When Does Improved Targeting Increase Revenue?

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

R. Preston McAfee
Microsoft Corp.
Redmond, WA, USA
preston@mcafee.cc

ABSTRACT

In second price auctions with symmetric bidders, we find that improved targeting via enhanced information disclosure decreases revenue when there are two bidders and increases revenue if there are at least four bidders. With asymmetries, improved targeting increases revenue if the most frequent winner wins less than 30.4% of the time, but can decrease revenue otherwise. We derive analogous results for position auctions. Finally, we show that revenue can vary non-monotonically with the number of bidders who are able to take advantage of improved targeting.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Economics, Theory

Keywords

Targeting, Revenue, Online Auctions, Position Auctions, Advertising

1. INTRODUCTION

There has been substantial concern in the Internet advertising business over whether improvements in targeting technology will reduce revenue from online advertising. The intuition underlying these concerns runs as follows. Improvements in targeting enable advertisers to more accurately identify the interests of consumers. If a consumer's interests are so accurately identified that advertisers know there is only one product this consumer would ever buy, wouldn't this process result in only a single advertiser who is willing to advertise to this consumer, meaning this advertiser can bid without competition? A more nuanced version of this argument relies on a quantity effect. Since advertisers will no longer purchase ads that reach consumers who will not be interested in their products, the total demand for advertisements will go down. If the supply of advertising opportunities remains unchanged, revenue from selling ads will decline.

The question of whether enhanced targeting increases revenue is important because of two powerful trends. First, media consumption is moving online, and print newspapers and radio have waned. The survival of much of the existing media appears to depend on the ability to monetize online content with advertising. Second, Internet advertising is increasingly using sophisticated targeting through extensive online databases on customers. Thus the likely survival of existing publishers turns on whether enhanced targeting will increase advertising revenue. Furthermore, since advertising exchanges typically take a constant fraction of a publisher's revenue, there is a direct correspondence between whether revenue increases for intermediaries and for publishers.

The argument that improvements in targeting will result in only a single relevant advertiser for each consumer is likely misplaced. The argument assumes that the purchase of the customer is a foregone conclusion, which ignores one of the main purposes of advertising: to influence the consumer's choice. While there is some advertising that is informational in nature—alerting consumers to the existence of a product and its features—the majority of advertising is intended to sway the consumer's perception of the product. This kind of advertising is commonly called emotional branding, and it is the most common kind by revenue. Coca-Cola advertises extensively to people already aware of its products. Similarly, how many American television watchers are unaware of Proctor and Gamble's Tide?

But the fact that the demand for advertisements to an individual consumer will not decline to one as a result of improved targeting does not invalidate the argument that targeting might reduce revenue. Enhanced targeting will typically increase advertiser welfare, by making advertising more effective, while reducing competition through specialization.[1] The effect of improving targeting in online advertising is exactly the reverse of pure bundling for a monopolist, where the monopolist requires consumers to purchase a bundle of objects or none at all. Targeting permits advertisers to distinguish unlike consumers, whereas pure bundling or the lack of targeting forces advertisers to treat different types of consumers as if they were the same.

[1][5] and [24] also note that such a trade-off is likely to arise as a result of improved targeting.

WWW 2015, May 18–22, 2015, Florence, Italy.
ACM 978-1-4503-3469-3/15/05.
http://dx.doi.org/10.1145/2736277.2741079.

Thus to analyze whether improvements in targeting technology increase revenue from auctions for advertisements, we can analyze whether enabling advertisers to learn more detailed information about their value before bidding would increase revenue from the auction. In particular, under targeting, we assume that the targeting information enables an advertiser to learn his exact value for advertising to a consumer before deciding how much to bid for an advertisement. By contrast, when an advertiser is unable to target, the advertiser only knows that his value for advertising to this consumer will be a random draw from some distribution, where the distribution reflects the different values the advertiser might place on advertising to different types of consumers. We compare a seller's expected revenue from auctions under these two different scenarios.

Throughout we consider a model in which bidders have private values and bidders' values are independently distributed. While this is not the only possible modeling choice, it is a natural one. There is empirical evidence that there is little correlation in bidder values within auctions on Microsoft's Ad Exchange [10], which [10] indicates implies that "bidder valuations are private, driven by idiosyncratic match quality, rather than a common component". Furthermore, if there is a common component to bidders' values that is not known to any participant and the bidders have private values that are independent conditional on the common value, then the results of this paper for the zero reserve price will continue to hold since the results are attained for each realization of the common component.[2]

We also frequently make use of the standard hazard rate condition on the cumulative distribution of the buyers' values. Although this assumption is not completely innocuous, it is satisfied by many distributions frequently encountered in empirical studies.

In this environment, when advertisements are being sold via second price auctions, we first demonstrate a result analogous to that in [8], [18], and [30] which illustrates that targeting decreases revenues when there are two bidders, even if there are asymmetries in the distributions of the bidders' values. We next show that when bidders' values are drawn from identical distributions, then improved targeting has an ambiguous effect on revenue when there are three bidders, but improved targeting increases revenue if there are at least four bidders. These results are virtually unaffected by the possibility that a seller can set reserve prices. Finally, we address the question of what happens when the bidders' values are drawn from different distributions. Here we find that if the strongest firm wins the auction less than 30.4% of the time, then improved targeting increases revenue, but targeting can reduce revenue when the two strongest bidders win a disproportionate percentage of the time.

While second price auctions for a single advertising opportunity are used by most publishers, we also consider position auctions, as these are frequently used by search engines as well as a smaller number of publishers. Here we find that targeting unambiguously decreases revenue when there are only a small number of bidders, increases revenue when there are a large number of bidders, and has an ambiguous effect on revenue when there are an intermediate number of bidders. When there are an intermediate number of bidders, improved targeting increases revenue if and only if the click-through rates of the top positions are sufficiently large compared to the click-through rates of the lower positions.

Finally, we address the question of how improved targeting affects revenue when only some advertisers are able to make use of the targeting information. In this setting we show that even when there are symmetric bidders whose values are drawn from a distribution satisfying standard regularity conditions, it could be the case that a seller's revenue may vary non-monotonically with the number of bidders who are able to make use of the targeting information. That is, the seller may be indifferent between targeting and bundling when only one bidder can target, prefer targeting to bundling when two bidders can target, and prefer bundling to targeting when three bidders can target. We also illustrate how improved targeting affects revenue when there is exactly one bidder who can make use of the targeting information. We find that this decreases revenue when the strongest bidder is making use of the targeting information, increases revenue when the weaker bidders are making use of the targeting information, and has an ambiguous effect on revenue for bidders of intermediate strength.

Our paper relates to two distinct strands of literature. First, our paper relates to the literature on whether a mechanism designer should provide information to bidders in a private value auction that would better help them assess their values for an object. Here [16] provides examples that illustrate that improving targeting may decrease revenue in a private value auction and [17] illustrates that an auctioneer may have an incentive to release less than full information to the bidders when the auctioneer has the ability to release partial information. [6] considers the optimal information structure in a joint design problem in which there may be a direct tie between the information the seller discloses and the mechanism the seller then uses to sell the object, [14] addresses the question of how much information the mechanism designer should provide under the optimal mechanism which may possibly involve charging the bidders in the auction for providing the information, and [12] considers questions related to how information disclosure affects the ultimate prices advertisers would charge for their products. Finally, [7] analyzes the value of targeting data to advertisers, [18] provides general methods of classifying the informativeness of signals to bidders in private value auctions, and [31] conducts field experiments analyzing the effect of information disclosure on wholesale auto auctions.

The second related strand of literature is work analyzing when sellers would want to bundle goods and sell them together. [4] and [15] both study a standard bundling framework in which a monopolist considers selling bundles of goods to buyers, but these papers do not consider situations in which the goods are sold via an auction. [2], [23], and [26] study mixed bundling in which a monopolist offers buyers both the option of buying various goods individually and the option of buying multiple goods at the same time, possibly for a discount. [11] studies a model in which a seller sells two objects via an auction and the seller must decide whether to sell them separately or via bundling. [20] studies how a seller's revenue from selling two objects separately or from only offering to sell the objects together compares to the seller's revenue from an optimal mechanism, which may

[2]In addition, we already know from [28] that if there is a common component to all bidders' values, then the seller has an incentive to reveal this common component. [1] further discusses when information asymmetries in common value auctions can lead to revenue losses.

not involve either of these approaches. [25] and [33] study questions related to when a monopolist's optimal mechanism involves take-it or leave-it mechanisms that set a price for each possible collection of goods. Finally, [3] studies questions related to optimal bundling with multiple sellers.

