

Online Ad Slotting With Cancellations

Florin Constantin^{*}

Jon Feldman[†]

S. Muthukrishnan[†]

Martin Pál[†]

ABSTRACT

Many advertisers (*bidders*) use Internet systems to buy advertisements on publishers' webpages or on traditional media such as radio, TV and newsprint. They seek a simple, online mechanism to *reserve* ad slots in advance. On the other hand, media publishers (*sellers*) represent a vast and varying inventory, and they too seek automatic, online mechanisms for pricing and allocating such reservations.

In this paper, we present and study a simple model for auctioning such ad slots in advance. A monopolist seller owns a set of slots that it will display at some point T in the future. Until T , bidders arrive sequentially and place a bid, demanding a reservation for one slot out of a set of desired slots. The seller must decide immediately whether or not to grant a reservation. Our model allows the seller to *cancel* at any time any reservation made earlier. If the seller cancels bidder i 's reservation, i incurs a utility loss amounting to a fraction of its value for the reservation.

Our main result is an online mechanism with many desirable properties. Winners have an incentive to be honest and bidding one's true value dominates any lower bid. Our mechanism's efficiency (respectively revenue) is within a constant fraction of the *a posteriori* optimally efficient solution (respectively VCG revenue). Our results make no assumptions about the arrival order or value distribution of bids. They still hold if the items for sale are elements of a matroid, a more general setting than slot allocation, or if bidders have linear (additive) value for a set of slots.

1. INTRODUCTION

Many advertisers now use Internet advertising systems. These take the form of advertisement (*ad*, henceforth) place-

^{*}SEAS, Harvard University, 33 Oxford St., Cambridge, MA 02138. Email: florin@eecs.harvard.edu. Most of this work was done while the author was visiting Google Research, New York.

[†]Google Research, 76 Ninth Ave., 4th Floor, New York, NY, 10011. Email: [jonfeld,muthu,mpal}@google.com](mailto:{jonfeld,muthu,mpal}@google.com).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

ments either in response to users' web search queries, or at predetermined slots on publishers' web pages. In addition, increasingly, advertisers use Internet systems that sell ad slots on behalf of offline publishers on TV, radio or newsprint. In sponsored search, and in some other cases, ad slots are typically sold via *real time* auctions, i.e., when a user poses a query or visits a web page, an auction is used to determine which ads will show and where they will be placed. On the other hand, traditionally, advertisers seek ad slots *in advance*, i.e. to reserve their slots. Product releases (such as movies, electronic gadgets, etc) and ad campaigns (e.g., creating and testing ads, budgets) are planned ahead of time and need to coordinate with future events that target suitable demographics. The advertisers then do not want to risk the vagrancies of real-time auctions and lose ad slots at critical events; they typically like a reasonable guarantee of ad slots at a specific time in the future.

Our motivation arises from systems that enable such advanced ad slotting. In particular, our focus is on automatic systems that have to manage ad slots in many different publishers' properties whose total supply may be massive. These properties differ wildly in their traffic, targeting, price and effectiveness. Not all publishers can estimate their inventory accurately: traffic to websites responds to time-dependent events, and sometimes webpages are generated dynamically so that even the availability of a slot in the future is not known *a priori*. Most web publishers are not able to estimate accurately a price for an ad slot, or provide sales agents to negotiate terms and would like automatic methods to price ad slots. Thus, what is desirable is a simple, automatic, online¹ market-based mechanism to enable advanced ad slotting over such varied, massive inventory.

Inspired by these considerations, we study the problem of mechanism design for advanced ad slotting. Our contribution is to propose a simple model, to design a suitable mechanism and to analyze its properties. In more detail, our contributions are as follows.

(i) We propose the following simple model for advanced ad slotting auctions. An auction starts at time 0; the seller has a set of heterogeneous *slots* for sale that will be published at time T . Bidder i arrives at some time $a_i < T$, having a private value $v(i)$ for exactly one slot out of a subset of slots $N(i)$. Upon his arrival, i places a bid $w(i)$ (which results in $N(i)$ becoming known to the seller) and requests an immediate response. Bidder i is either accepted or rejected; if

¹We use the word *online* as in *online algorithm*—i.e., the input arrives over time, and the algorithm makes sequential decisions—we do not mean “on the Internet.”

accepted, he may be removed (*bumped*) later by the seller. We assume that if bumped, i incurs a loss of $cv(i)$. At time T , each accepted bidder i that has not been bumped is published in one slot he was interested in (i.e. from $N(i)$).

This model lets the publisher accept a reservation at time t for a slot available at a later time T , and lets the advertiser get a reasonable guarantee. However, crucially, it lets the publisher *cancel* the reservation at a later time. Cancellation is necessary for publishers to take advantage of a spike in demand and rising prices for an item and not be forced to sell the slot below the market because of an *a priori* contract. In addition, in a pragmatic sense, cancellation is crucial: for example, a website might overestimate its inventory for a later date and accept ads, but as time progresses, its estimates may become smaller, and the publisher will not be able to honor all the accepted ads from the past. Finally, cancellations are very much part of the business with advance bookings, both within advertising and beyond. At the same time, it comes at a cost, which is the bumped bidders' utility loss. The publisher benefits from the reduction in uncertainty and can compensate bumped advertisers in various ways. We present our model formally in Section 3.

(ii) We present an efficiently implementable mechanism $M_\alpha(\gamma)$ for determining who is accepted, who is bumped and also the prices and bump payments. The parameter γ represents how much higher a new bid has to be in order to bump an older bid. A bumped bidder will be paid an α fraction of their bid, making up for their utility loss due to being bumped.

(iii) We show a number of important strategic as well as efficiency- and revenue-related properties of $M_\alpha(\gamma)$:

- $M_\alpha(\gamma)$ is individually rational and winners have an incentive to bid truthfully while losers should bid at least their true value.
- With respect to the bids received, the efficiency (value of assignment) of $M_\alpha(\gamma)$ is at least a constant factor (depending on γ and α) of the offline optimum. Under mild player rationality assumptions, we show that our mechanism is competitive with respect to the optimum offline efficiency on bidders' true *values*.
- We also consider the notion of *effective efficiency* which interprets social welfare as the sum of the winners' bids minus bumped (if any) bidders' losses. We show that for suitable $\gamma(\alpha)$, our mechanism's effective efficiency matches a numerically obtained upper bound on the effective efficiency of *any* deterministic algorithm.
- The revenue of $M_\alpha(\gamma)$ is at least a constant factor (dependent on γ and α) of that of the Vickrey-Clarke-Groves (VCG) mechanism on all received bids.
- We also study *speculators*, that is, ones who have no interest in the items for sale but who participate in order to earn the bump payment. We show several game theoretic properties about the behavior of the speculators, including bounding their overall profit.

To the best of our knowledge our results are the first about mechanisms with strong game-theoretic properties for advanced placement of ads (more generally, online bipartite matching) with a costly cancellation feature. We make no assumptions on the arrival order of the bidders or on their

values. Prior work has studied advanced sale of goods without cancellations, but only under a probabilistic distribution of bidders' values [7]. Under a worst case model like ours, no nontrivial results are possible without making additional assumptions; in our case, we overcome these impossibilities by allowing cancellations. In secretary problems [2], bids may be arbitrary but their order is assumed to be uniformly random (cannot be specified by an adversary).

There are specific examples of systems that implement advanced booking with cancellations. For example, this is common in the airline industry, where tickets may be booked ahead of time, and customers may be bumped later for a payment. In the airline case, the inventory is mostly fixed, sophisticated models are used to calculate prices over time, and often negotiations are involved in establishing the payment for bumping, just prior to time T . In some cases, the bump payments may even be larger than the original bid (price) of the customer. Likewise, in offline media such as TV or Radio, advanced prices are negotiated by humans, and often if the publisher does not respect the reservation due to inventory crunch, a payment is *a posteriori* arranged including possibly a better ad slot in the future. These methods are not immediately applicable to the auction-driven automatic settings like ours.

From a technical point of view, one can view our model as an online weighted bipartite matching problem (or more generally, an online maximum weighted independent set problem in a matroid). On one side we have slots known ahead of time. The other side comprises advertisers whose bids (weighted nodes) arrive online. Our goal is to find a "good" weighted matching in the eventual graph. Each time an advertiser appears we need to decide if we should retain it or discard it; retaining it may lead to discarding a previously retained bidder. Our mechanism builds on such an online matching algorithm [10] to determine a suitable bump payment and prices. It is, curiously, able to make use of such an online algorithm previously proposed in the *semi-streaming* model in the theoretical computer science literature.

