Int J Game Theory (2016) 45:1031–1052
DOI 10.1007/s00182-015-0501-y

ORIGINAL PAPER

Incentivizing advertiser networks to submit multiple bids

Patrick Hummel[1] · **R. Preston McAfee**[2] · **Sergei Vassilvitskii**[3]

Accepted: 6 September 2015 / Published online: 23 September 2015

Abstract This paper illustrates a method for making side payments to advertiser networks that creates an incentive for the advertiser networks to submit the second-highest bids they received to an ad exchange and simultaneously ensures that the publishers will make more money on average in the short run as a result of adopting this scheme. We also illustrate how this payment scheme affects publisher payoffs in the long run after advertisers have a chance to modify their strategies in response to the changed incentives of the mechanism.

Keywords Advertiser networks · Advertising exchanges · Auctions

1 Introduction

Online advertisements are commonly sold through ad exchanges that help bring advertisers and publishers together. Advertisers and publishers usually do not participate in the exchange directly, but instead are members of advertising networks, and ad networks are connected to each other through exchanges. There are hundreds of ad networks, though this number is diminishing, and a small number of major ad exchanges such as AppNexus, Facebook's Ad Exchange, Google's AdX, and OpenX. For a given

Electronic supplementary material The online version of this article (doi:10.1007/s00182-015-0501-y) contains supplementary material, which is available to authorized users.

✉ Patrick Hummel
phummel@google.com

[1] Google Inc., 1600 Amphitheatre Parkway, Mountain View, CA 94043, USA

[2] Microsoft Corp., One Microsoft Way, Redmond, WA 98052, USA

[3] Google Inc., 111 Eighth Avenue, New York, NY 10011, USA

opportunity to advertise, advertiser networks obtain bids from advertisers within their network, and these advertiser networks then submit a bid to the exchange. The advertiser with the highest bid runs an ad and pays the second-highest bid received by the exchange.

Typically advertiser networks on ad exchanges such as Google's AdX only submit the highest bid they have received from one of their advertisers to the exchange. By shielding other bids from the exchange, the networks hurt publisher revenue, as the winner only pays the second-highest *out of network* bid, instead of the second-highest bid overall. As concentration grows among ad networks, this lack of second bids becomes increasingly important.

The crux of the problem is that advertiser networks currently have little incentive to submit the second-highest bid they have received to the exchange, but they do have a strong incentive not to submit this bid. Since advertiser networks wish to provide a valuable service for their advertisers, these networks would like for their advertisers to be able to pay as low a price as possible for their advertisements. Submitting a second bid after the first bid can only increase the price that these advertisers might have to pay, so doing so would hurt the welfare of these advertisers. We address the question of whether one can modify the way payments are made on an ad exchange in such a way as to create incentives for each advertiser network to submit an additional bid beyond the highest bid that it received.

Specifically, we consider a payment scheme where the publisher shares some of the additional revenue obtained from the network submitting multiple bids with the network itself. If the network that supplies the highest bid submitted to the exchange is different from the network that supplies the second-highest bid submitted to the exchange, then no side payments are made. But if the network that supplies the highest bid is also the network that supplies the second-highest bid, then the winning advertiser pays the second-highest bid, the network that supplied the second-highest bid receives a payment equal to $\alpha(b_{(2)} - b_{(3)})$, and the publisher receives a payment equal to $(1-\alpha)b_{(2)} + \alpha b_{(3)}$, where $\alpha > 0$ is some positive constant, $b_{(2)}$ denotes the second-highest bid submitted to the exchange, and $b_{(3)}$ denotes the third-highest bid submitted to the exchange.

This payment scheme is advantageous because an advertiser network can now obtain a direct benefit from submitting the second-highest bid to the exchange. Networks may thus find it advantageous to submit the second-highest bids they have received to the exchange. But even if the networks fail to submit a second bid to the exchange, nothing will be lost by implementing this payment scheme. Whether a network submits a second bid to the exchange can only affect the payoffs of any of the players if the network supplies the highest bid submitted to the exchange. And if a network supplies the highest bid submitted to the exchange and the network does not submit a second bid to the exchange, then the outcomes and the payments for all the players will be the exact same as before this payment scheme was adopted.

The first purpose of this paper is to show theoretically that as long as each advertiser network values its own revenues more than it values the welfare of its advertisers, then it is possible to choose the value of α in such a way that all advertiser networks would prefer to submit multiple bids to the exchange and the publishers obtain greater payment in expectation for their inventory. This is a general theoretical result that holds for arbitrary distributions of advertiser values even if the distributions of advertiser values are different in the various networks and there may be arbitrary degrees of correlation in advertiser values within a network. However, this result does require the assumption that advertisers do not change their strategies in response to this change in the underlying mechanism.

In the second part of the paper, we address how the possibility that advertisers may want to change their strategies in response to the changes in the payments and actions of the advertiser networks affects the payoffs of the publishers. Here we show that once advertisers adjust their bidding strategies in response to the changed incentives of the mechanism, publishers may cease to earn more revenue as a result of the networks submitting multiple bids. The reason for this is that advertisers have an incentive to make higher bids when the advertiser networks are only submitting one bid to the advertising exchange, and the equilibrium reduction in advertiser bids that results from networks submitting multiple bids to the exchange may be sufficient to decrease the expected payoffs of the publishers. In fact, we show that the publishers will lose revenue as a result of adopting this payment scheme if the networks are all perfectly symmetric. However, if there are sufficient asymmetries in the networks, then the publishers may again increase their revenue as a result of this payment scheme, even if the advertisers strategically reduce their bids in response.

There has been relatively little work on the design of advertising exchanges. McAfee and Vassilvitskii (2012), Muthukrishnan (2009a, b), and Pai (2010) have all presented overviews of some of the challenges involved in design of advertising exchanges, and Mansour et al. (2012) presents an overview of the design of the Doubleclick Ad Exchange auction, but these surveys all neglect to mention the problem we address in this paper.

The few other existing papers on incentives in ad exchanges all address questions unrelated to those considered in this paper. For example, Cavallo et al. (2015) analyzes incentives of advertisers and arbitragers in a model in which advertisers prefer to pay per click but publishers prefer to be paid on a per impression basis and an arbitrager with a good estimate of click probabilities can absorb this risk and act as an intermediary. Feldman et al. (2010) analyzes equilibrium behavior in an auction with intermediaries and shows that revenue-maximizing intermediaries will use a randomized reserve price chosen from an interval. Stavrogiannis et al. (2013, 2014) analyze how advertisers would strategically choose which intermediary auction to participate in, and Balsiero et al. (2015) addresses questions related to optimal contracts for an advertiser network. But none of these papers analyze schemes for incentivizing advertiser networks to submit multiple bids to an ad exchange.

Other papers on ad exchanges do not address questions related to the incentives faced by ad networks in an ad exchange auction. Angel and Walfish (2013) presents methods for verifying that an auction on an ad exchange was conducted correctly. Balsiero et al. (2014) and Balsiero and Candogan (2015) address the problem of how a publisher can best meet the trade-off of how to allocate inventory between impressions filled by an ad exchange and traditional reservation based ad contracts. Chakraborty et al. (2010) considers the problem of which advertiser networks an ad exchange should solicit bids from when the ad exchange faces bandwith constraints. And Lang et al. (2011) discusses the problem of how to efficiently serve ads in an ad exchange when achieving low latency is important. These papers all differ significantly from our paper which instead focuses on the underlying incentives facing the advertisers and the ad networks in an ad exchange auction.

2 The model

There are $n \geq 2$ advertiser networks in the set $N = \{1, \ldots, n\}$, and each advertiser network i has a total of a_i advertisers. Advertisers submit bids for the right to place ads on a publisher through the networks, which in turn pass relevant information about bids they have received from advertisers to an advertising exchange. The advertising exchange keeps track of advertiser and bidding information submitted by the advertiser networks and ensures that an appropriate advertisement is displayed on the publisher's site.

The game proceeds as follows. First each network decides how many of the bids it receives that it will submit to the advertising exchange. In particular, each network i chooses an action s_i, where the action $s_i = k$ denotes an action in which network i submits all of the k highest bids it received to the exchange but does not submit any other bids to the exchange. The value of s_i that each network chooses must be an integer k satisfying $1 \leq k \leq a_i$.

