

# An Economic View of Prophet Inequalities

BRENDAN LUCIER

Microsoft Research

---

Over the past decade, an exciting connection has developed between the theory of posted-price mechanisms and the prophet inequality, a result from the theory of optimal stopping. This survey provides an overview of this literature, covering extensions and applications of the prophet inequality through the lens of an economic proof of this classic result. We focus on highlighting ways in which the economic perspective drives new advances in the theory of online stochastic optimization, and vice versa.

Categories and Subject Descriptors: J.4 [Social and Behavioral Science]: Economics

General Terms: Economics, Theory

Additional Key Words and Phrases: Prophet Inequalities, Pricing, Approximation

---

## 1. OVERVIEW

We are tasked with dividing up and selling a pool of resources among rational applicants, and our goal is to make the allocation as efficient (or, perhaps, as profitable) as possible. To what extent are optimal mechanisms approximated by simple sales protocols? This question stands at the forefront of algorithmic mechanism design, despite a lack of consensus on precisely what is meant by “simple.” While it’s admittedly difficult to move past “I’ll know it when I see it,” a natural contender for simplicity is sequential posted pricing. In a posted price mechanism, a seller or platform uses their knowledge of the market to design a menu of prices that are offered to customers as they arrive, who can then purchase whatever they like while supplies last. This is certainly a natural and practical approach to allocating resources, ubiquitous in real-world markets. But it is natural to wonder to what extent the posted-price paradigm approximates the performance of more complex market-resolution methods.

As it turns out, this question is closely related to the so-called *prophet inequality*. The prophet inequality was proven in the 70s in the context of optimal stopping theory, an offshoot of stochastic optimization. Since the connection between the prophet inequality and pricing was introduced to the economics and computation community a decade ago, numerous papers have extended the prophet inequality and used its insights to develop methods of constructing prices for increasingly complex markets.

In this survey, I give a brief overview of this literature through the lens of an economically-oriented proof of the prophet inequality. Beyond the direct applications to pricing and mechanism design, this economic perspective turns out to be useful for extending the reach of the prophet inequality as a tool for stochastic

---

This survey is based on a tutorial given by the author at WINE 2016.

Author’s address: B. Lucier, Microsoft Research, One Memorial Drive, Cambridge, MA 02142 USA. Email: [brlucier@microsoft.com](mailto:brlucier@microsoft.com).

optimization and online algorithms. We will survey these connections through a sequence of applications, beginning with the original prophet inequality and leading to inherently combinatorial and multi-dimensional settings.

**Outline of the Paper.** In Section 2 I review the history of the original prophet inequality in the stopping theory literature, and its connection to simple pricing problems. In Section 3 I present a proof of the prophet inequality that is motivated by the connection to pricing. Then in Section 4 I give an overview of results that extend the prophet inequality to more complex optimization tasks, using the economic proof as a guide but staying within the realm of single-parameter problems. After pausing in Section 5 to describe a unifying framework, I move on to multi-dimensional extensions in Section 6. Then in Section 7 I circle back to the implications for mechanism design. This survey will focus mainly on the connection between prophet inequalities and welfare maximization (i.e., allocating resources for economic efficiency), but at the end of the section I briefly describe applications to revenue maximization as well. Finally, Section 8 concludes and suggests some research directions for further exploration.

## 2. INTRODUCTION: A BRIEF HISTORY OF THE PROPHET INEQUALITY

Imagine that you are invited to play the following game. You are presented with a sequence of  $n$  treasure chests. Each chest contains a cash prize, but the chests are locked and you cannot see their contents. However, each chest has a distribution over non-negative values printed upon it. You are told that the value of the prize in each chest was drawn independently from its displayed distribution. The host running the game will open the chests for you, one at a time. When a chest is opened, you can see the prize and must make a choice. You can either accept the prize, ending the game immediately; or you can reject the prize, in which case it is lost to you forever and the game will proceed with the next chest. How should you play this game to maximize your expected winnings?

The optimal strategy for this game can be worked out by backward induction. If you reach the final chest, then it is certainly optimal to take its prize. This lets you calculate your expected winnings if you play optimally, starting at round  $n$ . Given this, the correct strategy for the second-to-last chest is to accept the prize if and only if its value is greater than your expected winnings if you continue to the final chest. This lets you calculate the expected winnings of the optimal strategy beginning at round  $n - 1$ . Iterating this reasoning yields an optimal strategy for the entire game, corresponding to a non-increasing sequence of acceptance thresholds.

This type of stochastic optimization problem (and its solution) was well-studied in the 70s, in the field of optimal stopping theory. Krengel and Sucheston [Krengel and Sucheston 1978] asked the following question: how does the expected value of the optimal strategy compare with the expected maximum prize? In other words, how well does an optimal game contestant perform relative to an omniscient prophet, who can see inside the chests and therefore trivially wins the best prize every time? Krengel and Sucheston proved a multiplicative bound: the gambler's expected winnings is always at least a quarter of the expected maximum prize. Garling then noted (in a private communication) that this bound could be improved to half the optimal prize [Krengel and Sucheston 1977]. The existence of a 2-approximate

policy for this problem became known as the *prophet inequality*.

In fact, a simple example shows that this approximation factor of 2 is best-possible. Suppose there are only two treasure chests. The first one deterministically contains a prize of 1. The second chest contains  $1/\epsilon$  with probability  $\epsilon$ , where  $\epsilon \in (0, 1]$  is arbitrarily small, and otherwise contains nothing. The maximum prize is  $1/\epsilon$  with probability  $\epsilon$ , and otherwise 1, so the expected maximum prize is  $2 - \epsilon$ . On the other hand, no strategy yields an expected value greater than 1. The first chest always contains 1, and the contestant can either (a) take it (getting 1 for sure) or (b) leave it, which means they must take a chance on the second chest whose expected prize is also 1. There is therefore no policy with approximation factor better than  $2 - \epsilon$  for any  $\epsilon > 0$ .

Later, Samuel-Cahn [Samuel-Cahn 1984] noted that one can achieve this same 2-approximation using a particularly simple, non-optimal strategy: accepting the first prize greater than a certain threshold  $\theta$  [Samuel-Cahn 1984]. Moreover, it suffices to set  $\theta$  equal to the median of the distribution of highest prizes.<sup>1</sup> It is notable that this choice of  $\theta$  is invariant to the order in which the prizes are revealed. This strategy therefore remains 2-approximate even if the order is chosen by an adaptive adversary who, in each round, can select the next chest to open based on the values revealed in the previous chests.

Following the original prophet inequality established by Krengel, Sucheston, and Garling, there was a significant line of work studying this optimality gap under restrictions on the distributions, relaxations of the independence assumption, and various classes of stopping rules and criteria. This literature is far too broad to cover here, but we recommend the survey [Hill and Kertz 1992] for highlights. One notable result from that line of work is that the factor of 2 can be improved if the distributions are identical [Hill and Kertz 1982]. The tight approximation factor of  $\approx 1.342$  for identical distributions was established quite recently in the computer science literature [Abolhassani et al. 2017; Correa et al. 2017].

## 2.1 The Pricing Connection

The prophet inequality was reintroduced to the economics and computation community by [Hajiaghayi et al. 2007], who noted a natural analogy to a simple pricing problem. In this analogy, there is a seller with a single indivisible item to sell, such as a used car. A sequence of  $n$  potential buyers, indexed  $1, \dots, n$ , will approach the seller one by one. Each buyer  $i$  has a private value  $v_i$  for the car, drawn from a distribution  $D_i$  that is known to the seller. These distributions can potentially vary between customers; for example, the seller might infer some information about the buyer's type from their apparel, the questions they ask about the vehicle, and so on. The buyer and seller are free to negotiate using an arbitrary (and possibly randomized) protocol, which ultimately leads to a decision to sell or not, and if so at what price. Once the car is sold, future buyers must be turned away; and if a

<sup>1</sup>In this solution, there is a subtlety surrounding tie-breaking in the event that a prize is precisely equal to  $\theta$ . One could imagine a policy that always selects such a prize, and another policy that always rejects such a prize. Samuel-Cahn shows that at least one of these two policies will yield a 2-approximation, but neither policy works all of the time. Later, we will present another choice of threshold that gives a 2-approximation regardless of the way such ties are broken.

customer decides not to purchase and leaves, they never return.

Viewing the customers as treasure chests, the prophet inequality establishes the existence of a sales protocol that guarantees at least half of the optimal gains from trade (that is, social welfare) in expectation. Taken literally, the prophet inequality would require that each buyer's value be fully revealed to the seller upon arrival, similar to observing the prize when a treasure chest is opened. But recall Samuel-Cahn's solution, which is to simply accept the first buyer whose value exceeds a threshold  $\theta$ . This protocol can be implemented even when buyer values are private, and has a very natural interpretation: post a take-it-or-leave-it price of  $\theta$  on the car, and simply allow each customer to purchase if they so choose, while supplies last.

We can conclude that a simple posted-price mechanism gives a 2-approximation to the optimal welfare, regardless of the order in which customers arrive. Moreover, no other mechanism can improve upon this guarantee, since the lower bound example described above extends directly to any equilibrium of any sales protocol: the first customer is known to the seller to have value 1, so any sales protocol simply reduces to the seller choosing whether or not to sell to the first buyer, and the lower bound follows.

This posted-price mechanism is simple, natural, and practical. It has many desirable properties, even among other pricing methods. Its choice of price is:

**Anonymous.** All customers are offered the same price, regardless of their type distribution.

**Static.** The choice of which price to offer which customer does not change as the mechanism progresses.

**Order-Oblivious.** The pricing rule does not depend on the order in which agents arrive, and in fact the order can be chosen by an adaptive adversary.

