

An Online Mechanism for Ad Slot Reservations with Cancellations

Florin Constantin*

Jon Feldman[†]

S. Muthukrishnan[‡]

Martin Pál[§]

Abstract

Many advertisers (*bidders*) use Internet systems to buy display 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.

We propose and study a simple model for auctioning such ad slot reservations in advance. A seller will display a set of slots at some point T in the future. Until T , bidders arrive sequentially and place a bid on the slots they are interested in. 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, in which case the holder of the reservation incurs a utility loss amounting to a fraction of her value for the reservation and may also receive a cancellation fee from the seller.

Our main result is an online mechanism for allocation and pricing in this model with many desirable game-theoretic properties. It is individually rational. Winners have an incentive to be honest and bidding one's true value dominates any lower bid. Further, it bounds the earnings of *speculators* who are in the game to obtain the cancellation fees. The mechanism in addition has optimization guarantees. Its revenue is within a constant fraction of the *a posteriori* revenue of the Vickrey-Clarke-Groves (VCG) mechanism which is known to be truthful (in the offline case). Our mechanism's efficiency is within a constant fraction of the *a posteriori* optimally efficient solution. If efficiency also takes into account the utility losses of bidders whose reservation was canceled, we show that our mechanism matches (for appropriate values of the parameters) an upper bound on the competitive ratio of *any* deterministic online algorithm.

Our mechanism's technical core is a variant of the online *weighted* bipartite matching problem where unlike prior variants in which one randomizes edge arrivals or bounds edge weights, we may revoke previously committed edges.

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

[†]Google, 76 9th Ave, New York, NY, 10011. jonfeld@google.com.

[‡]Google, 76 9th Ave, New York, NY, 10011. muthu@google.com.

[§]Google, 76 9th Ave, New York, NY, 10011. mpal@google.com.

Our results make no assumptions about bidders' arrival order or value distribution. They still hold if we replace items with elements of a matroid and matchings with independent sets, or if all bidders have additive value for a set of items.

1 Introduction

Many advertisers now use Internet advertising systems. These take the form of advertisement (*ad*, henceforth) placements 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 ads placed in response to users' web search queries and in some other cases, ad slots are typically sold via *spot* auctions, i.e., when a user poses a query, 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. Advertisers typically like a reasonable guarantee of ad slots at a specific time (e.g. critical events) in the future within their budget constraints today.

Our motivation arises from systems enabling such advance ad slot reservations. In particular, our focus is on automatic systems that have to manage ad slots in many different publishers' properties. These properties differ wildly in their traffic, targeting, price and effectiveness; their inventory levels are massive. 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, effective, automatic, online¹ market-based mechanism to enable advance ad slotting over such varied, massive inventory.

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.”

Inspired by these considerations, we study the problem of mechanism design for advance ad slot reservations. 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.

Model. We propose the following simple model for advance ad slot reservations. An auction starts at time 0; the seller has a set of *slots* for sale that will be published at time T . Bidder i arrives at some time $a_i < T$, having a 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 accepted, he may be removed (*bumped*) later by the seller, in which case he incurs a loss of an α fraction of his value and may be compensated with a *bump payment*. At time T , each accepted bidder i that has not been bumped is published in one slot from $N(i)$ (his slots of interest).

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 inventory may be 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 such as in airline bookings. 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 pays for this via bump payments. We present our model formally in Section 2.

Mechanism. 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 her bid, making up for her utility loss.

Properties of Our Mechanism. We show a number of important strategic as well as efficiency- and revenue-related properties of $M_\alpha(\gamma)$. First, the strategic properties:

- $M_\alpha(\gamma)$ is individually rational and winners have an incentive to bid truthfully while losers should bid at least their true value.
- We 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.

Next, optimization properties:

- 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 also show that our mechanism is competitive with respect to the optimum offline efficiency on bidders' true *values*.
- We prove similar bounds under the notion of *effective efficiency* which interprets social welfare as the sum of the winners' bids minus bumped bidders' losses. 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 famed Vickrey-Clarke-Groves (VCG) mechanism on all received bids.

To the best of our knowledge our results are the first about mechanisms with strong game-theoretic properties for advance ad reservations (more generally, online weighted bipartite matching) with a costly cancellation feature. We make no assumptions on the arrival order of the bidders or on their values. We note that all our results extend to a more general setting where the items for sale are elements of a matroid; we discuss this in Section 7.

There are specific examples of systems that implement advance booking with cancellations, not necessarily though an automatic mechanism. 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, advance 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 setting 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 [11] to determine a suitable bump payment and prices.

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

1.1 Related work. Several papers consider offline settings similar to ours. Bikhchandani et al. [4] present an ascending, efficient, truthful in equilibrium auction for selling elements of a matroid to patient bidders. Cary et al. [5] show that a random sampling profit extraction mechanism approximates a VCG-based target profit in an offline procurement setting on a matroid. Feige et al. [7] 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. They show that approximating the optimal solution with respect to effective efficiency (see Section 2) within any constant factor is NP-hard, but they provide a bi-criterion approximation result for an adaptive greedy algorithm.