While these papers are all interesting, they all differ from our work in significant ways. Our paper differs from the work on bundling goods together in that very few of these papers consider models of bundling in an auction setting, whereas we focus on auctions throughout our paper. Our paper also differs from the literature on information provision by a mechanism designer in that these papers do not attempt to derive the detailed results in this paper on how the number and sizes of the various bidders affects the suitability of bundling compared to targeting for a fixed auction format. Furthermore, none of the above papers consider the problem of whether to sell goods using targeting or bundling when the seller must use a position auction, and these papers also do not consider scenarios in which some of the buyers buy objects using bundling while other buyers buy the objects separately. Our paper thus makes a number of new contributions to the literature on information provision by a mechanism designer and bundling.

2. THE MODEL

Each buyer $i \in \{1, 2, \ldots, n\}$ has a value v_i that is an independent draw from the cumulative distribution function $F_i(v)$ with finite mean and variance, and a corresponding continuous density $f_i(v)$ on its support $[0, \overline{v}_i)$, where $\overline{v}_i$ may be infinite. These bidders compete in an auction, and bid either before their value is realized (bundling) or after their value is realized (targeting). The model also applies to situations in which bidders do not learn their exact values under targeting but instead learn estimates v_i that are correct in expectation. Each bidder i's expected value under bundling is $\int_0^\infty v f_i(v)\,dv = \int_0^\infty 1 - F_i(v)\,dv$. For convenience we name the bidders in decreasing order of their expected values so $\int_0^\infty 1 - F_i(v)\,dv \geq \int_0^\infty 1 - F_{i+1}(v)\,dv$ for all i.

Throughout this manuscript we consider two possible auction formats that the bidders may compete in. First we consider standard second price auctions in which there is one object for sale and the bidder who makes the highest bid wins the object and pays the second highest bid. The results with symmetric buyers for this format will also extend to first price auctions by the revenue equivalence theorem.

The second auction format we consider is a position auction. Position auctions differ from the setting considered above in that there are s positions, where s is a positive integer satisfying $1 \leq s < n$. Each position $k \leq s$ has a click-through-rate $c_k > 0$, where c_k is non-increasing in k for all $k \leq s$. Bidders compete by submitting bids per clicks. The top position then goes to the bidder with the highest bid, the second position goes to the bidder with the second highest bid, and so on, with ties broken randomly.

We consider two methods for setting prices in position auctions. The first pricing method we consider is a generalized second price (GSP) auction. In this setting, the k^{th} highest bidder pays a price per click that is equal to the bid submitted by the $k+1^{th}$ highest bidder. Thus if $v_{(k)}$ denotes the value of the k^{th} highest bidder and $b_{(k+1)}$ denotes the bid submitted by the $k+1^{th}$ highest bidder, then the final payoff of the k^{th} highest bidder is $c_k(v_{(k)} - b_{(k+1)})$. This is the same basic model of generalized second price auctions without clickability of ads that is considered in [13] and [32].

The second possibility we consider is the Vickrey-Clarke-Groves (VCG) mechanism. Under VCG pricing, each advertiser pays a total cost equal to the externality he imposes on other bidders by bidding in the auction. Thus under VCG pricing, the bidder who wins the k^{th} position pays a total cost of $\sum_{j=k}^{s}(c_j - c_{j+1})b_{(j+1)}$ and a total price per click equal to $\frac{1}{c_k}\sum_{j=k}^{s}(c_j - c_{j+1})b_{(j+1)}$, where we abuse notation by letting $c_{s+1} \equiv 0$.

Finally we also sometimes allow for the possibility of reserve prices. If there is a reserve price of r, then only bidders who bid at least r will be considered in the auction. Under standard second price auctions, if there is only one bidder who bids more than the reserve, then this bidder pays r for the object. Under GSP auctions, if there are only $k \leq s$ bidders who bid more than the reserve price, then the payoffs of the first $k-1$ of these bidders are unaffected by the reserve price, but the k^{th} highest bidder pays a price of r per click and obtains a payoff of $c_k(v_{(k)} - r)$.

Finally, under position auctions using VCG pricing, we introduce reserve prices in the following manner: If at least $s+1$ bidders submit a bid in the auction that is greater than the reserve price, then the reserve price has no effect on the outcome of the auction. If $K \leq s$ bidders submit a bid in the auction that is greater than the reserve price, then only the bidders who submitted a bid greater than the reserve price have their ads shown and these bidders pay a price per click equal to the price they would pay if there were exactly K positions available and there were an additional bidder who submitted a bid equal to the reserve price. [21] has noted in a more general setting that this method of introducing reserve prices into the VCG mechanism both preserves the incentive for advertisers to bid truthfully and also ensures that any advertisers who have their ads shown pay a price per click that is greater than or equal to the reserve price.

3. SECOND PRICE AUCTIONS WITHOUT RESERVE PRICES

We begin by comparing bundling to targeting in a standard second price auction setting with no reserve price. Under bundling, all bidders have a weakly dominant strategy of bidding their expected values. Thus the bidder with the highest expected value wins and pays the second highest bid, and the seller's revenue under bundling is the second highest expected value or $\int_0^\infty 1 - F_2(v)\,dv$.

Under targeting bidders bid their exact values after learning their values, and the seller's revenue is the second highest realized value. The second highest realized value is less than or equal to v when either the highest value is no greater than v, or the highest value exceeds v but all other values are less than or equal v. Thus if $v_{(2)}$ denotes the realization of the second highest value, the distribution of this realization is given by the following cumulative distribution function:

$$\begin{aligned} Pr(v_{(2)} \leq v) &= \prod_{j=1}^{n} F_j(v) + \sum_{j=1}^{n}(1 - F_j(v)) \prod_{i \neq j} F_i(v) \\ &= \sum_{j=1}^{n} \prod_{i \neq j} F_i(v) - (n-1) \prod_{j=1}^{n} F_j(v) \end{aligned}$$

From this it follows that the difference between the seller's expected revenue under targeting and bundling is $\Delta_n =$

$$\int_0^\infty 1 - \sum_{j=1}^n \prod_{i\neq j} F_i(v) + (n-1)\prod_{j=1}^n F_j(v) - (1-F_2(v))\, dv$$
$$= \int_0^\infty F_2(v) - \sum_{j=1}^n \prod_{i\neq j} F_i(v) + (n-1)\prod_{j=1}^n F_j(v)\, dv$$

First we illustrate that the insight in [18] and [30] that the seller prefers bundling to targeting when there are two bidders extends to cases where the values of the bidders are not drawn from identical distributions:

THEOREM 1. *Suppose there are $n = 2$ bidders. Then the seller prefers bundling to targeting.*

PROOF. The difference between the seller's expected revenue under targeting and bundling when there are $n = 2$ bidders is $\Delta_2 =$

$$\int_0^\infty F_2(v) - \sum_{j=1}^2 \prod_{i\neq j} F_i(v) + (2-1)\prod_{j=1}^2 F_j(v)\, dx$$
$$= \int_0^\infty -F_1(v) + F_2(v)F_1(v)\, dv = \int_0^\infty (F_2(v)-1)F_1(v)\, dv < 0.$$

Thus the seller prefers bundling to targeting when there are $n = 2$ bidders. $\square$

Next we consider cases in which the bidders' values are drawn from the same distributions. When $F_i(v) = F(v)$ for all v, the difference between the seller's expected revenue under targeting and bundling is

$$\begin{aligned}\Delta_n &= \int_0^\infty F_2(v) - \sum_{j=1}^n \prod_{i\neq j} F_i(v) + (n-1)\prod_{j=1}^n F_j(v)\, dv \\ &= \int_0^\infty F(v) - nF^{n-1}(v) + (n-1)F^n(v)\, dv\end{aligned}$$

Now we use this expression for the difference between the seller's expected revenue under targeting and bundling to show that the seller prefers targeting to bundling when there are $n \geq 4$ bidders. Throughout the remainder of this manuscript, we let $f(v)$ denote the density corresponding to the cumulative distribution function $F(v)$, and let $\overline{v}$ denote the upper bound of the support of $F(\cdot)$.

THEOREM 2. *Suppose $F_i(v) = F(v)$ for all i, $\frac{1-F(v)}{f(v)}$ is nonincreasing in v throughout its support, and $n \geq 4$. Then the seller prefers targeting to bundling.*

PROOF. First note that if $n \geq 4$, then $\phi(y) \equiv \frac{y^2}{2} + \frac{y^3}{3} + \ldots + \frac{y^{n-1}}{n-1} - \frac{n-1}{n}y^n \geq 0$ for all $y \in [0,1]$. $\phi(y) = y^n(\frac{y^{2-n}}{2} + \ldots + \frac{y^{n-n}}{n} - 1)$, so $\phi \geq 0$ if and only if $\frac{y^{2-n}}{2} + \ldots + \frac{y^{n-n}}{n} - 1 \geq 0$. And since $\frac{y^{2-n}}{2} + \ldots + \frac{y^{n-n}}{n} - 1$ is decreasing in y, $\phi \geq 0$ for all $y \in [0,1]$ if and only if $\frac{1}{2} + \ldots + \frac{1}{n} - 1 \geq 0$, which holds for all $n \geq 4$. Thus $\phi(y) \geq 0$ for all $y \in [0,1]$ if $n \geq 4$.