Moreover, since each customer is simply offered a take-it-or-leave-it price that they can choose to accept or not, the mechanism is ex post individually rational and incentive compatible in dominant strategies. In fact, it is *obviously strategyproof* [Li 2017], since each customer has at most one strategic decision to make (whether or not to purchase), and their payoff is not affected by anything that occurs after that decision is made. It is also weakly group strategyproof (i.e., resilient to collusion between customers), since no customer can, through his or her actions, reduce the price offered to any other customer. This solution is so appealing that it naturally begs the question: can we extend it to more complex allocation problems as well?

### 3. AN ECONOMIC PROOF OF THE PROPHET INEQUALITY

Before presenting extensions to other optimization problems, we will pause here to present a proof of the prophet inequality. This is a variation on an argument by Kleinberg and Weinberg [Kleinberg and Weinberg 2012], and mathematically the arguments are very similar. Nevertheless, I want to present the proof below because it inherently makes use of the connection between prophet inequalities and posted prices. The resulting economic interpretation will be convenient when we move on to extended settings that require more nuance.

We want to prove that there exists a threshold policy yielding, in expectation, at least half of the expected maximum value. To simplify notation, let us write  $V^*$  for the random variable whose value is  $\max_i v_i$ , the maximum of the  $n$  realized prizes. The threshold policy we will consider is the one that accepts the first prize whose value exceeds  $\frac{1}{2} \mathbb{E}[V^*]$ , if any, where the expectation is over the realization of the prizes.<sup>2</sup> We must show that the expected prize generated by this policy is at least  $\frac{1}{2} \mathbb{E}[V^*]$ .

**THEOREM 3.1** [KLEINBERG AND WEINBERG 2012]. *The policy that accepts the first prize that is at least  $\frac{1}{2} \mathbb{E}[V^*]$  has expected reward  $\frac{1}{2} \mathbb{E}[V^*]$ . This is true regardless of the decision made when a prize is equal to  $\frac{1}{2} \mathbb{E}[V^*]$ .*

To prove this, we will view the threshold policy as setting a price of  $p = \frac{1}{2} \mathbb{E}[V^*]$  on an indivisible good for sale, and then allowing  $n$  bidders, each with a value  $v_i$  drawn from  $D_i$ , to sequentially choose whether or not to purchase. This is certainly equivalent to the threshold policy, with the expected gains from trade (i.e., welfare) taking the role of the expected prize won. These gains are made up of two parts: the expected revenue, which forms the seller's utility, and the expected buyer surplus, which is the sum of buyer utilities. We will bound these two parts separately.

**Revenue:** The expected revenue of this policy is simply  $p$  times the probability that the item is sold. So, by the choice of  $p$ , the expected revenue is

$$\frac{1}{2} \mathbb{E}[V^*] \cdot \Pr [ \text{item is sold} ]. \quad (1)$$

**Buyer Surplus:** Note that if the item is not sold by the time the process reaches buyer  $i$ , then buyer  $i$  can choose to purchase the item if she so desires. The expected utility of buyer  $i$  is therefore at least  $(v_i - p)^+ := \max\{v_i - p, 0\}$ , in the event that the item is not sold before buyer  $i$  has a chance to purchase. But, crucially, whether or not the item is sold prior to  $i$ 's decision is independent of  $i$ 's value. The expected buyer surplus is therefore at least

$$\sum_i \mathbb{E}[(v_i - p)^+] \cdot \Pr [ i \text{ has a chance to purchase} ]. \quad (2)$$

Suppose that the item is left unsold at the end of the process, after all the buyers have come and gone. Then it must be that *every* customer had a chance to purchase. This means that (2) is at least

$$\left( \sum_i \mathbb{E}[(v_i - p)^+] \right) \cdot \Pr [ \text{item is unsold} ].$$

Then since

$$\begin{aligned} \sum_i \mathbb{E}[(v_i - p)^+] &\geq \mathbb{E}[\max_i (v_i - p)^+] \\ &\geq \mathbb{E}[\max_i v_i] - p \\ &= \frac{1}{2} \mathbb{E}[V^*], \end{aligned}$$

<sup>2</sup>Recall that this is not the policy originally proposed by Samuel-Cahn [Samuel-Cahn 1984].

we conclude that the expected buyer surplus is at least

$$\frac{1}{2} \mathbb{E}[V^*] \cdot \Pr[\text{item is not sold}]. \quad (3)$$

Adding (1) and (3), we conclude that the expected welfare generated by this sales process is at least  $\frac{1}{2} \mathbb{E}[V^*]$ , as desired.

### 3.1 Discussion

Why was  $\frac{1}{2} \mathbb{E}[V^*]$  the “right” price to use? We can view this threshold as balancing between two ways in which the prophet inequality can fail. If the price is too low, then a low-valued buyer might purchase, denying the chance to a higher-valued buyer later on. If the price is set too high, it becomes increasingly likely that no buyer purchases and the item is left unsold, generating no welfare.

The main idea behind the proof above is that revenue and surplus cover these two failure cases, respectively. The revenue of the pricing mechanism offsets the expected opportunity cost of serving a customer. That is to say, the price is high enough to discourage purchase by those whose value is too low relative to the expected maximum value.

On the other hand, the aggregate buyer surplus offsets the expected value “left on the table” — i.e., that could have been served in retrospect. Intuitively, if the item often goes unsold, then it must also be the case that each buyer often has an opportunity to purchase. The price  $\frac{1}{2} \mathbb{E}[V^*]$  is low enough that the buyers can successfully leverage these opportunities to obtain high expected utility.

**The Smoothness Connection.** Notably, the latter half of the argument above — i.e., the bound on buyer surplus — requires only that agents have the *opportunity* to purchase. It does not require us to analyze which customer actually buys the item. This is reminiscent of the recent literature on smoothness and the price of anarchy [Roughgarden 2012; Syrgkanis and Tardos 2013], where the goal is to analyze the efficiency of an auction at equilibrium but the nature of the game is too complex to fully characterize equilibrium outcomes. The smoothness paradigm bounds the performance of an auction (or, more generally, a game) using only the existence of profitable actions, regardless of whether or not those actions are actually taken by the participating agents. As it turns out, this connection will be helpful when we later extend the prophet inequality to more general settings. As we shall see, many price of anarchy results established via smoothness have corresponding prophet inequalities. We discuss this connection further in Section 5.

**Comparison with Market Prices:** It is instructive to compare the prophet inequality with the notion of a market-clearing price. Consider a deterministic, full-information version of the problem, where the customer values are known in advance. In this situation, the problem becomes trivial: setting price  $\max_i v_i$  yields the fully efficient outcome, assuming that customers buy in case of indifference. In fact, any price that lies strictly between the maximum value and the second-highest value would result in the efficient outcome, and causes each buyer to be “satisfied” in the sense of always being able to buy the item if they value it more than its price. This is a special case of Walrasian equilibrium, which characterizes a lattice of market-clearing prices that guide the market toward an optimal assignment of

goods [Kelso and Crawford 1982]. In comparison, the prophet inequality price guarantees only an approximation to the optimal welfare, and does not guarantee that each customer ends up with a utility-maximizing purchase, but it applies even when there is uncertainty in the market — that is, when buyer values are stochastic — and does not rely on the nature of tie-breaking.

**Robustness:** In addition to being robust to how ties are broken in case of indifference, this proof of the prophet inequality is also highly robust to market misspecification. Namely, if the price we choose is perturbed slightly — say,  $\frac{1}{2} \mathbb{E}[V^*] \pm \epsilon$  — then the resulting welfare guarantee degrades gracefully to  $\frac{1}{2} \mathbb{E}[V^*] - \epsilon$ .<sup>3</sup> This makes it possible to estimate an appropriate price, given a limited number of samples from the value distributions, with only a small loss in performance. This robustness to sampling errors is interesting in and of itself, and will be particularly useful when we extend to more complex settings and must address concerns of efficient computation. This robustness can be pushed further to derive prophet inequalities with even just a single sample from each distribution [Azar et al. 2014].

#### 4. VARIATIONS: ACCEPTING MULTIPLE PRIZES

We now turn to extensions of the prophet inequality to broader optimization tasks. In this section, we survey extensions of a particular form. There is still a sequence of treasure chests, each containing a single prize drawn independently from a chest-specific distribution, and acceptance or rejection decisions are still irrevocable. However, it may be possible to accept more than one prize. Such a problem is specified by some constraint on the sets of prizes that can be simultaneously accepted.

It will be helpful to fix some notation. We will index the chests by  $1, \dots, n$ . We will assume that the chests arrive in this order for notational convenience, although many results are order-oblivious and hold even when the arrival order can be chosen adversarially. We'll write  $v_i$  for the prize in chest  $i$ , which is a random variable drawn independently from distribution  $D_i$ . We will write vectors in boldface, so that  $\mathbf{v} = (v_1, \dots, v_n)$  is the profile of all prizes. An *allocation* is a choice of which chest(s) to accept. We will write  $\mathcal{F}$  for the set of feasible allocations. We will always assume that  $\mathcal{F}$  is downward-closed, meaning that any subset of a feasible allocation is also feasible. We will write  $\text{OPT}(\mathbf{v})$  for the maximum total value attainable by any feasible allocation, in retrospect.

For example, in the original prophet inequality,  $\mathcal{F}$  consists of all singletons plus the empty allocation, and for any profile of values  $\mathbf{v}$  we have  $\text{OPT}(\mathbf{v}) = \max_i v_i$ .

