Adaptive Limited-Supply Online Auctions

Mohammad T. Hajiaghayi
CSAIL
M.I.T.
Cambridge, USA
hajiagha@theory.lcs.mit.edu

Robert Kleinberg *
CSAIL
M.I.T.
Cambridge, USA
rdk@theory.lcs.mit.edu

David C. Parkes †
DEAS
Harvard University
Cambridge, USA
parkes@eecs.harvard.edu

ABSTRACT

We study a limited-supply online auction problem, in which an auctioneer has k goods to sell and bidders arrive and depart dynamically. We suppose that agent valuations are drawn independently from some unknown distribution and construct an adaptive auction that is nevertheless value- and time-strategyproof. For the $k = 1$ problem we have a strategyproof variant on the classic secretary problem. We present a 4-competitive (e-competitive) strategyproof online algorithm with respect to offline Vickrey for revenue (efficiency). We also show (in a model that slightly generalizes the assumption of independent valuations) that no mechanism can be better than $3/2$-competitive (2-competitive) for revenue (efficiency). Our general approach considers a *learning* phase followed by an accepting phase, and is careful to handle incentive issues for agents that span the two phases. We extend to the $k > 1$ case, by deriving strategyproof mechanisms which are constant-competitive for revenue and efficiency. Finally, we present some strategyproof competitive algorithms for the case in which adversary uses a distribution known to the mechanism.

Categories and Subject Descriptors

F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms

Algorithms, Economics

1. INTRODUCTION

Online auctions consider a setting in which agents arrive dynamically and require that an allocation and payment decision is made before they depart. As a motivating example, suppose that you are auctioning off last-minute theater tickets for *Proof* at 8pm. A traditional auction might set an auction deadline of 6pm and require that all bids received before that time are firm commitments and then determine the winners at 6pm. This can be suboptimal when bidders have different time constraints and decision horizons. In an online auction, a bidder could arrive at 12.10pm and state "I want one seat to Proof, and I will pay \$75." The same bidder could subsequently announce at 12.20pm that she needs to depart. The auctioneer is required to make dynamic allocation and payment decisions at or before the bidder's announced departure time, without the luxury of hindsight on bids yet to arrive.

The particular setting that we consider in this paper is a limited-supply online auction in which an auctioneer has k identical goods to sell. We assume that the number of agents is fixed, and known to the auctioneer. Each agent demands a single unit, and values are drawn independently from a fixed distribution which is stationary but unknown, i.e. the values are distributed i.i.d. Each agent also has an arrival and departure time, which may be arbitrary. The agent's departure time need not be known to the agent (or to the mechanism) until the moment of departure. We construct auctions that are both value- and time-strategyproof. In other words, an agent's dominant strategy is truthful reporting of its value for the good (*value-SP*), and an agent cannot benefit by misstating its arrival or departure time (*time-SP*). (This solution concept will be called simply *strategyproof* henceforth.) By the Revelation Principle [5], any outcome achievable in dominant-strategy equilibrium by any mechanism may be achieved by a strategyproof direct revelation mechanism, so we do not lose generality in restricting our attention to such mechanisms.

Our auctions compute an allocation whose efficiency is constant-competitive with the optimal allocation as determined by an omniscient auctioneer that knows the complete sequence of bids as it makes its decision. We also achieve constant-competitive revenue with such an omniscient auctioneer, according to definitions of revenue-competitiveness adopted by other authors in analogous works on auctions for identical goods [9, 11]. We note that without i.i.d., Lavi and Nisan [11] show that no constant-competitive auction is possible.

In the $k = 1$ case, our model generalizes offline auctions (the special case where everyone arrives at the beginning of time and departs at the end of time) and secretary problems (the special case where agents' arrival-departure intervals are disjoint). In both special cases, it is trivial to design a time-SP mechanism which is constant-competitive on efficiency. Vickrey's [14] second-price sealed-bid auction solves the offline case. The secretary algorithm [4] solves the disjoint interval online case. Our mechanism is constant-competitive and strategyproof in the general-case and actually reduces to the Vickrey and secretary algorithm in the special cases. Achieving time-strategyproofness in the secretary algorithm is nontrivial; the algorithm as ordinarily stated is manifestly not time-SP, because candidates with early arrival times always have an incentive to delay their participation.

We have an auction that achieves a competitive ratio of $e + o(1)$ for efficiency — thus matching the efficiency guarantee of

*Supported by a Fannie and John Hertz Foundation Fellowship.
†Supported in part by NSF grant IIS-0238147.

EC'04, May 17–20, 2004, New York, New York, USA.

the original, non-strategyproof secretary algorithm — and is $(e^2 + o(1))$-competitive for revenue. A modified version of this mechanism is $(4 + o(1))$-competitive for both efficiency and revenue. All of the aforementioned competitive ratios are in relation to the second-price sealed-bid auction, i.e. the offline Vickrey auction for a single item. These competitive ratio upper bounds are applicable in a somewhat stronger adversarial model than the independent-sampling hypothesis specified above (namely, an adversarial set of bids arriving in random order), and they are accompanied by lower bounds on the competitive ratios of any strategyproof mechanism in this stronger model. The lower bounds are 2 for efficiency and 3/2 for revenue; it is an open problem to close the gaps between these lower bounds and the upper bounds of e and 4. As an interesting counterpoint to these lower bounds, we provide a mechanism which is 1.724-competitive with respect to efficiency, in an alternative model where the bids are sampled from a distribution which is *known* to the seller. The mechanism is an adaptation of a secretary algorithm due to Gilbert and Mosteller [7].

There is a nice generalization of our method to the case of selling $k > 1$ items. We provide a strategyproof online auction whose efficiency is constant-competitive with respect to the offline efficiency-max-imizing allocation, and whose revenue is constant-competitive with $\mathcal{F}^{(2,k)}$, i.e. the best revenue one could achieve in an offline auction with a fixed price and selling at least two items and at most k items.

We believe this is the first discussion of adaptive online auctions that make no assumptions about the distribution from which agent valuations are drawn but provides both constant-competitiveness and time-SP. Our approach constructs two-phase online auctions, with a learning phase followed by a selling phase. We achieve strategyproofness even for agents with arrival-departure intervals that span both phases, by allowing allocations to those agents at the transition point with a decision criterion that is preferable to that available in the selling phase. We are also careful to make truthful reporting a strongly dominant strategy, even for early arrivals, through ensuring a small probability that the transition will occur at any time.

The outline of the paper is as follows. Section 2 formulates the problem we are discussing and enumerates the relevant variants. In Section 3, we discuss the relation of our problem to the secretary problem. Section 4 presents a general theorem to establish strategyproofness. Section 5 considers the single-item auction and Section 6 the $k > 1$ auction. Finally, we end with some open problems in Section 7.

1.1 Prior work

Lavi and Nisan [11] construct a value- and time-SP online auction for a multi-unit setting in which bidders can demand more than one copy of the good but have marginal-decreasing valuations. The bidder valuations are known to lie in an interval $[p, q]$ and the seller derives value p for each unit of good unsold. The auction is restricted to make an immediate decision for each arriving bid, and faces an adversary that can choose agents' values as well as their arrival order. Lavi and Nisan give an exact formula for the best-possible competitive ratio (versus Vickrey); asymptotically, their ratio grows like $\theta(\log(q/p))$.

Awerbuch et al. [2] retain the adversarial model of Lavi and Nisan and provide a method to convert any competitive online algorithm $\mathcal{A}$ into a value-SP online auction with a competitive ratio of $O(\rho + \log \mu)$ of the optimum offline profit obtained by an omniscient algorithm (that knows the true valuations of the users). Here ρ is the competitive ratio of $\mathcal{A}$ and μ is the ratio of the maximum to minimum valuation for a request. The auctions are time-SP for technical conditions on the optimization algorithm that provide increasing prices.

Prior work on learning in online auctions has not retained robustness against temporal manipulation. Blum et al. [3] adopt learning theory techniques to construct an online auction that is $(1 + \epsilon)$-competitive mechanism, for arbitrarily small positive ϵ, under a mild hypothesis. Their model is an unlimited supply, unit-demand setting, with bidders' valuations constrained to the interval $[1, h]$ and modeled adversarially. Kleinberg and Leighton [10] sharpen the Blum *et al.* results by determining the additive regret: it is $O(n^{2/3} \log(n)^{1/3})$, with a lower bound of $\Omega(n^{2/3})$, given n bidders. A revised lower bound of $\Omega(n^{1/2})$ is stated for a model in which bids are distributed according to some fixed but unknown distribution, along with an upper bound of $O((n \log n)^{1/2})$ with an additional technical hypothesis about the profit curve.