There has been extensive work in the field of revenue management for advance sale of goods (with or without cancellations), but only under a probabilistic distribution of bidders’ values or arrivals [13]. In particular, Gallien and Gupta [8] exhibit symmetric Bayes-Nash equilibria in online auctions with buyout prices.

Under a worst case model like ours, no nontrivial online results are possible without making additional assumptions; in our case, we overcome these impossibilities by allowing cancellations. In contrast, in secretary problems, bids may be arbitrary but their order is assumed to be uniformly random (cannot be specified by an adversary): Dimitrov and Plaxton [6] provide an algorithm with a constant efficiency competitive ratio. Generalizing their setting to matroids, Babaioff et al. [2] provide a $\log r$ -competitive algorithm where r is the rank of the matroid. Both these algorithms observe half of the input and then set a threshold price. A different assumption is that of bounded values: Lavi and Nisan [10] show that a simple online posted-price auction based on exponential scaling is optimal among online auctions for identical goods without cancellations.

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. They focus on effective efficiency, for which they analytically prove an upper bound on any deterministic algorithm’s competitive ratio (we only present numerical results 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.

2 Auction Model

We define an *online ad slot reservation auction* as follows. There is a seller who has a finite set of non-identical *slots*, which will be allocated at some future time T . The seller runs a continuous, online auction beginning at time 0, and ending at time T .

Each bidder i arrives online, at a unique time $a_i \in [0, T]$ and she reports a *choice set* $N(i)$ of items she is interested in, as well as a *bid* (positive amount) $w(i)$, demanding an immediate response (i is not allowed to bid again later). She is instantly *accepted* (i.e. promised an item from $N(i)$), or *rejected*. However, at any point between time a_i and T , the seller may choose to *bump* an accepted bidder i , in which case a *bump payment* \hat{p}_i is given to the bumped bidder. Any rejection, at arrival or by being bumped, is definitive. At time T , there must be a matching of items to accepted bidders that have not been bumped such that each such bidder i receives one item from her choice set $N(i)$; each such i is then charged a price p_i .

A *mechanism* for the online ad slotting problem defines the actions of the seller: whether to accept/reject incoming bidders, when to bump accepted bidders, and how to set bump payments and prices.

2.1 Bidder Valuation and Utility. We assume a *private value* model for the bidders: Each bidder i has a private value $v(i) \geq 0$ for being allocated (at time T) any single item from her choice set $N(i)$. The bidder does not necessarily need to report this value as her bid if it is in her interest not to do so. Additionally, we will model the cost incurred by a bumped bidder as a negative value $-\alpha v(i)$, that is, an α fraction of her value for being allocated. We will require² that any mechanism pays back $\alpha w(i)$ to a bumped bidder i , making up for her utility loss when bumped. The parameter $0 \leq \alpha < 1$ modeling the negative bump utility will play a central role in our mechanism and analysis. We formally model bidder i ’s *utility* as quasilinear in money:

$$(2.1) \quad \text{utility}(i) = f \cdot v(i) - x(i), \quad \text{where}$$

- $f = 0$ if i is rejected, $f = 1$ if i is accepted and granted an item from $N(i)$, and $f = -\alpha$ if i is bumped;
- $x(i)$ is i ’s *money transfer* to the seller: $x(i) = 0$ if i is rejected, $x(i) = p_i \geq 0$ if i is accepted and allocated, and $x(i) = -\alpha w(i) \leq 0$ if i is bumped.

²We impose this constraint primarily to ensure that honest bidders have non-negative utility. See Example 3 for further motivation.

For a mechanism run on bids \mathbf{w} , we will denote by $S = S(\mathbf{w})$ the set of *survivors* (bidders still accepted by time T) and by $R = R(\mathbf{w})$ the set of bumped bidders.

2.2 Efficiency and Revenue. How do we measure the quality of an outcome? The *efficiency* (or *social welfare*) of an auction is the total value derived by the bidders participating in the auction. Usually in a mechanism design setting, efficiency is the sum of the valuations of the bidders who were allocated items:

$$\text{efficiency} = \sum_{i \in S} v(i).$$

However there is another interpretation in our model since the bidders lose value if they are accepted then bumped; thus we will also consider the notion of *effective efficiency*:

$$\text{effective efficiency} = \sum_{i \in S} v(i) - \sum_{i \in R} \alpha v(i).$$

A mechanism's *revenue* is its total monetary gain/loss:

$$\text{revenue} = \sum_{i \in S} p_i - \sum_{i \in R} \alpha w(i)$$

We would like a mechanism that scores favorably in all of these metrics on all instances of the auction, under hopefully mild conditions on strategic behavior. Following the logic of *competitive analysis*, we will compare our mechanism to the standard offline solution: the VCG mechanism [9]. Mapped to our setting (but offline), this amounts to finding a maximum matching of bidders to items, and charging prices that induce truthfulness.

Finally, we note that it is essential to the novelty of this model that the bidders derive negative utility from having their allocation promise revoked. Indeed if $\alpha = 0$, one could simply accept all bidders as they arrive, and then at time T run the VCG mechanism (giving no bump payments) which makes being truthful a dominant strategy.