Now the difference between the seller's expected revenue under targeting and bundling when there are n bidders and $F_i(v) = F(v)$ for all i is $\Delta_n =$

$$\int_0^{\overline{v}} F(v) - nF^{n-1}(v) + (n-1)F^n(v)\, dv$$
$$= \int_0^{\overline{v}} \left(\frac{1-F(v)}{f(v)}\right)(F(v) + \ldots + F^{n-2}(v) - (n-1)F^{n-1}(v))f(v)\, dv$$
$$= \left(\frac{1-F(v)}{f(v)}\right)\left(\frac{F^2(v)}{2} + \ldots + \frac{F^{n-1}(v)}{n-1} - \frac{n-1}{n}F^n(v)\right)\Big|_0^{\overline{v}}$$
$$- \int_0^{\overline{v}} \left(\frac{1-F(v)}{f(v)}\right)'\left(\frac{F^2(v)}{2} + \ldots + \frac{F^{n-1}(v)}{n-1} - \frac{n-1}{n}F^n(v)\right)dv$$
$$= \left(\lim_{v\to\overline{v}} \frac{1-F(v)}{f(v)}\right)\left[\frac{1}{2} + \ldots + \frac{1}{n-1} - \frac{n-1}{n}\right]$$
$$- \int_0^{\overline{v}} \left(\frac{1-F(v)}{f(v)}\right)'\left[\frac{F^2(v)}{2} + \ldots + \frac{F^{n-1}(v)}{n-1} - \frac{n-1}{n}F^n(v)\right]dv$$

By the result in the first paragraph of this proof, both the terms in square brackets are positive for all v. Also, the term $\lim_{v\to\overline{v}} \frac{1-F(v)}{f(v)}$ is non-negative since $\frac{1-F(v)}{f(v)}$ is non-negative for all v, and the term $\left(\frac{1-F(v)}{f(v)}\right)'$ is non-positive by assumption. Thus $\Delta_n \geq 0$ and the seller prefers targeting to bundling. $\square$

Thus, in the symmetric case, when there are at least four bidders and the usual hazard rate condition is satisfied, targeting dominates bundling. The reason for this is that the second highest realized value (the revenue under targeting) when there are at least four bidders is greater than the average value of these bidders (the revenue under bundling) as long as the distribution does not have fat tails. The possibility of fat tails is ruled out by the standard hazard rate condition, so targeting dominates bundling when there are $n = 4$ bidders in this case.[3] A similar approach can illuminate the case of three bidders.

THEOREM 3. *Suppose $F_i(v) = F(v)$ for all i and there are $n = 3$ bidders. Then the seller prefers targeting to bundling if $f(v)$ is increasing in v on its support, prefers bundling to targeting if $f(v)$ is decreasing in v on its support, and is indifferent between targeting and bundling if $f(v)$ is constant on its support.*

PROOF. The difference between the seller's expected revenue under targeting and bundling when there are $n = 3$ bidders is $\Delta_3 =$

$$\int_0^{\overline{v}} F(v) - 3F^2(v) + 2F^3(v)\, dv$$
$$= \int_0^{\overline{v}} \frac{1}{f(v)}(F(v) - 3F^2(v) + 2F^3(v))f(v)\, dv$$
$$= \frac{1}{f(v)}\left(\frac{F^2(v)}{2} - F^3(v) + \frac{F^4(v)}{2}\right)\Big|_0^{\overline{v}}$$
$$- \int_0^{\overline{v}} \left(\frac{1}{f(v)}\right)'\left(\frac{F^2(v)}{2} - F^3(v) + \frac{F^4(v)}{2}\right)dv$$
$$= \frac{F^2(v)}{2f(v)}(1-F(v))^2\Big|_0^{\overline{v}} - \int_0^{\overline{v}} \left(\frac{1}{f(v)}\right)' \frac{F^2(v)}{2}(1-F(v))^2\, dv$$

Now if $f(v)$ is increasing in v on its support or $f(v)$ is constant on its support, then $\frac{F^2(v)}{2f(v)}(1-F(v))^2\big|_0^{\overline{v}} = 0$ and

[3]For any finite number of players, there exists a distribution with sufficiently fat tails such that the seller prefers bundling to targeting.

$\Delta_3 = -\int_0^{\overline{v}} (\frac{1}{f(v)})' \frac{F^2(v)}{2} (1-F(v))^2 \ dv$. Thus if $f(v)$ is increasing in v on its support, then $\Delta_3 > 0$, and if $f(v)$ is constant on its support, then $\Delta_3 = 0$. Similarly, if $f(v)$ is decreasing in v on its support and $\frac{F^2(v)}{2f(v)} (1-F(v))^2 \big|_0^{\overline{v}} = 0$, then $\Delta_3 = -\int_0^{\overline{v}} (\frac{1}{f(v)})' \frac{F^2(v)}{2} (1-F(v))^2 \ dv < 0$.

If $f(v)$ is decreasing in v and $\frac{F^2(v)}{2f(v)} (1-F(v))^2 \big|_0^{\overline{v}} \neq 0$ (which implies $\overline{v} = \infty$), then consider what Δ_3 would equal if the players' values were instead random draws from the distribution $F(v|\theta)$ satisfying $F(v|\theta) = \frac{F(v)}{F(\theta)}$ for $v \leq \theta$ and $F(v) = 1$ for $v > \theta$. For any finite $\theta > 0$, it is necessarily the case that $\frac{F^2(v|\theta)}{2f(v|\theta)} (1-F(v|\theta))^2 \big|_0^{\overline{v}|\theta} = 0$, where $f(v|\theta)$ denotes the density corresponding to $F(v|\theta)$ and $\overline{v}|\theta$ denotes the upper bound on the support of $F(v|\theta)$. Moreover, $f(v|\theta)$ is decreasing in v on its support, so $\Delta_3 < 0$ when the players' values are random draws from $F(v|\theta)$. Furthermore, Δ_3 must be bounded away from 0 for all $\theta \geq \theta^*$, where θ^* is some constant in the interior of the support of $F(\cdot)$.

But in the limit as θ becomes arbitrarily large, the value of Δ_3 when the values of the players are random draws from the distribution $F(v|\theta)$ becomes arbitrarily close to the value of Δ_3 when the values of the players are random draws from the distribution $F(v)$. From this it follows that if $f(v)$ is decreasing in v on its support, then $\Delta_3 < 0$ even if $\frac{F^2(v)}{2f(v)} (1-F(v))^2 \big|_0^{\overline{v}} \neq 0$. The result then follows. □

The fact that a seller may either prefer bundling or targeting when there are three bidders is intuitive. When there are three bidders, the second highest value is the median value, so the seller's expected revenue under targeting is just the median value of the bidders. The seller's expected revenue under bundling is the expected value of the bidders. Whether a seller prefers targeting to bundling simply depends on whether the mean or the median of a certain distribution is greater.

In summary, when the buyers' values are drawn from identical continuous distributions, the seller typically prefers targeting to bundling when there are four or more bidders, while the seller prefers bundling to targeting when there are two bidders. The seller's exact preferences in the case where there are three bidders depend on the distribution, but since most natural distributions of values have a density $f(v)$ that is decreasing on most of its support, the seller is also likely to prefer bundling to targeting when there are three bidders.

4. RESERVE PRICES

When there is more than one bidder, in symmetric settings appropriate reserve prices favor targeting. Adding reserve prices does not improve the seller's revenue under bundling since the seller's revenue is equal to the bidders' expected values regardless of whether the seller uses a reserve price, but reserve prices do increase the revenue from targeting. Thus when there are four or more bidders, targeting with reserve prices dominates bundling with reserve prices.