After networks make this election, advertisers' values for an advertisement are then drawn from a distribution specific to the network. In particular, the values of the advertisers $1, \ldots, a_i$ in network i, $(v_1^i, \ldots, v_{a_i}^i)$, are drawn from some joint cumulative distribution function $F_i(.)$ that may differ for different networks. After advertisers learn their values, advertisers submit bids for an impression to their network.

If an advertiser has the highest bid submitted to the exchange, then the advertiser wins the auction and pays the second-highest bid that was submitted to the exchange. If a network supplied both the highest and the second-highest bid submitted to the exchange, then the network obtains a payment of $\alpha(b_{(2)} - b_{(3)})$ and the publisher obtains a payment of $(1 - \alpha)b_{(2)} + \alpha b_{(3)}$, where $\alpha \geq 0$ is some non-negative constant, and $b_{(j)}$ denotes the jth-highest bid submitted to the exchange. Note that the case of $\alpha = 0$ corresponds to a standard second-price auction. If a network did not supply both the highest bid and the second-highest bid submitted to the exchange, then the network does not obtain any payment and the publisher obtains a payment of $b_{(2)}$.

The payoffs of the players are as follows. If an advertiser with value v for the object wins the auction, then the advertiser obtains a payoff of $v - b_{(2)}$, but the advertiser obtains a payoff of 0 otherwise. A publisher's payoff is simply the payment that the publisher receives as a result of the auction. And a network's payoff is divided into two components. The network cares not only about the revenue it obtains from the mechanism, but also about the surplus of the advertisers in its network. Thus the final payoff of the network is equal to 0 if the network did not supply the highest bid submitted to the exchange, and $\alpha(b_{(2)} - b_{(3)})I(b_{(2)}) + \beta(v_{(1)} - b_{(2)})$ if the network supplies the highest bid submitted to the exchange, where $I(b_{(2)})$ is an indicator function that equals 1 if the network also supplies the second-highest bid submitted to the exchange and 0 otherwise, $\beta > 0$ denotes the weight the network places on the welfare of its advertisers relative to its own revenues, and $v_{(1)}$ denotes the value of the advertiser who made the highest bid submitted to the exchange.[1]

3 Short-run effects of revenue-sharing mechanism

In this section we consider what happens if advertisers do not adjust their strategies in response to any changes to the side payments made to advertiser networks or the strategies of the advertiser networks. This is a realistic approximation of what would happen in the short run since most advertisers only change their bidding strategies very rarely.

The main question we are interested in addressing is whether it is possible to use this payment scheme to create incentives for each network to submit both the highest and the second-highest bid it receives to the exchange and simultaneously ensure that the publisher obtains at least as large a payment in expectation as it would without using this scheme. This may be challenging because on one hand, it is necessary to choose the value of α to be large enough so that the networks will obtain enough extra revenue that they will be willing to submit the second-highest bids they received to the advertising exchange even though this will hurt the welfare of their advertisers. On the other hand, we do not want to choose a value of α that is too large because large values of α may hurt the expected revenue of the publisher. It is unclear whether it is possible to choose a value of α that is simultaneously large enough to create incentives for the networks to submit multiple bids to the advertising exchange but not so large that the publisher obtains lower payments in expectation when the networks follow these strategies.

Our first result illustrates that one can choose an appropriate such value of α quite generally without making any assumptions about the distributions from which the advertisers' values are drawn. The main condition that is necessary in order for such a value of α to exist is that the networks value their own revenues more than they value creating surplus for their advertisers (i.e. $\beta < 1$) (Fig. 1).

Proposition 1 *Suppose that* $\beta < \alpha < 1$. *Then it is a dominant strategy for all networks to submit the two highest bids they received to the advertising exchange. Furthermore, the publisher's expected utility is strictly increasing in the number of networks that submit a second bid to the exchange.*

Proof First note that networks will never have an incentive to submit more than the two highest bids they received to the advertising exchange. If a network submits the three highest bids it received to the advertising exchange instead of just the two highest bids it received, then the only circumstance under which this might affect the network's payoff is the circumstance under which the network had both the highest bid and the second-highest bid that was submitted to the exchange and was receiving a

[1] Another possible formulation would be to assume that the advertiser network's payoff is tied to the additional value this network provides each advertiser beyond the value this advertiser could have achieved on its own. In this case, the network's utility would be equal $\alpha(b_{(2)} - b_{(3)})I(b_{(2)}) + \beta \max\{v_{(1)} - b_{(2)}, u\}$ for some $u \geq 0$ that reflects the utility an advertiser could have obtained without the network. Using this alternative formulation would not affect any of the results of the paper.

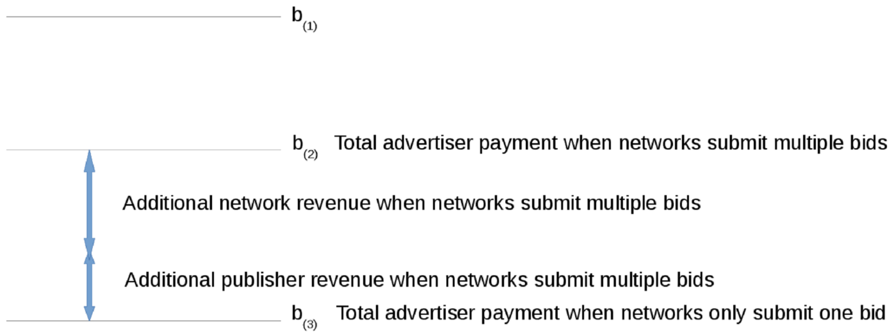

Fig. 1 Illustration of changes in advertiser payments, network revenue, and publisher revenue as a result of incentivizing advertiser networks to submit multiple bids when the highest and second-highest bids, $b_{(1)}$ and $b_{(2)}$, are both from the same network, and $b_{(3)}$ gives the highest bid from some other network

direct payment of $\alpha(b_{(2)} - b_{(3)})$. In this case, submitting a third bid to the advertising exchange can only lower the $\alpha(b_{(2)} - b_{(3)})$ payment that the network receives, so a network will never choose to submit a third bid to the exchange. From this it follows that submitting more than two bids to the exchange is always dominated by submitting exactly two bids to the exchange.

Next we note that if $\alpha > \beta$, then each network will always prefer to submit the two highest bids it received to the exchange than to only submit the highest bid it received. The only circumstance under which submitting a second bid to the exchange affects network i's payoff is the circumstance under which the second-highest bid submitted to network i is higher than any of the bids submitted to any of the other advertiser networks. Thus in deciding whether to submit a second bid to the exchange, network i can condition on the event that the second-highest bid submitted to network i is higher than any of the bids submitted to any of the other advertiser networks $j \neq i$.

But in this case, network i strictly prefers to submit a second bid to the exchange than to not submit this bid. Suppose b_j denotes the highest bid submitted to the exchange by some other advertiser network $j \neq i$, v denotes the value of the highest bidder in network i, and b_i denotes the second-highest bid submitted to network i. If network i does not submit a second bid to the exchange, then the network obtains a payoff of $\beta(v - b_j)$. And if network i does submit a second bid to the exchange, then the network obtains a payoff of $\alpha(b_i - b_j) + \beta(v - b_i)$. But since $b_i > b_j$, we have $(\alpha - \beta)(b_i - b_j) > 0$, meaning $\alpha(b_i - b_j) + \beta(v - b_i) > \beta(v - b_j)$. From this it follows that network i always prefers to submit a second bid to the exchange than to not submit this bid, and it is a dominant strategy for all networks to submit the two highest bids they received to the exchange.

Now we show that if $\alpha < 1$, then the publisher would prefer that network i submit a second bid to the exchange. To see this, note that the only circumstance under which network i's decision to submit a second bid to the exchange affects the publisher's payoff is the circumstance under which the second-highest bid submitted to network i is higher than any of the bids submitted to any of the other advertiser networks.

