MATHEMATICS OF OPERATIONS RESEARCH
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 0364-765X | EISSN 1526-5471 | 00 | 0000 | 0001

INFORMS
DOI 10.1287/xxxx.0000.0000

Minimization Fractional Prophet Inequalities for Sequential Procurement

Junjie Qin
Department of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA. qinj@berkeley.edu

Shai Vardi
Krannert School of Management, Purdue University, West Lafayette, IN, USA. svardi@purdue.edu

Adam Wierman
California Institute of Technology, Pasadena, CA, USA. adamw@caltech.edu

We consider a minimization variant on the classical prophet inequality with monomial cost functions. A firm would like to procure some fixed amount of a divisible commodity from sellers that arrive sequentially. Whenever a seller arrives, the seller's cost function is revealed and the firm chooses how much of the commodity to buy. We first show that if one restricts the set of distributions for the coefficients to a family of natural distributions that includes, e.g., the uniform and truncated normal distributions, then there is a thresholding policy that is asymptotically optimal in the number of sellers. We then compare two scenarios, based on whether the firm has in-house production capabilities or not. We precisely compute the optimal algorithm's competitive ratio when in-house production capabilities exist, and for a special case when they do not. We show that the main advantage of the ability to produce the commodity in-house is that it shields the firm from price spikes in worst-case scenarios.

1. Introduction For many divisible commodities, there exists a marketplace where sellers arrive in a sequential fashion and a decision maker must make contracts to procure some amount of the commodity without knowing the prices that future sellers will offer. In electricity markets, for example, power generation capacity must be procured by load serving entities (LSEs) in order to meet demand. The LSEs have limited information about the future prices of other options for obtaining power supply when making procurement decisions. They must make an initial long-term planning decision of how much to invest in their own generation capacity and then, afterwards, procure power from third-party producers as needed in order to meet their demand forecasts. These contracts are made with generators that arrive over a span of years or even decades (Sethi et al. 2005, Varaiya et al. 2011). The production cost functions are typically convex, due to capacity limitations of generators and varying marginal costs across different generation options (Wood and Wollenberg 2012); while the precise form of the cost functions may be complicated, it is often modeled as a quadratic function in analytic work (Bose et al. 2014, Low 2014, Wood and Wollenberg 2012). This form of sequential procurement is not limited to electricity markets. Other examples include natural gas supply markets (Rajagopal et al. 2013) and provisioning resources in cloud computing (Chaisiri et al. 2012). Typically, the decision maker has some knowledge of the distributions that the cost functions will be drawn from (for example, using distributional estimates derived from previous interactions), but does not know the cost functions ahead of time. The uncertainty about future cost functions leaves the decision maker vulnerable to variability in the costs of the sellers, and it is unclear how to make procurement decisions in the face of this uncertainty.

LSEs typically maintain some generation capability to help insulate themselves from price fluctuations and respond to emergencies; enterprise environments often make use of both external and on-premise data centers, where the on-premise data center offers a safety net that serves to protect them from uncertain price fluctuations while the cloud integration offers scalability. While it is clear that in many cases, the ability to produce the commodity can help protect the decision maker from uncertainty, it is unclear precisely what type of insurance it provides, or how effective this ability is at providing insulation against price fluctuations, especially when the decision to invest in production has to be made before the sellers' prices are revealed.

1.1. Prophet inequalities. The setting of sequential procurement described above is a variation on the well-studied prophet inequality problem (Krengel and Sucheston 1977, 1978), where a gambler sequentially observes realizations r_i, $i = 1, \dots, n$, of a series of independent random variables R_i, $i = 1, \dots, n$, and needs to accept one of them. After each realization, r_i, the gambler must irrevocably decide whether to accept r_i or not. The gambler's goal is to maximize the (expectation of the) reward. The renowned classical prophet inequality states that the gambler can guarantee a reward whose expectation is at least half of the optimal reward of a prophet who foresees all of the realizations of the random variables ahead of time (Hill and Kertz 1992, Krengel and Sucheston 1978). The prophet inequality can be rephrased as follows: Upon observing r_i, the gambler must irrevocably select the value of $x_i \in \{0,1\}$, with the goal of maximizing $\sum_i x_i r_i$, subject to $\sum_i x_i = 1$.

Many variations of the prophet inequality have been studied (see Section 1.4), but to our knowledge only two (Disser et al. 2020, Esfandiari et al. 2015) have considered minimization variations; Esfandiari et al. (2015) show that the competitive ratio (of the minimization problem) is unbounded, and give a lower bound that is exponential in the number of random variables. The role of in-house production has not been studied in the context of prophet inequalities, but it is easy to see that it does not convey any meaningful advantage to the decision maker in most previously-studied settings. We expand upon this later.

1.2. Problem formulation. In the problem we study, a decision maker faces a sequence of non-negative independent random variables C_i with known distributions D_i supported on $\mathcal{C}_i$ for $i \in \{1,\dots,n\}$. Except when noted otherwise, let $C = (C_1,\dots,C_n)$, $D = (D_1,\dots,D_n)$. In every stage, a realization c_i of C_i is drawn and and the decision maker needs to procure some amount $x_i \in [0,1]$, such that the total amount procured is 1; that is, $\sum_{i=1}^n x_i = 1$. We denote the decision algorithm by an ordered set of n functions $\pi = (\pi_1,\dots,\pi_n)$, where $\pi_i : \mathcal{C}_1 \times \cdots \times \mathcal{C}_i \to [0,1]$. For any realization $c = (c_1,\dots,c_n)$ and any $p \geq 1$, the cost of the decision maker is

$$\mathrm{ALG}^\pi(c,p) := \sum_{i=1}^n c_i x_i^p, \tag{1}$$

where $x_i = \pi_i(c_1,\dots,c_i)$. Equation (1) models the cost incurred by the decision maker when procuring from a sequence of suppliers, where the choice of p depends on the cost structure of the application under consideration. The decision maker's goal is to minimize $\mathbb{E}[\mathrm{ALG}^\pi(C,p)]$. We call this the *minimization prophet inequality problem*. We denote the input to this problem by a pair (D,p); problem family $F = (\mathcal{D},p)$ consists of a positive real number $p \geq 1$ and a set of distributions $\mathcal{D}$.

Let $\mathrm{OPT}(c,p)$ denote the optimal value of the following optimization problem:

$$\begin{aligned} \min_c \quad & \sum_{i=1}^n c_i x_i^p \\ \text{s.t.} \quad & \sum_{i=1}^n x_i = 1; \\ & 0 \leq x_i \leq 1; \quad i = 1,\dots,n. \end{aligned} \tag{2}$$

We define the *competitive ratio* for a problem family $F = (\mathcal{D},p)$ as

$$\mathrm{CR}(F,n) = \inf_\pi \sup_{D\in\mathcal{D}} \frac{\mathbb{E}[\mathrm{ALG}^\pi(C,p)]}{\mathbb{E}[\mathrm{OPT}(C,p)]},$$

where the expectations are taken over the random variables $C_i \sim D_i$, $i = 1,\dots,n$.

If the in-house production option is present, the decision maker can produce the commodity at a cost function whose coefficient we normalize to 1; we denote this by setting D_1 such that $\mathbb{P}[C_1 = 1] = 1$. We define the competitive ratio for a problem family $F = (\mathcal{D},p)$ when the decision maker has in-house production as follows:

$$\mathrm{CR}_{\text{in-house}}(F,n) = \inf_\pi \sup_{D_2,\dots,D_n\in\mathcal{D}} \frac{\mathbb{E}[\mathrm{ALG}^\pi(C,p)]}{\mathbb{E}[\mathrm{OPT}(C,p)]},$$

where D_1 is as above and the expectations are taken over random variables $C_2 \sim D_2, \ldots, C_n \sim D_n$.

Our setting differs from the classical prophet inequality setting in several aspects: (i) our goal is to minimize the cost, as opposed to maximizing the reward; (ii) we allow the decision maker to assign fractional values to x_i; (iii) we consider non-linear as well as linear cost functions. In addition, we take a close look at the role that the ability to produce the commodity in-house plays in reducing the vulnerability of the decision maker to the uncertainty of future costs.

1.3. Our results We first show that as long as the distributions are "well behaved", the competitive ratio is asymptotically 1. More formally, we prove the following.

THEOREM 1. *Fix $p \geq 1$, $0 < \ell \leq u < \infty$ and let $F = (\mathcal{D}, p)$, where $\mathcal{D}$ is the set of probability distributions whose cumulative distribution functions (CDFs) are invertible, supported on $[\ell, u]$, and with L-Lipschitz inverse for some fixed constant L. Then*

$$\mathrm{CR}(F, n) = 1 + O\left(\frac{1}{\log n}\right).$$

The family of distributions of Theorem 1 includes many natural distributions including the uniform and the truncated normal and exponential distributions. To prove the theorem, we describe a multi-thresholding algorithm for the minimization prophet inequality. Here, the algorithm 'guesses' the realization of the random variables and computes amounts to buy from the sellers based on this guess. The algorithm then partitions the support into intervals. Whenever a seller arrives, the algorithm buys some pre-allocated amount, which depends on which interval the cost realizes to. We note that a similar result (i.e., that the competitive ratio is asymptotically 1) can be shown for the maximization prophet inequality case, except that in the maximization case, the decision maker sets $x_i = 1$ for some i, whereas here $x_i > 0$ for every i. The asymptotic result holds whether or not production capabilities exist, implying that for certain natural distributions, in-house production is not beneficial when the decision maker has access to many sellers. When we relax the distribution constraints, we get the following result, for decision makers with in-house production capabilities.

THEOREM 2. *Fix $p > 1$ and let $F = (\mathcal{D}, p)$, where $\mathcal{D}$ is the set of distributions whose support is contained in $[\ell, \infty)$, $\ell < 1$.*

$$\mathrm{CR}_{\textit{in-house}}(F, n) = \left(1 + \frac{n-1}{\ell^{1/(p-1)}}\right)^{p-1}. \tag{3}$$

To prove Theorem 2, we give matching upper and lower bounds on the competitive ratio. For the upper bound, we compute the exact competitive ratio of a suboptimal, non-adaptive algorithm, which is simpler to analyze than the optimal algorithm. To determine the input for which this competitive ratio is the highest, we express the problem maximizing the competitive ratio as an optimization problem, and decompose it, allowing us to determine the value of each distribution in the worst-case input. This gives an upper bound on the competitive ratio of the optimal algorithm, as the cost of the optimal algorithm must be at most the cost of the suboptimal one, on any input. For the lower bound, we give an example of an input for which the competitive ratio is tight, by explicitly computing the expected cost of the optimal algorithm for this input.

Theorem 2 holds for $p > 1$. For the case $p = 1$, Esfandiari et al. (2015) showed that the competitive ratio is unbounded as well. Although their setting isn't identical to ours[1], it is straightforward to adapt their proof to our setting, to show that for $F = (\mathcal{D}, 1)$ where $\mathcal{D}$ is the set of distributions supported on $[0, \infty)$, $\mathrm{CR}_{\text{in-house}}(F, n) \geq \frac{1.11^n}{6}$. Taking the limit $p \to 1$ in (3), we obtain the following corollary to Theorem 2.

Corollary 1. *Let $F = (\mathcal{D}, 1)$, where $\mathcal{D}$ is the set of distributions supported on $[\ell, \infty)$, $\ell < 1$.*

$$\mathrm{CR}_{\textit{in-house}}(F, n) = \frac{1}{\ell}.$$

Corollary 1 shows that the competitive ratio is, in fact, independent of n (and u), and only depends on ℓ, subsuming the bound of Esfandiari et al. (2015).

