

Automated Online Mechanism Design and Prophet Inequalities

MohammadTaghi Hajiaghayi

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213, USA
Email: hajiagha@cs.cmu.edu.

Robert Kleinberg*

Department of Computer Science
Cornell University
Ithaca, NY 14850, USA
Email: rdk@cs.cornell.edu.

Tuomas Sandholm†

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213, USA
Email: sandholm@cs.cmu.edu.

Abstract

Recent work on online auctions for digital goods has explored the role of optimal stopping theory — particularly secretary problems — in the design of approximately optimal online mechanisms. This work generally assumes that the size of the market (number of bidders) is known a priori, but that the mechanism designer has no knowledge of the distribution of bid values. However, in many real-world applications (such as online ticket sales), the opposite is true: the seller has distributional knowledge of the bid values (e.g., via the history of past transactions in the market), but there is uncertainty about market size. Adopting the perspective of *automated mechanism design*, introduced by Conitzer and Sandholm, we develop algorithms that compute an optimal, or approximately optimal, online auction mechanism given access to this distributional knowledge. Our main results are twofold. First, we show that when the seller does not know the market size, no constant-approximation to the optimum efficiency or revenue is achievable in the worst case, even under the very strong assumption that bid values are i.i.d. samples from a distribution known to the seller. Second, we show that when the seller has distributional knowledge of the market size as well as the bid values, one can do well in several senses. Perhaps most interestingly, by combining dynamic programming with prophet inequalities (a technique from optimal stopping theory) we are able to design and analyze online mechanisms which are temporally strategyproof (even with respect to arrival and departure times) and approximately efficiency(revenue)-maximizing. In exploring the interplay between automated mechanism design and prophet inequalities, we prove new prophet inequalities motivated by the auction setting.

Introduction

Mechanism design has traditionally focused on the offline setting where all agents are present upfront. However, many electronic commerce applications do not fit that model because the agents can arrive and depart dynamically. This is characteristic, for example, of online ticket auctions (e.g., on Priceline, Expedia, and Travelocity), search keyword auctions (e.g., on Google, Yahoo!, and MSN), eBay-style Internet auctions, pricing access to a WiFi port (e.g., at Starbucks) and scheduling computing jobs on a shared server

*Supported by a National Science Foundation Mathematical Sciences Postdoctoral Research Fellowship. Portions of this work were completed while the author was a postdoctoral scholar at UC Berkeley.

†Supported by National Science Foundation grant IIS-0427858. Copyright © 2007, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

(e.g., in MoteLab and PlanetLab). The online aspect is characteristic of some important traditional applications as well, such as the sale of a house, where the buyers arrive and depart dynamically.

The design of online mechanisms has attracted research interest since 2000 (see, e.g., a survey by Parkes in (Nisan *et al.* 2007)). Designing such mechanisms is challenging because of the combination of mechanism design challenges (ensuring truthfulness — often not only about valuations but also about arrival and departure times) and online algorithm challenges (dealing with uncertainty about future inputs). For example, the most canonical technique for designing truthful offline mechanisms, the Vickrey-Clarke-Groves (VCG) scheme, is inapplicable in most online problems because it requires determination of an optimal allocation, which is generally impossible in the online setting.

Most prior work on designing online mechanisms has adopted a worst-case adversary model, as is common in the design of online algorithms. As a result, the guarantees that have been proven have usually been relatively weak. In most real applications, there is significant probabilistic information available about the future, and it seems counterproductive to ignore it. For example due to the history of past transactions in the market, in most real-world applications the auctioneer has distributional information about the valuations, and sometimes even the number, of bidders. Using such probabilistic information is crucial in practice, especially when we want to maximize revenue: for example, being within a factor 1.1 versus 2 is drastically different though both are constant in a theoretical sense. Our goal in this paper is to exploit the distributional information in online mechanism design.

Automated mechanism design, first introduced in (Conitzer & Sandholm 2002), plays a key role in this. In that approach, the mechanism is created automatically — using some optimization algorithm — for the specific problem instance (including the distributional information) and objective at hand. See, e.g., a survey in (Sandholm 2003). This has important advantages:

- It can be used in settings beyond the classes of problems that have been successfully studied in (manual) mechanism design to date. For example, it has been used to generate revenue-maximizing combinatorial auctions (Likhodedov & Sandholm 2005), a problem that eludes analytical characterization even in the 2-item case.
- It can circumvent the impossibility results: when the mechanism is designed for the setting (instance) at hand, it does not matter that it would not work on preferences

beyond those in that setting (*e.g.*, for a class of settings). Even when the automatically-created optimal mechanism does not circumvent the impossibility, it always minimizes the pain entailed by impossibility.

- It can yield better mechanisms (in terms of better outcomes and/or stronger nonmanipulability guarantees¹) than the canonical mechanisms because the mechanism capitalizes on the particulars of the setting, *i.e.* the probabilistic (or other) information that the mechanism designer has about the agents' preferences.
- It shifts the burden of mechanism design from humans to a machine. That is why one can afford to do the design anew for every instance.

Put together, automated mechanism design can exploit the specific model of the problem and the distributional knowledge about bidders' private information such as valuations. Work has been done both on general-purpose techniques for automated mechanism design as well as on automated mechanism design approaches for specific applications. While a bit of that research has studied multi-stage mechanisms (Sandholm & Gilpin 2006; Sandholm, Conitzer, & Boutilier 2007) (in order to reduce the agents' preference determination effort, communication costs, and privacy loss), all of the prior work on automated mechanism design has focused on offline settings where all the agents are present upfront. In this paper we present the first work on automated mechanism design for an online problem.