All our results extend to a setting where the items for sale are elements of a matroid, a more general setting than slot allocation. A bidder bids on exactly one element of the matroid, which is known ahead of time to the seller and may vary across bidders. A set S of bidders is then feasible if the set containing each bidder's element forms an independent set of the matroid. In the bipartite matching setting, the seller's matroid contains sufficiently many copies of one element for each subset of slots. A set of bidders (elements) is independent if the bidders can be matched to slots such that each one receives a slot from its subset. In a different direction, our results also extend to a (strictly bipartite matching) setting where a bidder's value for a set S of slots is the sum of values for each slot in S (no bidder can express substitute slots). In this setting, the multi-slot matching problem is actually a collection of single-slot matching problems since bids on two different slots can never interact. We prefer the bipartite matching setup where a bidder is interested in one out of a set of substitute slots for clarity of exposition.

We have initiated the study of mechanisms for advanced reservations with cancellations. A number of technical problems remain open, both within our model, as well as in its extensions, which we describe later for future study.

2. RELATED WORK

Babaioff et al. [2] address the *matroid secretary problem*: finding a competitive assignment when weighted elements of a matroid arrive online and no cancellations are allowed. As is common in secretary problems, while not making any assumption on bidder valuations, they assume that all orders of arrivals are equally likely. They present a $\log r$ -competitive algorithm for general matroids where r is the rank of the matroid (the size of the largest independent sets) and a $4d$ -competitive algorithm for our setting without cancellations (transversal matroids) but where each bidder can only be interested in at most d items. Both these algorithms observe half of the input and then set a threshold price: per item in the transversal case and uniform in the general case. The $4d$ -competitive algorithm ensures truthful bidding even when the items desired by an agent are private information. Dimitrov and Plaxton [5] extend the $4d$ -algorithm and provide an algorithm with a constant competitive ratio for any transversal matroid.

Bikhchandani et al. [3] present an ascending auction for selling elements of a matroid that ends with an optimal allocation (i.e. the auction is efficient). Truthful bidding is an equilibrium of the auction. They assume however that bidders are present throughout the auction.

Cary et al. [4] show that a random sampling profit extraction mechanism approximates a VCG-based target profit in a procurement setting on a matroid.

Gallien and Gupta [7] analyze players' strategies regarding buyout prices in online auctions. In their model, bidders' valuations are drawn from a known distribution and their utilities are time-discounted; furthermore there are no cancellations and arrivals are assumed to follow a Poisson process. They exhibit symmetric Bayes-Nash equilibria in which bidders follow threshold strategies, which can in turn be used by the seller (as they show) to optimize revenue.

Lavi and Nisan [9] consider online auctions for identical goods. In their model, bidders' values are arbitrary from the interval $[\rho, \bar{\rho}]$ and no cancellations are possible. They present a simple online posted-price auction based on exponential scaling. This auction is optimal among online auctions and achieves a $\Theta(\log(\bar{\rho}/\rho))$ approximation with respect to both efficiency and the VCG revenue.

Feige et al. [6] study an *offline* weighted bipartite matching problem where the seller can partially satisfy a bidder's request at the cost of paying a proportional penalty. Accepting bidder i but not providing any items results in the seller losing utility proportionally to i 's bid, similar to our definition of *effective efficiency* in Section 4.1. Recall however that in our model it is i and not the seller who incurs the utility loss. They show that it is NP-hard to approximate the optimal solution within any constant factor. They propose an adaptive greedy algorithm assigning one bidder at a time (but that inspects all unassigned bidders in deciding which bidder to allocate) and that may reassign bidders. They provide a lower bound on this algorithm's efficiency with respect to the optimal assignment.

Independently and concurrently, Babaioff et al. [1] study the same problem as this paper, but from an algorithmic perspective only, leaving incentives and revenue considerations aside. Their paper and ours present the same algorithm and efficiency results. Their focus is on effective efficiency, for which they analytically prove an upper bound on any deterministic algorithm's competitive ratio (we only present re-

sults of a numerical simulation in Fig. 1 strongly suggesting this bound). Unlike us, they go on to study costly cancellations ("buyback") in knapsack problems. They provide an algorithm similar to $M_\alpha(\gamma)$ and prove a bi-criterion approximation result, an informative bound since they also prove that no deterministic algorithm has a constant competitive ratio.

3. MODEL AND MECHANISM

We first define our model and present our mechanism which we will study in later sections.

3.1 Model Basics

There is a seller who has a finite set of *slots* (items), and starts sale at time 0 and ends sale at time T .

Each bidder i is interested in *exactly* one slot out of a set $N(i)$ called i 's *choice set*. We denote by $v(i)$ bidder i 's value for any slot in $N(i)$ and we assume that $v(i)$ is private information to i . Each bidder i places a bid $w(i)$ (that may be different than $v(i)$) as soon as it arrives, at time a_i . As a consequence of i bidding, i 's choice set becomes known to the seller. When i bids, he may be *accepted* (i.e. promised an item from $N(i)$), else rejected. If promised an allocation, he may get *bumped* later, losing the reservation. Any accepted bidder who is not bumped before time T is allocated. We model bidder i 's *utility* as quasilinear in money:

$$\lambda \cdot v(i) - x(i), \quad \text{where} \quad (1)$$

- λ equals 0 if i is rejected, 1 if i is accepted and granted an item from $N(i)$ and $-\alpha$ if i is accepted but bumped, where $0 \leq \alpha < 1$.
- $x(i)$ is i 's *transfer* to the seller (price). It is 0 if i is rejected, and some non-negative amount if i is accepted and allocated. $x(i)$ may be negative (e.g. the *bump payment* the seller makes in $M_\alpha(\gamma)$).

That is, a bidder is unaffected if rejected right away, has a value of $v(i)$ for being allocated, and incurs a loss amounting to an α fraction of its value if bumped.

From an algorithmic perspective, our problem is finding online a bipartite matching of "good" weight, when costly cancellations of acceptances are allowed. There is a bipartite graph with items (respectively bidders) as "right"- (respectively "left"-)hand side vertices. Initially, left-vertices are revealed while right-vertices are hidden. One by one, a right-vertex i is revealed together with its weight $w(i)$ and its edges (i.e. the set $N(i)$). A decision whether to accept it or not must be taken immediately, although i 's acceptance decision may be canceled at an $\alpha w(i)$ cost later.

Note that if the input was not manipulable by the bidders, given that cancellations are not detrimental to the seller, the seller could simply accept a new bidder whenever it improves the current matching. It is well-known that the resulting assignment would be optimal. If $\alpha = 0$, then by using the same algorithm and charging each bidder its VCG (see Section 5) price, truthfulness (a bid equal to one's true value) becomes a dominant strategy. However, in our setting bidders are self-interested and cancellations hurt them. Thus they may alter the input to the algorithm (their bids) if it is in their interest to do so. Therefore we aim for a mechanism that is competitive while bounding the bidders' manipulations.

3.2 Our $M_\alpha(\gamma)$ Mechanism

We present our advance-booking online mechanism $M_\alpha(\gamma)$ (allocation algorithm and payments). The allocation algorithm follows the *Find-Weighted-Matching* algorithm in [10]², that uses an unconstrained *improvement factor* $\gamma > 0$. We require $\alpha < \frac{\gamma}{1+\gamma}$ i.e. $\gamma \in (\frac{\alpha}{1-\alpha}, \infty)$ (recall that $0 < \alpha < 1$).

Denote the number of bidders by n ; our mechanism is independent of n . By relabeling bidders, assume that they bid in order $1, 2, \dots, n$; time is indexed likewise.

DEFINITION 1. *We say that a set of bidders B can be matched if for each $b \in B$ there exists an item $i_b \in N(b)$ such that $i_b \neq i_{b'}, \forall b \neq b'$. We say that B is a perfect matching if it can be matched and no item is left unmatched (B 's cardinality must equal the number of items).*

Algorithm 1 lists $M_\alpha(\gamma)$'s pseudocode. At a high level, $M_\alpha(\gamma)$ maintains a set of accepted bidders that form a perfect matching. For any new arriving bidder i bidding $w(i)$, the algorithm looks if there exists some bidder j in the accepted set with $w(j) < \frac{w(i)}{1+\gamma}$ such that i can be swapped in the matching if j is swapped out. If so, accept i and cancel the reservation of (*bump*) j^* , the lowest weight such j . Bidder j^* is paid the bump payment $\alpha w(j^*)$. Indeed, j^* makes no payment at all and in fact *gets money for free* from the seller for being bumped. An accepted bidder who is not bumped by time T is necessarily allocated a slot from his choice set and pays the seller an amount we define later (Eq.(2)).