However, while reserve prices increase the revenue from targeting but not the revenue from bundling, it is still the case that the seller typically prefers bundling to targeting when there are two bidders. This is illustrated below:

THEOREM 4. *Suppose there are $n = 2$ bidders, $F_i(v) = F(v)$ for all i, and the density $f(v)$ is non-increasing in v. Then the seller prefers bundling to targeting with the optimal reserve price.*

PROOF. We know from [9] that when $f(v)$ is non-increasing in v, the seller's expected revenue in an auction with two bidders and the optimal reserve price is lower than the seller's expected revenue in an auction with three bidders and no reserve price. However, we know from Theorem 3 that when $F_i(v) = F(v)$ for all i, and the density $f(v)$ is non-increasing in v, then the seller prefers bundling to targeting when there are $n = 3$ bidders and no reserve price. Since the seller's expected revenue under targeting when there are two bidders with the optimal reserve price is even lower than the seller's expected revenue under targeting when there are three bidders and no reserve price, it then follows that the seller prefers bundling to targeting when there are $n = 2$ bidders, even if the seller uses the optimal reserve price. □

In addition to bundling still typically being optimal in the case where there are two bidders, it is also the case that the seller will sometimes want to use bundling even when there are three bidders. Although the seller now prefers targeting to bundling in the case where the buyers' values are drawn from a uniform distribution, the seller still prefers bundling to targeting when the buyers' values are drawn from an exponential distribution, even if the seller uses the optimal reserve price:

THEOREM 5. *Suppose there are $n = 3$ bidders and $F_i(v) = F(v)$ for all i. Then the seller prefers targeting to bundling when the bidders' values are drawn from the uniform distribution, but the seller prefers bundling to targeting when the bidders' values are drawn from the exponential distribution.*

PROOF. When there are $n = 3$ bidders and each bidder's value is an independent and identically distributed draw from the uniform distribution, we know from Theorem 3 that the seller is indifferent between bundling and targeting when there is no reserve price. Since setting the optimal reserve price increases the seller's revenue under targeting but not under bundling, it then follows that when there are $n = 3$ bidders and the seller sets the optimal reserve price, the seller obtains greater revenue under targeting than under bundling when the bidders' values are drawn from the uniform distribution.

Now suppose there are $n = 3$ bidders and the bidders' values are independent and identically distributed draws from the exponential distribution with cumulative distribution function $F(v) = 1 - e^{-v}$. If there is no targeting, then all bidders bid their expected value of 1, and the seller's revenue will be 1. If there is targeting, then the seller's optimal reserve price r satisfies $r = \frac{1-F(r)}{f(r)} = 1$, and the seller's expected revenue is

$$\begin{aligned} & \int_r^\infty \left(v - \frac{1-F(v)}{f(v)} \right) nF(v)^{n-1} f(v) \ dv \\ = & \int_1^\infty (v-1) nF(v)^{n-1} f(v) \ dv \\ = & -(v-1)(1-F(v)^n) \Big|_1^\infty + \int_1^\infty 1 - F(v)^n \ dv \\ = & \int_1^\infty 1 - (1-e^{-v})^3 \ dv = \frac{2 - 9e + 18e^2}{6e^3} < 1. \end{aligned}$$

Thus the seller prefers bundling to targeting when the bidders' values are drawn from the exponential distribution. □

5. ASYMMETRIC BIDDERS WITHOUT RESERVE PRICES

We now consider a scenario in which the values of the bidders are not all drawn from the same distribution. In particular, we consider a scenario in which there is some cumulative distribution function $F(v)$ and some values $\alpha_1, \ldots, \alpha_n$ satisfying $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_n > 0$ such that $F_i(v) = F(v)^{\alpha_i}$ for all i. This formulation is useful because the values of α_i have a natural interpretation in terms of the bidders' probabilities of winning the auction. If $A \equiv \sum_{j=1}^{n} \alpha_j$, then the probability bidder j has the highest value is

$$Pr(v_j > v_i \; \forall \; i \neq j) = \int_0^\infty \prod_{i \neq j} F_i(v) \alpha_j F(v)^{\alpha_j - 1} f(v) \, dv = \frac{\alpha_j}{A}$$

We now give a result that expresses the circumstances under which targeting is preferred to bundling as a function of the probabilities the various bidders win the auction.

LEMMA 1. *Suppose that $F_i(v) = F(v)^{\alpha_i}$ for some $\alpha_1, \ldots, \alpha_n$ satisfying $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_n > 0$, α_2 is sufficiently large, $n \geq 3$, and $\frac{1-F(v)}{f(v)}$ is nonincreasing in v. Then if $\frac{\alpha_2}{A} \leq \left(1 - \frac{\alpha_1}{A}\right)\left(1 - \frac{\alpha_2}{A}\right)^{\frac{1-\alpha_1/A}{\alpha_2/A}}$, the seller prefers targeting to bundling.*

We defer the proof of this result to the full version of the paper [22]. The inequality in Lemma 1 is a function only of two variables, $\frac{\alpha_1}{A}$ and $\frac{\alpha_2}{A}$, and indicates that we can give conditions under which targeting is preferred to bundling solely in terms of the probabilities the strongest bidders will win the auction. We now seek to show when this inequality is satisfied given that α_1 and α_2 must meet the constraints $0 \leq \frac{\alpha_2}{A} \leq \min\{\frac{\alpha_1}{A}, 1 - \frac{\alpha_1}{A}\}$.

THEOREM 6. *The inequality in Lemma 1 is satisfied if $\frac{\alpha_1}{A} \leq 0.30366$. If $\frac{\alpha_1}{A} > 0.30366$, there exists some $y^* \in (0, \min\{\frac{\alpha_1}{A}, 1 - \frac{\alpha_1}{A}\})$ such that this inequality is satisfied if and only if $\frac{\alpha_2}{A} \leq y^*$. Furthermore, if this key value of y^* is taken as a function of $\frac{\alpha_1}{A}$, $y^*(\frac{\alpha_1}{A})/(1 - \frac{\alpha_1}{A})$ is increasing in $\frac{\alpha_1}{A}$ and as $\frac{\alpha_1}{A} \to 1$, $y^*(\frac{\alpha_1}{A})/(1 - \frac{\alpha_1}{A}) \to 1$.*

PROOF. First we show that there is some $y^* \in (0, 1 - \frac{\alpha_1}{A})$ such that the inequality in Lemma 1 is satisfied if and only if $\frac{\alpha_2}{A} \leq y^*$. Note that $\frac{\alpha_2}{A} \leq \left(1 - \frac{\alpha_1}{A}\right)\left(1 - \frac{\alpha_2}{A}\right)^{\frac{1-\alpha_1/A}{\alpha_2/A}}$ holds if and only if $\frac{\alpha_2/A}{1-\alpha_1/A} \leq \left(1 - \frac{\alpha_2}{A}\right)^{\frac{1-\alpha_1/A}{\alpha_2/A}}$. Now let $\beta \equiv \frac{\alpha_2}{A}$ and let $\gamma \equiv 1 - \frac{\alpha_1}{A}$. We can rewrite this inequality in terms of β and γ as $\frac{\beta}{\gamma} \leq (1-\beta)^{\gamma/\beta}$ or $(\beta/\gamma)^{(\beta/\gamma)} \leq 1 - \beta$. And if $x \equiv \frac{\beta}{\gamma}$, then we can further rewrite this inequality as $x^x \leq 1 - \gamma x$ or $x^x + \gamma x \leq 1$. Now $g(x; \gamma) \equiv x^x + \gamma x$ is a convex function of x that satisfies $g(0; \gamma) = 1$, $g(1; \gamma) = 1 + \gamma$ and $g'(0) = -\infty$. Thus $g(x; \gamma) \leq 1$ if x is sufficiently close to 0, $g(x; \gamma) > 1$ if x is sufficiently close to 1, and there is some $x^*(\gamma) \in (0, 1)$ such that $g(x; \gamma) \leq 1$ if and only if $x \leq x^*(\gamma)$. From this it follows that this inequality is satisfied if and only if $\frac{\alpha_2}{A} \leq y^*$ for some $y^* \in (0, 1 - \frac{\alpha_1}{A})$.