3 Main Results

In this section we will state our main results (without proof³) and highlight the significance of each. In the next section, we will define our mechanism $M_\alpha(\gamma)$. The mechanism is parametrized both by the model parameter α as well as an additional parameter $\gamma > 0$ that can be set arbitrarily as long as $0 < \alpha < \frac{\gamma}{1+\gamma}$. We will state our main results in terms of these two parameters.

Since we are in a game-theoretic setting, we must first reason about the strategic behavior of bidders in order to motivate the preconditions of our results. One basic property that any reasonable mechanism must have is that it is *individually rational*, which simply means that participating in

the auction is always a rational thing to do (i.e. participating is never worse than not participating). In our case we can define this in the following way: if a bidder reports her true value (sets $w_i = v_i$), then her utility is always non-negative.

Another desirable property of an auction mechanism is for it to be *truthful*, which means that the optimal strategy for participating bidders is always to report their true value. Unfortunately with bump payments (which we just argued were necessary) we cannot hope to have a truthful mechanism since anyone with no interest in any allocation (i.e., $v_i = 0$) can bid hoping to get a bump payment. So, given that we cannot assume bidders will be honest, the natural thing to do is analyze the efficiency and revenue of the mechanism in a Nash or other form of equilibrium. Unfortunately, a (pure-strategy) Nash equilibrium does not always exist, as we will argue in Section 6.2. However we can still argue that our mechanism has some strong incentive properties. We use the following standard game-theoretic terminology: a bid *dominates* another bid if it is at least as good a strategy given *any* bids by other players; a *best-response* is the best possible bid given a particular set of others' bids.

THEOREM 3.1. [Basic Incentive Properties]

- *The $M_\alpha(\gamma)$ mechanism is individually rational,*
- *Bidding truthfully dominates any lower bid,*
- *If truthful, any survivor is best-responding, and*
- *Bidding truthfully is a best-response unless a higher bump payment can be achieved with a higher bid.*

This theorem (proved in Section 6.1) establishes individual rationality, but more importantly it rules out the possibility that the bids will be lower than the values. To the best of our knowledge, this is a novel form of incentive compatibility: while it does not make truthfulness a dominant strategy, it ensures that competition is no less than if every bidder were truthful. Furthermore, it highlights truthfulness as a simple viable strategy from a practical point of view, or for unsophisticated bidders. The only reason for not bidding truthfully is the prospect of a higher bump payment. This motivates the following definition:

DEFINITION 1. [Speculators] A *speculator* is a bidder who bids above her true value. The auctioneer does not know if a bidder is a speculator or not.

By Theorem 3.1, speculators are bidders who aim for getting a (higher) bump payment.

We can now state the efficiency and revenue bounds for $M_\alpha(\gamma)$ under preconditions made reasonable by Theorem 3.1. From Theorem 3.1 and the speculator definition, we can assume that non-speculators bid truthfully. Speculators' bids, on the other hand, are aimed towards squeezing

³Missing proofs in this paper are presented in the full version, available from the first author's website.

as much money as possible from bump payments. The only other assumption we make for our efficiency bounds is that the set of speculators does not incur a loss, which is quite mild an assumption. Indeed, if speculators incur a loss and can collude they would be better off not participating at all.

For any vector $\mathbf{w} = (w(1), \dots, w(n))$ of bids, we let $\text{OPT}[\mathbf{w}]$ denote the weight of the optimal matching. Note that $\text{OPT}[\mathbf{v}]$ then gives the optimal efficiency and effective efficiency of an offline mechanism, achieved by VCG. On bids \mathbf{w} we denote the VCG revenue by $\text{REV}_{vcg}[\mathbf{w}]$.

THEOREM 3.2. [Efficiency] *Let \mathbf{w} be a set of bids such that all bidders bid at least their true value, and total utility among speculators is non-negative. Then $M_\alpha(\gamma)$ has*

$$\text{efficiency} \geq \frac{1 - \alpha - \frac{\alpha}{\gamma}}{(2 - \alpha - \frac{\alpha}{\gamma})(1 + \gamma)} \cdot \text{OPT}[\mathbf{v}] \text{ and}$$

$$\text{effective efficiency} \geq \frac{1 - \alpha - \frac{\alpha}{\gamma}}{(2 - \alpha)(1 + \gamma)} \cdot \text{OPT}[\mathbf{v}].$$

THEOREM 3.3. [Revenue] *Let \mathbf{w} be a set of bids such that all bidders bid at least their true value. Then $M_\alpha(\gamma)$ has*

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

In the full version we also show $\text{REV}_{vcg}[\mathbf{w}] \geq \text{REV}_{vcg}[\mathbf{v}]$, implying that $M_\alpha(\gamma)$'s revenue is also competitive with the offline VCG revenue on bidders' true values.

Note that a limit on manipulation is needed for a lower bound on true efficiency: if low value bidders bid really high, being allocated all the items and preventing the rightful winners from being allocated, the true efficiency of the mechanism is very low. We further discuss manipulations in Section 6.2, where we also give additional results on speculator strategies.

Efficiency bounds that leave incentives aside (Theorem 5.1 and Corollary 5.1 in Section 5) are tighter than the bounds in Theorem 3.2, and can be obtained more easily.