At time 0, A_0 is an arbitrary matching; we introduce r dummy bidders (each bidding 0) whose choice set is the whole set of items, arriving before all actual bidders. This will not affect our arguments below.³At time t , we call currently accepted bidders *alive*, and denote the set of alive bidders as A_t . Let $X_t = \{b \in A_{t-1} : A_{t-1} \cup \{t\} \setminus \{b\} \text{ can be matched}\}$; X_t is the set of alive bidders at $t-1$ that can be exchanged for t . Bidders still alive at the end (time T) are called *survivors* and $S = S(\mathbf{w})$ denotes this set. We denote the set of bumped bidders by $R = R(\mathbf{w})$.

Algorithm 1 Mechanism $M_\alpha(\gamma)$ for online ad slotting with cancellations.

```

Fix  $A_0$ , an arbitrary perfect matching on dummy bidders.
for each  $i \geq 1$  bidding  $w(i)$  (with choice set  $N(i)$ ) do
  Let  $X_i = \{j < i : A_{i-1} \cup \{i\} \setminus \{j\} \text{ can be matched}\}$ .
  Let  $j^* = \operatorname{argmin}_{j \in X_i} w(j)$ 
  if  $(1 + \gamma)w(j^*) < w(i)$  then
     $A_i = A_{i-1} \cup \{i\} \setminus \{j^*\}$ :  $i$  is accepted,  $j^*$  is bumped
     $j^*$  is paid  $\alpha w(j^*)$ , i.e. an  $\alpha$  fraction of its bid
  end if
end for
Charge any bidder  $i$  in  $S \triangleq A_n$  (i.e. survivor) as in Eq. (2).

```

DEFINITION 2. *Let i be a bidder and fix the bids of all other bidders. Let $w^{\text{ac}}(i)$ (i 's acceptance weight) be the infimum of all bids that i can make such that i is accepted*

²Unlike in [10], a bidder i 's value is the same for any slot (vertices as opposed to edges are weighted). The slot currently assigned to i may vary across stages in $M_\alpha(\gamma)$.

³When bidder t arrives, assume $A_{t-1} = A \cup D$ where D only contains dummy bidders and there exists a matching I_t of $A \cup \{t\}$ which matches t to some item i_t . By reassigning dummy bidders, we can assume that actual bidders are matched according to I_t . Then bidder t can bump at least the dummy bidder $d \in D$ that is matched to i_t in A_{t-1} .

given its arrival a_i and N_i . Similarly, let $w^{\text{sv}}(i)$ (i 's survival weight) be the infimum of all bids that i can make such that i is accepted and survives until time T (the end) given a_i and N_i . Clearly, $w^{\text{ac}}(i) \leq w^{\text{sv}}(i)$. Let $W^{\text{sv}} = \sum_{i \in S} w^{\text{sv}}(i)$.

Note that $w^{\text{sv}}(i)$ always exists since it suffices to bid $(1 + \gamma) \max_{j \neq i} w(j)$. Also, $w^{\text{ac}}(i)$ and $w^{\text{sv}}(i)$ are independent of i 's actual bid, but may depend on a_i and on the other bidders' bids or arrivals. $w^{\text{ac}}(i)$ can be computed by the seller at a_i whereas $w^{\text{sv}}(i)$ may depend on future bidders and thus can only be computed at time T .

In summary, i is $\begin{cases} \text{rejected,} & \text{if } w(i) < w^{\text{ac}}(i) \\ \text{bumped,} & \text{if } w^{\text{ac}}(i) \leq w(i) < w^{\text{sv}}(i) \\ \text{a survivor,} & \text{if } w^{\text{sv}}(i) \leq w(i) \end{cases}$

If i is a survivor, i 's price p_i is as follows:

$$p_i = \begin{cases} w^{\text{sv}}(i)(1 - \alpha) & \text{if } w^{\text{ac}}(i) < w^{\text{sv}}(i). \\ w^{\text{sv}}(i) & \text{if } w^{\text{ac}}(i) = w^{\text{sv}}(i). \end{cases} \quad (2)$$

The common case is when $w^{\text{ac}}(i) < w^{\text{sv}}(i)$: i gets a discount amounting to the highest refund it could have otherwise obtained: $\alpha w^{\text{sv}}(i)$. The special case of $w^{\text{ac}}(i) = w^{\text{sv}}(i)$ occurs when i 's acceptance is enough for its survival (in particular if i is the last bidder). When $w^{\text{ac}}(i) = w^{\text{sv}}(i)$, from the bidder's point of view, $M_\alpha(\gamma)$ posts a price of $w^{\text{sv}}(i)$.

This concludes the definition of $M_\alpha(\gamma)$. We first present an example and will then proceed to studying its properties.

Example. Let $\gamma = 0.5$ and an $\alpha < \frac{0.5}{0.5+1}$. Consider two items I_a, I_b and the following sequence of bidders: B_1 bidding 6 on any of I_a and I_b , B_2 bidding 4.4 on I_b , B_3 bidding 10 on I_a and B_4 bidding 7.5 on I_b . Bidder i arrives at time t_i and $T = t_4$. $M_\alpha(\gamma)$ accepts B_1 at t_1 , accepts B_2 at t_2 , accepts B_3 and bumps B_2 at t_3 and then rejects B_4 at t_4 . We write $w^{\text{ac}}(i)$ and $w^{\text{sv}}(i)$ for $w^{\text{ac}}(B_i)$ and $w^{\text{sv}}(B_i)$. We have $w^{\text{ac}}(1) = 0, w^{\text{ac}}(2) = 0, w^{\text{ac}}(3) = 1.5 \cdot 4.4 = 6.6$ (to bump B_2) and $w^{\text{ac}}(4) = 1.5 \cdot 6 = 9$ (to bump B_1). $w^{\text{sv}}(1) = \frac{7.5}{1.5} = 5$ (to prevent being bumped by B_4), $w^{\text{sv}}(2) = 6$ (to prevent being bumped by B_3 and B_4), $w^{\text{sv}}(3) = 6.6, w^{\text{sv}}(4) = w^{\text{ac}}(4) = 9$. B_1 and B_3 survive: B_1 pays $(1 - \alpha)w^{\text{sv}}(1)$ since $w^{\text{ac}}(1) < w^{\text{sv}}(1)$ and B_3 pays $w^{\text{sv}}(3)$ since $w^{\text{ac}}(3) = w^{\text{sv}}(3)$. Note that if a bidder B_5 were to arrive after B_4 bidding 10.5 on I_a then only $w^{\text{sv}}(3)$ would change to $\frac{10.5}{1.5} = 7$. In this case, B_3 's price becomes $(1 - \alpha) \cdot 7$ which may be lower than 6.6. Unless a bidder i 's $w^{\text{sv}}(i)$ price coefficient goes from 1 to $1 - \alpha$, i 's price cannot go down if new bidders arrive.

3.3 Basic Incentive Properties

The following theorem summarizes our mechanism's favorable incentive properties.

THEOREM 1. *The $M_\alpha(\gamma)$ mechanism is individually rational. Bidding one's true value (weakly) dominates any lower bid. If honest, any survivor is (weakly) best-responding.*

The proof follows from the two lemmas below.

Note that with bump payments ("money for nothing"), anyone with zero value for any slot can attempt to exploit the mechanism and get a bump payment by making an appropriate (untruthful) bid. It is still possible that other types of truthful competitive allocation mechanisms exist.

LEMMA 1. *Bidding less than one's true value is dominated by bidding one's true value. A survivor's (weakly) best-response is to be honest. A bumped bidder's best response may however be to bid more than its true value.*

PROOF. If $w^{\text{ac}}(i) < w^{\text{sv}}(i)$, bidder i 's highest possible bump payment is $\alpha w^{\text{sv}}(i)$. The price of $(1 - \alpha)w^{\text{sv}}(i)$ has been chosen such that i prefers winning to being paid $\alpha w^{\text{sv}}(i)$ if and only if $v(i) \geq w^{\text{sv}}(i)$. That is, the best bid i can make (i 's *best-response*) is to bid just below $w^{\text{sv}}(i)$ if $v(i) < w^{\text{sv}}(i)$ and to bid its true value otherwise.

If $w^{\text{ac}}(i) = w^{\text{sv}}(i)$, then i can never get a bump payment and i simply faces a take-it-or-leave-it offer of $w^{\text{sv}}(i)$. \square

$M_\alpha(\gamma)$ is individually rational:

LEMMA 2. *If bidder i bids its true value $v(i)$, then i 's utility after participating in the mechanism is non-negative.*

PROOF. When surviving, i pays at most $w^{\text{sv}}(i) \leq v(i)$. If i is not accepted then i 's utility is 0. If i is accepted and then bumped, i 's utility is $-\alpha v(i) + \alpha v(i) = 0$. \square

In the following two sections we show that apart from favorable incentive properties, $M_\alpha(\gamma)$ is also competitive with respect to revenue and efficiency.