Suppose that this case holds, b_j denotes the highest bid submitted to the exchange by some other advertiser network $j \neq i$, and b_i denotes the second-highest bid submitted to network i. If network i does not submit a second bid to the exchange, then the publisher obtains a payoff of b_j. But if network i does submit a second bid to the exchange, then the publisher obtains a payoff of $(1 - \alpha)b_i + \alpha b_j > b_j$ for $\alpha < 1$. From this it follows that the publisher's payoff is greater if network i submits a second bid to the exchange than if this publisher does not submit a second bid to the exchange. And since this holds regardless of the actions taken by the other networks, it follows that the publisher's expected utility is strictly increasing in the number of networks that submit their second-highest bids to the exchange. □

A few remarks are in order about this result. First it is worth noting that even if some of the networks fail to follow their dominant strategy of submitting two bids to the exchange, the publisher will still be better off than before we adopted this payment scheme. A network's decision about whether to submit a second bid to the exchange has no effect on the payoffs of any of the players unless that network has the highest bid submitted to the exchange. And if a network has the highest bid submitted to the exchange and the network does not submit a second bid to the exchange, then the outcomes and the payments for all the players will be the exact same as before this payment scheme was adopted. Thus if there are some networks that fail to follow the dominant strategy of submitting a second bid to the exchange, then the publisher simply obtains the same payoff as before this payment scheme was adopted in cases where one of these networks supplies the winning bid, but the publisher still earns more money in expectation in the cases where the winning bid came from one of the networks that chose to submit multiple bids to the exchange.

Also, it should be noted that we have implicitly assumed that each network will adopt a policy of either always submitting a second bid to the exchange or never submitting a second bid to the exchange and will not condition this decision on the bids of the advertisers. However, if the networks decided after learning the values of the bids they received from the advertisers whether they wanted to submit a second bid to the exchange, the networks would still have a dominant strategy of always submitting their second-highest bids to the exchange. None of the reasoning in the proof of Proposition 1 depends on any assumptions about the information networks have about the bids of their advertisers.

Finally, we note that the scheme we have devised is the cheapest scheme that can be used to ensure that advertiser networks will have an incentive to submit the two highest bids that they have received to the advertising exchange. An advertiser network requires an additional payment of $\beta(b_{(2)} - b_{(3)})$ to be indifferent between submitting a second bid and not submitting a second bid when the advertiser network has the two highest bids, and requires no additional payment when the advertiser network does not have both the highest bid and the second-highest bid. In the special case of our payment scheme where $\alpha = \beta$, our payment scheme succeeds in paying the advertiser networks the exact amount they need to be indifferent between submitting a second bid and not submitting a second bid, so there is no cheaper scheme that will incentivize these networks to submit multiple bids.

4 Long-run effects

While the results in the previous section hold in the short run before advertisers adjust their bidding strategies in response to the changed incentives, it is less clear whether these results will continue to hold when advertisers change their bidding strategies in response to the changes in the underlying incentives that they face. As this is something that we would expect advertisers to do in the long run, it is important to address how the changed incentives for advertisers would affect the payoffs of the various players in equilibrium. Arguably the long-term effects are also more appealing and informative from a practical standpoint.

Before addressing the question of how advertiser bidding incentives would change as a result of networks submitting a second bid, we first address another related question. Could advertisers ever have an incentive to switch advertiser networks if one network submits multiple bids while others do not? Moving from a network that submits multiple bids to some other network that only submits a single bid can only decrease the price the advertiser has to pay in an auction, so an advertiser may have an incentive to switch advertiser networks if their original network does not have any market power.

While this may be the case if an advertiser network does not have any market power, many advertiser networks have some market power over the advertisers within their network for several reasons. First, advertiser networks generally provide several useful services for the advertisers within their network such as help with designing ad creatives, figuring out what to bid in various auctions, and providing targeting data that helps the advertisers bid more effectively in the auctions. Advertiser networks frequently have some sort of specialization which enables them to help some types of advertisers more than others. As a result, the advertisers typically have idiosyncratic utilities for joining each of the ad networks. On top of this, there is typically some sort of lock-in that makes it difficult to for an advertiser to switch advertiser networks. Joining a new advertiser network is costly as it requires the advertiser to set up an entirely new advertising campaign, so once an advertiser has joined a network, that advertiser will be hesitant to leave.

The combination of these factors means that an advertiser may very well have an incentive to remain within a particular network even if their network is submitting a second bid while other advertiser networks are not. Thus it seems realistic to expect advertisers to remain within the same network regardless of whether the network is submitting multiple bids to the exchange.

Now we address the question of endogenous advertiser bidding strategies. First note that if the advertiser networks all submit multiple bids to the advertising exchange, then the payoffs of the advertisers will always be the same as they would be in a standard second-price auction. The advertisers would thus have a weakly dominant strategy of always bidding their values if the advertiser networks submit multiple bids to the advertising exchange.

However, if the advertiser networks all only submit one bid to the advertising exchange, then it is not an equilibrium for advertisers to simply bid their values.[2] To

[2] However, advertisers will never have an incentive to make a bid that is less than their value. This is proven in Lemma 1 in the appendix.

see this, consider what would happen if an advertiser with value v changed his bid from $b = v$ to $b = v + \epsilon$ for some small $\epsilon > 0$. There are two ways this could potentially affect the advertiser's payoff.

First, assuming the distribution of the competing bids is drawn from a distribution that is Lipschitz continuous, with probability $\Theta(\epsilon)$, this change will mean that the advertiser will now have the highest bid of all advertisers, whereas there was an advertiser in some other network that had the highest bid when this advertiser was only bidding $b = v$. In this case, the advertiser loses an amount $\Theta(\epsilon)$ in expectation since the advertiser spends an amount no more than ϵ greater than his value in order to win the object.

Second, again with probability $\Theta(\epsilon)$, this change could mean that the advertiser would now have the highest bid of all advertisers even though there was an advertiser in the same network that had the highest bid when this advertiser was only bidding $b = v$. In this case, the advertiser would gain an amount $\Theta(1)$ in expectation since the advertiser would now win the object for a price that would be significantly less than v on average. Combining these results indicates that an advertiser would obtain a larger profit in expectation by bidding slightly more than his value.

The fact that advertisers would have an incentive to bid more than their value if advertiser networks all only submit one bid to the advertising exchange indicates that it may not necessarily be the case that the publishers continue to make more money when advertisers submit multiple bids to the exchange if advertisers follow equilibrium strategies. When advertiser networks submit multiple bids to the exchange, advertisers will respond by bidding less in equilibrium, and these reduced bids may undermine the result in Proposition 1 that indicated that the publisher would make more money when advertisers did not modify their strategies in response to the networks submitting multiple bids. In fact, under the condition that the advertiser with the highest value would always win the auction in equilibrium when the advertiser networks only submit one bid to the exchange, the publishers now make less money as a result of adopting this scheme, as the following proposition indicates:

Proposition 2 *Suppose that advertisers follow equilibrium strategies and the advertiser with the highest value would always win the auction if advertiser networks only submitted one bid to the exchange. Then the publisher's expected payoff is greater if the advertiser networks only submit one bid to the exchange than if the advertiser networks submit multiple bids.*

Proof If the advertiser with the highest value would always win the auction in equilibrium when the advertiser networks only submit one bid to the exchange, then the winning advertiser is the same regardless of whether the advertisers submit multiple bids to the exchange. From the Revenue Equivalence Theorem (Myerson 1981), it then follows that the expected payments of the advertisers in equilibrium are the same regardless of whether the advertiser networks only submit one bid to the exchange. However, if the advertiser networks only submit one bid to the exchange, then no side payments are ever made to the advertiser networks. By contrast, if the advertiser networks submit multiple bids to the exchange, then the advertiser networks obtain positive side payments in expectation. Since the expected payments of the advertisers are the same in either case, it then follows that the expected payoffs of the publishers are greater in equilibrium if the advertiser networks only submit one bid to the exchange than if the advertiser networks submit multiple bids to the exchange. □

Thus while there is a short-run benefit to the publisher from inducing advertiser networks to submit multiple bids to the exchange before advertisers have a chance to modify their strategies, once advertisers have an opportunity to modify their strategies in response to the changes in the underlying mechanism, the publisher may cease to make more money in the long run. As long as the advertiser with the highest value would always win the auction, publishers will not be able to make more money as a result of this scheme.

