Operations Research Letters 47 (2019) 25–29

Contents lists available at ScienceDirect

Operations Research Letters

journal homepage: www.elsevier.com/locate/orl

From pricing to prophets, and back!

José Correa [a], Patricio Foncea [b], Dana Pizarro [a], Victor Verdugo [c,*]

[a] Departamento de Ingeniería Industrial, Universidad de Chile, Chile
[b] Operations Research Center, Massachusetts Institute of Technology, United States
[c] Instituto de Ciencias de la Ingeniería, Universidad de O'Higgins, Chile

ARTICLE INFO

Article history:
Received 27 May 2018
Received in revised form 21 November 2018
Accepted 21 November 2018
Available online 23 November 2018

Keywords:
Prophet inequalities
Posted-price mechanisms
Approximation

ABSTRACT

In this work we prove that designing PPMs is equivalent to finding stopping rules for prophets. This extends the connection that any prophet type inequality can be turned into a PPM with the same approximation guarantee (Hajiaghayi et al. 2007; Chawla et al. 2010). Our reduction is robust under multiple settings including matroid feasibility constraints, or different arrival orderings. One fundamental observation implied by this result is that designing PPMs in general is equally hard from an approximation perspective to designing PPMs when the valuations are regular.

1. Introduction

In the last few years online sales have been moving from an auction format, to posted price formats [15] and the basic reason for this trend switch seems to be that posted price mechanisms (PPM) are much simpler than optimal auctions, yet efficient enough. The way these mechanisms work is as follows. Suppose a seller has an item to sell. Customers arrive one at a time and the seller proposes to each customer a take-it-or-leave-it price. The first customer accepting the offer pays the price and takes the item. These types of mechanisms are flexible and adapt well to different scenarios; their simplicity and the fact that strategic behavior vanishes make them quite suitable for many applications [9]. Of course, PPMs are suboptimal and therefore the study of their approximation guarantees – where the benchmark is that given by the optimal Myerson's auction [28] – has been an extremely active area in the last decade, in particular in the computer science community.

Hajiaghayi et al. [20] and Chawla et al. [9] establish an interesting connection between (revenue maximizing) PPMs and prophet inequalities, a problem arising in optimal stopping theory. Here a gambler is faced to a sequence of random variables and has to pick a stopping time so that the expected value he gets is as close as possible to the expectation of the maximum of all random variables, interpreted as what a prophet, who knows the realizations in advance, could get. They implicitly show that any prophet type inequality can be turned into a PPM with the same approximation guarantee. This is obtained by noting that a PPM for revenue maximization can be seen as a (threshold) stopping rule for the gambler, but on the virtual values space, and later identify these virtual thresholds with prices. As a consequence, the follow up work in the field concentrated on prophet inequities and then applied the obtained results to sequential PPMs.

In this work we fill a gap in this line of research by proving the converse of the latter result, namely, that any posted price mechanism can be turned into a prophet type inequality with the same approximation guarantee. The core of the result is a method to go back from virtual values to arbitrary distributions which may find applications beyond the scope of this paper. This result amounts not only to apply approximation guarantees from prophet inequalities to PPMs, but also to carry over the lower bounds. We observe that through our reduction we can improve the best known lower bound for sequential PPMs (in which the arrival order is either random or selected by the seller) in the single item case, the k-uniform matroid case, the general matroid case, and the general downward-closed family case.

Posted price mechanisms. The recent survey by Lucier [26] is an excellent starting point in the area, where many variants of PPMs are described. For the specific scenario where only one item must be allocated, some pricing setting studied include anonymous (the offered price is the same for all customers) [2,9,13], static (the possibly different prices to offer do not evolve as the mechanism progresses) [11,16], and Order-Oblivious (the order in which agents arrive can be chosen by an adaptive adversary) [9].

Furthermore, PPMs may be used when selling multiple items or with constraints on the subsets of served customers. Typical side constraints include matroids constraints [23,33], downward-closed systems [5,29], combinatorial prophet inequalities [8,30], combinatorial auctions [1,17], and polymatroids constraints [14]. Attention has also been payed to settings with limited information

* Corresponding author.
E-mail addresses: correa@uchile.cl (J. Correa), foncea@mit.edu (P. Foncea), dana.pizarro@ug.uchile.cl (D. Pizarro), victor.verdugo@uoh.cl (V. Verdugo).

https://doi.org/10.1016/j.orl.2018.11.010
0167-6377/

or prior-independent, where the designer must learn the distribution in order to run the mechanism [3,4,7,10,12,27].

Prophet-inequalities. For fixed $n > 1$, let $X_1, \ldots, X_n$ be nonnegative, independent random variables and T_n their set of stopping rules. A classic result of Krengel and Sucheston, and Gairing [24,25] asserts that $\mathbb{E}(\max\{X_1, \ldots, X_n\}) \leq 2\sup\{\mathbb{E}(X_t) : t \in T_n\}$, and that 2 is the best possible bound. The study of this type of inequalities, known as *prophet inequalities*, was initiated by Gilbert and Mosteller [18] and attracted a lot of attention in the eighties [21–23,31,32]. In particular Samuel-Cahn [32] noted that rather than looking at the set of all stopping rules one can (quite naturally) only look at threshold stopping rules in which the decision to stop depends on whether the value of the currently observed random variable is above a certain threshold.