In Section B we give an *upper bound* on the effective efficiency (in terms of bids) of any deterministic algorithm, for which our allocation algorithm is tight (only when $\alpha < 0.618$ and for a certain γ that depends on α).

For example, suppose $\alpha = \frac{1}{4}$. Then we can set $\gamma = 1$, and be $\frac{1}{6}$ -competitive on efficiency, $\frac{1}{7}$ -competitive on effective efficiency, and $\frac{1}{4}$ -competitive on revenue.

We defer the proofs of Theorems 3.2 and 3.3 to the full version. As a revenue benchmark, we use the offline Vickrey-Clarke-Groves (VCG) mechanism [9], that charges each bidder her externality on the other bidders. Theorem 3.3's proof uses Lemma 5.1 and the following facts: a winning bidder's VCG payment is a losing bid and the VCG revenue can only increase if some bids are increased.

4 $M_\alpha(\gamma)$: an Online Mechanism

We present our advance-booking online mechanism $M_\alpha(\gamma)$ in this section. The allocation algorithm follows the *Find-Weighted-Matching* algorithm in [11]⁴, 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 $0 \leq \alpha < 1$) for non-negative lower bounds in Theorems 3.2 and 3.3.

Our mechanism $M_\alpha(\gamma)$ (given formally in Algorithm 1) maintains a set of accepted bidders for which there exists a matching of bidders to items. For each new arriving bidder i bidding $w(i)$, $M_\alpha(\gamma)$ adds i to the current matching if it can do so without bumping a currently accepted bidder. Otherwise, $M_\alpha(\gamma)$ looks for some bidder j in the accepted set with $w(j) < \frac{w(i)}{1+\gamma}$ such that replacing j by i maintains the existence of a matching. If such a bidder exists, the mechanism accepts i and bumps j^* , the lowest weight such bidder, who is paid the bump payment $\alpha w(j^*)$. At time T , after all bidders have been processed, the accepted bidders become the survivors. The survivors are allocated a slot from their choice set using an arbitrary matching, and they each make a payment that we define below.

Algorithm 1 $M_\alpha(\gamma)$: Allocation algorithm. A new bidder is accepted if she improves by at least a γ factor over her lowest-bidding indirect competitor. A bumped bidder receives a bump payment making up for her utility loss.

Let $A_0 := \emptyset$.

for each arriving bidder $i \geq 1$ bidding $w(i)$ **do**

if $A_{i-1} \cup \{i\}$ can be matched **then**

 grant i a reservation: $A_i := A_{i-1} \cup \{i\}$.

else let j^* be the lowest-bidding $j \in A_{i-1}$ such that $A_{i-1} \cup \{i\} \setminus \{j^*\}$ can be matched

if $w(i) < (1+\gamma)w(j^*)$ **then** reject i : $A_i := A_{i-1}$.

else cancel j^* 's reservation and pay her $\alpha w(j^*)$

 grant i a reservation: $A_i := A_{i-1} \cup \{i\} \setminus \{j^*\}$.

end for

Each bidder i in $S = A_n$ (i.e. survivors) is allocated an item from $N(i)$ and charged as in Eq. (4.2).

Eq. (4.2) below establishes a survivor's payment to the seller, and requires the following definitions.

DEFINITION 2. *Let i be a bidder and fix the bids of all other bidders. Let $w^{ac}(i)$ (i 's acceptance weight) be the infimum of all bids that i can make such that i is accepted given i 's arrival a_i and i 's choice set N_i . Similarly, let $w^{sv}(i) \geq w^{ac}(i)$ (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 .*

⁴Unlike in [11], a bidder i 's value is the same for any slot (vertices as opposed to edges are weighted). Our mechanism may then change the slot i is currently assigned to at various stages in the algorithm.

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 the time i arrives and on the other bidders' bids or arrivals.

We are now ready to define the prices charged by $M_\alpha(\gamma)$. If i is a survivor, we set i 's price p_i as follows:

$$(4.2) \quad p_i = \begin{cases} w^{\text{sv}}(i)(1 - \alpha) & \text{if } w^{\text{ac}}(i) < w^{\text{sv}}(i). \\ w^{\text{sv}}(i) & \text{if } w^{\text{ac}}(i) = w^{\text{sv}}(i). \end{cases}$$

These prices are designed with Theorem 3.1's conditions in mind, as we will see in its proof.

We present now an example run of $M_\alpha(\gamma)$. $w^{\text{sv}}(i)$ at some time step is i 's survival weight if $M_\alpha(\gamma)$ stopped then.

EXAMPLE 1. (A PARTICULAR INSTANCE OF $M_\alpha(\gamma)$)

Suppose $\alpha < \frac{0.5}{0.5+1}$ and let $\gamma = 0.5$. Consider two items I_a, I_b and bidders $B_{1..4}$ (B_i arrives at time t_i ; $T = t_4$): B_1 bids 6 on $N(1) = \{I_a, I_b\}$, B_2 bids 4.4 on $N(2) = \{I_b\}$, B_3 bids 10 on $N(3) = \{I_a\}$ and B_4 bids 7.5 on $N(4) = \{I_b\}$. $M_\alpha(\gamma)$ 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 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); we have $w^{\text{sv}}(1) = \frac{7.5}{1.5} = 5$ (to prevent being bumped by B_4), $w^{\text{sv}}(2) = 6$ (to prevent being bumped by B_3 and B_4), $w^{\text{sv}}(3) = 6.6$, and $w^{\text{sv}}(4) = w^{\text{ac}}(4) = 9$. 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 $w^{\text{ac}}(3) = w^{\text{sv}}(3)$.