Friedman and Parkes [5] consider a more general model in which a mechanism makes a sequence of decisions over time. Agents dynamically arrive and depart, with types drawn from a distribution known to the mechanism. VCG-based online mechanisms are proved value- and time-SP if the problem allows a perfectly-competitive online decision policy. Parkes and Singh [12] extend this model and implement optimal MDP-policies within a VCG-based mechanism that supports immediate and truthful bidding in a Bayesian-Nash equilibrium. Gallien [6] considers a limited-supply model with a known distribution over agent values and arrivals and unit-demand, and constructs value- and time-SP mechanisms that maximize the expected time-discounted revenue.

Goldberg et al. [9] consider *competitive auctions* in offline settings (thus without any timing) which are truthful and yield profits that are within a constant factor of the profit of optimal fixed pricing. In addition, they show why optimal fixed pricing is a good measure for evaluating competitive auctions. They also present several randomized auctions and justify the use of randomness instead of determinism for competitive auctions. We adopt their DSOT auction to provide the basis for our analysis of the multi-item extension of our online auction problem.

2. MODEL AND PROBLEM FORMULATION

We are considering on-line auctions for k identical goods, in which each bidder wants at most one copy of the good, and the seller derives no utility from retaining copies of the good. The only information available to the mechanism is the total number of agents n and the time-horizon $[0, T]$. Agents have an arrival time, a departure time and a value for the good. We assume that agents have *no* value for receiving an item outside of their arrival-departure interval. This information is all private to an agent, and constitutes its *type*, $\theta_i = (v_i, a_i, d_i)$, which consists of:

- v_i, a positive real number representing the agent's value for one item.
- a_i, a real number in $[0, T]$ representing the agent's arrival time.
- d_i, a real number in $[a_i, T]$ representing the agent's departure time.

and Θ denotes the set of agent types. We consider that the agent learns about the auction, or its value, at time a_i and needs a decision by time d_i.

We consider a setting in which the agents values v_i are sampled i.i.d. from some stationary distribution that is not necessarily known to the seller, and consider an adversary that chooses arrival-departure intervals arbitrarily. We are interested in mechanisms which learn this distributional information (either implicitly or explicitly) over the course of interacting with the bidders. Although motivated by this i.i.d. world, we find it convenient to work in the following adversarial model. Formally, we assume the agents' types are generated as follows: an adversary specifies a (potentially random) set $\{v_1, \ldots, v_n\}$ of valuations, and a set $\{[a_1, d_1], [a_2, d_2], \ldots, [a_n, d_n]\}$ of arrival-departure intervals, and the valuations are matched to the arrival-departure intervals using a random permutation. We call this the *random ordering* hypothesis. We note that we have such a random ordering in the case that all valuations are distributed i.i.d., and our constant-competitive upper bounds therefore hold for the i.i.d. setting. However, we acknowledge that it may be possible to strengthen our upper-bounds with a full i.i.d. analysis. We will generally assume that the valuations are all different, an assumption that is generic in the case that the distribution has a density function. In fact, most of our mechanisms work without this assumption, however we make this assumption mainly for the ease of presentation.

In the online auction setting we consider agents that are game-theoretic and seek to place bids to maximize their individual utility in equilibrium. Although our auctions do not require an agent to declare its departure time until the moment of its departure, we find it convenient to analyze our auctions as direct-revelation mechanisms (DRMs). The strategyspace in an online DRM allows an agent to declare some possibly untruthful type, $\hat{\theta}_i = (\hat{v}_i, \hat{a}_i, \hat{d}_i)$, subject to $\hat{a}_i \geqslant a_i$ and $\hat{d}_i \geqslant \hat{a}_i$. Note that we assume that an agent *cannot* announce an earlier arrival time than its true arrival. Given the reported types $\hat{\theta} = (\hat{\theta}_1, \ldots, \hat{\theta}_n)$, the mechanism computes an allocation $q_i(t, \hat{\theta})$ and payment $p_i(\hat{\theta})$, for $i = 1, 2, \ldots, n$, and $0 \leqslant t \leqslant T$. We interpret $q_i(t, \hat{\theta})$ as the number of items allocated to agent i at time t (hence $q_i(t, \hat{\theta})$ is always either 0 or 1), and $p_i(\hat{\theta})$ as the total payment by agent i at its departure $\hat{d}_i$. We place the following natural condidtions on the allocation and payments:

1. Allocations cannot be revoked, so $q_i(t, \hat{\theta})$ is a non-decreasing function of t.
2. The allocation and payment to agent i must be determined before it departs.
3. The allocation and payments must be online computable, in that they can only depend on information available at time t.

We make the standard assumption that agent i has a quasilinear utility function $u_i = v_i - p$ for an item allocated between $a_i \leqslant t \leqslant d_i$ at price p, and that agents seek to maximize their expected utility. We consider mechanisms that satisfy:

participation $p_i(\hat{\theta}) = 0$ if $q_i(T, \hat{\theta}) = 0$, and $p_i(\hat{\theta}) \leqslant \hat{v}_i$.

no-deficit $p_i(\hat{\theta}) \geqslant 0$

We are interested in mechanisms with *dominant-strategy* equilibrium, such that every agent i has a single optimal strategy whatever the strategies and types of other agents. This is a particularly robust solution concept. In particular, and without loss of generality given our goal of dominant-strategy equilibrium [5], we consider strategyproof mechanisms. A mechanism is value- and time-strategyproof (or simply *strategyproof*) if an agent's dominant strategy is to report its true value (*value-SP*) and true arrival-departure interval (*time-SP*).

DEFINITION 1 (STRATEGYPROOF). *Mechanism $M = (\Theta; q, p)$ is strategyproof if for any agent i with type $\theta_i = (v_i, a_i, d_i)$, and for any vector θ_{-i} of types for the remaining agents, the utility of agent i when playing strategy θ_i against θ_{-i} is at least as great as the expected utility when playing any other strategy $\hat{\theta}_i = (\hat{v}_i, \hat{a}_i, \hat{d}_i)$ satisfying $a_i \leqslant \hat{a}_i \leqslant \hat{d}_i$. (If M is randomized, the utility is defined as the expected utility with respect to the mechanism's random choices.)*

We define the allocative-efficiency of an outcome of an online auction as

$$Eff(\theta) = \left(\sum_i q_i(T, \theta) v_i\right) \Big/ Eff_v$$

and the revenue as

$$Rev(\theta) = \sum_i p_i(\theta).$$

The benchmark offline mechanisms that we adopt for the purpose of competitive analysis are:

Vickrey efficiency $Eff_v = \sum_{i \leqslant k} v^{(i)}$, where $v^{(m)}$ is the m-th highest value. This is the total value achieved in an offline Vickrey auction, allocating to the k highest bidders and breaking ties arbitrarily.

$\mathcal{F}^{(\mathbf{2},\mathbf{k})}$ **revenue** is defined, following Goldberg et al. [9], as

$$\mathcal{F}^{(2,k)} = \max_{2 \leqslant l \leqslant k} l \cdot v^{(l)}$$

where again $v^{(m)}$ is the m-th highest bid. This is the maximal revenue achievable by a fixed-price mechanism in which the number of items sold is between 2 and k.

In the special-case of single-item auctions, where $k = 1$ and $\mathcal{F}^{(2,k)}$ is consequently undefined, we adopt the Vickrey revenue (i.e. $v^{(2)}$ in the notation above) as our benchmark. For both efficiency and revenue competitive ratios, we consider the *expected* value, averaged over both the mechanism's random choices (if any) and the probability distribution on agent types (i.e. the random valuations, and the random matching between these valuations and the arrival-departure intervals).

3. RELATION TO SECRETARY PROBLEM

The auction mechanisms considered in this paper are closely related to the classical *secretary problem* from the theory of optimal stopping. In the secretary problem, an interviewer interacts with a pool of n job applicants in random order. The applicants are ranked in terms of quality, but the interviewer only knows the relative ordering of the applicants whom she has already interviewed. After meeting each applicant, the interviewer has a choice between hiring the current applicant or sending him away; either decision, once made, is irrevocable. The problem is to design a strategy which maximizes the probability of hiring the most qualified candidate. It is well-known [4] that the optimal policy is to interview the first $t-1$ applicants, then hire the next one whose quality exceeds that of the first $t-1$ applicants, where t is defined by

$$\sum_{j=t+1}^{n} \frac{1}{j-1} \leq 1 < \sum_{j=t}^{n} \frac{1}{j-1},$$

As $n \to \infty$, the probability of hiring the best applicant approaches $1/e$, as does the ratio t/n.