##### 4.1 Extension 1: Accepting up to $k$ Prizes

In addition to interpreting the prophet inequality in the context of posted prices, Hajiaghayi, Kleinberg, and Sandholm [Hajiaghayi et al. 2007] gave a natural extension to the case where up to  $k \geq 1$  prizes can be accepted. In the pricing analogy,

<sup>3</sup>While it's true that a small change in the price could increase or decrease the probability of sale dramatically, this only causes a shift between the revenue and buyer surplus, and the argument described above still goes through essentially unchanged. One simply uses  $p \geq \frac{1}{2} \mathbb{E}[V^*] - \epsilon$  and  $p \leq \frac{1}{2} \mathbb{E}[V^*] + \epsilon$  in the final lines of the revenue and surplus bounds, respectively.

there are  $k$  identical items for sale, and each buyer is interested in at most one. Not only does the prophet inequality extend, but in fact the achievable approximation factor improves, tending to 1 in the limit as  $k$  grows.

**THEOREM 4.1** [HAJIAGHAYI ET AL. 2007]. *In the prophet inequality setting where  $k \geq 1$  prizes can be accepted, there is a fixed threshold strategy achieving approximation factor at most  $1 + \sqrt{\frac{8 \ln(k)}{k}}$ .*

The idea behind their argument is a natural one: choose a price  $p$  so that the expected number of buyers who want to purchase — i.e., with value at least  $p$  — is a little less than  $k$ : say  $k - \delta$  for some appropriately-chosen value of  $\delta$ . We can interpret this  $p$  as the solution to a fractional relaxation of the problem, with a somewhat stricter feasibility constraint ( $k - \delta$  instead of  $k$ ). If  $\delta$  is chosen properly, standard concentration bounds imply that the actual number of buyers with realized value greater than  $p$  will lie between  $k - 2\delta$  and  $k$ , with high probability.

Let's prove the claim that, for any  $\mathbf{v}$  such that this event occurs, the welfare generated is at least  $(1 - \frac{2\delta}{k})\text{OPT}(\mathbf{v})$ . One way to see this is with a slight twist on the economic proof from Section 3, bounding separately the revenue and buyer surplus. This is very similar to the proof outlined by Hajiaghayi, Kleinberg, and Sandholm [Hajiaghayi et al. 2007].

**Revenue:** Since at least  $k - 2\delta$  customers will purchase, the revenue generated is at least  $(k - 2\delta)p$ .

**Buyer Surplus:** Since at most  $k$  buyers have value greater than  $p$ , each buyer who wanted to purchase at this price had an opportunity to do so. Similar to the single-item case from Section 3, this implies that the total buyer surplus is at least  $\sum_i (v_i - p)^+$ . Since the sum is at least the value of its top  $k$  terms, and since  $\text{OPT}(\mathbf{v})$  is precisely the sum of the  $k$  largest values in  $\mathbf{v}$ , we conclude that the expected buyer surplus is at least

$$\text{OPT}(\mathbf{v}) - kp.$$

Multiplying this bound on the buyer surplus by  $\frac{k-2\delta}{k}$  and adding the revenue completes the claim.

It turns out that setting  $\delta = \sqrt{2k \log k}$  provides an appropriate tradeoff between the probability that the number of customers wishing to purchase lies strictly between  $k - 2\delta$  and  $k$ , and the approximation guarantee of  $(1 - \frac{2\delta}{k})$  subject to that event. This ultimately leads to an approximation factor of  $1 + O(\sqrt{\log k/k})$ .

In addition to this upper bound, [Hajiaghayi et al. 2007] also established a lower bound of  $1 + \Omega(1/\sqrt{k})$  for any online protocol. Alaei [Alaei 2014] later provided an improved upper bound of  $(1 - 1/\sqrt{k+3})^{-1}$ , matching this lower bound. Notably, in addition to being asymptotically optimal, this improved bound recovers the tight bound of 2 for  $k = 1$ .

## 4.2 Extension 2: Matroid Constraints

Matroids are one of the most well-studied classes of downward-closed set systems. A downward-closed feasibility constraint  $\mathcal{F}$  forms a matroid if, for each  $S, T \in \mathcal{F}$ , if  $|S| < |T|$  then there exists some  $a \in T \setminus S$  such that  $S \cup \{a\} \in \mathcal{F}$ . We refer a set



that is feasible with respect to a matroid as *independent*. Matroids generalize the cardinality constraints from Section 4.1. An illustrative example is the graphical matroid: the buyers correspond to the edges in a fixed graph  $G$ , and it is feasible to sell to a subset  $S$  of buyers if and only if  $S$  does not contain a cycle.

Kleinberg and Weinberg [Kleinberg and Weinberg 2012] extended the prophet inequality to all matroid constraints. In particular, for any matroid  $\mathcal{F}$ , they presented a pricing policy that yields a approximation factor of 2. Their protocol is order-oblivious and allows an adaptive adversary to choose the arrival order, but the prices used in their policy are dynamic and personalized (as opposed to static and anonymous).

To define their pricing rule, we must first define what is meant by a residual value. Take any feasible set of buyers  $S$ , and suppose  $T$  is optimal feasible allocation from among those that contain  $S$ . That is,  $T$  maximizes  $\sum_{i \in T} v_i$  subject to  $T \in \mathcal{F}$  and  $S \subseteq T$ . Then the *residual value after  $S$*  is  $\text{OPT}(\mathbf{v} \mid S) = \mathbf{v}(T) - \mathbf{v}(S)$ . In other words, the residual value is the maximum total value remaining, in a world where the elements of  $S$  have already been taken.

The pricing rule of Kleinberg and Weinberg is now as follows. When buyer  $i$  arrives, suppose that  $S$  is the set of buyers that have already been served. If  $S \cup \{i\}$  is feasible, then buyer  $i$  will be offered price

$$p_i = \frac{1}{2} \mathbb{E}[\text{OPT}(\mathbf{v} \mid S) - \text{OPT}(\mathbf{v} \mid S \cup \{i\})]. \quad (4)$$

We can interpret this price as half the opportunity cost of providing service to buyer  $i$ , given the decisions that have already been made by the time buyer  $i$  arrives. If  $S \cup \{i\}$  is infeasible, then buyer  $i$  cannot purchase; we can think of this as setting  $p_i = \infty$ .

**THEOREM 4.2** [KLEINBERG AND WEINBERG 2012]. *Under the pricing scheme described above, the expected welfare generated is at least  $\frac{1}{2} \mathbb{E}[\text{OPT}(\mathbf{v})]$ .*

Why do these prices work? Again, we can take intuition from the economic proof of the prophet inequality. To convey this intuition, let's consider the full-information version of the problem, where all agent values are known in advance.<sup>4</sup> In this case, the prices and purchasing decisions are deterministic; we can write  $S$  for the set of buyers who ultimately purchase, and write  $S_{<i}$  for the set of buyers who purchase before  $i$  arrives. Then if  $S_{<i} \cup \{i\}$  is feasible, the price offered to agent  $i$  is

$$p_i = \frac{1}{2} (\text{OPT}(\mathbf{v} \mid S_{<i}) - \text{OPT}(\mathbf{v} \mid S_{<i} \cup \{i\})). \quad (5)$$

Again, let's bound the revenue and buyer surplus generated by these prices.

**Revenue:** The prices are defined precisely so that the revenue offsets the opportunity cost of serving the buyers who are accepted. Recall that  $S$  is the set of buyers who accept their prices. By a telescoping sum, the total revenue generated is equal

<sup>4</sup>The deterministic setting admits simpler proofs than the one we describe here. We are nevertheless interested in this economic proof in part because it extends directly to the general problem with stochastic values, and builds intuition for that more general argument.

to

$$\begin{aligned} \sum_{i \in S} p_i &= \sum_{i \in S} \frac{1}{2} (\text{OPT}(\mathbf{v} \mid S_{<i}) - \text{OPT}(\mathbf{v} \mid S_{<i} \cup \{i\})) \\ &= \frac{1}{2} (\text{OPT}(\mathbf{v}) - \text{OPT}(\mathbf{v} \mid S)). \end{aligned} \quad (6)$$

**Buyer Surplus:** We claim that the prices are set low enough that the total buyer surplus is at least half of the residual value left when the mechanism completes. The proof of this claim requires some facts about matroids, and we will not reproduce it here, but it is captured in the following proposition:

**PROPOSITION 4.3** [KLEINBERG AND WEINBERG 2012]. *For any sets  $S$  and  $T$  such that  $S \cup T \in \mathcal{F}$ , the sum of prices offered to  $T$  is at most  $\frac{1}{2} \text{OPT}(\mathbf{v} \mid S)$ .*

If  $S$  is the set of agents who purchase, and  $T$  is the set of maximum total value such that  $S \cup T$  is feasible, then (by downward-closedness of  $\mathcal{F}$ ) it must be that each customer in  $T$  had the opportunity to purchase. We therefore have that the total buyer surplus is at least  $\sum_{i \in T} (v_i - p_i)$ , which by Proposition 4.3 is at least  $\sum_{i \in T} v_i - \frac{1}{2} \text{OPT}(\mathbf{v} \mid S)$ . But from our choice of  $T$ ,  $\sum_{i \in T} v_i = \text{OPT}(\mathbf{v} \mid S)$ , and hence the total buyer surplus is at least

$$\frac{1}{2} \text{OPT}(\mathbf{v} \mid S). \quad (7)$$

Adding the revenue bound (6) to the buyer surplus bound (7) gives the desired welfare bound of  $\frac{1}{2} \text{OPT}(\mathbf{v})$ .