4. EFFICIENCY OF $M_\alpha(\gamma)$

DEFINITION 3. *For any vector $\mathbf{w}' = (w'(1), \dots, w'(n))$ of bids, we let $\text{OPT}[\mathbf{w}']$ be the optimal matching, or its weight when there is no confusion (recall that the seller is revealed any bidder's choice set). If B is a set of bidders, we will denote $w'(B) = \sum_{b \in B} w'(b)$. Unless specified otherwise, $\mathbf{w} = (w(1), \dots, w(n))$ denotes the input bids and \mathbf{v} denotes the true values. We will denote by $\text{OPT} = \text{OPT}[\mathbf{w}]$.*

From an algorithmic perspective, our mechanism is a $1 + \gamma$ approximation to the optimum assignment (Lemma 4). However, if incentives are not aligned then bidders may want to significantly alter the input to the algorithm (their bids). The allocation may then be a poor choice considering bidders' true values. We show that this is not the case, in Theorem 2, our main result regarding $M_\alpha(\gamma)$'s efficiency: the assignment output by our mechanism is a constant factor (depending on α and γ) approximation to the offline optimum on bidders' true values if bids are "reasonable".

THEOREM 2. *Let \mathbf{w} be a set of bids such that each bidder bids at least its true value, that is $w(i) \geq v(i) \forall i$, and the sum of all bidders' utilities is non-negative. When run on \mathbf{w} , $M_\alpha(\gamma)$'s efficiency with respect to the true values \mathbf{v} is:*

$$\sum_{i \in S(\mathbf{w})} v(i) \geq \frac{1 - \alpha - \frac{\alpha}{\gamma}}{(2 - \alpha - \frac{\alpha}{\gamma})(1 + \gamma)} \cdot \text{OPT}[\mathbf{v}].$$

Note that if all bidders are truthful then the right-hand side constant can be increased to $\frac{1}{1 + \gamma}$ (see Lemma 4 below).

Recall the impossibility of making truthfulness a dominant strategy due to the use of bump payments. Theorem 2's assumption allows for some bidders to have negative utility and therefore fail to best-respond (Lemma 2 shows that simply by being truthful, a bidder's utility is at least 0) as long as overall, gains outweigh losses in utility. The assumption fails when, for example, bidders with value 0 for any slot (the "speculators" of Section 6) grossly overestimate the actual bids and end up being allocated and having to pay due to bidding too high. In such a scenario, the true value of the allocation may be small, possibly 0. Section 6 provides a detailed discussion on how speculators affect incentives.

We now proceed to proving Theorem 2, establishing a few other important results along the way. The following bid vector will prove useful:

$$\tilde{\mathbf{w}}_S(i) = \begin{cases} w^{\text{sv}}(i), & \text{if } i \in S \\ w(i)/(1 + \gamma), & \text{if } i \notin S \end{cases}$$

Note that if $i \notin S$ then $w^{\text{sv}}(i) > w(i) > w(i)/(1 + \gamma)$.

Lemma 3 is our core technical result, being used for both efficiency and revenue claims. It states that the set of survivors given bids \mathbf{w} is the optimal offline assignment on $\tilde{\mathbf{w}}_S$. Its proof is more involved and we defer it to the Appendix.

LEMMA 3. $S(\mathbf{w}) = \text{OPT}[\tilde{\mathbf{w}}_S]$.

The following result provides an upper bound on the sum of bumped bidders' bids and shows that S is a $1 + \gamma$ approximation to the optimal offline matching given the same set of bids. Recall that R is the set of bumped bidders.

LEMMA 4. *We have $w(R) \leq W^{\text{sv}}/\gamma \leq w(S)/\gamma$. Also, $\text{OPT} \leq (1 + \gamma)w(S)$.*

PROOF. For each $s \in S$, let $d_1, \dots, d_J = s$ be a chain such that: d_{j+1} bumps d_j , $\forall 1 \leq j \leq J - 1$. To simplify notation, assume $d_1 = 1, \dots, d_{J-1} = J - 1$. We will show that

$$\sum_{j=1}^{J-1} w(j) \leq w^{\text{sv}}(s)/\gamma$$

The claim will follow since the set of bidders is the disjoint union of survivors' chains.

Since s bumped $J - 1$, we have $w_{J-1} \leq \frac{w^{\text{sv}}(s)}{1 + \gamma}$. Since $j + 1$ bumps j , $\forall 1 \leq j \leq J - 2$, $w_j \leq \frac{w_{j+1}}{1 + \gamma}$. Thus by induction, $w_j \leq w^{\text{sv}}(s)(1 + \gamma)^{j-J}$, $\forall 1 \leq j \leq J - 1$. We get

$$\sum_{j=1}^{J-1} w_j \leq w^{\text{sv}}(s) \sum_{j=1}^{J-1} (1 + \gamma)^{j-J} \leq w^{\text{sv}}(s)/\gamma$$

We have $w(S) \geq \tilde{\mathbf{w}}_S(S) = \text{OPT}[\tilde{\mathbf{w}}_S] \geq \text{OPT}[\mathbf{w}]/(1 + \gamma)$. Each inequality is implied by the fact that no bidder's contribution decreases when going from the left hand side to the right hand side. The equality follows from Lemma 3: S is an optimal assignment for $\tilde{\mathbf{w}}_S$, i.e. $\tilde{\mathbf{w}}_S(S) = \text{OPT}[\tilde{\mathbf{w}}_S]$. \square

An analogous lemma can be found in [10]. Our constants are tighter because in our model, a bidder's value for any slot is the same, and all edges incident to a bidder arrive simultaneously. These bounds are almost tight:

EXAMPLE 1. *Consider $k + 2$ truthful bidders competing on one item; bidder i is the i -th to arrive and has value $(1 + \gamma)^{i-1}$ unless $i = k + 2$, whose value is $(1 + \gamma)^{k+1} - \varepsilon$. Bidder $i + 1$ bumps i , $\forall 1 \leq i \leq k$. Only the $k + 1$ -st bidder survives. The bumped bidders have total weight $\sum_{i=0}^{k-1} (1 + \gamma)^i = ((1 + \gamma)^k - 1)/\gamma$. OPT is $(1 + \gamma)^{k+1} - \varepsilon$.*

PROOF OF THEOREM 2. For any $s \in S$, let $p(s)$ be s 's payment. By assumption, bidders' total utility is non-negative:

$$\sum_{s \in S} (v(s) - p(s)) + \sum_{r \in R} (\alpha w(r) - \alpha v(r)) \geq 0 \text{ or}$$

$$\sum_{s \in S} (v(s) - (1 - \alpha)w^{\text{sv}}(s)) + \sum_{r \in R} \alpha w(r) \geq 0$$

since $p(s) \geq (1 - \alpha)w^{\text{sv}}(s)$ (see Eq. (2)). We have from Lemma 4 that $\sum_{r \in R} w(r) \leq \sum_{s \in S} w^{\text{sv}}(s)/\gamma$. We prove that

$$\sum_{s \in S} (v(s) + w^{\text{sv}}(s)) \geq \text{OPT}[\mathbf{v}]/(1 + \gamma).$$

The theorem then follows by algebraic manipulation.

$$\text{Let } w'(x) := \begin{cases} \max(v(x), w^{\text{sv}}(x)), & \text{if } x \in S \\ w(x), & \text{if } x \notin S \end{cases}$$

Clearly, $v(s) + w^{\text{sv}}(s) \geq w'(s) \forall s \in S$. $S(\mathbf{w}) = S(\mathbf{w}')$ since only survivors in $S(\mathbf{w})$ change their bid, still bidding above their survival thresholds. By Lemma 4, $\sum_{s \in S} w'(s) \geq \text{OPT}[\mathbf{w}']/(1 + \gamma)$. The claim follows by noting that $\text{OPT}[\mathbf{w}'] \geq \text{OPT}[\mathbf{v}]$ since $w'(x) = w(x) \geq v(x), \forall x \notin S$. \square

4.1 Effective Efficiency

Efficiency is usually measured as the sum of winning bidders' values. An alternative definition which also takes into account the losses of bumped bidders and may be more appropriate when cancellations are allowed is the following:

DEFINITION 4. *Let w be a sequence of bids and A an online allocation algorithm, possibly with cancellations. Let $S_A(w)$ (resp. $R_A(w)$) be the set of winners (resp. bumped bidders) when A is run on w . We define the effective efficiency of A on w as*

$$u_A(w) = \sum_{s \in S_A(w)} w(s) - \alpha \sum_{r \in R_A(w)} w(r)$$

A 's effective efficiency competitive ratio is $\inf_w \frac{u(w)}{\text{OPT}[w]}$

We present an upper bound (obtained numerically) on the effective efficiency competitive ratio of *any* deterministic algorithm. For $\alpha < \frac{\sqrt{5}-1}{2} \simeq 0.618$ and a certain γ_α , $M_\alpha(\gamma_\alpha)$ matches this upper bound.