The secretary problem and its solution suggest the following mechanism for an online single-item auction: do nothing until at least $t = \lfloor n/e \rfloor$ agents have arrived and announced their bids; then, set a reservation price p equal to the highest bid seen so far, and sell to the first agent whose bid is greater than or equal to p. Although this mechanism is not time-SP in general — the earliest-arriving agents have an incentive to delay announcing their arrival — it is clear that the mechanism is both value- and time-strategyproof *in the special case where the arrival-departure intervals are disjoint*, because each agent faces a price which is determined before that agent's arrival and does not vary during the time when the agent is in the market. The analysis of the secretary problem indicates that the mechanism's competitive ratio (with respect to efficiency) approaches e as $n \to \infty$. Although our model provides the seller with a richer information set than the secretary problem — the agents reveal the numerical values of their bids, not just their relative ordering — it seems plausible that there is no way to capitalize on this numerical information, since we are making no assumptions at all about the distribution of bids. Thus we suspect that the competitive ratio of $e + o(1)$ is best-possible, though the best lower bound we have been able to prove is 2.

Several natural questions now present themselves:

1. If the arrival-departure intervals are allowed to overlap, can we maintain strategyproofness while preserving the same competitive ratio?

2. What can be said about the competitive ratio with respect to revenue, rather than efficiency?

3. What can be said if the supply is finite but greater than 1?

The following sections will answer all three questions. Addressing the first and second questions, we will show that the mechanism specified above may be modified so that it becomes strategyproof in the case of overlapping arrival-departure intervals, and that a similar mechanism is 4-competitive in terms of revenue. The third question appears to invite an examination of so-called "multiple secretary problems," in which the interviewer is allowed to hire $k > 1$ applicants. Although many such problems have been proposed and analyzed in the optimal stopping literature [1, 8, 13, 15], we found that none of these results is directly relevant to the problem of designing a constant-competitive online mechanism for selling k identical items. Instead, we present a constant-competitive strategyproof mechanism based on the DSOT auction of [9]. The analysis of this mechanism's competitive ratio closely follows that of [9]; our contribution lies in specifying an online implementation of this mechanism which maintains temporal strategyproofness.

4. ESTABLISHING STRATEGYPROOFNESS

The following provides a useful constructive approach to validate the incentive properties of an online auction, both for the single-item and multiple item auction problem.

LEMMA 2. *An online auction is value- and time-strategyproof if the allocation and payment rules satisfy the following properties:*

a) The auction constructs a price schedule $p_i(t)$ to agent i in period t that is independent of the agent's reported type $\hat{\theta}_i$.

b) Let $p_{\min,i} = \min_{\hat{a}_i \leqslant t \leqslant \hat{d}_i} p_i(t)$, and m_i denote the period in which this price first occurs. The auction should allocate an item to agent i when $p_{\min,i} < \hat{v}_i$, at price $p_{\min,i}$, and in a period $m_i \leqslant t \leqslant \hat{d}_i$.

c) The auction never allocates an item to agent i when $p_{\min,i} > \hat{v}_i$.

PROOF. First, consider a simpler auction with an agent-indepedent price schedule in which the rules are to sell the agent the good in any period in which the minimal price occurs whenever $p_{\min,i} < \hat{v}_i$ and not otherwise. An agent's bid determines the purchasing decision, but not the price, and the agent's dominant strategy would be to report its true type. We relax this rule, and allow the item to be sold in any period between m_i and $\hat{d}_i$ (but still at price $p_{\min,i}$). This delay is critical, because otherwise the online auction would need to correctly anticipate whether or not prices were going to fall again in the future. The delayed purchasing decision is utility-maximizing with respect to a *perturbed* price schedule $\tilde{p}_i(t)$, defined with price equal to $p_i(t)$ for $t < m_i$ but defaulting to $p_{\min,i}$ for all $t \geqslant m_i$. Notice that agent i is indifferent across all allocations in periods $m_i \leqslant t \leqslant \hat{d}_i$, and so all that is left to prove is that the new dependence between $\tilde{p}_i(t)$ and the agent's reported type has no incentive effect. The perturbed price schedule has the following two important properties: a) an agent reporting its true type is indifferent between $\tilde{p}_i(t)$ and $p_i(t)$; b) the perturbed schedule is independent of $\hat{v}_i$, and an agent cannot improve the price in $\tilde{p}_i(t)$ between its true a_i and d_i by misstating $\hat{a}_i$ or $\hat{d}_i$. To see this robustness to $\hat{a}_i$ and $\hat{d}_i$, first observe that price $p_{\min,i}$ is (weakly) increasing for later arrivals $\hat{a}_i$ and earlier departures $\hat{d}_i$. Our model rules out earlier $\hat{a}_i$. Finally, although a later announced departure can reduce $p_{\min,i}$, the effect is to leave prices before the previous m_i unchanged and (weakly) increase the price in periods between m_i and d_i because this can only occur when the new minimal price occurs in some period $t > d_i$. □

The characterization allows for randomized price schedules, as long as the randomness is independent of an agent's reported type. Furthermore, the price schedules for different agents can be dependent on correlated random variables. This is helpful when simulating an auction in which two bids arrive at the exact same time and with the same bid price and the auction needs to break a tie. We will use this technique to handle ties when establishing strategyproofness in this paper.

The current axioms are sufficient for *weak* truthfulness, but not *strong* truthfulness. We have not ruled out the possibility that there are some types of agents that cannot receive positive utility from the mechanism. We care about this because we are interested in learning information about the distribution on agent types from agent reports and without strong truthfulness we cannot be sure that all types will report information.

REMARK 3. *We achieve strong truthfulness with one additional axiom:*

d) for any type θ_i, there are some types of other agents, such that there is some non-zero probability that agent i's price schedule provides price $p_{\min,i} < v_i$.

This additional property is sufficient to provide strong truthfulness, whatever the type of an agent (and assuming that an agent cannot observe the types of other agents), because it retains the possibility that an agent might receive (strictly) positive utility from truthful participation.

5. SINGLE-ITEM AUCTION

In this section, we present simple strategyproof online mechanisms for the $k = 1$ special case which are constant-competitive for revenue and efficiency with respect to the offline Vickrey auction.

5.1 No-information case

In the no-information case we consider our adversarial model, in which an adversary specifies a set of valuations that are randomly matched with arrival-departure intervals. Our competitive ratios also hold in the relaxed model in which an adversary specifies a distribution on valuations.

Let $M(j)$ denote the following mechanism, which depends on a parameter j taking integer values between 1 and n. In the special case when all bidders arrive at once then this auction reduces to the standard Vickrey auction. Let $b_{\leqslant m}$ denote the first m bids to arrive, and $b_{\leqslant m \setminus i}$ denote the first m bids without the bid from agent i (if it occured during the first m bids). Let $b^{(s)}_{\leqslant m}$ denote the s-th highest bid during the first m.

1. **First Phase:** Let τ be the arrival time of the j-th bidder, and let j' be the number of bidders seen up to and including time τ. (In case of simultaneous arrivals it is possible $j' > j$.) Let $p_1 = b^{(1)}_{\leqslant j'}$ and $p_2 = b^{(2)}_{\leqslant j'}$. If any agent bidding $\geqslant p_1$ is still present then sell to that agent (breaking ties arbitrarily) at price p_2.

2. **Second Phase:** Otherwise, set a reserve price of p_1, and sell to the first agent who submits a bid $\geqslant p_1$, at price p_1, breaking ties between agents that arrive at the same time at random.

LEMMA 4. *The online mechanism $M(j)$ is strategyproof.*

PROOF. First, assume that there are no ties, either at the end of the first phase or during the second phase. In this case, construct the following price schedule for agent i. We refer to a bid from an agent $\neq i$ as an "outside bid." The price schedule changes whenever a new bid arrives, as follows:

$$p_i(t) = \begin{cases} \infty & \text{, for } < j-1 \text{ outside bids} \\ b^{(1)}_{\leqslant j-1 \setminus i} & \text{, for } j-1 \text{ outside bids} \\ b^{(1)}_{\leqslant j \setminus i} & \text{, for } \geqslant j \text{ outside bids, until some bid } \beta \\ & \quad \text{with } \hat{v}_\beta \geqslant b^{(1)}_{\leqslant j \setminus i} \text{ is present} \\ \infty & \text{, otherwise.} \end{cases}$$

Bid β represents a bid from an agent $\neq i$ that arrives as, or after, the j-th bidder and bids at least ask price p_1. This bid will win, unless agent i bids above its scheduled price before this event.

Let $p_{\min,i} = \min_{\hat{a}_i \leqslant t \leqslant \hat{d}_i} p_i(t)$, and m_i denote the period in which this price first occurs. Consider three cases.

Case 1. Agent i is present after $j-1$ outside bids and $b^{(1)}_{\leqslant j-1 \setminus i} = p_{\min,i} < \hat{v}_i$. So, agent i is present at the end of the first phase and the highest bidder and receives the item at price $p_2 = b^{(1)}_{\leqslant j-1 \setminus i}$.

Case 2. Agent i is present after $j-1$ outside bids but $b^{(1)}_{\leqslant j-1 \setminus i} > \hat{v}_i$. We have $p_{\min,i} > \hat{v}_i$, and $\hat{v}_i < p_1$ and the agent receives no allocation.

