

Buy-it-now or Take-a-chance: A Simple Sequential Screening Mechanism

L. Elisa Celis
Computer Science
University of Washington
ecelis@cs.washington.edu

Markus M. Mobius
Department of Economics
Harvard University
mobius@fas.harvard.edu

Gregory Lewis
Department of Economics
Harvard University
glewis@fas.harvard.edu

Hamid Nazerzadeh
Microsoft Research
New England Lab
hamidnz@microsoft.com

ABSTRACT

We present a simple auction mechanism which extends the second-price auction with reserve and is truthful in expectation. This mechanism is particularly effective in private value environments where the distribution of valuations are irregular. Bidders can “buy-it-now”, or alternatively “take-a-chance” where the top d bidders are equally likely to win. The randomized take-a-chance allocation incentivizes high valuation bidders to buy-it-now. We show that for a large class of valuations, this mechanism achieves similar allocations and revenues as Myerson’s optimal mechanism, and outperforms the second-price auction with reserve.

In addition, we present an evaluation of bid data from Microsoft’s AdECN platform. We find the valuations are irregular, and counterfactual experiments suggest our BIN-TAC mechanism would improve revenue by 11% relative to an optimal second-price mechanism with reserve.

Categories and Subject Descriptors

K.4.4 [Computing Milieux]: Electronic Commerce

General Terms

Algorithms, Economics, Theory

Keywords

Mechanism Design, Sequential Screening, Online Advertising, Ad Auctions, AdECN

1. INTRODUCTION

Many Internet companies generate revenue by selling the advertisement space on their webpages. Improved targeting technologies allow e-commerce firms to match advertisers and consumers with ever greater efficiency. While these technologies generate a lot of surplus for advertisers, they also tend to create thin markets with skewed value distributions. These environments pose special challenges for the predominant auction mechanisms that are used to sell online ads because they reduce competition among bidders, making

it difficult for the platform to extract the surplus generated by targeting; see [2, 9].

For example, a sportswear firm advertising on the New York Times website may be willing to pay much more for an advertisement placed next to a sports article than one next to a movie review. It might pay an additional premium for a local consumer who lives in New York City and an even higher premium if the consumer is known to browse websites selling sportswear. Each layer of targeting increases the sportswear firm’s valuation for the consumer but also dramatically narrows down the set of competitors to fellow sportswear firms in New York City. Without competition, revenue performance may be poor [9].

To get some intuition, consider a simple model: When advertisers “match” with users, they have high valuation; otherwise they have low valuation. Assume that match probabilities are independent across bidders, and sufficiently low that the probability that *any* bidder matches is relatively small. Then a second-price auction will typically get low revenue, since the probability of two “matches” occurring in the same auction is small. On the other hand, setting a high fixed price is not effective since the probability of zero “matches” occurring is relatively large and many impressions would go unallocated. Hence, allowing targeting creates asymmetries in valuations that can increase efficiency, but decrease revenue. In fact, because of this phenomenon, some have suggested that it is better to create thicker markets by bundling different impressions together [7, 6, 10].

Bundling may improve revenues, but reduces efficiency since the average quality of user-advertiser matches is degraded. In principle, one would like to allow targeting but still extract significant revenues. This paper outlines a new and simple mechanism which addresses this issue. We call it *buy-it-now or take-a-chance (BIN-TAC)*, and it works as follows. Goods are auctioned with a buy-it-now price p , set relatively high. If a single bidder is willing to pay the price, they get the good for price p . If more than one bidder takes the buy-it-now option, a second price auction is held between those bidders with reserve p . Finally, if no-one participates in buy-it-now, an auction is held in which the top d bidders are eligible to receive the good, and it is randomly awarded to one of them at the $(d + 1)$ -st price. Note that in the case where $d = 1$ this reduces to the second-price auction.

In this manner, we combine the advantages of an auction and a fixed price mechanism. When matches occur,

advertisers pay for the fixed-price buy-it-now option, allowing for revenue extraction. This is incentive compatible (i.e. truthful) because in the event that they “take-a-chance” on winning via auction, there is a significant probability that they will not win the impression. On the other hand, when no matches occur, the auction mechanism ensures the impression is still allocated.

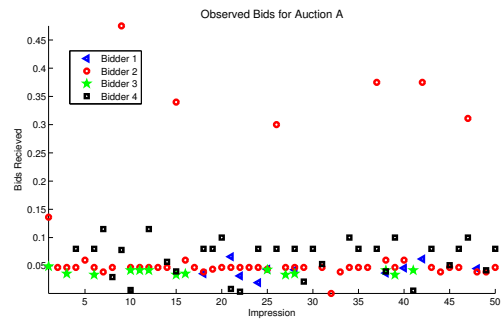
The BIN-TAC mechanism is simple, and requires relatively little input from the mechanism designer: a choice of buy-it-now price, take-a-chance parameter d and optionally a reserve in the take-a-chance auction. This makes it flexible across a wide range of environments. The tradeoff is that it is not the optimal mechanism analyzed by [13]. As it turns out, the downside is small. We show that when the valuations are drawn iid from a mixture of two regular distributions — a weighted combination of high and low valuation distributions with disjoint supports — our mechanism is “nearly optimal” in the sense that it has very similar allocation rules and transfer payments as the optimal mechanism. In this setting, the second price auction with reserve is rarely optimal, and is dramatically outperformed by the BIN-TAC mechanism. We also run simulations to show that BIN-TAC also outperforms the SPA by a large margin in environments where the supports overlap.

In the last part of the paper, we demonstrate our mechanism’s effectiveness using data from the AdECN platform. Since the current auction format is a second-price auction, which is incentive compatible, we can interpret bids as valuations, and therefore simulate how these bidders would counterfactually behave under a BIN-TAC format. We find that our mechanism generates 11% more revenue than the optimal second-price auction.


Related Work

Myerson [13] proposed a general approach to design optimal mechanisms when the private information of the agents is single-dimensional. However, if the distributions are not “well-behaved”, then characterizing the optimal mechanism can be challenging. The approach we take in this work is to look instead for a simple and “nearly optimal” mechanism. Hartline and Roughgarden [8] discuss the benefits of simple mechanisms, and show a variety of examples where they can approximate the optimal expected revenue.

Sequential screening models have been proposed for revenue maximization in dynamic environments. For instance, Courty and Li [5] consider a setting where the buyers themselves learn their type dynamically (first whether they are high or low, then their specific valuation). In this case, offering contracts after the first type revelation but before the second may be optimal; see [3] for a survey on dynamic mechanisms. In the static setting, sequential screening and posted-prices can be used to design optimal (or near-optimal) mechanisms when the bidders have multi-dimensional private information, e.g., see [16, 4]. Our model deals with the case where types are single-dimensional, have a mixture and buyers know their valuation from the outset. Additionally, our model considers only the private value setting. Abraham et al. [1], consider an adverse selection problem that arises in a common value setting when some bidders are privately informed; this is motivated by the display advertising and advertisement exchange markets when some advertiser are better able to utilize information obtained from cookies. They show that asymmetry of information



(a) Bids for 50 random impressions (ordered chronologically) of ad A.



(b) Bids for 50 random impressions (ordered chronologically) of ad B.

Figure 1: In these two examples we see that a bidder’s valuation fluctuates dramatically in a way that is uncorrelated with other bidders and uncorrelated over time.

can sometimes lead to low revenue in this market. For further discussion on advertisement exchange markets see [12].

