When Does Improved Targeting Increase Revenue?

PATRICK HUMMEL, Google Inc.

R. PRESTON MCAFEE, Microsoft Corp.

In second-price auctions, 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 symmetric bidders with values drawn from a distribution with a monotone hazard rate. With asymmetries, improved targeting increases revenue if the most frequent winner wins less than 30.4% of the time under a model in which shares are well defined, but can decrease revenue otherwise. We derive analogous results for position auctions. Finally, we show that revenue can vary nonmonotonically with the number of bidders who are able to take advantage of improved targeting.

CCS Concepts: • **Information systems** → **Online auctions**; • **Applied computing** → **Economics**; • **Theory of computation** → *Computational advertising theory*;

Additional Key Words and Phrases: Targeting, revenue, online auctions, position auctions, advertising

ACM Reference Format:

Patrick Hummel and R. Preston McAfee. 2016. When does improved targeting increase revenue? ACM Trans. Econ. Comput. 5, 1, Article 4 (October 2016), 29 pages.
DOI: <http://dx.doi.org/10.1145/2956586>

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 for these concerns runs as follows. Improvements in targeting enable advertisers to more accurately identify consumers' interests. If a consumer's interests are so accurately identified that advertisers know there is only one product that this consumer would ever buy, this process could result in only a single advertiser who is willing to advertise to this consumer. This means that this advertiser could conceivably 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 are not interested in their products, the total demand for advertisements will go down. If the supply of advertising opportunities remains unchanged, revenue from 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 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. 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,

Authors' addresses: P. Hummel, Google Inc., 1600 Amphitheatre Parkway, Mountain View, CA 94043; email: phummel@google.com; R. P. McAfee, Microsoft Corp., One Microsoft Way, Redmond, WA 98052; email: preston@mcafee.cc.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

2016 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 2167-8375/2016/10-ART4 \$15.00

DOI: <http://dx.doi.org/10.1145/2956586>

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 consumers' 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 product?

The fact that the demand for advertisements to an individual consumer will not decline to one does not invalidate the argument that targeting might reduce revenue, however. 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, in which 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.

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 advertisers to learn their exact value for advertising to a consumer before deciding how much to bid. By contrast, when advertisers are unable to target, they only know that their value will be a random draw from some distribution, the distribution reflecting the different values that the advertisers 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, which Celis et al. [2011] indicate 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 and the bidders have private values that are independent conditional on the common value, then the results of this article 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.

¹Bergemann and Bonatti [2011] and Levin and Milgrom [2010] also note that such a trade-off is likely to arise as a result of improved targeting.

²In addition, we know from Milgrom and Weber [1982] that if there is a common component to bidder values, then the seller has an incentive to reveal this information when bidders are symmetric. Abraham et al. [2014] further discuss when information asymmetries in common-value auctions can lead to revenue loss.

In this environment, when advertisements are being sold via standard second-price auctions, we first demonstrate a result analogous to that in Board [2009], Ganuza and Penalva [2010], and Palfrey [1983], 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, at least under a particular model of bidder values in which shares are well defined.³

While standard second-price auctions for a single advertising opportunity are used by most publishers, we also consider the possibility of position auctions, as these are frequently used by search engines as well as a smaller number of publishers. We first present a new characterization of the properties of equilibria in generalized second-price auctions when buyers have private information about their own values. We then use these results to compare revenue under targeting and bundling in position auctions. Here, we find that targeting unambiguously decreases revenue when there are 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 with symmetric bidders, it could be the case that a seller's revenue may vary nonmonotonically 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 for bidders of intermediate strength.

Our article relates to two distinct strands of literature. First, our article 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, Fu et al. [2012] provide examples that illustrate that improving targeting may decrease revenue in a private-value auction and Ganuza [2004] 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. Bergemann and Valimaki [2006] consider the optimal information structure in a joint design problem in which there may be a direct tie between the information that the seller discloses

³In particular, we work with a model in which firms have many products, but in any particular auction, a firm can only advertise its best product, and the firm's value for advertising a particular product is independent and identically distributed. In this model, each firm's probability of winning the auction (share) is proportional to the number of products that the firm has.

and the mechanism that the seller then uses to sell the object. Esó and Szentes [2007] address 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. de Cornière and de Nijs [2016] consider questions related to how information disclosure affects the ultimate prices that advertisers would charge for their products. Finally, Bhawalkar et al. [2014] analyze the value of targeting data to advertisers, Ganuza and Penalva [2010] provide general methods of classifying the informativeness of signals to bidders in private-value auctions, and Tadelis and Zettelmeyer [2015] conduct field experiments analyzing the effect of information disclosure on wholesale auto auctions.⁴

The second related strand of literature is work analyzing when sellers would want to bundle goods and sell them together. Bakos and Brynjolfsson [1999] and Fang and Norman [2006] both study a standard bundling framework in which a monopolist considers selling bundles of goods to buyers without using auctions. Adams and Yellen [1976], Jehiel et al. [2007], and McAfee et al. [1989] 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. Chakraborty [1999] 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. Hart and Nisan [2012] study 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 not involve either of these approaches.⁵ McAfee and McMillan [1988] and Vincent and Manelli [2007] 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, Armstrong [2013] 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 article differs from the work on bundling goods in that very few of these papers consider models of bundling in an auction setting. Our article 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 article 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 these papers considers 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 article thus makes a number of new contributions to the literature on information provision by a mechanism designer and bundling.

2. THE MODEL

Each buyer $i \in \{1, 2, \dots, n\}$ 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 $[0, \bar{v}_i]$, where $\bar{v}_i$ may be infinite. These bidders compete in an auction, and bid either before their value is realized (bundling) or

⁴In addition, Ghosh et al. [2015] analyze incentives to share information when there is a risk of information leakage that enables an advertiser to target a user on a different publisher; Johnson and Myatt [2006] analyze a framework for considering questions related to general transformations in buyer demand, and Bergemann and Valimaki [2006] further survey the literature on information provision in mechanism design.

⁵In addition, Yao [2015] studies bundling in the context of reducing the k -item n -bidder auction with additive valuations to k -item 1-bidder auctions.

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; thus, $\int_0^\infty 1 - F_i(v) dv \geq \int_0^\infty 1 - F_{i+1}(v) dv$ for all i .

Throughout this article, we consider two possible auction formats in which the bidders may compete. 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 that we consider is a position auction. Position auctions differ from the setting considered earlier 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 nonincreasing 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 that 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 $k + 1^{\text{th}}$ -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 $k + 1^{\text{th}}$ -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 GSP auctions without clickability of ads that is considered in Edelman et al. [2007] and Varian [2007].

The second possibility that we consider is the Vickrey-Clarke-Groves (VCG) mechanism. Under VCG pricing, each advertiser pays a total cost equal to the externality that the advertiser 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 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 only one bidder 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 that 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. Hummel [2016] 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 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. Throughout our analysis of second-price auctions, we assume that bidders follow their weakly dominant strategies of bidding their (expected) values. Thus, under bundling, all bidders bid their expected values, the bidder with the highest expected value wins and pays the second-highest bid, and the seller's revenue 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 to v . Thus, if $v_{(2)}$ denotes the realization of the second-highest value, the distribution of this realization is given by

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

From this, it follows that the difference between the seller's expected revenue under targeting and bundling is

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

First, we illustrate that the insight in Ganuza and Penalva [2010] and Palfrey [1983] that the seller prefers bundling to targeting when there are two bidders extends to cases in which the values of the bidders are not drawn from identical distributions.

THEOREM 3.1. *Suppose that there are $n = 2$ bidders. Then, the seller prefers bundling to targeting.*

All proofs are in the appendix. 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

$$\Delta_n = \int_0^\infty F(v) - nF^{n-1}(v) + (n-1)F^n(v) dv.$$

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 article, we let $f(v)$ denote the density corresponding to the cumulative distribution function $F(v)$.

When at least four values are independently drawn from the same distribution, the second-highest of these values will typically be higher than the average value under certain regularity conditions. This insight is useful in proving the following result.

THEOREM 3.2. *Suppose that $F_i(v) = F(v)$ 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.*

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 hazard rate condition; thus, targeting dominates bundling when there are $n = 4$ bidders in this case.⁶

Similarly, when there are three bidders, the second-highest value is the median value; thus, the seller's expected revenue under targeting is just the expectation of the median value of three samples. The seller's expected revenue under bundling is the expected value of the bidders. Since the mean (expectation of the median of three samples) of a given distribution is greater than the expectation of the median of three samples (mean) if the density corresponding to that distribution is decreasing (increasing), we obtain the following result.

THEOREM 3.3. *Suppose that $F_i(v) = F(v)$ for all i and there are $n = 3$ bidders. Then, the seller prefers targeting to bundling if $f(v)$ is nondecreasing in v (but not constant) on its support, prefers bundling to targeting if $f(v)$ is nonincreasing in v (but not constant) on its support, and is indifferent between targeting and bundling if $f(v)$ is constant on its support.*

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 in which 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.

We conclude this section by noting when targeting would be preferred to bundling in a model in which each bidder's value is either equal to 0 or 1. This case is useful to model scenarios under which an impression is either valuable to an advertiser (when it converts) or not valuable (when the impression fails to convert), but there is little heterogeneity in the value of an impression conditional on the impression being valuable. This model is also important in that the worst-case analyses of many famous problems, such as the secretary problem, achieve their worst case under such a distribution.

THEOREM 3.4. *Suppose that there are $n \geq 3$ bidders whose values are independent and identically distributed draws from the Bernoulli distribution that takes on the value 1 with probability p . Then, there is some $p^*(n) \in (0, 1)$ such that the seller prefers targeting to bundling if and only if $p > p^*(n)$. Furthermore, $\lim_{n \rightarrow \infty} n^2 p^*(n) = 2$.*

Theorem 3.4 indicates that, when buyers' values are drawn from the same binomial distribution, there are always some values of p for which the seller prefers targeting to bundling, and there are also values of p for which the seller prefers bundling to targeting. However, the seller is more likely to prefer targeting to bundling if p is large, and in the limit as the number of players becomes large, the set of values of p for which the seller prefers bundling to targeting becomes arbitrarily small. Intuitively, this arises because of the following: when there are a large number of bidders, the probability that there will be at least two bidders who have a value of 1 for the object becomes arbitrarily close to 1, and the seller is fairly certain to obtain 1 unit of revenue if the seller allows targeting.