To understand the role in-house production plays in protecting the decision maker from uncertainty, we would like to determine the competitive ratio for the case where no production capabilities exist. Unfortunately, computing the parameterized (by p and n) competitive ratio for this setting is more challenging than for the case with in-house production capabilities. We remark upon this shortly, but first, we give the exact competitive ratio for the case $n = p = 2$:

Theorem 3. *Let ℓ, u be real numbers such that $0 < \ell \leq u < \infty$. Let $F = (\mathcal{D}, 2)$, where $\mathcal{D}$ is the set of distributions supported on $[\ell, u]$. Then*

$$\mathrm{CR}(F, 2) = \frac{1}{4}\left(\sqrt{\frac{u}{\ell}} + \sqrt{\frac{\ell}{u}}\right) + \frac{1}{2}.$$

Computing the exact competitive ratio even for this simple case is more involved than when there is in-house production. To prove the theorem, we compute the 'worst case' input distributions D_1, D_2 by analytically solving a two level optimization in the space of probability distributions under moment constraints. This turns out challenging as (i) both the numerator and denominator in the competitive ratio depend on the distributions, and (ii) the optimal cost involves harmonic means.

We can compare Theorem 3 with the competitive ratio with in-house production for F as in Theorem 3 with $n = 2$: $\mathrm{CR}_{\text{in-house}}(F, 2) = 1 + \frac{1}{\ell}$. From this, we can see that in-house production

[1] Their construction is for identical distributions, i.e., does not include the setting where the decision maker can produce the commodity at a fixed price.

shields the decision maker from price spikes, as u appears on in the competitive ratio only when there are no production capabilities. In other words, if we set $u \to \infty$, the competitive ratio is unbounded when no in-house production exists, but it is finite with in-house production.

In the classical (linear, maximization) prophet inequality setting, Hill and Kertz (1981b) show that, for any instance in which $n > 2$, there exists an instance in which $n = 2$ for which the competitive ratio is at least as bad. Therefore, the exact competitive ratio for $n = 2$ gives an upper bound on the competitive ratio for any $n \geq 2$ in the classical case. We have already shown that this is not the case with in-house production; the competitive ratio grows with n. We show that, when there are no production capabilities, the case of $n = 2$ does not give a tight bound on the competitive ratio either (at least for $p = 2$), by showing that when $n = 3$ there exist values of ℓ and u for which the competitive ratio is greater than that of Theorem 3. We also demonstrate numerically that the competitive ratio is not monotone in n, unlike when there is in-house production. This helps give some intuition as to why computing the precise competitive ratio for all $n \geq 2, p > 1$ is a challenging task. We leave it as an open question to analytically resolve the competitive ratio.

Proofs that do not appear in the main body can be found in the appendix.

1.4. Related work Krengel, Sucheston, and Garling (Krengel and Sucheston (1977, 1978)) first showed that the gambler's expected reward is at least half of the prophet's by analyzing the optimal stopping rule. A simple example involving two distributions shows that this is tight. Hill and Kertz (1981b) simplified and tightened the analysis of Krengel and Sucheston (1978) by showing that the competitive ratio is maximized when $n = 2$ and then (explicitly) computed the worst case distribution for that case. Samuel-Cahn (1984) showed that a simple thresholding algorithm—where the threshold depends on the expectation of the random variables—can achieve the optimal bound, and later Kleinberg and Weinberg (2012) used a different threshold to show the same bounds, and extend it to a more general setting. In some sense, our techniques are more aligned with those of the earlier proofs, as we also explicitly compute the worst-case distributions in order to obtain the upper bound. In contrast, however, it does not appear possible to compute the worst case distribution in our case, as the competitive ratio is not minimized for $n = 2$, and it is not clear how to compute the value of the optimal non-clairvoyant solution for larger values of n. Instead, we compute the worst-case distribution for a suboptimal algorithm that is easier to analyze, and show a matching lower bound. Many variants on the classical prophet in equality have been studied (see Correa et al. (2019), Hill and Kertz (1992)). Prophet inequalities have recently gained more attention as Hajiaghayi et al. (2007) and later Chawla et al. (2010) made a connection between prophet inequalities and online ad-auctions, sparking broad interest in prophet inequalities within the operations research and computer science community, e.g., (Alaei 2014, Alaei et al. 2012, Correa et al. 2020, Düetting et al. 2017, Kleinberg and Weinberg 2012).

The minimization version of this problem is much less well studied. To our knowledge, only two papers study it, (Disser et al. 2020, Esfandiari et al. 2015). Esfandiari et al. (2015) show that the competitive ratio is at least exponential in n, even when the distributions are identical, and that it is unbounded even if there are only three sellers; the latter result is shown by setting $u \to \infty$. Their analysis does not shed light on the underlying reasons behind these bounds. We show that the bound is determined solely by the lower bound of the distribution, not the upper bound or the number of distributions. In addition to providing a tight bound, our results explain the driving factors of the bounds of Esfandiari et al. (2015). Disser et al. (2020) study a variant of the prophet inequality that, similarly to our setting, cannot be described as a simple stopping problem. In their setting, at each time step an applicant arrives, and their cost is revealed. The decision maker must immediately make a hiring decision, and if the applicant is hired, their duration of employment (the number of time steps) is fixed. The constraint is that in every time step, at least one candidate needs to be under contract, and the goal is to minimize the total hiring cost. This setting shares several similarities with ours—it is a minimization problem and there is some notion of 'fractional allocations' (in our case, setting $x_i < 1$ and in theirs, hiring an applicant for less than the entire duration). Other than these similarities however, the settings are very different and it is difficult to compare the results.

Bounded distributions have also been studied in (the maximization version of the) classical prophet inequalities. Unsurprisingly, the competitive ratio is strictly less than 2 if the distributions are bounded (Hill 1983, Hill and Kertz 1981a). This is similar to our case, where the closer the upper and lower bounds of the distributions, the better the competitive ratio.

While there is a large literature on prophet inequalities in both the classical setting and generalizations, almost all of the papers consider linear cost functions ($p = 1$) and integral allocations (i.e., $a_i \in \{0, 1\}$). Rubinstein and Singla (2017) consider submodular cost functions, where the total reward is a submodular function of the rewards of the individual prizes; as mentioned previously, Disser et al. (2020) allow a notion of fractional allocations.

Perhaps the most related body of work to the nonlinear prophet problem posed in this paper is the literature studying online[2] convex optimization, e.g., see the surveys of Shalev-Shwartz (2012) and Hazan (2016). Packing problems have received special attention in this literature, e.g., Agrawal et al. (2014), Azar et al. (2016), Buchbinder and Naor (2009). However, the goal in the online convex optimization literature is primarily to bound either the regret or the competitive ratio in

[2] "Online" here is used as accepted in the computer science literature, equivalent to the way we define "sequential".

an adversarial setting (without distributional assumptions on the cost functions), whereas in the case of prophet inequalities, it is to bound the ratio of the expectations when the distributions are known a-priori. In particular, we assume that the online decision maker has information about the distribution from which the cost functions are sampled and optimizes the *average* (expected) cost with respect to such distributions. As such, there is a stream of related literature on online convex optimization that focuses on online convex optimization with *predictions*, e.g., Andrew et al. (2013), Chen et al. (2015, 2016). However, the assumptions and analytic tools in these papers are very different than the assumptions and tools used in problem studied here.

In terms of techniques, our derivation of closed form competitive ratios depends on optimizing over the space of probability distributions under moment constraints. Optimizations of this type arise in the classical problem of identifying maximum entropy distributions under moment constraints (Cover and Thomas 2012) and more generally in the problem of finding robust Bayes distributions (Grünwald and Dawid 2004). In those contexts, due to the specialty of the functionals optimized (e.g., entropy), one can often show that the worst-case distribution has a density and thus can utilize methods such as information inequality to pin down the form of the probability density function (e.g., exponential family). Beyond these cases, analytical characterizations of solutions of worst case distribution problems are extremely rare, and other work has typically relied on numerical techniques for solving such problems. For example, Delage and Ye (2010) tackle the distributional robust optimization problem under moment constraints. They obtain numerical solutions using sums of squares techniques. A key step in their work is solving the moment problem of the form $\max_{f_\xi} \mathbb{E}_\xi[h(x,\xi)]$, where the distribution of f_ξ is optimized over a family of distributions whose first and second moments are fixed. We encounter a similar problem: to bound the competitive ratio we view it as an optimization problem, where we (analytically) optimize for the worse case distribution. The functional that is optimized involves the ratio between the power mean of expected values of a sequence of random variables and the expected value of their power mean. To bound the competitive ratio we view it as an optimization problem and convert the ratio into nested optimizations. In the innermost optimization, the numerator is fixed, and we optimize over the family of distributions with a fixed first moment. Delage and Ye (2010) assume that h is concave in ξ in order to use the dual of the problem to arrive at an efficient numerical method to solve their problem. As the functional involved in our analysis is neither convex nor concave with respect to the optimization variables (the probability distribution to be optimized), their approach cannot be directly applied to our setting.

The comparison of the competitive ratio with and without in-house production capabilities is novel. Conventional wisdom states that make-or-buy decisions are essentially a matter of comparing internal and external production costs and choosing the least costly alternative. As this is a straightforward insight, most of the research on make-or-buy decisions focuses on more strategic settings with competition between firms (Arya et al. 2008, Salop and Scheffman 1987). However, more recently, there has been an increased interest in the make-or-buy decision in the face of uncertainty, e.g., Chen and Guo (2014), Niu et al. (2019), although this line of research still focuses on strategic settings. In contrast, in this paper we put aside the strategic issues and focus solely on the impact of the decision to invest in production on a single decision maker facing a sequential procurement problem with price uncertainty.

2. A multi-thresholding algorithm

In this section we prove Theorem 1.

THEOREM 1. *Fix $p \geq 1$, $0 < \ell \leq u < \infty$ and let $F = (\mathcal{D}, p)$, where $\mathcal{D}$ is the set of probability distributions whose cumulative distribution functions (CDFs) are invertible, supported on $[\ell, u]$, and with L-Lipschitz inverse for some fixed constant L. Then*

$$\mathrm{CR}(F, n) = 1 + O\left(\frac{1}{\log n}\right).$$

To prove the theorem, we describe an algorithm that attains this competitive ratio. The algorithm is a multi-threshold algorithm: one such that the amount allocated at stage i does not depend on the exact value of x_i, only the range in which it falls (and, possibly, the number of $x_j : j < i$ such that x_j also falls in this interval).

To formally define the algorithm, we need some notation. Fix $r = \lfloor \log n \rfloor$, and set $\epsilon = \frac{u-\ell}{r}$. Partition $[\ell, u]$ into r intervals of size ϵ each: $[\ell, \ell+\epsilon), [\ell+\epsilon, \ell+2\epsilon), ..., [u-\epsilon, u]$. For each $j \in \{1, \ldots, r-1\}$, let I_j denote the interval $[\ell+(j-1)\epsilon, \ell+j\epsilon)$, and let I_r denote $[u-\epsilon, u]$. For each $j \in \{1, \ldots, r\}$ let $\mathbb{1}_i(j)$ denote the indicator function: $\mathbb{1}_i(j) = 1$ iff $c_i \in I_j$. Set $\mathbb{1}(j) = \sum_{i=1}^n \mathbb{1}_i(j)$: the random variable denoting the number of values c_i for which $c_i \in I_j$.