Organization

The paper proceeds in four parts. First, we describe the AdECN market, providing some interesting and (to our knowledge) novel observations about this market. In the second part we define the mechanism and an stylized environment inspired by the AdECN market, proving existence and characterization results, and solving for the revenue-maximizing parameter choices analytically. The third section consists of simulation results, comparing the performance of the BIN-TAC mechanism to the SPA and to the benchmark of full-surplus extraction, as the shape of the distributions, the probability of high valuation and the number of bidders vary. Finally, in the fourth part we estimate valuations and conduct counterfactual experiments using the AdECN data. All proofs are contained in the appendix.

2. THE ADECN MARKET

In this paper we focus on situations where bidder valuations fluctuate considerably. We first show evidence from a real-world market which drives this interest. Specifically, we examine data from AdECN, Microsoft’s real-time auction-based neutral exchange for online display advertising. On AdECN, advertisers, or firms acting on their behalf, may bid for display ads on various publishers. An *impression* is a single advertisement slot on a given webpage to a given user. An auction is held every time an individual browses




Figure 2: Kernel density estimate of the pdf of the (normalized) gap between the highest and second highest bids in auctions for ads A and B.

a webpage on one of the publishers. Consequently, a huge number of auctions are held each day. We examined the bids for a subset of products over a 24-hour period — a data set of over 2 million auctions (see Table 1 for an overview).

Impressions are grouped together into *products*, usually consisting of an advertising slot on a particular publisher (e.g., banner ad on the main New York Times sports page). This reduces the complexity of the market, by allowing bidders to express their bids in terms of products, rather than individual impressions. Yet AdECN provides bidders with some information about web page content, as well as demographic and historical information about the users, so that bidders can vary their bids with these characteristics in order to optimize their advertising to target audiences. The auction mechanism is a second-price auction with reserve. Since it is weakly dominant to bid one’s valuation in a SPA, we interpret bids as valuations.

Figure 1 shows 50 randomly selected impressions on two products. Looking at the figures, we see that there are relatively few bidders in the market, 4 on product A and 5 on product B, so the market is relatively thin. The highest bid varies markedly across auctions, consistent with bidders varying their bidding strategy based on observable information about the viewer. Most winning bids are quite low, but occasionally winning bids are much higher. Moreover, conditional on a high bid from one bidder, the other bids do not appear to be higher, which suggests that idiosyncratic advertiser-impression *matches* drive the high bids, rather than a commonly valued component. Additionally, the value of an impression does not vary depending on the time of day, suggesting the matches are driven by the user’s properties, not the page or advertisement content.

Given these observations, one might expect the gap between the winning bid and the price — the second highest bid — to be quite substantial. This is clear from Figure 2, a kernel density estimate of this gap. In Figure 3 we plot the *virtual valuations* $v - \frac{1-F(v)}{f(v)}$.

For both example products, the virtual valuations are non-increasing, which implies that the SPA with reserve is not the optimal mechanism. On the other hand, the repeated fluctuation in virtual valuation implies the optimal mechanism is quite complex, requiring “ironing” over several regions. This motivated our search for a middle ground: a mechanism that retains the simplicity of the SPA while getting nearly optimal revenue performance.




Figure 3: The virtual valuation (computed using a kernel density estimate of the pdfs).

3. BUY-IT-NOW OR TAKE-A-CHANCE

We start our analysis by formally defining the BIN-TAC mechanism. A *buy-it-now price* p is posted. Buyers simultaneously indicate whether they wish to *buy-it-now* (BIN). In the event that exactly one bidder elects to buy-it-now, that bidder wins the auction and pays p . If two or more bidders elect to BIN, a second-price sealed bid auction with reserve p is held between those bidders. Bidders who chose to BIN are obliged to participate in this auction. Finally, if no-one elects to BIN, a sealed bid *take-a-chance* (TAC) auction is held between all bidders, with a *reserve* r . In that auction, one of the top d bidders is chosen uniformly at random, and if that bidder’s bid exceeds the reserve, they win the auction and pay the maximum of the reserve and the $(d+1)$ -th bid.¹ We call d the *TAC-parameter*.

3.1 A Stylized Model

Motivated by the observations in Section 2, we define a stylized environment in which bidders generally have low valuations, but occasionally one or more bidders have much higher valuations. Assume n bidders participate in an auction for a single good which is valued at zero by the seller. Buyers are risk neutral, and draw their values V_i for the good independently and privately from some distribution F . This F is a mixture of F_L and F_H , and a valuation V takes the form $V = (1 - X)V_L + XV_H$ where X a Bernoulli random variable with parameter $0 \leq \alpha \leq 1$, $V_L \sim F_L$ and $V_H \sim F_H$. We assume F_L has support $[\underline{\omega}_L, \bar{\omega}_L]$ and F_H has support $[\underline{\omega}_H, \bar{\omega}_H]$, and that these supports are disjoint (so $\bar{\omega}_L < \underline{\omega}_H$). This formalizes the idea that there are two separate types: low valuation types (draws from F_L) and high valuation types (draws from F_H), although there is heterogeneity within these groups.

An important feature of this environment is that optimal mechanism design is not straightforward. Define the virtual valuations $\psi(v) \equiv v - \frac{1-F(v)}{f(v)}$. When $\psi(v)$ is strictly increasing, the optimal mechanism is a second-price auction with a reserve price [13]. Here, however, the virtual valuations are (infinitely) negative over the region $(\bar{\omega}_L, \underline{\omega}_H)$ since F is unsupported on this region. In this case the *ironing* of virtual values is required, and the optimal mechanism is relatively complicated and hard to compute. What we will later argue is that the BIN-TAC mechanism is much simpler and “nearly

¹Ties occur when multiple bidders bid the d -th highest bid: in that case, the price is the d -th highest bid, and all bidders who bid that amount jointly split a $1/d$ probability of allocation.

optimal” (see Section 3.3). First, however, we characterize equilibrium behavior.

3.2 Equilibrium Analysis

This is a sequential mechanism which we analyze by backward induction. The auctions that follow the initial BIN decision admit simple strategies. If multiple players choose to BIN, the allocation mechanism reduces to a second-price auction with reserve p . Thus, it is weakly dominant for players to bid their valuations. Since participation is obligatory at this stage, the minimum allowable bid is p . However, it is easy to show that an individually rational player will not choose to BIN unless her valuation is at least p , so this does not present a problem.

Likewise, in the TAC auction it is weakly dominant for the bidders to bid their valuations. The logic is standard: if a bidder with valuation v bids $b' > v$, it can only change the allocation when the maximum of the d -th highest rival bid and the reserve price is in $[v, b']$. But whenever this occurs, the resulting price of the object is above the bidder's valuation and if he wins he will regret his decision. Alternatively, if they bid $b' < v$, the price is not affected, and their probability of winning will decrease.

Taking these strategies as given, we now turn to the buy-it-now decision. Intuitively, one expects the BIN option to be more attractive to higher types: they have the most to lose from either random allocation (they may not get the good even if they are willing to pay the most) or from rivals taking the BIN option (they certainly do not get the good). This suggests that in equilibrium, the BIN decision takes a threshold form: $\exists \bar{v}$ such that types with $v \geq \bar{v}$ elect to BIN, and the rest do not. This is in fact the case.