⁶For any finite number of players, there exists a distribution with sufficiently fat tails such that the seller prefers bundling to targeting. In particular, if each bidder's value is a random draw from the lognormal distribution with mean $\mu < 0$ and variance $\sigma^2 = -2\mu$, then, for any n , one can show that for sufficiently negative μ , the seller prefers bundling to targeting.

4. RESERVE PRICES

This section illustrates that, with the exception of the case in which there is one bidder, in symmetric settings, appropriate reserve prices favor targeting. Thus, when there are four or more bidders, targeting with reserve prices dominates bundling with reserve prices when bidder values are drawn from a distribution with a monotone hazard rate. However, we also show that reserve prices are not enough to overturn the conclusion that bundling is preferred to targeting when there are two bidders or the conclusion that there is no general result as to whether targeting is preferred to bundling when there are three bidders.

With symmetric bidders, the only case in which reserve prices make bundling relatively more favorable is the case in which there is only one bidder. Without reserve prices, the seller's revenue when there is only one bidder is zero regardless of whether the seller uses targeting or bundling. However, when the seller uses reserve prices, the seller can extract the entire surplus under bundling by setting a reserve equal to the bidder's expected value. By contrast, the seller cannot extract the entire surplus under targeting; thus, the seller will prefer bundling to targeting when there is one buyer and the seller can set a reserve price.

Although the introduction of reserve prices makes targeting relatively less favorable compared to bundling in the case in which there is only one bidder, when there are multiple bidders whose values are all drawn from the same distribution, the introduction of reserve prices can only make the situation better for 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 under bundling regardless of whether the seller uses a reserve price. However, reserve prices do increase the revenue from targeting. Nonetheless, it is still the case that the seller typically prefers bundling to targeting when there are two symmetric bidders.

THEOREM 4.1. *Suppose that there are $n = 2$ bidders, $F_1(v) = F_2(v) = F(v)$, and the density $f(v)$ is nonincreasing in v . Then, the seller prefers bundling to targeting with the optimal reserve.*

In addition to bundling still typically being optimal with two bidders, it is also the case that the seller will sometimes want to use bundling with three bidders. Although the seller now prefers targeting to bundling in the case in which 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.

OBSERVATION 4.1. *Suppose that 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.*

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 $\alpha_1, \dots, \alpha_n$ satisfying $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_n > 0$ such that $F_i(v) = F(v)^{\alpha_i}$ for all i . This model has practical relevance because a firm typically has many products that it might wish to advertise, but in any particular auction, the firm will only have an opportunity to advertise its best product. If the firm's value for advertising a particular product is an independent and identically distributed draw from the distribution $F(v)$, then the value of the firm's best product is a random draw from the distribution $F(v)^{\alpha_i}$ if the firm has α_i products.

This formulation is also 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 that 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 present a result that expresses the circumstances under which targeting is preferred to bundling as a function of the probabilities with which the bidders have the highest values.

LEMMA 5.1. *Suppose that $F_i(v) = F(v)^{\alpha_i}$ for some $\alpha_1, \dots, \alpha_n$ satisfying $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_n > 0$, α_2 (and A) are sufficiently large, $n \geq 3$, and $\frac{1-F(v)}{f(v)}$ is nonincreasing in v . Then, if $\frac{\alpha_2}{A} \leq (1 - \frac{\alpha_1}{A})(1 - \frac{\alpha_2}{A})^{\frac{1-\alpha_1/A}{\alpha_2/A}}$, the seller prefers targeting to bundling.*

The inequality in Lemma 5.1 is a function only of two variables, $\frac{\alpha_1}{A}$ and $\frac{\alpha_2}{A}$. 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 5.2. *The inequality in Lemma 5.1 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 $\frac{\alpha_1}{A} \rightarrow 1$, $y^*(\frac{\alpha_1}{A})/(1 - \frac{\alpha_1}{A}) \rightarrow 1$.*

This 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 improves revenue. Furthermore, as $\frac{\alpha_1}{A} \rightarrow 1$ and the strongest firm becomes arbitrarily strong, $y^*(\frac{\alpha_1}{A})/(1 - \frac{\alpha_1}{A}) \rightarrow 1$, indicating that the second-strongest firm can also become arbitrarily strong relative to the weaker firms and still ensure that targeting improves revenue. This makes sense intuitively. When the strongest firm becomes more dominant, the expected second price becomes lower, and there is a greater need to allow targeting to increase the chances that the strongest firm will be given a substantial challenge.⁷

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 in which $F_i(v) = F(v)$ for all i , and $F(v)$ has compact support $[0, \bar{v}]$. Throughout, we also focus on symmetric pure-strategy equilibria.

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 $\int_0^\infty 1 - F(v) dv$, and the unique symmetric pure-strategy equilibrium is for all

⁷These results all also hold regardless of whether the bidders know the distribution from which other bidders' values are drawn. Regardless of whether bidders know this distribution, a bidder's dominant strategy is to bid truthfully; this is all that is needed to derive the results in this section.

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 $\sum_{k=1}^s c_k \int_0^\infty 1 - F(v) dv$.

Next, we consider the case in which bidders bid after learning the realizations of their values (targeting). We assume that when bidders bid, they know their own values but not the bids or the values of any of the other bidders. The assumption that bidders do not know the values of the other bidders is logical because, in practice, there are myriad dimensions in which auctions will differ from one another. There will be auctions for users of different ages, genders, or geographical areas, auctions that take place on different times of day, days of the week, or months of the year, auctions that take place on different devices such as mobile, tablet, or desktop, and so on. The myriad possible targeting dimensions means that it is unlikely that any exact auction will ever repeat itself; thus, bidders are unlikely to know the other bidders' exact values at the time that they bid in an auction.

We first address the question of when there exists an equilibrium to this game under GSP when there is a reserve price r . In addressing this question, we make the simplifying assumption that bidders are restricted to making bids in discrete increments of ϵ for some small $\epsilon > 0$. This assumption 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). Under this assumption, we obtain the following result.

LEMMA 6.1. *If bidders must make bids in discrete increments, then there exists a pure-strategy equilibrium in GSP auctions in which each bidder i with value $v \geq r$ follows a strategy of making some bid $b_i(v) \in [r, \bar{v}]$ that depends only on the bidder's value v . Moreover, in this equilibrium, it is necessarily the case that $b_i(v)$ must be nondecreasing in v for all i , and any equilibrium must be equivalent to a pure-strategy equilibrium.*

This result illustrates that, even in the asymmetric model, there exists a monotonic equilibrium in pure strategies. It is worth noting that this result can also be extended to cases in which the ads have different quality scores. In some GSP auctions, each advertiser i is ranked in part by the advertiser's quality score γ_i which reflects the overall likelihood that users will want to click on bidder i 's ad. In these auctions, a bidder with the k^{th} -highest value of $\gamma_i b_i$ obtains the k^{th} position and pays a price-per-click equal to $\gamma_{(k+1)} b_{(k+1)} / \gamma_{(k)}$, where $b_{(k+1)}$ and $\gamma_{(k+1)}$ denote the bid and quality score of the bidder with the $k + 1^{th}$ -highest value of $\gamma_i b_i$. Nothing in the overall proof strategy used to prove this result requires the assumption that the ads have the same quality scores; thus, a substantively identical proof can be used to show that there also exists a pure-strategy equilibrium in monotonic strategies in an analogous model in which ads have different quality scores.

Next, we consider the case in which the players' values are all drawn from the same distribution. In this case, we can go further by noting that there exists a symmetric equilibrium of the form given in the previous theorem. In analyzing the symmetric case, we let $\bar{v}$ denote the upper bound of the support of the distribution of the players' values.

LEMMA 6.2. *Suppose that $F_i(v) = F(v)$ for all i and bidders must make bids in discrete increments. Then, there exists a symmetric pure-strategy equilibrium in GSP auctions in which each bidder i with value $v \geq r$ follows the same bidding strategy $b(v) \in [r, \bar{v}]$ that depends only on the bidder's value v . Moreover, in this equilibrium, it is necessarily the case that $b(v)$ must be nondecreasing in v .*

We now return to the case in which players make bids on a continuous scale. Although the previous results do not guarantee existence of a symmetric monotonic equilibrium when players may make bids along a continuous scale, we have shown that symmetric monotonic equilibria always exist in GSP auctions if players must make

bids in arbitrarily fine discrete increments. Furthermore, Gomes and Sweeney [2014] have 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.

In order to address this question, we must first derive expressions for the seller's revenue in the case in which the bidders are allowed to target. This is done in the following lemma.

LEMMA 6.3. *Suppose that $F_i(v) = F(v)$ 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 $n \int_r^v \sum_{k=1}^s c_k \binom{n-1}{k-1} (1 - F(v))^{k-1} F(v)^{n-k} (v - \frac{1-F(v)}{f(v)}) f(v) dv$.*

This result illustrates that there is a natural correspondence between the seller's expected revenue in an auction for a single slot and the seller's revenue in a GSP auction. In a standard private-value auction, 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 exactly the same as the seller's revenue in position auctions using VCG pricing. Hummel [2016] has characterized the seller's revenue in a more general class of position auctions under VCG pricing. In the special case of Hummel [2016] corresponding to the model considered in this article, the seller's revenue under VCG pricing is exactly the same as the seller's revenue in Lemma 6.3.

Now, we use this result to address the question of whether the seller prefers targeting or bundling in position auctions. 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 single-object auctions. 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 an intermediate number of players, it is ambiguous as to whether the seller prefers targeting to bundling. This result is formalized in the following theorem.

THEOREM 6.4. *Suppose that $F_i(v) = F(v)$ 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 greater than or equal to the optimal reserve (but less than the bids under bundling).⁸ 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 preferred to targeting for all values of c_k if and only if $n \leq n^*$.*

⁸Our proof of Theorem 6.4 makes use of the fact that $E[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} | v_{(k)} \geq r] Pr(v_{(k)} \geq r)$ is decreasing in k , where $v_{(k)}$ denotes the k^{th} -highest value. This is immediate when either $r = 0$ (and $Pr(v_{(k)} \geq r) = 1$) or when r is greater than or equal to the optimal reserve (and $v - \frac{1-F(v)}{f(v)} \geq 0$ for all $v \geq r$). However, this may not hold for other values of r because smaller values of k may mean that it is more likely that $v_{(k)} \geq r$ and $v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} < 0$ will be satisfied. Thus, our proof breaks down without this assumption, though we conjecture that it may be possible to generalize the results to other suboptimal reserves.