We went through the reasoning above for the full-information version of the problem. In fact, the argument extends almost immediately to arbitrary value distributions, by considering the *expected* revenue and buyer surplus and making liberal use of the magic of linearity of expectation.<sup>5</sup> Indeed, we can (and should) interpret each price from (4) as precisely the expectation — over the distribution of buyer types — of the price (5) we chose to post in the full-information version of the problem. This has the flavor of an extension theorem: to solve the stochastic version of the problem, it suffices to produce a sufficiently well-behaved solution to the full-information version. The general solution then follows by taking expectations. As we will see, this is no accident — a similar extension result captures many prophet inequality proofs in the literature, and we will formalize this in Section 5.

**Beyond a Single Matroid.** In addition to proving a prophet inequality for matroid constraints, Kleinberg and Weinberg showed that this extends to a approximation factor of  $(4k - 2)$  for the intersection of  $k$  matroids, and that this linear dependence on  $k$  is necessary. Feldman, Svensson and Zenklusen later improved this upper bound to  $e(k+1)$  [Feldman et al. 2015]. Dütting and Kleinberg extended the matroid result to a 2-approximate prophet inequality for polymatroids, which are natural convex relaxation of matroid constraints [Dütting and Kleinberg 2015].

<sup>5</sup>There is an important subtlety when bounding the expected buyer surplus. In particular, one has to be careful about introducing correlation between  $S_{>i}$  and  $T$ . A common approach is to take expectations twice: once to determine  $S$ , and then a second time to determine  $T$ . We discuss this further when we present a more general framework in Section 5.

Finally, [Rubinstein and Singla 2017] extend the matroid prophet inequality to settings where the objective function is not necessarily additive over the elements selected, but can instead be an arbitrary submodular function; they establish a prophet inequality with constant approximation factor for this problem.

### 4.3 Extension 3: Knapsack Constraints

In a knapsack constraint, each customer  $i$  is associated with a size  $s_i \in [0, 1]$  in addition to their value  $v_i$ . The feasibility constraint is that a set  $S$  of customers can be simultaneously satisfied if and only if  $\sum_{i \in S} s_i \leq 1$ .

**Fractional Knapsack.** Feldman, Svensson, and Zenklusen considered prophet inequalities for knapsack constraints, but under a fractional relaxation of the problem [Feldman et al. 2015]. In this relaxed problem, we can allocate each agent a level of service  $x_i \in [0, 1]$ . The total value generated is then  $\sum_i x_i v_i$ , and the feasibility constraint is  $\sum_i x_i s_i \leq 1$ . We recover the original knapsack problem by restricting each  $x_i$  to lie in  $\{0, 1\}$ . Note that, unlike the previous examples, here we can serve each customer in multiple possible ways. The interpretation of a prophet inequality in this setting is that each customer, upon arrival, must be assigned a service level  $x_i \in [0, 1]$ . This decision is irrevocable and cannot be modified once made. Once each customer's level of service has been determined, the next customer arrives. The pricing analogy for this problem is natural: there is a single unit of a divisible resource for sale, and each customer  $i$  has value  $v_i/s_i$  per unit of good received, up to a maximum of  $s_i$  units.

A prophet inequality with constant approximation factor for fractional knapsack constraints was first established by [Feldman et al. 2015]. The approximation factor was subsequently improved to 2 by [Dütting et al. 2017]. Moreover, this 2-approximate prophet inequality applies even if, for each customer  $i$ , both the size  $s_i$  and value  $v_i$  are privately known to the customer and arbitrarily correlated with each other. That is, one can think of  $(v_i, s_i)$  as the *type* of customer  $i$ , and we allow types to be drawn from some (known) distribution over tuples.

The approach of Dütting et al. is to post a static, anonymous price  $p$  per unit of the resource, then offer this price to each arriving customer and allow them to purchase as much of the good as desired, while supplies last. If we write  $V^*$  for the random variable denoting the total value of the optimal allocation (over randomness in the type profile), the price posted by the seller is  $p = \frac{1}{2} \mathbb{E}[V^*]$ .

**THEOREM 4.4.** *For the knapsack problem described above, the policy that posts a per-unit price of  $\frac{1}{2} \mathbb{E}[V^*]$  and allows each buyer to purchase their desired quantity (up to the amount remaining) has approximation factor 2.*

Why is this the right price to post? As with matroid constraints, it is useful to first consider the full-information version of the problem, where the customer types are known to the seller in advance. In this case, write  $\mathbf{x}^*$  for the optimal allocation, which has total value  $V^* = \sum_i v_i x_i^*$ . Then the seller will post price  $\frac{1}{2} V^*$ . Again, we will bound separately the revenue and buyer surplus of this pricing policy.

**Revenue:** Suppose that a total of  $Y \leq 1$  units of the good are sold by the pricing policy. Then the revenue generated is simply  $p \cdot Y$ , which is equal to

$$\frac{1}{2} V^* \cdot Y. \tag{8}$$

**Buyer Surplus:** Note that since  $Y$  units are sold in total, there are at least  $1 - Y$  units available to be purchased by each customer. In particular, there is enough left over to provide each customer with their optimal allocation,  $s_i x_i^*$ , scaled down by a factor of  $(1 - Y)$ . So each customer must get at least as much utility as they would by purchasing  $s_i x_i^*(1 - Y)$  units. This yields a welfare bound of

$$\begin{aligned} \sum_i (v_i - p s_i) x_i^*(1 - Y) &= \left( \sum_i (v_i x_i^*) - p \cdot \sum_i (s_i x_i^*) \right) \cdot (1 - Y) \\ &\geq (V^* - p) \cdot (1 - Y) \\ &= \frac{1}{2} V^* (1 - Y) \end{aligned} \tag{9}$$

where in the inequality we used the fact that  $\sum_i s_i x_i^* \leq 1$ . Adding the revenue and buyer surplus bounds from (8) and (9), we have that the total welfare is at least  $\frac{1}{2} V^*$ .

As with the argument for matroid constraints, this bound extends directly to the general stochastic version of the problem by taking expectations and exploiting linearity. The expected revenue will be  $p = \frac{1}{2} \mathbb{E}[V^*]$  times the expected quantity sold, which is  $\frac{1}{2} \mathbb{E}[Y]$ . To bound the expected buyer surplus, we note that each buyer  $i$  obtains at least as much utility as they would if they used the following strategy: purchase the quantity they *expect* to receive in the optimal allocation (given their type, with expectation taken over the types of the other customers), scaled down by the amount of resource available. This gives an expected surplus that is at least  $(\mathbb{E}[V^*] - p) = \frac{1}{2} \mathbb{E}[V^*]$  times the expected quantity left over,  $\mathbb{E}[1 - Y]$ . Adding these two quantities gives an expected welfare bound of  $\frac{1}{2} \mathbb{E}[V^*]$ .

**Integral Knapsack.** The argument above applied to a fractional knapsack constraint. Dütting et al. further established a prophet inequality with approximation factor 5 for integral knapsack constraints [Dütting et al. 2017]. This result also employs a static, anonymous price.

Their argument follows the same general pattern as the arguments we have seen so far, but considers separately the contribution of buyers with  $s_i > 1/2$  (the “large” buyers) and buyers with  $s_i \leq 1/2$  (the “small” buyers) to the optimal welfare. The large buyers are easy to handle: since an optimal allocation can accept at most one large buyer, the original prophet inequality implies the existence of a single take-it-or-leave-it price, for the entire quantity of good, that gives a 2-approximation to the welfare contribution of large buyers. For small buyers, a slight variation on the fractional knapsack argument above, using  $\frac{2}{3} V^*$  for the per-unit price rather than  $\frac{1}{2} V^*$ , yields a 3-approximation to the optimal welfare attainable from small buyers. Combining these two bounds and considering the worst case over combinations of large and small buyers leads ultimately to an approximation factor of 5.

Finally, a brief note about computation. The prices for the “small-buyer” case were based on an optimal allocation  $\mathbf{x}^*$ , but recall that finding such an optimal allocation is an NP-hard task! Fortunately, this argument applies even if  $\mathbf{x}^*$  is only an approximately optimal allocation, in which case the approximation factor of the prophet inequality degrades by the approximation factor of the allocation method. In particular, we can use the allocation returned by an FPTAS for knapsack [Ibarra

and Kim 1975] and obtain a  $(5 + \epsilon)$ -approximate prophet inequality using prices that can be computed in polynomial time.

#### 4.4 A Lower Bound: General downward-closed constraints

Given our success so far at deriving prophet inequalities with constant approximation factors, it is tempting to ask whether there is a constant prophet inequality for arbitrary downward-closed feasibility constraints. As it turns out, this is not the case. Babaioff, Immorlica and Kleinberg established a lower bound for secretary problems, which also implies the following lower bound for prophet inequalities.

**THEOREM 4.5** [BABAI OFF ET AL. 2007]. *There exists a downward-closed feasibility constraint for which no online acceptance protocol can achieve approximation factor better than  $\Omega(\log n / \log \log n)$ . In this example, the values are drawn i.i.d. from a distribution supported on  $\{0, 1\}$ , and the arrival order is non-adaptive.*

**An Upper Bound.** A prophet inequality for arbitrary downward-closed feasibility constraints was established by [Rubinstein 2016]. The policy is randomized and yields an approximation factor of  $O(\log n \log r)$ , where  $r$  is the size of the largest feasible set. It remains open to determine whether this dependence on  $r$  can be removed, or whether a similar approximation can be obtained with a deterministic, price-based policy.