The algorithm INTERVALS is as follows. Fix $\delta = \Theta\left(\frac{1}{\log^3 n}\right)$. For each $j \in \{1, \ldots, r\}$, compute $\mathbb{E}[\mathbb{1}(j)]$ and define $s_j := \max\{\lfloor \mathbb{E}[\mathbb{1}(j)] - \delta n \rfloor, 0\}$. Assume that for $j \in \{1, \ldots, r-1\}$, s_j cost coefficients will be (exactly) $\ell + j\epsilon$, and that all other coefficients will be u. Compute the optimal allocation based on these assumptions; that is, the optimal allocations assuming that exactly s_j coefficients are $\ell + j\epsilon$, $j \in \{1, \ldots, r-1\}$ and the rest, denoted by ν are u. Denote this allocation by $\tau = (\tau_1, \ldots, \tau_r)$, where τ_j denotes the amount purchased per realization at price $\ell + j\epsilon$ (i.e., the total amount purchased at price $\ell + j\epsilon$ is $s_k\tau_k$). Note that $\nu = n - \sum_{j=1}^{r-1} s_j \leq r\delta n$ and $\sum_{k=1}^r s_k \tau_k = 1$.

In round i, C_i is realized and c_i falls in some interval I_k. If at least s_k of $c_1, \ldots, c_{i-1}$ fell in I_k, we say that slot k is *full.* If $k = r$ or slot k is full, we "place i in slot r", and set $x_i = \tau_r$. Otherwise, place i in slot k and set $x_i = \tau_k$.

To prove Theorem 1, we need the following lemma, whose proof appears in the appendix.

Lemma 1. *For any* $c=(c_1,\dots,c_n)$ *and* $p>1$,

$$\mathrm{OPT}(c,p)=\left(\sum_{i=1}^{n} c_i^q\right)^{\frac{1}{q}},$$

where $q=-\frac{1}{p-1}$.

Proof of Theorem 1. Let $q=-\frac{1}{p-1}$. For any probability distribution whose inverse is L-Lipschitz, the probability that it realizes to any interval is at least ϵ/L, hence $\mathbb{E}[\mathbb{1}(j)]\geq \epsilon n/L$, and $s_j\geq \epsilon n/L-\delta n$. Let $\Phi(j)$ denote the event $\mathbb{1}(j)\leq \mathbb{E}[\mathbb{1}(j)]-\delta n$. By the additive Chernoff bound, e.g., Alon and Spencer (2008), $\mathbb{P}[\Phi(j)]\leq e^{-2n\delta^2}$. Let $\epsilon'=re^{-2n\delta^2}$. By the union bound, the probability that some $\Phi(j)$ occurs ($j\in\{1,\dots,r\}$) is at most ϵ', hence with probability at least $1-\epsilon'$ none of the events $\Phi(j), j\in\{1,\dots,r\}$ occur. That is, with probability at least $1-\epsilon'$, at least s_j coefficients are revealed to be in I_j, for every j.

To bound the cost of Intervals, we first compute the cost if the assumptions of Intervals hold, using Lemma 1:

$$C=\left(s_1(\ell+\epsilon)^q+s_2(\ell+2\epsilon)^q+\cdots+s_ru^q+\nu u^q\right)^{\frac{1}{q}}.$$

For every i that has been placed in slot j, $\ell+j\epsilon$ is an upper bound on the realized price of C_i. Therefore

$$s_1(\ell+\epsilon)\tau_1^p+s_2(\ell+2\epsilon)\tau_2^p+\cdots+(s_r+\nu)u\tau_r^p \tag{4}$$

is an upper bound on the cost of the algorithm when no $\Phi(j)$ occurs. This is exactly the optimal solution of (2) under the assumptions of Intervals, hence equal to C. Therefore C is an upper bound on the cost incurred by the algorithm if no $\Phi(j)$ occurs. If some $\Phi(j)$ occurs, we take the trivial upper bound of u for the cost of Intervals.

Hence

$$\mathbb{E}\left[\mathrm{ALG}^{\textsc{Intervals}}(C,p)\right]\leq(1-\epsilon')\left(s_1(\ell+\epsilon)^q+s_2(\ell+2\epsilon)^q+\cdots+s_ru^q+\nu u^q\right)^{\frac{1}{q}}+\epsilon' u.$$

From Lemma 1 and using the trivial lower bound on the cost ℓn^{p-1},

$$\mathbb{E}[\mathrm{OPT}(C,p)]\geq(1-\epsilon')\left(s_1\ell^q+s_2(\ell+\epsilon)^q+\cdots+s_r(u-\epsilon)^q+\nu\ell^q\right)^{\frac{1}{q}}+\epsilon'\ell n^{p-1}.$$

Set

$$\begin{aligned}
B_{\textsc{Intervals}}&=\left(s_1(\ell+\epsilon)^q+s_2(\ell+2\epsilon)^q+\cdots+(s_r+\nu)u^q\right)^{\frac{1}{q}}\\
&\leq\left((s_1+\nu)(\ell+\epsilon)^q+(s_2+\nu)(\ell+2\epsilon)^q+\cdots+(s_r+\nu)u^q\right)^{\frac{1}{q}},\\
B_{\mathrm{OPT}}&=\left((s_1+\nu)\ell^q+s_2(\ell+\epsilon)^q+\cdots+s_r(u-\epsilon)^q\right)^{\frac{1}{q}}\\
&\geq\left(s_1\ell^q+s_2(\ell+\epsilon)^q+\cdots+s_r(u-\epsilon)^q\right)^{\frac{1}{q}},
\end{aligned}$$

$\epsilon_j = \frac{(s_j+\nu)(\ell+j\epsilon)}{s_j(\ell+(j-1)\epsilon)}$, $\epsilon^\dagger = \frac{\epsilon' u}{(\epsilon n)^{1/q}\ell}$ and $\epsilon^* = \frac{(r\delta n)(\ell+\epsilon)}{(\epsilon n/L-\delta n)\ell}$. Then

$$\begin{aligned}
\frac{\mathbb{E}\left[\mathrm{ALG}^{\textsc{Intervals}}(C,p)\right]}{\mathbb{E}\left[\mathrm{OPT}(C,p)\right]} &\leq \frac{(1-\epsilon')B_{\textsc{Intervals}}+\epsilon' u}{(1-\epsilon')B_{\mathrm{OPT}}} \\
&\leq \frac{B_{\textsc{Intervals}}+\epsilon' u}{B_{\mathrm{OPT}}} \\
&\leq \frac{B_{\textsc{Intervals}}}{B_{\mathrm{OPT}}}+\epsilon^\dagger \qquad (5)\\
&\leq \max_j \epsilon_j + \epsilon^\dagger \qquad (6)\\
&\leq \frac{(\epsilon n/L+(r-1)\delta n)(\ell+\epsilon)}{(\epsilon n/L-\delta n)\ell}+\epsilon^\dagger \qquad (7)\\
&= 1+\frac{\epsilon}{\ell}+\epsilon^*+\epsilon^\dagger \\
&= 1+\Theta\left(\frac{1}{\log n}\right).
\end{aligned}$$

Inequality (5) holds because at least one value of s_j will be at least ϵn. To see why (6) holds, note that $\frac{\sum_j c_j}{\sum_j b_j}$ is a weighted sum of $\frac{c_j}{b_j}$:

$$\frac{\sum_j c_j}{\sum_j b_j} = \sum_j \frac{c_j}{b_j}\frac{b_j}{\sum_i b_i},$$

and a weighted sum cannot be larger than the largest term. Inequality (7) holds as for every $j \geq 1, \epsilon_j \leq \frac{\ell+\epsilon}{\ell}$. □

3. Bounds for arbitrary functions (with in-house production)

In this section, we prove Theorem 2, which gives the competitive ratio when the set of price functions is constrained only by a lower bound. The competitive ratio is given as a function of both p, the power of the monomial in the cost function, and n, the number of sellers.

THEOREM 2. *Fix $p>1$ and let $F=(\mathcal{D},p)$, where $\mathcal{D}$ is the set of distributions whose support is contained in $[\ell,\infty)$, $\ell<1$.*

$$\mathrm{CR}_{\textit{in-house}}(F,n) = \left(1+\frac{n-1}{\ell^{1/(p-1)}}\right)^{p-1}. \qquad (3)$$

Theorem 2 follows from matching upper and lower bounds. To obtain the upper bound, consider the following non-adaptive algorithm, referred to as algorithm NA, for the problem. The algorithm is non-adaptive in that it decides all of the amounts ahead of time. The algorithm NA purchases the amounts as described by the solution to the following optimization problem.

$$\begin{aligned}
\min_x \quad & \mathbb{E}\left[\sum_i C_i x_i^p\right] \\
\text{s.t.} \quad & \sum_i x_i = 1; \qquad (8)\\
& 0 \leq x_i \leq 1; \quad i=1,\ldots,n.
\end{aligned}$$

First, we need an auxiliary lemma that is used to identify the 'worst-case' distributions for NA: the distribution for which the ratio of $\mathbb{E}[\mathrm{ALG}^{\mathrm{NA}}(C,p)]$ and $\mathbb{E}[\mathrm{OPT}(C,p)]$ is maximized.

LEMMA 2. *Let ℓ, μ and u be real numbers such that $0 < \ell \le \mu \le u < \infty$, $\mathcal{C}$ be the set of all random variables C whose distribution is supported on $[\ell, u]$, such that $\mathbb{E}[C] = \mu$, and f be a strictly concave function. Then a solution of*

$$\min_{C \in \mathcal{C}} \quad \mathbb{E}[f(C)] \tag{9}$$

if it exists, is

$$C = \begin{cases} u, & \text{with probability } p_u = \frac{\mu-\ell}{u-\ell}, \\ \ell, & \text{with probability } p_\ell = \frac{u-\mu}{u-\ell}. \end{cases}$$

We derive the competitive ratio of NA explicitly in terms of the power mean: for a sequence of positive real numbers $c = (c_1, \dots, c_n)$, the *power mean* of c with exponent q is defined as $M_q(c) = \left(\frac{1}{n}\sum_{i=1}^n c_i^q\right)^{\frac{1}{q}}$. ($M_{-1}(c)$ is the harmonic mean $H(c) = n\left(\frac{1}{c_1} + \cdots + \frac{1}{c_n}\right)^{-1}$.)

LEMMA 3. *Fix $p > 1$, and let $q = -\frac{1}{p-1}$. Let ℓ and u be real numbers such that $0 < \ell \le u < \infty$ and C be a n-dimensional random vector with independent elements each of which has its distribution supported on $[\ell, u]$. Then*

$$\frac{\mathbb{E}\left[[\mathrm{ALG}^{\mathrm{NA}}(C,p)\right]}{\mathbb{E}[\mathrm{OPT}(C,p)]} = \frac{M_q(\mathbb{E}[C])}{\mathbb{E}[M_q(C)]}.$$

Proof of Theorem 2. We prove the upper bound here; the proof of the lower bound is deferred to the appendix.

First note that the competitive ratio is non-increasing in u, as any distribution supported on $[\ell, u]$ is also supported on $[\ell, u']$, where $u' \ge u$. This means the supremum of the competitive ratio must occur for $u \to \infty$. We work with extended real numbers, $\mathbb{R} \cup \{-\infty, \infty\}$ with $0 \cdot \infty := 0$ (e.g., McShane 1983); this allows the seller to set the cost to ∞, which simplifies the asymptotic analysis.