Case 3. Agent i arrives after j outside bids and before another agent $j' \neq i$ with $\hat{v}_{j'} \geqslant b^{(1)}_{\leqslant j \setminus i}$ arrives. Agent i bids $\hat{v}_i > b^{(1)}_{\leqslant j \setminus i}$. So, agent i is the first to bid $\geqslant p_1$ in the second phase and receives the item at price $p_1 = b^{(1)}_{\leqslant j \setminus i} = p_{\min,i}$.

In all other cases it is clear that agent i's bid value $\hat{v}_i < p_{\min,i}$, and that the agent receives no allocation in the auction.

Ties at the end of phase one present no problem, because $p_2 = p_1 = \hat{v}_i = p_{\min,i}$ for all tied agents and agents are indifferent between winning or losing. We construct a randomized price schedule to handle ties in the second phase. Assume that each agent has a priority $\pi_i \in \{1, \ldots, n\}$, with priorities randomly permuted in each period. Assume that tie-breaking in $M(j)$ respects $\pi = (\pi_1, \ldots, \pi_n)$, such that the winning bidder is the one with the greatest priority in that period.

We can construct a corresponding price schedule for agent i as follows. In period t' in which one or more outside bids arrive with value $\geqslant p_1$, then if π_i is greater than all priorities of these outside bids then $p_i(t')$ remains at price p_1 and ticks up to ∞ at time $t' + \epsilon$ for some small $\epsilon > 0$ (smaller than the minimal interarrival time between agents). Otherwise, if an outside bid with value $\geqslant p_1$ has higher priority, agent i's price is immediately set to ∞. The effect is that (a) if only one bid arrives (no ties) with value $\geqslant p_1$ then its price remains at p_1 in its arrival period and it gets the good; (b) if multiple bids arrive with value $\geqslant p_1$ then the price stays low for the agent with the highest priority, and this agent gets the good. □

Theorem 5 shows the competitiveness of the mechanism. We assume for now that all arrivals are unique and can ignore ties. This case is handled after the main proof.

THEOREM 5. *The mechanism $M(\lfloor n/2 \rfloor)$ is a $4+o(1)$-competitive strategyproof online algorithm with respect to offline Vickrey for both revenue and efficiency.*

PROOF. We consider two cases as follows.

The item sells at time τ: In this case, the highest-bidding agent among the first $\lfloor n/2 \rfloor$ is still in the market at time τ. (Because the mechanism is defined so that it only sells the item at time τ if the highest-bidding agent among the first $n/2$ is still present.) With probability $1/2 - O(1/n)$, this agent is the highest-bidding among the entire population. (In fact, the probability is exactly $1/2$ unless n is an odd number.) Conditional on this event, the probability that the second-highest bidder is also among the first half is $1/2 - O(1/n)$. (Here we have term $O(1/n)$, since there are only $\lfloor n/2 \rfloor - 1$ agents remaining in the first half, and $\lceil n/2 \rceil$ in the second half.) So the probability of selling to the highest bidder at the second-highest price is at least $1/4 - O(1/n)$.

The item does not sell at time τ: In this case, the probability that the highest-bidding agent is among the second half of the population is at least $1/2 - O(1/n)$ (again the probability is exactly $1/2$ unless n is an odd number.) Conditional on this event, the probability that the second-highest bid was received at or before time τ is at least $1/2 - O(1/n)$. If both these events occur (as happens with probability at least $1/4 - O(1/n)$) the item will be sold to the highest bidder at the second-highest price.

We note that in both cases we use the fact that the event *the item sells at time τ* is independent of the pair of events *the highest bidder arrives at or before τ* and *the second highest bidder arrives at or before τ*. This is because the first event is entirely determined by the relative ordering of the bidders in the first half of the population, while the other two events do not depend on this ordering but only on how the population is partitioned into the first half and second half. Therefore in expectation, our revenue (efficiency) is at least $1/4 - O(1/n)$ fraction of that of the offline Vickrey mechanism. □

We need our assumption of unique valuations, generic for value distributions with a density, for efficiency-competitiveness. Without this assumption, there are examples (e.g. one highest bid and $n-1$ second-highest bids) in which the item, if sold in the second phase, is very likely to be sold to a second-highest bidder rather than to the highest bidder.

REMARK 6. *Introducing ties in arrival times, the analysis in the second phase remains valid because the case in which both the event that the second-highest bid is in the first phase and the highest bid is in the second phase already precludes the ability for joint arrivals with bids $\geqslant p_1$. A simple argument shows that the efficiency and revenue for the case of multiple simultaneous arrivals at the end of phase one is at least that of $M(\lfloor n/2 \rfloor)$ for unique arrivals and any arbitrary ordering of the arrivals. If the auction sells to the highest bid in these simultaneous arrivals then given any arbitrary ordering the same bid or some earlier bid wins, and for a price that is dominated by the second-price in the simultaneous case. A similar argument shows that if no sale is made at the end of phase one with simultaneous arrivals, then no sale is made at the end of phase one or to any of these bids for any arbitary reordering.*

It is worth mentioning that the current mechanism $M(\lfloor n/2 \rfloor)$ is not strongly truthful, e.g., the people who leave before time τ are guaranteed not to win the auction, and thus being truthful or not has the same outcome. We can make the mechanism strongly truthful as follows.

COROLLARY 7. *There is a mechanism which is a $4 + o(1)$-competitive strongly strategyproof online algorithm with respect to offline Vickrey for both revenue and efficiency.*

PROOF. Consider the following mechanism M' in which with probability $1 - 1/n$ we run the mechanism $M(\lfloor n/2 \rfloor)$ and with uniform probability $1/n^2$ we run mechanism $M(j)$, $1 \leqslant j \leqslant n$. We note that since with positive probability every agent is present at time τ (and thus participates in the Vickrey auction at that time), which confirms the additional axiom specified in Remark 3 and establishes that the mechanism is strongly truthful. $\square$

In fact, we can improve the competitive ratio for efficiency by increasing the competitive ratio for revenue.

THEOREM 8. *The mechanism $M(\lfloor n/e \rfloor)$ is a strategyproof online algorithm which is $e + o(1)$-competitive for efficiency and $e^2 + o(1)$-competitive for revenue with respect to offline Vickrey.*

PROOF. The proof is the same as the proof of Theorem 5. For the competitiveness again we consider two cases. If the item sells at time τ, then with probability $\lfloor n/e \rfloor / n = 1/e - o(1)$ the agent that buys the item is the highest-bidding among the entire population. If the item does not sell at time τ, according to the secretary problem, with probability $1/e - o(1)$, we sell in the second phase to the highest-bidding agent. Hence, our expected efficiency is at least $1/e - O(1/n)$ of the efficiency of the offline Vickrey mechanism. For revenue, with probability $(1/e)^2 - O(1/n)$, we sell the item to the highest bidder at the second-highest price in the first case. We probability $(1/e)(1-1/e) - O(1/n)$, we sell the item to the highest bidder at the second-highest price in the second case. Hence our expected revenue is at least $(1/e)^2 - O(1/n)$ of the offline Vickrey revenue. $\square$

The mechanism specified in Theorem 8 is not strongly truthful, but we can make it strongly truthful using the same idea as above in Corollary 7.

5.2 Lower bounds in the random-ordering model

In the secretary problem, it is well-known that no stopping rule achieves asymptotic success probability greater than $1/e$. However, this does not immediately translate into a lower bound of e for the competitive ratio of online single-item auction mechanisms, because:

1. The mechanism has richer feedback — it sees the numerical values of the bids, not just their relative ordering.
2. Allocations in which the item is awarded to a bidder who is not the highest, but still has a positive valuation for the item, contribute positively to the expected efficiency of the mechanism.

It turns out that the second problem is easy to fix — one considers a bid set in which the highest bid exceeds the others by such a large factor that the expected efficiency ratio of the allocation is determined, to first order, by the probability of awarding the item to the highest bidder. However, addressing the first issue appears technically challenging, and the best lower bound we know of is:

THEOREM 9. *For any constant $\varepsilon > 0$, there is no online single-item auction mechanism which is $(2-\varepsilon)$-competitive for efficiency, in the random-ordering model.*

The foregoing discussion implies nothing about lower bounds on competitive ratios for revenue. (The trouble is that revenue competitive ratio is determined by comparing with the second-highest bid rather than the highest.) We can modify the proof of Theorem 9 to obtain:

THEOREM 10. *For any constant $\varepsilon > 0$, there is no strategyproof online single-item auction mechanism which is $(3/2 - \varepsilon)$-competitive for revenue, in the random-ordering model.*

It is interesting to note that the hypothesis of strategyproofness is required in the lower bound for revenue (Theorem 10) but not for efficiency (Theorem 9).