- (2) *There exists some $n^{**} > n^*$ such that targeting is preferred to bundling for all values of c_k 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 .*

In position auctions with symmetric bidders, the number of players never has any effect on seller revenues under bundling because the players always bid their (identical) expected values under bundling. However, 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 6.4.

The comparative statics results in Part (3) of Theorem 6.4 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 revenue compared to the bottom positions make targeting a better choice, whereas situations in which the bottom positions contribute a substantial percentage of revenue may make bundling a better choice. This gives the comparative statics results given in Part (3) of Theorem 6.4.

Finally, we present an example to give a sense of the values n^* and n^{**} that arise in Theorem 6.4. When the players' values are drawn from the uniform distribution and there is no reserve price, we obtain the following result.

OBSERVATION 6.1. *Suppose that 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 6.4 are $n^* = 3$ and $n^{**} = 2s + 1$.*

We close with one remark about the robustness of these results to modeling assumptions. Throughout this section, we have assumed that bidders do not know each other's values when they bid under targeting. However, the results of this section for the zero reserve price will hold even if bidders are able to learn the other bidders' values before bidding. If bidders know each other's values, then we know from Edelman et al. [2007] that, even under GSP, there exists an envy-free equilibrium in which the players' payoffs are the same as they would be in the dominant strategy equilibrium of the VCG mechanism. Under this equilibrium, the seller's revenue would be the same as it is under VCG for any realization of the targeting data. Thus, the seller's expected revenue unconditional on the realization of the targeting data is also the same as it would be under VCG. However, the expression we have given for the seller's expected revenue under targeting in Lemma 6.3 is equal to the seller's expected revenue under VCG. Thus, even if bidders learn each other's values before bidding under targeting, the substantive conclusions in Theorem 6.4 will hold if bidders follow the main equilibrium strategies considered in Edelman et al. [2007].

7. WHAT IF NOT ALL BIDDERS CAN TARGET?

So far in this article, we have compared scenarios in which all bidders can target with scenarios in which no bidders can target. While this an important baseline, there may also be settings in which 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 allowing only 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, i learns i 's value before placing a bid, but if bidder i is not able to target, then the bidder simply knows that the bidder's 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 . We consider both standard second-price auctions and generalized second-price auctions with no reserve price.

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 can 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 this will continue to hold if at least two bidders cannot target, this will not hold in general, as the following result illustrates.

THEOREM 7.1. *If bidders' values are drawn from the same distribution, 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. However, a seller's expected revenue from targeting will be monotonic in the number of bidders that can target if at least two bidders cannot target.*

When at least two bidders cannot target, the second-highest bid will always be at least as large as the expected value, but could be strictly larger if the second-highest bid of the bidders that can target is greater than this expected value. This second-highest bid of the bidders that can target will be larger, on average, when more bidders can target. Thus, a seller's expected revenue from targeting will be monotonic in the number of bidders that can target if at least two bidders cannot target.

However, the seller's revenue from targeting may be lower if only one bidder can target than if two bidders can target. This is especially likely to arise if the bidders' values are drawn from distributions with fat tails. In this case, if there are four bidders and only two bidders can target, the seller's expected revenue under targeting is larger than that of bundling. However, if three bidders can target, then it is very likely that these bidders will all learn that 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 affect only one bidder's estimate of the bidder's value for advertising to a certain user. First, we consider auctions for a single object.

THEOREM 7.2. *Suppose that the bidders' values are drawn from different distributions and 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.*

Next, we address how allowing just one bidder to target would affect seller revenues in position auctions. In analyzing how this would affect seller revenues in position auctions, we make the assumption that the seller's revenue is the same as it would be in the dominant-strategy equilibrium of the VCG mechanism. Edelman et al. [2007] have noted that, even in the GSP auction, there are settings in which there is always an envy-free equilibrium in which the players' payoffs are the same as they would be in the dominant-strategy equilibrium of the VCG mechanism. Since this is a natural equilibrium to focus on, we consider how allowing only one bidder to target would affect the seller's revenue in such an equilibrium.

THEOREM 7.3. *Suppose that only one bidder will be able to make use of certain targeting information. Then, if bidders follow an equilibrium of the position auction that results in the same revenue as the dominant-strategy equilibrium of the VCG mechanism, 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 $k \geq s + 2$ is the only bidder that can target.*
- (3) *If a bidder with the k^{th} -highest expected value for some $k \in [2, s + 1]$ is the only bidder that can target, there is no general result as to whether the seller prefers targeting to bundling.*

Together with Theorem 7.2, Theorem 7.3 suggests that it is not in a seller's interest to enable targeting if only the strongest bidder will be able to use 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 the insights on optimal auctions by Myerson [1981]. Myerson [1981] finds that, when asymmetric bidders are competing in an auction, a seller can improve its revenue by giving an artificial bonus to the weaker bidders. In this setting, the seller can likewise improve revenue when the weaker bidders have the advantage of being able to target.

8. EXPLORING ADS

In this section, we explore a connection between the circumstances under which improved targeting increases revenue and the circumstances under which exploring ads with unknown click-through rates would increase revenue. Often in online advertising, ads are ranked on the basis of the product of the bid that an advertiser has placed per click as well as a predicted click-through rate, which we refer to as an expected cost-per-1000-impressions or eCPM bid. While the predicted click-through rates are likely to be quite accurate for ads for which there is a lot of evidence about the click-through rate of the ad because the ad has been shown a large number of times, they may be less accurate for ads for which there is little that is known about the click-through rate of the ad because the ad has hardly been shown at all.

In this case, if the system always ranks the ads on the basis of their eCPM bids, then ads ranked below the top ad will never be shown and we will never learn the click-through rates of these ads. On the other hand, the system could try to explore ads for which the click-through rates of the ads are not known by sometimes showing these ads to learn more about their click-through rates. If one does enough exploration, then over the course of many auctions, one will eventually learn the click-through rates of all the ads with arbitrary precision.

There is a connection between the circumstances under which exploring the click-through rates of ads will increase long-run revenue relative to not doing any exploration in a Bayesian model for uncertain eCPMs and when improved targeting increases revenue. If one does not explore the click-through rates of the ads, then one will simply always show the ad that one expects to be best, and this advertiser will pay an average price equal to the expected second-highest eCPM bid. If one systematically explores the click-through rates of the ads, then one eventually learns the true eCPMs of all the ads and revenue will ultimately equal the actual second-highest eCPM bid. From this, we have the following result.

Remark 8.1. The circumstances under which exploring ads increases revenue in the long run are isomorphic to the circumstances under which improved targeting increases revenue.

An implication of this result is that if it is beneficial to improve one's estimates of the predicted click-through rates of the ads, then it is also beneficial to improve targeting. Similarly, improving targeting is beneficial if it is beneficial to improve one's estimates of the predicted click-through rates of the ads.

9. CONCLUSION

This article has analyzed when 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 close by discussing the robustness of our results to one possible modeling assumption. Throughout this article, we have assumed that the number of bidders in the auction is the same regardless of whether the seller uses bundling or targeting, but one might imagine that the number of bidders in the auction could change as a result of improved targeting. However, this possibility would have no effect on most of the results in this article. When bidders' values are all drawn from the same distribution, the seller's revenue under bundling is independent of the number of bidders in the auction, assuming that at least two bidders bid in the auction or there is a reserve price. Thus, if we interpret n to be the number of bidders in the auction under targeting, whether targeting is preferred to bundling is independent of the number of bidders under bundling. We therefore can immediately extend all of our results for the symmetric setting to a model in which the number of bidders may change as a result of improved targeting.

APPENDIX

PROOF OF THEOREM 3.1. The difference between the seller's expected revenue under targeting and bundling when there are $n = 2$ bidders is

$$\begin{aligned}\Delta_2 &= \int_0^\infty F_2(v) - \sum_{j=1}^2 \prod_{i \neq j} F_i(v) + (2-1) \prod_{j=1}^2 F_j(v) dv \\ &= \int_0^\infty -F_1(v) + F_2(v)F_1(v) dv = \int_0^\infty (F_2(v) - 1)F_1(v) dv < 0.\end{aligned}$$

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

Throughout the proofs of Theorems 3.2 and 3.3, we let $\bar{v}$ denote the upper bound of the support of $F(\cdot)$.

PROOF OF THEOREM 3.2. First, note that if $n \geq 4$, then $\phi(y) \equiv \frac{y^2}{2} + \frac{y^3}{3} + \dots + \frac{y^{n-1}}{n-1} - \frac{n-1}{n}y^n \geq 0$ for all $y \in [0, 1]$. $\phi(y) = y^n(\frac{y^{2-n}}{2} + \dots + \frac{y^{n-n}}{n} - 1)$; thus, $\phi \geq 0$ if and only if $\frac{y^{2-n}}{2} + \dots + \frac{y^{n-n}}{n} - 1 \geq 0$. Since $\frac{y^{2-n}}{2} + \dots + \frac{y^{n-n}}{n} - 1$ is decreasing in y , $\phi \geq 0$ for all $y \in [0, 1]$ if and only if $\frac{1}{2} + \dots + \frac{1}{n} - 1 \geq 0$, 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

$$\begin{aligned} \Delta_n &= \int_0^{\bar{v}} F(v) - nF^{n-1}(v) + (n-1)F^n(v) dv \\ &= \int_0^{\bar{v}} \left(\frac{1-F(v)}{f(v)} \right) (F(v) + F^2(v) + \dots + F^{n-2}(v) - (n-1)F^{n-1}(v)) f(v) dv \\ &= \left(\frac{1-F(v)}{f(v)} \right) \left(\frac{F^2(v)}{2} + \frac{F^3(v)}{3} + \dots + \frac{F^{n-1}(v)}{n-1} - \frac{n-1}{n}F^n(v) \right) \Big|_0^{\bar{v}} \\ &\quad - \int_0^{\bar{v}} \left(\frac{1-F(v)}{f(v)} \right)' \left(\frac{F^2(v)}{2} + \frac{F^3(v)}{3} + \dots + \frac{F^{n-1}(v)}{n-1} - \frac{n-1}{n}F^n(v) \right) dv \\ &= \left(\lim_{v \rightarrow \bar{v}} \frac{1-F(v)}{f(v)} \right) \left[\frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{n-1} - \frac{n-1}{n} \right] \\ &\quad - \int_0^{\bar{v}} \left(\frac{1-F(v)}{f(v)} \right)' \left[\frac{F^2(v)}{2} + \frac{F^3(v)}{3} + \dots + \frac{F^{n-1}(v)}{n-1} - \frac{n-1}{n}F^n(v) \right] dv. \end{aligned}$$