For a fixed α , let $n \geq 2$ a positive integer and $c \in (0, 1)$: we aim for n bidders and a competitive ratio of c . Consider one item and a sequence of bids $\{a_k(c)\}_{1 \leq k \leq n}$ on it (bidder k bids a_k) such that $a_1 = 1, a_2 = \frac{1}{c} > 1$ and

$$ca_{k+1}(c) = a_k(c) - \alpha \sum_{j=1}^{k-1} a_j(c) \forall k \geq 2 \text{ implying} \\ ca_{k+1} = (1 + c)a_k - (1 + \alpha)a_{k-1} \forall k \geq 2 \quad (3)$$

We will look for a $c = c_n$ such that

$$a_n(c) - \alpha \sum_{j=1}^{n-1} a_j(c) = ca_n(c) \iff a_n = (1 + \alpha)a_{n-1} \quad (4)$$

E.g. $c_2 = \frac{1}{1 + \alpha} > c_3 = \frac{1}{1 + 2\alpha} > c_4 = \frac{2}{1 + 3\alpha + \sqrt{(1 + 5\alpha)(1 + \alpha)}}$.

Unfortunately, c_n does not have a nice closed form for $n \geq 4$ (in addition, c_n may be not be unique - the smallest $c_n \in [0, 1]$ is then of interest). Furthermore, for certain c and n no such sequence may exist.

THEOREM 3. *Fix n and α . Let c_n be the lowest number (if any) in $[0, 1]$ for which Eqs. (3) and (4) simultaneously hold. Then no deterministic algorithm can have an effective efficiency competitive ratio higher than c_n .*

PROOF. On any input, the offline optimum with respect to effective efficiency is simply the highest weight assignment, and it results in bumping no bidders.

Suppose towards a contradiction that there was a deterministic algorithm A with a competitive ratio $c' > c_n$. Assume that the bids that arrive are a_1, \dots, a_{k_0} for some $1 \leq k_0 \leq n$. Then at each k , the algorithm A must accept a_k , or its competitive ratio will be smaller than c_n when $k = k_0$. This is clear for $k = 1$. Fix $k \in [2, n - 1]$. Let M_k be the highest (i.e. the offline optimum) of a_1, \dots, a_k . If A does not accept k then the competitive ratio on input a_1, \dots, a_k will be at most

$$\frac{a_{k-1}(c_n) - \alpha \sum_{j=1}^{k-2} a_j(c_n)}{M_k(c_n)} = \frac{c_n a_k(c_n)}{M_k(c_n)} \leq c_n$$

where the equality follows from Eq. (3). Now we claim that *whether or not* A accepts a_n , its competitive ratio will be at most c_n . If a_n is accepted, $\alpha \sum_{j=1}^{n-1} a_j$ has been lost due to bumping bidders $1, \dots, n - 1$; if a_n is rejected the effective efficiency is $a_{n-1} - \alpha \sum_{j=1}^{n-2} a_j$. By Eqs. (3) and (4), both quantities are a c_n fraction of a_n , which in turn is at most M_n , the optimal (effective) efficiency. \square

Figure 1 strongly suggests that the competitive ratio of any algorithm cannot be higher than $2\alpha + 1 - 2\alpha^{0.5}(\alpha + 1)^{0.5}$, shown as squares in the figure. Note that for this c the characteristic equation of Eq. (3) has a double root.

The triangles plot the minimum c found for the corresponding α for different values of n (we used Fibonacci values up to rank 12, i.e. largest n was 144). The c values were found via binary search. It was true in general, although not always, that the higher n , the lower c_n . We suspect that one can always find an increasing sequence of integers $\{n_i\}_{i \geq 1}$ such that a solution c_{n_i} to Eqs. (3) and (4) converges from above to $2\alpha + 1 - 2\alpha^{0.5}(\alpha + 1)^{0.5}$ as $i \rightarrow \infty$.

Lemma 4 implies that for our algorithm $M_\alpha(\gamma)$

$$u = w(S) - \alpha w(R) \geq w(S) - \alpha w(S)/\gamma \geq \frac{\text{OPT}}{1 + \gamma} \left(1 - \frac{\alpha}{\gamma}\right).$$

Let $\underline{u}(\gamma) = \frac{1}{1 + \gamma} \left(1 - \frac{\alpha}{\gamma}\right)$. Subject to the constraint $\alpha \leq \frac{\gamma}{\gamma + 1}$, $\underline{u}(\gamma)$ is maximized for $\gamma_0 = \max\{\alpha + \sqrt{\alpha^2 + \alpha}, \frac{\alpha}{1 - \alpha}\}$. $\underline{u}(\gamma_0)$ is displayed in Fig. 1 by circles. The value 0.618 (the golden ratio) is where $\frac{\alpha}{1 - \alpha}$ becomes higher than $\alpha + \sqrt{\alpha^2 + \alpha}$. If $\alpha < 0.618$, $\underline{u}(\gamma_0) = 2\alpha + 1 - 2\alpha^{0.5}(\alpha + 1)^{0.5}$, which matches the numerical upper bound. Recall that this is just a worst-case lower bound on the effective efficiency, but likewise, so are the upper bounds. The top curve plots $c_3 = 1/(1 + 2\alpha)$.

Recall that when $\alpha = 0$ all bidders can be tentatively accepted (by letting $\gamma = 0$) since they incur no loss and do not have to be refunded. Then, the optimal matching can be found via a one-shot (offline) algorithm at time T .

5. REVENUE OF $M_\alpha(\gamma)$

We show that apart from favorable incentive and efficiency properties, $M_\alpha(\gamma)$ is also competitive with respect to revenue.

As a revenue benchmark, we consider the offline Vickrey-Clarke-Groves (VCG) mechanism because it generates the highest revenue among truthful efficient allocation mechanisms [8]. We show that our mechanism is competitive with respect to revenue with VCG on bidders' true values.

Let \mathbf{w}' be a sequence of bids - when defining VCG on \mathbf{w}' we will assume that all bids are received at once by VCG. Let \mathbf{w}'_{-i} denote the set of all bids in \mathbf{w}' except bidder i 's. The VCG mechanism implements an efficient allocation and

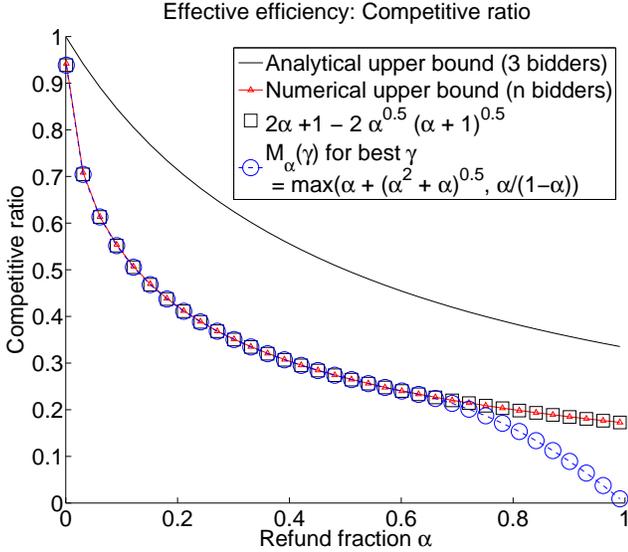


Figure 1: Effective efficiency competitive ratio (EECR) bounds as a function of α . The top curve is $c_3 = 1/(1 + 2\alpha)$. The middle curve is a numerical upper bound ($c = c_n$ of Eq. (4)) on any deterministic algorithm’s EECR. The bottom curve shows (a lower bound on) our algorithm’s EECR for the best γ_α : it matches the upper bound for $\alpha < 0.618$. Our choice of γ is constrained by $\alpha < \gamma/(\gamma + 1)$; if it were not, the bounds would match for all α .

thus the matching it outputs is optimal. If $i \in \text{OPT}[\mathbf{w}']$ then VCG charges bidder i its externality on the other bidders:

$$\sum_{k \in \text{OPT}[\mathbf{w}'_{-i}]} \mathbf{w}'(k) - \sum_{j \neq i, j \in \text{OPT}[\mathbf{w}']} \mathbf{w}'(j) \quad (5)$$

We will use the following known (see e.g. [4], Fact 3.2) combinatorial property of our setting: $\forall i \neq x$, if $x \in \text{OPT}[\mathbf{w}']$ then $x \in \text{OPT}[\mathbf{w}'_{-i}]$.

LEMMA 5. *A winning bidder’s VCG payment is a losing bid. Also, the VCG revenue can only increase if some bids in \mathbf{w}' are increased.*

On bids \mathbf{w}' we denote the VCG revenue by $\text{REV}_{\text{vcg}}[\mathbf{w}']$ and the net revenue of $M_\alpha(\gamma)$ (payments from survivors minus bump payments) by $\text{REV}_{\gamma, \alpha}[\mathbf{w}']$.