Since the key condition needed for this result is that the advertiser with the highest value would always win the auction when the advertiser networks submit one bid to the exchange, whether adopting this scheme will hurt publishers is likely to hinge on whether this assumption is likely to hold in practice. In general, there are at least some plausible circumstances under which this assumption will hold. For example, if there is a symmetry in the various advertiser networks, then there will indeed be an equilibrium in which the advertisers follow the same strictly monotonic bidding strategy, which in turn guarantees that the advertiser with the highest value will win the auction even if the advertiser networks only submit one bid to the exchange. This is illustrated in the following proposition:

Proposition 3 *Suppose each network has the same number of advertisers, and each advertiser's value is an independent and identically distributed draw from the cumulative distribution function $F(v)$ with corresponding probability density function $f(v)$ for which $f(v)$ is continuous in v and $\frac{f(v)}{F(v)}$ is decreasing in v. Then there is an equilibrium in which each advertiser with value v follows a strategy of bidding some amount $b(v)$, where $b(v)$ is a strictly increasing function of v, if the advertiser networks all submit one bid to the exchange.*

All remaining proofs are in the appendix. While Proposition 3 guarantees that the advertiser with the highest value will win the auction whenever there is symmetry in the advertisers in the various networks, this need not hold when the networks are not symmetric. The case when advertiser networks are asymmetric is important since, as we have noted, advertiser networks frequently have some specialization that enables them to provide better data or matching for certain types of advertising opportunities than others, and this is likely to induce asymmetries in the various networks. We thus address the case where the advertiser networks are asymmetric next.

When the networks are asymmetric, in general the advertisers in the different networks will follow different bidding strategies, which in turn implies that there will be at least some circumstances under which the advertiser with the highest value will not win the auction. In fact, when the networks are asymmetric, efficiency will always be greater when the networks submit multiple bids to the exchange than when the networks only submit one bid to the exchange. When the networks submit multiple bids to the exchange, all advertisers have an incentive to bid truthfully, so the advertiser with the highest value will win the auction. But when the networks only submit one bid to the exchange, the advertisers follow asymmetric bidding strategies, and there will be some circumstances in which the advertiser with the highest value fails to win the auction, thus resulting in an efficiency loss.[3] It is interesting to ask how large these efficiency losses are in some natural examples.

To address this, we consider a simple setting in which there is one advertiser network with m advertisers and a second advertiser network with exactly one advertiser. We also suppose that the advertisers' values are all independent and identically distributed draws from the same distribution, and for simplicity we focus on two distributions, an exponential distribution and a generalized Pareto distribution. In this case, if the first network does not submit multiple bids, then advertisers in the first network will follow a strategy of bidding according to some function $b(v)$ satisfying $b(v) \geq v$ for all v, while advertisers in the second network will bid their value.

Proposition 4 *Suppose there is one advertiser network with m advertisers, a second advertiser network with one advertiser, and the advertisers' values are all independent and identically distributed draws from the distribution $F(\cdot)$. Then if the first network only submits one bid to the exchange, the advertisers in the first network bid according to a function $b(v)$ satisfying $b(0) = 0$, $b'(0) = \frac{2m}{m+1}$ and*

$$b'(v) = \frac{(m-1)((v-1)e^{b(v)} + b(v) - v + 1)}{(e^v - 1)(b(v) - v)}$$

if $F(\cdot)$ corresponds to the exponential distribution $F(v) = 1 - e^{-v}$ and

$$b'(v) = \frac{(m-1)b(v)(1+b(v))[v(2+b(v)) - b(v)]}{v(1+v)(2+v)(b(v) - v)}$$

if $F(\cdot)$ corresponds to the generalized Pareto distribution $F(v) = 1 - \frac{1}{(1+v)^2}$ for values of $v > 0$.

The differential equations in Proposition 4 do not admit analytic solutions, though we prove in the electronic supplementary material that a solution to each of these differential equations must exist. However, we can use the characterizations in Proposition 4 to approximate the solution to $b(v)$ numerically by using the Euler method. To do this, we consider discrete increments of 0.001 and recursively approximate the value of $b(v+0.001)$ as $b(v)+0.001b'(v)$ for all v.[4] Once we have this approximation to the equilibrium bidding strategies of the advertisers in the first network when the first network only submits one bid to the exchange, we can then compare efficiency and revenue when this network submits multiple bids to the exchange to when it does not.

[3] This in turn implies that there will be no revenue equivalence between the situations where the networks do or do not submit multiple bids. The fact that there is no revenue equivalence between different mechanisms when there are asymmetries in the bidders also holds more generally. See Krishna (2010) for more discussion on this point and Gavious and Minchuk (2014) and Mares and Swinkels (2014) for two recent treatments.

[4] We also prove in the electronic supplementary material that the method of recursively approximating the value of $b(v + \delta)$ as $b(v) + \delta b'(v)$ converges to the true solution to the differential equation in the limit as $\delta \to 0$.

To do this, we conducted 100,000 simulations in which for each simulation, we considered draws of the values of the advertisers in the various advertiser networks. We computed the outcomes of the auction that would result both when the first advertiser network only submits one bid to the exchange (and the advertisers in this network thus follow the strategy of bidding according to the function $b(v)$), and when the first advertiser network submits its two highest bids to the exchange (and the advertisers in this network thus bid truthfully). For each simulation, we computed both whether the bidder with the highest value won the auction, the resulting efficiency or value of the winning advertiser, and the amount that this advertiser had to pay. We used these results to then compute the average increase in efficiency, advertiser payments, and the fraction of time the bidder with the highest value wins the auction as a result of the first advertiser network submitting multiple bids to the exchange. The results of the simulations are given in Tables 1 and 2.

Several patterns in these tables are worth noting. First, both the size of the efficiency lift and the percentage increase in the fraction of the time that the bidder with the high value wins the auction vary non-monotonically with the number of bidders in the first advertiser network. This makes sense intuitively. When the number of bidders in the first advertiser network is small, these bidders only have a relatively modest incentive to increase their bids above their values, and the inefficiencies resulting from the high bidder failing to win the auction will be relatively modest. When there are a larger number of bidders in the first network, these bidders have a greater incentive to increase their bids above their values, thereby increasing the chances that the bidder in the second network will fail to win the auction even if this bidder has the high value. However, when the number of bidders in the first network becomes exceedingly large, the highest bidder in the first network usually has a higher value than the bidder in the

Table 1 Percentage increase in efficiency, advertiser payments, and the fraction of time the bidder with the highest value wins the auction from the first advertiser network submitting multiple bids when there is one network with multiple bidders and a second network with one bidder (standard errors in parentheses)

Number of bidders	High-bidder win-rate lift	Efficiency lift	Payment lift	Maximum α
2	16.50 % (0.20 %)	2.27 % (0.03 %)	7.39 % (0.20 %)	0.345 (0.007)
3	19.35 % (0.13 %)	4.28 % (0.05 %)	21.40 % (0.23 %)	0.571 (0.005)
4	18.41 % (0.09 %)	4.87 % (0.06 %)	35.66 % (0.27 %)	0.727 (0.003)
5	16.79 % (0.06 %)	4.91 % (0.06 %)	48.59 % (0.30 %)	0.772 (0.003)
6	15.07 % (0.04 %)	4.50 % (0.05 %)	62.11 % (0.31 %)	0.828 (0.002)
7	13.43 % (0.03 %)	4.16 % (0.05 %)	73.07 % (0.33 %)	0.860 (0.002)
8	12.00 % (0.02 %)	3.76 % (0.05 %)	83.80 % (0.33 %)	0.885 (0.002)
9	10.83 % (0.02 %)	3.29 % (0.05 %)	93.72 % (0.34 %)	0.906 (0.001)
10	9.86 % (0.01 %)	2.88 % (0.04 %)	103.71 % (0.34 %)	0.922 (0.001)

The last column reports the maximum fraction of the revenue increase that can be shared with the advertiser networks without decreasing the publisher's revenue. These results are based on 100,000 simulations in which all advertiser values are independent and identically distributed draws from the exponential distribution $F(v) = 1 - e^{-v}$

