

Convex Prophet Inequalities

Junjie Qin^{*}, Ram Rajagopal[†], Shai Vardi[‡], Adam Wierman[§]

ABSTRACT

We introduce a new class of prophet inequalities—convex prophet inequalities—where a gambler observes a sequence of convex cost functions $c_i(x_i)$ and is required to assign some fraction $0 \leq x_i \leq 1$ to each, such that the sum of assigned values is exactly 1. The goal of the gambler is to minimize the sum of the costs. We provide an optimal algorithm for this problem, a dynamic program, and show that it can be implemented in polynomial time when the cost functions are polynomial. We also precisely characterize the competitive ratio of the optimal algorithm in the case where the gambler has an outside option and there are polynomial costs, showing that it grows as $\Theta(n^{p-1}/\ell)$, where n is the number of stages, p is the degree of the polynomial costs and the coefficients of the cost functions are bounded by $[\ell, u]$.

Keywords

Prophet inequality, online algorithms, resource allocation, stochastic control

1. INTRODUCTION

Consider an online decision maker tasked with procuring $C > 0$ units of a divisible commodity at minimal cost from n suppliers that arrive online. Supplier i arrives with a strictly positive, real-valued, independent convex cost function c_i that is drawn from a known distribution that may be chosen adversarially. The decision maker must decide on a contract with supplier i before supplier $i + 1$ arrives. As a motivating example, consider electricity markets. When procuring generation capacity in order to meet demand, load serving entities (LSEs) make contracts with generators. Most of these contracts are made months or even years in advance, when future availability of other options for generation is not known to the LSE [1, 2, 3, 4]. Thus, LSEs face an online decision problem where, given a forecast for the ca-

^{*}University of California, Berkeley, CA, USA. E-mail: qinj@berkeley.edu.

[†]Stanford University, Stanford, CA, USA. E-mail: ramr@stanford.edu.

[‡]California Institute of Technology, Pasadena, CA, USA. E-mail: svardi@caltech.edu.

[§]California Institute of Technology, Pasadena, CA, USA. E-mail: adamw@caltech.edu.

capacity needed, they must make contracts with generators that arrive online over a span of years. Contracts with generators are typically strictly convex. While the precise form of the cost functions may be complicated, it is often modeled as quadratic in analytic work [5, 6, 7], which motivates us to focus on quadratic (or more generally, *polynomial*) cost functions. Beyond electricity markets, online packing problems with convex cost functions are also highly relevant for other procurement problems, e.g., optimal control [8, 9], cloud computing [10, 11], and inventory management [12, 13].

If the cost functions are linear, the problem is a generalization of the minimization version of the classical *prophet inequalities* setting. In the classical (maximization) setting, a decision maker must choose one of n items that arrive online and the reward of each item is drawn from a known, non-negative, real-valued distribution that may be chosen adversarially, with the goal of maximizing the reward. The prophet inequality bounds the reward obtainable by the decision maker as a function of the reward that can be obtained by a *prophet* who can foresee the entire sequence and stop at the maximal value. In the minimization version, the rewards are replaced by costs, and the goal is to minimize the cost. This problem and ours are analogous in the case of linear costs: the optimal strategies of both the decision maker and prophet in the setting above is to set $x_i = 1$ for some i and $x_j = 0$ for all $j \neq i$, which correspond to the decision maker and prophet choosing a single item. Classical prophet inequalities are *stopping problems*, as the decision maker sees the items in sequence, and chooses when to stop and accept the current item.

When the cost functions are strictly convex and positive, this is not a stopping problem since the prophet always procures a strictly positive amount of the commodity at each stage, and thus never “stops”. Instead, this problem is better thought of as an online packing problem. More specifically, this problem differs from the setting of the classical prophet inequality in three ways: (i) the cost functions are convex, (ii) the decision maker can make real-valued (non-integral) decisions, and (iii) the decision maker seeks to minimize cost instead of maximize reward.

2. MODEL AND PRELIMINARIES

We consider an online decision maker (gambler) tasked with procuring a single unit of a divisible commodity at minimal cost from n suppliers that are arriving online. Supplier i arrives with a non-negative, real-valued convex cost function c_i , which is drawn from a known distribution D_i .

The distributions D_1, \dots, D_n are independent and possibly chosen adversarially. The gambler must decide on a contract with supplier i before supplier $i + 1$ arrives.

