

Convex Prophet Inequalities

Extended Abstract

Junjie Qin
UC Berkeley
qinj@berkeley.edu

Shai Vardi
California Institute of Technology
svardi@caltech.edu

Ram Rajagopal
Stanford University
ramr@stanford.edu

Adam Wierman
California Institute of Technology
adamw@caltech.edu

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 $[\ell, u]$.

CCS CONCEPTS

• Theory of computation → Online algorithms;

KEYWORDS

Prophet inequality, online algorithms, resource allocation, stochastic control

ACM Reference Format:

Junjie Qin, Ram Rajagopal, Shai Vardi, and Adam Wierman. 2018. Convex Prophet Inequalities: Extended Abstract. In *Proceedings of ACM SIGMETRICS 2018 (SIGMETRICS'18)*. ACM, Irvine, CA, USA, 2 pages.

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 [5, 10–12]. Thus, LSEs face an online decision problem where, given a forecast for the 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 [4, 9, 13], 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 [3, 8], cloud computing [1, 6], and inventory management [2, 7].

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 inequality 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.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
SIGMETRICS'18, June 2018, Irvine, California, USA
© 2018 Copyright held by the owner/author(s).

2 MODEL

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 $\mathbb{E}[\text{ALG}^\pi] = \mathbb{E} \left[\sum_{i=1}^n c_i(\pi_i(\mathcal{H}_i)) \right]$. 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: $\inf_{\pi \in \Pi} \mathbb{E}[\text{ALG}^\pi]$. For deriving convex prophet inequalities, we are interested in obtaining the *exact competitive ratio*

$$\sup_D \inf_{\pi \in \Pi} \frac{\mathbb{E}[\text{ALG}^\pi]}{\mathbb{E}[\text{OPT}]}, \quad (1)$$

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 (1) 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.

3 MAIN RESULTS

Consider the setting with 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* of the model where $a_1 = 1$.

Our main results are summarized below:

Optimal algorithm. We derive optimal online algorithm that is polynomial time implementable based on explicitly solving dynamic programming.

Exact competitive ratio (prophet inequalities). We obtain the exact value of (1) by characterizing worst case distributions:

THEOREM 1. Consider polynomial costs 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 $[\ell, u]$,

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

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

Simple thresholding algorithm. We develop a thresholding algorithm (BALANCED THRESHOLD) that is much simpler than dynamic program and has the following guarantees:

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 $k = np_\theta - \sqrt{n \log \left(\frac{2nu}{\ell} \right)}$. Then, the asymptotic competitive ratio of BALANCED THRESHOLD (i.e., for $n \rightarrow \infty$) is at most $\frac{\theta}{\ell p_\theta} + o(1)$.

REFERENCES

- [1] Jonatha Anselmi, Danilo Ardagna, John CS Lui, Adam Wierman, Yunjian Xu, and Zichao Yang. 2013. The Economics of the Cloud: Price Competition and Congestion. *Proceedings of NetEcon* (2013).
- [2] Alain Bensoussan, Metin AĖakanyÅsldÅsrÅsm, and Suresh P. Sethi. 2007. A Multiperiod Newsvendor Problem with Partially Observed Demand. *Mathematics of Operations Research* 32, 2 (2007), 322–344.
- [3] Dimitri P Bertsekas. 1995. *Dynamic programming and optimal control*. Vol. 1. Athena Scientific Belmont, MA.
- [4] Subhmesh Bose, Desmond Cai, Steven Low, and Adam Wierman. 2014. The role of a market maker in networked cournot competition. In *Proceedings of IEEE CDC*.
- [5] M. Carrion, A. B. Philpott, A. J. Conejo, and J. M. Arroyo. 2007. A Stochastic Programming Approach to Electric Energy Procurement for Large Consumers. *IEEE Transactions on Power Systems* 22, 2 (May 2007), 744–754.
- [6] Sivadon Chaisiri, Bu-Sung Lee, and Dusit Niyato. 2012. Optimization of resource provisioning cost in cloud computing. *IEEE Transactions on Services Computing* 5, 2 (2012), 164–177.
- [7] Lynwood A Johnson and Douglas C Montgomery. 1974. *Operations research in production planning, scheduling, and inventory control*. Vol. 6. Wiley New York.
- [8] Panqanamala Ramana Kumar and Pravin Varaiya. 2015. *Stochastic systems: Estimation, identification, and adaptive control*. SIAM.
- [9] S. H. Low. 2014. Convex Relaxation of Optimal Power Flow, Part I: Formulations and Equivalence. *IEEE Transactions on Control of Network Systems* 1, 1 (March 2014), 15–27. <https://doi.org/10.1109/TCNS.2014.2309732>
- [10] Jayakrishnan Nair, Sachin Adlakha, and Adam Wierman. 2014. Energy procurement strategies in the presence of intermittent sources. In *Proceedings of ACM Sigmetrics*.
- [11] Ram Rajagopal, Eilyan Bitar, Pravin Varaiya, and Felix Wu. 2013. Risk-limiting dispatch for integrating renewable power. *International Journal of Electrical Power & Energy Systems* 44, 1 (2013), 615 – 628.
- [12] Suresh Sethi, Houmin Yan, J Houzhong Yan, and Hanqin Zhang. 2005. An analysis of staged purchases in deregulated time-sequential electricity markets. *Journal of Industrial and Management Optimization* 1, 4 (2005), 443–463.
- [13] Allen J Wood and Bruce F Wollenberg. 2012. *Power generation, operation, and control*. John Wiley & Sons.