THEOREM 4. *Assume $w(i) \geq v(i) \forall i$, i.e. no one bids below their true value, since that would be dominated. Then*

$$\text{REV}_{\gamma, \alpha}[\mathbf{w}] \geq \frac{1 - \alpha - \alpha/\gamma}{1 + \gamma} \text{REV}_{\text{vcg}}[\mathbf{w}]$$

This theorem shows the tradeoff between γ , the improvement factor required for bumping an accepted bidder and α , the fraction returned as the bump payment. For instance, for $\alpha = 0.25$ (refund of 1/4th the bid), if we choose $\gamma = 1$ then the constant in Theorem 4 becomes 0.25, i.e. our mechanism obtains at least a quarter of the VCG revenue.

We will now prove Theorem 4.

Lemma 4 implies that payments received by $M_\alpha(\gamma)$ are at least $W^{\text{sv}}(1 - \alpha)$ (since only survivors pay) and that bump

payments sum to at most $W^{\text{sv}} \frac{\alpha}{\gamma}$. It will suffice to show that $W^{\text{sv}} \geq \text{REV}_{\text{vcg}}[\mathbf{w}]/(1 + \gamma)$.

$$\text{Let } \mathbf{u}(i) = \begin{cases} \max(w^{\text{sv}}(i), w(i)/(1 + \gamma)), & \text{if } i \in S \\ \frac{w(i)}{1 + \gamma}, & \text{if } i \notin S \end{cases}$$

Lemma 3 states that $\forall s \in S, s \in \text{OPT}[\tilde{\mathbf{w}}_S]$. Since on $\tilde{\mathbf{w}}_S$ VCG payments cannot be higher than its efficiency, $W^{\text{sv}} \geq \text{REV}_{\text{vcg}}[\tilde{\mathbf{w}}_S]$. Also, Lemma 5 implies

$$\text{REV}_{\text{vcg}}[\tilde{\mathbf{w}}_S] = \text{REV}_{\text{vcg}}[\mathbf{u}] \geq \text{REV}_{\text{vcg}}[\mathbf{w}]/(1 + \gamma),$$

since when going from $\tilde{\mathbf{w}}_S$ to \mathbf{u} only VCG winners may increase their bid, and for all i , $\mathbf{u}(i) \geq w(i)/(1 + \gamma)$.

Note that unlike the analogous efficiency result (Theorem 2), this result makes no assumption on bidders’ utilities.

6. SPECULATORS

Since money is given away, speculators, that is, bidders without interest in any item are likely to enter the mechanism looking for bump payments. For a speculator i , utility is also given by Eq. (1), but with $v(i) = 0$. Speculators may bid (under false identities) more than once or collude. Their bids can effectively induce reserve prices, since actual bidders will have to bid a $1 + \gamma$ factor higher than a competing speculator. If speculators bid judiciously on high-demand items, they can garner payment from the auctioneer, who gets even more revenue via larger prices for high-value bidders. So, it is clear that speculators are impactful.

We now address the question of how speculators affect the mechanism. In Section 6.1, we show two positive results

- We show that the $M_\alpha(\gamma)$ algorithm has good overall efficiency, as long as the speculators have positive overall surplus and the survivors are best-responding.
- We prove a bound on the overall revenue that speculators can obtain.

In Section 6.2 we give a more detailed discussion of speculator strategy. Along the way, we show that many natural simplifying assumptions regarding speculators’ or bidders’ strategies are unfortunately false. Specifically, we show that:

- the profits available for speculators may depend on the arrival order of actual bidders (Example 2);
- there may be no pure Nash equilibrium for actual bidders or speculators (Example 2);
- speculators may prefer to induce a suboptimal perfect matching of actual bidders (Example 3);
- a colluding set of speculators may be able to get higher bump payments if some of them survive (Example 4).

6.1 Impact on Efficiency and Revenue

The following result gives a competitive ratio of our algorithm’s efficiency with respect to the optimum efficiency given bidders’ true values. It only requires that total speculator utility is non-negative: this is particularly applicable if speculators are coordinated and can make monetary transfers between them.

PROPOSITION 1. *Let \mathbf{w} be a set of bids such that actual survivors are best-responding and total speculator surplus,*

i.e. the sum of speculators' payments minus the sum of speculators' prices, is non-negative. Then the $M_\alpha(\gamma)$ algorithm with true values \mathbf{v} has efficiency

$$\sum_{i \text{ actual}, i \in S} v(i) \geq \frac{1 - \alpha - \frac{\alpha}{\gamma}}{(1 - \alpha)(1 + \gamma)} \cdot \text{OPT}[\mathbf{v}]$$

The proof is mostly algebraic and deferred to the Appendix. This result is a strengthening of Theorem 2: Prop. 1's constant is larger and its preconditions are less general. Note that Prop. 1 requires that *actual* survivors best-respond; a speculator's best response cannot induce it to survive.

Next we prove an upper bound on speculators' profit:

PROPOSITION 2. *Speculators' total profit is at most $\alpha \text{OPT}/\gamma$.*

PROOF. Let Σ be the sum of survival weights for speculators that have survived. Denote speculators' profit by $\Pi \leq -(1 - \alpha)\Sigma + \alpha w(R)$, where R are the participants who obtain bump payments (some may be true bidders).

By Lemma 4, $w(R) \leq (\Sigma + A)/\gamma$, where A is the total weight of survivors that are actual bidders. We get

$$\Pi \leq -(1 - \alpha - \frac{\alpha}{\gamma})\Sigma + \frac{\alpha}{\gamma}A$$

The claim follows since $(1 - \alpha - \frac{\alpha}{\gamma})\Sigma \geq 0$ and $A \leq \text{OPT}$. \square

6.2 Speculator Strategies

At first glance, it would seem that it is in the speculators' best interest to induce an assignment of actual bidders of weight as high as possible in the survivor set, since then overall bump payments would be maximized. This is true in some cases but not always (Example 3). The reason for such a distinction is that the *order* of bidders arriving also influences the maximum refunds attainable by speculators as shown below.

EXAMPLE 2. *Consider two bidders, one bidding 1, the other $C > 1$, on two items and assume that speculators cannot collude. If C arrives first, no speculator can have higher revenue if bumped than when bidding $1/(1 + \gamma)$ on both items: this is actually a Nash equilibrium (NE) for them. If 1 however arrives first, then speculators could participate with two identities bidding $1/(1 + \gamma)$ and $C/(1 + \gamma)$ on both items, both being bumped. One can show via a case analysis that there is no pure strategy NE for speculators.*

This example also shows that there may not be a pure strategy NE when only actual bidders participate: if two bidders with low values arrive, followed by the 1 bidder and after that the C bidder, then the two low value bidders are essentially speculators and the argument in the example applies.

Observe that a speculator who is bumped with a bid of x could have obtained more bump payment by entering an earlier bid of at most $x/(1 + \gamma)$; likewise, he could have obtained yet more by bidding earlier $x/(1 + \gamma)^2$; and so on. This is formalized as follows.

DEFINITION 5. *Let $x > 0$. We say that the speculator σ is an x -geometric speculator with choice set $N(i)$ if σ places bids as follows on choice set $N(i)$. Let ε be the minimum strictly positive bid that can be made and*

$$l = 1 + \left\lfloor \frac{\log(x/\varepsilon)}{\log(1 + \gamma)} \right\rfloor \text{ i.e. } l \in \mathbb{Z} \ \& \ \frac{x}{(1 + \gamma)^l} \geq \varepsilon > \frac{x}{(1 + \gamma)^{l+1}}$$

Then σ participates with $l + 1$ different identities, placing consecutive bids of $\frac{x}{(1 + \gamma)^l}, \frac{x}{(1 + \gamma)^{l-1}}, \dots, \frac{x}{(1 + \gamma)}, x$ on $N(i)$.

If speculators have full information on bidders' values and bidders in OPT arrive in increasing order of their values, the outcome has many desirable properties:

LEMMA 6. *Fix a set of actual bids such that $\text{OPT}[\mathbf{v}]$ bids arrive in increasing order. Suppose that speculators collude and want to maximize their joint revenue. Then optimal speculator bidding has the following consequences:*

- *no speculator survives, no actual bidder is bumped; all OPT bidders and only them are accepted.*
- *speculators can achieve the highest payoff possible as given by Proposition 2.*
- *truthful bidding is a NE for all actual bidders.*

This is a further strengthening of Theorem 2. Optimal speculator bidding in this case is as follows. For each bidder $i \in \text{OPT}$ with choice set $N(i)$ there will be one $w(i)/(1 + \gamma)$ -geometric speculator σ_i with the same choice set. The proof is deferred to the full version.