Recall that D_1 is a point-mass distribution with support $\{1\}$. We would like to determine the worst-case distributions $D_2, \dots, D_n$ for NA with $u = \infty$. Let $D = D_2, \dots, D_n$ and $C = C_2, \dots, C_n$, where $C_i \sim D_i$, $i \in \{2, \dots, n\}$. By Lemma 3, the competitive ratio of NA is

$$\sup_{D} \frac{M_q(1, \mathbb{E}[C])}{\mathbb{E}[M_q(1, C)]}, \tag{10}$$

where $q = -1/(p-1)$. This optimization can be written as

$$\sup_{(D_3,\dots,D_n)} \sup_{D_2} \frac{M_q(1,\mathbb{E}[C])}{\mathbb{E}[M_q(1,C)]} = \sup_{(D_3,\dots,D_n)} \sup_{\mu_2} \sup_{D_2 \in \mathcal{D}(\mu_2)} \frac{M_q(1,\mathbb{E}[C])}{\mathbb{E}[M_q(1,C)]},$$

where $\mathcal{D}(\mu)$ is the set of all distributions whose mean is μ. Consider the inner optimization $\left(\sup_{D_2 \in \mathcal{D}(\mu_2)} \frac{M_q(1,\mathbb{E}[C])}{\mathbb{E}[M_q(1,C)]}\right)$. Here, $D_3, \dots, D_n$ and μ_2 are fixed, hence $M_q(1, \mathbb{E}[C])$ is a constant. The inner optimization is therefore is equivalent to:

$$\begin{aligned} \inf_{D_2} \quad & \mathbb{E}[M_q(1, C_2, \dots, C_n)] = \mathbb{E}_{C_2}[f(C_2)] \\ \text{s.t.} \quad & \mathbb{E}[C_2] = \mu_2, \end{aligned} \tag{11}$$

where $f(c)=\mathbb{E}_{C_3,\dots,C_n}[M_q(1,C)|C_2=c]$ is a strictly concave function in c. By Lemma 2, the solution to (11) is a long-shot distribution, with

$$C_2=\begin{cases}u, & \text{with probability } p_u=\frac{\mu-\ell}{u-\ell},\\ \ell, & \text{with probability } p_\ell=\frac{u-\mu}{u-\ell}.\end{cases}$$

.

Note that while p_u and p_ℓ are defined explicitly for every $\mu_2<\infty$, they are not necessarily well-defined when $\mu_2=\infty$. In this case, there are (exactly) two possible values of μ_2. Either $\mu_2=\infty$ (implying $p_u>0$), or $\mu_2=\infty$, or $\mu_2=\ell$ (implying $p_u=0$). However, if $\mu_2=\infty$, $p_u=\frac{\infty}{\infty}$, which is undefined. We can nevertheless use this characterization to determine the value of p_u for which the supremum is attained by considering all possible values of p_u when $\mu_2=\infty$. Set $\epsilon_2=p_u$, therefore $p_\ell=(1-\epsilon_2)$. To simplify the notation, let $X=1+\sum_{j=3}^n\mu_j^q$, $Y=\mathbb{E}_{C_3,\dots,C_n}[M_q(1,C_2,C_3,\dots,C_n)|C_2=\ell]$ and $Z=\mathbb{E}_{C_3,\dots,C_n}[M_q(1,C_2,C_3,\dots,C_n)|C_2=u]$. From the characterization C_2, we have that $M_q(1,\mathbb{E}[C])=n^{-1/q}\left(X+\mu_2^q\right)^{\frac{1}{q}}$ and $\mathbb{E}[M_q(1,C)]=(1-\epsilon_2)Y+\epsilon_2 Z$.

If $\mu_2=\ell$,

$$\sup_{D_2\in\mathcal{D}(\mu_2=\ell)}\frac{M_q(1,\mathbb{E}[C])}{\mathbb{E}[M_q(1,C)]}=\frac{(X+\ell^q)^{\frac{1}{q}}}{Y}. \tag{12}$$

If $\mu_2=\infty$,

$$\sup_{D_2\in\mathcal{D}(\mu_2=\infty)}\frac{M_q(1,\mathbb{E}[C])}{\mathbb{E}[M_q(1,C)]}=\frac{\left(n^{-1/q}X\right)^{\frac{1}{q}}}{(1-\epsilon_2)Y+\epsilon_2 Z},$$

as $q<0$. As $Y<Z$, the supremum occurs for $\epsilon_2\to 0$, and as (12) is less than $\frac{\left(n^{-1/q}X\right)^{\frac{1}{q}}}{Y}$, we have that

$$\sup_{\mu_2}\sup_{D_2\in\mathcal{D}(\mu_2)}\frac{M_q(1,\mathbb{E}[C])}{\mathbb{E}[M_q(1,C)]}=\frac{\left(n^{-1/q}X\right)^{\frac{1}{q}}}{Y},$$

and $\mu_2=\infty$. As this result holds for *any* $D_3,\dots,D_n$, we can apply the same argument to each of them (keeping the remaining distributions constant), and conclude that the supremum $\sup_D\frac{M_q(1,\mathbb{E}[C])}{\mathbb{E}[M_q(1,C)]}$ occurs when, $i\in\{2,\dots,n\}$, D_i is a long-shot distribution, with $\mu_i=\infty$, where $C_i=\ell$ with probability $1-\epsilon_i$, $\epsilon_i\to 0$.

We therefore obtain

$$\begin{aligned}\sup_D\frac{M_q(1,\mathbb{E}[C])}{\mathbb{E}[M_q(1,C)]}&=\lim_{\epsilon_i\to 0,\forall i}\frac{n^{-1/q}}{\mathbb{E}[M_q(1,C)]}\\&=n^{-1/q}\big(M_q(1,\underbrace{\ell,\dots,\ell}_{(n-1)\text{ terms}})\big)^{-1}\\&=(1+(n-1)\ell^q)^{-1/q}.\end{aligned}$$

□

4. Bounds for arbitrary functions (without in-house production) We now turn to the case when the decision maker cannot produce the commodity. We prove Theorem 3, which gives the competitive ratio for $p = n = 2$ when the decision maker has no production capabilities. In addition, we show that—similarly to when the decision maker has production capabilities, but unlike the classical prophet inequality scenario—the case $n = 2$ does not set an upper bound on the competitive ratio for arbitrary n. We further show that—unlike when the decision maker has production capabilities—the competitive ratio is not monotone in n. These results give some intuition as to why obtaining tight results for $n > 2$ appears to be a challenging task.

4.1. Proof of Theorem 3

THEOREM 3. *Let ℓ, u be real numbers such that $0 < \ell \leq u < \infty$. Let $F = (\mathcal{D}, 2)$, where $\mathcal{D}$ is the set of distributions supported on $[\ell, u]$. Then*

$$\mathrm{CR}(F, 2) = \frac{1}{4}\left(\sqrt{\frac{u}{\ell}} + \sqrt{\frac{\ell}{u}}\right) + \frac{1}{2}.$$

We first show that the cost of both the prophet and the optimal algorithm can be written in terms of harmonic means. We note that this result holds for any number of sellers.

LEMMA 4. *Consider the minimization prophet inequality problem with input $(D, 2)$, $D = \{D_1, D_2, \ldots, D_n\}$. For any realization $c_1, \ldots, c_n$ of the random variables $C_1, \ldots, C_n$, the prophet's cost is $\mathrm{OPT} = \frac{H(c_1,\ldots,c_n)}{n}$.*

The decision maker's expected cost, i.e., the cost achieved by the optimal algorithm DP*, is*

$$\mathbb{E}[\mathrm{ALG}^{\mathrm{DP}}(C, p)] = \mathbb{E}[b_1],$$

and

$$x_i = \frac{b_i}{c_i}\prod_{j=1}^{i-1}\left(1 - \frac{b_j}{c_j}\right), \quad i = 1, \ldots, n, \tag{13}$$

where $b_n = C_n$ and

$$b_i = \frac{H(c_i, \mathbb{E}[b_{i+1}])}{2}, \quad i = 1, \ldots, n-1.$$

We use this result with $n = 3$ in Subsection 4.2 to show that the competitive ratio for the case of $n = 2$ is not the worst for all n. For $n = 2$, we can more succinctly phrase Lemma 4 as follows:

LEMMA 5. *Consider the minimization prophet inequality problem with input $(D, 2)$, $D = \{D_1, D_2\}$. The prophet's cost is $\mathrm{OPT} = \frac{H(c_1,c_2)}{2}$, and the expected cost of the optimal online algorithm conditioned on C_1 realizing to c_1 is $\frac{H(c_1,\mathbb{E}[C_2])}{2}$, where $H(\cdot,\cdot)$ denotes the Harmonic mean.*

We now express the competitive ratio in terms of the prophet and decision maker's costs, and decompose this expression as follows:

$$\inf_{\pi} \sup_{D_1,D_2} \frac{\mathbb{E}[\mathrm{ALG}^{\pi}(C,2)]}{\mathbb{E}[\mathrm{OPT}(C,2)]} = \sup_{D_1,D_2} \frac{\mathbb{E}_{D_1}[H(C_1,\mathbb{E}_{D_2}[C_2])]}{\mathbb{E}_{D_1,D_2}[H(C_1,C_2)]}. \tag{14}$$

We can treat (14) as an optimization problem, however the objective function is not concave, and in addition, (14) is an infinite-dimensional optimization (i.e., a semi-infinite program with an infinite number of variables).

We address these issues by decomposing (14) as follows:

$$\sup_{D_1,D_2} \frac{\mathbb{E}_{D_1}[H(C_1,\mathbb{E}_{D_2}[C_2])]}{\mathbb{E}_{D_1,D_2}[H(C_1,C_2)]} = \sup_{D_1} \sup_{\mu_2} \sup_{D_2:\mathbb{E}[C_2]=\mu_2} \frac{\mathbb{E}_{D_1}[H(C_1,\mathbb{E}_{D_2}[C_2])]}{\mathbb{E}_{D_1,D_2}[H(C_1,C_2)]}. \tag{15}$$

The proof of Theorem 3 now proceeds in three parts: (i) computing the 'worst-case' D_1, for any D_2, (ii) given D_1, finding the 'worst-case' mean of D_2, $\mu_2 = \mathbb{E}[C_2]$, and (iii) given D_1 and μ_2, finding the 'worst-case' D_2. We take a bottom-up approach, solving (iii) first.

Computing the worst-case D_2. Fixing D_1 and $\mu_2 \in [\ell, u]$, numerator of the inner optimization of (15) is constant. Computing the supremum is equivalent to computing the infimum of the denominator; hence the problem of finding the worst-case D_2 can be formally stated as

$$\begin{aligned} \inf_{D_2} \quad & \int_{\ell}^{u}\int_{\ell}^{u} H(c_1,c_2)\, dD_1(c_1)\, dD_2(c_2) \\ \text{s.t.} \quad & \int_{\ell}^{u} c_2\, dD_2(c_2) = \mu_2, \end{aligned} \tag{16}$$

where D_2 belongs to the set of all probability measures and Lebesgue integration is used to express expectations (Dudley 2002).

