.3%
7	15.39	30.8%	0.0%	-36.2%
8	17.51	45.1%	0.0%	-36.7%

Table 3: Optimal α and percentage improvement in revenue from using the optimal bid-limited targeting mechanism over no targeting, full targeting, and full targeting with the optimal reserve price as a function of the number of bidders (n) when each bidder's value is an independent and identically distributed draw from the distribution $F(\cdot) = (1 - \pi)F_L(\cdot) + \pi F_H(\cdot)$ with $\pi = 0.05$, where $F_L(\cdot)$ represents a lognormal distribution with mean 1 and variance $\frac{1}{2}$ and $F_H(v) = F_L(v - \Delta)$ for $\Delta = 10$.

Finally, we note that while bid-limited targeting can lead to higher revenue than the optimal *ex post* individually rational mechanism, it will not lead to higher revenue than the optimal *ex ante* individually rational mechanism. When bidders' values are independent draws from the same distribution, the optimal *ex ante* individually rational mechanism will involve charging bidders an entry fee equal to a bidder's *ex ante* expected payoff from participating in a second-price auction with n bidders and then running a standard second-price auction. Nonetheless, there is still a sense in which bid-limited targeting compares favorably to the optimal *ex ante* individually rational mechanism. If the mechanism designer uses the optimal *ex ante* individually rational mechanism for a fixed number of players, and additional bidders arrive at some future time, these bidders will have an incentive to not pay the entry fee required to participate in the auction, because doing so would lead to a negative payoff. By contrast, if additional bidders arrive at some future time under bid-limited targeting, these bidders will still be able to profitably participate in the auction. Thus bid-limited targeting is more robust to not discouraging additional bidders to participate in the auction in the future should such bidders arrive.

5 OTHER PROPERTIES

In this section we address questions related to the optimal parameters in bid-limited targeting. Normally the value to choosing the optimal auction is less than the value of getting an additional bidder to participate, as [7] has noted that choosing the optimal reserve price increases the auctioneer's revenue by less in expectation than getting another bidder to participate in the auction. However, we have seen previously that revenue under bid-limited targeting can exceed revenue under the Myerson optimal auction [25], so it seems possible that choosing the optimal bid multiplier could increase

revenue by more than adding another bidder to the auction. This is in fact the case, as the following theorem illustrates:

THEOREM 5.1. *Adding an additional bidder to the auction may increase revenue by less than switching from full targeting to bid-limited targeting with the optimal α .*

PROOF. Suppose the bidders' values are independent draws from the exponential distribution with probability density function $f(v) = e^{-v}$ and cumulative distribution function $F(v) = 1 - e^{-v}$. To prove the result it suffices to show that revenue under bid-limited targeting when there are $n = 2$ bidders in the limit as $\alpha \rightarrow 1$ is greater than revenue under full targeting when there are $n = 3$ bidders.

We have seen previously that if $n = 2$ bidders participate under bid-limited targeting, then in the limit as $\alpha \rightarrow 1$, the equilibrium minimum bid b is given by the solution to $F(b)^2[b - E[v|v < b]] = (1 - F(b))^2[E[v|v > b] - b]$. Now for the exponential distribution, we have $F(b)^2 = (1 - e^{-b})^2$, $E[v|v < b] = \frac{\int_0^b x e^{-x} dx}{1 - e^{-b}} = \frac{1 - (b+1)e^{-b}}{1 - e^{-b}}$, $(1 - F(b))^2 = e^{-2b}$, and $E[v|v > b] - b = 1$. Thus for the exponential distribution, $F(b)^2[b - E[v|v < b]] = (1 - F(b))^2[E[v|v > b] - b]$ is satisfied if and only if $(1 - e^{-b})(b - 1 + e^{-b}) = e^{-2b}$. Computationally it follows that this equation is satisfied when $b = 0.87942$. And since revenue under bid-limited targeting in the limit as $\alpha \rightarrow 1$ equals b , it follows that the auctioneer's revenue under bid-limited targeting with $n = 2$ bidders in the limit as $\alpha \rightarrow 1$ approaches 0.87942.