Our proof of Theorem 9 hinges on constructing a probability distribution on bid sets, such that no online stopping rule has probability better than $1/2 - \varepsilon$ of stopping on the highest bid. (The stopping rule is allowed to base its decisions on the numerical values of the bids it has seen, not just on their relative ranking.) Given an arbitrary positive integer n and a real number $\varepsilon > 0$, put $N = \lceil (10/\varepsilon) n \rceil$, and let D denote the uniform distribution on sets of $n+1$ *consecutive* integers in $\{0, 1, \ldots, N\}$. In other words, a random sample S from D is generated by picking an element u of the set $\{n, n+1, \ldots, N\}$ uniformly at random, and putting $S = \{u-n, u-n+1, \ldots, u\}$.

PROPOSITION 11. *Let R be any online stopping rule. If a bid set S is sampled according to D and the elements of S are presented to R in random order, then*

$$\Pr(R \text{ stops on } \max(S)) < \frac{1}{2} + \frac{\varepsilon}{10}.$$

PROOF. Let $S = \{u-n, u-n+1, \ldots, u\}$. At any point in time as the bids in S are being presented to R, define the *state* of the system to be the triple (x_*, x^*, β), where x_* is the lowest bid presented so far, x^* is the highest bid presented so far, and $\beta = 1$ if x^* was presented more recently than x_*, 0 otherwise. The pair (x_*, x^*) captures all the information influencing the distribution of S; in other words, if $x_1, \ldots, x_j$ are the first j bids and (x_*, x^*, β) is the state at the end of step j, then for any $y \in \{n, n+1, \ldots, N\}$,

$$\Pr(u = y \mid x_1, \ldots, x_j) = \Pr(u = y \mid x_*, x^*). \tag{1}$$

We can think of the states (x_*, x^*, β) as states of a Markov chain. States in which $x^* - x_* = n$ are absorbing states, i.e. there are no state transitions out of such a state. An absorbing state in which $\beta = 1$ is called a winning state, and an absorbing state in which $\beta = 0$ is called a losing state.

LEMMA 12. *If $x_* > 0$, $x^* < N$, and $x^* - x_* < n$, then* $\Pr(u = x^* \mid x_*, x^*) \leq 1/2$.

PROOF. If $x^* < n$ then $\Pr(u = x^* \mid x_*, x^*) = 0$. If $n \leq x^* < N$, then the events $u = x^*$ and $u = x^* + 1$ are both possible, and are therefore equiprobable. Thus $\Pr(u = x^* \mid x_*, x^*) \leq 1/2$. $\square$

LEMMA 13. *If $x^* > n$, $x_* < N - n$, and $x^* - x_* < n$, then*

$$\Pr(\textit{next state is a winning state} \mid x_*, x^*) = \Pr(\textit{next state is a losing state} \mid x_*, x^*).$$

PROOF. The next state is a winning state if and only if $u = x_* + n$ and the next sample outside $[x_*, x^*]$ is u. The next state is a losing state if and only if $u = x^*$ and the next sample outside $[x_*, x^*]$ is $u - n$. The hypotheses of the lemma guarantee that both events are possible; by symmetry, they have the same probability. $\square$

LEMMA 14. *If $x^* > n, x_* < N - n$, and $x^* - x_* < n$, then* $\Pr(R \textit{ stops on } u \mid x_*, x^*) \leq 1/2$.

PROOF. The proof is by downward induction on $x^* - x_*$. If $x^* - x_* = n$ then the statement is vacuously true. Otherwise, if R stops in the current state, then Lemma 12 ensures that the probability of stopping on u is at most 1/2. If R does not stop in the current state but stops before the next state change, then it is assured of not stopping on u. (Because $u \geq x^*$, and all the bids seen before the next state change are $< x^*$.) So assume that there is at least one more state change after (x_*, x^*, β). If it is an absorbing state, then the probability that it is a winning state is only 1/2. (Lemma 13.) If it is not an absorbing state, then it is another state $(\hat{x}_*, \hat{x}^*, \hat{\beta})$ satisfying the criteria of Lemma 14, but with $\hat{x}^* - \hat{x}_* > x^* - x_*$. By the induction hypothesis, $\Pr(R \text{ stops on } u \mid x_*, x^*) \leq 1/2$ in this case. $\square$

The proof of Proposition 11 is now easy. Let x_1 denote the first bid. Then

$$\Pr(n < x_1 < N - n) > 1 - 2n/N \geq 1 - \varepsilon/5,$$

and conditional on $n < x_1 < N - n$, Lemma 14 tells us that R has at best a 1/2 chance of stopping on $\max(S)$. Even if R is assured of stopping on $\max(S)$ in the other cases, its overall probability of success is at best $1/2(1 - \varepsilon/5) + \varepsilon/5 = 1/2 + \varepsilon/10$.

PROOF OF THEOREM 9. Let $S = \{x_1, \ldots, x_n\}$ be sampled according to the distribution D specified in Proposition 11. Transform this into a bid set $b = (v_1, \ldots, v_n)$ by putting

$$v_i = (10/\varepsilon)^{x_i}.$$

Assume, as usual, that $v_1 > v_2 \ldots > v_n$; in this particular case, it implies that all bids besides v_1 are $\leq (\varepsilon/10)v_1$. Pick disjoint arrival-departure intervals for the agents, so that each time the mechanism encounters a new agent, it must decide whether or not to allocate the item to this agent without the benefit of seeing any future bids. In other words, the mechanism must determine the allocation according to an online stopping rule.

Let Y be the random variable which measures the mechanism's efficiency divided by v_1; note that $0 \leq Y \leq 1$. No online mechanism has better than $1/2 + \varepsilon/10$ chance of allocating the item to the highest bidder. (Otherwise such a mechanism could easily be transformed into a stopping rule R which violates the conclusion of Proposition 11.) If the item is allocated to any other bidder, then $Y \leq \varepsilon/10$. Therefore the expectation of Y, over the random choice of bid set S as well as the random ordering of that set, is at most

$$(1/2 + \varepsilon/10) \cdot 1 + (1/2 - \varepsilon/10) \cdot (\varepsilon/10) < 1/2 + \varepsilon/5.$$

There must be some particular bid set S such that $\mathbb{E}[Y \mid S] < 1/2 + \varepsilon/5$. On this bid set, the competitive ratio of the mechanism is at least $(1/2 + \varepsilon/5)^{-1} > 2 - \varepsilon$. $\square$

5.3 Full-information case

We find it interesting to take a brief detour and consider the *full-information* case, in which an adversary picks a distribution on valuations that is *known* to the mechanism. In this case there is a strategyproof online mechanism that approaches 1.724-competitive for efficiency as $n \to \infty$, in contrast to the e-competitive mechanism in the no-information case. This raises an interesting question, that we leave unanswered in this paper, as to whether it is possible to design an adaptive mechanism in the no-information case that can perform better than the $M(j)$ family of mechanisms against an adversary that chooses a distribution on valuations.

We adapt the methods presented by Gilbert and Mosteller [7] for the *full-information secretary* problem. Our innovation is to introduce a method to make their schemes value- and time-SP,[1] without diminishing the probability of allocating the good to the highest bidder. (It is proved in [7] that no online stopping rule can achieve a higher success probability than their scheme.)

We can assume that the distribution on values v_i is uniform on $[0, 1]$. This assumption is without loss of generality for distributions having a continuous cumulative distribution function $F(x)$, because the largest element in a sample remains the largest under all monotonic transformations, and the mapping $x \mapsto F(x)$ transforms the given distribution into the uniform distribution on $[0, 1]$.

Let $M(r_1, \ldots, r_n)$ denote the following mechanism, which is parameterized with a fixed *price schedule*, $(r_1, \ldots, r_n)$, with $0 \leqslant r_i \leqslant 1$:

1. **Allocation Phase:** In each period in which one or more new bidders arrive, let $p(t) = \max(r_j, b^{(2)}_{\leqslant j})$, where j bids have arrived in total. Continue until an agent, j^*, bidding $> p$ is present. (Note, we allow multiple agents to arrive at the same time, but can ignore ties in value given uniform [0,1] values.) Agent j^* will be the winner. Initialize the price to this agent, $p^* = p(t)$.

2. **Pricing Phase.** Keep the auction open until j^* departs. In each period in which one or more new bidders arrive, adjust price p^* down to $\max(r_{j'}, b^{(2)}_{\leqslant j'})$ (if this new price is smaller than the current price), where j' bids have arrived in total.

The auction closes when bidder j^* leaves, at which point this agent is allocated the good for the adjusted price p^*.

The analysis of Gilbert and Mosteller [7] provides a method to set the price schedule $r_1, \ldots, r_n$ to maximize the expected probability of selecting the winner. Optimal schedules have *decreasing* prices, since the idea is to accept the current draw with value v_j if the probability of winning with this draw is more likely than the probability that a draw with a larger value will occur in the future and be chosen. This latter probability decreases with the number of remaining draws. Intuitively, in the final round we should take any bid that is the highest bid seen, while in the penultimate round we should take a bid if it is the highest seen so far and $\geqslant 0.5$ because that is the probability that the highest bid will be seen in the final round. For instance, given $n = 5$, the sequence of optimal prices r^* is, $r_1^* = 0.825, r_2^* = 0.776, r_3^* = 0.690, r_4^* = 0.500, r_5^* = 0$. It is these decreasing prices that make time-SP an interesting question.