Our results and techniques. The main insight we derive is a *valuation mapping lemma* stating that for any distribution F there is another distribution G whose virtual value distributes according to F. It is surprising that this basic result was missing from the auction theory literature and we believe that it may prove useful in settings beyond PPMs.

Our result is robust to different settings. It applies to random, adversarial, or best possible orders, as well as when there are multiple items and constraints on the allowed allocation sets. As already mentioned before, the sufficiency condition of the theorem is a known fact and although it has never appeared explicitly, it is implicit in previous work [9,20]. The necessary condition, however, is novel and not obvious. The main difficulty comes from taking an arbitrary distribution in the prophet inequality problem and mapping it back to a PPM. Here is where the valuation mapping lemma, that holds for arbitrary distributions, comes into play. Consider the operator that picks an arbitrary probability distribution over the nonnegative reals and returns the distribution of the ironed virtual valuation function. The Valuation Mapping Lemma states that this operator is surjective over the space of distributions. Interestingly, the lemma gives an explicit construction so we can easily interpret the thresholds as prices in the PPM.

A remarkable feature of the Valuation Mapping Lemma is that when mapping a distribution F into another distribution G whose virtual value follows F, G turns out to be *regular* (i.e., it has a monotone non-decreasing virtual value). Although in principle there may be many functions G satisfying the statement of the lemma, we can identify one explicitly with this appealing property. Together with our main theorem these imply that the posted price problem can be reduced to a prophet inequality problem, which can in turn be reduced to a posted price problem with regular distributions. Therefore, designing PPMs in general is equally hard from an approximation perspective to designing PPMs when the valuations are regular.

Another consequence of our results is that we can translate all known upper and lower bounds from PPMs into prophet inequalities and back. One example which we will further analyze in Section 4 is the case of sequential posted price mechanisms (SPM, [9]). The current best known lower bound for this setting is $\sqrt{\pi/2} \approx 1.253$ [6]. This is also the best known when the feasibility constraint is a general matroid, and even the intersection of two matroids. Our result implies an improvement on this bound to 1.341 by using the lower bound for the i.i.d. prophet inequality designed by Hill and Kertz [21]. Although our results are presented in the context of single-parameter mechanism design, they can be generalized to multi-parameter settings [9].

Organization. In Section 2 we introduce formally the online selection problem and the auction problem. In Section 3 we prove our main result – formally stated in Theorem 9 –, that is, the reduction from PPM to prophet inequalities. In Section 4 we show the improved lower bound for SPMs in more detail.

2. Preliminaries

Online selection problem. An instance of this problem corresponds to a tuple $(X, \mathcal{F}, \mathcal{T})$, where X is the ground set of n elements and each set in $\mathcal{T} \subseteq 2^X$ is called *feasible selection*. For each $x \in X$ there is a random variable w_x, called *weight*, distributed according to F_x with compact support contained in $\mathbb{R}_+$, and $\mathcal{F} = \{F_x : x \in X\}$. We assume them to be independent. The random variables are presented in an order $\sigma : [n] \to X$, and an algorithm for the problem has to decide whether to select or not an element of X when arrived. An algorithm is correct if it outputs a feasible selection.

An algorithm is an α*-approximation* if the expected weight of the output selection is at least $\alpha \cdot \mathbb{E}\left(\max_{A\in\mathcal{T}} \sum_{x\in A} w_x\right)$, that is, an α fraction of the expectation of the maximum weight over feasible selections. In the latter, the expectation is taken over $\mathcal{F}$ and the (possibly) algorithm internal randomness.

Multi-item mechanism design. Consider a single seller who provides a set of n items given by $\mathcal{I}$. For each item $i \in \mathcal{I}$, there exists a buyer having a random valuation v_i for that item. We denote by G_i the distribution of the valuation v_i, and we assume this to have a compact support contained in $\mathbb{R}_+$. We denote by $\mathcal{G} = \{G_i : i \in \mathcal{I}\}$ the set of valuation distributions. There exists a set of *feasibility constraints* for the seller, $\mathcal{T} \subseteq 2^{\mathcal{I}}$, and every set in $\mathcal{T}$ is called a *feasible allocation*. Therefore, an instance for this problem is given by a tuple $(\mathcal{I}, \mathcal{G}, \mathcal{T})$.

This setting is known to be the *single-parameter domain*. We assume the valuation distributions to be independent, and they are known by the seller. Buyers arrive in an arbitrary order $\sigma : [n] \to \mathcal{I}$.

Distributions. Throughout this work we only consider distributions with bounded support. In general, a distribution F is not invertible but we work with its *generalized inverse*, given by $F^{-1}(y) = \inf\{t \in \mathbb{R} : F(t) \geq y\}$. In particular, the derivative F' exists almost everywhere and it is called the *density* function. In what follows we consider distributions with strictly positive density.