But when there are $n = 3$ bidders, the auctioneer's revenue under full targeting is $\int_0^\infty (v - \frac{1 - F(v)}{f(v)}) n F(v)^{n-1} f(v) dv = \int_0^\infty (v - 1) n F(v)^{n-1} f(v) dv = -(v - 1)(1 - F(v)^n)|_0^\infty + \int_0^\infty 1 - F(v)^n dv = -1 + \int_0^\infty 1 - (1 - e^{-v})^3 dv = \frac{5}{6} < 0.87942$. Thus revenue under full targeting when there are $n = 3$ bidders is less than revenue under bid-limited targeting when there are $n = 2$ bidders in the limit as $\alpha \rightarrow 1$. $\square$

While optimally selecting the value of α can be quite important, there is little benefit to using a reserve price in conjunction with bid-limited targeting. If the reserve price is lower than the minimum bid the bidders use in equilibrium, then the reserve price has no effect. And if the reserve price is higher than the minimum bid the bidders would have used in the absence of a reserve price, then bidders will instead use a minimum bid that is infinitesimally lower than the reserve price, and the seller will only obtain a lower revenue than she would have by running a standard second-price auction with a reserve price. Thus using reserve prices in conjunction with bid-limited targeting is uninteresting.

6 ASYMMETRIC BIDDERS

Throughout the paper so far, we have focused on cases in which the bidders' values are all random draws from the same distribution, but it is also important to address what would happen if the bidders' values were drawn from different distributions. Specifically, we consider what would happen if each bidder i has a value that is an independent random

draw from the cumulative distribution function $F_i(v)$, where the distributions $F_i(v)$ are different for the different players. We also suppose that each distribution $F_i(v)$ has an interval support of the form $[v_i, \bar{v}_i]$ for some v_i and $\bar{v}_i$. Unfortunately, in this setting pure-strategy equilibria need not exist:

THEOREM 6.1. *Suppose the auctioneer uses a value of α under bid-limited targeting such that there are at least two bidders i for which $\frac{\bar{v}_i}{v_i} > \alpha$ and exactly two such bidders i and j for which $\bar{v}_i > v_j$ and $\bar{v}_j > v_i$ and both of these players have a strictly positive probability of having the highest value amongst all the bidders. Then generically there is no pure-strategy equilibrium under bid-limited targeting.*

The reason for this result is that it is never possible for there to be a pure-strategy equilibrium in which two different bidders use different minimum bids because if different bidders use different minimum bids, the bidder with the lower minimum bid would be able to profitably deviate by raising his minimum bid. Such a deviation would not increase this bidder's risk of being forced to win advertising opportunities that he would not want to win, since his minimum bid is already sufficient to be able to guarantee that he can lose any auction that he wants to lose. But the increased minimum bid would enable the bidder to win more advertising opportunities where he has a high value. Thus this would be a profitable deviation.

From this it follows that if there exists a pure-strategy equilibrium, it must be the case that the different asymmetric bidders use the same minimum bids. But generically it cannot be the case that players whose values are drawn from different distributions simultaneously have their incentive constraints satisfied when they use the same minimum bid. Thus pure-strategy equilibria generically fail to exist when there are multiple asymmetric bidders. Because of this, it is difficult to make analytic statements about revenue under bid-limited targeting with asymmetric bidders.¹¹

7 DROPPING OUT OF AUCTIONS

In this section we analyze an extension of the main bid-limited targeting mechanism that we have presented. Throughout we have assumed that once a bidder submits a minimum bid, this bidder is required to submit a bid in all auctions. However, another plausible model would be one in which a bidder is allowed to not bid in certain auctions. We analyze the consequences of such a model in this section.

As in the original model, once a bidder has chosen a minimum bid b_i , the best strategy for the bidder in the second stage of the auction will take on a specific form. If the bidder's value v_i satisfies $v_i \in [b_i, \alpha b_i]$, then the bidder has a dominant strategy of bidding his value. If the bidder's value v_i satisfies $v_i > \alpha b_i$, then the bidder would like to make a bid of v_i but is required to make a bid no greater than αb_i , so the bidder does the next best thing and makes a bid of αb_i . And if the bidder's value v_i satisfies $v_i < b_i$, then the

bidder would like to submit a bid of v_i but is required to either make a bid of at least b_i or to not bid in the auction. As a result, a bidder with value $v_i < b_i$ will follow a strategy of bidding b_i if $v_i \in [v_i^*, b_i]$ and not bidding at all if $v_i < v_i^*$, where the value of v_i^* is chosen so that a bidder with value v_i^* is indifferent between bidding b_i and not bidding at all.

