

Buy-it-now or Take-a-chance: A New Pricing Mechanism for Online Advertising*

L. Elisa Celis Gregory Lewis Markus M. Mobius
Hamid Nazerzadeh

October 3, 2011

Abstract

Increasingly sophisticated tracking technology offers publishers the ability to offer targeted advertisements to advertisers. Such targeting enhances advertising efficiency by improving the match quality between advertisers and users, but also thins the market of interested advertisers. Using bidding data from Microsoft’s Ad Exchange (AdECN) platform, we show that there is often a substantial gap between the highest and second highest willingness to pay. This motivates our new BIN-TAC mechanism, which is effective in extracting revenue when such a gap exists. Bidders can “buy-it-now”, or alternatively “take-a-chance” in an auction, where the top $d > 1$ 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 distributions, this mechanism achieves similar allocations and revenues as Myerson’s optimal mechanism, and outperforms the second-price auction with reserve. For the AdECN data, we use structural methods to estimate counterfactual revenues, and find that our BIN-TAC mechanism improves revenue by 11% relative to an optimal second-price auction.

Keywords: Advertising, Auctions, Mechanism Design

JEL Classification Numbers: D44, L86

*Department of Computer Science, University of Washington (ecelis@cs.washington.edu), Department of Economics, Harvard University (glewis@fas.harvard.edu), Microsoft Research New England (mobius@microsoft.com), and USC Marshall School of Business (nazerzad@marshall.usc.edu) respectively. All authors are grateful to Microsoft Research New England for their hospitality, and Greg would also like to thank NET Institute (www.NETinst.org) for financial support. We thank Susan Athey, Alessandro Bonatti, Michael Grubb, Mallesh Pai and Lee Zen for fruitful discussions that improved this work. Finally we thank Lee Zen and members of Microsoft AdECN team — in particular Ian Ferreira, Keith Hawley, and Brandon Zirkle — for providing the data.

1 Introduction

Many Internet companies generate revenue by selling the advertisement space on their web-pages. 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 where perhaps only a single advertiser has a high willingness to pay. 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 (Bergemann and Bonatti 2010, Levin and Milgrom 2010).

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 participating bidders to fellow sportswear firms in New York City. Without competition, revenue performance may be poor (Levin and Milgrom 2010).

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 (Ghosh, Nazerzadeh and Sundararajan 2007, Even-Dar, Kearns and Wortman 2007, McAfee, Papineni and Vassilvitskii 2010).

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.

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 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 Myerson (1981). 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 continues to outperform the second price auction when the supports overlap.

In the last part of the paper, we demonstrate our mechanism’s effectiveness using data from Microsoft’s Ad Exchange (AdECN) platform for selling display advertising. Since the current auction format is a second-price auction, and it is weakly dominant for an advertiser to bid their valuation, we can interpret bids as valuations. It becomes relatively easy to then 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 (1981) 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 (2009) discuss the benefits of simple mechanisms, and show a variety of examples where they approximate the optimal expected revenue.

The question of whether sellers should provide information that allows buyers to “target” their bids is a question that arises in the analysis of optimal seller disclosure (see for example Bergemann and Pesendorfer (2001)). Here we specialize to a mechanism that treats all bidders symmetrically, and proceeds sequentially. Sequential screening models have been proposed for revenue maximization in dynamic environments. For instance, Courty and Li (2000) 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 Bergemann and Said (2010) 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 (see for example Rochet and Chone (1998) and Chawla, Hartline, Malec and Sivan (2010)). Our model deals with the static case where types are single-dimensional and have a mixture form and buyers know their valuation from the outset. Additionally, our model considers only the private value setting. Abraham, Athey, Babioff and Grubb (2010), 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 can sometimes lead to low revenue in this market. For further discussion on advertisement exchange markets see Muthukrishnan (2010).

Organization The paper proceeds in four parts. First, we describe the AdECN market, providing some interesting and (to our knowledge) novel observations about this display advertising 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 Display Advertising 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 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 the top panel of Figure 2, a kernel density estimate of this gap.