Myerson's optimal mechanism. In his seminal work, Myerson [28] characterizes the mechanism maximizing the revenue for single-parameter domains. In order to analyze the optimization problem, he introduces a quantity called *virtual valuation*, that allows to solve the problem in an equivalent and simpler maximization setting.

Definition 1. For a random variable v with distribution G and density g, the *virtual valuation of* v is the function $\phi_G(t) = t - (1 - G(t))/g(t)$. We say that G is *regular* if ϕ_G is monotone non-decreasing.

In the regular case, the optimal mechanism computes the virtual valuation for each buyer and then it allocates to a subset of them maximizing its total virtual value. Recall that a mechanism is called *incentive-compatible* if each player has a weakly dominant strategy of truthful reporting.

Theorem 2 ([28])**.** *If the distributions in $\mathcal{G}$ are regular, the expected revenue of any incentive-compatible single-parameter mechanism $\mathcal{M}$ is equal to its expected virtual surplus, $\mathbb{E}\left(\sum_{x\in M} \phi_x^+(v_x)\right)$, where M is the allocation provided by the mechanism, and $\phi_x^+ = \max\{0, \phi_x\}$.*

In particular, when the distributions are regular, Myerson's optimal mechanism is incentive-compatible and so it satisfies the above conditions in the theorem. We introduce a technical lemma about virtual valuations that is used along the reductions. When the distributions are not regular, Myerson considered an *ironed* virtual valuation for its analysis [28], denoted by $\bar{\phi}_G$ when the valuation distribution is G. More specifically, take $Q(\theta) = \theta G^{-1}(1-\theta)$ and let R be the concave hull of Q, namely, $R(\theta)$ is given by

$$\min_{x,\theta_1,\theta_2\in[0,1]} \{xQ(\theta_1) + (1-x)Q(\theta_2) : x\theta_1 + (1-x)\theta_2 = \theta\}.$$

The ironed virtual valuation is $\bar{\phi}_G(t) = R'(1 - G(t))$. In particular, when the valuation is regular, the ironed virtual valuation corresponds to the virtual valuation, $\bar{\phi}_G = \phi_G$. If the context is clear, we omit the subscript on the notation for the (ironed) virtual valuation.

Lemma 3. *Let v be a random variable with distribution G. Let $\tau \geq 0$ and $q = \mathbb{P}(\bar{\phi}(v) \geq \tau)$. Then, there exist $q_1, q_2, x \in [0, 1]$ such that $xq_1 + (1-x)q_2 = q$ and*

$$\mathbb{E}(\bar{\phi}(v) \mid \bar{\phi}(v) \geq \tau) = \frac{1}{q}\Big(xq_1G^{-1}(1-q_1) + (1-x)q_2G^{-1}(1-q_2)\Big).$$

In particular, when the distribution G is regular we have $\mathbb{E}(\phi(v) \mid \phi(v) \geq \tau) = G^{-1}(1-q)$.

Proof of Lemma 3. By expanding the conditional expectation and changing the integration domain,

$$\mathbb{E}(\bar{\phi}(v) \mid \bar{\phi}(v) \geq \tau)\mathbb{P}(\bar{\phi}(v) \geq \tau) = \int_{\bar{\phi}^{-1}(\tau)}^{\infty} \bar{\phi}(u)dG(u) = \int_0^q R'(\theta)d\theta = R(q),$$

where $\bar{\phi}^{-1}$ corresponds to the generalized inverse of $\bar{\phi}$. Since R is the concave-hull of $Q(\theta) = \theta G^{-1}(1-\theta)$, there exists $x, q_1, q_2 \in [0, 1]$ such that $xq_1 + (1-x)q_2 = q$ and $R(q) = xQ(q_1) + (1-x)Q(q_2) = q_1G^{-1}(1-q_1) + (1-x)q_2G^{-1}(1-q_2)$. When G is regular, Q is a concave function and therefore $R(q) = q \cdot G^{-1}(1-q)$. □

Posted-price mechanisms. In a *posted-price mechanism* (PPM), once a buyer arrives the seller offers a price in a take-it-or-leave-it fashion. The posted price mechanism, $\mathcal{M}$, upon arrival of a buyer preferring item $i \in \mathcal{I}$, computes a price p_i. This price is a function of the *history* $\mathcal{H}_t = (\sigma_t, A_{t-1}, \mathcal{V}_{t-1})$ at time t, where σ_t is the order in which the buyers arrived up to time t, A_{t-1} denotes the current allocation and $\mathcal{V}_{t-1} = \{v_{\sigma(j)} : j \in \{1, \ldots, t-1\}\}$ is the set of valuation realizations for the buyers so far arrived. The *expected revenue* of the mechanism is just $\mathbb{E}(\sum_{i\in\mathcal{I}} p_i(1 - G_i(p_i)))$.