Setting

In our setting there are k identical indivisible goods (a.k.a. *units*) for sale, and there are n agents (a.k.a. bidders), each of whom wants to purchase one unit. The *type* of an agent i , $1 \leq i \leq n$, is defined by an ordered triple (a_i, d_i, v_i) , whose three components are called *arrival time*, *departure time*, and *value*, respectively. We assume quasi-linear utilities throughout this paper, so if an agent receives one or more units during the time interval $[a_i, d_i]$ with a payment p_i , her utility for this allocation is $v_i - p_i$; for all other allocations her utility is 0. As is common in online auction design (Hajiaghayi, Kleinberg, & Parkes 2004; Hajiaghayi *et al.* 2005), we adopt a *restricted misreporting* model. Throughout this paper, unless stated otherwise, we assume *no early arrivals* but *unrestricted departures*, *i.e.*, an agent of type (a_i, d_i, v_i) may report any triple $(\hat{a}_i, \hat{d}_i, \hat{v}_i)$ satisfying $\hat{a}_i \geq a_i$. The no early arrival assumption is motivated by the view that in practice an agent would not participate in an auction before she knows that she desires a unit, and we can view the time a_i as the time when her desire arises. When we assume $a_i = d_i$, we call the agents *instantaneous*; otherwise we call them *patient*.

An online direct revelation mechanism consists of an *allocation rule* $q_i(t, \vec{\theta})$ (where t denotes the time and $\vec{\theta}$ denotes the vector of reported types) along with a *payment rule* $p_i(\vec{\theta})$ such that $q_i(t, \vec{\theta})$ is a monotonically non-decreasing $\{0, 1\}$ -valued function of t , $\sum_{i=1}^n q_i(t, \vec{\theta}) \leq k$ for all $t, \vec{\theta}$,

¹*E.g.*, satisfaction of *ex post* incentive compatibility and/or individual rationality constraints rather than their *ex interim* variants.

and $q_i(t, \vec{\theta})$ depends only on the types reported by agents who report an arrival time \hat{a}_j which is less than or equal to t . We seek a direct-revelation mechanism which is truthful (a.k.a. strategyproof) in dominant strategies (*i.e.*, the utility of agent i is maximized if she bids truthfully, regardless of what other agents report). When it is a dominant strategy for agents to truthfully report not only their value, but also their arrival and departure time, we say that the mechanism is *temporally strategyproof*. To evaluate the performance of an online mechanism, we will use two measures of solution quality: efficiency and revenue. The *efficiency* of an outcome is the combined welfare of all agents, *i.e.* $\sum_i q_i v_i$. The *revenue* of an outcome is the sum of the payments made by the agents, *i.e.* $\sum_i p_i$. We say a mechanism is ρ -competitive with respect to efficiency (resp. revenue) if the expected efficiency (resp. revenue) of the outcome computed by the mechanism is at least $1/\rho$ times the expectation of the maximum efficiency over all outcomes (resp., the maximum revenue that can be obtained by setting a single fixed price p and selling to all agents whose value is at least p).

We will assume throughout that the seller has some distributional information about the sequence of bids to be received. More precisely, let σ be a permutation of $\{1, 2, \dots, n\}$ such that $a_{\sigma(1)} \leq a_{\sigma(2)} \leq \dots \leq a_{\sigma(n)}$ and let \mathbf{x} be the infinite sequence $(v_{\sigma(1)}, v_{\sigma(2)}, \dots, v_{\sigma(n)}, 0, 0, \dots)$. (In a truthful direct revelation mechanism, the first n entries of the sequence \mathbf{x} represent the bid values in the order they are received in dominant strategy equilibrium.) A probability distribution on agent populations (*i.e.*, finite sets of ordered triples (a_i, d_i, v_i)) induces a probability distribution on bid sequences \mathbf{x} . We will assume that there is an infinite sequence $\mathbf{y} = (y_1, y_2, \dots)$, whose distribution is known to the seller, such that $x_i = y_i z_i$, where $z_i = 1$ if $1 \leq i \leq n$, 0 otherwise. (One may think of \mathbf{y} as the sequence of bids which would be received if the agent population were inexhaustible.) We make two different types of assumptions about the seller's distributional information:

- **Full information:** The seller knows the distribution of \mathbf{x} . That is, the seller has distributional information about valuations and market size.
- **Unknown n :** The seller knows the distribution of \mathbf{y} but not \mathbf{x} . That is, the seller has distributional information about valuations but not the market size.

The following special cases will be of interest:

- **Independent bids:** The random variables y_i ($1 \leq i < \infty$) are independent.
- **Independent bids, fixed n :** The random variables x_i ($1 \leq i < \infty$) are independent. Equivalently, the random variables y_i are independent and the value of n is fixed.
- **I.i.d. bids:** The random variables y_i ($1 \leq i < \infty$) are independent and identically distributed.

Our contributions

This paper makes three main contributions to the theory of online mechanism design. First, we raise the issue of designing online mechanisms when the number of bidders is not known in advance, and develop the basic possibility and

impossibility theorems which accompany this notion. Second, we present an automated mechanism design algorithm (and other results) for the setting in which the number of bidders is known at least probabilistically. Third, we reveal the power of prophet inequality techniques as a toolkit for solving problems in automated online mechanism design. Below, we elaborate on each of these contributions.

Our first set of results concerns the case in which the number of bidders, n , is specified adversarially but the bid values are randomly sampled from a known distribution. We prove that no mechanism can be constant-competitive in this setting, and we exhibit mechanisms whose competitive ratio is nearly logarithmic in n , or logarithmic in h , the ratio between the maximum and minimum possible bids. Next we turn to the case when the value of n is unknown, but its distribution is known to the mechanism designer. For settings where an upper bound on n is known, we specify a dynamic program to compute, in polynomial time, the best non-decreasing price sequences for revenue and for efficiency. We insist on non-decreasing prices to ensure temporal strategyproofness. For settings where n is unbounded, we prove that the revenue-maximizing price sequence is non-decreasing if and only if the distribution of n has non-increasing hazard rate. So, under that condition, requiring temporal strategyproofness does not compromise revenue.