Table 2 Percentage increase in efficiency, advertiser payments, and the fraction of time the bidder with the highest value wins the auction from the first advertiser network submitting multiple bids when there is one network with multiple bidders and a second network with one bidder (standard errors in parentheses)

Number of bidders	High-bidder win-rate lift	Efficiency lift	Payment lift	Maximum α
2	13.03 % (0.20 %)	1.78 % (0.08 %)	3.83 % (0.24 %)	0.166 (0.017)
3	16.64 % (0.14 %)	4.19 % (0.15 %)	13.63 % (0.45 %)	0.347 (0.015)
4	16.29 % (0.10 %)	5.88 % (0.18 %)	23.64 % (0.62 %)	0.450 (0.013)
5	15.21 % (0.07 %)	6.18 % (0.16 %)	37.92 % (0.63 %)	0.569 (0.010)
6	13.91 % (0.06 %)	6.14 % (0.15 %)	51.95 % (0.66 %)	0.652 (0.008)
7	12.66 % (0.04 %)	6.23 % (0.17 %)	63.52 % (0.77 %)	0.693 (0.008)
8	11.49 % (0.03 %)	6.10 % (0.17 %)	74.77 % (0.85 %)	0.728 (0.007)
9	10.53 % (0.03 %)	5.56 % (0.15 %)	88.72 % (0.85 %)	0.769 (0.006)
10	9.59 % (0.02 %)	4.99 % (0.14 %)	101.92 % (0.83 %)	0.801 (0.005)

The last column reports the maximum fraction of the revenue increase that can be shared with the advertiser networks without decreasing the publisher's revenue. These results are based on 100,000 simulations in which all advertiser values are independent and identically distributed draws from the generalized Pareto distribution $F(v) = 1 - \frac{1}{(1+v)^2}$

second network anyway, so inducing these bidders to bid truthfully by submitting a second bid to the exchange has little potential to change which ad wins the auction. Thus the efficiency lift from submitting multiple bids to the exchange initially increases in the number of bidders in the first advertiser network and then decreases.

It is also notable that the percentage increase in advertiser payments from submitting a second bid to the exchange is substantial for all numbers of bidders, and increases considerably as the number of advertisers in the first advertiser network increases. This again makes sense intuitively. When there are a large number of bidders in the first advertiser network, the second-highest value of these advertisers is likely to be considerably higher than the value of the advertiser in the second advertiser network, and considerable revenue gains can be achieved by inducing this advertiser network to submit multiple bids.[5] By contrast, smaller increases in advertiser payments will result from submitting multiple bids if there are just a small number of bidders in the first advertiser network because the second-highest value of these advertisers typically will not exceed that of the advertiser in the second advertiser network by quite as much.

Our simulations also indicate that a substantial fraction of the revenue lift that is achieved as a result of submitting a second bid to the advertising exchange can be shared with the advertiser network without causing this increase in advertiser payments to no longer increase publisher revenue. Thus when there are asymmetries in the

[5] For any distribution $F(\cdot)$, if the number of advertisers in the first advertiser network, m, is large enough, then the fractional increase in revenue that results from the first advertiser network submitting a second bid will be arbitrarily close to $E[v_{(2)}|m]/E[v|v \sim F]$, where $E[v_{(2)}|m]$ denotes the expected value of the second-highest of m draws from the distribution $F(\cdot)$. This becomes arbitrarily large in the limit as $m \to \infty$ if $F(\cdot)$ has unbounded support, so for any fixed α, if the number of advertisers in the first advertiser network is large enough, using this incentive scheme will increase publisher revenue.

networks, using this scheme to incentivize advertiser networks to submit multiple bids can increase efficiency, publisher revenue, and the welfare of the advertiser networks by judiciously choosing the fraction α of the revenue lift that will be shared with the advertiser networks.

5 Other strategies for exchanges and networks

In this section we discuss some other possible strategies for ad exchanges and advertiser networks and how these possibilities would affect the results of this paper. One possible alternative strategy for an ad exchange would be to require that a network submit a minimum of two bids to the auction in order for any of this advertiser network's bids to be considered. However, making this change without also making side payments would not incentivize the advertiser networks to submit the two highest bids they have received to the exchange. Instead an advertiser network would now have an incentive to submit the lowest bid it received in addition to also submitting the highest bid the network received. The reason the network would choose to submit the lowest bid rather than the second-highest bid is that if there are no side payments, then the advertiser network solely wishes to maximize the utility of its advertisers, and the advertiser network will minimize the price the advertiser in its network will have to pay if this network submits the lowest bid rather than the second-highest bid. Thus requiring the advertiser networks to submit multiple bids to the exchange would not incentivize the networks to submit their two highest bids.

The fact that an advertiser network could sometimes have an incentive to submit a second bid that does not equal the second-highest bid also raises the question of whether an advertiser network might submit a bid that does not correspond to any advertiser within its network at all. For instance, in the mechanism we have proposed, if the top bid within an advertiser network equals b, this network could maximize the side payment made to the network by submitting a second bid equal to $b - \epsilon$ for some arbitrarily small $\epsilon > 0$. Might an advertiser network have an incentive to submit a second bid that is higher than the second-highest bid it actually received?

While following such a strategy might enable the network to increase the side payment it receives in some auctions, such a strategy is fraught with considerable legal risk. When an advertiser network submits a bid to the advertising exchange, there is some small exogenous probability that the advertisement corresponding to this bid will be rejected for exogenous reasons even if this is the highest bid submitted to the advertising exchange.[6] This is because there is always some chance that the exchange may determine that an ad does not fit the formatting requirements for ads on the page. In these cases, the advertiser with the second-highest bid may actually win the auction. If the advertiser network submits a bid for this advertiser that is greater

[6] The fact that a bid can be rejected for exogenous reasons may create an incentive for the advertiser networks to submit multiple bids to the exchange. However, we show in Proposition 5 and Corollary 1 in the appendix that it would be necessary for bids to be rejected with some probability $p \geq \frac{1}{2}$ in order for networks to have an incentive to submit multiple bids as a result of this. Since these probabilities are much less than $\frac{1}{2}$ in practice, this alone will not create an incentive for the advertiser networks to submit multiple bids.

than this advertiser's actual bid, then this advertiser may be forced to pay more than his actual bid for the advertising opportunity. This would typically be a contract violation for the advertiser network that could easily expose the network to a fraud lawsuit. As such, it seems unlikely that an advertiser network would submit a second bid that is greater than the second-highest bid of one of its advertisers.

Another possible strategy for an advertiser network would be to set an internal reserve price such that the network will only submit an advertiser's bid to the advertising exchange if the advertiser's bid exceeds the reserve price. While this may be a reasonable strategy in some settings, it is worth noting that such a strategy is available to advertiser networks regardless of whether one uses the mechanism we have considered to incentivize advertiser networks to submit multiple bids. And even if the advertiser networks make use of this strategy, the networks will still have an incentive to submit multiple bids to the advertising exchange as a result of our payment scheme. None of the reasoning in the proof of Proposition 1 would be affected if some networks only submitted bids that exceeded some minimum reserve.

Finally, advertiser networks may set fees for its advertisers differently than they typically do today. Rather than solely charge some fixed percentage for helping advertisers bid on an advertising exchange, these networks may instead charge winning advertisers in part on the basis of the difference between the second-highest bid within the network and the second-highest bid that was submitted to the exchange. That is, if $b_{(2)}$ denotes the second-highest bid that was submitted to the advertiser network and $b_{(3)}$ denotes the second-highest bid that was submitted to the advertising exchange, then the advertiser network may charge an advertiser from its network who won the advertising opportunity a fee of $\alpha'(b_{(2)} - b_{(3)})$ for some $\alpha' \in [\alpha, 1]$ (if $b_{(2)} > b_{(3)}$) while still also charging some additional amount or percentage for helping advertisers bid on an advertising exchange.