We approach problem (16) by first considering a simpler version where D_1 is taken to be a point mass, therefore C_1 is some fixed constant c_1:

$$\begin{aligned} \inf_{D_2} \quad & \int_{\ell}^{u} H(c_1,c_2)\, dD_2(c_2) \\ \text{s.t.} \quad & \int_{\ell}^{u} c_2\, dD_2(c_2) = \mu_2. \end{aligned} \tag{17}$$

Note that Problem (17) is exactly of the form given in Lemma 2, and hence the worst-case distribution is as follows:

$$C_2 = \begin{cases} \ell, & \text{with probability } \frac{u-\mu_2}{u-\ell}, \\ u, & \text{with probability } \frac{\mu_2-\ell}{u-\ell}. \end{cases} \tag{18}$$

As the solution of this problem does not depend on the value of c_1, the same result holds for any distribution D_1.

Finding the worst-case D_1 and μ_2. Before attempting to compute μ_2, we turn to the problem of identifying the worst-case D_1 as a function of D_2. We show that D_1 does not depend on the value of μ_2, hence D_1 can be determined without knowing μ_2.

LEMMA 6. *Fix $0<\ell\leq\mu_2\leq u<\infty$ and suppose that D_2 takes the value ℓ with probability $\frac{u-\mu_2}{u-\ell}$ and u with probability $\frac{\mu_2-\ell}{u-\ell}$. The distribution D_1 for which (15) is maximized is such that $C_1=\sqrt{\ell u}$ with probability* 1*.*

Proof of Lemma 6. First consider the case that D_1 is supported on a finite set of m arbitrary points $\{c_1^{(1)},\ldots,c_1^{(m)}\}$ with $\ell=c_1^{(1)}<\cdots<c_1^{(m)}=u$, where $m\in\mathbb{N}$ is also arbitrary. Assume that for some $k^\star\in[m], c_i^{(k^\star)}=\sqrt{\ell u}$; we will show that this is without loss of generality. The problem of identifying the worst-case D_1 can be written as follows, where for convenience, we use min instead of max,

$$\begin{aligned}\min_{p\in\mathbb{R}^m}\quad & \frac{q^\top p}{h^\top p}\\ \text{s.t.}\quad & \mathbf{1}^\top p=1,\\ & p\geq 0,\end{aligned}\tag{19}$$

where $q\in\mathbb{R}^m$ has entries $q_k=\frac{u-\mu_2}{u-\ell}H(c_1^{(k)},\ell)+\frac{\mu_2-\ell}{u-\ell}H(c_1^{(k)},u)$ and $h\in\mathbb{R}^m$ has entries $h_k=H(c_1^{(k)},\mu_2)$. This is a linear-fractional program which is equivalent to the following LP

$$\begin{aligned}\min_{y\in\mathbb{R}^m,z\in\mathbb{R}}\quad & q^\top y\\ \text{s.t.}\quad & \mathbf{1}^\top y=z,\\ & h^\top y=1,\\ & y,z\geq 0,\end{aligned}$$

with the change of variable $y=p/h^\top p$, $z=1/h^\top p$. Simplifying, we get

$$\begin{aligned}\min_{y\in\mathbb{R}^n}\quad & q^\top y\\ \text{s.t.}\quad & h^\top y=1,\\ & y\geq 0.\end{aligned}\tag{20}$$

Since $h_k>0$ and $q_k>0$ for all k, a solution of (20) is clearly to place all mass on the coordinate with the smallest ratio q_k/h_k. Consider the ratio as a function of c_1

$$r(c_1)=\frac{\frac{u-\mu_2}{u-\ell}H(c_1,\ell)+\frac{\mu_2-\ell}{u-\ell}H(c_1,u)}{H(c_1,\mu_2)}=\frac{(c_1+\mu_2)(c_1\mu_2+u\ell)}{(c_1+u)(c_1+\ell)\mu_2}.$$

Compute

$$(\log r(c_1))'=\frac{(c_1^2-u\ell)(u-\mu_2)(\mu_2-\ell)}{(c_1+\mu_2)(c_1\mu_2+u\ell)(c_1+u)(c_1+\ell)},$$

which has a single root $c_1 = \sqrt{\ell u}$ on $c_1 \in \mathbb{R}_+$. Furthermore, $(\log r(c_1))' < 0$ for $c_1 < \sqrt{\ell u}$ and $(\log r(c_1))' > 0$ for $c_1 > \sqrt{\ell u}$, so $c_1 = \sqrt{\ell u}$ is a minimizer for $\log(r(c_1))$ as well as $r(c_1)$. Thus the solution to (19) is $p_{k^\star} = 1$ and $p_k = 0$ for all $k \neq k^\star$. Note that as this solution minimizes (19) when $\sqrt{\ell u} \in \{c_1^{(1)}, \ldots, c_1^{(m)}\}$ *regardless of the choice of* m, it must hold that this is optimal for completely arbitrary choice of $\{c_1^{(1)}, \ldots, c_1^{(m)}\}$: assume that there is a better choice of $\{c_1^{(1)}, \ldots, c_1^{(m)}\}$, where $\sqrt{\ell u} \notin \{c_1^{(1)}, \ldots, c_1^{(m)}\}$, and denote its solution by X (i.e., X is the vector of allocations to each coordinate). Then X is also a solution for $\{c_1^{(1)}, \ldots, c_1^{(m)}\} \cup \sqrt{\ell u}$, a contradiction to the optimality of placing all of the mass on $\sqrt{\ell u}$. Therefore, the worst case distribution among any distributions supported on a finite number of points in $[\ell, u]$ is $C_1 = \sqrt{\ell u}$ with probability 1.

We proceed to show that the same solution structure holds when we consider any distributions supported on the continuous interval $[\ell, u]$. Suppose otherwise, i.e., there exists a different distribution $\tilde{D}_1$ supported on $[\ell, u]$ that leads to a strictly better objective value for optimization (15). Denote the random variable with distribution $\tilde{D}_1$ by $\tilde{C}_1$. Denote the distribution that places all the probability mass at $\sqrt{\ell u}$ by $D_1^\star$ and the corresponding random variable by $C_1^\star$. Let

$$J(D_1) := \frac{\mathbb{E}_{D_1, D_2}[H(C_1, C_2)]}{\mathbb{E}_{D_1}[H(C_1, \mathbb{E}_{D_2}[C_2])]},$$

where D_2 is fixed as (18). Note here J is the objective function in (15) with the numerator and denominator flipped (as we are working with minimization in this proof) and D_2 and μ_2 fixed. Under our hypothesis here, there exists an $\epsilon > 0$ such that

$$J(D_1^\star) > J(\tilde{D}_1) + \epsilon.$$

Consider a discrete approximation of $\tilde{D}_1$ supported on $\{\ell, \ell + \frac{u-\ell}{T}, \ldots, u - \frac{u-\ell}{T}, u\}$, denoted by $\tilde{D}_1(T)$ with the corresponding random variable denoted by $\tilde{C}_1(T)$, where

$$\mathbb{P}\left(\tilde{C}_1(T) = \ell + \frac{t(u-\ell)}{T}\right) = \mathbb{P}\left(\ell + \frac{t(u-\ell)}{T} \leq \tilde{C}_1 < \ell + \frac{(t+1)(u-\ell)}{T}\right), \quad t = 0, \ldots, T-1,$$

and $\mathbb{P}\left(\tilde{C}_1(T) = u\right) = \mathbb{P}(\tilde{C}_1 = u)$. Since $\tilde{C}_1(T)$ converges to $\tilde{C}_1$ in distribution by construction, $J(\tilde{D}_1(T)) \to J(\tilde{D}_1)$ as $T \to \infty$. Thus there exists a T_1 large enough such that for all $T > T_1$,

$$J(D_1^\star) > J(\tilde{D}_1(T)) + \epsilon/2.$$

Since $\tilde{D}_1(T)$ is supported on a finite number of points, this contradicts with our earlier conclusion that $D_1^\star$ is the worst case distribution for any distributions with a finite support. □

Finding the worst-case μ_2. Lemma 6 gives D_1, and (18) describes D_2. It remains to identify the value of $\mu_2 \in [\ell, u]$ that maximizes (15). Simple calculus gives that $\mu_2^\star = \sqrt{\ell u}$.

Proof of Theorem 3. Plugging in the values of D_1 and D_2 into (15) completes the proof. □

4.2. More than two sellers In the classical (linear, maximization) prophet inequality setting, Hill and Kertz (1981b) show that, for any instance in which $n > 2$, there exists an instance in which $n = 2$ for which the competitive ratio is at least as bad. Therefore, the exact competitive ratio for $n = 2$ gives an upper bound on the competitive ratio for any $n \geq 2$. We show that, the competitive ratio for the case of $n = 2$ is not the worst for all n. Specifically, we show that there exist values of ℓ, u for which the competitive ratio for $n = 3$ is strictly higher. Note that it is not the case that the competitive ratio is strictly higher for *all* ℓ, u, as when $\ell = u$, the competitive ratio is 1 for any value of n.

Lemma 7. *There exist ℓ, u such that $0 < \ell < u < \infty$ for which*

$$\mathrm{CR}(F,3) > \frac{1}{4}\left(\sqrt{\frac{u}{\ell}} + \sqrt{\frac{\ell}{u}}\right) + \frac{1}{2},$$

where $F = (\mathcal{D}, 2)$, and $\mathcal{D}$ is the set of distributions supported on $[\ell, u]$.

To prove Lemma 7, we consider a specific sequence of distributions.

Proof. Let $C_1 = (u+\ell)/2$ w.p. 1 and C_2 and C_3 be such that $C_i = u$ w.p. 0.5, $C_i = \ell$ w.p. 0.5 for $i \in \{2,3\}$. For $C = (C_1, C_2, C_3)$, we can calculate $\mathbb{E}[\mathrm{ALG}^{\mathrm{OPT}}(C,2)]$ using Lemma 4 and $\mathbb{E}[\mathrm{OPT}(C,2)]$ using Lemma 1:

$$\frac{\mathbb{E}[\mathrm{ALG}^{\mathrm{OPT}}(C,2)]}{\mathbb{E}[\mathrm{OPT}(C,2)]} = \frac{(u^2+\ell^2+6u\ell)(2u+\ell)(2\ell+u)}{u^4+\ell^4+38u^2\ell^2+16u^3\ell+16\ell^3u}.$$

For $u = 4$, $\ell = 1$, we have the ratio for $n = 3$ is approximately 1.1336; the competitive ratio for $n = 2$ with the given u, ℓ values is 1.125. □

As $n = 2$ is not the worst case, one might conjecture that the competitive ratio is monotone in n, as it is in the case where firms invest in production (Theorem 2). In Appendix C.2 we show empirically that this does not appear to be the case. the competitive ratio is not monotone in n. In fact, the number of sellers that exhibit the worst-case competitive ratio depends on the distribution domain. When the domains are bounded by $[1,2]$ and $[1,4]$, the worst-case competitive ratios transpire at three and four sellers respectively. Computing the exact competitive ratio as a function of the problem parameters (p, n, ℓ, u) is a challenging open problem.

5. Concluding Remarks We study a sequential procurement problem where there exists a marketplace with sellers arriving over time and investigate the impact of having in-house production capacity. We first consider a natural class of cost distributions. In this case, we show that it is possible to asymptotically achieve the optimal cost as the number of sellers grows. There is no advantage to in-house production capabilities, which corroborates the intuition that, with an abundance of sellers, there is less need for a ‘safety net’. Further, the decision maker can be asymptotically optimal using a multi-threshold policy that is much less complex than the optimal adaptive policy. More generally (with unrestricted distributions and a small number of sellers), the ability to produce the commodity oneself provides insurance from price spikes, i.e., the competitive ratio remains finite even if the upper bound on the cost is infinite when the decision maker has the ability to produce the commodity but increases with the upper bound when it does not.