Prior to stating a formal theorem, we introduce the following notation. Let the random variable Y^j be the j -th highest valuation from $n - 1$ iid samples from F (the j -th highest rival bid); and let Y^* be the maximum of Y^d and the TAC reserve r .

THEOREM 1 (EQUILIBRIUM CHARACTERIZATION).

Assume $p \leq \frac{d-1}{d}\bar{\omega}_H + \frac{1}{d}E[Y^*]$. Then there exists a unique pure strategy Bayes-Nash equilibrium of the game, characterized by a unique threshold \bar{v} satisfying:

$$\bar{v} = p + \frac{1}{d}E[\bar{v} - Y^* | Y^1 < \bar{v}] \quad (1)$$

Types with $v \geq \bar{v}$ take the BIN option; and all types bid their valuation in any auction that may occur.

Equation (1) is intuitive: At what point is a bidder indifferent between the BIN and TAC options? The only time the choice is relevant is when there are no higher valuation bidders (since they would win the BIN auction). So if a bidder has the highest value and chooses to BIN, they get a surplus of $\bar{v} - p$. Choosing to TAC gives $\frac{1}{d}E[\bar{v} - Y^* | Y^1 < \bar{v}]$, since they only win with probability $\frac{1}{d}$, although their payment of Y^* is on average much lower. Equating these two yields Equation (1). The assumption that $p \leq \frac{d-1}{d}\bar{\omega}_H + \frac{1}{d}E[Y^*]$ rules out uninteresting cases where the BIN price is so high that no-one ever chooses BIN.

Now we consider the revenue-maximizing choices of the design parameters: the BIN price p , the TAC reserve r and the TAC parameter d . One way to think about the BIN price is as a reserve, where bidders who fail to meet the

reserve still have some chance of participation. Perhaps unsurprisingly, we get some familiar looking equations for the optimal reserves. Again, we must introduce some notation. For $k = 0, 1$, let $R_k^{TAC}(d, r)$ be the expected revenue from the TAC auction when all bidders participate, but only k bidders have high valuation. Then we have the following theorem.

ASSUMPTION 1. $v - \frac{1-F(v)}{f(v)}$ is continuous and increasing in $[\underline{\omega}_L, \bar{\omega}_L]$ and in $[\underline{\omega}_H, \bar{\omega}_H]$.

THEOREM 2 (OPTIMAL BIN-TAC RESERVES).

Under Assumption 1, the revenue-maximizing TAC reserve r is independent of d and p , and satisfies:

$$r^* = \frac{1 - F(r^*)}{f(r^*)} \quad (2)$$

If a solution exists with $\bar{v}(p^*, d, r) \in [\underline{\omega}_H, \bar{\omega}_H]$, then the optimal BIN price is given by:

$$p^*(d, r) = R_1^{TAC}(d, r) + \left(\frac{d-1}{d}\right) \frac{1 - F(\bar{v}(p^*, d, r))}{f(\bar{v}(p^*, d, r))}. \quad (3)$$

Otherwise, p^* solves $\bar{v}(p^*, d, r) = \underline{\omega}_H$:

$$p^*(d, r) = \frac{d-1}{d}\underline{\omega}_H + \frac{1}{d}E[Y^* | Y^1 < \underline{\omega}_H] \quad (4)$$

Equation 2 is somewhat surprising; the optimal TAC reserve is exactly the standard reserve in [13], ensuring that no types with negative virtual valuation are ever awarded the object. This is despite the fact that our BIN-TAC mechanism is not the optimal mechanism.

The key insight is that the TAC reserve is relevant for the BIN choice. Raising the TAC reserve lowers the surplus from participating in the TAC auction, and so one can also raise the BIN price while keeping the indifferent type \bar{v} constant. So the trade-off is exactly the usual one: raising the TAC reserve extracts revenue from types above r^* — even those above \bar{v} — at the cost of losing revenue from the marginal type. This is why we get the usual solution.

On the other hand, the optimal BIN price is non-standard. To get some intuition, notice that the BIN price is only relevant when there is a single bidder with valuation above \bar{v} . Then the first term in the RHS can be interpreted as an “outside option” — if the good doesn't sell by buy-it-now, the seller gets to hold a TAC auction with expected revenue $R_1^{TAC}(d, r)$. To this outside option, the seller adds a “markup” term of $\frac{1-F(\bar{v})}{f(\bar{v})}$. This markup is weighted down by $\frac{d-1}{d}$ since a small increase in \bar{v} doesn't increase revenue one-for-one.

We note that in many cases, there is no interior solution for p^* . Whenever the high valuations are substantially larger than the low valuations (i.e. $\underline{\omega}_H \gg \bar{\omega}_L$) it is not profitable to randomize the allocation for high types by setting $\bar{v}(p, d, r) \in [\underline{\omega}_H, \bar{\omega}_H]$, since the efficiency loss would be large. In this case p^* is set so that the lowest high type at $\underline{\omega}_H$ is indifferent between TAC and BIN.

3.3 Performance Comparisons

We would like to compare our mechanism to two benchmark mechanisms, the optimal mechanism, and the second price auction with reserve r^* . Yet as we noted in section 3.1, we need to use an ironing procedure to get the optimal

mechanism. This ironing procedure, outlined in the appendix, yields the *ironed virtual valuations* $\phi(v)$ defined below. We observe that if $\alpha\omega_H \geq r^*(1 - F(r^*))$, then, the optimal mechanism is a second-price auction with reserve ω_H . Now, assume $\alpha\omega_H < r^*(1 - F(r^*))$. Then, there exists v^* , $r^* \leq v^* \leq \bar{\omega}_L$, such that

$$(2 - \alpha - F(v^*))F(v^*) + \alpha(\omega_H - v^*)f(v^*) = 1 - \alpha \quad (5)$$

where r^* is defined in Eq. (2). This defines the ironed virtual valuations as follows:

$$\phi(v) = \begin{cases} 0 & v \in [\omega_L, r^*) \\ \psi(v) & v \in [r^*, v^*) \\ \psi(v^*) & v \in (v^*, \omega_H) \\ \psi(v) & v \in [\omega_H, \bar{\omega}_H], \end{cases} \quad (6)$$

Ironed Mechanism: Award the good to the bidder with the highest ironed virtual valuation, breaking ties uniformly at random, provided the virtual valuation is positive. Let w_1 be the valuation of the winning bidder; let w_2 be the valuation of the bidder with the second-highest virtual valuation (again break ties randomly). The payments are determined as follows. Let k be the number of bidders with bid in $[v^*, \omega_H)$. Then the price for the winning bidder is computed as follows:

$$w_1 \in [r^*, v^*) : \quad p = \max\{w_2, r^*\}.$$

$$w_1 \in [v^*, \omega_H) : \quad p = \begin{cases} \max\{w_2, r^*\} & \text{If } k = 1, \\ v^* & \text{otherwise.} \end{cases}$$

$$w_1 \in [\omega_H, \bar{\omega}_H) : \quad p = \begin{cases} \max\{w_2, r^*\} & \text{If } w_2 \leq v^*, \\ \frac{1}{k+1}(kw_H + v^*) & \text{If } w_2 \in [v^*, \omega_H), \\ w_2 & \text{otherwise.} \end{cases}$$

Over the “ironed” region $[v^*, \omega_H)$ allocation probabilities are constant, and consequently the expected payment of a winning bidder is also constant.