Our next set of results reveals a relationship between these mechanism design problems and the subject of *prophet inequalities* from optimal stopping theory. Specifically, prophet inequalities (along with the constructive proofs of these inequalities) allow us to design online mechanisms with two especially desirable features: the mechanisms are temporally strategyproof, and they satisfy a provable performance guarantee which relates their efficiency to the optimum efficiency in hindsight. These two features, in turn, constitute an analysis technique which allows us to show that the mechanism produced by the dynamic program discussed above achieves a constant-factor approximation to the omniscient surplus when n is known.

Prior work on online auctions

While mechanism design has traditionally focused on a static problem where all the agents are present up front, there has been significant recent work on designing online mechanisms where the agents arrive and depart over time. The first paper on online auctions was in 2000 (Lavi & Nisan 2000); the first on online double auctions (a.k.a. exchanges) was in 2002 (Blum, Sandholm, & Zinkevich 2006).

A recent overview article by Parkes surveys the field of online mechanisms (Nisan *et al.* 2007). Much of the work (e.g. (Blum *et al.* 2003; Kleinberg & Leighton 2003)) assumes that the agents arrive in a predetermined order which is not under their control, and that an agent's only private information is her value. This makes it much easier to design truthful mechanisms. Some online mechanisms (e.g. (Awerbuch, Azar, & Meyerson 2003; Lavi & Nisan 2000)) are strategyproof against agents misstating their arrival or departure time because they are based on prices which do not decrease over time.

VCG-based online mechanisms were introduced

in (Friedman & Parkes 2003). Such mechanisms are (dominant-strategy) truthful in the rare cases where the underlying allocation problem admits an online algorithm with competitive ratio 1. (Parkes & Singh 2003) have studied VCG-based online mechanisms also under a weaker notion of incentive compatibility, Bayes-Nash equilibrium, adopting the framework of Markov Decision Processes. The setting for this work is quite general.

(Hajiaghayi, Kleinberg, & Parkes 2004) present constant-competitive online mechanisms for auctioning identical goods when each agent is assumed to arrive and depart dynamically. Unlike all previous papers, they assume each agent has three pieces of private information: her value, her arrival time and her departure time. However, they assume that the agents arrive in random order and that the value of n (the total number of agents) is known to the mechanism designer in advance. This work has led to several subsequent papers (e.g., (Babaioff, Immorlica, & Kleinberg 2007; Bredin & Parkes 2005; Hajiaghayi *et al.* 2005; Kleinberg 2005; Lavi & Nisan 2005; Ng *et al.* 2005)). In the same setting in which each bidder has three pieces of private information (not necessarily random order of bidders), (Hajiaghayi *et al.* 2005) study the case of re-usable goods such as processor time in which goods can be allocated to different bidders at different time slots. They also present general characterizations for the class of truthful online allocation rules, which extend beyond the typical single-parameter settings and formalize the role of *restricted misreporting* in reversing existing price-based characterizations. (Hajiaghayi *et al.* 2005) mainly consider truthful online auctions for unit-length jobs; (Porter 2004) presents a truthful mechanism for the variation with different length jobs in which an agent derives positive utility if she is granted the resource for a total duration equal to its job length. (Lavi & Nisan 2005) also study an online auction setting which is closely related to that of (Hajiaghayi *et al.* 2005). Assuming unrestricted misreports, they prove strong negative results for deterministic truthful auctions (no such mechanism can achieve a competitive ratio better than the number of units) and this leads them to consider a weaker solution concept (*set-Nash equilibrium*) which admits constant-competitive mechanisms.

Parallel and independent to our work on automated online mechanism design, (Pai & Vohra 2006) also recently consider dynamic online auctions in the same setting as (Hajiaghayi, Kleinberg, & Parkes 2004) (but without the random ordering assumption), for which they derive the revenue maximizing Bayesian incentive compatible selling mechanism. They observe the failure of the revelation principle in this setting and together with (Gallien 2006), they extend Myerson's truthful optimal auctions (Myerson 1981) to dynamic environments. Like our results, they assume distributional knowledge of the bid valuations, but unlike our results their algorithm has exponential running time and their focus is on characterizing dynamic optimal auctions rather than running time of the algorithm. All of our algorithms in this paper have polynomial running times, and much of our work focuses on the case of unknown number of bidders, which is different from the assumption by Pai and Vohra.

In the online mechanism design setting, to the best of our

knowledge, the only previous work to incorporate the aspect of unknown market size is (Mahdian & Saberi 2006). Motivated by search keyword auctions on, for example, Yahoo!, Google, and MSN, where the supply (search queries by users, each of which results in a search results page on which ads can be displayed) is not known in advance, they study a multi-unit auction for a perishable item where an unknown number of units (instead of bidders) arrive online. They design approximately revenue-maximizing auctions using an online algorithm similar to an algorithm used for the ski-rental problem. However, they do not use any distributional knowledge of the bid values — which is usually available via the history of past transactions.

Prior work on prophet inequalities

The full-information case of our online auction problem is closely related to *prophet inequalities*, a topic which has been studied in optimal stopping theory since the 1970’s. For an instantaneous agent, the mechanism is required to decide whether or not the agent will receive a unit at the moment that agent’s bid is revealed, and no later. In this special case, the problem of designing an allocation rule which maximizes (or approximately maximizes) efficiency is equivalent to the following problem: given the distribution of a sequence of random variables x_1, x_2, \dots , design a sequence of k stopping rules² $\tau_1 < \tau_2 < \dots < \tau_k$ to maximize the expectation of the sum $x_{\tau_1} + \dots + x_{\tau_k}$. A great deal is known about the solution of this problem in the case $k = 1$, and comparatively little is known for $k > 1$. The most basic prophet inequality, discovered by Krengel, Sucheston, and Garling, concerns the case in which $k = 1$ and the random variables x_1, x_2, \dots are independent (but not necessarily identically distributed). If V denotes the supremum of $\mathbb{E}(x_\tau)$ over all stopping rules τ , and $M = \mathbb{E}(\sup_n x_n)$, then assuming $M < \infty$ we have $M \leq 2V$, and the constant 2 is the best possible in that setting (Krengel & Sucheston 1977; 1978). The inequality has been interpreted as meaning that a prophet who can see the future has only a bounded advantage over a gambler who observes the random variables one by one, and this explains the name “prophet inequality”.