From prophets to pricing. The idea of constructing PPM from existing prophet inequalities has been exploited extensively the last decade starting with the work of Hajiaghayi et al. [20] and that of Chawla et al. [9]. An algorithm for an online selection problem is *based on thresholds* if every time that an element arrives, it is included to the current solution if its weight is above a certain threshold.

Theorem 4 ([9,20]). *Suppose there exists an online selection algorithm based on thresholds that is an α-approximation for $(\mathcal{I}, \mathcal{G}^{\phi}, \mathcal{T})$ presented in order σ. Then, there exists a posted-price mechanism that is an α-approximation for $(\mathcal{I}, \mathcal{G}, \mathcal{T})$ presented in order σ.*

3. From pricing to prophets

Reduction overview. Consider an instance $(X, \mathcal{F}, \mathcal{T})$ for the optimal stopping problem, and suppose we have access to a single-parameter PPM $\mathcal{M}$ that provides a guarantee over the ground set X and feasibility constraints $\mathcal{T}$. If we were able to find valuation distributions $\mathcal{G} = \{G_x : x \in X\}$ such that $\phi^+_{G_x}(v_x)$ has distribution F_x, where v_x has distribution G_x, then we could feed the mechanism $\mathcal{M}$ with the instance $(X, \mathcal{G}, \mathcal{T})$ of a multi-item auction problem, using the same order σ in which the elements of the ground set X are output in the optimal stopping problem. In particular, since the weight of x is distributed according to $\phi^+_{G_x}(v_x)$, and they are all independent, by Theorem 2, the revenue of the mechanism on the instance $(X, \mathcal{G}, \mathcal{T})$ equals the sum of the weights of the elements selected, and so our online stopping algorithm provides a prophet inequality that preserves the approximation given by mechanism $\mathcal{M}$ in the multi-item auction instance.

3.1. Valuation mapping lemma

In this section we introduce the key lemma that allows us to map from a weight distribution F to a valuation distribution G with virtual valuation distributed according to F. W.l.o.g. we restrict ourselves to the case where F has support $[0, 1]$. The result extends to compact support in $\mathbb{R}_+$ via an affine transformation. Formally, we prove the following lemma.

Lemma 5 (*Valuation Mapping Lemma*). *Let w be a random variable with distribution F and support in $[0, 1]$. Then, there exists a distribution G with non-negative support such that if v is distributed according to G, then $\phi^+_G(v)$ is distributed according to F.*

The proof is constructive and we provide an explicit expression for the distribution G: we define G to be the generalized inverse of H, where

$$H(q) = \frac{1}{1-q}\int_q^1 F^{-1}(y)dy \tag{1}$$

if $q \in [0, 1)$, $H \equiv 0$ in $(-\infty, 0)$ and $H \equiv 1$ in $[1, +\infty)$. Observe that $H(0) = \mathbb{E}(w)$, where w follows distribution F, and therefore H might be discontinuous in 0.

We now present two intermediate results that will be of use in the proof of Lemma 5.

Proposition 6. *H is continuous in $(0, 1]$. Furthermore, there exists $T \in [0, 1]$ such that H is strictly increasing in the interval $[0, T)$, and is constant equal to 1 in the interval $[T, 1]$. In particular, H is a distribution.*

In fact, we show that $T = 1$ if F is continuous by the left in $t = 1$. Otherwise, if F is discontinuous in $t = 1$ then $T < 1$. This behavior of H in the interval $(0, 1]$ translates to G, in the sense that H is invertible in the whole $(0, 1]$ except when $T < 1$ and so G has a discontinuity at $t = 1$. In other words, G is also strictly increasing and continuous in $(H(0), 1)$.

Proof of Proposition 6. The continuity of H in $(0, 1]$ comes from the fundamental theorem of calculus. To study the monotonicity of H let us compute the first derivative of H and analyze its sign. Observe that

$$H'(q) = \frac{1}{(1-q)^2}\left(-F^{-1}(q)(1-q) + \int_q^1 F^{-1}(y)dy\right) = \frac{1}{(1-q)^2}\int_q^1 (F^{-1}(y) - F^{-1}(q))dy,$$

and $F^{-1}(y) - F^{-1}(q) \geq 0$, since F^{-1} is non-decreasing, and y is at least q. Therefore, $H'(q) = 0$ if and only if F^{-1} equals $F^{-1}(q)$ almost everywhere in $[q, 1]$, which in turn happens if and only if $\lim_{s\to F^{-1}(q)^-} F(s) = q$ and $F(F^{-1}(q)) = 1$. Taking $T = q$ the proof follows. □

Proposition 7. *Let G be defined as in (1). Then:*

1. *G is a distribution with support $[\mathbb{E}(w), 1]$, where w is a random variable with distribution F.*
2. *For all t in the support of G, $\phi_G(t) = F^{-1}(G(t))$.*

3. *The virtual valuation ϕ_G is non-decreasing. In particular, ϕ_G is non-negative and therefore $\phi_G^+ = \phi_G$.*

Proof of Proposition 7.