THEOREM 3 (OPTIMAL MECHANISM). *Suppose Assumption 1 holds and $\psi(\bar{\omega}_L) \leq \psi(\omega_H)$. If $\alpha\omega_H \geq r^*(1 - F(r^*))$, then the optimal mechanism is the second-price auction with reserve ω_H . If $\alpha\omega_H < r^*(1 - F(r^*))$, then the ironed-mechanism described above is optimal.*

The main challenge in proving this theorem is computing v^* . The difficulty in even this relatively simple case lends force to our claim that BIN-TAC is a useful mechanism for these kinds of environments.


Having obtained this characterization, we can compare the BIN-TAC mechanism with the optimal mechanism. Under Assumption 1, it is easy to prove that as either $n \rightarrow \infty$, $\alpha \rightarrow 1$ or $\omega_H/\omega_L \rightarrow \infty$, the BIN-TAC mechanism converges to the optimal mechanism. This, however, is not particularly interesting (a second-price mechanism will also converge to optimal). The interesting cases, both theoretically and in practice, occur for small values of the above parameters. It is here that BIN-TAC simulates OPT much better than the optimal second price auction.

For concreteness, we assume F_L is the uniform distribution over $[0, 1]$, and F_H is the uniform distribution over $[\tau, \tau + 1]$, $\tau \geq 3$. By Theorem 3, we have

$$r^* = \frac{1}{2(1 - \alpha)} \quad \text{and} \quad v^* = \left(1 - \sqrt{\frac{\alpha(\tau - 1)}{1 - \alpha}}\right).$$



(a) Allocation probabilities.



(b) Expected payments.

Figure 4: Characterization of OPT, SPA and BIN-TAC mechanisms when the distributions F_L and F_H are uniform. The x -axis corresponds to the bid.

Also, recall that the optimal second-price auction is equivalent to a BIN-TAC mechanism with $d = 1$. The table below, compares the expected revenue and welfare obtain by these mechanisms for $n = 5$ and $\tau = 3$. In addition, Figure 3.3, depicts the probability of allocation and expected payment of a bidder assuming the value of the other 4 bidders distributed according to the distribution described above.

Mechanism:	OPT	SPA	BIN-TAC, $d=2$	BIN-TAC, $d=3$
$\mathbb{E}[\text{Revenue}]$	0.89	0.76	0.85	0.83
$\mathbb{E}[\text{Welfare}]$	1.40	1.43	1.33	1.23

4. SIMULATIONS

The BIN-TAC mechanism can be applied in a much wider context than considered thus far. Specifically, we no longer make the assumption that F_L and F_H have disjoint support. Since the BIN-TAC mechanism is uniquely determined by the three parameters r^* , d^* , and p^* , it is much easier to calculate than the optimal mechanism. The reserve price r^* can be calculated according to the analysis presented in Section 3 as long as Assumption 1 holds. Unfortunately, the analysis for p^* and d^* do not follow when F_L and F_H have disjoint support. However, for any fixed d^* , numerically finding p^* is a one dimensional optimization problem. Since there are linearly many possible d^* , this makes the problem very tractable. Additionally, the performance of BIN-TAC is still good.

We first demonstrate this by simulating the mechanism on two distributions where F_L and F_H do not have disjoint support. The optimal mechanism is much more complex, we can no longer easily compare to the optimal mechanism. Instead, we do the following: Let MAX be the maximum amount of revenue extraction possible; i.e. the revenue acquired if the bidder with the highest valuation wins and pay exactly his valuation. MAX, though unattainable, dominates the optimal revenue, and gives us a useful and computable baseline. To show the effectiveness of BIN-TAC, we compare it to the optimal second price auction, and report the revenue of both as a percentage of MAX.




Figure 5: Our simulations show BIN-TAC outperforming a second price auction, often significantly.

For our simulations, we restrict ourselves to location families where the distribution $F_H(\cdot) = F_L(\cdot - \Delta)$ for some shift-parameter Δ . This Δ is the difference in mean valuation between the high and low groups. We consider two different location families; $F_L \sim \mathcal{N}$, $F_L \sim \log \mathcal{N}$, where both have mean 1 and variance 0.5. We allow Δ , n and α to vary across experiments, and compute r^* , p^* and d^* as discussed above. The results are presented in Figure 5.

The default parameters we consider are $n = 10$, $\Delta = 10$, and $\alpha = .05$, and we vary one parameter at a time. Each experiment is repeated for 1000 impressions, and we report the average. Recall that BIN-TAC generalizes the second price auction, so its performance is always at least as good, and often significantly better. Figures 5(a) and 5(b) show how as either n or α increases, we approach optimal. This is because the expected number of bidders that can target is αn . As this increases, the lower distribution becomes irrelevant, and the second price auction is once again a good approximation of optimal—i.e., there is no room for improvement. The same phenomenon can be seen for small α ; here, the high distribution becomes irrelevant and again the a second price auction approximates the optimal mechanism. However, in between the two extremes, our mechanism performs significantly better. Figure 5(c) shows the dependence on the gap Δ . As expected, the performance of BIN-TAC increases while that of a second price auction decreases as Δ gets larger. Since there is more revenue to be gained from high-valued bidders, BIN-TAC can only perform better with a large Δ . However, a second price auction would have to find a tradeoff between losing low-valued impressions and extracting revenue from high-valued impressions, hence hurting its performance.

5. ANALYSIS OF ADECN DATA

We now test our mechanism’s performance in a real-world setting. Specifically, we recover the valuations of advertisers in the AdECN market introduced in Section 2 from their observed bids, and then simulate their counterfactual bidding behavior under our BIN-TAC mechanism. This allows us to see whether our mechanism has the potential to improve platform revenues, and tests it in a less stylized environment than that of our theoretical model.

Our dataset consists of all bids submitted on all products sold by a single publisher over a 24-hour period. We restricted analysis to the subset of products that averaged at least two bidders per impression, since with zero or one bidders the BIN-TAC approach is not viable (the threat of randomization is meaningless). This left us with ten products (placements), with bidding patterns summarized in Table 1. Over 1M impressions were sold, with participation ranging

from 3-6 bidders per auction. Bids vary widely: the average bid below the 95th percentile is 0.07 while the average bid above it is 0.8, over 10 times greater. Sample skewness is consistently high, even when disaggregated by product. We note two other facts. First, the correlation of bids within an auction is consistently small, no higher than 0.09, and often negative. This suggests that bidder valuations are private, perhaps driven by idiosyncratic match quality, rather than a common component. Second, the autocorrelation within bids for a given bidder is also small, no higher than 0.02, again suggesting that there are no dynamic patterns in the evolution of bidder valuations, and the bids do not correlate with time of day.

Since the current auction format is a second price auction with reserve, we can infer the distribution of valuations directly from the bids, since they should be equal (we observe bids even when they fall below the reserve). We first normalize the bids on each product by the mean bid on that product, calculating this mean using the first 10% of our data, which was randomly selected for training purposes. Then we can estimate the density of normalized valuations.