The result in the previous paragraph implies that if the inequality in Lemma 1 is satisfied when $\frac{\alpha_2}{A} = \frac{\alpha_1}{A}$, then this inequality is also satisfied for any values of $\frac{\alpha_2}{A} \leq \frac{\alpha_1}{A}$. Now the inequality in Lemma 1 is satisfied when $\frac{\alpha_2}{A} = \frac{\alpha_1}{A}$ if and only if $\frac{\alpha_1}{A} \leq \left(1 - \frac{\alpha_1}{A}\right)\left(1 - \frac{\alpha_1}{A}\right)^{\frac{1-\alpha_1/A}{\alpha_1/A}}$, which holds if and only if $\frac{\alpha_1/A}{1-\alpha_1/A} \leq \left(1 - \frac{\alpha_1}{A}\right)^{\frac{1-\alpha_1/A}{\alpha_1/A}}$. Thus if $\delta \equiv \frac{\alpha_1}{A}$, then this holds if and only if $\frac{\delta}{1-\delta} \leq (1-\delta)^{\frac{1-\delta}{\delta}}$ or $(\frac{\delta}{1-\delta})^{\delta/(1-\delta)} + \delta \leq 1$. Now $h(\delta) \equiv (\frac{\delta}{1-\delta})^{\delta/(1-\delta)} + \delta$ satisfies $h(0) = 1$, $h(1) = \infty$, and $h'(0) = -\infty$. Thus $h(\delta) \leq 1$ for δ sufficiently close to 0, $h(\delta) > 1$ if δ is sufficiently close to 1, and there is some $\delta^* \in (0, 1)$ such that $h(\delta^*) \leq 1$ if and only if $\delta \leq \delta^*$. Thus the inequality in Lemma 1 is satisfied for all $\frac{\alpha_2}{A} \leq \frac{\alpha_1}{A}$ if and only if $\frac{\alpha_1}{A} \leq \delta^*$, where δ^* is the unique $\delta \in (0, 1)$ satisfying $h(\delta) = 1$. Computationally it follows that $\delta^* = 0.30366$, so the inequality in Lemma 1 is satisfied if $\frac{\alpha_1}{A} \leq 0.30366$, and if $\frac{\alpha_1}{A} > 0.30366$, then there is some $y^* \in (0, \min\{\frac{\alpha_1}{A}, 1 - \frac{\alpha_1}{A}\})$ such that this inequality is satisfied if and only if $\frac{\alpha_2}{A} \leq y^*$.

Furthermore, since $g(x; \gamma)$ is increasing in γ, the critical value of $x^*(\gamma)$ given in the first paragraph of this proof is decreasing in γ, meaning x^* is increasing in $\frac{\alpha_1}{A}$. From this it follows that for the key value of y^* above, it must be the case that $y^*(\frac{\alpha_1}{A})/(1 - \frac{\alpha_1}{A})$ is increasing in $\frac{\alpha_1}{A}$. Furthermore, when $\gamma = 0$, $x^*(\gamma) = 1$ since $g(1; 0) = 1$. Thus as $\frac{\alpha_1}{A} \to 1$, $y^*(\frac{\alpha_1}{A})/(1 - \frac{\alpha_1}{A}) \to 1$ as well. □

The above result indicates that when $\frac{\alpha_1}{A} \leq 0.30366$ and the strongest firm wins the auction less than 30.366% of the time, there is automatically enough competition in the auction that targeting will increase revenue. When the largest firm is larger than this, then improved targeting will increase revenue if and only if the second largest firm is sufficiently small and there is enough competition from other firms.

Interestingly, the result that $y^*(\frac{\alpha_1}{A})/(1 - \frac{\alpha_1}{A})$ is increasing in $\frac{\alpha_1}{A}$ indicates that as the strongest firm becomes more dominant, the second strongest firm can be relatively stronger compared to the weaker firms without changing the result that targeting increases revenue. Furthermore, as $\frac{\alpha_1}{A} \to 1$ and the strongest firm becomes arbitrarily strong, $y^*(\frac{\alpha_1}{A})/(1 - \frac{\alpha_1}{A}) \to 1$, indicating that the second strongest firm can also become arbitrarily strong relative to the weaker firms and still ensure that targeting increases revenue. This makes sense intuitively since when the strongest firm becomes more dominant, there is a greater need to allow targeting to increase the chances that the strongest firm will be given a substantial challenge in the auction.

6. POSITION AUCTIONS

Having discussed whether targeting is preferred to bundling in the case of a single object auction, we now consider how the results would be affected by using position auctions. We focus on the case where $F_i(v) = F(v)$ for all i, and $F(v)$ has compact support $[0, \overline{v}]$.

In this symmetric case, if the seller uses bundling, and all bidders bid before learning the realizations of their values, then all bidders have an expected value for a click that equals $\int_0^{\overline{v}} 1 - F(v) \, dv$ and it is an equilibrium for all bidders to bid this expected value regardless of whether we use GSP or VCG pricing. Total revenue for the seller under bundling is therefore equal to $\sum_{k=1}^{s} c_k \int_0^{\overline{v}} 1 - F(v) \, dv$.

Next we consider the case in which bidders bid after learning the realizations of their values. We assume that when bidders bid that they know their own values but they do not know the bids or the values of any of the other bidders.

Throughout our analysis of targeting, we assume that there is a symmetric pure strategy equilibrium in which bidders follow monotonic bidding strategies. This assumption is a mild assumption for two reasons. First, we show in [22] that when bidders are restricted to making bids in discrete increments of ϵ for some small $\epsilon > 0$ that there exists a symmetric pure strategy equilibrium in which bidders follow the same monotonic bidding strategies. This assumption that bidders bid in discrete increments is realistic in situations in which bidders cannot adjust their bids by less than some very small amount (such as a small fraction of a penny).

Second, [19] has already illustrated that for a wide variety of cases, there exists a symmetric pure strategy equilibrium in which bidders follow strictly monotonic bidding strategies, even if the players may submit bids along a continuous scale. We thus take symmetric monotonic equilibria as a starting point and use this to address the question of whether targeting is preferred to bundling for the seller.

When bidders follow a symmetric monotonic bidding strategy in equilibrium, the seller's expected revenue in the GSP auction is pinned down by the following lemma:

LEMMA 2. *Suppose $F_i(v) = F(v)$ for all i and the bidders use a symmetric and strictly monotonic bidding strategy $b(v)$ in equilibrium. Then expected revenue in GSP auctions equals $n \int_r^{\overline{v}} \sum_{k=1}^{s} c_k \binom{n-1}{k-1}(1-F(v))^{k-1}F(v)^{n-k}(v - \frac{1-F(v)}{f(v)})f(v)\, dv$.*

We defer the proof of this result to the full version of the paper [22]. The above result illustrates that there is a natural correspondence between the seller's expected revenue in a single auction setting and the seller's revenue in a generalized second price auction with private values. In a standard private values auction setting, the seller's expected revenue is just the expectation of the highest virtual valuation $v - \frac{1-F(v)}{f(v)}$. In a GSP auction, the only difference is that the seller's expected revenue is now the sum of the expectations of the j^{th} highest virtual valuations $v - \frac{1-F(v)}{f(v)}$ weighted by the various click-through rates.

The seller's revenue in GSP auctions also turns out to be the exact same as the seller's revenue in position auctions using VCG pricing. [21] has characterized the seller's revenue in a more general class of position auctions when prices for advertising are set according to VCG pricing. In the special case of [21] corresponding to the model considered in the present paper, the seller's revenue under VCG pricing is the exact same as the seller's revenue in Lemma 2.

Now we use the above result to address the question of whether the seller prefers targeting or bundling when there are no more than s positions in the auction for some fixed s. Our characterization of the circumstances under which the seller prefers targeting to bundling illustrates that there are some natural similarities between the situations in which the seller prefers targeting to bundling in position auctions and standard auctions for a single object. When there is a relatively small number of players, the seller prefers bundling to targeting, and when there is a larger number of players, the seller prefers targeting to bundling. For intermediate numbers of players, it is ambiguous whether the seller prefers targeting to bundling. This result is formalized in the following theorem:

THEOREM 7. *Suppose $F_i(v) = F(v)$ for all i, $v - \frac{1-F(v)}{f(v)}$ is increasing in v, the bidders use a symmetric and strictly monotonic bidding strategy $b(v)$ in equilibrium under GSP pricing, and the reserve price is either zero or the optimal reserve. Then the following hold regardless of whether the seller uses GSP or VCG pricing:*

(1). There exists some $n^* \geq 2$ such that bundling is always preferred to targeting if and only if $n \leq n^*$.

(2). There exists some $n^{**} > n^*$ such that targeting is always preferred to bundling if and only if $n \geq n^{**}$.

(3). For values of $n \in (n^*, n^{**})$, there exists some positive integer $k^* < s$ such that targeting is preferred to bundling if and only if the values of c_k for $k \leq k^*$ are sufficiently large compared to the values of c_k for $k > k^*$. Moreover, this k^* is nondecreasing in n.

PROOF. The seller's expected revenue from targeting under the conditions of the theorem can be rewritten as $\sum_{k=1}^{s} c_k E\left[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} \middle| v_{(k)} \geq r\right] Pr(v_{(k)} \geq r)$, where $v_{(k)}$ denotes the k^{th} highest value of n draws from the distribution F. We use this to prove each of the three results.