5 Online Matching with Cancellations

For our incentive-aware bounds from Section 3, $M_\alpha(\gamma)$ must find a good matching given the declared bids \mathbf{w} , regardless of those bidders' true values. This is a pure online-algorithms question (i.e., no game theory), which we treat in this section, providing theorems which not only are key for the Section 3 bounds, but are also of independent interest.

Recall that $\text{OPT}[\mathbf{w}]$ denotes the optimal offline matching on the bids \mathbf{w} . For bids \mathbf{w} and a set of bidders B , we let $w(B) = \sum_{i \in B} w(i)$ and $w^{\text{sv}}(B) = \sum_{i \in B} w^{\text{sv}}(i)$.

Theorem 5.1 shows a competitive ratio for efficiency (the difficult part of the proof is deferred to Section A):

THEOREM 5.1. Mechanism $M_\alpha(\gamma)$ is a $\frac{1}{1+\gamma}$ -approximation to the optimal offline matching: $w(S) \geq \frac{1}{1+\gamma} \text{OPT}[\mathbf{w}]$.

Proof. Our key technical lemma shows that if non-survivor bids are essentially given a late-arriving penalty, then the optimal offline matching coincides with $M_\alpha(\gamma)$'s solution:

LEMMA 5.1. $S = \text{OPT}[\tilde{\mathbf{w}}]$ for $\tilde{\mathbf{w}}(i) = \begin{cases} w^{\text{sv}}(i), & \text{if } i \in S \\ \frac{w(i)}{1+\gamma}, & \text{if } i \notin S \end{cases}$

We prove this Lemma in Section A. To finish the theorem, let $\hat{\mathbf{w}}(i) = \max(w^{\text{sv}}(i), \frac{w(i)}{1+\gamma})$ if $i \in S$, and $\frac{w(i)}{1+\gamma}$ otherwise. We

have $w(S) \geq \hat{\mathbf{w}}(S) = \text{OPT}[\hat{\mathbf{w}}] \geq \text{OPT}[\mathbf{w}]/(1 + \gamma)$: each inequality is implied by the fact that no bidder's contribution decreases when going from the left to the right hand side. Lemma 5.1 yields the equality: when going from $\tilde{\mathbf{w}}$ to $\hat{\mathbf{w}}$ only bids already in the optimum (i.e. S) can increase.

The following bound assures us that not too much utility (of bumped bidders) is sacrificed for high efficiency:

THEOREM 5.2. Total bumped bidder weight $w(R) \leq \frac{w^{\text{sv}}(S)}{\gamma}$.

Proof. For an $r \in R$, let $s^*(r) \in S$ be the survivor at the end of the sequence of bumps that starts from r . For an $s \in S$, let R_s be the refunded bidders in s 's sequence of bumps: $R_s = \{r \in R : s^*(r) = s\}$. As R is the disjoint union of R_s for all $s \in S$, the theorem follows by showing:

$$(5.3) \quad \text{For all } s \in S, \quad w(R_s) \leq w^{\text{sv}}(s)/\gamma.$$

To show Eq. (5.3), fix $s \in S$, and, to simplify notation, assume that the elements in R_s are $1, 2, \dots, s-1, s$: $j+1$ bumps j , $\forall 1 \leq j < s$. We now show that: $\sum_{j=1}^{s-1} w(j) \leq w^{\text{sv}}(s)/\gamma$. As s bumps $s-1$, $w_{s-1} \leq \frac{w^{\text{sv}}(s)}{1+\gamma}$. As $j+1$ bumps j , $\forall 1 \leq j \leq s-2$, $w_j \leq \frac{w_{j+1}}{1+\gamma}$. Thus by induction, $w_j \leq w^{\text{sv}}(s)(1+\gamma)^{j-s}$, $\forall 1 \leq j < s$. We get $\sum_{j=1}^{s-1} w_j \leq w^{\text{sv}}(s) \sum_{j=1}^{s-1} (1+\gamma)^{j-s} \leq w^{\text{sv}}(s)/\gamma$.

Let the effective weight of a solution be the weight of the matching minus a penalty amounting to the total utility loss by bumped bidders ($\alpha w(i)$ for each $i \in R$). Theorem 5.2 implies $w(S) - \alpha w(R) \geq w(S)(1 - \alpha/\gamma)$, implying (by Theorem 5.1) the following lower bound on effective weight:

COROLLARY 5.1. The $M_\alpha(\gamma)$ mechanism is a $\frac{1-\alpha/\gamma}{1+\gamma}$ -approximation to the optimal offline matching in terms of effective weight: $w(S) - \alpha w(R) \geq \frac{1-\alpha/\gamma}{1+\gamma} \text{OPT}[\mathbf{w}]$.