Figure 1: **Bids over Time.** The figure shows the bids made by the four leading bidders in our data on 50 randomly chosen impressions for a given product.

The bottom panel shows 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 motivates our search for a middle ground: a mechanism that retains the simplicity of the SPA while getting nearly optimal revenue performance.

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




Figure 2: **Bidding Gap and Virtual Valuations.** The top panel shows a kernel density estimate of the pdf of the (normalized) gap between the highest and second highest bids in auctions for ads A and B. The bottom panel shows the virtual valuations for all bids on ads A and B.

maximum of the reserve and the $(d + 1)$ -th bid.¹ We call d the *TAC-parameter*.

3.1 The Environment

Motivated by the observations in Section 2, we define a stylized environment in which the platform allows targeting. As a result, bidders generally do not “match” with the specific user or publisher characteristics and have low valuation, but occasionally one or more bidders “match” and have high valuation. This hardens in the idea that targeting may make markets thin. Assume n bidders participate in an auction for a single good which is valued at zero by

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.

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 $V_H \sim F_H$. The event $X = 1$ indicates that a “match” has occurred, and we are generally interested in the case where α is close to 0. 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. Here, the bidders know their exact valuation, but one could imagine a case in which the platform does not disclose X to the advertisers. In that case they may have draws from F_L and F_H but be uncertain as to which of them is their valuation. We shall analyze such a “no targeting” case later in the paper. 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 (Myerson 1981). We assume that $\psi(v)$ is continuous, increasing and single-crosses zero over the regions $[\underline{\omega}_L, \bar{\omega}_L]$ and $[\underline{\omega}_H, \bar{\omega}_H]$. But 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 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 $b' < v$, when they win 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 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. Let $R(v, d, r)$ be the conditional expected revenue from a TAC auction when the highest valuation is exactly equal to $\bar{v}$. Then we have the following theorem.

Theorem 2 (Optimal Buy Price and Reserve) *The revenue-maximizing TAC reserve r satisfies:*

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

There are exactly two solutions to this equation. For $r_1^ \in [\underline{\omega}_H, \bar{\omega}_H]$, $d = 1$ and any $p \geq r_1^*$ is optimal. For $r_2^* \in [\underline{\omega}_L, \bar{\omega}_L]$, then if a solution exists with $\bar{v}(p^*, d, r) \in [\underline{\omega}_H, \bar{\omega}_H]$, the optimal BIN price is given by:*

$$p^* = \frac{R(\bar{v}(p^*), r, d)F(\bar{v}(p^*)) - (n-1)(1 - F(\bar{v}(p^*)))\bar{v}(p^*)}{F(\bar{v}(p^*)) - (n-1)(1 - F(\bar{v}(p^*)))} + \frac{(1 - F(\bar{v}(p^*)))F(\bar{v}(p^*))/\bar{v}_p(p^*)}{f(\bar{v}(p^*))} \quad (3)$$

where $\bar{v}_p = \frac{\partial \bar{v}(p, d, r)}{\partial p}$. If no interior solution exists, p^* solves $\bar{v}(p^*, d, r) = \underline{\omega}_H$.

Equation 2 is somewhat surprising; the optimal TAC reserve is exactly the standard reserve in Myerson (1981), 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 in some sense sets a reserve at $\bar{v}$. If two bidders meet the reserve, he gets the second highest bid; if only one, the BIN price; and if none, he gets the TAC revenue. So a marginal increase in the “reserve” has three effects. First, if the highest bidder has valuation exactly equal to the reserve, following an increase he will shift from BIN to TAC. This costs the seller $p - R(\bar{v}, d, r)$. Second, if the second highest bidder has valuation equal to the reserve, an increase will knock him out of the BIN auction, and the seller’s revenue falls by $\bar{v}(p) - p$. Finally, if the highest bidder is above the reserve and the second highest is below, an increase gains the seller $p'(\bar{v})$. Working out the probabilities of these various

events, and inverting $p'(\bar{v})$ by the implicit function theorem, we get the result.