We compare the cost of the gambler to the cost of a *prophet*, who can foresee the sequence of realized cost functions. Since the prophet knows the cost functions, it can simply solve the following convex program to identify the optimal allocation

$$\begin{aligned} \text{OPT} &= \min_x \sum_{i=1}^n c_i(x_i) \\ \text{s.t.} \quad &\sum_{i=1}^n x_i = 1; \quad 0 \leq x_i \leq 1; \quad i = 1, \dots, n. \end{aligned}$$

Since the cost functions are random, we are interested in the expected performance, $\mathbb{E}[\text{OPT}]$.

Without the ability to foresee the future, the *gambler* has to base decisions on the information that is available at stage i , denoted by \mathcal{H}_i . This information includes all the distributions, the realized cost functions up to c_i , and the amount of commodity that has already been obtained (denoted by s_i). That is, the decision of the gambler has to be *causal*, i.e., $x_i = \pi_i(\mathcal{H}_i)$, $i = 1, \dots, n$, where π_i is the gambler's policy for determining the amount to procure from supplier i given the information at stage i .

Let $\pi = (\pi_1, \dots, \pi_n)$ be the sequence of policies the gambler may use and let Π be the set of admissible policies that includes all the policies generating x_i 's such that $x_i \geq 0$ and $\sum_i x_i = 1$. Then, the expected cost of the gambler using an admissible policy $\pi \in \Pi$ is $\mathbb{E}[\text{ALG}^\pi] = \mathbb{E}[\sum_{i=1}^n c_i(\pi_i(\mathcal{H}_i))]$.

For convenience, we write $\mathbb{E}[\text{ALG}]$ when the policy used is clear from context. The *gambler's problem* is then to design an admissible policy π that minimizes his expected cost:

$$\inf_{\pi \in \Pi} \mathbb{E}[\text{ALG}^\pi]. \quad (1)$$

For deriving convex prophet inequalities, we are interested in obtaining the *exact competitive ratio*

$$\sup_D \inf_{\pi \in \Pi} \frac{\mathbb{E}[\text{ALG}^\pi]}{\mathbb{E}[\text{OPT}]}, \quad (2)$$

where the maximization is over the space of all sequences of independent distributions $D = (D_1, \dots, D_n)$.

Our goals are (i) to obtain explicit expressions of (2) for a wide class of cost functions which are useful for bounding the algorithm performance for suboptimal policy or for arbitrary distributions, and (ii) to identify simple and efficient algorithms for solving the online allocation problem with provably good competitive ratios.

2.1 An optimal algorithm

We formulate the optimal policy for the gambler using dynamic programming. This policy is the admissible policy that minimizes the expected cost, i.e., the solution of (1). In particular, (1) is a stochastic control problem, and so the optimal policy can be characterized via dynamic programming (DP) using the backward recursion described below. First, define the *cost-to-go functions* J_i on $s_i \in [0, 1]$ as follows:

$$J_n(s_n) = c_n(1 - s_n),$$

$$J_i(s_i) = \min_{0 \leq x_i \leq 1 - s_i} c_i(x_i) + \mathbb{E}[J_{i+1}(s_i + x_i)], \quad i = 1, \dots, n - 1.$$

Equipped with these cost-to-go functions, the optimal causal policy is a mapping from the state $s_i \in [0, 1]$ to the action

$x_i \in [0, 1]$ that takes the form

$$x_n(s_n) = 1 - s_n,$$

$$x_i(s_i) \in \arg \min_{0 \leq x_i \leq 1 - s_i} c_i(x_i) + \mathbb{E}[J_{i+1}(s_i + x_i)], \quad i = 1, \dots, n - 1.$$

When the cost function is polynomial $c_i(x_i) = a_i x_i^p$, recursions above admit explicit solutions and can be implemented in polynomial time.

3. POLYNOMIAL PROPHET INEQUALITIES

In this section, we show that we can obtain precise prophet inequalities for polynomial cost functions $c_i(x_i) = a_i x_i^p$ when there is an outside option. The existence of an outside option is common in real world applications, as it is typically possible to procure part of the desired quantity of the commodity through channels outside of the online procurement process. In our model, an outside option corresponds to a *normalization*¹ of the model where $a_1 = 1$.

Our main results are given below.

THEOREM 1. *Consider polynomial cost functions of the form $c_i(x) = a_i x^p$, $p > 1$, with $\ell < a_1 = 1 < u$. For any independent distributions $D = (D_2, \dots, D_n)$ with D_i supported on $[\ell, u]$,*

$$\inf_{\pi \in \Pi} \frac{\mathbb{E}[\text{ALG}^\pi]}{\mathbb{E}[\text{OPT}]} \leq \left(1 + \frac{n}{\ell^{1/(p-1)}}\right)^{p-1}.$$

An interesting feature of the upper bound is that it is achieved with a *non-adaptive algorithm* which only uses $\mathbb{E}a_i$, $i = 1, \dots, n$, not their realized values.