1. By Proposition 6, G is a distribution. The result for the support comes from $H(0) = \int_0^1 F^{-1}(y)dy = \mathbb{E}(w)$ and $H(1) = 1$, and G being strictly increasing in the interior of the interval.
2. Let t be in $[\mathbb{E}(w), 1)$. By Proposition 6, G is strictly increasing and continuous on this interval, and therefore invertible. It is then sufficient to show that $\phi_G(H(q)) = F^{-1}(q)$ where $H(q) = t$. In particular, $q \in [0, T)$, with T as in the statement of Proposition 6. Since G is also differentiable and $G^{-1} = H$ in this interval, it follows that $\phi_G(H(q)) = H(q) - (1 - q)/G'(H(q)) = H(q) - (1 - q)H'(q)$. On the other hand, from the definition of H,
$$H'(q) = \frac{1}{(1-q)^2}\left(-F^{-1}(q)(1-q) + \int_q^1 F^{-1}(y)dy\right) = \frac{1}{1-q}\left(-F^{-1}(q) + H(q)\right),$$
and therefore $H(q) - (1 - q)H'(q) = F^{-1}(q)$.
3. It follows from (2) that ϕ_G is non-decreasing, as both F^{-1} and G are non-decreasing. Since $\phi_G(\mathbb{E}(w)) = F^{-1}(0) = 0$, it holds that ϕ_G is non-negative which in turn implies that $\phi_G^+ = \phi_G$. □

Proof of Lemma 5. Recall that v is a random variable distributed according to F. Let U be a random variable uniformly distributed in $[0, 1)$. Then, by Proposition 7(1), G is a distribution and therefore the generalized inverse of U, namely $v = G^{-1}(U)$, has distribution G. By Proposition 7(3), $\phi_G^+ = \phi_G$ and therefore it remains to study the distribution of ϕ_G. By Proposition 7(2), we observe that for $t < 1$, $\mathbb{P}(\phi_G(v) \le t) = \mathbb{P}(F^{-1}(G(v)) \le t) = \mathbb{P}(G(v) \le F(t)) = F(t)$, since by Proposition 6, G is invertible in $[\mathbb{E}(w), 1)$. Thus, $G(v)$ is uniformly distributed in $[0, 1)$. □

To prove the reduction, stated formally in Theorem 9, we need other small technical ingredient. Given non-decreasing functions $\eta : [0, 1] \to [0, 1]$ and $\nu : [a, 1] \to [0, 1]$ for $0 \le a \le 1$, we say that ν is a *non-linear stretching* of η if there exists $\xi : [a, 1] \to [0, 1]$ strictly increasing and continuous in $[a, 1)$ such that $\nu = \eta \circ \xi$ in $[a, 1)$.

Proposition 8. *If ν is constant in the interval $[c, r)$, then η is constant in the interval $[\xi(c), \xi(r))$.*

Observe that since G is strictly increasing and continuous in $[\mathbb{E}(w), 1)$, it follows that ϕ_G is a non-linear stretching of F^{-1}. In other words, if ϕ_G is constant in an interval $[c, r)$, then F^{-1} is constant over $[G(c), G(r))$.

Proof of Proposition 8. Suppose η is not constant over the interval $[\xi(c), \xi(r))$, i.e., there exists $s \in (\xi(c), \xi(r))$ such that $\eta(s) > \eta(\xi(r))$. Since ξ is strictly increasing and continuous in (c, r), there exists $z \in (c, r)$ such that $\xi(z) = s$. Thus, $\nu(z) = \eta(\xi(z)) = \eta(s) > \eta(\xi(r)) = \nu(r)$, which contradicts the fact that ν is constant in $[c, r)$. □

From posted prices to online selection. Given an instance for the online selection problem, we feed a single-parameter mechanism, $\mathcal{M}$, by constructing a set of valuations $\mathcal{G}$ using the Valuation Mapping Lemma. For ease of notation, let ϕ_x be the virtual valuation of G_x. We perform a randomized tie-breaking to determine the thresholds. More specifically, consider the *boundary prices* given by $p_x^- = \inf\{p \in \mathbb{R} : \phi_x(p) = \phi_x(p_x)\}$ and $p_x^+ = \sup\{p \in \mathbb{R} : \phi_x(p) = \phi_x(p_x)\}$.

Algorithm 1 From posted prices to thresholds.