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 three benchmark mechanisms, the second price auction with optimal reserve and targeting, the second price auction without allowing targeting, and the fully optimal Myerson (1981) mechanism.

Second Price Auctions The second price auction (SPA) is widely used in practice, which owes something to the fact that it is both strategy-proof (bidding one's valuation is optimal), and in many cases revenue maximizing. In the case where the platform allows targeting, the SPA with reserve r is just a special case of BIN-TAC for parameters $p = \bar{\omega}_H$ and $d = 1$. This is clear: no one takes the BIN, and since $d = 1$ the TAC auction is just an SPA. More generally, any BIN-TAC auction with parameters $p \geq r$ and $d = 1$ generates the same revenue. To see this, notice that for any BIN price $p \geq r$, there is as usual some threshold type that is indifferent between BIN and TAC. But that type will get the object with certainty in TAC if $d = 1$, and so the BIN price must be equal to his expected payment in the TAC auction. So whenever two types have valuation above the threshold, BIN-TAC generates the SPA revenue; whenever one type is above the threshold, BIN-TAC generates the expected second highest bid, equal to the SPA revenue; and when all types are below the threshold, it is exactly an SPA.

As we have already noted, the SPA mechanism with targeting may perform badly if the valuations are much higher when a match occurs, and matches occur rarely. An alternative that has been widely suggested is to continue using the second price mechanism, but hide the information that reveals matches. To formalize this, we temporarily assume that F_L and F_H are degenerate with all mass on v_L and v_H respectively. Learning X reveals to the advertiser whether he has high valuation v_H or low valuation v_L . In the absence of this information, his expected valuation is $(1 - \alpha)v_L + \alpha v_H$.

Let the revenue from the SPA without revealing any info be denoted $R_{\text{No Info}}^{SPA}$, the revenue from the SPA with info be denoted R_{Info}^{SPA} and the revenue from the optimal BIN-TAC mechanism be denoted $R^{BIN-TAC}$. We know already that $R^{BIN-TAC} \geq R_{\text{Info}}^{SPA}$.

Theorem 3 *Suppose F_L and F_H are degenerate. Then the revenue from the optimal BIN-TAC model is higher than the optimal second price auction without targeting, strictly for $n \geq 3$.*

Optimal Mechanism We now compare the BIN-TAC mechanism to the optimal mechanism. As argued earlier, we need to use an ironing procedure to do this. We show in the appendix that whenever $\alpha \underline{\omega}_H \geq r^*(1 - F(r^*))$, the optimal mechanism is a second-price auction with reserve $\underline{\omega}_H$. So assume the opposite.

Then, there exists v^* , $r^* \leq v^* \leq \bar{\omega}_L$, such that

$$(2 - \alpha - F(v^*))F(v^*) + \alpha(\underline{\omega}_H - v^*)f(v^*) = 1 - \alpha \quad (4)$$

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

$$\phi(v) = \begin{cases} 0 & v \in [\underline{\omega}_L, r^*) \\ \psi(v) & v \in [r^*, v^*] \\ \psi(v^*) & v \in (v^*, \underline{\omega}_H) \\ \psi(v) & v \in [\underline{\omega}_H, \bar{\omega}_H], \end{cases} \quad (5)$$

The allocation procedure is as follows: award the good to the bidder with the highest ironed virtual valuation, breaking ties uniformly at random, provided the virtual valuation is positive. Notice that all types between v^* and $\underline{\omega}_H$ get the same virtual valuations, and therefore if they tie, the winner is selected at random. Like the BIN-TAC mechanism, this creates the potential for inefficiency, but allows additional revenue extraction from higher types.

The payments are determined as follows. Whenever the virtual valuation of the second highest bidder has a unique inverse (i.e. outside of the ironed region between v^* and $\underline{\omega}_H$), the winning bidder pays the maximum of the reserve and the valuation of the second highest bidder (as in a second price auction with reserve). Whenever both the highest and second highest bidder have virtual valuations in the ironed region, the required payment is v^* .