When the gambler has multiple choices (i.e., $k > 1$), then much less is known about prophet inequalities. (As-saf & Samuel-Cahn 2000) studied the problem of designing k stopping rules $\tau_1 \leq \dots \leq \tau_k$ to optimize the quantity $\mathbb{E}(\max_i \{x_{\tau_i}\})$. They proved that there exists a sequence of k stopping rules such that the expected maximum of the k choices is within a factor $(k + 1)/k$ of the prophet’s payoff. However, in the auction setting, the natural objective is to maximize the expected sum of the k choices rather than their expected maximum. Surprisingly, only one prior paper considers this objective (Kennedy 1987). It compares the sum of

²We define a *stopping rule* for a sequence x_1, x_2, \dots to be a random variable τ taking values in $\mathbb{N} \cup \{\infty\}$, such that for all i , the event $\tau = i$ depends only on the values of x_1, \dots, x_i . Our convention differs from the convention adopted by most authors, in that we allow the value $\tau = \infty$. In auction applications, this corresponds to allowing units to be unsold at the end of the auction. We also adopt the conventions that $x_\tau = 0$ when $\tau = \infty$, and that the inequality $\tau_j < \tau_{j+1}$ is valid when $\tau_j = \tau_{j+1} = \infty$.

the k values chosen by the gambler with a *single* value chosen by the prophet. Letting $M = \mathbb{E}(\max_n x_n)$ as before, and letting V_k denote the supremum of $\mathbb{E}(x_{\tau_1} + \dots + x_{\tau_k})$ over all sequences of stopping rules τ_1, \dots, τ_k , Kennedy seeks inequalities of the form $M \leq \alpha_k V_k$ which are valid for all sequences of independent non-negative random variables x_1, x_2, \dots . He gives a recursive formula for the best possible constant α_k in such an inequality. This type of “apples to oranges” comparison (one choice for the prophet, k choices for the gambler) is much less natural than a comparison between a prophet with k choices and a gambler with the same number of choices. In other words, letting M_k denote the expectation of the sum of the k largest samples in the sequence x_1, x_2, \dots , we desire inequalities of the form $M_k \leq \beta_k V_k$, valid for all sequences of independent non-negative random variables x_1, x_2, \dots . Kennedy notes that this objective seems much more difficult than determining the best possible constants α_k , because the recursions involved are much harder to manipulate. In this paper we make progress on this problem by proving that

$$1 + \sqrt{\frac{1}{512k}} \leq \beta_k \leq 1 + \sqrt{\frac{8 \ln(k)}{k}}.$$

We also supply nearly tight bounds for a similar question concerning the *additive* difference between V_k and M_k when x_1, x_2, \dots are uniformly bounded but not necessarily independent.

Unknown n : lower and upper bounds

In this section, we consider online auctions in which all bidders’ valuations are drawn i.i.d. from a demand distribution which is known by the algorithm designer. As argued above, assuming advance knowledge of such valuation distributions is well-motivated by several real-world applications such as online ticket sales, search keyword auctions, eBay-style auctions, pricing access to a WiFi port, scheduling computer jobs on a shared server, and house sales. The main focus of this section is the case in which n , the number of bidders, is unknown due to the online nature of the problem.

We first show that in the case of one unit to sell and instantaneous agents — and even with a highly concentrated (aka. light-tailed) demand distribution — we cannot be constant-competitive for efficiency or revenue. Next, we show that we can obtain logarithmic competitive ratios even in a much more general setting with arrivals and departures.

Theorem 1. *Suppose we have only one unit to sell ($k = 1$), all bidders are instantaneous, and their valuations are drawn i.i.d. from a demand distribution which is given to the algorithm in advance. Without knowledge of n (the number of bidders), it is impossible for a mechanism to achieve a constant competitive ratio with respect to revenue or efficiency.*

Theorem 2. *There are truthful mechanisms which are $\min\{O(\log h), O(\log n(\log \log n)^2)\}$ competitive (with respect to both efficiency and revenue) for online auctions of k units with bidders of arbitrary arrival-departure intervals even when n is unknown to the auctioneer in advance. Here h is the ratio of the maximum bid to the minimum bid.*

For space reasons, we present only a sketch of the proofs of Theorems 1 and 2. To prove Theorem 1, for any constant $C > 1$ one proves the non-existence of C -competitive mechanisms by constructing a distribution of bids which is supported on the set $\{0, 1, C, C^2, C^3, \dots\}$, and a sequence of numbers t_1, t_2, \dots , such that the maximum of t_r independent samples from the given distribution is with high probability equal to C^r . One then finds a value of r such that the mechanism has probability less than $1/C$ of stopping and selling the item at a time t between t_{r-1} and t_r , and one sets $n = t_r$. This ensures that, with high probability, the optimum efficiency and revenue are equal to C^r . By construction, with high probability the mechanism either doesn't sell the item at all, or it sells the item at a time earlier than t_{r-1} ; in the latter case, the price can not be greater than the maximum of the first t_{r-1} bids, which by construction is at most C^{r-1} with high probability. To prove Theorem 2, one considers a mechanism which guesses the value of n , either by choosing a uniformly random power of 2 between 1 and h or else by choosing a random power of 2 from the distribution which assigns probability proportional to $k \log^2(k)$ to the number 2^k . Having guessed the value of n , it then uses the mechanism from Section 6, Corollary 23 of (Hajiaghayi, Kleinberg, & Parkes 2004), which is constant-competitive for efficiency and revenue provided that the guessed value of n is within a constant factor of the true value of n .