As before we search for a symmetric equilibrium in which all bidders make the same minimum bid of b and then, conditional on making a minimum bid of b , they then use the same cutoff v^* such that a bidder with value $v \leq b$ follows a strategy of bidding b if $v \in [v^*, b]$ and not bidding in the auction otherwise. Unfortunately, unlike with the original game, it is no longer the case that there will always be a pure-strategy equilibrium. For some natural distributions it could be the case that all equilibria are in mixed strategies.

THEOREM 7.1. *If bidders have the option of not bidding in an auction, then pure-strategy equilibria need not exist.*

The reason for this result is that, as before, in order for it to be an equilibrium for all bidders to make a minimum bid of b , it must be the case that no bidder can profitably deviate by instead using a lowest possible bid of $b - \epsilon$ or $b + \epsilon$ for some $\epsilon > 0$. In the original model, if the conditions needed to ensure that a bidder could not profitably deviate by using a lowest possible bid of $b + \epsilon$ were satisfied with equality, then this also guaranteed that a bidder would not be able to profitably deviate by using a lowest possible bid of $b - \epsilon$. However, in the modified model this is no longer the case. As such, it is possible to construct examples where any equilibrium must be in mixed strategies.

8 CONCLUSION

This paper has presented the first theoretical analysis of a bid-limited targeting mechanism that can be used in allowing advertisers to target different users. The mechanism imposes a cap on the maximum bid multiplier that advertisers may use in multiplying some base bid that they make for generic advertising opportunities to arrive at the bid they may make for advertising opportunities that satisfy their preferred targeting criteria. Using this mechanism has the effect of tying the maximum bid that bidders may ever make in an auction to the minimum bid that they ever make in an auction.

When bidders' values are independent and identically distributed, bid-limited targeting typically leads to higher revenue than both no targeting and full targeting. Moreover, bid-limited targeting may further lead to higher revenue than setting the optimal reserve price, and choosing the optimal parameters in this mechanism can be more important than inducing an additional bidder to bid in the auction, in contrast to standard findings in mechanism design such as those in [7]. However, when bidders' values are not identically distributed, pure-strategy equilibria need not exist. A more thorough analysis of this case is left to future work.

ACKNOWLEDGMENTS

We thank Andrei Broder, Preston McAfee, Michael Schwarz, and D. Sivakumar for helpful comments and discussions.

¹¹Nonetheless, we show in the full version of the paper [20] that a mixed-strategy equilibrium will exist in this setting.