Before running the counterfactual simulations, we must choose the optimal TAC reserve r , TAC parameter d and BIN price p . In principle, we could do this product-by-product. Instead, we use a single set of parameters for all the different products, “un-normalizing” our chosen normalized reserve r and BIN price p by multiplying by the product means to get something more individual specific. This provides a much stronger test of our approach, since we could certainly do better by conditioning our parameter choices on the individual product valuation densities. In addition, it has the advantage of being simpler, allowing a way to calculate parameters for very thin or new markets, and increasing incentive compatibility in practice.

Following our theory, we choose the reserve price r as the first time the virtual valuations are positive, as calculated from our training data. Note that this may not be optimal. We also fix $d = 2$, since the market is relatively thin. Since the data does not literally follow a mixture model, the optimal BIN-TAC price must be calculated numerically using the training data. Our counterfactual simulations — the procedure for which is outlined below — are run on the remaining 90% of the data, thus avoiding a potential overfitting problem in our parameter choices.

The simulation procedure is as follows. For some fixed parameter choices (d, r, p) , we calculate the indifferent type \bar{v}_j for each product $j = 1 \dots 10$ numerically. This requires solving for a solution to the implicit Equation 1 by iterative methods. As an input into this calculation we need the distribution of Y^* conditional on $Y^1 < \bar{v}$; we take this dis-

Products:	All	1	2	3	4	5	6	7	8	9	10
# Bids	2592025	30038	482190	406181	224417	2913	5917	711368	717	241397	486887
# Impressions	1145871	18505	223028	256084	103034	876	4048	215570	193	107791	216742
# Bidders	6	3	4	3	4	5	3	5	5	3	3
% Total Bid Value	100	1.10	11.2	15.5	4.98	0.28	0.16	51.9	0.08	3.95	10.8
Avg Bid	0.114	0.108	0.069	0.113	0.066	0.281	0.078	0.216	0.340	0.049	0.066
σ Bid	0.46	0.097	0.075	0.15	0.049	0.46	0.074	0.34	0.50	0.52	0.88
5th percentile	0.011	0.013	0.010	0.015	0.020	0.001	0.017	0.021	0.010	0.010	0.10
95th percentile	0.341	0.193	0.230	0.572	0.115	1.500	0.150	1.186	1.500	0.098	0.120
Avg Bid above 95th	0.796	0.349	0.353	0.677	0.200	1.500	0.341	1.472	1.500	0.232	0.172
Avg Bid below 95th	0.070	0.094	0.054	0.082	0.057	0.137	0.064	0.150	0.210	0.038	0.033
Sample Skew	164.8	10.4	3.59	3.10	5.72	2.08	3.48	2.75	1.72	183.5	107.2
Correlation	-0.012	-0.033	0.055	0.043	0.010	0.038	-0.044	0.084	-0.021	0.004	0.001
Autocorrelation	0.003	0.012	0.0002	-0.002	-0.004	0.013	0.020	-0.001	-0.009	-0.001	0.001

Table 1: Monetary units are cents. Statistics for the data set used in our experiment, which consisted of one publisher and all products with at least two bidders per impression on average over a 24-hr time period.

Mechanism:	AdECN	Opt SPA	BIN-TAC
Total Rev	761.8	851.8	945.6
% from BIN	0	0	53.6
% Imp Unallocated	0.001	0.014	0.017

Table 2: Counterfactual revenue results (in dollars) for the mechanisms in question. AdECN is the mechanism currently used by AdECN. Opt SPA is the second-price auction with optimal reserve ($r = 0.067$). BIN-TAC uses this same reserve, $d = 2$, and the optimal price $p = 3.8$.

tribution straight from the data. The main assumption we are making here is that bidders believe the environment to be symmetric and iid, since then our calculated \bar{v}_j correctly summarizes their incentives. This appears to be a reasonable assumption since there is little bid correlation and autocorrelation, although the symmetry assumption is probably too strong. To get the simulated BIN-TAC outcomes, we re-run the auctions in turn, assuming the highest bidder takes the BIN option if their valuation is above \bar{v}_j , and otherwise the object allocation is randomized between the top two highest bidders. We run this procedure on the training data for various p in order to determine p^* .

Once these parameters have been determined, we run the mechanism on the remaining 90% of the data to calculate counterfactual revenues. For comparison purposes, we also look at the optimal SPA, the second-price auction with reserve r^* . We find r^* numerically, and somewhat surprisingly r is very close r^* (the first time the virtual valuations are positive).² Thus, the optimal reserve price for the SPA for our data is the first time the virtual valuation is non-negative. The results are shown in Table 2. Notice that a large fraction of the revenue in the BIN-TAC mechanism is coming from the BIN prices: this right tail of valuations contributes 53.6%. This reflects the skewness in the observed valuations. The main finding is that the BIN-TAC mechanism increases revenues by 11% relative to the optimal SPA, which in turn improves on the current AdECN mechanism by 11%. This demonstrates the BIN-TAC mechanism is effective in extracting revenue, yet still allows targeting.

²See [14] for a further discussion on reserve prices.

Acknowledgments

We would like to thank Lee Zen, Susan Athey, and Mallesh Pai for fruitful discussions that improved this work. We also thank Lee Zen and members of Microsoft AdECN team, in particular Ian Ferreira, Keith Hawley, and Brandon Zirkle, for providing the data.

6. REFERENCES

- [1] Ittai Abraham, Susan Athey, Moshe Babioff, and Michael Grubb. Peaches, lemons and cookies: The impact of information structure on revenue in common value auctions. working paper, Microsoft Research, 2010.
- [2] Dirk Bergemann and Alessandro Bonatti. Targeting in advertising markets: Implications for offline vs. online media. Cowles Foundation Discussion Paper 1758R, 2010.
- [3] Dirk Bergemann and Maher Said. Dynamic auctions: A survey. *Wiley Encyclopedia of Operations Research and Management Science*, 2010.
- [4] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC)*, pages 311–320, 2010.
- [5] Pascal Courty and Hao Li. Sequential screening. *The Review of Economic Studies*, 67(4):pp. 697–717, 2000.
- [6] Eyal Even-Dar, Michael J. Kearns, and Jennifer Wortman. Sponsored search with contexts. In *Internet and Network Economics, Third International Workshop (WINE)*, pages 312–317, 2007.
- [7] Arpita Ghosh, Hamid Nazerzadeh, and Mukund Sundararajan. Computing optimal bundles for sponsored search. In *Internet and Network Economics, Third International Workshop (WINE)*, pages 576–583, 2007.
- [8] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, 2009.
- [9] Johathan Levin and Paul Milgrom. Online advertising: Heterogeneity and conflation in market design. *American Economic Review, Papers and Proceedings*, 100(2), 2010.

- [10] R. Preston McAfee, Kishore Papineni, and Sergei Vassilvitskii. Maximally representative allocations for guaranteed delivery advertising campaigns. Yahoo! Research Working Paper, 2010.
- [11] Paulo Klinger Monteiro and Benar Fux Svaiter. Optimal auction with a general distribution: Virtual valuation without densities.
- [12] S. Muthukrishnan. Ad exchanges: Research issues. manuscript, 2010.
- [13] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [14] Michael Ostrovsky and Michael Schwarz. Reserve prices in internet advertising auctions: A field experiment. Working Paper, 2009.
- [15] Mallesh Pai and Rakesh Vohra. Optimal auctions with financially constrained bidders. Working paper, 2009.
- [16] Jean-Charles Rochet and Philippe Chone. Ironing, sweeping, and multidimensional screening. *Econometrica*, 66(6):788–826, July 1998.