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

PROOF OF THEOREM 3.3. The difference between the seller's expected revenue under targeting and bundling when there are $n = 3$ bidders is

$$\begin{aligned} \Delta_3 &= \int_0^{\bar{v}} F(v) - 3F^2(v) + 2F^3(v) dv = \int_0^{\bar{v}} \frac{1}{f(v)} (F(v) - 3F^2(v) + 2F^3(v)) f(v) dv \\ &= \frac{1}{f(v)} \left(\frac{F^2(v)}{2} - F^3(v) + \frac{F^4(v)}{2} \right) \Big|_0^{\bar{v}} - \int_0^{\bar{v}} \left(\frac{1}{f(v)} \right)' \left(\frac{F^2(v)}{2} - F^3(v) + \frac{F^4(v)}{2} \right) dv \\ &= \frac{F^2(v)}{2f(v)} (1-F(v))^2 \Big|_0^{\bar{v}} - \int_0^{\bar{v}} \left(\frac{1}{f(v)} \right)' \frac{F^2(v)}{2} (1-F(v))^2 dv. \end{aligned}$$

Now, if $f(v)$ is nondecreasing in v on its support, then $\frac{F^2(v)}{2f(v)}(1-F(v))^2 \Big|_0^{\bar{v}} = 0$ and $\Delta_3 = - \int_0^{\bar{v}} \left(\frac{1}{f(v)} \right)' \frac{F^2(v)}{2} (1-F(v))^2 dv$. Thus, if $f(v)$ is nondecreasing in v (but not constant) 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 nonincreasing in v (but not constant) on its support and $\frac{F^2(v)}{2f(v)}(1-F(v))^2 \Big|_0^{\bar{v}} = 0$, then $\Delta_3 = - \int_0^{\bar{v}} \left(\frac{1}{f(v)} \right)' \frac{F^2(v)}{2} (1-F(v))^2 dv < 0$.

If $f(v)$ is nonincreasing in v and $\frac{F^2(v)}{2f(v)}(1-F(v))^2 \Big|_0^{\bar{v}} \neq 0$ (which implies that $\bar{v} = \infty$), 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 $\theta > 0$ such that $f(v)$ is not constant for all $v \leq \theta$, it is necessarily the case that $\frac{F^2(v|\theta)}{2f(v|\theta)}(1 - F(v|\theta))^2|_0^\theta = 0$, where $f(v|\theta)$ denotes the density corresponding to $F(v|\theta)$ and $\bar{v}|\theta$ denotes the upper bound on the support of $F(v|\theta)$. Moreover, $f(v|\theta)$ is nonincreasing in v (but not constant) on its support; thus, $\Delta_3 < 0$ 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 $F(\cdot)$.

However, 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 nonincreasing in v (but not constant) on its support, then $\Delta_3 < 0$ even if $\frac{F^2(v)}{2f(v)}(1 - F(v))^2|_0^{\bar{v}} \neq 0$. The result follows. $\square$

PROOF OF THEOREM 3.4. Under bundling, all bidders have an expected value of p ; thus, all bidders bid p and the seller obtains a revenue of p . Under targeting, if at least two bidders learn that they have a value of 1, these bidders will all bid 1, and the seller's revenue will be 1. However, if no more than one bidder learns that it has a value of 1, then all other bidders make a bid of 0, the second-highest bid will be 0, and the seller's revenue will be 0. Thus, under targeting, a seller's expected revenue is just the probability there will be at least two bidders who learn that they have a value of 1, or $1 - (1 - p)^n - np(1 - p)^{n-1}$.

From this, it follows that the seller prefers targeting to bundling if and only if $1 - (1 - p)^n - np(1 - p)^{n-1} \geq p \Leftrightarrow 1 - p \geq (1 - p)^n + np(1 - p)^{n-1} \Leftrightarrow 1 \geq (1 - p)^{n-1} + np(1 - p)^{n-2}$. Now, $g(p) = (1 - p)^{n-1} + np(1 - p)^{n-2} = (1 + (n - 1)p)(1 - p)^{n-2}$ satisfies $g(0) = 1$ and $g(1) = 0$ and $\frac{dg}{dp} = (n - 1)(1 - p)^{n-2} - (n - 2)(1 + (n - 1)p)(1 - p)^{n-3} = (n - 1)(1 - p)(1 - p)^{n-3} - (n - 2)(1 + (n - 1)p)(1 - p)^{n-3} = [(n - 1) - (n - 1)p - (n - 2) - (n - 2)(n - 1)p](1 - p)^{n-3} = [1 - (n - 1)^2 p](1 - p)^{n-3}$. However, the fact that $\frac{dg}{dp} = [1 - (n - 1)^2 p](1 - p)^{n-3}$ means that $\frac{dg}{dp} > 0$ if $p < \frac{1}{(n - 1)^2}$ and $\frac{dg}{dp} < 0$ if $p > \frac{1}{(n - 1)^2}$, meaning that $g(p)$ is initially increasing in p and then decreasing in p .

Combining this with the fact that $g(0) = 1$ and $g(1) = 0$ means that there is some $p^* \in (0, 1)$ for which $g(p^*) = 1$, and at this p^* , it must be the case that $g'(p^*) < 0$, $g(p) \geq 1$ for $p \leq p^*$, and $g(p) < 1$ for $p > p^*$. Thus, there is some $p^*(n) \in (0, 1)$ such that the seller prefers targeting to bundling if and only if $p > p^*(n)$.

Furthermore, since this $p^*(n)$ must satisfy $g(p^*(n)) = 1$, we must have that $(1 - p^*(n))^{n-1} + np^*(n)(1 - p^*(n))^{n-2} = 1$. This, in turn, implies that $\lim_{n \rightarrow \infty} p^*(n) = 0$ because, for any fixed $p^* \in (0, 1)$, $\lim_{n \rightarrow \infty} (1 - p^*)^{n-1} + np^*(1 - p^*)^{n-2} = 0$. Thus, it must be that $\lim_{n \rightarrow \infty} p^*(n) = 0$ in order for $(1 - p^*(n))^{n-1} + np^*(n)(1 - p^*(n))^{n-2} = 1$ to hold for all n .

It also must be the case that $\lim_{n \rightarrow \infty} np^*(n) = 0$. To see this, note that if there is some subsequence of $\{n\}_{n=3}^\infty$ for which $\lim_{n \rightarrow \infty} np^*(n) = \alpha > 0$ along this subsequence, then $\lim_{n \rightarrow \infty} (1 - p^*(n))^{n-1} + np^*(n)(1 - p^*(n))^{n-2} = \lim_{n \rightarrow \infty} (1 - \frac{\alpha}{n})^{n-1} + \alpha(1 - \frac{\alpha}{n})^{n-2} = e^{-\alpha} + \alpha e^{-\alpha} = (1 + \alpha)e^{-\alpha} \neq 1$ for $\alpha \neq 0$. Thus, in order for $(1 - p^*(n))^{n-1} + np^*(n)(1 - p^*(n))^{n-2} = 1$ to hold for all n , it must be the case that $\lim_{n \rightarrow \infty} np^*(n) = 0$.

Finally, it must be the case that $\lim_{n \rightarrow \infty} n^2 p^*(n) = 2$. To see this, note that

$$\begin{aligned} & (1 - p^*(n))^{n-1} + np^*(n)(1 - p^*(n))^{n-2} \\ &= 1 - (n - 1)p^*(n) + \frac{(n - 1)(n - 2)}{2} p^*(n)^2 + O(n^3 p^*(n)^3) \\ & \quad + np^*(n)(1 - (n - 2)p^*(n) + O(n^2 p^*(n)^2)) \\ &= 1 + p^*(n) - \frac{(n + 1)(n - 2)}{2} p^*(n)^2 + O(n^3 p^*(n)^3). \end{aligned}$$

Thus, in order for $(1 - p^*(n))^{n-1} + np^*(n)(1 - p^*(n))^{n-2} = 1$ to hold, it must be the case that $1 + p^*(n) - \frac{(n+1)(n-2)}{2} p^*(n)^2 + O(n^3 p^*(n)^3) = 1$ for large n . However, a necessary condition for this to hold is that $\lim_{n \rightarrow \infty} \frac{p^*(n)}{\frac{(n+1)(n-2)}{2} p^*(n)^2} = 1$ or $\lim_{n \rightarrow \infty} \frac{2}{n^2 p^*(n)} = 1$. Since this can only hold if $\lim_{n \rightarrow \infty} n^2 p^*(n) = 2$, we know that $\lim_{n \rightarrow \infty} n^2 p^*(n) = 2$. $\square$

PROOF OF THEOREM 4.1. We know from Bulow and Klemperer [1996] that when $f(v)$ is nonincreasing 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 3.3 that when $F_i(v) = F(v)$ for all i , and the density $f(v)$ is nonincreasing 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 follows that the seller prefers bundling to targeting when there are $n = 2$ bidders, even if the seller uses the optimal reserve price. $\square$

PROOF OF OBSERVATION 4.1. 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 3.3 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 that 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

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

Thus, the seller prefers bundling to targeting when the bidders' values are drawn from the exponential distribution. $\square$

PROOF OF LEMMA 5.1. The revenue gain from targeting is

$$\begin{aligned} \Delta_n &= \int_0^\infty F_2(v) - \sum_{j=1}^n \prod_{i \neq j} F_i(v) + (n-1) \prod_{j=1}^n F_j(v) dv \\ &= \int_0^\infty F^{\alpha_2}(v) - \sum_{j=1}^n \prod_{i \neq j} F^{\alpha_i}(v) + (n-1) \prod_{j=1}^n F^{\alpha_j}(v) dv \\ &= \int_0^\infty F^{\alpha_2}(v) - \sum_{j=1}^n F^{A-\alpha_j}(v) + (n-1) F^A(v) dv. \end{aligned}$$

Now, define $h(z)$ to be $h(z) \equiv z^{\alpha_2} - \sum_{j=1}^n z^{A-\alpha_j} + (n-1)z^A$. Note that $h(0) = h(1) = 0$. Also, since $n \geq 3$, we have $\alpha_2 < A - \alpha_j$ for all j ; thus, for very small values of $z > 0$, we have that $h(z) > 0$. Furthermore, $h'(1) = \alpha_2 - \sum_{j=1}^n (A - \alpha_j) + (n-1)A = \alpha_2 > 0$; thus, for values of z near 1, $h(z) < 0$. Combining these facts shows that there must be at least one value of $z \in (0, 1)$ for which $h(z) = 0$.