THEOREM 2. *The bound $\left(1 + \frac{n}{\ell^{1/(p-1)}}\right)^{p-1}$ in Theorem 1 is tight.*

Note that Theorem 1 and Theorem 2 also hold for the linear case, $p = 1$. In this case, the gambler will take $x_1 = 1$ and the prophet will pick any $a_i = \ell$ for $i > 1$ in the worst case distribution where $a_i = \ell$ with probability $1 - \epsilon$ with $\epsilon \rightarrow 0$. The competitive ratio is

$$\sup_D \inf_{\pi \in \Pi} \frac{\mathbb{E}[\text{ALG}^\pi]}{\mathbb{E}[\text{OPT}]} = \lim_{p \rightarrow 1} \left(1 + \frac{n}{\ell^{1/(p-1)}}\right)^{p-1} = \frac{1}{\ell}.$$

4. A SIMPLE THRESHOLDING ALGORITHM

Motivated by the complexity of the dynamic program, we propose the following threshold-based algorithm, which we denote by **BALANCED THRESHOLD**: choose some threshold that strikes a balance between being low enough so that we can allocate a reasonable share to any producer whose cost is below it, and high enough so that sufficiently many producers' cost is below the threshold. Then, allocate $1/k$ to the first k producers whose cost is below the threshold, where k is chosen to minimize the cost. Intuitively, we would like k to be such that, with high probability, at least k producers

¹Without outside option (see Section 5), the competitive ratio scales with $\sqrt{u/\ell}$ even for 2 stages which grows unbounded for $u \rightarrow \infty$. This is because the first stage worst case distribution for this case is a point mass at $\sqrt{u\ell} \rightarrow \infty$ for $u \rightarrow \infty$, resulting in an infinite cost for the gambler. Normalizing $a_1 = 1$ removes this effect and guarantees a finite cost for the gambler.

have costs below the threshold; but not many more than k producers do. More concretely, the algorithm is parameterized by a threshold $\theta \in [\ell, u]$ and an integer $k \in [n]$. Let i be the current stage, and let s be the number of producers that have been allocated a non-zero share thus far (i.e., $s = |\{j : x_j \neq 0, j < i\}|$). Set $x_i = \frac{1}{k}$ if $s < k$ and either (1) $a_i \leq \theta$ or (2) $i \geq n - k + s$. Otherwise, set $x_i = 0$. Note that there are always exactly k producers from which the consumer buys $1/k$.

It is not clear how to choose “good” θ and k for general convex functions, however in the case of quadratic i.i.d. cost functions, the following intuition guides us. Let $p_{\theta,k}$ be the probability that the realization of at least k coefficients is at most θ . Then $\frac{p_{\theta,k}\theta + (1-p_{\theta,k})u}{k}$ is an upper bound on the expected cost for k and θ . To see this, note that, with probability $p_{\theta,k}$ we expect to have k coefficients at most θ , and so the cost is at most $\sum_{i=1}^k \frac{\theta}{k^2} = \frac{\theta}{k}$. Otherwise, with probability $(1 - p_{\theta,k})$, there is no non-trivial bound on the cost, and the best bound we can prove is u/k . This gives the following optimization for determining θ and k :

$$\min_{k,\theta} \frac{p_{\theta,k}\theta + (1-p_{\theta,k})u}{k},$$

The following theorem bounds the asymptotic competitive ratio of BALANCED THRESHOLD policies that are parameterized by the probability p_θ that $a_i \leq \theta$.

THEOREM 3. *Consider i.i.d. quadratic cost functions of the form $c_i(x) = a_i x^2$ and BALANCED THRESHOLD parameterized by $p_\theta \in (0, 1)$ such that θ satisfies $\Pr[a_i \leq \theta] = p_\theta$ and $k = np_\theta - \sqrt{n \log(\frac{2nu}{\ell})}$. Then, the asymptotic competitive ratio of BALANCED THRESHOLD (i.e., for $n \rightarrow \infty$) is at most $\frac{\theta}{\ell p_\theta} + o(1)$.*

5. NO OUTSIDE OPTION

Without outside option (relaxing the normalization that $a_1 = 1$), characterizing the prophet inequalities for general n stages and polynomial costs becomes much more complicated. To gain insight, we consider the two-stage quadratic setting.