If advertiser networks set fees for its advertisers in the manner described in the previous paragraph, then advertiser networks would have no incentive to submit a second bid to the exchange as a result of the mechanism we have proposed because submitting a second bid would only lower the payments it receives from $\alpha'(b_{(2)} - b_{(3)})$ to $\alpha(b_{(2)} - b_{(3)})$ while also decreasing advertiser welfare. Thus the scheme we have proposed would be ineffective at inducing advertiser networks who set fees this way to submit multiple bids to the exchange.

But while any such advertiser network will not be induced to submit multiple bids to the exchange, the mechanism we have proposed can still be effective. Advertiser networks currently do not set fees this way, and it would likely be costly for an advertiser network to renegotiate its contracts with the advertisers to a more complex contract in which fees are based in part on the difference between $b_{(2)}$ and $b_{(3)}$, especially since there is a strong desire for simple contracts.[7] It thus seems likely that

[7] The standard contracts in which advertiser networks charge advertisers a fixed percentage of their advertising spend as a fee are already suboptimal in several ways. For instance, it would be more efficient for an advertiser network to charge a fixed cost for designing ad creatives and then charge a percentage of the advertiser's spend as a fee on top of this since designing ad creatives is a one-time cost that does not vary with the amount of inventory an advertiser network serves an advertiser. Nonetheless, contracts in which advertisers are charged a fixed percentage of their spend as a fee are standard because of a desire to keep the fee structure as simple as possible.

at least some advertiser networks would continue to use their existing fee structures, and any such advertiser networks will have an incentive to submit multiple bids to the exchange.

Furthermore, even if some advertiser networks charge fees of $\alpha'(b_{(2)} - b_{(3)})$ (if $b_{(2)} > b_{(3)}$) to their advertisers, and thus do not submit multiple bids to the exchange, the ad exchange will not be hurt by these advertiser networks, as these networks will simply be following the same strategy of submitting one bid to the exchange that they follow today. And even if only some advertiser networks submit multiple bids to the exchange, the ad exchange may still benefit from employing this mechanism, as it is not necessary for all advertiser networks to submit multiple bids in order to obtain benefits from this mechanism. Thus the system of side payments that we have proposed continues to be viable even if some networks charge fees of $\alpha'(b_{(2)} - b_{(3)})$ (if $b_{(2)} > b_{(3)}$) to their advertisers.

6 Conclusion

Advertiser networks typically only submit the highest bid they have received to an advertising exchange, and this hurts publisher revenue by potentially reducing the second-highest bid and thus the seller's revenue from the auction. This paper has investigated whether it is possible to incentivize these advertiser networks to also submit the second-highest bid they received to the advertising exchange by using a scheme that would share the resulting revenue increase from the auction with the advertiser networks.

Our results have indicated that it is indeed possible to induce these advertiser networks to submit a second bid to the advertising exchange by using such a scheme, and this will increase publisher revenue as long as advertisers do not modify their bids. Making use of such a scheme creates incentives for advertisers to reduce their bids, and if the advertiser networks are perfectly symmetric, publisher revenue will decline as a result. However, when the networks are asymmetric, this change in advertiser bidding strategies will necessarily result in more efficient allocations, and it may then be feasible to increase publisher revenue by incentivizing the advertiser networks to submit a second bid to the advertising exchange. Thus while the scheme considered in this paper is an effective scheme for incentivizing advertiser networks to submit multiple bids, whether it enhances publisher revenue in the long run will hinge on the extent to which the networks differ from one another.

7 Appendix

Lemma 1 *For any fixed strategies of the other bidders, an individual advertiser can always achieve at least as high an expected payoff by making a bid $b = v$ rather than a bid $b < v$.*

Proof Suppose an advertiser with value v changes his bid b from some $b = b' < v$ to $b = v$. Note that this change in an advertiser's bidding strategy only affects the advertiser's payoff if the advertiser would have the highest bid in the auction when

$b = v$ but not when $b = b' < v$. In this case, the advertiser's payoff changes from zero to $v - p$ for some $p \leq v$ that reflects the amount the advertiser has to pay for the advertising opportunity. Since the advertiser's payoff increases from such a change, an advertiser can always achieve at least as high a payoff from making a bid $b = v$ than by making a bid $b < v$. □

Proof of Proposition 3 First we derive necessary and sufficient conditions for there to exist an equilibrium in which all advertisers follow the same strictly monotonic bidding strategy $b(v)$. Suppose there are m advertisers in each advertiser network. Note that if all other advertisers follow the same strictly monotonic bidding strategy $b(v)$ and an advertiser with value v makes a bid of $b(\hat{v})$, then this advertiser obtains an expected payoff of $F(\hat{v})^{m-1} \int_0^{\hat{v}} (v - b(y))(nm - m) f(y) F(y)^{nm-m-1}\, dy$. In order for it to be an equilibrium for all advertisers to follow the bidding strategy $b(v)$, this expected payoff must be maximized when $\hat{v} = v$.

We first show that a sufficient condition for this expected payoff to be maximized when $\hat{v} = v$ is that the derivative of this expected payoff with respect to $\hat{v}$ is zero when $\hat{v} = v$. To see this, note that we can rewrite an advertiser's expected payoff from making a bid of $b(\hat{v})$ when the advertiser's value is v as $U(v, \hat{v}) = vq(\hat{v}) - p(\hat{v})$, where $q(\hat{v})$ represents the probability that an advertiser who makes a bid of $b(\hat{v})$ wins the object and $p(\hat{v})$ represents the expected payment for an advertiser who makes a bid of $b(\hat{v})$. $q(\hat{v})$ and $p(\hat{v})$ are both continuously differentiable and strictly increasing in $\hat{v}$. Thus if the derivative of the advertiser's expected payoff with respect to $\hat{v}$ is zero when $\hat{v} = v$, then $vq'(v) - p'(v) = 0$. Now $\frac{\partial U(v,\hat{v})}{\partial \hat{v}} = vq'(\hat{v}) - p'(\hat{v})$ so if $vq'(v) - p'(v) = 0$ for all v, then $\frac{\partial U(v,\hat{v})}{\partial \hat{v}} > 0$ if $v > \hat{v}$ and $\frac{\partial U(v,\hat{v})}{\partial \hat{v}} < 0$ if $v < \hat{v}$. From this it follows that $U(v, \hat{v})$ is increasing in $\hat{v}$ for $\hat{v} < v$ and decreasing in $\hat{v}$ for $\hat{v} > v$. Thus if the derivative of this expected payoff with respect to $\hat{v}$ is zero when $\hat{v} = v$, then this expected payoff is maximized when $\hat{v} = v$.

Since the first-order conditions are sufficient for a global maximum, it is an equilibrium for all advertisers to follow the bidding strategy $b(v)$ if and only if the derivative of $U(v, \hat{v}) = F(\hat{v})^{m-1} \int_0^{\hat{v}} (v - b(y))(nm - m) f(y) F(y)^{nm-m-1}\, dy$ with respect to $\hat{v}$ is zero when $\hat{v} = v$ or if and only if

$$
\begin{aligned}
&(m-1) f(v) F(v)^{m-2} \int_0^v (v - b(y)) f(y) F(y)^{nm-m-1}\, dy \\
&\quad + (v - b(v)) f(v) F(v)^{nm-2} = 0 \\
&\quad \Leftrightarrow (m-1) \int_0^v (v - b(y)) f(y) F(y)^{nm-m-1}\, dy + (v - b(v)) F(v)^{nm-m} = 0
\end{aligned}
$$

By differentiating this expression with respect to v, we further see that it is an equilibrium for all advertisers to follow the bidding strategy $b(v)$ if and only if

$$
\begin{aligned}
&(m-1)(v - b(v)) f(v) F(v)^{nm-m-1} + (m-1) \int_0^v f(y) F(y)^{nm-m-1}\, dy \\
&\quad + (1 - b'(v)) F(v)^{nm-m} + (nm - m)(v - b(v)) f(v) F(v)^{nm-m-1} = 0 \\
&\quad \Leftrightarrow (nm-1)(v - b(v)) f(v) F(v)^{nm-m-1} + \frac{m-1}{nm-m} F(v)^{nm-m}
\end{aligned}
$$

$$+ (1 - b'(v))F(v)^{nm-m} = 0$$
$$\Leftrightarrow b'(v) = (nm-1)(v - b(v))\frac{f(v)}{F(v)} + \frac{nm-1}{nm-m}.$$

Thus there exists an equilibrium in which all bidders follow the same strictly monotonic bidding strategy $b(v)$ if and only if there exists a solution to the differential equation given above in which $b(0) = 0$ and $b(v)$ is strictly increasing in v. Standard results on differential equations guarantee that there exists a solution to this differential equation in which $b(0) = 0$, so it only remains to prove that any such function $b(v)$ that solves this differential equation is strictly increasing in v.