REFERENCES

- [1] Ittai Abraham, Susan Athey, Moshe Babaioff, and Michael Grubb. 2016. *Peaches, Lemons, and Cookies: Designing Auction Markets with Dispersed Information*. Typescript. Microsoft Research, Mountain View, CA.
- [2] Mohammadhosen Bateni, Jon Feldman, Vahab Mirrokni, and Sam Chiu-Wai Wong. 2014. Multiplicative Bidding in Online Advertising. In *Proceedings of the 15th ACM Conference on Electronic Commerce (EC '14)*. 715–732.
- [3] Dirk Bergemann and Martin Pesendorfer. 2007. Information Structures in Optimal Auctions. *Journal of Economic Theory* 137, 1 (2007), 580–609.
- [4] Dirk Bergemann and Achim Wambach. 2015. Sequential Information Disclosure in Auctions. *Journal of Economic Theory* 159, B (2015), 1074–1095.
- [5] Kshipra Bhawalkar, Patrick Hummel, and Sergei Vassilvitskii. 2014. Value of Targeting. In *Proceedings of the 7th International Symposium on Algorithmic Game Theory (SAGT '14)*. 194–205.
- [6] Simon Board. 2009. Revealing Information in Auctions: The Allocation Effect. *Economic Theory* 38, 1 (2009), 125–135.
- [7] Jeremy I. Bulow and Paul Klemperer. 1996. Auctions Versus Negotiations. *American Economic Review* 86, 1 (1996), 180–194.
- [8] Elisa L. Celis, Gregory Lewis, Markus M. Mobius, and Hamid Nazerzadeh. 2011. Buy-It-Now or Take-A-Chance: A Simple Sequential Screening Mechanism. In *Proceedings of the 20th International Symposium on Algorithmic Game Theory (WWW '11)*. 147–156.
- [9] Indranil Chakraborty. 1999. Bundling Decisions for Selling Multiple Objects. *Economic Theory* 13, 3 (1999), 723–733.
- [10] Alexandre De Cornière and Romain De Nijs. 2016. Online Advertising and Privacy. *RAND Journal of Economics* 47, 1 (2016), 48–72.
- [11] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars of Keywords. *American Economic Review* 97, 1 (2007), 242–259.
- [12] Yuval Emek, Michal Feldman, Iftah Gamzu, Renato Paes Leme, and Moshe Tennenholtz. 2014. Signaling Schemes for Revenue Maximization. *ACM Transactions of Economics and Computation* 2, 2, Article 5 (2014).
- [13] Péter Eso and Balázs Szentes. 2007. Optimal Information Disclosure in Auctions and the Handicap Auction. *Review of Economic Studies* 74, 3 (2007), 705–731.
- [14] Hu Fu, Patrick Jordan, Mohammad Mahdian, Uri Nadav, Inbal Talgam-Cohen, and Sergei Vassilvitskii. 2012. Ad Auctions with Data. In *Proceedings of the 5th International Symposium on Algorithmic Game Theory (SAGT '12)*. 184–189.
- [15] Juan-José Ganuza. 2004. Ignorance Promotes Competition: An Auction Model with Endogenous Private Valuations. *RAND Journal of Economics* 35, 3 (2004), 583–598.
- [16] Arpita Ghosh, Hamid Nazerzadeh, and Mukund Sundarajan. 2007. Computing Optimal Bundles for Sponsored Search. In *Proceedings of the 3rd International Workshop on Internet and Network Economics (WINE '07)*. 576–583.
- [17] Arpita Ghosh and Aaron Roth. 2015. Selling Privacy at an Auction. *Games and Economic Behavior* 91 (2015), 334–346.
- [18] Renato Gomes and Kane Sweeney. 2014. Bayes-Nash Equilibria of the Generalized Second-Price Auction. *Games and Economic Behavior* 86 (2014), 421–437.
- [19] Patrick Hummel and R. Preston McAfee. 2016. When Does Improved Targeting Increase Revenue? *ACM Transactions of Economics and Computation* 5, 1, Article 4 (2016).
- [20] Patrick Hummel and Uri Nadav. 2017. *Bid-Limited Targeting*. Typescript. Google Inc., Mountain View, CA.
- [21] Pooya Jalaly and Eva Tardos. 2015. *Quality of Outcomes in Online Advertisement with Multiplicative Bid-Adjustments*. Typescript. Cornell University, Ithaca, NY.
- [22] Chinmay Karande, Aranyak Mehta, and Ramakrishnan Srikant. 2013. Optimizing Budget Constrained Spend in Search Advertising. In *Proceedings of the 6th ACM International Conference on Web Search and Data Mining (WSDM '13)*. 697–706.
- [23] Sébastien Lahaie and R. Preston McAfee. 2011. Efficient Ranking in Sponsored Search. In *Proceedings of the 7th International Workshop on Internet and Network Economics (WINE '11)*. 254–265.
- [24] Paul Milgrom and Robert J. Weber. 1982. A Theory of Auctions and Competitive Bidding. *Econometrica* 50, 5 (1982), 1089–1122.
- [25] Roger B. Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
- [26] Michael Ostrovsky and Michael Schwarz. 2016. *Reserve Prices in Internet Advertising Auctions: A Field Experiment*. Typescript. Stanford Graduate School of Business, Stanford, CA.
- [27] Thomas Palfrey. 1983. Bundling Decisions by a Multiproduct Monopolist with Incomplete Information. *Econometrica* 51, 2 (1983), 463–483.
- [28] Yang Sun, Yunhong Zhou, and Xiaotie Deng. 2014. Optimal Reserve Prices in Weighted GSP Auctions. *Electronic Commerce Research and Applications* 13, 3 (2014), 178–187.
- [29] Hal R. Varian. 2007. Position Auctions. *International Journal of Industrial Organization* 127, 1 (2007), 1–35.