We now show that there is, in fact, only one value of $z \in (0, 1)$ for which $h(z) = 0$. To see this, note that, if $h(z) = 0$, $\sum_{j=1}^n z^{A-\alpha_j} = z^{\alpha_2} + (n-1)z^A$. Thus, for any such z , we have that

$$\begin{aligned}
 zh'(z) &= \alpha_2 z^{\alpha_2} - \sum_{j=1}^n (A - \alpha_j) z^{A-\alpha_j} + (n-1)Az^A \\
 &= \alpha_2 z^{\alpha_2} - n \sum_{j=1}^n \frac{1}{n} (A - \alpha_j) z^{A-\alpha_j} + (n-1)Az^A \\
 &\geq \alpha_2 z^{\alpha_2} - n \sum_{j=1}^n \frac{1}{n} (A - \alpha_j) \sum_{j=1}^n \frac{1}{n} z^{A-\alpha_j} + (n-1)Az^A \\
 &= \alpha_2 z^{\alpha_2} - \sum_{j=1}^n \frac{1}{n} (A - \alpha_j) \sum_{j=1}^n z^{A-\alpha_j} + (n-1)Az^A \\
 &= \alpha_2 z^{\alpha_2} - \frac{n-1}{n} A(z^{\alpha_2} + (n-1)z^A) + (n-1)Az^A \\
 &= \left(\alpha_2 - \frac{n-1}{n} A \right) z^{\alpha_2} + (n-1) \left(1 - \frac{n-1}{n} \right) Az^A \\
 &= \left(\alpha_2 - \frac{n-1}{n} A \right) z^{\alpha_2} + \frac{n-1}{n} Az^A = z^{\alpha_2} \left[\left(\alpha_2 - \frac{n-1}{n} A \right) + \frac{n-1}{n} Az^{A-\alpha_2} \right],
 \end{aligned}$$

where the inequality follows from the fact that $A - \alpha_j$ is increasing in j and $z^{A-\alpha_j}$ is decreasing in j . Therefore, the covariance between these terms, $\text{Cov}(A - \alpha_j, z^{A-\alpha_j})$, is nonpositive. Thus, $E[(A - \alpha_j)(z^{A-\alpha_j})] = E[A - \alpha_j]E[z^{A-\alpha_j}] + \text{Cov}(A - \alpha_j, z^{A-\alpha_j}) \leq E[A - \alpha_j]E[z^{A-\alpha_j}]$, meaning that $\sum_{j=1}^n \frac{1}{n} (A - \alpha_j) z^{A-\alpha_j} \leq \sum_{j=1}^n \frac{1}{n} (A - \alpha_j) \sum_{j=1}^n \frac{1}{n} z^{A-\alpha_j}$. The fifth line in these equations invokes the fact that $\sum_{j=1}^n z^{A-\alpha_j} = z^{\alpha_2} + (n-1)z^A$ when $h(z) = 0$.

Now, the sign of the final expression in this inequality for $zh'(z)$ is nondecreasing in z for $z > 0$. From this, it follows that if there is some $\hat{z} \in (0, 1)$ for which $h(\hat{z}) = 0$ and $h'(\hat{z}) > 0$ (thus, $h(z) > 0$ for values of z slightly greater than $\hat{z}$), then it is necessarily the case that, for any other values of $z > \hat{z}$ for which $h(z) = 0$, we must have that $h'(z) > 0$ as well. However, if $h(z) > 0$ for values of z slightly greater than $\hat{z}$, it is necessarily the case that the next smallest value of $z > \hat{z}$ satisfying $h(z) = 0$ also satisfies $h'(z) < 0$. Since there is at least one other value of $z > \hat{z}$ satisfying $h(z) = 0$ (when $z = 1$), this contradicts the fact that $h'(z) > 0$ for any values of $z > \hat{z}$ for which $h(z) = 0$. Thus, if there is some $\hat{z} \in (0, 1)$ for which $h(\hat{z}) = 0$, it must be the case that $h'(\hat{z}) < 0$.

However, if $h'(\hat{z}) < 0$ for any $\hat{z} \in (0, 1)$ satisfying $h(\hat{z}) = 0$, then it must be the case that $h(z) < 0$ for values of z slightly greater than $\hat{z}$, and the next smallest value of $z > \hat{z}$ satisfying $h(z) = 0$ also satisfies $h'(z) > 0$. Combining this with the results in the previous paragraph shows that if $h'(\hat{z}) < 0$ for any $\hat{z} \in (0, 1)$ satisfying $h(\hat{z}) = 0$, then $z = 1$ must be the next smallest value of $z > \hat{z}$ satisfying $h(z) = 0$. From this, it follows that there is exactly one value of $z \in (0, 1)$ for which $h(z) = 0$. Furthermore, if $\hat{z}$

denotes the unique $z \in (0, 1)$ for which $h(z) = 0$, then $h(z) > 0$ for values of $z \in (0, \hat{z})$ and $h(z) < 0$ for values of $z \in (\hat{z}, 1)$.

Now, by construction,

$$\begin{aligned}\Delta_n &= \int_0^\infty F^{\alpha_2}(v) - \sum_{j=1}^n F^{A-\alpha_j}(v) + (n-1)F^A(v) dv = \int_0^\infty h(F(v)) dv \\ &= \int_0^\infty \frac{1-F(v)}{f(v)} \frac{h(F(v))}{1-F(v)} f(v) dv \\ &= \frac{1-F(v)}{f(v)} \int_0^v \frac{h(F(y))}{1-F(y)} f(y) dy \Big|_0^\infty - \int_0^\infty \left(\frac{1-F(v)}{f(v)} \right)' \int_0^v \frac{h(F(y))}{1-F(y)} f(y) dy dv.\end{aligned}$$

Since there is some $\hat{z} \in (0, 1)$ for which $h(z) > 0$ for values of $z \in (0, \hat{z})$ and $h(z) < 0$ for values of $z \in (\hat{z}, 1)$, it follows that if $\int_0^v \frac{h(F(y))}{1-F(y)} f(y) dy > 0$ when $v = \infty$, then $\int_0^v \frac{h(F(y))}{1-F(y)} f(y) dy > 0$ also holds for all $v < \infty$. Combining this with the fact that $\frac{1-F(v)}{f(v)}$ is nonincreasing in v shows that if $\int_0^\infty \frac{h(F(y))}{1-F(y)} f(y) dy = \int_0^1 \frac{h(v)}{1-v} dv > 0$, we have that $\Delta_n > 0$.

Now,

$$\begin{aligned}\phi &= \int_0^1 \frac{h(v)}{1-v} dv = \int_0^1 \frac{v^{\alpha_2} - \sum_{j=1}^n v^{A-\alpha_j} + (n-1)v^A}{1-v} dv \\ &= \int_0^1 \frac{v^{\alpha_2} - v^A - \sum_{j=1}^n (v^{A-\alpha_j} - v^A)}{1-v} dv = \int_0^1 \sum_{i=\alpha_2}^{A-1} v^i - \sum_{j=1}^n \sum_{i=A-\alpha_j}^{A-1} v^i dv \\ &= \sum_{i=\alpha_2}^{A-1} \frac{1}{i+1} - \sum_{j=1}^n \sum_{i=A-\alpha_j}^{A-1} \frac{1}{i+1} = \sum_{i=\alpha_2+1}^A \frac{1}{i} - \sum_{j=1}^n \sum_{i=A-\alpha_j+1}^A \frac{1}{i}.\end{aligned}$$

Note that $\log(\frac{m+1}{k}) < \sum_{i=k}^m \frac{1}{i} < \log(\frac{m}{k-1})$. Thus, $\phi = \sum_{i=\alpha_2+1}^A \frac{1}{i} - \sum_{j=1}^n \sum_{i=A-\alpha_j+1}^A \frac{1}{i} > \log(\frac{A+1}{\alpha_2+1}) - \sum_{j=1}^n \log(\frac{A}{A-\alpha_j})$; therefore, if $\log(\frac{A+1}{\alpha_2+1}) - \sum_{j=1}^n \log(\frac{A}{A-\alpha_j}) > 0$, then $\phi > 0$ holds as well. Furthermore, $\log(\frac{A+1}{\alpha_2+1}) - \sum_{j=1}^n \log(\frac{A}{A-\alpha_j}) > 0 \Leftrightarrow \frac{A+1}{\alpha_2+1} > \prod_{j=1}^n \frac{A}{A-\alpha_j} \Leftrightarrow \frac{\alpha_2+1}{A+1} < \prod_{j=1}^n (1 - \frac{\alpha_j}{A})$. However, in the limit as α_2 (and A) become large, the difference between $\frac{\alpha_2}{A}$ and $\frac{\alpha_2+1}{A+1}$ becomes vanishingly small. Thus, if α_2 is sufficiently large and $\frac{\alpha_2}{A} < \prod_{j=1}^n (1 - \frac{\alpha_j}{A})$, it follows that $\Delta_n > 0$, and the seller prefers targeting to bundling.