Finally, when the winning bidder has valuation above $\underline{\omega}_H$, but k other bidders have valuations in the ironed region, the winner pays $\frac{1}{k+1}(k\underline{\omega}_H + v^*)$. This last condition is easily derived from incentive compatibility: the bidder on the top margin of the ironed region, type $\underline{\omega}_H$, gets a payoff of $v - \frac{1}{k+1}(k\underline{\omega}_H + v^*)$ when the highest bidder; but could alternately pretend to be in the ironed region, with payoff $\frac{1}{k+1}(v - v^*)$ — these two are identical.

Theorem 4 (Optimal Mechanism) *Suppose $\psi(\bar{\omega}_L) \leq \psi(\underline{\omega}_H)$. If $\alpha\underline{\omega}_H \geq r^*(1 - F(r^*))$, then the optimal mechanism is the second-price auction with reserve $\underline{\omega}_H$. If $\alpha\underline{\omega}_H < r^*(1 - F(r^*))$, 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. It is easy to prove that as either $n \rightarrow \infty$, $\alpha \rightarrow 1$ or $\underline{w}_H/\bar{w}_L \rightarrow \infty$, the BIN-TAC mechanism converges to the optimal mechanism. This, however, is not particularly interesting (a second-price mechanism will also converges 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 4, we have

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

Also, recall that the optimal second-price auction is equivalent to a BIN-TAC mechanism with $d = 1$. Table 2 below compares the expected revenue and welfare obtained by these mechanisms for $n = 5$ and $\tau = 3$ and $\alpha = 0.05$. As you can see from the table, the performance of BIN-TAC is close to OPT (about 96%), much better than the optimal SPA (85%). Figures 3 helps explain this. The top panel depicts the probability of allocation and the bottom panel the expected payment of a bidder, assuming the values of the other 4 bidders are distributed according to the distribution described above. As you can see, the BIN-TAC mechanism approximates the discontinuous increase in allocation probability at v^* with a smooth curve, whereas the SPA increases the probability of allocation to a much




Figure 3: **Comparison of Allocations and Payments.** Allocation probabilities (top panel) and expected payments (bottom panel) for the OPT, SPA and BIN-TAC mechanisms when the distributions F_L and F_H are uniform. The x -axis corresponds to the bid.

higher level. As a result, the SPA cannot extract revenue from the high types (who could easily pretend to be a lower type without losing much), while the BIN-TAC mechanism has similar revenue performance to OPT at the top.

4 Simulations

We would like to test our mechanism against the benchmarks in a wider setting than those considered thus far. We drop the assumption that F_L and F_H have disjoint support. The optimal BIN-TAC mechanism is reasonably easy to calculate. Nothing in the proof of Theorem 2 required the disjoint supports for determining r^* and p^* , and so these can be solved for numerically for each d . Thus the optimization problem reduces to a one dimensional discrete optimization problem, which can be quickly solved. By contrast, finding the optimal mechanism requires solving for the ironing region, a 2-dimensional optimization problem in continuous controls (although here too there are numeric approaches which may be preferable).

To show the effectiveness of BIN-TAC, we report the expected revenue relative to the optimal mechanism OPT, as well as the optimal second price auction without targeting, and the

optimal second price auction with targeting.

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; $F_L \sim N$, $F_L \sim \log N$, where both have mean 1 and standard deviation 0.5. We allow Δ , n and α to vary across experiments, and compute r^* , p^* and d^* as discussed. The results are presented in Figures 4, 5 and 6.

The default parameters we consider are $n = 10$, $\Delta = 5$, and $\alpha = .05$, and we vary one parameter at a time. Each experiment is repeated for 100000 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 4 and 5 show how as either n or α increases, we approach the performance of the optimal mechanism. 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 6 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 performs 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 Empirical Analysis

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 shows whether our mechanism has the potential to improve platform revenues in a less stylized environment than that of our theoretical model.

5.1 Data

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.

5.2 Estimation Approach

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 over-fitting 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 $j = 1 \cdots 10$ 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 distribution 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^* .