First note that in the limit as $n \to \infty$, $E\left[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} \middle| v_{(k)} \geq r\right] Pr(v_{(k)} \geq r) \to \overline{v}$ for all k since in the limit as $n \to \infty$, $v_{(k)} \to \overline{v}$ and $\frac{1-F(v_{(k)})}{f(v_{(k)})} \to 0$ with probability arbitrarily close to 1 for all k. Thus in the limit as $n \to \infty$, the expected revenue from the mechanism under targeting approaches $\sum_{k=1}^{s} c_k \overline{v}$. By contrast, under bundling, all bidders bid $w \equiv \int_0^{\overline{v}} 1 - F(v)\ dv < \overline{v}$, and the total expected revenue under bundling is $\sum_{k=1}^{s} c_k w < \sum_{k=1}^{s} c_k \overline{v}$. From this it follows that for sufficiently large values of n, the expected revenue from targeting exceeds the expected revenue from bundling for all values of c_k.

Also note that $E\left[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} \middle| v_{(k)} \geq r\right] Pr(v_{(k)} \geq r)$ is increasing in n for all k since the distribution of the k^{th} highest of $n+1$ draws from the cumulative distribution function F first order stochastically dominates the distribution of the k^{th} highest of n draws from the cumulative distribution function F and the k^{th} highest virtual valuation $v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})}$ is strictly increasing in the k^{th} highest value $v_{(k)}$. From this it follows that the expected revenue from the mechanism under targeting is strictly increasing in n. But we have seen the expected revenue from the mechanism under bundling is $\sum_{k=1}^{s} c_k w$, which is independent of n. Combining this with the results in the previous paragraph shows that there is some n^{**} such that targeting is always preferred to bundling if and only if $n \geq n^{**}$.

Next note that if $n = 2$, then bundling is strictly preferred to targeting. If $n = 2$, then it must be the case that $s = 1$ and the position auction is equivalent to a standard second price auction. But we have already seen that under the standard second price auction that bundling is strictly preferred to targeting when $n = 2$. Thus bundling is also preferred to targeting in position auctions when $n = 2$. And we have also seen that the seller's expected revenue from targeting is strictly increasing in n, while the seller's expected revenue from bundling is independent of n. Combining these facts shows that there is some $n^* \geq 2$ such that bundling is always preferred to targeting if and only if $n \leq n^*$.

Finally consider values of $n \in (n^*, n^{**})$ for which it is neither the case that targeting is always preferred to bundling or that bundling is always preferred to targeting. The seller's expected revenue under targeting is $\sum_{k=1}^{s} c_k E\left[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} \middle| v_{(k)} \geq r\right] Pr(v_{(k)} \geq r)$, whereas the seller's expected revenue under bundling is $\sum_{k=1}^{s} c_k w$, where $w \equiv \int_0^{\overline{v}} 1 - F(v)\, dv < \overline{v}$. Thus the difference between the seller's expected revenue under targeting and the seller's expected revenue under bundling is $\sum_{k=1}^{s} c_k(E\left[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})}|v_{(k)} \geq r\right] Pr(v_{(k)} \geq r) - w)$. But $E\left[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})}|v_{(k)} \geq r\right] Pr(v_{(k)} \geq r)$ is decreasing in k since the distribution of the k^{th} highest of n draws from the cumulative distribution function F first order stochastically dominates the distribution of the $k+1^{th}$ highest of n draws from the cumulative distribution function F and the k^{th} highest virtual valuation $v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})}$ is strictly increasing in the k^{th} highest value $v_{(k)}$. Thus there is some $k^* \in [1, s)$ such that $E\left[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})}|v_{(k)} \geq r\right] Pr(v_{(k)} \geq r) > w$ if and only if $k \leq k^*$.

But from this it follows that the difference between the seller's expected revenue under targeting and the seller's expected revenue under bundling, $\sum_{k=1}^{s} c_k(E\left[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})}|v_{(k)} \geq r\right] Pr(v_{(k)} \geq r) - w)$, is strictly decreasing in c_k for all $k > k^*$ and strictly increasing in c_k for all $k \leq k^*$. From this it follows that for values of $n \in (n^*, n^{**})$, there is some $k^* \in [1, s)$ such that targeting is preferred to bundling if and only if the values of c_k for $k \leq k^*$ are sufficiently large compared to the values of c_k for $k > k^*$. Moreover, since $E\left[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})}|v_{(k)} \geq r\right] Pr(v_{(k)} \geq r)$ is increasing in n for all k, the relevant value of $k^* \in [1, s)$ for which $E\left[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})}|v_{(k)} \geq r\right] Pr(v_{(k)} \geq r) > w$ if and only if $k \leq k^*$ is nondecreasing in n. Thus the $k^* \in [1, s)$ for which targeting is preferred to bundling if and only if the values of c_k for $k \leq k^*$ are sufficiently large compared to the values of c_k for $k > k^*$ is also nondecreasing in n. $\square$

To understand the intuition behind this result, note that the number of players in position auctions never has any effect on seller revenues under bundling because the players always bid their (identical) expected values under bundling. But seller revenues are increasing in the number of bidders in position auctions because the expectations of the k^{th} highest virtual valuations $v - \frac{1-F(v)}{f(v)}$ are all increasing in the number of players. This explains the observed comparative statics results with respect to the number of players in Theorem 7.

The comparative statics results in part (3) of Theorem 7 follow from the difference between the seller's expected revenue from each position in the position auction under targeting and bundling. Under targeting, the seller's expected revenue per click from the top positions is greater than the seller's expected revenue per click from the bottom positions, but the seller's expected revenue per click is independent of position under bundling. Thus situations in which the top positions contribute a disproportionate percentage of revenues compared to bottom positions make targeting a better choice, whereas situations in which the bottom positions contribute a substantial percentage of revenues may make bundling a better choice. This gives the comparative statics results given in part (3) of Theorem 7.

Finally we give an example to give a sense of the values n^* and n^{**} that arise in Theorem 7. When the players' values are drawn from the uniform distribution and there is no reserve price, we obtain the following result:

THEOREM 8. *Suppose the bidders' values are independent draws from the uniform distribution on* $[0, 1]$ *and there is no reserve price. Then the appropriate values for* n^* *and* n^{**} *in Theorem 7 are* $n^* = 3$ *and* $n^{**} = 2s + 1$.

PROOF. Under position auctions, the seller's expected revenue from targeting equals the seller's expected revenue under the VCG mechanism, which is $\sum_{k=1}^{s} k(c_k - c_{k+1})E[v_{(k+1)}]$, where $v_{(k)}$ denotes the value of the bidder with the k^{th} highest value, and $c_{s+1} \equiv 0$. Now when the bidders' values are draw from the uniform distribution on $[0, 1]$, it is necessarily the case that $E[v_{(k+1)}] = 1 - \frac{k+1}{n+1} = \frac{n-k}{n+1}$, so the seller's revenue under targeting is $\sum_{k=1}^{s} \frac{k(n-k)}{n+1}(c_k - c_{k+1})$. Also, since the bidders all make a bid of $\frac{1}{2}$ under bundling, the seller's revenue under bundling is $\frac{1}{2}\sum_{k=1}^{s} c_k$.

Now if $n = 3$, then $s \leq 2$, and the seller's revenue under targeting reduces to $\frac{1}{2}(c_1 - c_2) + \frac{1}{2}c_2 = \frac{1}{2}c_1$, while the seller's revenue under bundling is $\frac{1}{2}(c_1 + c_2)$. From this it follows that if $n = 3$, then bundling dominates targeting. If $n = 4$, then $s \leq 3$, and the seller's revenue under targeting reduces to $\frac{3}{5}(c_1 - c_2) + \frac{4}{5}(c_2 - c_3) + \frac{3}{5}c_3 = \frac{3}{5}c_1 + \frac{1}{5}(c_2 - c_3)$, but the seller's revenue under bundling is $\frac{1}{2}(c_1 + c_2 + c_3)$. From this it follows that if $c_2 = c_3 = 0$, then the seller's revenue under targeting is greater than the seller's revenue under bundling, but if $c_2 = c_3 = c_1$, then the seller's revenue under bundling is greater than the seller's revenue under targeting. Thus the key value of n^* in Theorem 7 is $n^* = 3$.

Now by part (3) of Theorem 7, we know that if the seller's revenue under targeting is greater than the seller's revenue under bundling when $c_2 = \ldots = c_s = c_1$, then the seller's revenue under targeting is greater than the seller's revenue under bundling for all possible values of the click-through rates. Now when $c_1 = c_2 = \ldots = c_s$, the seller's revenue under targeting is $\sum_{k=1}^{s} \frac{k(n-k)}{n+1}(c_k - c_{k+1}) = \frac{s(n-s)}{n+1}c_1$, and the seller's revenue under bundling is $\frac{s}{2}c_1$, so the seller's revenue under targeting is greater than the seller's revenue under bundling if and only if $\frac{n-s}{n+1} \geq \frac{1}{2}$, which holds if and only if $n \geq 2s + 1$. From this it follows that the seller's revenue under targeting is only guaranteed to be greater than the seller's revenue under bundling if $n \geq 2s + 1$.