Now, the minimizer of the function $\prod_{j=2}^n (1 - \frac{\alpha_j}{A})$ subject to the constraints that $\sum_{j=2}^n \alpha_j = A - \alpha_1$ and $\alpha_j \leq \alpha_2$ for all $j \leq 2$ is the same as the minimizer of the function $\log \prod_{j=2}^n (1 - \frac{\alpha_j}{A}) = \sum_{j=2}^n \log(1 - \frac{\alpha_j}{A})$, which is minimized when $\alpha_2 = \alpha_3 = \dots = \alpha_m$ for the largest value of m satisfying $\sum_{j=2}^m \alpha_j \leq A - \alpha_1$, $\alpha_{m+1} = A - \alpha_1 - (m-1)\alpha_2$, and $\alpha_j = 0$ for all $j > m+1$. Thus, the minimum possible value of $\prod_{j=2}^n (1 - \frac{\alpha_j}{A})$ subject to the constraints that $\sum_{j=2}^n \alpha_j = A - \alpha_1$ and $\alpha_j \leq \alpha_2$ for all $j \leq 2$ is greater than or equal to $(1 - \frac{\alpha_2}{A})^{\frac{A-\alpha_1}{\alpha_2}}$. From this, it follows that if $\frac{\alpha_2}{A} \leq (1 - \frac{\alpha_1}{A})(1 - \frac{\alpha_2}{A})^{\frac{A-\alpha_1}{\alpha_2}} = (1 - \frac{\alpha_1}{A})(1 - \frac{\alpha_2}{A})^{\frac{1-\alpha_1/A}{\alpha_2/A}}$, then the seller prefers targeting to bundling. $\square$

PROOF OF THEOREM 5.2. First, we show that there is some $y^* \in (0, 1 - \frac{\alpha_1}{A})$ such that the inequality in Lemma 5.1 is satisfied if and only if $\frac{\alpha_2}{A} \leq y^*$. Note that $\frac{\alpha_2}{A} \leq (1 - \frac{\alpha_1}{A})(1 - \frac{\alpha_2}{A})^{\frac{1-\alpha_1/A}{\alpha_2/A}}$ holds if and only if $\frac{\alpha_2/A}{1-\alpha_1/A} \leq (1 - \frac{\alpha_2}{A})^{\frac{1-\alpha_1/A}{\alpha_2/A}}$. 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 $(\beta/\gamma)^{(\beta/\gamma)} \leq 1 - \beta$. 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 5.1 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 5.1 is satisfied when $\frac{\alpha_2}{A} = \frac{\alpha_1}{A}$ if and only if $\frac{\alpha_1}{A} \leq (1 - \frac{\alpha_1}{A})(1 - \frac{\alpha_1}{A})^{\frac{1-\alpha_1/A}{\alpha_1/A}}$, which holds if and only if $\frac{\alpha_1/A}{1-\alpha_1/A} \leq (1 - \frac{\alpha_1}{A})^{\frac{1-\alpha_1/A}{\alpha_1/A}}$. 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 5.1 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 $\delta^* = 0.30366$; thus, the inequality in Lemma 5.1 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 that x^* is increasing in $\frac{\alpha_1}{A}$. Thus, for the key value of y^* earlier, 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 $\frac{\alpha_1}{A} \rightarrow 1$, $y^*(\frac{\alpha_1}{A})/(1 - \frac{\alpha_1}{A}) \rightarrow 1$ as well. $\square$

PROOF OF LEMMA 6.1. Note that an equilibrium to the related game in which each bidder i with value $v_i \geq r$ is restricted to making bids in the interval $[r, \bar{v}_i]$ is also an equilibrium of the original game. This is because of the fact that any strategy that ever involves making bids $b > \bar{v}_i$ is weakly dominated by a strategy that replaces all of these bids with a bid equal to $\bar{v}_i$. Thus, it suffices to demonstrate that there exists a symmetric equilibrium in the related game in which bidders are restricted to making bids in the interval $[r, \bar{v}_i]$.

Now, consider another game that differs from the just-mentioned game only in that players may now use distributional strategies, as defined in Milgrom and Weber [1985]. In this context, a distributional strategy for bidder i is a joint distribution over values and bids with the property that the marginal density over his values is f_i . Note that the set of distributional strategies for each bidder is a convex and compact metric space. Furthermore, for any given distributional strategies of the players, x and y , a player's payoff when the players use the distributional strategy $\alpha x + (1 - \alpha)y$ is linear in α for all $\alpha \in [0, 1]$. Thus each bidder's payoff function is quasiconcave in one's distributional strategy.

Further note that this game is also better-reply secure in the sense defined in Reny [1999] and Reny [2008]. The condition of better-reply security is trivially satisfied when bidders are restricted to making bids in discrete increments since a player's payoff is never discontinuous in the strategy choices of the players in the game in this case. Thus, this game is better-reply secure.

We know from Theorem 3.1 of Reny [1999], however, that for any game in which the players' strategy spaces are convex and compact metric spaces, the player's payoff function is quasiconcave, and the game is better-reply secure, there must exist a pure-strategy equilibrium. From this, it follows that there exists a pure-strategy equilibrium in the game in which players use distributional strategies.

To prove the result, it thus suffices to show that any equilibrium in distributional strategies must be equivalent to a pure-strategy equilibrium in which each bidder i follows a strategy of making some bid $b_i(v) \in [r, \bar{v}_i]$ that depends only on the bidder's value v . To see this, suppose by means of contradiction that there were a positive measure of values of v for which a bidder i randomized among bid choices when the bidder had a value of v . Let $G_i(b)$ denote the cumulative distribution function corresponding to the distribution of bids that i uses (unconditional on the realization of i 's value). Consider an alternative bidding strategy in which the distribution of bids that i uses (unconditional on the realization of i 's value) is still $G_i(b)$, but i instead uses a pure-strategy bidding function $b_i(v)$ that is nondecreasing in v . If bidder i uses this alternative strategy, then the probability bidder i finishes in the k^{th} position is the same as before for any k , but the average value per click that the bidder obtains in the circumstances under which the bidder obtains a click is greater than before since the bidder is now making higher bids when the bidder has a higher value. Since this would be a profitable deviation for bidder i , it follows that any equilibrium in distributional strategies must be equivalent to a pure-strategy equilibrium in which each bidder i follows a strategy of making some bid $b_i(v) \in [r, \bar{v}_i]$ that depends only on the bidder's value v .

Identical reasoning to that just given shows that any pure-strategy equilibrium in which each bidder i follows a strategy of making some bid $b_i(v) \in [r, \bar{v}_i]$ that depends only on the bidder's value v must be monotonic in the sense that $b_i(v)$ must be nondecreasing in v . If $G_i(b)$ denotes the cumulative distribution function corresponding to the distribution of bids that i uses (unconditional on the realization of i 's value), and the bidder employs an alternative bidding strategy in which the distribution of bids that i uses (unconditional on the realization of i 's value) is still $G_i(b)$ but i instead uses a pure-strategy bidding function $b_i(v)$ that is nondecreasing in v , then this would be a profitable deviation for bidder i by the same reasoning in the previous paragraph. From this, it follows that bidders must use monotonic bidding strategies in any pure-strategy equilibrium. $\square$

PROOF OF LEMMA 6.2. Consider the same related game described in the proof of Lemma 6.1 in which bidders are restricted to making bids in the interval $[r, \bar{v}]$ and players use distributional strategies, in which a distributional strategy for bidder i is a joint distribution over values and bids with the property that the marginal density over the bidder's values is f . We know from the proof of Lemma 6.1 that this is a game in which the players' strategy spaces are convex and compact metric spaces, the player's payoff function is quasiconcave, and the game is better-reply secure. From this, it follows that the game also satisfies the weaker conditions of being diagonally quasiconcave and diagonally better-reply secure that are defined in Reny [1999]. Thus, we know from Theorem 4.1 of Reny [1999] that this game possesses a symmetric pure-strategy equilibrium in which players use distributional strategies.

The same argument used in the proof of Lemma 6.1 to show that any equilibrium in distributional strategies must be a pure-strategy equilibrium in which each bidder follows a monotonic bidding strategy also applies in this less general setting. From this, it follows that there exists a symmetric pure-strategy equilibrium in which each bidder i follows the same bidding strategy $b(v) \in [r, \bar{v}]$ that depends only on the bidder's value v and that bidders necessarily use monotonic bidding strategies in this equilibrium. $\square$

PROOF OF LEMMA 6.3. If bidders are following a symmetric and strictly monotonic bidding strategy $b(v)$ in equilibrium, then a bidder with value $v \geq r$ wins the k^{th} position if and only if the bidder has the k^{th} -highest value, which happens with probability $\binom{n-1}{k-1}(1 - F(v))^{k-1}F(v)^{n-k}$. Now, we know from the Integral Form Envelope Theorem in Milgrom [2004] that if $u(b, v)$ denotes the expected utility that a bidder with value $v \geq r$

obtains from making a bid of b and $U(v) \equiv \sup_{b \in [r, \bar{v}]} u(b, v)$, then $U(v) = u(b(v), v) = u(b(r), r) + \int_r^v u_2(b(x), x) dx$, where $u_2(b, x)$ denotes the derivative of $u(b, v)$ with respect to v evaluated at $v = x$.

Now, if $p(v)$ denotes the expected payment that a bidder with value v makes in equilibrium, then we know that $u(b(v), v) = \sum_{k=1}^s c_k \binom{n-1}{k-1} (1 - F(v))^{k-1} F(v)^{n-k} v - p(v)$. We also know that $u(b(r), r) = 0$. Since $u_2(b(x), x) = \frac{d}{dv} \sum_{k=1}^s c_k \binom{n-1}{k-1} (1 - F(x))^{k-1} F(x)^{n-k} v - p(x)$ evaluated at $v = x$ or $\sum_{k=1}^s c_k \binom{n-1}{k-1} (1 - F(x))^{k-1} F(x)^{n-k}$, we have that $\int_r^v u_2(b(x), x) dx = \int_r^v \sum_{k=1}^s c_k \binom{n-1}{k-1} (1 - F(x))^{k-1} F(x)^{n-k} dx$. Combining these facts with the result in the previous paragraph shows that $\sum_{k=1}^s c_k \binom{n-1}{k-1} (1 - F(v))^{k-1} F(v)^{n-k} v - p(v) = \int_r^v \sum_{k=1}^s c_k \binom{n-1}{k-1} (1 - F(x))^{k-1} F(x)^{n-k} dx$.