Known distribution over n

Until now, we have assumed that the number of bidders, n , was unknown to the designer, and we have adopted worst-case competitive analysis with respect to the choice of n . If instead the distribution of n is known, then we can make several observations about the design of mechanisms to maximize expected revenue or efficiency. We split the discussion into two subsections based on whether or not the distribution over n has bounded support, i.e., whether or not there exists a finite maximum number of bidders that may arrive.

Bounded number of buyers

For the setting where the distribution of n has bounded support, a mechanism (sequence of prices) can be designed automatically.

Theorem 3. *Assume a bounded-support distribution \mathcal{N} over n and bounded-support distributions of bidders' valuations. If we insist (e.g., in order to achieve strategyproofness of reporting arrival times³) that the price sequence be non-decreasing (i.e., that the price never drops as a new agent enters the system), then a mechanism that maximizes expected efficiency and revenue (among all mechanisms with non-decreasing prices) can be automatically constructed.*

Proof. The algorithm is based on a dynamic program that uses backward induction on time, starting from the arrival

³In order to achieve strategyproofness of reporting arrival times, it is not quite necessary for the posted prices to be increasing. One could post prices that sometimes decrease, as long as rebates (or additional charges) are made afterwards so that the prices are *effectively* non-decreasing. See, e.g., (Hajiaghayi *et al.* 2005).

time of the last bidder. The array that we fill has three dimensions: n' is the number of bidders seen so far, k' is the number of remaining units to sell, and q' is the last (i.e., greatest) price at which we sold a unit so far. $D[n', k', q']$ stores the best expected revenue (efficiency) that we can obtain at this stage. The pseudo-code is given below.

Algorithm OPTMech

Input: distribution \mathcal{N} and an upper bound N on n ,
distributions \mathcal{Q}_i , $1 \leq i \leq n$,
an upper bound Q on bidders' valuations, and
 k , the number of units to sell

Output: Table D and a corresponding table R of reserve prices

- 1 **initialize** table D by 0 when $n' = N + 1$, $k' = 0$, or $q' = Q + 1$
- 2 **for** $n' = N$ down to 0
- 3 **for** $k' = 1$ to k
- 4 **for** $q' = Q$ down to 0
- 5 **let** $p_{>n'}$ be the probability (from distribution \mathcal{N}) of having more than n' bidders, conditional on having seen n' of them so far
- 6 **let** $p_{\geq q''}$ be the probability (from distribution $\mathcal{Q}_{n'}$) that the n' th bidder has valuation q'' or higher
- 7 **set** reserve price $R[n', k', q'] = q''$ where q'' , $q' \leq q'' \leq Q$ maximizes
($1 - p_{>n'}$) $p_{\geq q''} B_{q''} + p_{>n'}((1 - p_{\geq q''}) D[n + 1, k, q''] + p_{\geq q''}(B_{q''} + D[n' + 1, k - 1, q'']))$.
Here, $B_{q''} = q''$ if we want to maximize revenue and $B_{q''} = \sum_{x=q''}^Q p_x x$ if we want to maximize efficiency (p_x is the probability that the n' th bidder has valuation x).

The generated mechanism starts with $n' = 0$, $k' = k$, and $q' = 0$.

The proof of correctness of the dynamic program can be easily seen from the pseudo-code above. The proof of strategyproofness of the resulting mechanism mainly follows from (Hajiaghayi, Kleinberg, & Parkes 2004; Hajiaghayi *et al.* 2005). Essentially since the prices never decrease, there is no incentive for bidders to arrive late in the system. Since the prices are independent of bidders' announced valuations, the mechanism is also strategyproof with respect to reporting valuations. \square

In OPTMech, different bidders can have different demand distributions, though in the simplest case they are all i.i.d. from the same distribution \mathcal{Q} . We will discuss the case of different demand distributions further when we discuss algorithms that use prophet inequalities.

Unbounded number of buyers

We now shift attention to the setting where n has unbounded support, i.e., the number of agents may end up being arbitrarily large. We start by showing that in a 1-unit auction, if n is drawn from a distribution with *non-increasing hazard rate*, then the optimal price sequence is inherently non-decreasing. (Intuitively, if the option of continuing to take future bids increases in value over time, then it is best to increase the prices over time as well.) This means that first-best revenue is achievable even while requiring temporal strategyproofness (non-decreasing prices). On the other hand, with increasing hazard rate, first-best revenue may not be achievable under temporal strategyproofness.

Theorem 4. *Let there be one unit to sell, and let f be the distribution from which each bidder's valuation is independently drawn. Assume that f is bounded and that the support of f is bounded from above. Now, if the number of buyers, n , is drawn from a distribution with non-increasing hazard rate (i.e., a distribution such that the function $\phi(t) = \Pr(n = t \mid n \geq t)$ is a non-increasing function of t), then the prices in the optimal price sequence do not decrease over time. (Analogously, if $\phi(t)$ is increasing, the prices in the revenue-maximizing price sequence decrease over time.)*

In the special case where n is drawn from a geometric distribution, the hazard rate, $\phi(t)$, is constant. Using the technique of the proof of Theorem 4 (omitted for space), it follows that it is optimal to keep posting the same price across time. That optimal price can be found via binary search.

Theorem 4 can also be applied to any tail of a distribution (i.e., all arrivals after some given number of arrivals), if the tail satisfies non-increasing (increasing) hazard rate, even if the entire distribution does not satisfy that condition.

If the seller has multiple units to sell, the state includes also the number of units that are still unsold, i.e., the MDP value function is a two-variable function $v^*(t, u)$ instead of a single-variable function $v^*(t)$. For any fixed u , Theorem 4 still applies. In other words, if the hazard rate is non-increasing (increasing), the price is non-decreasing (decreasing). Then, any time that a unit is sold, u is decremented by one. With less supply left, the optimal price naturally does not decrease. Therefore, we can say that with non-increasing hazard rate, price is nondecreasing even in the multi-unit case; thus first-best revenue is achievable even under the requirement of temporal strategyproofness.