5.3 Results

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 nonnegative. The results are shown in Table 3. 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

²See (Ostrovsky and Schwarz 2009) for a further discussion on reserve prices.

BIN-TAC mechanism is effective in extracting revenue, yet still allows targeting.

6 Conclusion and Future Work

We presented the BIN-TAC mechanism, particularly suited for environments where the distribution of valuations is irregular. We showed this mechanism closely approximates Myerson’s optimal mechanism with ironing, achieving similar allocations and revenues performs at least as well — and often much better — than the optimal second-price-auction. In addition, in a sequence of simulations we showed that our mechanism is flexible and applicable in many settings. We demonstrated this further by applying our mechanism to data from Microsoft’s AdECN platform, and attained a marked increase in revenue.

References

- Abraham, Ittai, Susan Athey, Moshe Babioff, and Michael Grubb**, “Peaches, Lemons and Cookies: The Impact of Information Structure on Revenue in Common Value Auctions,” 2010. working paper, Microsoft Research.
- Bergemann, Dirk and Alessandro Bonatti**, “Targeting in Advertising Markets: Implications for Offline vs. Online Media,” 2010. Cowles Foundation Discussion Paper 1758R.
- and **Maher Said**, “Dynamic Auctions: A Survey,” *Wiley Encyclopedia of Operations Research and Management Science*, 2010.
- and **Martin Pesendorfer**, “Information Structures in Optimal Auctions,” 2001. Cowles Foundation Discussion Paper 1323.
- Chawla, Shuchi, 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)” 2010, pp. 311–320.
- Courty, Pascal and Hao Li**, “Sequential Screening,” *The Review of Economic Studies*, 2000, 67 (4), pp. 697–717.

- Even-Dar, Eyal, Michael J. Kearns, and Jennifer Wortman**, “Sponsored Search with Contexts,” in “Internet and Network Economics, Third International Workshop (WINE)” 2007, pp. 312–317.
- Ghosh, Arpita, Hamid Nazerzadeh, and Mukund Sundararajan**, “Computing Optimal Bundles for Sponsored Search,” in “Internet and Network Economics, Third International Workshop (WINE)” 2007, pp. 576–583.
- Hartline, Jason D. and Tim Roughgarden**, “Simple versus Optimal Mechanisms,” in “Proceedings of the 10th ACM Conference on Electronic Commerce” 2009.
- Levin, Johathan and Paul Milgrom**, “Online Advertising: Heterogeneity and Conflation in Market Design,” *American Economic Review, Papers and Proceedings*, 2010, 100 (2).
- McAfee, R. Preston, Kishore Papineni, and Sergei Vassilvitskii**, “Maximally Representative Allocations for Guaranteed Delivery Advertising Campaigns,” 2010. Yahoo! Research Working Paper.
- Monteiro, Paulo Klinger and Benar Fux Svaiter**, “Optimal auction with a general distribution: Virtual valuation without densities,” *Journal of Mathematical Economics*, January 2010, 46 (1), 21–31.
- Muthukrishnan, S.**, “Ad Exchanges: Research Issues,” 2010. manuscript.
- Myerson, Roger B.**, “Optimal Auction Design,” *Mathematics of Operations Research*, 1981, 6 (1), 58—73.
- Ostrovsky, Michael and Michael Schwarz**, “Reserve Prices in Internet Advertising Auctions: A Field Experiment,” 2009. Working Paper.
- Pai, Mallesh and Rakesh Vohra**, “Optimal Auctions with Financially Constrained Bidders,” 2009. Working paper.
- Rochet, Jean-Charles and Philippe Chone**, “Ironing, Sweeping, and Multidimensional Screening,” *Econometrica*, July 1998, 66 (6), 788–826.

7 Appendix

7.1 Proof of Theorem 1

Fix an equilibrium of the form in the theorem, and let the payoffs to taking taking BIN be $\pi_B(v)$ and to TAC be $\pi_T(v)$. They are given by:

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

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

$$\begin{aligned}\pi_B(\bar{v}) &= \mathbb{E} [1(Y^1 < \bar{v})(\bar{v} - p)] \\ &= \mathbb{E} \left[1(Y^1 < \bar{v}) \frac{1}{d}(\bar{v} - Y^\star) \right] = \pi_T(\bar{v}).\end{aligned}\tag{6}$$

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

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

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

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

Combining this with Equation (6), 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.