We show that, similarly to the classical case, the worst case distributions are either a *point mass* (concentrated at a single point) or a *long-shot* distribution, which takes values only at the extreme points of the support. Specifically, we give an exact characterization of the worst case distributions: a_1 is point mass taking the value $\sqrt{u\ell}$, a_2 follows a long-shot distribution, taking the value ℓ w.p. $\frac{u-\sqrt{u\ell}}{u-\ell}$ and u w.p. $\frac{\sqrt{u\ell}-\ell}{u-\ell}$. This characterization yields the following bound on the competitive ratio:

THEOREM 4. *The competitive ratio for the two-stage convex prophet problem with quadratic cost functions is*

$$\sup_{D_1, D_2} \inf_{\pi \in \Pi} \frac{\mathbb{E}[\text{ALG}^\pi]}{\mathbb{E}[\text{OPT}]} = \frac{1}{4} \left(\sqrt{\frac{u}{\ell}} + \sqrt{\frac{\ell}{u}} \right) + \frac{1}{2}.$$

This bound highlights the improvement of the dynamic program over the proportionate partition scheme described earlier: $\frac{u}{\ell}$ improves to $O\left(\sqrt{u/\ell}\right)$. However, unlike in the classical prophet inequality setting, the two-stage convex prophet inequality is not rich enough to provide a worst

case example for convex prophet inequalities. More generally, we provide a bound on the competitive ratio in the n -stage setting.

THEOREM 5. *The competitive ratio for n -stage convex prophet problem with i.i.d. quadratic cost functions satisfies*

$$\sup_D \inf_{\pi \in \Pi} \frac{\mathbb{E}[\text{ALG}^\pi]}{\mathbb{E}[\text{OPT}]} \geq \frac{n-1}{n^2} \left(\sqrt{\frac{u}{\ell}} + \sqrt{\frac{\ell}{u}} \right) + \frac{(n-1)^2 + 1}{n^2}.$$

6. REFERENCES

- [1] M. Carrion, A. B. Philpott, A. J. Conejo, and J. M. Arroyo. A stochastic programming approach to electric energy procurement for large consumers. *IEEE Transactions on Power Systems*, 22(2):744–754, May 2007.
- [2] Ram Rajagopal, Eilyan Bitar, Pravin Varaiya, and Felix Wu. Risk-limiting dispatch for integrating renewable power. *International Journal of Electrical Power & Energy Systems*, 44(1):615 – 628, 2013.
- [3] Suresh Sethi, Houmin Yan, J Houzhong Yan, and Hanqin Zhang. An analysis of staged purchases in deregulated time-sequential electricity markets. *Journal of Industrial and Management Optimization*, 1(4):443–463, 2005.
- [4] Jayakrishnan Nair, Sachin Adlakha, and Adam Wierman. Energy procurement strategies in the presence of intermittent sources. In *Proceedings of ACM Sigmetrics*, 2014.
- [5] Allen J Wood and Bruce F Wollenberg. *Power generation, operation, and control*. John Wiley & Sons, 2012.
- [6] S. H. Low. Convex relaxation of optimal power flow, part i: Formulations and equivalence. *IEEE Transactions on Control of Network Systems*, 1(1):15–27, March 2014.
- [7] Subhmesh Bose, Desmond Cai, Steven Low, and Adam Wierman. The role of a market maker in networked cournot competition. In *Proceedings of IEEE CDC*, 2014.
- [8] Dimitri P Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena Scientific Belmont, MA, 1995.
- [9] Panqamala Ramana Kumar and Pravin Varaiya. *Stochastic systems: Estimation, identification, and adaptive control*. SIAM, 2015.
- [10] Sivadon Chaisiri, Bu-Sung Lee, and Dusit Niyato. Optimization of resource provisioning cost in cloud computing. *IEEE Transactions on Services Computing*, 5(2):164–177, 2012.
- [11] Jonatha Anselmi, Danilo Ardagna, John CS Lui, Adam Wierman, Yunjian Xu, and Zichao Yang. The economics of the cloud: Price competition and congestion. *Proceedings of NetEcon*, 2013.
- [12] Alain Bensoussan, Metin akany ld r m, and Suresh P. Sethi. A multiperiod newsvendor problem with partially observed demand. *Mathematics of Operations Research*, 32(2):322–344, 2007.
- [13] Lynwood A Johnson and Douglas C Montgomery. *Operations research in production planning, scheduling, and inventory control*, volume 6. Wiley New York, 1974.