Require: $(X, \mathcal{F}, \mathcal{T})$ of the online selection problem.
1: Initialize $A_0 \leftarrow \emptyset$.
2: **for** $t = 1$ to n **do**
3: Let $x = \sigma(t)$, and set price $p_x = \mathcal{M}(\mathcal{H}_{t-1}, \mathcal{G}, x)$,
4: **if** $p_x^- = p_x^+$ and $w_x \ge \phi_x(p_x)$ **then**
5: select x, $A_t \leftarrow A_{t-1} \cup \{x\}$;
6: **else if** $p_x^- < p_x^+$ **then**
7: **if** $w_x > \phi_x(p_x)$ **then**
8: select x, $A_t \leftarrow A_{t-1} \cup \{x\}$;
9: **else if** $w_x = \phi_x(p_x)$ **then**
10: set $\theta_x = \frac{G(p_x^+) - G(p_x)}{G(p_x^+) - G(p_x^-)}$
11: select x w.p. θ_x, $A_t \leftarrow A_{t-1} \cup \{x\}$,
12: reject x w.p. $1 - \theta_x$, $A_t \leftarrow A_{t-1}$.
13: **else** reject x, $A_t \leftarrow A_{t-1}$.
14: Return Alg $= A_n$.

Theorem 9. *Let $(X, \mathcal{F}, \mathcal{T})$ be an instance of the online selection problem, and $(X, \mathcal{G}, \mathcal{T})$ the instance of the multi-item auction obtained by the Valuation Mapping Lemma. If the mechanism $\mathcal{M}$ is an α-approximation for $(X, \mathcal{G}, \mathcal{T})$ presented in order σ, then Algorithm 1 is an α-approximation for $(X, \mathcal{F}, \mathcal{T})$ presented in order σ.*

Proof of Theorem 9. Let Q_t be the event that item $\sigma(t)$ is selected by Algorithm 1. We denote by $\mathbb{P}_{t-1}$ the probability distribution conditional on the history $\mathcal{H}_{t-1}$, and the notation extends to the expectation. We denote by $\chi(Q_t)$ the indicator function of event Q_t. By conditioning on the history, we have that $\mathbb{E}\left(\sum_{x\in \text{Alg}} w_x\right) = \sum_{t=1}^n \mathbb{E}\left(w_{\sigma(t)}\chi(Q_t)\right) = \sum_{t=1}^n \mathbb{E}\left(\mathbb{E}_{t-1}\left(w_{\sigma(t)}\chi(Q_t)\right)\right)$. For $t \in \{1, \dots, n\}$, let $x = \sigma(t)$ and $p_x = \mathcal{M}(\mathcal{H}_{t-1}, \mathcal{G}, x)$ be the price computed by $\mathcal{M}$. We claim that Algorithm 1 satisfies $\mathbb{E}_{t-1}(w_x\chi(Q_t)) = p_x(1 - G(p_x))$, where G_x is the distribution of v_x. Before proving this, we see how to conclude the theorem using the equality above. Since $\mathcal{M}$ is an α-approximation and using Theorem 2, we have that $\sum_{x\in X} p_x(1 - G_x(p_x))$ is at least
$$\alpha \cdot \mathbb{E}\left(\max_{A\in\mathcal{T}} \sum_{x\in A} \phi_x^+(v_x)\right) = \alpha \cdot \mathbb{E}\left(\max_{A\in\mathcal{T}} \sum_{x\in A} w_x\right),$$
where in the last equality we used the fact that the valuations are obtained from the Valuation Mapping Lemma, and that the distributions in $\mathcal{F}$ are independent. This proves that Algorithm 1 is an α-approximation.

It remains to prove $\mathbb{E}_{t-1}(w_x\chi(Q_t)) = p_x(1-G_x(p_x))$. To this end, we condition on whether we are in line 5 or 6 of Algorithm 1. If the condition in line 5 holds, then $\phi_x(v_x) > \phi_x(p_x)$ if and only if $v_x > p_x$. In particular, $\mathbb{P}_{t-1}(Q_t) = \mathbb{P}_{t-1}(\phi_x(v_x) > \phi_x(p_x)) = \mathbb{P}_{t-1}(v_x > p_x) = 1 - G_x(p_x)$. By Lemma 3 and Proposition 7, $\mathbb{E}_{t-1}\left(w_{\sigma(t)}|Q_t\right) = \mathbb{E}_{t-1}(\phi_x(v_x)|\phi_x(v_x) > \phi_x(p_x)) = p_x$, so setting $R(t) = p(1 - G_x(t))$, we conclude that $\mathbb{E}_{t-1}\left(w_{\sigma(t)}\chi(Q_t)\right) = \mathbb{E}_{t-1}\left(w_{\sigma(t)}|Q_t\right)\mathbb{P}_{t-1}(Q_t) = p_x(1 - G_x(p_x)) = R(p_x)$. Suppose now that the condition in line 6 is satisfied. By Proposition 8, the function F^{-1} is constant in the interval $[G(p_x^-), G(p_x^+))$.