7.2 Proof of Theorem 2

By assumption, $\psi(v)$ single-crosses zero from below over both $[\underline{\omega}_L, \bar{\omega}_L]$ and $[\underline{\omega}_H, \bar{\omega}_H]$, so the implicit equation for r^* has exactly two solutions. Also $p < r$ cannot be better than weakly optimal, since everyone takes the BIN price, turning it into an SPA with reserve p . We next show that whenever $p \geq r$, the equation must hold at an optimum. So fix d and $\bar{v} > p \geq r$ and define $p(r)$ implicitly as the BIN price that holds $\bar{v}$ constant as r changes. Then there are two effects of increasing the reserve r slightly: first, you can raise the BIN price without changing $\bar{v}$, increasing revenues when there is a single bidder who takes BIN; second, if all bidders TAC, increasing the reserve raises the expected payment of some types, while decreasing the probability of sale. The marginal increase in revenue due to the first effect is:

$$nF(\bar{v})^{n-1}(1 - F(\bar{v}))\frac{1}{d}\Pr(Y^d \leq r)$$

With probability $F(\bar{v})^n$ there are no BIN bidders. Writing $F_{\bar{v}}$ for $F(v|v < \bar{v})$:

$$F(\bar{v})^n \frac{1}{d} \sum_{k=1}^d \left[\sum_{j=k}^d \binom{n}{j} (1 - F_{\bar{v}}(r))^j F_{\bar{v}}(r)^{n-j} r \right. \\ \left. + \int_r^{\bar{v}} \frac{n!}{d!(n-1-d)!} f_{\bar{v}}(s) F_{\bar{v}}(s)^{n-d-1} (1 - F_{\bar{v}}(s))^d ds \right]$$

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

$$F(\bar{v})^n \frac{1}{d} \sum_{k=1}^d \binom{n}{k} k (1 - F_{\bar{v}}(r))^{k-1} F_{\bar{v}}(r)^{n-k} (1 - F_{\bar{v}}(r) - r f_{\bar{v}}(r))$$

Summing both marginal effects and expanding $P(Y^d \leq r)$:

$$n(1 - F(\bar{v})) \left(\sum_{k=0}^{d-1} \binom{n-1}{k} (1 - F_{\bar{v}}(r))^k F_{\bar{v}}(r)^{n-1-k} \right) + \\ F(\bar{v}) \sum_{k=1}^d \binom{n}{k} k (1 - F_{\bar{v}}(r))^{k-1} F_{\bar{v}}(r)^{n-k} (1 - F_{\bar{v}}(r) - r f_{\bar{v}}(r))$$

Changing summation limits, factorizing, eliminating constants and setting the FOC = 0:

$$(1 - F(\bar{v})) + (1 - F_{\bar{v}}(r) - r f_{\bar{v}}(r)) F(\bar{v}) = 0$$

Now since $F_{\bar{v}} = F(v|v < \bar{v}) = F(v)/F(\bar{v})$, we can simplify and solve to get $r^* = \frac{1-F(r^*)}{f(r^*)}$. For the higher solution $r_1^* \in [\underline{\omega}_H, \bar{\omega}_H]$, the virtual valuations are strictly increasing above the reserve, and so a second-price auction is optimal, implying $d = 1$ and $p \geq r_1^*$.