## 5. INTERLUDE: BALANCED PRICES

The arguments from the previous section all follow a similar general paradigm. For any given deterministic set of values, exhibit prices that satisfy two properties: (a) the prices are high enough that the revenue generated from any outcome offsets the opportunity cost of making those selections, and (b) the prices are low enough that the buyer surplus offsets any residual value left on the table when the process ends. Posting these prices leads to an approximation result for the deterministic optimization problem. The more general stochastic result then follows by taking expectations. In this section, we make this extension theorem explicit.

Recalling that prices might be dynamic (as for matroids) and the outcome space for each buyer might not be binary (as for fractional knapsack), we'll need to introduce some notation. We will represent an allocation by a vector  $\mathbf{x} = (x_1, \dots, x_n)$ , with  $x_i$  representing the decision made in round  $i$  of the game. As before, we will tend to assume that customers arrive in the order they are indexed for convenience, although most results we discuss are order-oblivious. We'll say that a *pricing rule* assigns a price  $p_i(x \mid \mathbf{y})$  to each potential outcome  $x$  for agent  $i$ , given that partial allocation  $\mathbf{y}$  has already been made and  $x$  is feasible given  $\mathbf{y}$ . We'll write  $\mathbf{x}_{<i}$  for the allocation made to agents prior to  $i$ , so that  $p_i(x_i \mid \mathbf{x}_{<i})$  is the price offered to agent  $i$ , in order to get outcome  $x_i$ , given that the previous buyers made selections  $x_1, \dots, x_{i-1}$ .

Given some allocation  $\mathbf{x}$ , we'll write  $\mathcal{F}_{\mathbf{x}}$  for the set of “residual” allocations that are disjoint from  $\mathbf{x}$ , and can be combined with  $\mathbf{x}$  to form a feasible allocation. For example, if  $\mathcal{F}$  is the set of all allocations that accept at most 5 prizes, and  $\mathbf{x}$  is the allocation that accepts prizes 4 and 7, then  $\mathcal{F}_{\mathbf{x}}$  contains all allocations that accept at most 3 prizes, none of which are 4 or 7. Recalling the notion of marginal value

from the matroid extension, we will write  $\text{OPT}(\mathbf{v} \mid \mathbf{x})$  for the maximum possible residual value over all allocations in  $\mathcal{F}_{\mathbf{x}}$ . That is,  $\text{OPT}(\mathbf{v} \mid \mathbf{x}) = \max_{\mathbf{y} \in \mathcal{F}_{\mathbf{x}}} \sum_i v_i(y_i)$ .

**Balanced Prices.** The following definition describes what it means for a set of prices to be “balanced” for a deterministic set of values  $\mathbf{v}$ . This is a variation on a notion of balanced thresholds due to [Kleinberg and Weinberg 2012]. The main difference is that the definition below is extended to allow a more general set of outcomes, which will be useful when we consider multi-dimensional settings in Section 6. We also note that Kleinberg and Weinberg considered balancedness with respect to a profile of distributions, rather than a fixed valuation profile.

*Definition 5.1.* For  $\alpha, \beta > 0$ , a pricing rule  $\mathbf{p}$  is  $(\alpha, \beta)$ -balanced for valuation profile  $\mathbf{v}$  if for all  $\mathbf{x} \in \mathcal{F}$ ,

- (1)  $\sum_{i \in N} p_i(x_i \mid \mathbf{x}_{<i}) \geq \frac{1}{\alpha} \cdot (\text{OPT}(\mathbf{v}) - \text{OPT}(\mathbf{v} \mid \mathbf{x}))$ , and
- (2) for all  $\mathbf{x}' \in \mathcal{F}_{\mathbf{x}}$ :  $\sum_{i \in N} p_i(x'_i \mid \mathbf{x}_{<i}) \leq \beta \cdot \text{OPT}(\mathbf{v} \mid \mathbf{x})$ .

If in the second condition we replace the right hand side with  $\beta \cdot \text{OPT}(\mathbf{v})$  (a weaker condition), we say  $\mathbf{p}$  is *weakly*  $(\alpha, \beta)$ -balanced.

The definition of  $(\alpha, \beta)$ -balancedness captures sufficient conditions for a posted-price mechanism to guarantee high welfare when agents have a known valuation profile  $\mathbf{v}$ . Condition 1 precisely states that the revenue generated by allocating some  $\mathbf{x}$  is at least some fraction of the value lost due to allocating  $\mathbf{x}$ . Condition 2 states that, for any residual allocation  $\mathbf{x}'$  that is possible after  $\mathbf{x}$  has been allocated, the total price of  $\mathbf{x}'$  is not too large relative to the optimal residual welfare. As we’ve now shown in multiple examples, these two conditions are sufficient to derive a welfare bound when the values are known in advance. In the relaxation to weak balancedness, we require only that the total price of any allocation is not too large relative to the optimal (non-residual) welfare. Since any residual welfare is bounded by the optimum total welfare, this is indeed a weaker requirement.

**An Extension Theorem.** Our interest in  $(\alpha, \beta)$ -balanced pricing rules comes from the fact that their implied welfare bounds for deterministic instances extend to stochastic settings as well.

**THEOREM 5.2** [DÜTTING ET AL. 2017]. *Suppose that, for each valuation profile  $\mathbf{v}$  in the support of the buyers’ distributions, the pricing rule  $\mathbf{p}^{\mathbf{v}}$  is  $(\alpha, \beta)$ -balanced for  $\mathbf{v}$ . Then for  $\delta = \frac{\alpha}{1+\alpha\beta}$  the posted-price mechanism with pricing rule  $\delta\mathbf{p}$ , where  $p_i(x_i \mid \mathbf{y}) = \mathbb{E}_{\tilde{\mathbf{v}}}[p_i^{\tilde{\mathbf{v}}}(x_i \mid \mathbf{y})]$ , generates welfare at least  $\frac{1}{1+\alpha\beta} \cdot \mathbb{E}_{\mathbf{v}}[\text{OPT}(\mathbf{v})]$ . If instead  $\mathbf{p}^{\mathbf{v}}$  is weakly  $(\alpha, \beta)$ -balanced, then we can take  $\delta = \min\{\alpha, 1/2\beta\}$  and the welfare generated is at least  $\frac{1}{4\alpha\beta} \cdot \mathbb{E}_{\mathbf{v}}[\text{OPT}(\mathbf{v})]$ .*

The intuition behind Theorem 5.2 follows the economic prophet inequality arguments we have seen so far. The definition of balancedness almost directly implies a welfare result for a deterministic value profile, with a revenue bound following directly via telescoping sum and a surplus bound following by having each buyer consider purchasing  $x'_i$ , where  $\mathbf{x}' \in \mathcal{F}$  is the welfare-optimal allocation in  $\mathcal{F}_{\mathbf{x}}$ . The extension result then follows by taking expectations over value profiles. Some care

is needed to maintain independence between purchasing decisions when taking expectations in this way; we refer the interested reader to [Dütting et al. 2017] for further details.<sup>6</sup>

*Example 5.3.* As a simple example of how to apply Theorem 5.2, consider again the standard prophet inequality from Section 3. Here, each  $x_i$  is either 1 (accept) or 0 (reject). Consider the pricing rule given by  $p_i(1 | \mathbf{x}) = \max_i v_i$ . That is, each customer is offered a price of  $\max_i v_i$  for purchasing, as long as it is feasible to do so. This price rule is static and anonymous. It is also easy to check that it is (1, 1)-balanced: for the first condition, the sum of prices for any non-empty allocation is  $\max_i v_i$ , which is equal to the optimal outcome. For the second condition,  $\mathbf{x}'$  is non-trivial only if  $\mathbf{x}$  does not allocate the item, in which case the price of allocating to any agent is  $\max_i v_i$  which, again, is equal to  $\text{OPT}(\mathbf{v} | \mathbf{x})$ . Taking  $\delta = \frac{\alpha}{1+\alpha\beta} = \frac{1}{2}$ , our theorem directly implies that setting price  $\frac{1}{2} \mathbb{E}[\max_i v_i]$  leads to the original prophet inequality with approximation factor 2.

Theorem 5.2 is stated for general pricing rules. Note, however, that if  $\mathbf{p}^{\mathbf{v}}$  is static, anonymous, and/or order-oblivious for each  $\mathbf{v}$ , then the prices suggested by Theorem 5.2 inherit these properties as well.

The strength of Theorem 5.2 comes from the fact that it's sufficient to consider only deterministic instances when constructing prices. As it turns out, this simplification will be particularly helpful when we consider broader, multi-dimensional scenarios in the next section.

**The Smoothness Connection Revisited.** A crucial aspect of Theorem 5.2 is that, when bounding the performance of the price-based policy, it is not necessary to explicitly analyze which items are purchased by which buyers. When bounding the buyer surplus, it suffices to consider only potential actions of the buyers, and in particular to argue that prices are low enough to enable certain “high-utility” buying strategies. This trick has the most bite when buyers have many possible actions available, such as when there are multiple items for sale; we consider such multi-item scenarios in the next section. As we noted in Section 3, this argument is reminiscent of the so-called smoothness method for bounding the welfare of a mechanism at equilibrium [Roughgarden 2012; Syrgkanis and Tardos 2013]. In a smoothness argument, one derives a bound on the utility of buyer by considering a simple “deviation” strategy that they could have played. Similarly, our proof bounds each buyer’s utility by considering a certain canonical item — or set of items — that they could have purchased. Mathematically these arguments have similar structure, and indeed the prophet inequalities we will establish in Section 6 using Theorem 5.2 have approximation factors that are similar to known smoothness bounds for related auctions [Christodoulou et al. 2016; Syrgkanis and Tardos 2013].