From this, it follows that a bidder's expected payment if the bidder has a value of v is equal to $\sum_{k=1}^s c_k \binom{n-1}{k-1} (1 - F(v))^{k-1} F(v)^{n-k} v - \int_r^v \sum_{k=1}^s c_k \binom{n-1}{k-1} (1 - F(x))^{k-1} F(x)^{n-k} dx$, and a bidder's expected payment unconditional on the precise realization of the bidder's value is $\sum_{k=1}^s c_k \binom{n-1}{k-1} [\int_r^{\bar{v}} (1 - F(v))^{k-1} F(v)^{n-k} v f(v) dv - \int_r^{\bar{v}} \int_r^v (1 - F(x))^{k-1} F(x)^{n-k} f(v) dx dv] = \sum_{k=1}^s c_k \binom{n-1}{k-1} [\int_r^{\bar{v}} (1 - F(v))^{k-1} F(v)^{n-k} v f(v) dv - \int_r^{\bar{v}} \int_x^{\bar{v}} (1 - F(x))^{k-1} F(x)^{n-k} f(v) dv dx] = \sum_{k=1}^s c_k \binom{n-1}{k-1} [\int_r^{\bar{v}} (1 - F(v))^{k-1} F(v)^{n-k} v f(v) dv - \int_r^{\bar{v}} (1 - F(x))^{k-1} F(x)^{n-k} dx] = \int_r^{\bar{v}} \sum_{k=1}^s c_k \binom{n-1}{k-1} (1 - F(v))^{k-1} F(v)^{n-k} (v - \frac{1-F(v)}{f(v)}) f(v) dv$. Since the n bidders all make the same expected payments unconditional on the precise realizations of their values, it then follows that the seller's expected revenue in the generalized second-price auction equals $n \int_r^{\bar{v}} \sum_{k=1}^s c_k \binom{n-1}{k-1} (1 - F(v))^{k-1} F(v)^{n-k} (v - \frac{1-F(v)}{f(v)}) f(v) dv$. $\square$

PROOF OF THEOREM 6.4. The seller's expected revenue from targeting under the conditions of the theorem can be rewritten as $\sum_{k=1}^s c_k E[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} | v_{(k)} \geq r] Pr(v_{(k)} \geq r)$, 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 $n \rightarrow \infty$, $E[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} | v_{(k)} \geq r] Pr(v_{(k)} \geq r) \rightarrow \bar{v}$ for all k since in the limit as $n \rightarrow \infty$, $v_{(k)} \rightarrow \bar{v}$ and $\frac{1-F(v_{(k)})}{f(v_{(k)})} \rightarrow 0$ with probability arbitrarily close to 1 for all k . Thus, in the limit as $n \rightarrow \infty$, the expected revenue from the mechanism under targeting approaches $\sum_{k=1}^s c_k \bar{v}$. By contrast, under bundling, all bidders bid $w \equiv \int_0^{\bar{v}} 1 - F(v) dv < \bar{v}$, and the total expected revenue under bundling is $\sum_{k=1}^s c_k w < \sum_{k=1}^s c_k \bar{v}$. 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 $E[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} | v_{(k)} \geq r] Pr(v_{(k)} \geq r)$ 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 . However, we have seen that 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 preferred to bundling for all values of c_k 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. However, we have already seen 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 $n = 2$. 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 preferred to targeting for all values of c_k 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 preferred to bundling nor that bundling is preferred to targeting for all values of c_k . The seller's expected revenue under targeting is $\sum_{k=1}^s c_k E[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} | v_{(k)} \geq r] Pr(v_{(k)} \geq r)$, whereas the seller's expected revenue under bundling is $\sum_{k=1}^s c_k w$, where $w \equiv \int_0^{\bar{v}} 1 - F(v) dv < \bar{v}$. Thus, the difference between the seller's expected revenue under targeting and the seller's expected revenue under bundling is $\sum_{k=1}^s c_k (E[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} | v_{(k)} \geq r] Pr(v_{(k)} \geq r) - w)$. $E[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} | v_{(k)} \geq r] Pr(v_{(k)} \geq r)$, however, 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 $v - \frac{1-F(v)}{f(v)}$ is strictly increasing in v . Thus, there is some $k^* \in [1, s)$ such that $E[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} | v_{(k)} \geq r] Pr(v_{(k)} \geq r) > w$ if and only if $k \leq k^*$.

From this, it follows that the difference between the seller's expected revenue under targeting and the seller's expected revenue under bundling, $\sum_{k=1}^s c_k (E[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} | v_{(k)} \geq r] Pr(v_{(k)} \geq r) - w)$, is strictly decreasing in c_k for all $k > k^*$ and strictly increasing in c_k for all $k \leq k^*$. It then 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 $E[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} | v_{(k)} \geq r] Pr(v_{(k)} \geq r)$ is increasing in n for all k , the relevant value of $k^* \in [1, s)$ for which $E[v_{(k)} - \frac{1-F(v_{(k)})}{f(v_{(k)})} | v_{(k)} \geq r] Pr(v_{(k)} \geq r) > w$ 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 . $\square$

PROOF OF OBSERVATION 6.1. 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 drawn from the uniform distribution on $[0, 1]$, it is necessarily the case that $E[v_{(k+1)}] = 1 - \frac{k+1}{n+1} = \frac{n-k}{n+1}$; thus, 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 $\frac{1}{2}(c_1 + c_2 + c_3)$. Thus, 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. Therefore, the key value of n^* in Theorem 6.4 is $n^* = 3$.

Now, by Part (3) of Theorem 6.4, we know that if the seller's revenue under targeting is greater than the seller's revenue under bundling when $c_2 = \dots = c_s = c_1$, 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 $c_1 = c_2 = \dots = c_s$, 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 $\frac{s}{2} c_1$. Thus, 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 6.4 are $n^* = 3$ and $n^{**} = 2s + 1$, respectively. $\square$

PROOF OF THEOREM 7.1. We first show that if at least two bidders are unable to target, then the seller's expected revenue will be monotonic in the number of bidders that can target. To see this, let w denote the common bid that is made by the bidders that are unable to target, and let $v_{(2)}$ denote the second-highest bid of the bidders that are able to target. If at least two bidders are unable to target, then the seller's revenue will be $\max\{v_{(2)}, w\}$. But if $G(v_{(2)}|m)$ denotes the distribution of $v_{(2)}$ conditional on the number of bidders who can target m , then $G(v_{(2)}|m+1)$ first order stochastically dominates $G(v_{(2)}|m)$ for all m . Thus, $E[\max\{v_{(2)}, w\}]$ is increasing in m , and if at least two bidders are unable to target, the seller's expected revenue will be monotonic in the number of bidders that can target.

Now, we show a seller's expected revenue need not be monotonic in the number of bidders that can target in general. To see this, suppose that there are $n = 4$ bidders and each bidder's value is drawn from the lognormal distribution with parameters $\mu < 0$ and $\sigma^2 = -2\mu$. Note that if no bidders can 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 can 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 bidders that are able to target will learn that their values are greater than 1, these bidders 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 can target in the limit as $\mu \rightarrow -\infty$ and $\sigma^2 = -2\mu$. Note that if exactly one of the three bidders who can target learns that its value is greater than 1 and the other bidders who can target learn that their values are less than or equal to 1, then the seller's revenue will be exactly the 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 $\mu \rightarrow -\infty$ 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 $\mu \rightarrow -\infty$ 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

$\lim_{\mu \rightarrow -\infty} p(\mu) = 0$. Thus, when exactly three bidders are able to target, in the limit as $\mu \rightarrow -\infty$ and $\sigma^2 = -2\mu$, the probability that all three bidders who 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 $(1 - p(\mu))E[v|v \leq 1] + p(\mu)E[v|v > 1]$. We know that in the limit as $\mu \rightarrow -\infty$ when $\sigma^2 = -2\mu$, we have that $p(\mu) \rightarrow 0$ and $E[v|v \leq 1] \rightarrow 0$. Thus, since each bidder has an expected value of 1, it follows that in the limit as $\mu \rightarrow -\infty$ when $\sigma^2 = -2\mu$, we must have that $p(\mu)E[v|v > 1] \rightarrow 1$, meaning that $E[v|v > 1] = \Theta(\frac{1}{p(\mu)})$. However, 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 $\mu \rightarrow -\infty$. 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 $\mu \rightarrow -\infty$ when $\sigma^2 = -2\mu$.

We have seen, however, 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 $\mu \rightarrow -\infty$ 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. $\square$

PROOF OF THEOREM 7.2. If the bidder with the highest expected value is the only bidder that can target and this bidder learns that its value exceeds the second-highest expected value, then the seller's revenue is unaffected by targeting. But if this bidder learns that its 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 its 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 its 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 $w_{(3)} \rightarrow w_{(2)}$, where $w_{(2)}$ denotes the second-highest expected value, and $w_{(1)} \rightarrow \infty$, $\int_0^{w_{(3)}} w_{(3)} f_2(v) dv + \int_{w_{(3)}}^{w_{(1)}} v f_2(v) dv + \int_{w_{(1)}}^{\infty} w_{(1)} f_2(v) dv$ approaches $\int_0^{w_{(2)}} w_{(2)} f_2(v) dv + \int_{w_{(2)}}^{\infty} v f_2(v) dv > \int_0^{\infty} v f_2(v) dv = w_{(2)}$. In the limit as

$w_{(3)} \rightarrow 0$ and $w_{(1)} \rightarrow w_{(2)}$, $\int_0^{w_{(3)}} w_{(3)} f_2(v) dv + \int_{w_{(3)}}^{w_{(1)}} v f_2(v) dv + \int_{w_{(1)}}^\infty w_{(1)} f_2(v) dv$ 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. $\square$

PROOF OF THEOREM 7.3. If bidders follow an equilibrium of a position auction that results in the same payoffs as the VCG mechanism, then the seller's total revenue from the auction is $\sum_{k=1}^s k(c_k - c_{k+1})v_{(k+1)}$, where $v_{(k)}$ denotes the value of the bidder with the k^{th} -highest value and $c_{s+1} \equiv 0$. From this, it follows that if the bidder with the highest expected value is the only bidder that can target and this bidder learns that it has a value that is still greater than or equal to the second-highest expected value, then the value of $\sum_{k=1}^s k(c_k - c_{k+1})v_{(k+1)}$ is the same regardless of whether the seller allows targeting, and the seller's revenue is unaffected by targeting. But if the bidder with the highest expected value learns that it has a value that is less than the second-highest expected value, then the value of $\sum_{k=1}^s k(c_k - c_{k+1})v_{(k+1)}$ is lower under targeting than it would be under bundling, and allowing targeting decreases the seller's revenue. Thus, the seller strictly 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 $k \geq s+2$ is the only bidder that can target, and this bidder learns that it has a value that is still less than or equal to the $s+1^{\text{th}}$ -highest expected value, then the value of $\sum_{k=1}^s k(c_k - c_{k+1})v_{(k+1)}$ is the same regardless of whether the seller allows targeting, and the seller's revenue is unaffected by targeting. But if a bidder with the k^{th} -highest expected value for some $k \geq s+2$ instead learns that it has a value that is greater than the $s+1^{\text{th}}$ -highest expected value, then the value of $\sum_{k=1}^s k(c_k - c_{k+1})v_{(k+1)}$ is greater under targeting than it would be under bundling, and allowing targeting increases the seller's revenue. Thus, the seller strictly prefers targeting to bundling if a bidder with the k^{th} -highest expected value for some $k \geq s+2$ is the only bidder that can target.