Prophet inequalities

In this section we will develop online mechanisms based on a technique from optimal stopping theory called *prophet inequalities* (described in the introduction). Prophet inequalities concern the design of one or more stopping rules for a sequence of random variables which are often, but not always, assumed to be independent. In the online auction setting, the random variables correspond to bids, and the stopping rules specify the times when units are sold. Our consideration of these stopping rules yields several benefits:

- We will be able to derive competitive ratios on these mechanisms, i.e., comparisons between expected efficiency (revenue) of the algorithm's allocation and the expected efficiency (revenue) of the optimum allocation that can be obtained with perfect foresight and without considering any incentive compatibility (regarding reporting values and times). The competitive ratios hold under the assumption that the bid values are independent random variables and that the value of n is known.
- Each mechanism here uses constant non-discriminatory prices. Therefore, the same mechanism could have been generated by the dynamic program of the previous section. It follows that the competitive ratios apply to mechanisms generated by the dynamic program as well, which implies, *a fortiori*, that the dynamic program gen-

erates a mechanism which is a constant-factor approximation to the efficiency (or revenue) of an optimal mechanism. This result strengthens Theorem 3, which only asserts optimality within the class of mechanisms with non-discriminatory non-decreasing prices. (We reiterate, however, that the competitive ratios — and hence the approximation guarantee — rely on the assumption that the value of n is known.) To summarize, the prophet inequality based mechanisms can serve as an analysis tool for the mechanisms generated by the dynamic program.

- Unlike the mechanisms generated using the dynamic program, each mechanism here uses a constant price. Therefore, these mechanisms are temporally strategyproof against arbitrary temporal deviations, including early arrivals.

Example: the case $k=1$

To illustrate these ideas, consider the problem of stopping a sequence of independent random variables x_1, x_2, \dots, x_n , with known (not necessarily identical) distributions supported on the nonnegative reals, at a stopping time τ , to maximize the expectation of x_τ . This corresponds to the online automated mechanism design problem with known n and $k = 1$, and with the objective of maximizing efficiency. As mentioned in the introduction, Kregel, Sucheston, and Garling proved (Kregel & Sucheston 1977; 1978) that there exists a stopping rule τ such that

$$2\mathbb{E}(x_\tau) \geq \mathbb{E}(\max_{1 \leq i \leq n} x_i). \quad (1)$$

In fact, the bound (1) is achieved by at least one of the following two stopping rules σ, σ' . First, let $x^* = \max_{1 \leq i \leq n} x_i$. Next let m be median of the distribution of x^* , i.e., choose m such that $\Pr(x^* < m) \leq 1/2$ and $\Pr(x^* > m) \leq 1/2$. Finally, let σ be the minimum value of i such that $x_i > m$ (or $\tau_0 = n$ if there is no such i) and let σ' be the minimum value of i such that $x_i \geq m$ (or $\tau_1 = n$ if there is no such i). The following theorem is due to Ester Samuel-Cahn (Samuel-Cahn 1984).

Theorem 5. *At least one of the numbers $2\mathbb{E}(x_\sigma), 2\mathbb{E}(x_{\sigma'})$ is greater than or equal to $\mathbb{E}(x^*)$.*

In fact, there is a simple criterion for determining which of the stopping rules σ, σ' approximates the expectation of $\mathbb{E}(x^*)$. Let $\beta = \sum_{i=1}^n \mathbb{E}(\max\{0, x_i - m\})$. If $m \leq \beta$ then $2\mathbb{E}(x_\sigma) \geq \mathbb{E}(x^*)$. If $m \geq \beta$ then $2\mathbb{E}(x_{\sigma'}) \geq \mathbb{E}(x^*)$. Note that both m and β can be efficiently computed if the distribution of each random variable x_i has finite support and is given explicitly as part of the input: the value of m can be found by binary search, and the value of β can be found by directly evaluating the formula which defines β .

In the online auction setting, suppose there are n bidders whose bids are independent random variables with known distributions. Let x_1, x_2, \dots, x_n be the n bids, in the order they are received, and let $x_{n+1} = 0$. One can define stopping rules σ, σ' , as above, for the sequence x_1, \dots, x_{n+1} . Each of these stopping rules corresponds to an online allocation rule which sells the unit to the bidder who arrives at stopping time σ (resp. σ') unless the stopping time is $n + 1$,

in which case the unit is unsold. Note that both of these allocation rules can be implemented, in dominant strategy equilibrium, by a posted price mechanism which sells the unit to the first bidder whose bid value is strictly greater than m (in the case of σ) or greater than or equal to m (in the case of σ'). A posted-price mechanism with a price that does not vary over time is temporally strategyproof (even against arbitrary misreporting of arrival and departure times) so although we designed the mechanism by reasoning about the instantaneous agents setting, we obtained a stronger form of incentive compatibility “for free” due to the constructive proof of the prophet inequality. This is a theme which will be repeated in future sections.

Selling more than one unit

Generalizing the foregoing discussion to the case when $k > 1$, n is known, and the bids are independent, we arrive at the following question about prophet inequalities.

Question 6. *If x_1, x_2, \dots, x_n is a finite sequence of random variables, let $OPT_k(x_1, \dots, x_n)$ denote the random variable which is the sum of the k largest elements of the set $\{x_1, \dots, x_n\}$. For a given natural number k , what is the smallest constant β_k such that for every finite sequence of independent nonnegative random variables x_1, x_2, \dots, x_n , there exists a sequence of k stopping rules $\tau_1 < \tau_2 < \dots < \tau_k$ satisfying $\beta_k \mathbb{E}(x_{\tau_1} + \dots + x_{\tau_k}) \geq \mathbb{E}(OPT_k(x_1, \dots, x_n))$?*

We have seen that $\beta_1 = 2$. Surprisingly, the problem of determining, or estimating, the value of β_k for $k > 1$ has not been explicitly considered in the literature on prophet inequalities. Here we present upper and lower bounds for β_k and then discuss their implications for automated online mechanism design.