APPENDIX

A. PROOF OF THEOREM 1

Consider the payoffs to taking BIN, denoted by $\pi_B(v)$, and to TAC, denoted by $\pi_T(v)$, in this equilibrium. They are given by:

$$\pi_B(v) = E \left[\underbrace{1(v > Y^1 > \bar{v})(v - Y^1)}_{\text{Rival chooses BIN}} \right] + E \left[\underbrace{1(Y^1 < \bar{v})(v - p)}_{\text{No rival chooses BIN}} \right]$$

$$\pi_T(v) = \mathbb{E} \left[1(Y^1 < \bar{v}) 1(Y^* < v) \frac{1}{d} (v - Y^*) \right]$$

The threshold type \bar{v} must be indifferent, so

$$\begin{aligned} \pi_B(\bar{v}) &= \mathbb{E} [1(Y^1 < \bar{v})(\bar{v} - p)] \\ &= \mathbb{E} \left[1(Y^1 < \bar{v}) \frac{1}{d} (\bar{v} - Y^*) \right] = \pi_T(\bar{v}). \end{aligned} \quad (7)$$

We next show that no other type wants to deviate. Suppose $v > \bar{v}$. Then:

$$\begin{aligned} \pi_B(v) &= \mathbb{E}[1(v > Y^1 > \bar{v})(v - Y^1)] \\ &\quad + \mathbb{E}[1(Y^1 < \bar{v})(v - \bar{v})] + \mathbb{E}[1(Y^1 < \bar{v})(\bar{v} - p)] \\ &\geq \mathbb{E}[1(Y^1 < \bar{v})(v - \bar{v})] + \pi_T(\bar{v}) \\ &= \mathbb{E}[1(Y^1 < \bar{v})(v - \bar{v})] + \mathbb{E}[1(Y^1 < \bar{v}) \frac{1}{d} (\bar{v} - Y^*)] \\ &\geq \frac{1}{d} (\mathbb{E}[1(Y^1 < \bar{v})(v - \bar{v})] + \mathbb{E}[1(Y^1 < \bar{v})(\bar{v} - Y^*)]) \\ &= \pi_T(v) \end{aligned}$$

Similarly, for $v < \bar{v}$, we have:

$$\begin{aligned} \pi_T(v) &= \mathbb{E} \left[1(Y^1 < \bar{v}) 1(Y^* < v) \frac{1}{d} (v - Y^*) \right] \\ &\geq \mathbb{E} \left[1(Y^1 < \bar{v}) \frac{1}{d} (v - Y^*) \right] \\ &= \frac{1}{d} \mathbb{E} [1(Y^1 < \bar{v})(\bar{v} - Y^*)] \\ &\quad - \frac{1}{d} \mathbb{E} [1(Y^1 < \bar{v})(\bar{v} - v)] \end{aligned}$$

By combining this with Eq. (7), we get,

$$\begin{aligned} \pi_T(v) &\geq \mathbb{E} [1(Y^1 < \bar{v})(\bar{v} - p)] - \mathbb{E} [1(Y^1 < \bar{v})(\bar{v} - v)] \\ &= \pi_B(v) \end{aligned}$$

Next, we show a \bar{v} satisfying Eq. (1) exists and is unique. Suppose $d > 1$. Then the right hand side of Eq. (1) is a function of \bar{v} with first derivative $\frac{1}{d}(1 - \frac{\partial}{\partial \bar{v}} \mathbb{E}[Y^* | Y^1 < \bar{v}]) < 1$. Since at $\bar{v} = 0$ it has value $p > 0$ and globally has slope less than 1, it must cross the 45° line exactly once. Thus there is exactly one solution to the implicit Eq. (1). On the other hand, suppose $d = 1$; then by assumption $p < E[Y^1]$. Hence, Eq. (1) simplifies to $\mathbb{E}[Y^1 | Y^1 < \bar{v}] = p$, which has a solution since $\mathbb{E}[Y^1 | Y^1 < \bar{v}] = p < \mathbb{E}[Y^1]$.

Finally, we need to argue there are no other pure strategy equilibria. Let A be the set of types who elect BIN, v_A be the infimum of this set and v_B be the supremum of its complement. Since \bar{v} is uniquely defined, any such equilibrium cannot have a threshold form, so $v_B > v_A$. Then reasoning similar to the above shows that $v_A - p \geq \frac{1}{d} \mathbb{E}[v_A - Y^* | Y_{-i} \notin A]$ but $v_B - p < \frac{1}{d} \mathbb{E}[v_B - Y^* | Y_{-i} \notin A]$, which implies $v_B < v_A$, a contradiction.

B. PROOF OF THEOREM 2

TAC Reserve: Fix (r, d) and $\bar{v} \in (\bar{\omega}_L, \bar{\omega}_H)$. Then define $p(r)$ implicitly as the BIN price that holds \bar{v} constant as r changes. The effects of increasing the reserve r slightly are two: first, if only one person has valuation above \bar{v} , you can charge them a slightly higher BIN price; second, if all bidders have low valuation, increasing the reserve may raise the expected payment of some types, while decreasing the probability the object will be sold.

The marginal increase in revenue due to the first effect is:

$$n\alpha(1 - \alpha)^{n-1} \frac{1}{d} \Pr(Y^d \leq r)$$

If there are no BIN bidders, we look at the expected payment from each of the k highest bidders:

$$\begin{aligned} &(1 - \alpha)^n \frac{1}{d} \sum_{k=1}^d \left[\sum_{j=k}^d \binom{n}{j} (1 - F_L(r))^j F_L(r)^{n-j} r \right. \\ &\quad \left. + \int_r^{\bar{v}} \frac{n!}{d!(n-1-d)!} f(s) F_L(s)^{n-d-1} (1 - F_L(s))^d ds \right] \end{aligned}$$

Taking a first order condition in r , canceling telescoping terms and simplifying the double summations:

$$(1 - \alpha)^n \frac{1}{d} \sum_{k=1}^d \binom{n}{k} k (1 - F_L(r))^{k-1} F_L(r)^{n-k} (1 - F_L(r) - r f_L(r))$$

Summing both marginal effects, expanding $P(Y^d \leq r)$ we set the following equation to zero:

$$\begin{aligned} &n\alpha \left(\sum_{j=0}^{d-1} \binom{n-1}{j} (1 - F_L(r))^j F_L(r)^{n-1-j} \right) + \\ &(1 - \alpha) \sum_{k=1}^d \binom{n}{k} k (1 - F_L(r))^{k-1} F_L(r)^{n-k} (1 - F_L(r) - r f_L(r)) \end{aligned}$$

Changing the summation limits, factorizing and eliminating constants:

$$\alpha + (1 - F_L(r) - r f_L(r)) (1 - \alpha) = 0$$

Since $F_L(r) = \frac{1}{1-\alpha}F(r)$ and $f_L(r) = \frac{1}{1-\alpha}f(r)$ we get $r^* = \frac{1-F(r^*)}{f(r^*)}$. Note that by Assumption 1, this point is unique.