LEMMA 15. *The mechanism family $M(r_1, \ldots, r_n)$ is strategyproof.*

PROOF. Let $b_{\leqslant j}(t)$ denote any of the first j bids to arrive that are *still present* at time t, and $b^{(1)}_{\leqslant j}(t)$ include any bid that is maximal across the bids seen in the first j bids and still present at time t. We'll construct the price schedule for agent i, and refer to bids from $\neq i$ as "outside" bids. Let bid β denote the first outside bid (if any) for which:

$$\hat{v}_\beta > \max(r_{j\setminus i}, b^{(2)}_{\leqslant j\setminus i})$$

where j bids have arrived in total in a period while bid β is present. Call this the "trigger" bid. The price schedule to agent i is constructed as:

$$p_i(t) = \begin{cases} r_1 & \text{, before any bid} \neq i \\ \max(r_{(j\setminus i)+1}, b^{(1)}_{\leqslant j\setminus i}) & \text{, after } j \text{ outside bids but no trigger} \\ \hat{v}_\beta & \text{, at trigger, until period } \hat{a}_\beta + \epsilon \\ \infty & \text{, otherwise} \end{cases}$$

where $r_{(j\setminus i)+1}$ is the j-th price from schedule r when i is in the first j-th arrivals, and the $(j + 1)$-st price otherwise. Let $p_{\min,i} = \min_{\hat{a}_i \leqslant t \leqslant \hat{d}_i} p_i(t)$.

First, we show that if $p_{\min,i} < \hat{v}_i$ that agent i is the winner. Consider three cases: **Case (i).** Schedule price, $r_1 < \hat{v}_i$, and agent i is the first to arrive. This is also the price the agent faces in the auction, and the agent is allocated the item upon arrival. **Case (ii).** Suppose m bids have been received, but there was no trigger bid. If $\max(r_m, b^{(1)}_{\leqslant m\setminus i}) < \hat{v}_i$ while agent i is present then $p_{\min,i} < \hat{v}_i$. Together with the lack of a trigger, this condition implies that auction price $p_i(t) = \max(r_m, b^{(2)}_m) < \hat{v}_i$, and agent i is allocated the item. **Case (iii).** Suppose m bids have been received in total, including a trigger bid and agent i is present on arrival of the trigger β, and that $\hat{v}_\beta < \hat{v}_i$. Again, $p_{\min,i} < \hat{v}_i$ and we have $p_i(t) = \max(r_m, b^{(2)}_m) < \hat{v}_i$ in this period since the trigger implies $\hat{v}_i > \hat{v}_\beta > r_m$.

[1] In fact, Gilbert and Mosteller also present a 1+ϵ-competitive scheme in the limit as $n \to \infty$ for the variant of the full-information secretary's problem in which the goal is to maximize the expected value of the choice and not the probability of selecting the winner. We omit this scheme because it is unclear to us whether this result holds against an adversary able to select arbitrary distributions.

Second, we show that if $p_{\min,i} > \hat{v}_i$ then the item is not allocated. Assume initially that agent i is not the first to bid and that the trigger occurs (if at all) after agent i departs. Then, $p_{\min,i} > \hat{v}_i$ requires $\hat{v}_i < \max(r_j, b^{(1)}_{\leqslant j \setminus i})$ in all periods $\hat{a}_i \leqslant t \leqslant \hat{d}_i$, where j denotes the number of bids that have arrived. By contradiction, if agent i is a winner then $\hat{v}_i \geqslant \max(r_j, b^{(2)}_{\leqslant j})$ in one of these periods, but this implies $\hat{v}_i \geqslant \max(r_j, b^{(1)}_{\leqslant j \setminus i})$. By a similar argument, we need $r_1 > \hat{v}_i$ if agent i is the first to bid and $\hat{v}_\beta > \hat{v}_i$ if the trigger occurs while agent i is present, from which agent i is not a winner in the auction.

Third, we show that if agent i is allocated the item then its payment is $p_{\min,i}$. Price $p_i(t)$ is equal to the price schedule before the trigger, and p^* is initially set to the minimal price between its arrival $\hat{a}_i$ and the end of the allocation phase. Then, p^* in the auction tracks the minimal price in schedule $p_i(t)$ because it follows the price agent i would have faced if it had delayed its arrival but still submitted a bid early enough to beat a potentially winning bid from an outside bidder. □

We show the competitiveness of the mechanism by demonstrating its performance relative to the scheme $M_{\mathrm{GM}}(r^*_1, \ldots, r^*_n)$ described by Gilbert and Mosteller, which assumes disjoint arrival-departure intervals, defines an optimal price schedule r^*, and works as follows:

1. At the arrival time of the j-th bidder, let $p = \max(r^*_j, b^{(2)}_{\leqslant j})$. Continue until a bid $> p$ is received, and make this agent the winner.

THEOREM 16. *[7] Algorithm $M_{\mathrm{GM}}(r^*)$ selects the item with the highest value with an asymptotic probability 0.5802..., for large numbers of items n and an appropriately chosen price schedule.*

COROLLARY 17. *Mechanism $M(r^*_1, \ldots, r^*_n)$ is a 1.724- competitive strategyproof online algorithm with respect to offline-Vickrey for efficiency.*

PROOF. The proof is by comparison with mechanism $M_{\mathrm{GM}}(r^*)$, which is 1.724-competitive but not strategyproof. Equivalently, we show that the probability of $M(r^*)$ selecting the highest bid is at least as good as that of $M_{\mathrm{GM}}(r^*)$. Note first that $M(r^*)$ will never accept a later bid than $M_{\mathrm{GM}}(r^*)$, because any time a bid is successful in $M_{\mathrm{GM}}(r^*)$, it is the highest bid so far and satisfies the price schedule test in $M(r^*)$. Next, in any period t in which both mechanisms take the same decision, the probability of success is the same. Finally, consider the case in which $M(r^*)$ selects a bid when $M_{\mathrm{GM}}(r^*)$ passes. We claim that this bid is more likely to be the highest bid than the probability of subsequently selecting the highest bid by following the $M_{\mathrm{GM}}(r^*)$ strategy. To see this, let j denote the period in which this occurs, and r^*_j denote the GM threshold. Let $x_j > r^*_j$ denote the value of the bid that is accepted by $M(r^*)$. The threshold value, r^*_j, is set so that if there are $(n - j)$ additional draws to make i.i.d. from $U[0, 1]$ and the GM mechanism is holding a sample x_j of value exactly r^*_j and x_j is the maximal value seen so far, then the probablity that x_j will be maximal across all n samples is equal to the probability that $M_{\mathrm{GM}}(r^*)$ will select the maximal bid in the future. Thus, if $x_j > r^*_j$ then monotonicity argues that the probability that x_j will be the highest bid is *greater* than the probability that $M_{\mathrm{GM}}(r^*)$ will select the highest bid in the future. Notice that it does not matter how we come to hold the current best sample (i.e. it does not matter whether we have been able to hold onto this sample for a number of periods). □

Finally, we note that the optimal price schedule as defined by Gilbert and Mosteller is set assuming the adversarial setting of disjoint arrival-departure intervals and leaves open the possibility of stronger upper-bounds for a more nuanced analysis of the degree of overlap between agents' intervals.

6. MULTI-ITEM AUCTION $(K > 1)$

In this section, we consider the case in which the auctioneer has k identical items. We will assume without loss of generality that $k \leqslant n$, since it is impossible to sell more than n items when each agent demands only one. Note that it is trivial to obtain a competitive ratio of ek for efficiency or $4k$ for revenue: simply throw away $k - 1$ of the items, and sell the remaining one using one of the single-item mechanisms described in Section 5.

On the other hand, it is not obvious how to obtain a competitive ratio which is constant (independent of k) as $k \to \infty$. Mechanisms which achieve this goal of constant competitive ratio are presented below. Although these mechanisms generalize the single-item mechanisms presented above, the analysis will be quite different. While our analysis of the single-item mechanisms was based on the secretary problem, the basis for our analysis of multi-item mechanisms will be the DSOT auction of [9].

First, we present a mechanism which is C-competitive (for a constant $C < 6338$) with respect to $\mathcal{F}^{(2,k)}$ for revenue. It is worth mentioning that we consider $\mathcal{F}^{(2,k)}$ instead of $\mathcal{F}^{(2)}$ (the optimal revenue for a fixed-price *unlimited-supply* mechanism which sells at least two items), since otherwise the competitive ratio is at least n/k, e.g., when all valuations are identical. The reason that we use $\mathcal{F}^{(2,k)}$ (instead of $\mathcal{F}^{(1,k)}$, the optimal revenue for a fixed-price mechanism) is the same as the reason mentioned in [9] and hence omitted.