<sup>6</sup>Roughly speaking, when bounding the buyer surplus, it is important to maintain independence between what agent  $i$  considers purchasing and what subsequent agents actually purchase. This can be handled by taking expectations twice: once to determine the actual agent values, and hence what agents will actually purchase; and then a second time to define the  $x'_i$  that agent  $i$  considers purchasing. This can be interpreted as having the agent “hallucinate” random valuations  $\mathbf{v}'$  for the other agents, and consider purchasing according to the residual allocation  $\mathbf{x}'$  that would be welfare-optimal for  $\mathbf{v}'$ . See [Dütting et al. 2017] for further details.

## 6. MULTI-DIMENSIONAL PROPHET INEQUALITIES

So far we have focused on situations where a single value  $v_i$  is revealed in each round, and the decision-maker chooses whether or not to accept the value (or, in some cases, a level of fractional acceptance). In this section, we turn to scenarios where the decision-maker must make a more involved decision in each round. For example, there may be multiple chests opened each round, and the optimizer can accept some subset of them. In terms of our pricing analogy, this might correspond to a seller with multiple heterogeneous products to sell. As in the case of a single item, we will be particularly interested in *price-based* solutions: the seller assigns a price to each potential choice, and then for each buyer the allocation chosen is the one that maximizes the buyer's utility at those prices.

This section is based heavily on [Feldman et al. 2015], which studies high-welfare posted-price mechanisms for combinatorial auctions, and [Dütting et al. 2017], which applies the concept of balanced prices to a variety of stochastic optimization problems.

### 6.1 Matching Markets

Let's begin with a simple extension of the standard prophet inequality to a multi-item setting. In this scenario, there are multiple items for sale, and each arriving customer  $i$  has some value  $v_{ij}$  for each item  $j$ . The values are assumed to be independent across customers, but a single customer's values can be arbitrarily correlated across items. The customers are unit-demand, which is to say that each customer wants at most a single item. Here we let  $x_i$  be the item allocated to agent  $i$ , or  $\emptyset$  if agent  $i$  obtains no item.

This scenario was studied by [Alaei et al. 2012], who presented a 2-approximate prophet inequality for this setting. Here we will present a different argument, which establishes a 2-approximate prophet inequality by way of static item prices. That is, the seller will choose a price for each item, and offer this menu of prices to each customer. Each customer, upon arrival, can purchase whichever item they prefer most at the given prices. The prices will be static, anonymous, and order-oblivious.

What price should be set on item  $j$ ? Guided by Theorem 5.2, we will focus on constructing balanced prices for a given value profile  $\mathbf{v}$ . Write  $V_j^*$  for the value generated by item  $j$  in the optimal assignment. That is, if the optimal allocation assigns item  $j$  to buyer  $i$ , then  $V_j^*$  will be equal to  $v_{ij}$ . We will then take the (static and anonymous) price of item  $j$  to be  $p_j = V_j^*$ .

We claim that these prices are  $(1, 1)$ -balanced. For the first condition of balancedness, note that the revenue generated by selling some set  $S$  of items is equal to  $\sum_{j \in S} V_j^*$ , which is at least the loss in optimal value suffered if those items were removed from the market. For the second condition, note that if some set  $S$  of items were sold, then the total residual welfare available in the remaining items is at least  $\sum_{j \notin S} V_j^*$ , which is equal to the sum of prices of all the remaining items.

Theorem 5.2 then implies that the static, anonymous prices  $p_j = \frac{1}{2} \mathbb{E}[V_j^*]$  imply a prophet inequality with approximation factor 2. In other words, in order to determine the price of item  $j$ , we should take the expected contribution of item  $j$  to the optimal welfare, over all value realizations, and divide this by 2.

*Example 6.1.* Suppose there are two items,  $\{a, b\}$ , and two buyers,  $\{1, 2\}$ . Buyer



1's values are deterministic:  $v_{1a} = 5$  and  $v_{1b} = 6$ . Buyer 2's value for item  $a$  is deterministically  $v_{2a} = 3$ , but her value for  $b$  is stochastic and equally likely to be either 0 or 10. If  $v_{2b} = 0$ , then the optimal assignment gives  $a$  to buyer 2 and  $b$  to buyer 1, yielding  $V_a^* = 3$  and  $V_b^* = 6$ . If  $v_{2b} = 10$ , then the optimal assignment gives  $a$  to buyer 1 and  $b$  to buyer 2, leading to  $V_a^* = 5$  and  $V_b^* = 10$ . Taking half the average of each item's welfare contribution, the pricing scheme above would post item prices  $p_a = \frac{1}{2} \cdot \frac{3+5}{2} = 2$  and  $p_b = \frac{1}{2} \cdot \frac{6+10}{2} = 4$ .

**Submodular Combinatorial Auctions.** The argument above extends directly to a richer space of combinatorial auctions with submodular valuations. In this setting there are still multiple items for sale, but now each buyer can purchase multiple items. We can then think of each  $x_i$  as an arbitrary subset of the set  $M$  of all items, and the feasibility constraint is that  $x_i \cap x_k = \emptyset$  for any two buyers  $i$  and  $k$ . The items exhibit decreasing marginal values for each buyer.

To extend our 2-approximate prophet inequality to this setting, the idea is to consider the so-called XOS representation of a valuation [Lehmann et al. 2001], which expresses the valuation function as a maximum over linear functions. This allows us to define the contribution of an item  $j$  to the welfare generated by an allocation  $\mathbf{x}^*$ : if  $j$  is assigned to buyer  $i$ , then  $V_j^*$  is simply the weight of  $j$  in the additive function that defines the value  $v_i(x_i^*)$ . As in the unit-demand case, pricing each item  $j$  at half of its expected contribution to the social welfare leads to a 2-approximate prophet inequality. We refer the interested reader to [Feldman et al. 2015] for more details.

**Multiple Copies of Each Item.** In addition to establishing a 2-approximate prophet inequality for matching markets, [Alaei et al. 2012] showed that the approximation factor improves to  $(1 - 1/\sqrt{k} + 3)^{-1}$  when there are at least  $k$  identical copies of each item. This directly extends the result of [Alaei 2014] for cardinality constraints discussed in Section 4.1, yielding the same approximation factor in a multidimensional setting. This analysis was then extended further by [Alaei et al. 2013] to obtain a similar approximation factor for the generalized assignment problem.

## 6.2 Multiple-Choice Matroid Constraints

When we sketched the argument for the matroid prophet inequality in Section 4.2, we essentially established that a certain dynamic pricing policy is  $(1, 1)$ -balanced. As it turns out, this argument extends directly to a multi-dimensional version of the matroid prophet inequality as well. Instead of taking a matroid constraint over buyers directly, imagine that there is a universe of elements  $E$  partitioned into disjoint subsets  $E_1, \dots, E_n$ . The decision maker must allocate to each buyer  $i$  a subset  $x_i \subseteq E_i$ . The feasibility constraint is defined by a matroid over  $E$ , and an allocation is feasible if and only if  $\cup_i x_i$  is independent according to that matroid. In other words, there is a matroid feasibility constraint over elements, and multiple elements can be presented to the optimizer simultaneously. Each element has an associated value, and the value of a subset  $x_i$  is simply the sum of the individual item values. The proof we sketched in Section 4.2 extends directly to this multi-dimensional setting, leading again to a prophet inequality with approximation factor

2 via posted prices. As for single-dimensional matroid constraints, the prices applied here will be dynamic and personalized, though the approximation is order-oblivious.

Recall that it suffices to define a balanced pricing policy for a fixed valuation profile  $\mathbf{v}$ . The pricing policy we use prices each collection of elements at half of the opportunity cost they impose on the residual value. That is, for each  $i$ , if  $\mathbf{x}_{<i}$  is the collection of allocations made previously, then the price of some feasible allocation  $x_i$  (which, recall, is a subset of  $E_i$ ) is set to

$$p_i(x_i \mid \mathbf{x}_{<i}) = \frac{1}{2} (\text{OPT}(\mathbf{v} \mid \mathbf{x}_{<i}) - \text{OPT}(\mathbf{v} \mid \mathbf{x}_{\leq i})).$$

Note that these prices may not take the form of prices on individual elements, but rather might be arbitrary bundle prices, as illustrated in the following example.

*Example 6.2.* Suppose the constraint is that at most 2 elements can be accepted. In the first round, the optimizer is presented with two elements whose values are both deterministically 2, and in the second round the optimizer is presented with a single element whose value is deterministically 10. In this case, the price assigned to any single element in the first round is 1 (half of 2) but the price assigned to the pair of both elements is 6 (half of 12, the value of the optimal solution).

As with matching markets, this pricing strategy extends naturally to cases where the value of agent  $i$  for a set of elements  $x_i$  is not necessarily additive, but can be an arbitrary submodular valuation over the set of elements selected. This yields a 2-approximate prophet inequality for this setting as well [Dütting et al. 2017].

### 6.3 Combinatorial Auctions

We next consider a combinatorial auction problem. There are  $m$  items for sale and  $n$  buyers. Each buyer  $i$  has a valuation function  $v_i$  that assigns non-negative value to every subset of at most  $d$  items. There is an  $O(d)$ -approximate algorithm for maximizing welfare in the offline version of this problem, and a lower bound of  $\Omega(d/\log d)$  assuming  $P \neq NP$  [Trevisan 2001]. Our goal is to match this  $O(d)$  approximation with a prophet inequality using posted item prices.