Theorems 5.1 and 5.2 have analogs in [11]. 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. Our bounds are almost tight:

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

We also show an upper bound on how well any deterministic algorithm can approximate the effective weight:

THEOREM 5.3. Fix n (the number of bidders). No deterministic online algorithm can approximate the optimal offline matching in terms of effective weight with a factor better than c_n , where c_n is the lowest number (if any) in $[0, 1]$

for which Eqs. (2.4) and (2.5) simultaneously hold (see Section B). Based on computing c_n numerically, we conjecture that c_n approaches $2\alpha + 1 - 2\alpha^{0.5}(\alpha + 1)^{0.5}$ as $n \rightarrow \infty$.

For $\alpha < \frac{\sqrt{5}-1}{2} \simeq 0.618$ and the best γ given α , the approximation ratio in Corollary 5.1 for $M_\alpha(\gamma)$ matches this upper bound. (See Section B for further discussion.)

6 Detailed game-theoretic analysis

We now focus on the game-theoretic properties induced by our mechanism. We first offer some more intuition on survival and acceptance weights and then prove Theorem 3.1.

Recall from Section 4 that i is rejected if $w(i) < w^{ac}(i)$, bumped if $w^{ac}(i) \leq w(i) < w^{sv}(i)$ and a survivor if $w^{sv}(i) \leq w(i)$. If i bumps j^* , $w^{ac}(i) = (1 + \gamma)w(j^*)$ but $w^{sv}(i)$ can either be $(1 + \gamma)w(j^*)$, $w(k)$ for a (past or future) bumped bidder k or $\frac{w(k)}{1+\gamma}$ for a future rejected k . Thus, $w^{ac}(i)$ can be computed when i arrives but $w^{sv}(i)$ may depend on future bidders and can only be computed at T .

Let us focus now on a survivor i 's ($w(i) \geq w^{sv}(i)$) payment in Eq. (4.2). The common case is when $w^{ac}(i) < w^{sv}(i)$: i gets a discount amounting to the highest bump payment she could have otherwise obtained: $\alpha w^{sv}(i)$. The special case of $w^{ac}(i) = w^{sv}(i)$ occurs when i 's acceptance is enough for her survival (in particular if i is the last bidder). When $w^{ac}(i) = w^{sv}(i)$, from the bidder's point of view, $M_\alpha(\gamma)$ posts a price of $w^{sv}(i)$.

Say that in Example 1 a bidder B_5 were to arrive after B_4 bidding 10.5 on I_a . Only $w^{sv}(3)$ would change to $\frac{10.5}{1.5} = 7$ and B_3 's price would now be $(1 - \alpha) \cdot 7$ which may be lower than 6.6. Unless a bidder i 's $w^{sv}(i)$ price coefficient goes from 1 to $1 - \alpha$ (like in Example 1 for B_3 if B_5 arrives), i 's price cannot go down if new bidders arrive.

6.1 Proof of Theorem 3.1. If bidder i bids her true value $v(i)$, then her utility is either $v(i) - p_i \geq v(i) - w^{sv}(i) \geq 0$ if she survives, 0 if she is rejected, or $\alpha v(i) - \alpha v(i) = 0$ if she is bumped. This establishes (1).

If $w^{ac}(i) < w^{sv}(i)$, bidder i 's highest possible bump payment is $\alpha w^{sv}(i)$. The price of $(1 - \alpha)w^{sv}(i)$ has been chosen such that i prefers winning to being paid $\alpha w^{sv}(i)$ if and only if $v(i) \geq w^{sv}(i)$. That is, i 's best response is to bid just below $w^{sv}(i)$ if $v(i) < w^{sv}(i)$ and to bid her true value otherwise. This establishes (2), (3) and (4) for this case.

If $w^{ac}(i) = w^{sv}(i)$, then i can never get a bump payment and i simply faces a take-it-or-leave-it offer of $w^{sv}(i)$. Bidding truthfully is a best response in this case, and (2), (3) and (4) follow.

6.2 Speculator Strategies. In this section we focus on speculators' strategic behavior, assuming speculators have more manipulations available: speculators may collude with each other, lie about their choice sets $N(i)$ or arrival times, and may also have value $v(i) = 0$ (i.e., they may have

no value for an item).

Prop. 6.1 bounds speculators' total monetary profit:

PROPOSITION 6.1. *Speculators' profit is at most $\alpha \text{OPT}/\gamma$.*

At first glance, it would seem that speculators' best strategy is 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, and indeed colluding speculators may even want some of them to survive:

THEOREM 6.1. *Under the $M_\alpha(\gamma)$ mechanism,*

- *there exist input instances such that optimal speculator bidding induces optimal efficiency, and truthful bidding is a Nash Equilibrium for all non-speculators;*
- *there exist input instances where optimal speculator bidding induces sub-optimal efficiency,*
- *there exist input instances where there is no pure-strategy Nash equilibrium,*
- *there exist input instances where speculators achieve higher profit if they "sacrifice," i.e. some of them intentionally survive so that others obtain high refunds.*

We may conclude from this theorem that it is unreasonable to expect stronger incentive properties than Theorem 3.1, such as truthfulness or a Nash Equilibrium. However, despite all the game-theoretic complexity that can arise from speculators, their effect on efficiency and revenue can still be bounded: implicitly via the results in Section 3 or explicitly in Prop. 6.1.