The description of the mechanism requires the following notation. Suppose b is any set of real-valued bids. Let $v_1 \geq v_2 \geq \ldots \geq v_n$ be the bids in b, arranged in descending order. For any price p, let $n_p(b)$ denote the number of bids in b which are greater than or equal to p. For a positive integer s, define $F^s_p(b)$ to be the revenue obtained by selling at most s items at price p, i.e.

$$F^s_p(b) = p \cdot \min\{s, n_p(b)\}.$$

Define

$$opt^s(b) = \arg\max_p (F^s_p(b)),$$

i.e. $opt^s(b)$ is the price which maximizes revenue subject to the constraint that at most s items are sold.

Let $M_k(j)$ denote the following mechanism, generalizing the one mentioned in the previous section.

1. Let $s = \lceil k/2 \rceil$. Let τ be the arrival time of the j-th bidder, and let j' be the number of bidders seen up to and including time τ. (In case of simultaneous arrivals, it is possible that $j' > j$.) Of the bidders who arrived at time τ, pick a random subset of $j' - j$ of them, and mark these bidders as *inactive*. Define every other bidder as either *active* if they arrived at or before τ and have not yet departed, or *departed* otherwise; there are exactly j agents who are active or departed. Set p equal to the $(s+1)$st price over the bids from these j agents. (Or $p = 0$ if $j \leqslant s$.) Sell at price p to any bidder still present. If fewer than s items are sold at time τ, discard items until the remaining supply is $\lfloor k/2 \rfloor$.

2. Let b' denote the set of j active bids or departed bids that we have seen so far. Set a reserve price $p' = opt^s(b')$. For bidders arriving after the j-th (breaking ties due to simultaneous arrivals at random), sell at price p' if their bid is $\geqslant p'$, until the mechanism runs out of items or bidders.

Let RM_k be $M_k(j)$ where j is number of heads in n tosses of a fair coin, i.e., j is drawn from the binomial distribution $B(n, 1/2)$.

THEOREM 18. *RM_k is a strategyproof online mechanism which is C-competitive with $\mathcal{F}^{(2,k)}$ for revenue, for a constant $C < 6338$.*

PROOF. To show that mechanism $M_k(j)$ is strategyproof for any j we'll again construct an agent-independent price schedule and show utility-maximization with respect to the schedule. To keep things simple we'll assume that no two agents arrive at the same time and provide a short proof sketch. Construct the price schedule to agent i as:

$$p_i(t) = \begin{cases} \infty & \text{, for } < j-1 \text{ outside bids} \\ b^{(s)}_{\leqslant j-1\setminus i} & \text{, for } j-1 \text{ outside bids} \\ opt^s(b_{\leqslant j\setminus i}) & \text{, after } j \text{ outside bids,} \\ & \text{ before } \lfloor k/2 \rfloor \text{ additional outside bids} \\ & \text{ with value } \geqslant opt^s(b_{\leqslant j\setminus i}) \text{ arrive} \\ \infty & \text{, otherwise.} \end{cases}$$

Again, let $p_{\min,i} = \min_{\hat{a}_i \leqslant t \leqslant \hat{d}_i} p_i(t)$, with m_i denoting the period in which this minimal price occurs. Consider the following cases.

Case 1. Agent i is present after $j-1$ outside bids, and its bid price $\hat{v}_i > b^{(s)}_{\leqslant j-1\setminus i}$. The price at the end of phase one is also equal to $b^{(s)}_{\leqslant j-1\setminus i}$ in this case, and the agent wins an item. Moreover, this is the minimal price over $[\hat{a}_i, \hat{d}_i]$ because $opt^s(b_{\leqslant j\setminus i}) \geqslant b^{(s)}_{\leqslant j\setminus i} \geqslant b^{(s)}_{\leqslant j-1\setminus i}$, since $\text{opt}^s(b) \geqslant b^{(s)}$ for any set of bids because price $b^{(s)}$ sells every item, and a lower price than $b^{(s)}$ can only reduce the revenue received for every item sold, but without increasing the quantity sold.

Case 2. Agent i is present after $j-1$ outside bids but its bid price $\hat{v}_i < b^{(s)}_{\leqslant j-1\setminus i}$. In this case the agent neither wins at the end of the first phase or during the second phase.

Case 3. Agent i's bid arrives after j outside bids and before an additional $\lfloor k/2 \rfloor$ outside bids with value $\geqslant opt^s(b_{\leqslant j\setminus i})$, and with bid value $\hat{v}_i \geqslant opt^s(b_{\leqslant j\setminus i})$. In this case, the price in the auction matches the price in the schedule and the agent wins an item. Moreover, this price is the minimal price it will face because it arrived after j outside bids (the price can only increase to ∞ in future periods).

It is straightforward to show that the agent does not win an item in the auction in the other cases, for which its bid price is less than the price schedule.

The proof of competitiveness closely parallels the proof of competitiveness of DSOT in [9]. Let b'' denote the set of all bids which do not appear in b'. The multiset b' (multiset b'') is a random sample from the uniform distribution on all 2^n sub-multisets of b; this follows from the random ordering hypothesis together with the fact that j is distributed according to the binomial distribution, which is the distribution of the cardinality of a random sub-multiset of b. Thus b', b'' have the same distribution as a uniformly-sampled random partition of b into two sub-multisets. We also recall a definition and a lemma from [9].

DEFINITION 19. *For any $1 \leqslant j \leqslant n$, define B_j to be a sub-multiset of b containing the top j bids, i.e. $B_j = \{v_1, \dots, v_j\}$, and define $n'(B_j)$ to be the number of bids in B_j that are in b'. Given $c \in (0,1)$, we say that j is* c-good *if $\lceil cj \rceil \leqslant n'(B_j) \leqslant j - \lceil cj \rceil$.*

LEMMA 20 ([9]).

$$\Pr[j \textit{ is not } c\textit{-good for some } j > t] \leqslant \frac{2\exp\left(-(1-2c)^2 t/4\right)}{1-\exp\left(-(1-2c)^2/4\right)}.$$

COROLLARY 21.

$$\Pr[j \textit{ is not } 1/24\textit{-good for some } j > 24] < 0.06823.$$

Suppose that $opt^s(b)$ is the m-th largest bid, so that $\mathcal{F}^{(2,s)} = mv_m$. Let G denote the event that $v_1 \in b''$, $v_2 \in b'$. Clearly, $\Pr(G) = 1/4$. Let H denote that event that all $j > 1$ are $1/24$-good and that event G occurs. If event G occurs, all $j \in \{2, 3, \dots, 24\}$ are $1/24$-good. Moreover, by the corollary above, the (unconditional) probability that some $j > 24$ is not $1/24$-good is $<$ 0.06823. Therefore, $\Pr(H) > 1 - \Pr(\neg G) - 0.06823 = 0.18177$. Let $p' = opt^s(b')$; note that this is the price which RM_k charges to bidders in b''. Conditional on event H, $p' n_{p'}(b') \geqslant v_m n_{v_m}(b')$, since $m \leqslant s$. Thus, conditional on H,

$$Rev(RM_k) \geqslant p' n_{p'}(b'') \geqslant \left(\frac{n_{p'}(b'')}{n_{p'}(b')}\right) v_m n_{v_m}(b').$$

Because all $j > 1$ are $1/24$-good, we have

$$\frac{n_{p'}(b'')}{n_{p'}(b')} > \frac{1}{24}\ , \quad \frac{n_{v_m}(b')}{m} \geqslant \frac{1}{24}$$

Hence $\mathbf{Exp}[Rev(RM_k) \mid H] > \frac{1}{24} \cdot \frac{1}{24} \cdot mv_m = \frac{1}{576}\mathcal{F}^{(2,s)}$. To relate this to $\mathcal{F}^{(2,k)}$, we make the following simple observation. Suppose that $opt^k(b)$ is the ℓ-th largest bid, so that $\mathcal{F}^{(2,k)} = \ell v_\ell$. Then $\mathcal{F}^{(2,s)} \geqslant F^s_{v_\ell}(b) = v_\ell \cdot \min\{s, \ell\} \geqslant v_\ell \cdot \frac{\ell}{2} = \frac{1}{2}\mathcal{F}^{(2,k)}$, so

$$\begin{aligned} \mathbf{Exp}[Rev(RM_k)] &> 0.18177 \cdot \mathbf{Exp}[Rev(RM_k) \mid H] \\ &> 0.18177 \cdot \frac{1}{576} \cdot \frac{1}{2}\mathcal{F}^{(2,k)}. \end{aligned}$$

□

To obtain a mechanism which is constant-competitive for efficiency in the $k > 1$ case, we introduce the following mechanism $M'_k(j)$:

1. Let $s = \lceil k/3 \rceil$. At the arrival time of the j-th bidder, let p be the $(s+1)$-th highest bid seen so far. Sell at most s items, at price p, to any bidder in the market at this time whose valuation is greater than or equal to the s-th highest bid. Ties are broken in favor of earlier arrivals, and arbitrarily in the case of simultaneous arrivals.
2. Sell an item to each subsequent bidder whose valuation is greater than or equal to p (break ties due to simultaneous arrivals at random), until the mechanism runs out of items or bidders.