We will again consider the simpler full information case where all valuations are known in advance. Note that the pricing problem is still non-trivial in the deterministic case, and in fact there may not exist prices that lead to the optimal allocation.<sup>7</sup> We will construct balanced prices as follows: given valuation profile  $\mathbf{v}$ , consider the welfare-maximizing allocation  $\mathbf{x}^*$  (which we can assume allocates all items). Then for each item  $j$ , say  $j \in x_i^*$ , set the price of  $j$  to  $p_j = v_i(x_i^*)/|x_i^*|$ . These prices are set so that the total price of all items is precisely  $\text{OPT}(\mathbf{v})$ . At the same time, prices are high enough that, for any set of goods  $S$ , the total price of  $S$  is at least  $1/d$  of the value of allocations in the optimal allocation  $\mathbf{x}^*$  that intersect  $S$ , which bounds the total loss of welfare if the goods in  $S$  were removed. This establishes

<sup>7</sup>For example, suppose there are three items and four single-minded bidders. The first three bidders each have value 2 for a different pair of items, and the last bidder has value 3 for the set of all three items, so at most one bidder can get positive value, and it is optimal to allocate all items to the last bidder. However, at any item prices where the last bidder is willing to purchase, one of the other bidders will purchase first. This leads to a  $3/2$  approximation.

that these prices are weakly  $(d, 1)$ -balanced. Theorem 5.2 then immediately implies that these prices lead to a prophet inequality with approximation factor  $4d$ .

This argument can be extended to a broader class of MPH combinatorial auctions, which restricts the extent of complementarities between items in the preferences of any buyer [Feige et al. 2015]. It turns out that a natural extension of the price scheme described above is weakly  $(d, 1)$ -balanced for MPH- $d$  combinatorial auctions, leading again to a price-based prophet inequality with an  $O(d)$  approximation factor, using static and anonymous item prices.

#### 6.4 Discussion: Computation

The arguments above establish existential prophet inequalities. But in these inherently combinatorial settings, we must consider whether the appropriate prices can be computed. There are two barriers here. First, it is necessary to compute the required balanced prices for any given fixed valuation. Second, it is necessary to take expectations with respect to the distribution over all valuation profiles. For the second issue, it is helpful that the welfare implication of balanced prices is inherently robust as discussed in Section 3: perturbing prices by some small amount  $\epsilon$  will in turn degrade the welfare guarantee by a small amount. It is therefore possible to estimate the necessary prices by sampling a polynomial number of value profiles, computing balanced prices for each, and using the average of the prices computed.

The first issue is more fundamental. In each of the pricing schemes described above, the first step is to find an optimal allocation; this is then transformed into appropriate prices by “dividing” the welfare among the individual elements. This approach is inherently intractable when the allocation problem is NP-hard. One way to approach this issue is to base prices instead on an approximately optimal allocation, as we did for the knapsack prophet inequality. It turns out that this is consistent with the balanced-price approach; it simply degrades the final approximation factor of the prophet inequality by the approximation factor of the allocation algorithm being used. This approach was used to construct an  $O(1)$ -approximate prophet inequality for submodular combinatorial auctions, and an  $O(d^2)$  approximation for combinatorial auctions with set size  $d$  [Feldman et al. 2015].

A different approach is to consider instead a fractional relaxation of the allocation problem, which can be solved exactly in polynomial time. Many of the pricing schemes discussed so far extend naturally to fractional allocations as well, leading to balanced prices for those relaxed problems. The result is a prophet inequality, based on posted prices, that applies when buyers are allowed to make fractional purchases. A natural thought, then, is to simply keep the same prices, but remove all fractional allocations from the menu. In other words, use the fractional version of the problem to compute prices, then use those prices as a solution to the original, integral problem. As it turns out, for some problems this restriction does not degrade the approximation factor at all! When this works, we end up with a prophet inequality with the same approximation factor we would have enjoyed if we could have solved the original (NP-hard) problem exactly. This approach has been used to construct a 2-approximate prophet inequality for submodular combinatorial auctions, and an  $O(d)$  approximation for combinatorial auctions with set size  $d$  [Dütting et al. 2017].

## 7. MECHANISM DESIGN IMPLICATIONS

Having now built up a portfolio of prophet inequalities, we can consider their implications for mechanism design. Each of the prophet inequalities described above can be implemented as a sequential posted-price mechanism that is ex post individually rational and incentive compatible in dominant strategies. This yields, in particular, a 2-approximate truthful mechanism for submodular combinatorial auctions, an  $O(d)$ -approximate truthful mechanism for combinatorial auctions with maximum allocation size  $d$ , and others. Importantly, these mechanisms assume a Bayesian setting, and their performance guarantees hold in expectation over the realized buyer preferences. They calculate prices using agent value distributions, and therefore require some degree of market knowledge. However, this dependence affects only the welfare guarantees, rather than the incentive properties, and (as discussed in Section 3) the price computation can be made robust to errors.

Let’s compare these posted-price mechanisms to the literature on simple auctions, which are not incentive compatible but have good performance guarantees at equilibrium. For example, a natural mechanism for combinatorial auctions (and submodular combinatorial auctions) is the simultaneous item auction, where each item is sold separately but simultaneously in a sealed bid auction, and the winner of each item pays their bid for it. Bayes-Nash equilibria of this auction format have worst-case performance guarantees that are similar to the prophet inequalities established above [Christodoulou et al. 2016; Syrgkanis and Tardos 2013; Feige et al. 2015]. Moreover, this auction format can be implemented without any prior knowledge of agent preferences. That said, simultaneous item auctions offload the problem of actually finding an equilibrium (and forming beliefs about the preferences of others) to the bidders, and this is known to present computational challenges [Cai and Papadimitriou 2014]. In contrast, a posted-price mechanism pushes the computational heavy lifting to the price designer, and as we’ve seen this (approximate) price-selecting task is computationally tractable for a variety of interesting settings.

**Revenue Maximization.** To this point, I have discussed the mechanism design implications of prophet inequalities exclusively with respect to welfare maximization. As it turns out, prophet inequalities also inform the design of posted prices for approximate revenue maximization. While a comprehensive treatment of Bayesian (revenue) optimal mechanism design is beyond the scope of this survey, it would be remiss of me not to provide at least a high-level overview of this line of literature. I recommend [Hartline 2013] for an overview of the concepts from Bayesian optimal mechanism design used below.

For single-parameter settings, such as those covered in Section 4, the connection to revenue maximization makes use of Myerson’s theorem that equates expected revenue with *virtual value* [Myerson 1981]. Roughly speaking, for each buyer  $i$  we can define a virtual value function that depends on  $D_i$ , the distribution over agent  $i$ ’s value, and maps each value to a (possibly negative) virtual value. A standard result in Bayesian mechanism design equates the revenue of a mechanism with its expected virtual welfare. One can therefore approximate the revenue of the optimal mechanism by applying the prophet inequality policy to the virtual values, rather than the original values. E.g., for a single item, one could accept the first

prize whose virtual value is greater than half the expected maximum virtual value.<sup>8</sup> This yields an order-oblivious posted-price mechanism, albeit one with potentially personalized prices. We refer the interested reader to [Chawla et al. 2010] for more details on this approach. An interesting question is how well one can approximate the optimal revenue using an anonymous price, rather than personalized prices; [Alaei et al. 2015] show that an  $\epsilon$ -approximation is possible using a single posted price, under a standard regularity assumption on the value distributions.

Myerson’s characterization does not hold in general multi-dimensional settings, so this direct application of prophet inequalities to revenue maximization does not apply. However, there have been significant advances in applications of prophet inequalities to specific revenue-maximization problems. One of the first such innovations was an approximately revenue-optimal sequential posted-price mechanism for matching markets, under an additional assumption that each agent’s values for the items are independent of each other [Chawla et al. 2007; Chawla et al. 2010]. This work provides a direct argument that bounds the optimal revenue in the unit-demand case by the optimal revenue of a related single-parameter problem. Among other techniques, [Chawla et al. 2010] show how to invoke the prophet inequality, using virtual values in the related single-parameter problem, to obtain an approximation result for revenue. Similar techniques have been used to extend beyond the unit-demand case to scenarios with, e.g., matroid constraints [Kleinberg and Weinberg 2012], under the assumption of independent values across items.

More recently, a line of work in algorithmic mechanism design has extended the connection between virtual welfare and revenue for multi-dimensional problems. These connections interpret virtual values in the context of marginal revenue and dual solutions in an associated allocation program [Alaei et al. 2013; Cai et al. 2013; Cai et al. 2016]. This has led to improved upper bounds on the optimal revenue in multi-item mechanism design problems, opening the door for approximation results and the application of multi-dimensional prophet inequalities to revenue maximization. This line of inquiry is very nascent at the time of this survey, but has been used to generate constant approximations to the optimal revenue, using posted prices, in broader classes of multi-item problems. In particular, [Cai and Zhao 2017] establish the existence of sequential posted price mechanisms that  $O(1)$ -approximate the optimal revenue in submodular combinatorial auctions (and others), under an independence assumption across items. These advances appear to mark the beginning of an exciting direction for future study.

## 8. CONCLUSIONS AND FUTURE DIRECTIONS

In this survey, I described a simple economically-motivated proof of the prophet inequality. As I hope to have convinced you, this form of argument is convenient for establishing prophet inequalities for broad classes of combinatorial markets, and thereby build simple, incentive compatible posted-price mechanisms for various resource allocation problems. A natural direction to pursue is to try to apply this technique to other allocation problems. For example: scheduling traffic in a fixed

<sup>8</sup>This brief description omits a step known as ironing, which is necessary when virtual value functions are non-monotone; for a more thorough treatment of such issues we again suggest [Hartline 2013]

network, assigning cloud resources to computing jobs, combinatorial auctions with *subadditive* valuations, and others. In particular, recalling the connection with smoothness discussed in Section 5, settings for which price of anarchy bounds are known (e.g., subadditive combinatorial auctions [Feldman et al. 2013]) seem like promising candidates for new price-based prophet inequalities.