6.3 Other game-theoretic considerations. We now ask a couple of "what if?" questions, whose answers further help motivate our model choices.

For stronger incentive properties, a standard modification to $M_\alpha(\gamma)$ is to pay a bumped bidder i $\alpha w^{sv}(i)$ (a bid-independent amount) instead of $\alpha w(i)$. Example 3 shows why this may result in a deficit:

EXAMPLE 3. (Alternate bump payments) *Consider two bidders on one item: B_1 arrives first, bidding 1, followed by B_2 bidding $L > (1 + \gamma)^2/\alpha$. Bidder B_2 survives and pays $1 + \gamma$. If B_1 's bump payment were $\alpha w^{sv}(e) = \alpha L/(1 + \gamma)$ then the choice of L ensures that B_1 is paid more than B_2 pays, i.e. the mechanism runs a deficit.*

We assumed throughout that as soon as a bidder arrives, her choice set is known. If however that is private information as well, incentives become weaker: in an example in the full version, no bid by a certain bidder on her true choice set is a best-response if bidding on different item(s) instead is allowed. We have however the following conjecture: if a bidder prefers surviving to being refunded, she is better off bidding on her true choice set.

One can show that, if bidders myopically and simultaneously best-respond (over sequences of instances of $M_\alpha(\gamma)$), they may reach bid vectors with a negative sum of utilities.

Given the instantaneous nature of bidders, temporal manipulations (such as changes in the order of arrival or multiple bids) become less motivated. Note however that in $M_\alpha(\gamma)$ a certain item's price is never decreasing with a new bidder which makes more attractive only temporal manipulations that improve bump payments.

6.4 Implementation in undominated strategies

In this section, we position our results with respect to a related algorithmic game theory concept.

DEFINITION 3. (BABAIOFF ET AL. [3]) *A mechanism M is an algorithmic implementation in undominated strategies (AIUS) of a c -approximation algorithm if there exists a set D of dominating strategies such that*

- *M obtains a c -approximation whenever all players play strategies from D*
- *for any strategy $s \notin D$, there exists $s' \in D$ (computable in polynomial time from s) that dominates s .*

Let D be the set of strategies representing truthful bids and overbids. Fix a c and a domain V such that VCG's revenue is a c fraction of the offline optimum for any set of bids from V . Then $M_\alpha(\gamma)$ is an AIUS of the VCG mechanism on V with respect to revenue. $M_\alpha(\gamma)$ only satisfies Def. 3 for efficiency under the mild assumption of Theorem 3.2. Otherwise, for D , the true value of winners may be small (if winners overbid to the extreme) even though any strategy not in D is dominated by being truthful. On the other hand, being truthful does not dominate a higher bid: a higher bump payment may be achieved by a losing truthful bidder.

However, being truthful is an equilibrium in undominated strategies of $M_\alpha(\gamma)$ and we argue that it is a *focal point* (Schelling [12]), being highlighted by Theorem 3.1.

7 Extensions

All our results extend to a setting where the items for sale are elements of a matroid, which is more general 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 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 her subset.

In a different direction, our results also extend to a (strictly bipartite matching) setting where a bidder's value for a set I of items is the sum of values for each item in

I (no bidder can express substitute items). In this setting, the multi-item matching problem is actually a collection of single-item matching problems since bids on two different items can never interact. We preferred the basic setup (bipartite matching where a bidder is interested in one out of a set of substitute items) for clarity of exposition.

By starting out with dummy bidders with non-zero values, the seller can set *reserve prices* on items.

8 Concluding Remarks

Advertisers seek a mechanism to reserve ad slots in advance, while publishers present a large inventory of ad slots and seek automatic, online 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 bids' order of arrival or their value distribution.

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 a bidder 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. Finally, advertisers may want a bundle of slots, say many impressions at multiple websites simultaneously, which will require a combinatorial extension of our model and mechanism.

We believe that there is a rich collection of such mechanism design and analysis issues of interest that will need to inform any online system for advance ad slot reservations.

Acknowledgments

We thank Stanislav Angelov for pointing us to [11] and the EconCS research group at Harvard for helpful discussions.

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 *SODA*, 2007.
- [3] M. Babaioff, R. Lavi, and E. Pavlov. Single-value combinatorial auctions and implementation in undominated strategies. In *SODA*, 2006.
- [4] S. Bikhchandani, J. Schummer, S. de Vries, and R. Vohra. An ascending Vickrey auction for selling bases of a matroid. November 2006.
- [5] M. C. Cary, A. D. Flaxman, J. D. Hartline, and A. R. Karlin. Auctions for structured procurement. In *SODA*, 2008.
- [6] N. B. Dimitrov and C. G. Plaxton. Competitive weighted matching in transversal matroids. UT Austin, Technical Report T-08-03, January 2008.
- [7] 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.
- [8] J. Gallien and S. Gupta. Temporary and permanent buyout prices in online auctions. *Mgmt. Science*, 53:814–833, 2007.
- [9] V. Krishna. *Auction Theory*. Academic Press, 2002.
- [10] R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. In *ACM Conf. on Electronic Commerce*, pages 233–241, 2000.
- [11] A. McGregor. Finding graph matchings in data streams. In *Proceedings of APPROX-RANDOM*, pages 170–181, 2005.
- [12] T. C. Schelling. *The Strategy of Conflict*. Harvard University Press, 1960.
- [13] K. Talluri and G. V. Ryzin. *The theory and practice of revenue management*. Kluwer Academic Publ, 2004.