We have made several simplifying assumptions that need to be addressed in future work. The first assumption is that the decision maker has unlimited production capacity. That is, it can produce all of the required commodity in-house. This is made with some loss of generality, although we note that the decision maker can compute ahead of time the amount that it needs to produce, so this can be incorporated into the decision whether or not to invest in production. Another simplifying assumption is that the production decision is binary, whereas in reality, different initial investments would lead to different production capabilities and costs. In this paper we abstracted both of these issues by normalizing the in-house cost function coefficient to 1.

The results of Sections 2 and 3 hold for general values of p (i.e., monomial cost functions), but the results of Section 4 hold only for $p=2$. It would be interesting to extend the results of Section 4 to other values of p, and indeed, to study the sequential procurement problem with more general costs than monomials; it is not clear how to generalize our results. Additionally, we have introduced a policy, Intervals, that is asymptotically optimal, but it would be interesting to understand if other policies converge at a faster rate than Intervals. Further, it would be interesting to understand the tradeoff between the rate of convergence to optimality and the complexity of the policy. For example, can a single-threshold policy match the convergence rate of Intervals?

Appendix A: Supporting Material for Section 2

LEMMA 1. *For any $c = (c_1, \dots, c_n)$ and $p > 1$,*

$$\mathrm{OPT}(c,p) = \left(\sum_{i=1}^{n} c_i^q\right)^{\frac{1}{q}},$$

where $q = -\frac{1}{p-1}$.

Proof. To obtain the prophet's cost, we consider the Lagrangian of optimization (2).

$$\sum_{i=1}^{n} c_i x_i^p + \lambda\left(1 - \sum_{i=1}^{n} x_i\right).$$

Taking the derivative yields

$$\lambda = p c_i x_i^{p-1}, \quad \text{i.e.,} \quad x_i = \left(\frac{\lambda}{p}\right)^{-q} c_i^q.$$

Substituting this into $\sum_i x_i = 1$, we have

$$\left(\frac{\lambda}{p}\right)^{-q} \sum_{i=1}^{n} c_i^q = 1.$$

Substituting $(\lambda/p)^{-q} = \left(\sum_{i=1}^{n} c_i^q\right)^{-1}$, we get

$$x_i = \frac{c_i^q}{\sum_{j=1}^{n} c_j^q}.$$

The prophet's cost is then

$$\mathrm{OPT} = \sum_{i=1}^{n} \frac{c_i^{pq+1}}{\left(\sum_{j=1}^{n} c_j^q\right)^p} = \frac{\sum_{i=1}^{n} c_i^q}{\left(\sum_{j=1}^{n} c_j^q\right)^p} = \left(\sum_{j=1}^{n} c_j^q\right)^{-(p-1)}.$$

□

Appendix B: Supporting Material for Section 3

LEMMA 2. *Let ℓ, μ and u be real numbers such that $0 < \ell \le \mu \le u < \infty$, $\mathcal{C}$ be the set of all random variables C whose distribution is supported on $[\ell, u]$, such that $\mathbb{E}[C] = \mu$, and f be a strictly concave function. Then a solution of*

$$\min_{C \in \mathcal{C}} \quad \mathbb{E}[f(C)] \tag{9}$$

if it exists, is

$$C = \begin{cases} u, & \text{with probability } p_u = \frac{\mu-\ell}{u-\ell}, \\ \ell, & \text{with probability } p_\ell = \frac{u-\mu}{u-\ell}. \end{cases}$$

Proof. Denote the long-shot distribution in Lemma 2 by $D^\star$ and let C be a random variable sampled from some distribution D supported on $[\ell, u]$ with mean μ. Then

$$\begin{aligned}
\mathbb{E}_D[f(C)] &= \int f(c)\,\mathrm{d}\,D(c) \\
&= \int f\left(\frac{u-x}{u-\ell}\ell + \frac{x-\ell}{u-\ell}u\right)\mathrm{d}\,D(c) \\
&\leq \int \left[\frac{u-c}{u-\ell}f(\ell) + \frac{c-\ell}{u-\ell}f(u)\right]\mathrm{d}\,D(c) \\
&= \frac{u-\mu}{u-\ell}f(\ell) + \frac{\mu-\ell}{u-\ell}f(u) \\
&= \mathbb{E}_{D^\star}[f(C)],
\end{aligned}$$

where the inequality follows from Jensen's inequality. □

Lemma 3. *Fix $p > 1$, and let $q = -\frac{1}{p-1}$. Let ℓ and u be real numbers such that $0 < \ell \leq u < \infty$ and C be a n-dimensional random vector with independent elements each of which has its distribution supported on $[\ell, u]$. Then*

$$\frac{\mathbb{E}\left[[\mathrm{ALG}^{\mathrm{NA}}(C,p)\right]}{\mathbb{E}[\mathrm{OPT}(C,p)]} = \frac{M_q(\mathbb{E}[C])}{\mathbb{E}[M_q(C)]}.$$

Proof of Lemma 3. From Lemma 1, $\mathrm{OPT}(c,p) = (\sum_{i=1}^n c_i^q)^{\frac{1}{q}}$. As $M_q(c) = \left(\frac{1}{n}\sum_{i=1}^n c_i^q\right)^{\frac{1}{q}}$, the expected prophet cost is simply $\mathbb{E}[\mathrm{OPT}(C,p)] = n^{\frac{1}{q}}\mathbb{E}\left[M_q(c)\right]$. Using similar reasoning, it is easy to show that the cost of the optimal non-adaptive algorithm is

$$\mathbb{E}[\mathrm{ALG}^{\mathrm{NA}}(C,p) = n^{\frac{1}{q}} M_q(\mathbb{E}[C]),$$

as the optimization (8) is equivalent to (2) with $\mathbb{E}[C_i]$ in place of C_i. □

B.1. Proof of Theorem 2 (lower bound). It is straightforward to verify that, similarly to the classical prophet inequality, the optimal policy can be characterized as a dynamic program. We describe such an algorithm, denoted DP, and show that the cost-to-go functions of DP share the same functional form as the cost functions.

To start, let us first define the *cost-to-go functions* J_i on $s_i \in [0,1]$ as follows:

$$\begin{aligned}
J_n(s_n) &= c_n(1-s_n), \\
J_i(s_i) &= \min_{0\leq x_i\leq 1-s_i} c_i(x_i) + \mathbb{E}\left[J_{i+1}(s_i+x_i)\right], \quad i=1,\ldots,n-1.
\end{aligned} \tag{21}$$

The optimal causal policy is a mapping from the state (i.e., the accumulative amount of commodity that was obtained up to time step i) $s_i \in [0,1]$ to the action $x_i \in [0,1]$ that takes the form

$$\begin{aligned}
x_n(s_n) &= 1-s_n, \\
x_i(s_i) &\in \operatorname*{arg\,min}_{0\leq x_i\leq 1-s_i} c_i(x_i) + \mathbb{E}\left[J_{i+1}(s_i+x_i)\right], \quad i=1,\ldots,n-1.
\end{aligned} \tag{22}$$

The optimality of the dynamic programming recursion above is a classical result in the literature of stochastic control (cf. Proposition 1.3.1 of Bertsekas (1995)) which relies on the simple intuition that if the truncated policy $\{\pi_i, \pi_{i+1}, \ldots, \pi_n\}$ does not implement optimal solutions of (22), then it would be possible to reduce the cost for the "tail problem" starting from stage i with state s_i further by switching to an optimal policy. In Lemma 8, we show that the minimization in the DP recursion can be solved analytically and therefore DP is implementable in polynomial time.

LEMMA 8. *Let D be a distribution supported on $[\ell, u]$, where $0 < \ell \le u < \infty$ and $p > 1$. An optimal algorithm for the minimization prophet inequality problem for (D, p) has cost-to-go functions*

$$J_i(s_i) = b_i(1-s_i)^p, \quad i = 1, \ldots, n, \tag{23}$$

where b_i is defined recursively as

$$b_n = C_n, \tag{24}$$

$$b_i = \frac{c_i \mathbb{E}[b_{i+1}]}{\left(c_i^{1/(p-1)} + (\mathbb{E}[b_{i+1}])^{1/p-1}\right)^{p-1}}, \quad i = 1, \ldots, n-1. \tag{25}$$

The corresponding allocation is given by

$$x_i = \frac{(\mathbb{E}[b_{i+1}])^{1/(p-1)}}{c_i^{1/(p-1)} + (\mathbb{E}[b_{i+1}])^{1/(p-1)}}(1-s_i), \quad i = 1, \ldots, n. \tag{26}$$

Proof. We prove that the cost-to-go functions are of the form (23) using backward induction. The base case, $J_n(s_n) = b_n(1-s_n)^p = c_n(1-s_n)^p$ holds by the definition of the cost-to-go function. Suppose that $J_{i+1}(s_{i+1}) = b_{i+1}(1-s_{i+1})^p$. Then we have

$$\begin{aligned} J_i(s_i) &= \min_{x_i \in [0,1]} c_i x_i^p + \mathbb{E}[J_{i+1}(s_i + x_i)], \\ &= \min_{x_i \in [0,1]} c_i x_i^p + \mathbb{E}[b_{i+1}(1-s_i-x_i)^p]. \end{aligned}$$

The first order optimality condition of the optimization above is

$$c_i x_i^{p-1} - (1-s_i-x_i)^{p-1}\mathbb{E}[b_{i+1}] = \left(c_i^{1/(p-1)} x_i - (\mathbb{E}[b_{i+1}])^{1/(p-1)}(1-s_i-x_i)\right) f(s_i, x_i) = 0,$$

where $f(s_i, x_i)$ is a function of s_i and x_i satisfying $f(s_i, x_i) > 0$ for $x_i \ge 0$, $s_i \ge 0$ and $s_i + x_i > 0$. The unconstrained minimizer is

$$x_i = \frac{(\mathbb{E}[b_{i+1}])^{1/(p-1)}}{c_i^{1/(p-1)} + (\mathbb{E}[b_{i+1})^{1/(p-1)}]}(1-s_i),$$

which satisfies $0 \le x_i \le (1-s_i) \le 1$ for $s_i \in [0,1]$. It follows that

$$\begin{aligned} J_i(s_i) &= \frac{c_i(\mathbb{E}[b_{i+1}])^{p/(p-1)}}{\left(c_i^{1/(p-1)} + (\mathbb{E}[b_{i+1}])^{1/(p-1)}\right)^p}(1-s_i)^p + \frac{c_i^{p/(p-1)}\mathbb{E}[b_{i+1}]}{\left(c_i^{1/(p-1)} + (\mathbb{E}[b_{i+1}])^{1/(p-1)}\right)^p}(1-s_i)^p \\ &= \frac{c_i \mathbb{E}[b_{i+1}]}{\left(c_i^{1/(p-1)} + (\mathbb{E}[b_{i+1}])^{1/p-1}\right)^{p-1}}(1-s_i)^p. \end{aligned}$$

Thus, the claim holds for any i. □

Proof of Theorem 2 (lower bound).