BIN Price: Let H be the total number of bidders with high valuations. Also, let V^j be a random variable corresponding to the j -th highest valuations among the bidders. Define $t_H(d, r, j)$ to be the expected revenue obtained from the bidder with the j -th highest valuations among the bidders, conditioned that H agents have high valuations. Namely,

$$t_0(d, r, j) = \frac{1}{d} \Pr(V^j > r | H = 1) \mathbb{E}[\max(V^{d+1}, r) | V^j > r, H = 1]$$

$$t_1(d, r, j) = \frac{1}{d} \Pr(V^j > r | H = 0) \mathbb{E}[\max(V^{d+1}, r) | V^j > r, H = 0]$$

The revenue function when $H \leq 1$ is:

$$\pi(d, p, r) = n\alpha(1-\alpha)^{n-1} \left((1 - F_H(\bar{v}))p + F_H(\bar{v}) \sum_{j=1}^d t_1(d, r, j) \right)$$

$$+ (1-\alpha)^n \sum_{j=1}^d t_0(d, r, j)$$

Now since $p = \frac{d-1}{d}\bar{v} + \frac{1}{d}E[Y^* | Y^1 < \bar{v}] = \frac{d-1}{d}\bar{v} + t_1(d, r, 1)$, we can substitute to get:

$$\pi(d, p, r) = n\alpha(1-n\alpha) \left[(1 - F_H(\bar{v})) \frac{d-1}{d}\bar{v} + t_1(d, r, 1) \right]$$

$$+ F_H(\bar{v}) \sum_{j=2}^d t_1(d, r, j) + (1-n\alpha) \sum_{j=1}^d t_0(d, r, j)$$

Taking a first order condition of the right hand side in $\bar{v} \in [\underline{\omega}_H, \bar{\omega}_H]$:

$$\frac{d-1}{d} ((1 - F_H(\bar{v})) - f_H(\bar{v})\bar{v}) + f_H(\bar{v}) \left(\sum_{j=2}^d t_1(d, r, j) \right) = 0$$

Re-arranging terms:

$$\bar{v} - \frac{(1 - F_H(\bar{v}))}{f_H(\bar{v})} = \frac{d}{d-1} \left(\sum_{j=2}^d t_1(d, r, j) \right)$$

Note that by Assumption 1, if there exists a solution, it is unique. Solving for \bar{v} :

$$\bar{v} = \frac{d}{d-1} \sum_{j=2}^d t_1(d, r, j) + \frac{(1 - F_H(\bar{v}))}{f_H(\bar{v})}$$

or in terms of p :

$$p^* = \sum_{j=1}^d t_1(d, r, j) + \frac{(d-1)(1 - F_H(\bar{v}))}{df_H(\bar{v})}$$

$$= R_1^{TAC}(d, r) + \frac{(d-1)(1 - F(\bar{v}))}{df(\bar{v})}$$

If there exists no interior solution that satisfies the first order condition, then $\bar{v} = \tau$.

C. PROOF OF THEOREM 3

We prove the theorem using the approach proposed by Myerson [13]. Some definitions are in order to describe the ironing approach; for the sake of consistency, we would

use the same notation as [13]. This approach requires that the distribution of the values, F , to be strictly increasing.³ Hence, we consider the following distribution of the values.

$$f_\varepsilon(x) = \begin{cases} \beta f_L(x) & x \in [\underline{\omega}_L, \bar{\omega}_L] \\ \varepsilon & x \in (\bar{\omega}_L, \underline{\omega}_H) \\ f_H(x)\alpha & x \in [\underline{\omega}_H, \bar{\omega}_H] \end{cases}$$

$$F_\varepsilon(x) = \begin{cases} \beta F_L(x) & x \in [\underline{\omega}_L, \bar{\omega}_L] \\ \beta + \varepsilon(x - \bar{\omega}_L) & x \in (\bar{\omega}_L, \underline{\omega}_H) \\ (1-\alpha) + \alpha F_H(x - \underline{\omega}_H) & x \in [\underline{\omega}_H, \bar{\omega}_H] \end{cases}$$

where $\beta + \varepsilon(\underline{\omega}_H - \bar{\omega}_L) + \alpha = 1$. Hence, as ε tends to 0 we get the original model back. Myerson [13] showed that if $\psi_\varepsilon(x) = x - \frac{1-F_\varepsilon(x)}{f_\varepsilon(x)}$ is strictly increasing, then, the optimal mechanism would always allocate the item to the bidder with the maximizes $\psi(\cdot)$. However, if this function is not increasing, then one needs to “iron” the virtual values. Note that for $x \in (\bar{\omega}_L, \underline{\omega}_H)$, and small enough ε ,

$$\psi_\varepsilon(x) = x - \frac{1 - (\beta + \varepsilon(x - 1))}{\varepsilon} < 0 \leq \psi_\varepsilon(r^*)$$

For $q \in [0, 1]$, let $F_\varepsilon^{-1}(q)$ be the inverse of $F_\varepsilon(\cdot)$. Define:

$$h(q) = F_\varepsilon^{-1}(q) - \frac{1-q}{f_\varepsilon(F_\varepsilon^{-1}(q))}$$

$$H(q) = \int_0^q h(y) dy$$

$$G(q) = \text{conv}H(q)$$

$$= \min_{\lambda, r_1, r_2 \in [0, 1], \lambda r_1 + (1-\lambda)r_2 = q} \{\lambda H(r_1) + (1-\lambda)H(r_2)\}$$

This implies that $G(\cdot)$ is the highest convex function on $[0, 1]$ such that $G(q) \leq H(q)$ for every q .

Define $\phi(v) = G'(F(v))$ as the *virtual value* of type v . By Theorem 6.1 [13], the optimal mechanism randomly allocates the item to one of the bidders with the highest positive virtual value. In the following, we compute the virtual values.

LEMMA 1. Let $q^* = (1-\alpha)v^*$ and v^* be the solution of

$$-F^2(v^*) + (2-\alpha)F(v^*) + \alpha(\underline{\omega}_H - v^*)f(v^*) = 1 - \alpha.$$

Under the assumption of Theorem 3, as $\varepsilon \rightarrow 0$,

$$G'(q) = \begin{cases} h(q) & q \in [0, q^*] \\ h(q^*) & q \in (q^*, 1-\alpha) \\ h(q) & q \in [1-\alpha, 1] \end{cases}$$

PROOF. First note that $H(q)$ is convex in $[0, \beta]$ because of the assumption that $x - \frac{1-F(x)}{f(x)}$ is increasing in $[\underline{\omega}_L, \bar{\omega}_L]$. It is also decreasing in $[0, q_0]$ and increasing in $[q_0, \beta]$, where $q_0 = F(r^*)$ is the minimum of $H(\cdot)$ in this range. Also, observe that $H(q)$ is decreasing in $[\beta, 1-\alpha]$ because $h(q) < 0$ in this interval. In addition, by Assumption 1, $H(q)$ is increasing and convex in $[1-\alpha, 1]$. Therefore, $G(\cdot)$ includes the tangent line from the point $(1-\alpha, H(1-\alpha))$ to $H(q)$ in

³ See [11, 15] for optimal mechanisms when distributions have discrete support.