A Proof of Lemma 5.1.

In this section it will be simpler to have as initial matching A_0 for $M_\alpha(\gamma)$ an arbitrary perfect matching on dummy bidders instead of the empty matching. We introduce dummy bidders (one per item, each bidding 0) whose choice set is the whole set of items, arriving before all actual bidders. This will ensure that a perfect matching is maintained by $M_\alpha(\gamma)$, but will not affect other arguments below⁵

At 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 swapped for t and $j^* = \operatorname{argmin}_{j \in X_t} w(j)$ (see Algorithm 1).

Assume wlog that bidder i arrives at time i . We 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 4. *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*

⁵When 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} .

that B γ -dominates a bidder i at time t if B is tight for i at t and $\forall b \in B$, we have $w_{\leq t}^{\text{sv}}(b) \geq w(i)/(1 + \gamma)$.

LEMMA A.1. *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.

Let i^* be the time step when i ceases to be alive (i.e. $i^* = i$ if 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 A.2. *B_t γ -dominates i at time t .*

Proof. By definition, all bidders in B_t are alive at t . We proceed by induction starting with the 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 A.1. 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: otherwise i^* would not bump $i \in X_{i^*}$.

In the inductive step, we assume that B_{t-1} γ -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 A.1). 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 $w(t) \geq (1 + \gamma)w(b)$.

Proof. [Proof of Lemma 5.1] Let V be the $\text{OPT}[\tilde{w}]$ 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 A.2 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{w} 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. $\tilde{w}(b) \geq \tilde{w}(i)$. 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.

B Proof of Theorem 5.3

For $c \in \mathbf{R}_+$, consider one item and a sequence of bids $\{a_k(c)\}_{1 \leq k \leq n}$ on it such that $a_1 = 1, a_2 = \frac{1}{c} > 1$ and $ca_{k+1}(c) = a_k(c) - \alpha \sum_{j=1}^{k-1} a_j(c) \forall k \geq 2$, implying

$$(2.4) \quad ca_{k+1} = (1+c)a_k - (1+\alpha)a_{k-1} \forall k \geq 2$$

For a fixed $n \geq 1$, we will look for a $c = c_n$ such that

$$(2.5) \quad a_n(c) - \alpha \sum_{j=1}^{n-1} a_j(c) = ca_n(c) \iff \frac{a_n}{1+\alpha} = a_{n-1}$$

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.

Proof. [Proof of Theorem 5.3] 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. (2.4). Now we claim that *whether or not* A accepts a_n , the competitive ratio will be at most c_n , which contradicts our assumption. If a_n is accepted, $\alpha \sum_{j=1}^{n-1} a_j$ has been lost due to bumping bidders

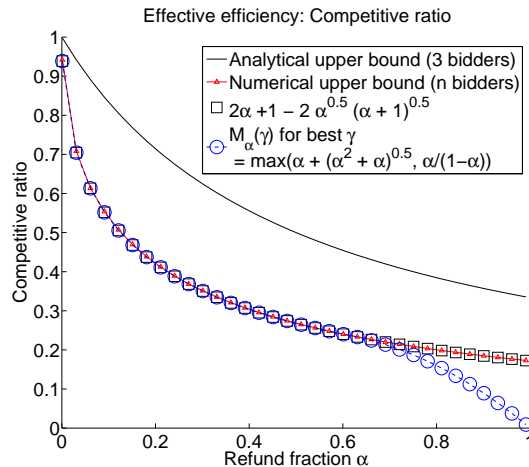


Figure 1: Effective efficiency competitive ratio (EECR) bounds as a function of α . The top curve is $c_3 = 1/(1+2\alpha)$. The middle curve is a numerical upper bound ($c = c_n$ of Eq. (2.5)) on any deterministic algorithm's EECR. The bottom curve shows (a lower bound on) our algorithm's EECR for the best γ_α : it matches the upper bound for $\alpha < 0.618$. γ is constrained by $\alpha < \gamma/(\gamma+1)$; if it were not, the bounds would match for all α .

$1, \dots, n-1$; if a_n is rejected the effective efficiency is $a_{n-1} - \alpha \sum_{j=1}^{n-2} a_j$. By Eqs. (2.4) and (2.5), both quantities are a c_n fraction of a_n , which in turn is at most M_n , the optimal (effective) efficiency.

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. (2.4) 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 there exists an increasing sequence of integers $\{n_i\}_{i \geq 1}$ such that a solution c_{n_i} to Eqs. (2.4) and (2.5) converges from above to $2\alpha+1-2\alpha^{0.5}(\alpha+1)^{0.5}$ as $i \rightarrow \infty$.

Let $\underline{u}(\gamma) = \frac{1}{1+\gamma} \left(1 - \frac{\alpha}{\gamma}\right)$, the competitive ratio from Corollary 5.1. 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. The top curve plots $c_3 = 1/(1+2\alpha)$.