Finally, if a bidder with the k^{th} -highest expected value for some $k \in [2, s+1]$ is the only bidder that can target, there is no general result as to whether the seller prefers targeting to bundling. If $c_j = c_{j+1}$ for all $j \leq k$, but $c_j > c_{j+1}$ for all other values of j , then the seller strictly prefers bundling to targeting by the same reasoning in the first paragraph of this proof. If $k \geq 3$ and $c_j = c_{j+1}$ for all $j \geq k-1$, but $c_j > c_{j+1}$ for all other values of j , then the seller strictly prefers targeting to bundling by the reasoning in the previous paragraph. If $k = 2$ and $c_j = c_{j+1}$ for all $j \geq 2$, but $c_1 > c_2$, then we know from Theorem 7.2 that there is no general result as to whether the seller prefers targeting to bundling. Thus, if $k \in [2, s+1]$, there is no general result as to whether the seller prefers targeting to bundling. $\square$

ACKNOWLEDGMENTS

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

REFERENCES

- Ittai Abraham, Susan Athey, Moshe Babaioff, and Michael Grubb. 2014. *Peaches, Lemons, and Cookies: Designing Auction Markets with Dispersed Information*. Typescript. Microsoft Research.
- William James Adams and Janet L. Yellen. 1976. Commodity bundling and the burden of monopoly. *Quarterly Journal of Economics* 90, 3, 475–498. DOI: <http://dx.doi.org/10.2307/1886045>
- Mark Armstrong. 2013. A more general theory of commodity bundling. *Journal of Economic Theory* 148, 2, 448–472. DOI: <http://dx.doi.org/10.1016/j.jet.2012.12.004>

- Yannis Bakos and Erik Brynjolfsson. 1999. Bundling information goods: Pricing, profits, and efficiency. *Management Science* 45, 12, 1613–1630. DOI: <http://dx.doi.org/10.1287/mnsc.45.12.1613>
- Dirk Bergemann and Alessandro Bonatti. 2011. Targeting in advertising markets: Implications for offline versus online media. *RAND Journal of Economics* 42, 3, 417–443. DOI: <http://dx.doi.org/10.1111/j.1756-2171.2011.00143.x>
- Dirk Bergemann and Juuso Valimäki. 2006. *Advances in Economics and Econometrics*. Cambridge University Press, Cambridge, UK, 186–221.
- Kshipra Bhawalkar, Patrick Hummel, and Sergei Vassilvitskii. 2014. Value of targeting. In *Proceedings of the 7th International Symposium on Algorithmic Game Theory (SAGT'14)*, Ron Lavi (Ed.), Vol. 7. Springer, Haifa, Israel, 194–205. DOI: http://dx.doi.org/10.1007/978-3-662-44803-8_17
- Simon Board. 2009. Revealing information in auctions: The allocation effect. *Economic Theory* 38, 1, 125–135. DOI: <http://dx.doi.org/10.1007/s00199-006-0177-4>
- Jeremy I. Bulow and Paul D. Klemperer. 1996. Auctions versus negotiations. *American Economic Review* 86, 1, 180–194.
- L. Elisa Celis, Gregory Lewis, Markus M. Möbius, and Hamid Nazerzadeh. 2011. 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'11)*, Elisa Bertino and Ravi Kumar (Eds.), Vol. 20. ACM, Hyderabad, India, 147–156. DOI: <http://dx.doi.org/10.1145/1963405.1963429>
- Indranil Chakraborty. 1999. Bundling decisions for selling multiple objects. *Economic Theory* 13, 3, 723–733. DOI: <http://dx.doi.org/10.1007/s001990050278>
- Alexandre de Cornière and Romain de Nijs. 2016. Online advertising and privacy. *RAND Journal of Economics* 47, 1, 48–72. DOI: <http://dx.doi.org/10.1111/1756-2171.12118>
- 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, 242–259. DOI: <http://dx.doi.org/10.1257/aer.97.1.242>
- Péter Eső and Balázs Szentes. 2007. Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies* 74, 3, 705–731. DOI: <http://dx.doi.org/10.1111/j.1467-937X.2007.00442.x>
- Hanming Fang and Peter Norman. 2006. To bundle or not to bundle. *RAND Journal of Economics* 37, 4, 946–963. DOI: <http://dx.doi.org/10.1111/j.1756-2171.2006.tb00065.x>
- 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)*, Maria Serna (Ed.), Vol. 5. Springer, Barcelona, Spain, 184–189. DOI: http://dx.doi.org/10.1007/978-3-642-33996-7_15
- Juan-José Ganuza. 2004. Ignorance promotes competition: An auction model with endogenous private valuations. *RAND Journal of Economics* 35, 3, 583–598.
- Juan-José Ganuza and José S. Penalva. 2010. Signal orderings based on dispersion and the supply of private information in auctions. *Econometrica* 78, 3, 1007–1030. DOI: <http://dx.doi.org/10.3982/ECTA6640>
- Arpita Ghosh, Mohammad Mahdian, R. Preston McAfee, and Sergei Vassilvitskii. 2015. To match or not to match: Economics of cookie matching in online advertising. *ACM Transactions of Economics and Computation* 3, 2, Article 7, 19 pages. DOI: <http://dx.doi.org/10.1145/2745801>
- Rentao Gomes and Kane Sweeney. 2014. Bayes-Nash equilibria of the generalized second-price auction. *Games and Economic Behavior* 86, 421–437. DOI: <http://dx.doi.org/10.1016/j.geb.2012.09.001>
- Sergiu Hart and Noam Nisan. 2012. *Approximate Revenue Maximization with Multiple Items*. Typescript. Hebrew University of Jerusalem.
- Patrick Hummel. 2016. Position auctions with dynamic resizing. *International Journal of Industrial Organization* 45, 38–46. DOI: <http://dx.doi.org/10.1016/j.ijindorg.2015.12.002>
- Phillipe Jehiel, Moritz Meyer ter Vehn, and Benny Moldovanu. 2007. Mixed bundling auctions. *Journal of Economic Theory* 134, 1, 494–512. DOI: <http://dx.doi.org/10.1016/j.jet.2006.02.001>
- Justin P. Johnson and David P. Myatt. 2006. On the simple economics of advertising, marketing, and product design. *American Economic Review* 96, 3, 756–784. DOI: <http://dx.doi.org/10.1257/aer.96.3.756>
- Jonathan Levin and Paul Milgrom. 2010. Online advertising: Heterogeneity and conflation in market design. *American Economic Review: Papers and Proceedings* 100, 2, 603–607. DOI: <http://dx.doi.org/10.1257/aer.100.2.603>
- R. Preston McAfee and John McMillan. 1988. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory* 46, 2, 335–354. DOI: [http://dx.doi.org/10.1016/0022-0531\(88\)90135-4](http://dx.doi.org/10.1016/0022-0531(88)90135-4)
- R. Preston McAfee, John McMillan, and Michael D. Whinston. 1989. Multiproduct monopoly, commodity bundling, and correlation of values. *Quarterly Journal of Economics* 104, 2, 371–383. DOI: <http://dx.doi.org/10.2307/2937852>

- Paul R. Milgrom. 2004. *Putting Auction Theory to Work*. Cambridge University Press, New York, NY.
- Paul R. Milgrom and Robert J. Weber. 1982. A theory of auctions and competitive bidding. *Econometrica* 50, 5, 1089–1122.
- Paul R. Milgrom and Robert J. Weber. 1985. Distributional strategies for games with incomplete information. *Mathematics of Operations Research* 10, 4, 619–632. DOI: <http://dx.doi.org/10.1287/moor.10.4.619>
- Roger B. Myerson. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1, 58–73. DOI: <http://dx.doi.org/10.1287/moor.6.1.58>
- Thomas Palfrey. 1983. Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica* 51, 2, 463–483. DOI: <http://dx.doi.org/10.2307/1912001>
- Philip J. Reny. 1999. On the existence of pure and mixed strategy equilibria in discontinuous games. *Econometrica* 67, 5, 1029–1056. DOI: <http://dx.doi.org/10.1111/1468-0262.00069>
- Philip J. Reny. 2008. *New Palgrave Dictionary of Economics* (2nd ed.). Palgrave Macmillan, New York, NY, 186–221.
- Steven Tadelis and Florian Zettelmeyer. 2015. Information disclosure as a matching mechanism: Theory and evidence from a field experiment. *American Economic Review* 105, 2, 886–905. DOI: <http://dx.doi.org/10.1257/aer.20110753>
- Hal Varian. 2007. Position auctions. *International Journal of Industrial Organization* 25, 6, 1163–1178. DOI: <http://dx.doi.org/10.1016/j.ijindorg.2006.10.002>
- Daniel R. Vincent and Alejandro Manelli. 2007. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory* 127, 1, 1–35. DOI: <http://dx.doi.org/10.1016/j.jet.2005.08.007>
- Andrew Chi-Chih Yao. 2015. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'15)*, Piotr Indyk (Ed.), Vol. 26. SIAM, San Diego, 92–109. DOI: <http://dx.doi.org/10.1137/1.9781611973730.8>

Received May 2015; revised September 2015; accepted December 2015