This result has an appealing interpretation. If very well informed, speculators can overcome the efficiency loss due to late bidders not being able to improve by a $1 + \gamma$ factor over their earlier competitors.

In general however, speculators may prefer to induce a suboptimal perfect matching:

EXAMPLE 3. *Consider two items $\{i_1, i_2\}$ and three bidders b_1, b_2, b_3 arriving in this order; bidder k is interested in item $i_k, k = 1, 2$, while bidder 3 is interested in any of i_1 or i_2 . Note that any matching that does not match all three bidders is valid. Assume that $w(b_1) < w(b_3) < (1 + \gamma)w(b_1)$ and $w(b_2) > 2w(b_3)$. The following analysis shows that speculators prefer the suboptimal set of actual bidders b_1 and b_2 to the optimal one with b_2 and b_3 .*

- *If both b_2 and b_3 survive, then speculators' profit is at most $2w(b_3)/\gamma$: the speculator bumped by b_2 must have a lower weight than the one bumped by b_3 , which is at most $w(b_3)/(1 + \gamma)$. Even if speculators are geometric, speculator profit can only go as high as $2w(b_3)/\gamma$.*
- *If however b_1 and b_2 are alive when b_3 arrives, b_3 cannot bump b_1 . By simply having one geometric $w(b_2)/(1 + \gamma)$ -speculator which is bumped by b_2 , speculator profit is $w(b_2)/\gamma > 2w(b_3)/\gamma$.*

The following example shows that speculators may be able to make more money if they "sacrifice", i.e. some of them intentionally survive so that others obtain high refunds:

EXAMPLE 4. *Let there be k items, $k - 1$ actual bidders bidding $C > 1$ all arriving before an actual bidder bidding 1; all k bidders bid on all the items. If speculators coordinate and participate with k identities as $C/(1 + \gamma)$ -geometric speculators on all the items then total speculator payoff is*

$$(k - 1)\alpha C/\gamma - (1 - \alpha)C/\gamma = (k\alpha - 1)C/\gamma$$

since $k - 1$ will be bumped, but one will survive. If no speculator survives, the most money speculators can make is k/γ , by participating as k $1/(1 + \gamma)$ -geometric speculators. For any $\alpha > 1/k$, for a large enough C , speculators' profit is higher when one of them is sacrificed.

7. OTHER GAME-THEORETIC CLAIMS

We will now show that several appealing statements regarding incentives in our algorithm are false.

The algorithm may be more appealing for incentive purposes if we paid a bumped bidder $\alpha w^{sv}(i)$ instead of $\alpha w(i)$ as bump payment. This, however, may result in a deficit:

EXAMPLE 5. Fix α and consider an early bidder e bidding 1 and a late bidder ℓ bidding L on one item where $L > (1 + \gamma)^2/\alpha$. Bidder ℓ survives and pays $(1 + \gamma)$. If we were to refund e an α fraction of $w^{sv}(e)$, e would get $\alpha L/(1 + \gamma)$. The choice of L ensures that e is paid more than ℓ pays, i.e. the mechanism runs a deficit.

We assumed throughout that as soon as a bidder arrives, its choice set is known. If however that is private information as well, incentives become weaker: in Example 6, no bid by B^* on its true choice set $\{i_1, i_2\}$ is a best-response if bidding on different item(s) instead is allowed. This example also suggests why a naive generalization of $M_\alpha(\gamma)$ to the setting where bidders have a different value for each of several items would not be able to incentivize bidders to bid at least their true value for each item.

EXAMPLE 6. Consider two items i_1, i_2 and the following set of three bidders (arriving in this order): $B^{-3/2}$ with value $(1 + \gamma)^{-3/2}$ for any of i_1, i_2 (only demanding one of them), B^* who has value $x < \alpha(1 + \gamma)^{-3/2}$ for item i_2 and B^1 bidding 1 on item i_1 . Assume $B^{-3/2}$ and B^1 bid truthfully. We will show that, whenever B^* bids on $\{i_1, i_2\}$, it can do strictly better by bidding on i_1 only.

We claim that if B^* bids on $\{i_1, i_2\}$ then its utility is at most $\alpha(1 + \gamma)^{-3/2}$. This is clear if it survives. If it is bumped by B^1 , then its bid cannot be higher than $(1 + \gamma)^{-3/2}$ ($B^{-3/2}$'s bid), since B^1 can replace any of $B^{-3/2}$ and B^* . But then B^* 's compensation is at most $\alpha(1 + \gamma)^{-3/2}$. Let $0 < \varepsilon < 1/2$. By bidding $(1 + \gamma)^{-1-\varepsilon}$ on i_1 only and being bumped by B^1 , B^* can get utility $\alpha(1 + \gamma)^{-1-\varepsilon} > \alpha(1 + \gamma)^{-3/2}$.

We have however the following conjecture: if a bidder prefers surviving to being refunded, they are better off bidding on their true choice set.

We have an example (omitted for space) showing that, if bidders myopically and simultaneously best-respond (over sequences of instances of $M_\alpha(\gamma)$), then bid vectors where the sum of utilities is negative may be obtained. This example does not preclude good performance for other dynamics.

8. CONCLUDING REMARKS

Advertisers seek a mechanism to reserve ad slots in advance, while the publishers present a large inventory of ad slots with varying characteristics and seek automatic, online methods for pricing and allocation of reservations. In this paper, we present a simple model for auctioning such ad slots in advance, which allows canceling allocations at the cost of a bump payment. We present an efficiently implementable online mechanism to derive prices and bump payments that has many desirable properties of incentives, revenue and efficiency. These properties hold even though we may have speculators who are in the game for earning bump payments only. Our results make no assumptions about order of arrival of bids or the value distribution of bidders.

Our work leaves open several technical and modeling directions to study in the future. From a technical point of

view, the main questions are about designing mechanisms with improved revenue and efficiency, perhaps under additional assumptions about value distributions and bid arrivals. Also, mechanisms that limit further the role of speculators will be of interest. In addition, there are other models that may be applicable as well. Interesting directions for future research include allowing bidders to pay more for higher γ (making it harder for future bidders to displace this bidder) or higher α (being refunded more in case of being bumped). Other mechanisms may allow α to be a function of time between the acceptance and bumping. Accepted advertisers may be allowed to withdraw their bid at any time. There may be a secondary market where bidders may buy insurance against cancellations. Finally, advertisers may want a bundle of slots, say many impressions at multiple websites simultaneously, which will result in combinatorial extension of the auctions we study here. We believe that there is a rich collection of such mechanism design and analysis issues of interest which will need to inform any online system for advanced ad slotting with cancellations.

Acknowledgments. We would like to thank Stanislav Angelov for pointing us to [10] and the EconCS research group at Harvard for helpful discussions.

9. REFERENCES

- [1] M. Babaioff, J. Hartline, and R. Kleinberg. Selling banner ads: Online algorithms with buyback. In *Fourth Workshop on Ad Auctions*, 2008. To appear.
- [2] M. Babaioff, N. Immorlica, and R. Kleinberg. Matroids, secretary problems and online algorithms. In *Proceedings of SODA*, 2007.
- [3] S. Bikhchandani, J. Schummer, S. de Vries, and R. Vohra. An ascending Vickrey auction for selling bases of a matroid. November 2006.
- [4] M. C. Cary, A. D. Flaxman, J. D. Hartline, and A. R. Karlin. Auctions for structured procurement. In *Proceedings of SODA*, 2008.
- [5] N. B. Dimitrov and C. G. Plaxton. Competitive weighted matching in transversal matroids. UT Austin, Technical Report T-08-03, January 2008.
- [6] U. Feige, N. Immorlica, V. Mirrokni, and H. Nazerzadeh. A combinatorial allocation mechanism for banner advertisement with penalties. In *World Wide Web Conference (WWW)*, 2008.
- [7] J. Gallien and S. Gupta. Temporary and permanent buyout prices in online auctions. *Management Science*, 53(5):814–833, May 2007.
- [8] V. Krishna. *Auction Theory*. Academic Press, 2002.
- [9] R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. In *ACM Conf. on Electronic Commerce*, pages 233–241, 2000.
- [10] A. McGregor. Finding graph matchings in data streams. In *Proceedings of APPROX-RANDOM*, pages 170–181, 2005.