Claim 10. *For every $p \in [p_x^-, p_x^+)$, we have that*
$$G(p) = \frac{p - \phi_x(p_x^+)G(p_x^+) - \int_{G(p_x^+)}^1 F^{-1}(y)dy}{p - \phi_x(p_x^+)}.$$

We postpone the proof of the claim to the end of this section. Using the expression of G shown in the claim, it follows that $R(p) = \phi_G(p_x^+)[G(p_x^+) - G(p)] + \int_{G(p_x^+)}^1 F^{-1}(y)dy$. By definition of θ_x, $G(p) = \theta_x G(p_x^-) + (1 - \theta_x)G(p_x^+)$ and therefore $R(p) = \theta_x R(p_x^-)$

$+(1-\theta_x)R(p_x^+)$. Note that $\phi_x(v_x) \geq \phi_x(p_x)$ if and only if $v_x \geq p_x^-$, and $\phi_x(v_x) > \phi_x(p_x)$ if and only if $v_x \geq p_x^+$. By conditioning on whether line 8 or line 11 is satisfied, we have that $\mathbb{E}_{t-1}\left(w_{\sigma(t)}\chi(Q_t)\right)$ equals

$$\theta_x\mathbb{E}_{t-1}(\phi_x(v_x)|v_x \geq p_x^-)\mathbb{P}_{t-1}(v_x \geq p_x^-) + (1-\theta_x)\mathbb{E}_{t-1}(\phi_x(v_x)|v_x > p_x^+)\mathbb{P}_{t-1}(v_x > p_x^+).$$

Using Lemma 3 we conclude that $\mathbb{E}_{t-1}\left(w_{\sigma(t)}\chi(Q_t)\right)$ equals $\theta_x R(p_x^-) + (1-\theta_x)R(p_x^+) = R(p_x)$. □

Proof of Claim 10. Given $q \in [G(p_x^-), G(p_x^+)]$, we have

$$H(q) = \frac{1}{1-q}\left[\int_q^{G(p_x^+)} F^{-1}(y)dy + \int_{G(p_x^+)}^1 F^{-1}(y)dy\right] = \frac{G(p_x^+)-q}{1-q}\phi_x(p_x) + \frac{1}{1-q}\int_{G(p_x^+)}^1 F^{-1}(y)dy,$$

where the last equality follows from F^{-1} being equal to $\phi_x(p_x)$ on the interval $[G(p_x^-), G(p_x^+))$. □

4. Implications

A direct consequence of Theorem 9 and the Valuation Mapping Lemma is that we obtain lower bounds for the guarantees of PPMs by considering lower bound instances of the online selection problem. We improve the previous known lower bounds for SPM when constraints are on the form of downward closed families, from $\log n/(3\log\log n)$ to $\log n/(2\log\log n)$ [9,29], and in the k-uniform matroid setting from 1.253 to 1.341 [6,21] (this holds for general matroids or even intersection of matroids). Additionally, using the results from Göbel et al. [19], it is possible to derive a new lower bound for PPM where the feasibility is given by stable sets in graphs, of $\Omega(\log n/\log^2\log n)$.

Consider the single item sequential posted price problem as defined by Chawla et al. [9], where one seller can also choose the order in which the buyers are offered the price. The current best known lower bound for this problem was obtained a decade ago by Blumrosen and Holenstein [6] considering an instance where all buyers have i.i.d. valuations distributed according to $F(v) = 1 - 1/v^2$. They show that the expected revenue of the optimal mechanism is $\Gamma(1/2)\sqrt{n}/2$, while that of the optimal SPM is $\sqrt{n/2}$. Thus, the ratio is approximately 1.253. Rather surprisingly, this lower bound is the best known when the feasibility constraint is a k-uniform matroid, a general matroid, and even the intersection of two matroids.

To see that the lower bound can be improved we consider the lower bound for the i.i.d. prophet inequality designed by Hill and Kertz [21] over three decades ago. They considered the problem of finding the best constant a_n such that for n i.i.d. random variables the expected gambler's gains are within a factor a_n of that of the gambler. They were able to characterize a_n through a recursion and also to find the instance that exactly achieves this gap of a_n. In follow-up work, Kertz [22] proves that a_n converges to $\beta \approx 1.341$, the unique solution to the integral equation

$$\int_0^1 1/(y(1-\ln(y)) + (\beta-1))dy = 1. \quad (2)$$

With the Valuation Mapping Lemma we map back the distributions of the instances of Hill and Kertz to distributions for the PPM. Since the distributions used in Hill and Kertz's instances are i.i.d., those for the sequential posted price problem are also i.i.d. and Theorem 9 guarantees the gap of $\beta \approx 1.341$ to be preserved.

Acknowledgment

Dana Pizarro was supported by CONICYT/Doctorado Nacional/2016-21161440, Chile.