Theorem 7.

$$1 + \sqrt{\frac{1}{512k}} \leq \beta_k \leq 1 + \sqrt{\frac{8 \ln(k)}{k}}$$

for all sufficiently large k .

For the applications to online mechanism design, it is of course necessary to understand the algorithm which achieves the upper bound in Theorem 7. As in the $k = 1$ case, the algorithm corresponds to a posted-price mechanism with a fixed posted price which does not vary over time. The price, which we denote by m_k , is the infimum of the set of numbers a satisfying $\sum_{i=1}^n \Pr(x_i > m_k) \leq k - \sqrt{2k \ln(k)}$. As above, it is easy to compute m_k in polynomial time using a binary search, provided that the distributions of the variables x_i have finite support and are explicitly specified as part of the input. Large deviation inequalities imply that with high probability the number of bids x_i exceeding the threshold value m_k will be between $k - \sqrt{4k \ln(k)}$ and k , and when this happens the revenue and efficiency will both be within a factor of $1 - O(\sqrt{\ln(k)/k})$ of optimal. This constitutes a proof sketch of the upper bound in Theorem 7. The lower bound (whose proof is omitted for space reasons) arises from considering the following

sequence of independent random variables: x_i is deterministically equal to 1 if $1 \leq i \leq \sqrt{k/8}$, uniformly distributed in $\{0, 2\}$ if $\sqrt{k/8} < i \leq 2k + \sqrt{k/8}$, and deterministically equal to 0 for all larger values of i .

The fact that the price m_k does not vary over time and does not depend on bids received implies, as before, that the allocation rule may be implemented by a posted-price mechanism that satisfies temporal strategyproofness, even when agents can lie arbitrarily about arrival and departure times.

Unknown n and dependent bids

We now turn to cases in which the elements of the sequence x_1, x_2, \dots are not necessarily independent, but their joint distribution is known to the mechanism designer. A special case of this problem arises in the case of independent bids but an unknown value of n which is drawn from some known distribution. For arbitrarily distributed sequences of non-negative random variables, it is impossible to obtain non-trivial prophet inequalities with a multiplicative bound such as (1). However, if we assume that the random variables x_1, x_2, \dots, x_n are uniformly bounded (taking values in $[0, 1]$, without loss of generality), then there are non-trivial additive bounds relating $x_{\tau_1} + \dots + x_{\tau_k}$ to $OPT_k(x_1, \dots, x_n)$. For example, in the case $k = 1$ there always exists a stopping rule τ such that $\mathbb{E}(x_{\tau}) + 1/e \geq OPT_1(x_1, \dots, x_n)$, and the constant $1/e$ is asymptotically the best possible as n tends to infinity (Hill & Kertz 1983). For the case $k > 1$, we present here a simple additive $k/2$ -approximation.

Theorem 8. *For any sequence of random variables x_1, x_2, \dots, x_n taking values between 0 and 1, there is a sequence of k stopping times $\tau_1 < \tau_2 < \dots < \tau_k$ such that*

$$\mathbb{E}(x_{\tau_1} + \dots + x_{\tau_k}) + k/2 \geq \mathbb{E}(OPT_k(x_1, \dots, x_n)).$$

In fact τ_{i+1} may be taken to be the smallest $j > \tau_i$ satisfying $x_j \geq 1/2$, or $\tau_{i+1} = \infty$ if no such j exists.

As in the preceding two subsections, the stopping rules achieving the bound in Theorem 8 correspond to a posted-price mechanism with a price that does not vary over time (namely, a price of $1/2$), so the allocation rule can be implemented by a mechanism which satisfies temporal strategyproofness.

The following theorem demonstrates that, unfortunately, the $O(k)$ additive error term in Theorem 8 can not be improved by more than a constant factor, even when bids are i.i.d. samples from a known distribution and n is randomly sampled from a two-element set.

Theorem 9. *Suppose that n is sampled uniformly at random from the set $\{k, k^3\}$ and that the bids x_1, x_2, \dots, x_n are i.i.d. random samples from the distribution which assigns probability $1 - 1/k$ to the value $1/2$ and probability $1/k$ to the value 1. For any sequence of stopping rules $\tau_1 < \dots < \tau_k$, we have*

$$\mathbb{E}(x_{\tau_1} + \dots + x_{\tau_k}) + \frac{k}{8} - 1 \leq \mathbb{E}(OPT_k(x_1, \dots, x_n)). \quad (2)$$

Conclusions and open problems

In this paper for the first time, we designed automated mechanism design techniques for designing online mechanisms — in order to exploit distributional information about valuations of bidders who arrive online. Sometimes even the number of bidders is not known in advance. Along the way, we identified a rich interplay between these problems and prophet inequalities from statistics. We also proved new prophet inequalities motivated by the auction setting.

This is also a fertile area for future research. Suppose that we have patient (not instantaneous) bidders who have arrival and departure times as well as valuations. Suppose moreover that the mechanism designer knows the joint distribution of the entire input — the number of bidders, and all three parameters of their types. What is the complexity of designing the optimal mechanism? Is there still a dynamic program of subexponential size (or even PSPACE)? If the answer is no, what about approximation? (Pai & Vohra 2006) have a dynamic program for designing an optimal auction, but it has an exponential number of states. Even if bidders cannot lie about their arrival and departure times (only about their valuations) and all valuations are i.i.d. from the same distribution of polynomially-bounded support, it is unclear whether there is an efficient algorithm to design an optimal mechanism. The same questions can be asked in the reusable good setting, i.e., when at every time slot there is one unit which may be allocated at that time but not at any other time.