Consider the optimal algorithm, DP, on the particular distribution where $C_i = \ell$ with probability $1-\epsilon$, and $C_i = u \to \infty$ with probability $\epsilon > 0$ for all $i \in \{1,\ldots,n\}$. From the DP recursion and by Lemma 8, we have $\mathbb{E}[b_n] = \mathbb{E}[c_n] \to \infty$ as $u \to \infty$,

$$\mathbb{E}[b_{n-1}] = \mathbb{E}\left[\frac{1}{\left(c_{n-1}^{-1/(p-1)} + (\mathbb{E}[b_n])^{-1/(p-1)}\right)^{p-1}}\right] = \mathbb{E}[c_{n-1}] \to \infty, \text{ as } u \to \infty,$$

and thus $\mathbb{E}[b_i] \to \infty$ as $u \to \infty$ inductively for all i. It follows that $\mathbb{E}[\mathrm{ALG}^{\mathrm{DP}}(C,p)] = 1$. Following the same lines of arguments as in the proof of the upper bound, we can calculate $\mathbb{E}[\mathrm{OPT}(C,p)]$ for $\epsilon \to 0$, and obtain the competitive ratio for DP with the given distribution as

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

□

Appendix C: Supporting Material for Section 4

C.1. Proof of Lemma 5

LEMMA 4. *Consider the minimization prophet inequality problem with input $(D,2)$, $D = \{D_1, D_2, \ldots, D_n\}$. For any realization $c_1,\ldots,c_n$ of the random variables $C_1,\ldots,C_n$, the prophet's cost is* $\mathrm{OPT} = \frac{H(c_1,\ldots,c_n)}{n}$.

The decision maker's expected cost, i.e., the cost achieved by the optimal algorithm DP, *is*

$$\mathbb{E}[\mathrm{ALG}^{\mathrm{DP}}(C,p)] = \mathbb{E}[b_1],$$

and

$$x_i = \frac{b_i}{c_i}\prod_{j=1}^{i-1}\left(1 - \frac{b_j}{c_j}\right), \quad i = 1,\ldots,n, \tag{13}$$

where $b_n = C_n$ and

$$b_i = \frac{H(c_i, \mathbb{E}[b_{i+1}])}{2}, \quad i = 1,\ldots,n-1.$$

We prove a more general result; Lemma 5 is a special case.

Proof of Lemma 4. Concerning the first part of Lemma 4 (the prophet's cost and allocations): the Lagrangian for (2) with $p = 2$ can be written as (keeping the positivity constraints):

$$\frac{1}{2}\sum_{i=1}^{n} c_i x_i^2 - \lambda(x_1 + \ldots + x_n - 1).$$

Differentiating gives that, for all i, $x_i = \lambda / c_i$ with $\sum_{i=1}^n x_i = 1$. Substituting gives $\lambda = \left(\sum_{i=1}^n c_i^{-1}\right)^{-1}$, and rearranging then completes the proof.

For the optimal online algorithm's cost and allocations: By Lemma 8, we have

$$x_i = \frac{\mathbb{E}[b_{i+1}]}{c_i + \mathbb{E}[b_{i+1}]}(1 - s_i), \quad i = 1, \ldots, n,$$

with $b_n = C_n$ and

$$b_i = \frac{c_i \mathbb{E}[b_{i+1}]}{c_i + \mathbb{E}[b_{i+1}]} = \frac{H(c_i, \mathbb{E}[b_{i+1}])}{2}, \quad i = 1, \ldots, n-1.$$

It follows that

$$x_i = \frac{b_i}{c_i}(1 - s_i) = \frac{b_i}{c_i}\left(1 - \sum_{j=1}^{i-1} x_j\right).$$

Expression (13) then holds by induction: The base case holds as (13) is equivalent to the expression above for x_1. Suppose (13) holds for x_i, then

$$1 - s_i = \prod_{j=1}^{i-1}\left(1 - \frac{b_j}{c_j}\right),$$

and

$$1 - s_{i+1} = 1 - s_i - x_i = \prod_{j=1}^{i-1}\left(1 - \frac{b_j}{c_j}\right) - \frac{b_i}{c_i}\prod_{j=1}^{i-1}\left(1 - \frac{b_j}{c_j}\right) = \prod_{j=1}^{i}\left(1 - \frac{b_j}{c_j}\right).$$

Thus we have

$$x_{i+1} = \frac{b_{i+1}}{c_{i+1}}(1 - s_{i+1}) = \frac{b_{i+1}}{c_{i+1}}\prod_{j=1}^{i}\left(1 - \frac{b_j}{c_j}\right).$$

□

C.2. Non-monotonicity in n In this section, we explore the competitive ratio of the optimal online algorithm in the case where the firm cannot produce the commodity. The numerics highlight the complex behavior of this setting, illustrating the difficulty in precisely characterizing the competitive ratio in this case and showing a variety of interesting non-monotonicities.

We partition the experiments into two sections. First we consider a general class of discrete distributions over $[\ell, u]$ and then explore specific cases with long-shot and point mass distributions.

General discrete distributions To explore the competitive ratio we create examples by discretizing the interval $[\ell, u]$ and considering distributions $D_1, \ldots, D_n$ supported on S points in $\{\ell, \ell + \Delta, \ldots, u\}$ where $\Delta = (u - \ell)/(S - 1)$. The problem of identifying the worst case distribution (with the given support) that maximizes the competitive ratio is a non-linear and non-convex program over the probability mass functions of $D_1, \ldots, D_n$. We solve this non-linear program numerically using an interior point method for $n \in \{2, \ldots, 6\}$ and $S \in \{3, \ldots, 6\}$. While these solutions are local, they are likely also global, as we obtain consistent results for a number of runs with random initializations.

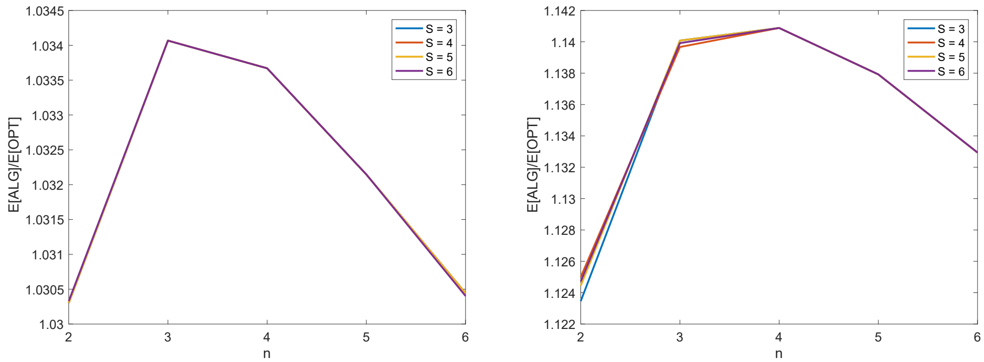

(a) $(\ell, u) = (1, 2)$ (b) $(\ell, u) = (1, 4)$

FIGURE 1. Competitive ratios under the worst-case discretized distributions.

We restrict to small values of n and S in these experiments due to (i) the complexity of evaluating the expected optimal prophet's cost is $O(S^n)$, (ii) the non-convex and nonlinear nature of the optimization problem at hand, and (iii) the fact that no significant change in the results has been observed when we increase the S value.

Fig. 1 presents the results of our experiments. The competitive ratio $\mathbb{E}[\text{ALG}]/\mathbb{E}[\text{OPT}]$ for $(\ell, u) \in \{(1,2),(1,4)\}$ is shown in Fig. 1 over the range of n values considered. The most important observation is that the competitive ratio is not monotonic – we observe that the competitive ratio as a function of n is unimodal; the peak is at value of n that increases with u.

Fig. 2 depicts the worst-case distributions for $n \in \{3,\dots,6\}$, $u = 4$ (structurally similar results are obtained for $u = 2$), and $S = 6$. In each panel of Fig. 2, we show the probability mass function of D_i for $i = 1,\dots,n$. Consistent with the intuition from the results in the paper, we see that across for any n value, the resulting D_i, $i = 1,\dots,n$ is either point-mass or long-shot. This motivates us to focus on these distributions in the following section.

Point-mass/long-shot distributions Motivated by the results above, we focus on long-shot and point mass distributions in this section. This allows us to explore a wider variety of settings with considerably less computational effort. Specifically, we focus on $S = 3$ as D_i is supported on $\{\ell, m, u\}$ for some intermediate point $m \in (\ell, u)$. We include the locations of the intermediate points as optimization variables and re-run the simulations for larger values of n.

The resulting competitive ratios are depicted in Fig. 3, which confirms the observations that we made for Fig. 1 with a larger range of n values.

Since the worst case distributions we consider are either point-mass or long-shot, we can summarize them using an indicator (point-mass or long-shot) and the mean $\mathbb{E}[C_i]$'s. Fig. 4 shows the mean values of the worst case distributions for $(\ell, u) = (1,4)$. Structurally similar results are obtained

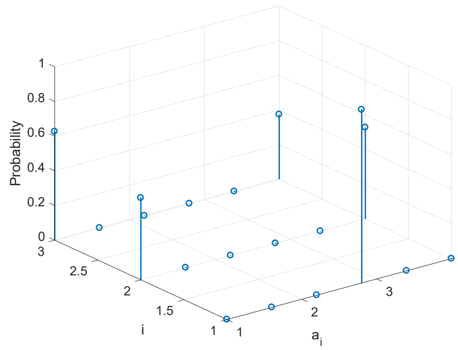

(a) $(\ell, u) = (1,4), n = 3$

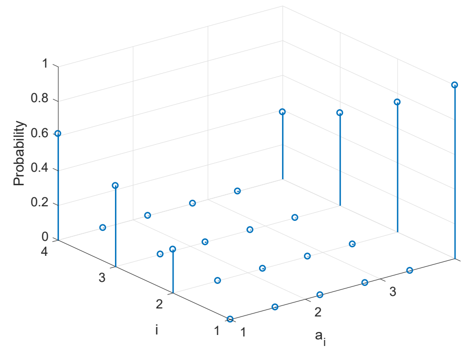

(b) $(\ell, u) = (1,4), n = 4$

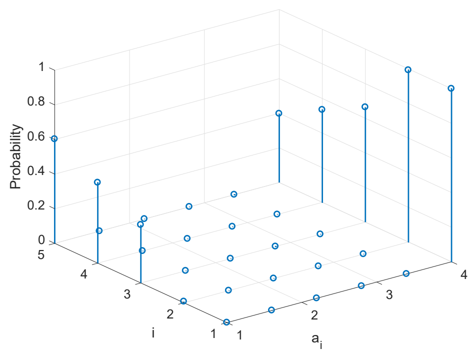

(c) $(\ell, u) = (1,4), n = 5$

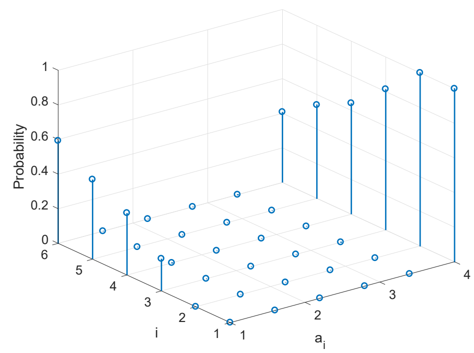

(d) $(\ell, u) = (1,4), n = 6$

FIGURE 2. Probability mass functions of the worst case discretized distributions

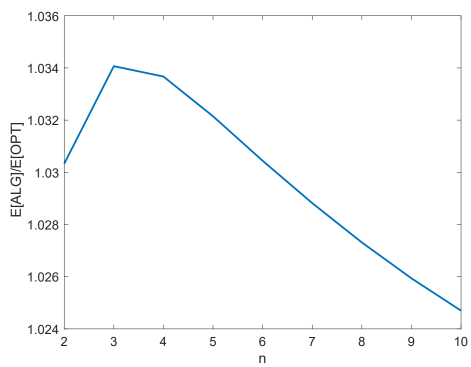

(a) $(\ell, u) = (1,2)$

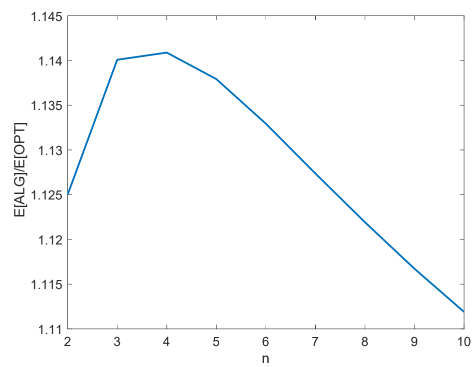

(b) $(\ell, u) = (1,4)$

FIGURE 3. Competitive ratio under the worst case long-shot/point mass distributions.