$[0, \beta]$. Let q^* be the tangent point. We have

$$G(q) = \begin{cases} H(q) & q \in [0, q^*] \\ \frac{(q-q^*)H(q^*)+(1-\alpha-q)H(1-\alpha)}{1-\alpha-q^*} & q \in (q^*, 1-\alpha) \\ H(q) & q \in [1-\alpha, 1] \end{cases}$$

which immediately leads to the claim. In the rest we compute q^* . For $q \in [0, \beta]$,

$$\begin{aligned} H(q) &= \int_0^q \left(F_\varepsilon^{-1}(y) - \frac{1-y}{f_\varepsilon(F_\varepsilon^{-1}(y))} \right) dy \\ &= \int_{\underline{\omega}_L}^{F_\varepsilon^{-1}(q)} \left(x - \frac{1-F_\varepsilon(x)}{f_\varepsilon(x)} \right) f_\varepsilon(x) dx \\ &= \int_{\underline{\omega}_L}^{F_\varepsilon^{-1}(q)} ((x f_\varepsilon(x) + F_\varepsilon(x)) - 1) dx \\ &= (q-1)F_\varepsilon^{-1}(q) + \underline{\omega}_L \end{aligned}$$

In particular,

$$H(\beta) = (\beta-1)\bar{\omega}_L + \underline{\omega}_L$$

For $q \in (\beta, 1-\alpha)$, because $h(q) = \frac{2q-(1+\beta)}{\varepsilon} + \bar{\omega}_L$, we get

$$\begin{aligned} H(q) &= H(\beta) + \left[\frac{x^2 - (1+\beta - \varepsilon\bar{\omega}_L)x}{\varepsilon} \right]_\beta^q \\ &= (\beta-1)\bar{\omega}_L + \underline{\omega}_L + \frac{q^2 - \beta^2 - (q-\beta)(1+\beta - \varepsilon\bar{\omega}_L)}{\varepsilon} \\ &= (q-1)\bar{\omega}_L + \underline{\omega}_L + (q-\beta)\frac{q-1}{\varepsilon} \end{aligned} \quad (8)$$

In particular,

$$H(1-\alpha) = -\alpha\bar{\omega}_L + \underline{\omega}_L + (1-\alpha-\beta)\frac{-\alpha}{\varepsilon} = -\alpha\omega_H + \underline{\omega}_L$$

To iron the distribution, we compute the tangent from $H(1-\alpha)$ to the $H(q)$, for $q \in [0, 1-\alpha]$. Note that if q^* is the tangent point then

$$h(q^*) = \frac{H(1-\alpha) - H(q^*)}{1-\alpha-q^*} \quad (9)$$

Observe that by Eq. (8) we have

$$\begin{aligned} &\frac{H(1-\alpha) - H(q^*)}{1-\alpha-q^*} \\ &= \frac{(-\alpha\omega_H + \underline{\omega}_L) - ((q^*-1)F_\varepsilon^{-1}(q^*) + \underline{\omega}_L)}{1-\alpha-q^*} \\ &= \frac{-\alpha\omega_H - (q^*-1)F_\varepsilon^{-1}(q^*)}{1-\alpha-q^*} \end{aligned}$$

Let $v^* = F_\varepsilon^{-1}(q^*)$, i.e., $q^* = F_\varepsilon(v^*) = \beta F_L(v^*)$. Therefore,

$$\frac{H(1-\alpha) - H(q^*)}{1-\alpha-q^*} = \frac{-\alpha\omega_H - (F_\varepsilon(v^*) - 1)v^*}{1-\alpha - F_\varepsilon(v^*)}$$

$$h(q^*) = v^* - \frac{1 - F_\varepsilon(v^*)}{f_\varepsilon(v^*)}$$

As $\varepsilon \rightarrow 0$, the $F_\varepsilon(\cdot) \rightarrow F(\cdot)$. Plugging into Eq. (9) we get

$$\begin{aligned} &(v^* f(v^*) - 1 + F(v^*))(1-\alpha - F(v^*)) \\ &= f(v^*)(-\alpha\omega_H - (F(v^*) - 1)v^*) \end{aligned}$$

Hence, rearranging the terms,

$$-F^2(v^*) + (2-\alpha)F(v^*) + \alpha(\omega_H - v^*)f(v^*) = 1 - \alpha$$

Observe that only if $H(1-\alpha) > H(q_0)$, then $h(q^*)$ is positive.

$$-\alpha\omega_H + \underline{\omega}_L \geq (q_0 - 1)F^{-1}(q_0) + \underline{\omega}_L = (F(r^*) - 1)r^* + \underline{\omega}_L$$

This is equivalent to $\alpha\omega_H \leq (1 - F(r^*))r^*$.

Finally, observe that

$$h(q^*) \leq h(\beta) = \bar{\omega}_L \leq \omega_H - \frac{1 - F_\varepsilon(\omega_H)}{f_\varepsilon(\omega_H)}$$

which shows that $G(\cdot)$ is convex and completes the proof. \square

By the above lemma, if $\alpha\omega_H \leq (1 - F(r^*))r^*$, the (ironed) virtual value of a bidder with value v by $\phi(v)$ is defined in Eq. (6). Note that for $v \in [0, r^*)$ the virtual value is in fact negative, but because the mechanism only allocates to positive virtual values, it is equivalent to let the virtual value equal to 0. Also, note that if $v^* < r^*$, then the mechanism is equivalent to the second-price mechanism with reserve ω_H . Now we describe the payments, assuming $v^* > r^*$.

Let v denoted the vector of valuations (bids) of bidders. By [13], the expected payment of each bidder i , p_i , is:

$$p_i(v) = v_i - \int_0^{v_i} \rho_i(x, v_{-i}) dx \quad (10)$$

where $\rho_i(x, v_{-i})$ is the probability of allocation to bidder i with value x , when the values of other bidders is v_{-i} .

Recall that the optimal mechanism awards the good to the bidder with the highest valuation, breaking ties *randomly*, provided the virtual valuation is positive. Let w_1 be the valuation of the winning bidder; let w_2 be the valuation of the bidder with the second-highest virtual valuation (again break ties randomly).

Also let k be the number of bidders with bid in $[v^*, \omega_H)$. Then, by Eq. (10), the price for the winning bidder is computed as follows:

$w_1 \in [r^*, v^*)$: By Eq. (10),

$$p = w_1 - \int_0^{w_1} \rho_i(x, v_{-i}) dx = w_1 - \int_{\max\{w_2, r^*\}}^{w_1} 1 dx$$

Because the probability of the allocation before $\max\{w_2, r^*\}$ is 0 and after that is 1,

$$p = w_1 - \int_{\max\{w_2, r^*\}}^{w_1} 1 dx$$

Therefore, $p = \max\{w_2, r^*\}$.

$w_1 \in [v^*, \omega_H)$: If $k = 1$, similar to the previous case, $p = \max\{w_2, r^*\}$. Otherwise, by Eq. (10), the expected payment of each bidder with $v_i \in [v^*, \omega_H)$ is $\frac{1}{k}v^*$. This is equivalent to charging the winner v^* .

$w_1 \in [\omega_H, \bar{\omega}_H]$: If $w_2 \leq v^*$ or $w_2 \geq \omega_H$ then similar to the first case, $p = \max\{w_2, r^*\}$. If $w_2 \in [v^*, \omega_H)$, then $p = \frac{1}{k+1}(k\omega_H + v^*)$.

$$\begin{aligned} p &= w_1 - \int_0^{w_1} \rho_i(x, v_{-i}) dx \\ &= \omega_H - \int_{v^*}^{\omega_H} \frac{1}{k+1} dx \\ &= \frac{1}{k+1}(k\omega_H + v^*). \end{aligned}$$