APPENDIX

A.1 Proof of Prop. 1

Assume that actual survivors bid honestly: $w(i) = v(i) \forall i$; at the end of the proof we will eliminate this assumption.

$$\text{Let } w'(i) = \begin{cases} w^{sv}(i), & \text{if } i \in S \text{ and } i \text{ is a speculator} \\ w(i), & \text{otherwise} \end{cases}$$

As $w^{\text{sv}}(i)$ can only be $(1+\gamma)w(k)$ if i bumps k or $w(i')$ if i' is bumped instead of i , the set of survivors will still be S if the mechanism is run on \mathbf{w}' instead of \mathbf{w} . We have

$$w'(S) = \sum_{i \in S} w'(i) \geq \frac{\text{OPT}[\mathbf{w}']}{1+\gamma} \geq \frac{\text{OPT}[\mathbf{v}]}{1+\gamma} \quad (6)$$

The first inequality follows from Lemma 4. Each actual bidder i bids at least its true value: $w'(i) \geq v(i)$ and there is the additional competition of speculators; this fact yields the second inequality. Also,

$$\gamma w(R) \leq W^{\text{sv}} \leq \sum_{\substack{i \text{ actual} \\ i \in S}} w(i) + \sum_{\substack{i \text{ speculator} \\ i \in S}} w^{\text{sv}}(i) = w'(S)$$

Again, the first inequality follows from Lemma 4. By assumption, speculator payments (a $(1-\alpha)$ fraction of their survival weights) cannot be higher than speculator refunds:

$$(1-\alpha) \sum_{\substack{i \text{ spec.} \\ i \in S}} w^{\text{sv}}(i) \leq \sum_{\substack{i \text{ spec.} \\ i \in R}} \alpha w(i) \leq \sum_{i \in R} \alpha w(i) = \alpha w(R)$$

By combining the last two relationships, we get

$$(1-\alpha) \sum_{i \text{ spec.}, i \in S} w^{\text{sv}}(i) \leq \frac{\alpha}{\gamma} w'(S)$$

Adding $(1-\alpha) \sum_{i \text{ actual}, i \in S} w(i)$ to both sides we get

$$(1-\alpha)w'(S) \leq \frac{\alpha}{\gamma}w'(S) + (1-\alpha) \sum_{i \text{ actual}, i \in S} w(i) \quad \text{i.e.}$$

$$\frac{1-\alpha-\frac{\alpha}{\gamma}}{1-\alpha}w'(S) \leq \sum_{i \text{ actual}, i \in S} w(i) = \sum_{i \text{ actual}, i \in S} v(i)$$

The last equality follows since actual survivors are bidding truthfully. This last inequality, together with Eq. (6) imply the proposition's claim.

Lemma 1 shows that bidding truthfully is a (weak) best-response for survivors. Let i be an actual survivor. Any bid above its true value is also a best-response for i . In the claim, as we increase i 's bid, the right-hand side quantity increases (if at all) at a constant rate which is less than 1, the left-hand side quantity's increase rate.

A.2 Proof of Lemma 3

Recall that bidder i arrives at time i . We will denote by $w_{\leq t}^{\text{sv}}(b)$ the minimum bid bidder b must make in order to survive up to and including time t . Then $w^{\text{ac}}(b) = w_{\leq b}^{\text{sv}}(b)$ and $w^{\text{sv}}(b) = w_{\leq T}^{\text{sv}}(b)$. It is clear that $w_{\leq t}^{\text{sv}}(b) \leq w_{\leq t+1}^{\text{sv}}(b)$.

DEFINITION 6. Let B be a set of bidders. We say that B is tight for a bidder i at time t if all bidders in B are alive at t , B can be matched but $B \cup \{i\}$ cannot be matched. We say that B $\tilde{\mathbf{w}}_S$ -dominates a bidder i at time t in the algorithm if B is tight for i at t and $\forall b \in B$, $w_{\leq t}^{\text{sv}}(b) \geq w(i)/(1+\gamma)$.

LEMMA 7. X_t is tight for t at t .

PROOF. X_t can be matched since $X_t \subseteq A_{t-1}$.

Suppose for a contradiction that $X_t \cup \{t\}$ can be matched. Then $X_t \neq A_{t-1}$ since A_{t-1} is a perfect matching by assumption. Therefore there exists $X \subset A_{t-1} \setminus X_t$, $|X| = |A_{t-1}| - |X_t| - 1$ such that $X_t \cup \{t\} \cup X$ can be matched. There exists exactly one bidder $\{y\} = A_{t-1} \setminus (X_t \cup X)$ and we have that $X_t \cup \{t\} \cup X = A_{t-1} \cup \{t\} \setminus \{y\}$ is a perfect matching, implying $y \in X_t$, contradiction. \square

Let i^* be the time step when i ceases to be alive (i.e. i is not accepted or the time i is bumped if i was accepted). We inductively construct a sequence $\{B_t\}_{i^* \leq t \leq n}$ as follows: if i is not accepted, $B_i = X_i$; if i is bumped by i^* then $B_{i^*} = X_{i^*} \cup \{i^*\} \setminus \{i\}$. At time $t \geq i^* + 1$,

- if no bidder in B_{t-1} is bumped, then we let $B_t = B_{t-1}$.
- if t bumps some $b \in B_{t-1}$ then we let $B_t = (B_{t-1} \cup X_t \cup \{t\}) \setminus \{b\}$

We will prove inductively on t that

LEMMA 8. B_t $\tilde{\mathbf{w}}_S$ -dominates i at time t .

PROOF OF LEMMA 8. Note that by definition, all bidders in B_t are alive at t .

Base case $t = i^*$:

If i is not accepted ($i^* = i$), i cannot bump any bidder in X_i : therefore $\forall b \in X_i$, $w_{\leq i}^{\text{sv}}(b) \geq w(i)/(1+\gamma)$. X_i is tight for i at i by Lemma 7.

If i is bumped, then $w(i) \leq w_{\leq i^*}^{\text{sv}}(r)$, $\forall r \in X_{i^*}$. $B_{i^*} = X_{i^*} \cup \{i^*\} \setminus \{i\}$ can be matched since they are all alive at i^* . $X_{i^*} \cup \{i^*\}$ cannot be matched since otherwise i^* would not need to bump $i \in X_{i^*}$.

Inductive step: Assume that B_{t-1} $\tilde{\mathbf{w}}_S$ -dominates i at $t-1$. If at time t , no bidder in B_{t-1} is bumped, then the claim obviously holds by the induction hypothesis. Otherwise, let $b \in B_{t-1}$ be the bidder that is bumped by t .

Clearly, $(B_{t-1} \cup X_t \cup \{t\}) \setminus \{b\}$ can be matched since they are alive at t . Suppose for a contradiction that $B_t \cup \{i\} = (B_{t-1} \cup X_t \cup \{t\}) \cup \{i\} \setminus \{b\}$ could be matched. $i \notin B_t$ since i is no longer alive. $B_{t-1} \cup X_t$ can be matched since they are all alive at $t-1$. As $|B_{t-1} \cup X_t| = |B_t \cup \{i\}| - 1$, either $B_{t-1} \cup X_t \cup \{i\}$ or $B_{t-1} \cup X_t \cup \{t\}$ can be matched. The first case is not possible since a subset, $B_{t-1} \cup \{i\}$, cannot be matched (by the induction hypothesis); the second case is not possible since $X_t \cup \{t\}$ cannot be matched (Lemma 7). We have reached a contradiction, so B_t must be tight for i .

By the induction hypothesis, $\forall b' \in B_{t-1}$, $w^{\text{sv}}(b')_{\leq t-1} \geq w(i)/(1+\gamma)$. As noted before, survival thresholds can only increase from $t-1$ to t and $w(t) \geq (1+\gamma)w(b)$. \square

We are ready for

PROOF OF LEMMA 3. Let V be the $\text{OPT}[\tilde{\mathbf{w}}_S]$ assignment (where ties are broken in favor of bidders in S). Suppose for a contradiction that there exists a non-survivor $i \in V$. By Lemma 8 for time n , i is dominated by a set $B_n \subseteq S$ at time n . Since $i \notin S$, but $B_n \subseteq S$, in $\tilde{\mathbf{w}}_S$ any bidder in B_n has a higher weight than i .

Since V is a perfect matching and B_n can be matched there must exist $V' \subset V \setminus B_n$, $|V'| = |V| - |B_n|$ ($V' = \emptyset$ if B_n is a perfect matching) such that $B_n \cup V'$ is a (perfect) matching. We know that $B_n \cup \{i\}$ cannot be matched, therefore $i \notin V'$. However, $i \in V$ therefore $i \in V \setminus V'$. $V \setminus \{i\}$ can be matched and has size $|V| - 1$. Therefore there $\exists b \in B_n \cup V'$, $b \notin V \setminus \{i\}$ such that $V \cup \{b\} \setminus \{i\}$ can be matched. That implies $b \in B_n \subseteq S$, i.e. $\tilde{\mathbf{w}}_S(b) \geq \tilde{\mathbf{w}}_S(i)$. But then $V \cup \{b\} \setminus \{i\}$ is a perfect matching of higher weight than V , contradiction.

That is, $V \setminus S = \emptyset$, i.e. $V = S$ since both are perfect matchings. \square