To see this, first note that $b(v)$ is increasing in v for values of v close to zero since $b'(0) > 0$. Thus it suffices to show that there cannot exist some $\hat{v} > 0$ for which $b(v)$ is increasing in v for values of $v < \hat{v}$ but non-increasing in v in some interval $[\hat{v}, \tilde{v}]$ with $\tilde{v} > \hat{v}$. Suppose by means of contradiction that such a $\hat{v}$ exists. This implies that in the limit as v approaches $\hat{v}$ from below, $b'(v)$ becomes close to zero, which in turn implies that $(nm-1)(v-b(v))\frac{f(v)}{F(v)}$ will be becoming less negative as v increases for values of v just below $\hat{v}$ and thus $(nm-1)(v-b(v))\frac{f(v)}{F(v)} + \frac{nm-1}{nm-m}$ will be increasing in v for values of v just below $\hat{v}$. This contradicts the fact that $b'(v) > 0$ for values of v just below $\hat{v}$ but $b'(v) \leq 0$ for values of v just above $\hat{v}$, and proves that $b(v)$ must be strictly increasing in v. Thus there is an equilibrium in which each advertiser with value v follows a strategy of bidding some amount $b(v)$, where $b(v)$ is a strictly increasing function of v. □

Proof of Proposition 4 First we derive necessary and sufficient conditions for there to exist an equilibrium in which all advertisers in the first advertiser network follow the same strictly monotonic bidding strategy $b(v)$. Let $f(\cdot)$ denote the probability density function corresponding to $F(\cdot)$. Note that if all the advertisers in the first advertiser network follow the same strictly monotonic bidding strategy $b(v)$ and an advertiser with value v makes a bid of $b(\hat{v})$, then this advertiser obtains an expected payoff of $F(\hat{v})^{m-1}\int_0^{b(\hat{v})}(v-y)f(y)\,dy$. In order for it to be an equilibrium for these advertisers to follow the bidding strategy $b(v)$, this expected payoff must be maximized when $\hat{v} = v$, meaning it is an equilibrium for these advertisers to follow the bidding strategy $b(v)$ if and only if

$$(m-1)f(v)F(v)^{m-2}\int_0^{b(v)}(v-y)f(y)\,dy + F(v)^{m-1}b'(v)(v-b(v))f(b(v)) = 0$$
$$\Leftrightarrow (m-1)f(v)\int_0^{b(v)}(v-y)f(y)\,dy + F(v)b'(v)(v-b(v))f(b(v)) = 0. \quad (1)$$

For the exponential distribution, Eq. (1) holds if and only if

$$\Leftrightarrow (m-1)e^{-v}\int_0^{b(v)}(v-y)e^{-y}\,dy + (1-e^{-v})b'(v)(v-b(v))e^{-b(v)} = 0$$
$$\Leftrightarrow (m-1)e^{-v}[v(1-e^{-b(v)}) - 1 + (1+b(v))e^{-b(v)}]$$
$$+ (1-e^{-v})b'(v)(v-b(v))e^{-b(v)} = 0, \quad (2)$$

which reduces to the condition given in the statement of the proposition for $v > 0$.

For the generalized Pareto distribution in which $F(v) = 1 - \frac{1}{(1+v)^2}$ and $f(v) = \frac{2}{(1+v)^3}$, Eq. (1) holds if and only if

$$\begin{aligned}
&\frac{2(m-1)}{(1+v)^3}\int_0^{b(v)} \frac{2(v-y)}{(1+y)^3}\,dy + \left(1 - \frac{1}{(1+v)^2}\right) b'(v)(v-b(v))\frac{2}{(1+b(v))^3} = 0 \\
&\Leftrightarrow (m-1)(1+b(v))^3 \int_0^{b(v)} \frac{2(v-y)}{(1+y)^3}\,dy + (1+v)(2v+v^2)b'(v)(v-b(v)) = 0 \\
&\Leftrightarrow (m-1)(1+b(v))^3 \left[\frac{(v+1)(2b(v)+b(v)^2)}{(1+b(v))^2} - \frac{2b(v)}{1+b(v)}\right] \\
&\quad + v(1+v)(2+v)b'(v)(v-b(v)) = 0 \\
&\Leftrightarrow (m-1)b(v)(1+b(v))[(v+1)(2+b(v)) - 2(1+b(v))] \\
&\quad + v(1+v)(2+v)b'(v)(v-b(v)) = 0 \\
&\Leftrightarrow (m-1)b(v)(1+b(v))[v(2+b(v)) - b(v)] \\
&\quad + v(1+v)(2+v)b'(v)(v-b(v)) = 0, \qquad (3)
\end{aligned}$$

which also reduces to the condition given in the statement of the proposition for $v > 0$.

To prove the result, it only remains to show that $b'(0) = \frac{2m}{m+1}$. To see this, first note that $b(v) = \Theta(v)$ when v is arbitrarily close to 0. We already know from Lemma 1 that $b(v) \geq v$ for all v, so it suffices to show that there cannot be some sequence of values of v with limit 0 for which $b(v) = \omega(v)$ in this sequence. To see this, note that since $e^{-b(v)} = 1 - b(v) + \frac{b(v)^2}{2} + o(b(v)^2)$, the left-hand side of Eq. (2) reduces to

$$\begin{aligned}
&(m-1)e^{-v}\left[vb(v) - 1 + (1+b(v))\left(1 - b(v) + \frac{b(v)^2}{2}\right)\right] \\
&\quad + (1-e^{-v})b'(v)(v-b(v))e^{-b(v)} + o(b(v)^2) \\
&\quad = (m-1)\left[vb(v) - \frac{b(v)^2}{2}\right] + (1-e^{-v})b'(v)(v-b(v))e^{-b(v)} + o(b(v)^2)
\end{aligned}$$

But if $b(v) = \omega(v)$ in some sequence, then $(m-1)\left[vb(v) - \frac{b(v)^2}{2}\right] < 0$ for values of v close to 0 in this sequence. And we know from Lemma 1 that $b(v) \geq v$, so it follows that $(1-e^{-v})b'(v)(v-b(v))e^{-b(v)} \leq 0$ for all v. Thus if $b(v) = \omega(v)$ along some sequence, then the left-hand side of Eq. (2) is negative for sufficiently small values of v in this sequence. Thus $b(v) = \Theta(v)$ must hold for the exponential distribution.

Similarly, if there is some sequence of values of v with limit 0 for which $b(v) = \omega(v)$ in this sequence, then the left-hand side of Eq. (3) reduces to $-(m-1)b(v)^2 + v(1+v)(2+v)b'(v)(v-b(v)) + o(b(v)^2)$, which is negative for sufficiently small values of v in this sequence. Thus $b(v) = \Theta(v)$ must hold for the generalized Pareto distribution.

Since $b(v) = \Theta(v)$, in the limit as $v \to 0$, we have $e^{-b(v)} = 1 - b(v) + \frac{b(v)^2}{2} + o(b(v)^2) = 1 - b(v) + \frac{b(v)^2}{2} + o(v^2)$. Thus the left-hand side of equation (2) reduces to

$$(m-1)e^{-v}\left[vb(v)-1+(1+b(v))\left(1-b(v)+\frac{b(v)^2}{2}\right)\right]$$
$$+vb'(v)(v-b(v))e^{-b(v)}+o(v^2)$$
$$=(m-1)\left[vb(v)-\frac{b(v)^2}{2}\right]+vb'(v)(v-b(v))+o(v^2) \quad (4)$$

Now if $b'(0)=k$, then the expression in Eq. (4) will equal zero for small v if and only if $(m-1)(k-\frac{k^2}{2})+k(1-k)=0 \Leftrightarrow (m-1)(1-\frac{k}{2})+1-k=0 \Leftrightarrow k=\frac{2m}{m+1}$. Thus $b'(0)=\frac{2m}{m+1}$ for the exponential distribution.