for $(\ell, u) = (1,2)$. Many observations can be made regarding Fig. 4. First, note that the structure of the sequence of worst case distributions for any n is always such that there exists a $k(n)$ for which D_i is point-mass for $i \le k(n)$ and D_i is long-shot for $k(n) < i \le n$. Second, for each fixed n, the sequence of means for which D_i's are long-shot, i.e. $\{\mathbb{E}[C_i] : i > k(n)\}$, is a decreasing sequence. Finally, the transition point $k(n)$ is nondecreasing in n. Formally establishing these properties may pave the way to obtaining exact quadratic prophet inequalities for general $n > 2$.

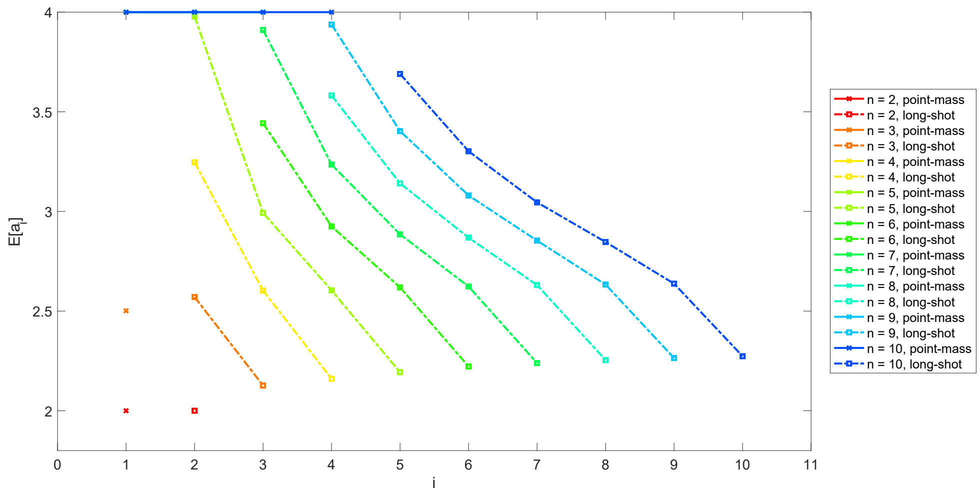

Figure 4. Mean values of the worst case long-shot/point mass distributions for $(\ell, u) = (1, 4)$. For each n values, there are two line segments, the first (solid) line segment corresponds to i values for which D_i's are point-mass. When there is only one i has a point-mass distribution, the line-segments becomes a point. The second (dashed) line segment corresponds to i values for which D_i's are long-shot.

References

Agrawal S, Wang Z, Ye Y (2014) A dynamic near-optimal algorithm for online linear programming. *Operations Research* 62(4):876–890.

Alaei S (2014) Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM J. Comput.* 43(2):930–972.

Alaei S, Hajiaghayi M, Liaghat V (2012) Online prophet-inequality matching with applications to ad allocation. *ACM Conference on Electronic Commerce, EC '12*, 18–35.

Alon N, Spencer J (2008) *The Probabilistic Method* (John Wiley), 3rd edition.

Andrew L, Barman S, Ligett K, Lin M, Meyerson A, Roytman A, Wierman A (2013) A tale of two metrics: Simultaneous bounds on competitiveness and regret. *Conference on Learning Theory*, 741–763.

Arya A, Mittendorf B, Sappington DEM (2008) The make-or-buy decision in the presence of a rival: Strategic outsourcing to a common supplier. *Management Science* 54(10):1747–1758, URL http://dx.doi.org/10.1287/mnsc.1080.0896.

Azar Y, Buchbinder N, Chan TH, Chen S, Cohen IR, Gupta A, Huang Z, Kang N, Nagarajan V, Naor J, Panigrahi D (2016) Online algorithms for covering and packing problems with convex objectives. *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS*, 148–157.

Bertsekas DP (1995) *Dynamic programming and optimal control*, volume 1(2) (Athena Scientific Belmont, MA).

Bose S, Cai DWH, Low SH, Wierman A (2014) The role of a market maker in networked Cournot competition. *53rd IEEE Conference on Decision and Control, CDC 2014, Los Angeles, CA, USA, December 15-17, 2014*, 4479–4484.

Buchbinder N, Naor JS (2009) The design of competitive online algorithms via a primal–dual approach. *Foundations and Trends® in Theoretical Computer Science* 3(2–3):93–263.

Chaisiri S, Lee BS, Niyato D (2012) Optimization of resource provisioning cost in cloud computing. *IEEE Transactions on Services Computing* 5(2):164–177.

Chawla S, Hartline JD, Malec DL, Sivan B (2010) Multi-parameter mechanism design and sequential posted pricing. *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC '10*, 311–320.

Chen J, Guo Z (2014) Strategic sourcing in the presence of uncertain supply and retail competition. *Production and Operations Management* 23(10):1748–1760, URL http://dx.doi.org/10.1111/poms.12078.

Chen N, Agarwal A, Wierman A, Barman S, Andrew LLH (2015) Online convex optimization using predictions. *Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, Portland, OR, USA, June 15-19, 2015*, 191–204.

Chen N, Comden J, Liu Z, Gandhi A, Wierman A (2016) Using predictions in online optimization: Looking forward with an eye on the past. *Proceedings of the 2016 ACM SIGMETRICS International Conference* *on Measurement and Modeling of Computer Science, Antibes Juan-Les-Pins, France, June 14-18, 2016*, 193–206.

Correa J, Foncea P, Hoeksma R, Oosterwijk T, Vredeveld T (2019) Recent developments in prophet inequalities. *SIGecom Exch.* 17(1):61–70, URL http://dx.doi.org/10.1145/3331033.3331039.

Correa J, Foncea P, Hoeksma R, Oosterwijk T, Vredeveld T (2020) Posted price mechanisms and optimal threshold strategies for random arrivals. *Mathematics of Operations Research* 0(0):null, URL http://dx.doi.org/10.1287/moor.2020.1105.

Cover TM, Thomas JA (2012) *Elements of information theory* (John Wiley & Sons).

Delage E, Ye Y (2010) Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research* 58(3):595–612.

Disser Y, Fearnley J, Gairing M, Göbel O, Klimm M, Schmand D, Skopalik A, Tönnis A (2020) Hiring secretaries over time: The benefit of concurrent employment. *Mathematics of Operations Research* 45(1):323–352, URL http://dx.doi.org/10.1287/moor.2019.0993.

Dudley R (2002) *Real Analysis and Probability*, volume 74 (Cambridge University Press).

Düetting P, Feldman M, Kesselheim T, Lucier B (2017) Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*, 540–551 (IEEE).

Esfandiari H, Hajiaghayi M, Liaghat V, Monemizadeh M (2015) Prophet secretary. *Algorithms - ESA 2015 - 23rd Annual European Symposium, '15, Proceedings*, 496–508.

Grünwald PD, Dawid AP (2004) Game theory, maximum entropy, minimum discrepancy and robust bayesian decision theory. *Ann. Statist.* 32(4):1367–1433.

Hajiaghayi MT, Kleinberg RD, Sandholm T (2007) Automated online mechanism design and prophet inequalities. *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, '07*, 58–65.

Hazan E (2016) Introduction to online convex optimization. *Foundations and Trends® in Optimization* 2(3-4):157–325.

Hill TP (1983) Prophet inequalities and order selection in optimal stopping problems. *Proceedings of the American Mathematical Society* 88(1):131–137.

Hill TP, Kertz RP (1981a) Additive comparisons of stop rule and supremum expectations of uniformly bounded independent random variables. *Proceedings of the American Mathematical Society* 83(3):582–585.

Hill TP, Kertz RP (1981b) Ratio comparisons of supremum and stop rule expectations. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* 56(2):283–285, ISSN 1432-2064.

Hill TP, Kertz RP (1992) A survey of prophet inequalities in optimal stopping theory. *Contemporary Mathematics* 125:191–207.

Kleinberg R, Weinberg SM (2012) Matroid prophet inequalities. *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC '12*, 123–136.

Krengel U, Sucheston L (1977) Semiamarts and finite values. *Bull. Amer. Math. Soc.* 83(4):745–747.

Krengel U, Sucheston L (1978) On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, volume 4 of *Adv. Probab. Related Topics*, 197–266 (Dekker, New York).

Low SH (2014) Convex relaxation of optimal power flow, part i: Formulations and equivalence. *IEEE Transactions on Control of Network Systems* 1(1):15–27.

McShane EJ (1983) *Unified Integration* (Academic Press).

Niu B, Li J, Zhang J, Cheng HK, Tan YR (2019) Strategic analysis of dual sourcing and dual channel with an unreliable alternative supplier. *Production and Operations Management* 28(3):570–587, URL http://dx.doi.org/10.1111/poms.12938.

Rajagopal R, Bitar E, Varaiya P, Wu F (2013) Risk-limiting dispatch for integrating renewable power. *International Journal of Electrical Power & Energy Systems* 44(1):615 – 628, ISSN 0142-0615.

Rubinstein A, Singla S (2017) Combinatorial prophet inequalities. Klein PN, ed., *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, 1671–1687 (SIAM), URL http://dx.doi.org/10.1137/1.9781611974782.110.

Salop SC, Scheffman DT (1987) Cost-raising strategies. *The Journal of Industrial Economics* 36(1):19–34, ISSN 00221821, 14676451, URL http://www.jstor.org/stable/2098594.

Samuel-Cahn E (1984) Comparison of threshold stop rules and maximum for independent nonnegative random variables. *The Annals of Probability* 12(4):1213–1216, ISSN 00911798.

Sethi S, Yan H, Yan JH, Zhang H (2005) An analysis of staged purchases in deregulated time-sequential electricity markets. *Journal of Industrial and Management Optimization* 1(4):443–463.

Shalev-Shwartz S (2012) Online learning and online convex optimization. *Foundations and Trends® in Machine Learning* 4(2):107–194.

Varaiya PP, Wu FF, Bialek JW (2011) Smart operation of smart grid: Risk-limiting dispatch. *Proceedings of the IEEE* 99(1):40–57.

Wood AJ, Wollenberg BF (2012) *Power generation, operation, and control* (John Wiley & Sons).