By combining the results in the previous two paragraphs, it follows that, under the conditions of the theorem, the critical values n^* and n^{**} in Theorem 7 are $n^* = 3$ and $n^{**} = 2s + 1$ respectively. $\square$

7. WHAT IF NOT ALL BIDDERS CAN TARGET?

So far in this manuscript we have compared scenarios in which all bidders can target with scenarios in which not all bidders can target. While this an important baseline, there may also be some important scenarios in which certain targeting information would only help some bidders more accurately assess the values they have for a particular advertisement. Additionally, a seller may want to experiment with making targeting information available to certain advertisers but not to others. This section explores the consequences of only allowing certain bidders to target.

As before, we consider a model in which there are n bidders, and bidder i's value, v_i, is an independent draw from the cumulative distribution function F_i with corresponding density f_i. If bidder i is able to target, he learns his value before placing a bid, but if bidder i is not able to target, then the bidder simply knows that his expected value for a click equals $\int_0^\infty v f_i(v)\, dv = \int_0^\infty 1 - F_i(v)\, dv$. For notational convenience, we assume throughout that $\int_0^\infty 1 - F_i(v)\, dv \geq \int_0^\infty 1 - F_{i+1}(v)\, dv$ for all i.

First we address whether the types of comparative statics results that we obtained in the previous sections continue to hold when only some of the bidders are able to target. Previously we obtained results that suggested that targeting is more likely to be preferred to bundling when there are more bidders who can target. While one might expect this result to continue to hold when only some bidders can target, this is not the case, as the following result illustrates:

THEOREM 9. *A seller's expected revenue from targeting need not be monotonic in the number of bidders that can target in an auction for a single object.*

PROOF. Suppose there are $n = 4$ bidders and each bidder's value is an independent and identically distributed draw from the lognormal distribution with parameters $\mu < 0$ and $\sigma^2 = -2\mu$. Note that if no bidders are able to target, then each bidder has an expected value of $e^{\mu+\sigma^2/2} = 1$, each bidder bids this amount, and the seller's revenue is 1. If exactly one bidder is able to target, then three of the bidders only know that they have an expected value equal to 1, these three bidders all bid this amount, and the seller's revenue is again 1.

If exactly two bidders are able to target, then the two bidders that are not able to target both only know that they have an expected value equal to 1, these two bidders both bid this amount, and the seller's revenue is always at least 1. At the same time, there is a strictly positive probability that both sellers that are able to target will learn that their values are greater than 1, these sellers will both bid more than 1, and the seller's revenue will be greater than 1. Thus if exactly two bidders are able to target, then the seller's expected revenue in the auction is strictly greater than 1.

Now consider what happens when exactly three bidders are able to target in the limit as $\mu \to -\infty$ and $\sigma^2 = -2\mu$. Note that if exactly one of the three bidders who is able to target learns that his value is greater than 1 and the other bidders who are able to target learn that their values are less than or equal to 1, then the the seller's revenue will be the exact same as it would be if no bidders were able to target. Thus whether it is beneficial for the seller to allow targeting depends on the relative costs and benefits from circumstances in which all three bidders who are able to target learn that their values are less than 1 with the circumstances under which at least two bidders learn that their values are greater than or equal to 1.

Note that the probability that a given bidder has a value less than c for any $c > 0$ goes to one in the limit as $\mu \to -\infty$ when $\sigma^2 = -2\mu$. From this it follows that, conditional on a buyer having a value less than 1, the expectation of the buyer's value goes to zero in the limit as $\mu \to -\infty$ when $\sigma^2 = -2\mu$. Similarly, if $p(\mu)$ denotes the probability that a buyer has a value greater than 1 for a given $\mu < 0$ when $\sigma^2 = -2\mu$, it follows that $\lim_{\mu \to -\infty} p(\mu) = 0$. Thus when exactly three bidders are able to target, in the limit as $\mu \to -\infty$ and $\sigma^2 = -2\mu$, the probability all three bidders that are able to target learn that their values are less than or equal 1 goes to 1, and conditional on this event taking place, the expectation of the highest of these three bidders' values goes to 0.

Now a bidder's expected value is $(1 - p(\mu))E[v|v \leq 1] + p(\mu)E[v|v > 1]$. And we know that in the limit as $\mu \to -\infty$ when $\sigma^2 = -2\mu$, we have $p(\mu) \to 0$ and $E[v|v \leq 1] \to 0$. Thus since each bidder has an expected value of 1, it follows that in the limit as $\mu \to -\infty$ when $\sigma^2 = -2\mu$, we must have $p(\mu)E[v|v > 1] \to 1$, meaning $E[v|v > 1] = \Theta(\frac{1}{p(\mu)})$. But the probability that at least two of the bidders who are allowed to target learn that their values are greater than or equal to 1 is $O(p(\mu)^2)$ in the limit as $\mu \to -\infty$. And the expectation of the second highest of these bidders' values given that at least two of these bidders have values greater than 1 is no greater than $E[v|v > 1] = \Theta(\frac{1}{p(\mu)})$. From this it follows that the total expected benefit to allowing exactly three bidders to target from the circumstances in which at least two bidders learn that their values are greater than or equal to 1 is $O(p(\mu)^2 \frac{1}{p(\mu)}) = O(p(\mu))$, which goes to zero in the limit as $\mu \to -\infty$ when $\sigma^2 = -2\mu$.

But we have seen that the total expected costs to the seller from allowing exactly three bidders to target that result from the circumstances in which all three bidders who are able to target learn that their values are less than 1 is roughly 1 unit of revenue in expectation in the limit as $\mu \to -\infty$ when $\sigma^2 = -2\mu$. It thus follows that for sufficiently negative values of μ and $\sigma^2 = -2\mu$, a seller's expected revenue from allowing exactly three bidders to target is lower than the seller's expected revenue from not allowing any bidders to target. From this it follows that a seller's expected revenue from targeting need not be monotonic in the number of bidders that can target. $\square$

Theorem 9 illustrates that a seller's expected revenue from targeting need not increase in the number of bidders who are able to target when we fix the number of bidders and vary the number of bidders that are allowed to target. In fact, our proof illustrates a stronger result. It is possible that whether a seller will want to allow targeting may vary nonmonotonically with the number of bidders that are able to target. It may be the case that the seller is indifferent between targeting and bundling when exactly one bidder can target, the seller strictly prefers targeting to bundling when exactly two bidders can target, and the seller strictly prefers bundling to targeting when exactly three bidders can target.

These non-monotonicities are especially likely to arise in cases where the values of the bidders are drawn from distributions with fat tails. Then if there are four bidders but only one bidder can target, the second highest price in the auction is always the same as it would be if there were no targeting. If two bidders can target, this second highest price is always at least this high and sometimes strictly larger, so targeting is preferred to bundling. But if three bidders can target, then it is very likely that these bidders will all learn they have a very small value, the seller's revenue is likely to be very small, and bundling will be preferred to targeting.

Now we turn to the question of how allowing just one bidder to target would affect seller revenues when the buyers' values are drawn from different distributions. This situation is important because some targeting information may only affect one bidder's estimate of the bidder's value for advertising to a certain user.

THEOREM 10. *Suppose that only one bidder will be able to make use of certain targeting information in an auction for a single object. Then the following results hold:*

(1). The seller strictly prefers bundling to targeting if the bidder with the highest expected value is the only bidder that can target.

(2). The seller strictly prefers targeting to bundling if a bidder with the k^{th} highest expected value for some $k \geq 3$ is the only bidder that can target.

(3). If a bidder with the second highest expected value is the only bidder that can target, then the seller prefers targeting to bundling if and only if the values of the highest and third highest expected bids are sufficiently high.

PROOF. If the bidder with the highest expected value is the only bidder that can target and this bidder learns that his value exceeds the second highest expected value, then the seller's revenue is unaffected by targeting. But if this bidder learns that his value is lower than the second highest expected value, then allowing targeting decreases the seller's revenue. Thus the seller prefers bundling to targeting if the bidder with the highest expected value is the only bidder that can target.

Similarly, if a bidder with the k^{th} highest expected value for some $k \geq 3$ is the only bidder that can target and this bidder learns that his value is less than or equal to the second highest expected value, then the seller's revenue is unaffected by targeting. But if a bidder with the k^{th} highest expected value for some $k \geq 3$ learns that his value is greater than the second highest expected value, then targeting increases the seller's revenue. Thus the seller prefers targeting to bundling if a bidder with the k^{th} highest expected value for some $k \geq 3$ is the only bidder that can target.