Some of the prophet inequalities described here are established using static, anonymous item prices. These are all desirable properties, for various reasons. Other results use more general pricing schemes that are dynamic, personalized, and/or price bundles rather than individual items. For what problems is this extra power necessary? For a given problem of interest, what is the gap in approximation power between, for example, static and dynamic pricing schemes?

All of the results considered here assume that the designer wishes to maximize the sum of rewards gained over the rounds. One might instead consider other ways of aggregating the prizes selected. For example, [Assaf and Samuel-Cahn 2000] established a variant of the original prophet inequality where the goal was to select up to  $k$  prizes that maximize the best of the  $k$  prizes chosen. Relatedly, [Rubinstein and Singla 2017] consider a prophet inequality over a matroid constraint where a single item can be selected (or not) each round, and the goal is to maximize a submodular (or subadditive) objective function over the set of items chosen. To what extent can this be extended to other feasibility constraints, or broader classes of aggregation functions?

For some of the problems discussed, the achievable approximation factor improves as markets grow large, in the sense of having multiple copies of items (or, more generally, needing to accept multiple prizes before any one prize becomes infeasible). Both the original prophet inequality and the matching-market prophet inequality have approximation factors tending to 1 as the number of copies of each item grows. Is there a sense in which this generalizes to broader classes of allocation problems, such as combinatorial auctions? In many cases, in the continuous large-market limit, aggregate market uncertainty is eliminated entirely and the pricing exercise reduces to the computation of market equilibrium prices. This question therefore has the most bite in large-but-finite markets, where a primary concern is the rate of convergence to efficient outcomes.

We’ve seen that in many instances, appropriate prices can be computed efficiently by sampling instances of the market, finding balanced prices, and then averaging over these prices. This assumes the ability to sample from the distribution over market instances. Can we instead learn prices efficiently through more natural feedback mechanisms, such as the revealed preference from buyer purchasing decisions? Is there a sense in which natural (or, at least, local) price-adjustment methods converge robustly to approximately balanced (or otherwise “good”) prices in the face of market uncertainty, avoiding the cycling behavior of traditional tâtonnement in the absence of market equilibria?

## REFERENCES

- ABOLHASSANI, M., EHSANI, S., ESFANDIARI, H., HAJIAGHAYI, M. T., KLEINBERG, R., AND LUCIER, B. 2017. Beating  $1-1/e$  for ordered prophets. In *Proceedings of the 49th ACM Symposium on Theory of Computing*. to appear.
- ACM SIGecom Exchanges, Vol. 16, No. 1, August 2017, Pages 26–49

- ALAEI, S. 2014. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing* 43, 2, 930–972.
- ALAEI, S., FU, H., HAGHPANAH, N., AND HARTLINE, J. 2013. The simple economics of approximately optimal auctions. In *Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. FOCS '13. 628–637.
- ALAEI, S., HAJIAGHAYI, M., AND LIAGHAT, V. 2012. Online prophet-inequality matching with applications to ad allocation. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. 18–35.
- ALAEI, S., HAJIAGHAYI, M., AND LIAGHAT, V. 2013. The online stochastic generalized assignment problem. In *Proceedings of the 16th/17th APPROX-RANDOM Workshop*. 11–25.
- ALAEI, S., HARTLINE, J. D., NIAZADEH, R., POUNTOURAKIS, E., AND YUAN, Y. 2015. Optimal auctions vs. anonymous pricing. In *Proceedings of 56th IEEE Symposium on Foundations of Computer Science*. 1446–1463.
- ASSAF, D. AND SAMUEL-CAHN, E. 2000. Simple ratio prophet inequalities for a mortal with multiple choices. *J. Appl. Probab.* 37, 4 (12), 1084–1091.
- AZAR, P. D., KLEINBERG, R., AND WEINBERG, S. M. 2014. Prophet inequalities with limited information. In *Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms*. 1358–1377.
- BABAIOFF, M., IMMORLICA, N., AND KLEINBERG, R. 2007. Matroids, secretary problems, and online mechanisms. In *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms*. 434–443.
- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2013. Understanding incentives: Mechanism design becomes algorithm design. In *Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. FOCS '13. 618–627.
- CAI, Y., DEVANUR, N. R., AND WEINBERG, S. M. 2016. A duality based unified approach to bayesian mechanism design. In *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing*. STOC '16. 926–939.
- CAI, Y. AND PAPADIMITRIOU, C. H. 2014. Simultaneous bayesian auctions and computational complexity. In *Proceedings of the 15th ACM Conference on Economics and Computation*. 895–910.
- CAI, Y. AND ZHAO, M. 2017. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th ACM Symposium on Theory of Computing*. to appear.
- CHAWLA, S., HARTLINE, J. D., AND KLEINBERG, R. D. 2007. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*. 243–251.
- CHAWLA, S., HARTLINE, J. D., MALEC, D. L., AND SIVAN, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*. 311–320.
- CHRISTODOULOU, G., KOVÁCS, A., AND SCHAPIRA, M. 2016. Bayesian combinatorial auctions. *Journal of the ACM* 63, 2, 11.
- CORREA, J., FONCEA, P., HOEKSMÁ, R., OOSTERWIJK, T., AND VREDEVELD, T. 2017. Posted price mechanisms for a random stream of customers. In *Proceedings of the 18th ACM Conference on Economics and Computation*. to appear.
- DÜTTING, P., FELDMAN, M., KESSELHEIM, T., AND LUCIER, B. 2017. Prophet inequalities made simple: Stochastic optimization by pricing non-stochastic inputs. *CoRR abs/1612.03161*.
- DÜTTING, P. AND KLEINBERG, R. 2015. Polymatroid prophet inequalities. In *Proceedings of the 23rd European Symposium on Algorithms*. 437–449.
- FEIGE, U., FELDMAN, M., IMMORLICA, N., IZSAK, R., LUCIER, B., AND SYRGKANIS, V. 2015. A unifying hierarchy of valuations with complements and substitutes. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*. 872–878.
- FELDMAN, M., FU, H., GRAVIN, N., AND LUCIER, B. 2013. Simultaneous auctions are (almost) efficient. In *Proceedings of the 45th ACM Symposium on Theory of Computing*. 201–210.
- FELDMAN, M., GRAVIN, N., AND LUCIER, B. 2015. Combinatorial auctions via posted prices. In *Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms*. 123–135.

- FELDMAN, M., SVENSSON, O., AND ZENKLUSEN, R. 2015. A simple  $o(\log \log(\text{rank}))$ -competitive algorithm for the matroid secretary problem. In *Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms*. 1189–1201.
- HAJIAGHAYI, M., KLEINBERG, R., AND SANDHOLM, T. W. 2007. Automated mechanism design and prophet inequalities. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*. 58–65.
- HARTLINE, J. D. 2013. Bayesian mechanism design. *Foundations and Trends in Theoretical Computer Science* 8, 3, 143–263.
- HILL, T. AND KERTZ, R. P. 1992. A survey of prophet inequalities in optimal stopping theory. *Contemporary Mathematics* 125, 191–207.
- HILL, T. P. AND KERTZ, R. P. 1982. Comparisons of stop rule and supremum expectations of i.i.d. random variables. *Ann. Probab.* 10, 2 (05), 336–345.
- IBARRA, O. H. AND KIM, C. E. 1975. Fast approximation algorithms for the knapsack and sum of subset problems. *Journal of the ACM* 22, 463–468.
- KELSO, ALEXANDER S, J. AND CRAWFORD, V. P. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica* 50, 6 (November), 1483–1504.
- KLEINBERG, R. AND WEINBERG, S. M. 2012. Matroid prophet inequalities. In *Proceedings of the 44th ACM Symposium on Theory of Computing*. 123–136.
- KRENGEL, U. AND SUCHESTON, L. 1977. Semiamarts and finite values. *Bulletin of the American Mathematical Society* 83, 745–747.
- KRENGEL, U. AND SUCHESTON, L. 1978. On semiamarts, amarts, and processes with finite value. *Advances in Probability and Related Topics* 4, 197–266.
- LEHMANN, B., LEHMANN, D., AND NISAN, N. 2001. Combinatorial auctions with decreasing marginal utilities. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*. 18–28.
- LI, S. 2017. Obviously strategy-proof mechanisms. *Working Paper*.
- MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1, 58–73.
- ROUGHGARDEN, T. 2012. The price of anarchy in games of incomplete information. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. 862–879.
- RUBINSTEIN, A. 2016. Beyond matroids: Secretary problem and prophet inequality with general constraints. In *Proceedings of the 48th ACM Symposium on Theory of Computing*. 324–332.
- RUBINSTEIN, A. AND SINGLA, S. 2017. Combinatorial prophet inequalities. In *Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms*. 1671–1687.
- SAMUEL-CAHN, E. 1984. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Annals of Probability* 12, 1213–1216.
- SYRGKANIS, V. AND TARDOS, É. 2013. Composable and efficient mechanisms. In *Proceedings of the 45th ACM Symposium on Theory of Computing*. 211–220.
- TREVISAN, L. 2001. Non-approximability results for optimization problems on bounded degree instances. In *Proceedings of the 33rd ACM Symposium on Theory of Computing*. 453–461.