So fix the reserve at the second solution $r_2^* \in [\underline{\omega}_L, \bar{\omega}_L]$. Since the virtual valuations are increasing on $[\underline{\omega}_L, \bar{\omega}_L]$, it must be that p is set so that $\bar{v} \geq \underline{\omega}_H$. Write $R(\bar{v})$ for the expected revenue from the TAC mechanism when $V^1 = \bar{v}$. There are three effects of a marginal increase in $\bar{v}$. First, the second highest bidder may have valuation $\bar{v}$ and choose not to take BIN, which decreases revenue by $\bar{v} - p(\bar{v})$. The probability of $V^2 = \bar{v}$ is given by $n(n-1)f(\bar{v})(1-F(\bar{v}))F(\bar{v})^{n-2}$. The second is that that highest bidder may have valuation $\bar{v}$ and choose not to take BIN, reducing revenue by $p(\bar{v}) - R(\bar{v})$. This happens with probability $nf(\bar{v})F(\bar{v})^{n-1}$. Finally, the highest bidder may have valuation above $\bar{v}$ and the second highest below it, in which case this raises revenue by $p'(\bar{v})$. This happens with probability

$$n(1 - F(\bar{v}))F(\bar{v})^{n-1}.$$

Setting the sum of these effects equal to zero, and eliminating common factors we get:

$$f(\bar{v})(((n-1)(1-F(\bar{v}))(\bar{v}-p(\bar{v}))+F(\bar{v})(p(\bar{v})-R(\bar{v})))=(1-F(\bar{v}))F(\bar{v})p'(\bar{v})$$

Solving for the optimal $p(v^*)$:

$$p(v^*) = \frac{R(v^*)F(v^*) - (n-1)(1-F(v^*))v^*}{F(v^*) - (n-1)(1-F(v^*))} + \frac{(1-F(v^*))F(v^*)p'(v^*)}{f(v^*)}$$

Equivalently, the optimal BIN price is given by:

$$p^* = \frac{R(v(p^*), r, d)F(v(p^*)) - (n-1)(1-F(v(p^*)))v(p^*)}{F(v(p^*)) - (n-1)(1-F(v(p^*)))} + \frac{(1-F(v(p^*)))F(v(p^*))/\bar{v}_p(p^*)}{f(v(p^*))}$$

7.3 Proof of Theorem 3

The optimal BIN-TAC mechanism in this case threatens to randomize among all bidders to induce a high BIN price, so $d = n$, and has a reserve of v_L . It follows that the optimal BIN price is $v_H - \frac{(v_H - v_L)}{n} = \frac{n-1}{n}v_H + \frac{1}{n}v_L$. Then we can expand the revenue of BIN-TAC:

$$\begin{aligned} R^{BIN-TAC} &= (1 - (1 - \alpha)^n - n\alpha(1 - \alpha)^{n-1})v_H + n\alpha(1 - \alpha)^{n-1}p^{BIN} + (1 - \alpha)^nv_L \\ &= (1 - (1 - \alpha)^n - n\alpha(1 - \alpha)^{n-1})v_H + n\alpha(1 - \alpha)^{n-1}\left(\frac{n-1}{n}v_H + \frac{1}{n}v_L\right) + (1 - \alpha)^nv_L \\ &= (1 - (1 - \alpha)^n - \alpha(1 - \alpha)^{n-1})v_H + (\alpha(1 - \alpha)^{n-1} + (1 - \alpha)^n)v_L \end{aligned}$$

and similarly for the SPA without info:

$$\begin{aligned} R_{\text{no Info}}^{SPA} &= (\alpha v_H + (1 - \alpha)v_L)^{(2:n)} \\ &= (\alpha v_H + (1 - \alpha)v_L) \text{ (degeneracy)} \end{aligned}$$

Since these are both probability distributions with two point support, to rank revenues it suffices to show that the mass on v_L is lower under BIN-TAC. So we must show that $(\alpha(1 - \alpha)^{n-1} + (1 - \alpha)^n) < (1 - \alpha)$. After a bit of simple algebra, this is equivalent to showing $(1 - \alpha)(1 - \alpha(1 - \alpha)^{n-2} - (1 - \alpha)^{n-1}) \geq 0$, which holds by binomial expansion of 1 with equality for $n = 2$ and strictly for $n > 2$. This proves the claim.