References

[1] S. Alaei, Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers, SIAM J. Comput. 43 (2) (2014) 930–972.
[2] S. Alaei, J. Hartline, R. Niazadeh, E. Pountourakis, Y. Yuan, Optimal auctions vs. Anonymous pricing, in: Proceedings of the IEEE 56th Annual Symposium on Foundations of Computer Science, in: FOCS'15, 2015.
[3] P.D. Azar, R. Kleinberg, S.M. Weinberg, Prophet inequalities with limited information, in: Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms, in: SODA'14, 2014.
[4] M. Babaioff, L. Blumrosen, S. Dughmi, Y. Singer, Posting prices with unknown distributions, ACM Trans. Econ. Comput. 5 (2) (2017) 13:1–13:20.
[5] M. Babaioff, N. Immorlica, R. Kleinberg, Matroids, secretary problems, and online mechanisms, in: Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms, in: SODA'07, 2007.
[6] L. Blumrosen, T. Holenstein, Posted prices vs. negotiations: An asymptotic analysis, in: Proceedings of the 9th ACM Conference on Electronic Commerce, in: EC'08, 2008.
[7] Y. Cai, C. Daskalakis, Learning multi-item auctions with (or without) samples, in: Proceedings of the IEEE 58th Annual Symposium on Foundations of Computer Science, in: FOCS'17, 2017.
[8] Y. Cai, M. Zhao, Simple mechanisms for subadditive buyers via duality, in: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, in: STOC 2017, 2017.
[9] S. Chawla, J.D. Hartline, D.L. Malec, B. Sivan, Multi-parameter mechanism design and sequential posted pricing, in: Proceedings of the 42th ACM Symposium on Theory of Computing, in: STOC'10, 2010.
[10] R. Cole, T. Roughgarden, The sample complexity of revenue maximization, in: Proceedings of the 46th ACM Symposium on Theory of Computing, in: STOC'14, 2014.
[11] J. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, T. Vredeveld, Posted price mechanisms for a random stream of customers, in: Proceedings of the ACM Conference on Economics and Computation, in: EC'17, 2017.
[12] P. Dhangwatnotai, T. Roughgarden, Q. Yan, Revenue maximization with a single sample, in: Proceedings of the 11th ACM Conference on Electronic Commerce, in: EC'10, 2010.
[13] P. Dütting, F. Fischer, M. Klimm, Revenue Gaps for Discriminatory and Anonymous Sequential Posted Pricing, 2016.
[14] P. Dütting, R. Kleinberg, Polymatroid prophet inequalities, in: Algorithms-ESA 2015, Springer, 2015.
[15] L. Einav, C. Farronato, J. Levin, N. Sundaresan, Auctions versus posted prices in online markets, J. Polit. Econ. 126 (1) (2018) 178–215.
[16] H. Esfandiari, M. Hajiaghayi, V. Liaghat, M. Monemizadeh, Prophet secretary, in: Algorithms-ESA 2015, Springer, 2015, pp. 496–508.
[17] M. Feldman, N. Gravin, B. Lucier, Combinatorial auctions via posted prices, in: Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms, in: SODA'15, 2015.
[18] J.P. Gilbert, F. Mosteller, Recognizing the maximum of a sequence, J. Amer. Statist. Assoc. 61 (1966) 35–76.
[19] O. Göbel, M. Hoefer, T. Kesselheim, T. Schleiden, B. Vöcking, Online independent set beyond the worst-case: secretaries, prophets, and periods, in: Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, 2014.
[20] M. Hajiaghayi, R. Kleinberg, T. Sandholm, Automated Online Mechanism Design and Prophet Inequalities, Vol. 7, AAAI, 2007, pp. 58–65.
[21] T.P. Hill, R.P. Kertz, Comparisons of stop rule and supremum expectations of i.i.d. random variables, Ann. Probab. 10 (2) (1982) 336–345.
[22] R.P. Kertz, Stop rule and supremum expectations of i.i.d. random variables: A complete comparison by conjugate duality, J. Multivariate Anal. 19 (1986) 88–112.
[23] R. Kleinberg, S.M. Weinberg, Matroid prophet inequalities, in: Proceedings of the 44th ACM Symposium on Theory of Computing, in: STOC'12, 2012.
[24] U. Krengel, L. Sucheston, Semiamarts and finite values, Bull. Amer. Math. Soc. 83 (1977) 745–747.
[25] U. Krengel, L. Sucheston, On semiamarts, amarts, and processes with finite value, Adv. Probab. 4 (1978) 197–266.
[26] B. Lucier, An economic view of prophet inequalities, ACM SIGECOM Exch. 16 (1) (2017) 26–49.
[27] J. Morgenstern, T. Roughgarden, Learning simple auctions, in: 29th Annual Conference on Learning Theory, 2016, pp. 1298–1318.
[28] R.B. Myerson, Optimal auction design, Math. Oper. Res. 6 (1) (1981) 58–73.
[29] A. Rubinstein, Beyond matroids: Secretary problem and prophet inequality with general constraints, in: Proceedings of the 48th ACM Symposium on Theory of Computing, in: STOC'16, 2016.
[30] A. Rubinstein, S. Singla, Combinatorial prophet inequalities, in: Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms, in: SODA'17, 2017.
[31] U. Saint-Mont, A simple derivation of a complicated prophet region, J. Multivariate Anal. 80 (2002) 67–72.
[32] E. Samuel-Cahn, Comparisons of threshold stop rule and maximum for independent nonnegative random variables, Ann. Probab. 12 (4) (1983) 1213–1216.
[33] Q. Yan, Mechanism design via correlation gap, in: Proceedings of the 22th ACM-SIAM Symposium on Discrete Algorithms, in: SODA'11, 2011.