Finally, if a bidder with the second highest expected value is the only bidder that can target, then the second highest bid is the value of the bidder with the second highest expected value (if this value is between the highest expected value and the third highest expected value), the highest expected value (if this value is less than the value of the bidder with the second highest expected value), or the third highest expected value (if this value is greater than the value of the bidder with the second highest expected value). Thus the seller's expected revenue is $\int_0^{w_{(3)}} w_{(3)} f_2(v)\, dv + \int_{w_{(3)}}^{w_{(1)}} v f_2(v)\, dv + \int_{w_{(1)}}^{\infty} w_{(1)} f_2(v)\, dv$, where $w_{(1)}$ denotes the highest expected value and $w_{(3)}$ denotes the third highest expected value. This expression is increasing in both $w_{(1)}$ and $w_{(3)}$. Furthermore, in the limit as $w_{(3)} \to w_{(2)}$, where $w_{(2)}$ denotes the second highest expected value, and $w_{(1)} \to \infty$, $\int_0^{w_{(3)}} w_{(3)} f_2(v)\, dv + \int_{w_{(3)}}^{w_{(1)}} v f_2(v)\, dv + \int_{w_{(1)}}^{\infty} w_{(1)} f_2(v)\, dv$ approaches $\int_0^{w_{(2)}} w_{(2)} f_2(v)\, dv + \int_{w_{(2)}}^{\infty} v f_2(v)\, dv > \int_0^{\infty} v f_2(v)\, dv = w_{(2)}$. And in the limit as $w_{(3)} \to 0$ and $w_{(1)} \to w_{(2)}$, $\int_0^{w_{(3)}} w_{(3)} f_2(v)\, dv + \int_{w_{(3)}}^{w_{(1)}} v f_2(v)\, dv + \int_{w_{(1)}}^{\infty} w_{(1)} f_2(v)\, dv$ approaches $\int_0^{w_{(2)}} v f_2(v)\, dv + \int_{w_{(2)}}^{\infty} w_{(2)} f_2(v)\, dv < \int_0^{\infty} v f_2(v)\, dv = w_{(2)}$. Combining these results shows that if a bidder with the second highest expected value is the only bidder that can target, then the seller prefers targeting to bundling if and only if the values of the highest and third highest expected bids are sufficiently high. $\square$

Theorem 10 suggests that it is not in a seller's interest to enable targeting if only the strongest bidder will be able to make use of the targeting information, it is in a seller's interest to improve targeting if only the weakest bidders will be able to make use of the targeting information, and it may or may not be in a seller's interest to improve targeting if only intermediate-strength bidders will be able to make use of the targeting. These results are somewhat related to Myerson's [29] insights on optimal auctions. [29] finds that when asymmetric bidders are competing in an auction, a seller can increase revenue by giving an artificial bonus to the weaker bidders. In this setting, the seller can likewise increase revenue when the weaker bidders have the advantage of being able to target. We show in the full version of the paper [22] that this insight extends to position auctions as well.

8. CONCLUSION

This paper has analyzed circumstances under which improved targeting increases revenue. We have generally found that improved targeting increases revenue when there are a sufficiently large number of serious bidders, but targeting can hurt revenue when there are just a few dominant bidders. These types of results tend to hold regardless of whether we are in a standard second price auction or a position auction, and regardless of whether the seller uses reserve prices. We now discuss several possible avenues for future research.

In this paper, we have considered multiple scenarios in which the bidders' values are drawn from asymmetric distributions, but have not considered what happens under position auctions with asymmetric bidders when each bidder has private information about his or her value. A natural question to ask is how the results would extend to position auctions when bidders' values are drawn from different distributions. In addition, throughout this paper we have restricted attention to cases in which the bidders have private values, and not considered models of common value auctions. There is little that is known about the equilibria of position auctions when there is a common component to bidders' values, so it is natural ask whether a seller would want to reveal information that affects a common component to bidder preferences in position auctions. We leave this question for future work.

Acknowledgments

We thank Pierre Grinspan, Vahab Mirrokni, Uri Nadav, Prabhakar Raghavan, and seminar attendees at Google and University of British Columbia for helpful discussions.

9. REFERENCES

[1] I. Abraham, S. Athey, M. Babaioff, and M. Grubb. Peaches, lemons, and cookies: designing auction markets with dispersed information. *Microsoft Research Typescript*, 2014.

[2] W. J. Adams and J. L. Yellen Commodity bundling and the burden of monopoly. *Quarterly Journal of Economics*, 90(3): 475–498, 1976.

[3] M. Armstrong A more general theory of commodity bundling. *Journal of Economic Theory*, 148(2): 448–472, 2013.

[4] Y. Bakos and E. Brynjolfsson. Bundling information goods: pricing, profits, and efficiency. *Management Science*, 45(12): 1613–1630, 1999.

[5] D. Bergemann and A. Bonatti. Targeting in advertising markets: implications for offline versus online media. *RAND Journal of Economics*, 42(3): 417–443, 2011.

[6] D. Bergemann and M. Pesendorfer. Information structures in optimal auctions. *Journal of Economic Theory*, 137(1): 580–609, 2007.

[7] K. Bhawalkar, P. Hummel, and S. Vassilvitskii. Value of targeting. In *Proceedings of the 7th International Symposium on Algorithmic Game Theory (SAGT)*, pages 194–205, 2014.

[8] S. Board Revealing information in auctions: the allocation effect. *Economic Theory*, 38(1): 125–135, 2009.

[9] J. I. Bulow and P. D. Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1): 180–194, 1996.

[10] L. E. Celis, G. Lewis, M. M. Mobius, and H. Nazerzadeh. Buy-it-now or take-a-chance: a simple sequential screening mechanism. In *Proceedings of the 20th International Conference on the World Wide Web (WWW)*, pages 147–156, 2011.

[11] I. Chakraborty. Bundling decisions for selling multiple objects. *Economic Theory*, 13(3): 723–733, 1999.

[12] A. De Cornière and R. De Nijs. Online advertising and privacy. *RAND Journal of Economics* (forthcoming).

[13] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: selling billions of dollars of keywords. *American Economic Review*, 97(1): 242–259, 2007.

[14] P. Eső and B. Szentes. Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies*, 74(3): 705–731, 2007.

[15] H. Fang and P. Norman. To bundle or not to bundle. *RAND Journal of Economics*, 37(4): 946–963, 2006.

[16] H. Fu, P. Jordan, M. Mahdian, U. Nadav, I. Talgam-Cohen, and S. Vassilvitskii. Ad auctions with data. In *Proceedings of the 5th International Symposium on Algorithmic Game Theory (SAGT)*, pages 184–189, 2012.

[17] J. Ganuza. Ignorance promotes competition: an auction model with endogenous private valuations. *RAND Journal of Economics*, 35(3): 583–598, 2004.

[18] J. Ganuza and J. S. Penalva. Signal orderings based on dispersion and the supply of private information in auctions. *Econometrica*, 78(3): 1007–1030, 2010.

[19] R. Gomes and K. Sweeney. Bayes-Nash equilibria of the generalized second price auction. *Games and Economic Behavior*, 86: 421–437, 2014.

[20] S. Hart and N. Nisan. Approximate revenue maximization with multiple items. *Hebrew University of Jerusalem Typescript*, 2012.

[21] P. Hummel. Position auctions with dynamic resizing. *Google Inc. Typescript*, 2014.

[22] P. Hummel and R. P. McAfee. When does improved targeting increase revenue? *Google Inc. Typescript*, 2015.

[23] P. Jehiel, M. Meyer-ter-Vehn, and B. Moldovanu. Mixed bundling auctions. *Journal of Economic Theory*, 134(1): 494–512, 2007.

[24] J. Levin and P. Milgrom. Online advertising: heterogeneity and conflation in market design. *American Economic Review: Papers and Proceedings*, 100(2): 603–607, 2010.

[25] R. P. McAfee and J. McMillan. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory*, 46(2): 335–354, 1988.

[26] R. P. McAfee, J. McMillan, and M. D. Whinston. Multiproduct monopoly, commodity bundling, and correlation of values. *Quarterly Journal of Economics*, 104(2): 371–383, 1989.

[27] P. R. Milgrom. *Putting Auction Theory to Work.* New York, Cambridge University Press, 2004.

[28] P. R. Milgrom and R. J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5): 1089–1122, 1982.

[29] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1): 58–73, 1981.

[30] T. Palfrey. Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica*, 51(2): 463–483, 1983.

[31] S. Tadelis and F. Zettelmeyer. Information disclosure as a matching mechanism: theory and evidence from a field experiment. *American Economic Review*, 105(2): 886–905, 2015.

[32] H. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6): 1163–1178, 2007.

[33] D. R. Vincent and A. Manelli. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory*, 127(1): 1–35, 2006.