7.4 Proof of Theorem 4

Since the payment structure is well-known given the ironed virtual valuations, the challenge is to compute the ironed virtual values. We follow the approach proposed by Myerson (1981). This approach requires the distribution of 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 $\beta + \varepsilon(\underline{\omega}_H - \bar{\omega}_L) + \alpha = 1$. As ε tends to 0 we get the original model back. We need to “iron” the virtual values. For $q \in [0, 1]$, let $F_\varepsilon^{-1}(q)$ be the inverse of $F_\varepsilon(\cdot)$. Define:

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

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

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

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 (Myerson 1981), the optimal mechanism randomly allocates the item to one of the bidders with the highest positive virtual value. We first show that the ironed virtual values are the same as the original virtual valuations, except for a set of quantiles between q^* and $(1 - \alpha)$:

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

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

³See (Monteiro and Svaiter 2010, Pai and Vohra 2009) for optimal mechanisms when distributions have discrete support.

Under the assumption of Theorem 4, 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 ??, $H(q)$ 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 $[0, \beta]$. 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

$$\begin{aligned} H(q) &= H(\beta) + \left[\frac{x^2 - (1 + \beta - \varepsilon \bar{\omega}_L)x}{\varepsilon} \right]_{\beta}^q \\ &= (\beta - 1)\bar{\omega}_L + \underline{\omega}_L + \frac{q^2 - \beta^2 - (q - \beta)(1 + \beta - \varepsilon \bar{\omega}_L)}{\varepsilon} \\ &= (q - 1)\bar{\omega}_L + \underline{\omega}_L + (q - \beta) \frac{q - 1}{\varepsilon} \end{aligned} \quad (7)$$

$$H(1 - \alpha) = -\alpha \bar{\omega}_L + \underline{\omega}_L + (1 - \alpha - \beta) \frac{-\alpha}{\varepsilon} = -\alpha \underline{\omega}_H + \underline{\omega}_L \quad (8)$$

To iron the distribution, we compute the tangent from $H(1 - \alpha)$ to $H(q)$, 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. (7) we have

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

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

$$\begin{aligned} \frac{H(1 - \alpha) - H(q^*)}{1 - \alpha - q^*} &= \frac{-\alpha \underline{\omega}_H - (F_{\varepsilon}(v^*) - 1)v^*}{1 - \alpha - F_{\varepsilon}(v^*)} \\ h(q^*) &= v^* - \frac{1 - F_{\varepsilon}(v^*)}{f_{\varepsilon}(v^*)} \end{aligned}$$

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

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

Hence, rearranging the terms,

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

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

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

This is equivalent to $\alpha\underline{\omega}_H \leq (1 - F(r^*))r^*$. If this fails, the optimal reserve r^* is above the ironed region, and so a second price auction is optimal.

Finally, observe that

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

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




Figure 4: **Relative Performance 1: Number of bidders.** This figure shows simulated expected revenues for different mechanisms as the number of bidders n varies, in an environment where F_L has mean 1 and standard deviation 0.5, the match probability is 0.05 and the match increment is 5.




Figure 5: **Relative Performance 2: Match Probability.** This figure shows simulated expected revenues for different mechanisms as the probability of a match α varies, where F_L has mean 1 and standard deviation 0.5, the number of bidders is 10 and the match increment is 5.




Figure 6: **Relative Performance 3: Match Increment.** We show simulated expected revenues for different mechanisms as the size of the gap between the low and high distributions Δ varies, where F_L has mean 1 and standard deviation 0.5, the match probability is 0.05 and the number of bidders is 10.

Products:	All	1	2	3	4	5	6	7	8	9	10
# Bids (000's)	2592	30	482	406	224	3	6	711	1	241	489
# Impressions (000's)	1146	19	223	256	103	1	4	215	1	108	217
# 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	-0.002	-0.004	0.013	0.020	-0.001	-0.009	-0.001	0.001

Table 1: 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. Summary statistics for each product are shown separately. Monetary units are cents.

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

Table 2: Expected revenue and welfare under different mechanisms, for the uniform environment with $\tau = 3$, $\alpha = 0.05$ and the number of bidders $n = 5$.

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 3: 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 $p = 3.8$.