Similarly, if $b'(0)=k$, the left-hand side of Eq. (3) reduces to

$$(m-1)kv(1+kv)[v(2+kv)-kv]+v(1+v)(2+v)k(v-kv)+o(v^2)$$
$$=(m-1)kv^2(2-k)+v^2 2k(1-k)+o(v^2), \quad (5)$$

so the expression in Eq. (5) will equal zero for small v if and only if $(m-1)k(2-k)+2k(1-k)=0 \Leftrightarrow (m-1)(2-k)+2(1-k)=0 \Leftrightarrow k=\frac{2m}{m+1}$. Thus $b'(0)=\frac{2m}{m+1}$ for the Pareto distribution as well. □

Proposition 5 *Suppose any bid submitted by an advertiser network is rejected with probability* p *and no side payments are made to the advertiser networks. Then if* $p \in [\frac{k-1}{k}, \frac{k}{k+1})$ *for some positive integer* k*, each advertiser network* i *has a dominant strategy of submitting exactly* $\min\{k, a_i\}$ *bids to the advertising exchange.*

Proof To prove this result it suffices to prove that an advertiser network would prefer to submit the top $k+1$ bids the advertiser network has received rather than only submitting the top k bids if and only if $p \geq \frac{k}{k+1}$. Consider the incentives faced by a generic advertiser network i, let b_j denote the highest bid submitted by some other advertiser network $j \neq i$ that is not rejected, and let $b_{(k)}$ denote the kth-highest bid that advertiser network i has received.

Note that whether advertiser network i submits the $k+1$th-highest bid that this network has received can only affect this network's payoff if $b_{(k+1)} > b_j$ because if this inequality does not hold, then the advertiser with bid $b_{(k+1)}$ will never win the auction and will also never affect the price paid by any of the advertisers in network i if one of these advertisers wins the auction. Also note that whether advertiser network i submits the $k+1$th-highest bid that this network has received can only affect this advertiser network's payoff if this bid is not rejected since if this bid is rejected, the advertiser network is in the same situation that would arise if this network never submitted this bid in the first place. Finally, note that whether advertiser network i submits the $k+1$th-highest bid that this network has received can only affect this advertiser network's payoff if no more than one of the top k bids that this advertiser network submitted are accepted because if at least two of the other k bids that this advertiser network submitted are accepted, then submitting the $k+1$th-highest bid will not affect either who wins the auction or the price paid by this advertiser.

Thus in analyzing whether advertiser network i prefers to submit the $k+1$th-highest bid that this network has received, we can condition on the possibility that $b_{(k+1)} > b_j$, the $k + 1$th-highest bid will not be rejected if it is submitted, and no more than one of the top k bids that this advertiser network submits will be accepted. Note that the probability the rth-highest bid that this advertiser network submits will be accepted and all of the other top k bids that this network submits will be rejected is $(1 - p)p^{k-1}$ for all $r \leq k$. And the probability that all of the top k bids that this network submits will be rejected is p^k.

If the rth-highest bid that this advertiser network submits is accepted and all of the other top k bids that this network submits are rejected, then the advertiser network's payoff is $\beta(b_{(r)} - b_{(k+1)})$ if this network submits the $k + 1$th-highest bid to the advertising exchange and $\beta(b_{(r)} - b_j)$ otherwise, meaning the advertiser network's payoff decreases by $\beta(b_{(k+1)} - b_j)$ as a result of submitting the $k + 1$th-highest bid to the advertising exchange. And if all of the top k bids that the network submits are rejected, then the advertiser network's payoff is $\beta(b_{(k+1)} - b_j)$ if this network submits the $k + 1$th-highest bid to the advertising exchange and 0 otherwise, meaning the advertiser network's payoff increases by $\beta(b_{(k+1)} - b_j)$ as a result of submitting the $k + 1$th-highest bid to the advertising exchange.

By combining the results in the previous two paragraphs, it follows that if the advertiser network submits the $k + 1$th-highest bid that this network has received to the advertising exchange, then conditional on $b_{(k+1)} > b_j$ and the $k + 1$th-highest bid not being rejected, this advertiser network's payoff changes by

$$p^k\beta(b_{(k+1)} - b_j) - \sum_{r=1}^{k}(1 - p)p^{k-1}\beta(b_{(k+1)} - b_j)$$
$$= p^{k-1}\beta(b_{(k+1)} - b_j)(p - k(1 - p)),$$

which is non-negative if and only if $p - k(1 - p) \geq 0$ or $p \geq \frac{k}{k+1}$. Thus advertiser networks prefer to submit the top $k + 1$ bids they have received to the advertising exchange rather than submitting only the top k bids if and only if $p \geq \frac{k}{k+1}$. The result then follows. □

Corollary 1 *Suppose any bid submitted by an advertiser network is rejected with probability $p < \frac{1}{2}$ and no side payments are made to the advertiser networks. Then each advertiser network has a dominant strategy of submitting exactly one bid to the advertising exchange.*

References

Angel S, Walfish M (2013) Verifiable auctions for online ad exchanges. In: Proceedings of the ACM Special Interest Group Data Communication Conference (SIGCOMM), pp 195-206

Balsiero SR, Besbes O, Weintraub GY (2015) Repeated auctions with budgets in ad exchanges: approximations and design. Manage Sci 61(4):864–884

Balsiero SR, Candogan O (2015) Optimal contracts for intermediaries in online advertising. Duke University, Durham

Balsiero SR, Feldman J, Mirrokni V, Muthukrishnan S (2014) Yield optimization of display advertising with ad exchange. Manage Sci 86(12):2886–2907

Cavallo R, McAfee RP, Vassilvitskii S (2015) Display advertising auctions with arbitrage. ACM Trans Econ Comput 3(3):1–23, Article No. 15

Chakraborty T, Even-Dar E, Guha S, Manour Y, Muthukrishnan S (2010) Selective call out and real time bidding. Proc Worksh Internet Netw Econ (WINE) 6:145–157

Feldman J, Mirrokni V, Muthukrishnan S, Pai MM (2010) Auctions with intermediaries. Proc ACM Conf Electron Com (EC) 11:23–32

Gavious A, Minchuk Y (2014) Ranking asymmetric auctions. Int J Game Theory 43(2):369–393

Krishna V (2010) Auction theory. Elsevier, Burlington

Lang K, Delgado J, Jiang D, Ghosh B, Das S, Gajewar A, Jagadish S, Seshan A, Botev C, Binderberger-Ortega M, Nagaraj S, Stata R (2011) Efficient online ad serving in a display advertising exchange. Proc ACM Int Conf Web Search Data Mining (WSDM) 4:307–316

Mansour Y, Muthukrishnan S, Nisan N (2012) Doubleclick ad exchange auction. Tel Aviv University, Tel Aviv-Yafo

Mares V, Swinkels JM (2014) Comparing first and second price auctions with asymmetric bidders. Int J Game Theory 43(3):487–514

McAfee RP, Vassilvitskii S (2012) An overview of practical exchange design. Curr Sci 103(9):1056–1063

Muthukrishnan S (2009a) Ad exchanges: research issues. Proc Int Worksh Internet Netw Econ (WINE) 5:1–12

Muthukrishnan S (2009b) AdX: a model for ad exchanges. ACM SIGecom Exch 8(2):1–2

Myerson RB (1981) Optimal auction design. Math Op Res 6(1):58–73

Pai MM (2010) Competition in mechanisms. ACM SIGecom Exch 9(1):1–5

Stavrogiannis LC, Gerding EH, Polukarov M (2013) Competing intermediary auctions. Proc Int Conf Auton Agents Multiagent Syst (AAMAS) 12:667–674

Stavrogiannis LC, Gerding EH, Polukarov M (2014) Auction mechanisms for demand-side intermediaries in online advertising exchanges. Proc Int Conf Auton Agents Multiagent Syst (AAMAS) 13:1037–1044