It would also be desirable to obtain tight upper and lower bounds in Theorem 7 (possibly $\beta_k = 1 + \Theta(\sqrt{1/k})$), Theorem 8 (possibly k/e), and Theorem 2 (possibly $\tilde{O}(\log n)$). Another open (but not practically important) problem is the case of unit supply ($k = 1$) when bidders can have only two possible valuations (say, 1 and h) or when there are only two possible values of n . In the former case, there is a simple 2-competitive randomized algorithm (via the approach of Theorem 2) and in the latter case there is an obvious 2-competitive randomized algorithm which guesses one of the two values uniformly at random, and then applies an optimal pricing policy assuming that this guess is correct. But can we get a competitive ratio better than 2 in either case?

References

- Assaf, D., and Samuel-Cahn, E. 2000. Simple ratio prophet inequalities for a mortal with multiple choices. *J. Appl. Prob.* 37(4):1084–1091.
- Awerbuch, B.; Azar, Y.; and Meyerson, A. 2003. Reducing truth-telling online mechanisms to online optimization. In *ACM Symposium on Theory of Computing*, 503–510.
- Babaioff, M.; Immorlica, N.; and Kleinberg, R. 2007. Matroids, secretary problems, and online mechanisms. In *ACM-SIAM Symposium on Discrete Algorithms*, 434–443.
- Blum, A.; Kumar, V.; Rudra, A.; and Wu, F. 2003. Online learning in online auctions. In *ACM-SIAM Symposium on Discrete Algorithms*, 202–204.
- Blum, A.; Sandholm, T.; and Zinkevich, M. 2006. Online algorithms for market clearing. *Journal of the ACM* 53:845–879.
- Bredin, J., and Parkes, D. C. 2005. Models for truthful online double auctions. In *Conference on Uncertainty in Artificial Intelligence*, 50–59.
- Conitzer, V., and Sandholm, T. 2002. Complexity of mechanism design. In *Conference on Uncertainty in Artificial Intelligence*.
- Friedman, E. J., and Parkes, D. C. 2003. Pricing WiFi at starbucks: issues in online mechanism design. In *ACM Conference on Electronic Commerce*, 240–241.
- Gallien, J. 2006. Dynamic mechanism design for online commerce. *Operations Research* 54:291–310.
- Hajiaghayi, M.; Kleinberg, R.; Mahdian, M.; and Parkes, D. 2005. Online auctions with re-usable goods. In *ACM Conference on Electronic Commerce*, 165–174.
- Hajiaghayi, M. T.; Kleinberg, R.; and Parkes, D. C. 2004. Adaptive limited-supply online auctions. In *ACM conference on Electronic Commerce*, 71–80.
- Hill, T. P., and Kertz, R. P. 1983. Stop rule inequalities for uniformly bounded sequences of random variables. *Trans. Am. Math. Soc.* 278:197–207.
- Kennedy, D. P. 1987. Prophet-type inequalities for multi-choice optimal stopping. *Stoch. Proc. Applic.* 24:77–88.
- Kleinberg, R., and Leighton, T. 2003. The value of knowing a demand curve: Bounds on regret for on-line posted-price auctions. In *Symposium on Foundations of Computer Science*.
- Kleinberg, R. 2005. A multiple-choice secretary algorithm with applications to online auctions. In *ACM-SIAM Symposium on Discrete Algorithms*, 630–631.
- Krengel, U., and Sucheston, L. 1977. Semiamarts and finite values. *Bull. Am. Math. Soc.* 83:745–747.
- Krengel, U., and Sucheston, L. 1978. On semiamarts, amarts, and processes with finite value. In Kuelbs, J., ed., *Probability on Banach Spaces*.
- Lavi, R., and Nisan, N. 2000. Competitive analysis of incentive compatible on-line auctions. In *ACM Conference on Electronic Commerce*, 233–241.
- Lavi, R., and Nisan, N. 2005. Online ascending auctions for gradually expiring items. In *ACM-SIAM Symposium on Discrete Algorithms*, 1146–1155.
- Likhodedov, A., and Sandholm, T. 2005. Approximating revenue-maximizing combinatorial auctions. In *AAAI*.
- Mahdian, M., and Saberi, A. 2006. Multi-unit auctions with unknown supply. In *ACM Conference on Electronic Commerce*, 243–249.
- Myerson, R. 1981. Optimal auction design. *Mathematics of Operation Research* 6:58–73.
- Ng, C.; Buonadonna, P.; Chun, B. N.; Snoeren, A. C.; and Vahdat, A. 2005. Addressing strategic behavior in a deployed microeconomic resource allocator. In *ACM SIGCOMM workshop on Economics of peer-to-peer systems*, 99–104.
- Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. 2007. *Algorithmic Game Theory*. To appear.
- Pai, M., and Vohra, R. 2006. Optimal dynamic auctions. Technical report, Kellogg School of Mgmt., Northwestern University.
- Parkes, D. C., and Singh, S. 2003. An MDP-based approach to online mechanism design. *Conference on Neural Information Processing Systems*.
- Porter, R. 2004. Mechanism design for online real-time scheduling. In *ACM conference on Electronic Commerce*, 61–70.
- Samuel-Cahn, E. 1984. Comparisons of optimal stopping values and prophet inequalities for independent non-negative random variables. *Ann. Prob.* 12:1213–1216.
- Sandholm, T., and Gilpin, A. 2006. Sequences of take-it-or-leave-it offers: Near-optimal auctions without full valuation revelation. In *AAMAS*.
- Sandholm, T.; Conitzer, V.; and Boutilier, C. 2007. Automated design of multistage mechanisms. In *IJCAI*.
- Sandholm, T. 2003. Automated mechanism design: A new application area for search algorithms. In *International Conference on Principles and Practice of Constraint Programming*.