THEOREM 22. *Let RM'_k be the mechanism which selects a random $j \in \{1, \dots, n\}$ according to the binomial distribution $B(n, 1/2)$, and then runs the mechanism $M'_k(j)$. Then RM'_k is a strategyproof online algorithm which is C-competitive with respect to the Vickrey auction for both efficiency and revenue, for a constant $C < 48$.*

PROOF. The proof of strategyproofness is the same as that of Theorem 18 and hence omitted. The proof of competitiveness proceeds as follows. Number the agents $1, 2, \dots, n$ in order of decreasing valuation. As before, the transition point divides the set of agents into two subsets b', b'', and this partition has the same distribution as a random partition sampled by placing each agent independently and equiprobably in b' or b''. Let E denote the event

E : each of the sets b', b'' contains at least $\lceil k/3 \rceil$ of the top k agents.

One proves using Chernoff bounds, as in [9], that the probability of E is at least $1 - 2e^{-k/36}$, which tends to 1 as $k \to \infty$. In fact, for $k > 1$, $\Pr(E)$ attains its minimum value of $3/8$ at $k = 4$.

Conditional on E, we claim that the expected efficiency of RM'_k is at least $1/18$ of the expected efficiency of the optimal allocation, i.e. the allocation which assigns the k items to agents $1, 2, \dots, k$. We will prove this by showing that for $1 \leq i \leq \lceil k/3 \rceil$, agent i has probability at least $1/6$ of winning an item. Let C_i, D_i denote the events

C_i : $i \in b''$

D_i : Of the agents in b'' with the k highest bids, i is among the $\lceil k/3 \rceil$ with the earliest arrival times.

Event C_i is independent of E, and has probability $1/2$. Conditional on C_i and E, the relative ordering of the agents in b'' is random, so D_i holds with probability at least $1/3$. This verifies that $\Pr(C_i \wedge D_i|E) \geq 1/6$. Finally, it is easy to see, from the definition of RM'_k, that agent i wins an item if C_i and D_i hold.

This gives an upper bound of $48 = [(3/8) \cdot (1/18)]^{-1}$ on the competitive ratio for efficiency. The competitive ratio for revenue is bounded above by 8, using a similar proof. Conditional on the event E defined above, RM'_k sells at least $k/3$ items at a price which is greater than or equal to the $(k+1)$-st price, which is the price set by the k-item Vickrey auction. Thus, the expected revenue of RM'_k conditional on E is at least $1/3$ of the Vickrey revenue. Recalling that $\Pr(E) \geqslant 3/8$, we obtain that RM'_k is 8-competitive with Vickrey for revenue. □

By combining mechanisms RM_k and RM'_k we can obtain the following corollary.

COROLLARY 23. *There is a strategyproof online mechanism M''_k which is simultaneously constant-competitive with $\mathcal{F}^{(2,k)}$ for revenue and constant-competitive with Vickrey for efficiency.*

PROOF. We can obtain mechanism M''_k by combining RM_k and RM'_k as follows. At the outset of the auction, toss a fair coin. If it comes up heads, use the revenue-competitive mechanism RM_k; if tails, use the efficiency-competitive mechanism RM'_k. The proof of strategyproofness and competitiveness follows immediately. □

7. CONCLUSIONS AND OPEN PROBLEMS

In this paper, we considered online auctions with a limited supply. We presented value- and time-strategyproof mechanisms with constant efficiency- and revenue-competitiveness when values of the bidders are drawn independently from some unknown distribution. In all of these mechanism, we had a learning phase in which we set a fixed price for the rest of the auction. Following are some open problems.

The main open problem is closing the gaps between the current upper bounds and lower bounds for both efficiency and revenue. More precisely, we believe that our lower bounds 2 for efficiency and 1.5 for revenue in the $k = 1$ case are not tight, and can be improved using different distributions than those considered in the proofs of lower bounds. Solving this problem for the multi-item case seems even more challenging.

The next open question is: what are the lower bounds for efficiency and revenue when the valuations are independent random samples from an unknown probability distribution? (Recall that our lower bounds were for the *random-ordering* model only, not for the independent-sampling model in which it is tougher to prove such lower bounds.)

It would be interesting to know whether for the multi-item case there is an online allocation rule whose efficiency-competitive ratio approaches 1 as k goes to infinity. (If so, is there a time–strategyproof online mechanism implementing this allocation rule?)

Considering the random ordering hypothesis in the multi-unit demand model of Lavi and Nisan [11], in which we have an auction of k identical indivisible goods to a set of bidders where the bidders can demand more than one unit, is an interesting open problem. By treating the k items as a single bundle and running a single-item auction for this bundle, we can obtain strategyproof mechanisms with competitive ratios of ek for efficiency or $4k$ for revenue with respect to the offline Vickrey mechanism. Can we obtain constant competitive ratios for this case? Is the Vickrey mechanism the appropriate benchmark for revenue in this case, or is there an appropriate analogue of $\mathcal{F}^{(2,k)}$ for multi-unit auctions?

Finally, it is very interesting to consider online auctions when we have re-usable goods e.g. in pricing WiFi. Even for the single-item case, it appears very challenging to obtain competitive online mechanisms in this setting.

8. ACKNOWLEDGEMENTS

We would like to thank Jason D. Hartline, Anna Karlin, and Ron Lavi for fruitful discussions.

9. REFERENCES

[1] M. AJTAI, N. MEGIDDO, AND O. WAARTS, *Improved algorithms and analysis for secretary problems and generalizations*, SIAM J. Discrete Math., 14 (2001), pp. 1–27.

[2] B. AWERBUCH, Y. AZAR, AND A. MEYERSON, *Reducing truth-telling online mechanisms to online optimization*, in Proceedings of the thirty-fifth ACM symposium on Theory of computing, ACM Press, 2003, pp. 503–510.

[3] A. BLUM, V. KUMAR, A. RUDRA, AND F. WU, *Online learning in online auctions*, in Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (Baltimore, MD, 2003), New York, 2003, ACM, pp. 202–204.

[4] E. B. DYNKIN, *The optimum choice of the instant for stopping a markov process*, Sov. Math. Dokl., 4 (1963), pp. 627–629.

[5] E. J. FRIEDMAN AND D. C. PARKES, *Pricing WiFi at starbucks: issues in online mechanism design*, in The 4nd ACM Conference on Electronic Commerce, 2003, pp. 240–241.

[6] J. GALLIEN, *Dynamic mechanism design for online commerce*, tech. rep., Sloan School of Management, MIT, 2002.

[7] J. P. GILBERT AND F. MOSTELLER, *Recognizing the maximum of a sequence*, J. Amer. Statist. Assoc., 61 (1966), pp. 35–73.

[8] K. S. GLASSER, R. HOLZSAGER, AND A. BARRON, *The d choice secretary problem*, Comm. Statist. C—Sequential Anal., 2 (1983), pp. 177–199.

[9] A. V. GOLDBERG, J. D. HARTLINE, A. R. KARLIN, M. SAKS, AND A. WRIGHT, *Competitive auctions and digital goods*, preliminary version in Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms (Washington, DC, 2001), Philadelphia, PA, 2001, SIAM, pp. 735–744. Journal version to appear in Games and Economic Behavior.

[10] R. KLEINBERG AND T. LEIGHTON, *The value of knowing a demand curve: Bounds on regret for on-line posted-price auctions*, in 44th Annual Symposium on Foundations of Computer Science (Cambridge, 2003), IEEE Computer Soc., Los Alamitos, CA, 2003. To appear.

[11] R. LAVI AND N. NISAN, *Competitive analysis of incentive compatible on-line auctions*, in The 2nd ACM Conference on Electronic Commerce, 2000, pp. 233–241.

[12] D. C. PARKES AND S. SINGH, *An MDP-based approach to online mechanism design*. In Proc. 17th Annual Conference on Neural Information Processing Systems, 2003.

[13] R. J. VANDERBEI, *The optimal choice of a subset of a population*, Math. Oper. Res., 5 (1980), pp. 481–486.

[14] W. VICKREY, *Counterspeculations, auctions, and competitive sealed tenders*, Journal of Finance, 16 (1961), pp. 8–37.

[15] J. G. WILSON, *Optimal choice and assignment of the best m of n randomly arriving items*, Stochastic Process. Appl., 39 (1991), pp. 325–343.