Delegated Search Approximates Efficient Search

JON KLEINBERG, Cornell University, USA
ROBERT KLEINBERG, Cornell University, USA

There are many settings in which a principal performs a task by *delegating* it to an agent, who searches over possible solutions and proposes one to the principal. This describes many aspects of the workflow within organizations, as well as many of the activities undertaken by regulatory bodies, who often obtain relevant information from the parties being regulated through a process of delegation. A fundamental tension underlying delegation is the fact that the agent's interests will typically differ – potentially significantly – from the interests of the principal, and as a result the agent may propose solutions based on their own incentives that are inefficient for the principal. A basic problem, therefore, is to design mechanisms by which the principal can constrain the set of proposals they are willing to accept from the agent, to ensure a certain level of quality for the principal from the proposed solution.

In this work, we investigate how much the principal loses – quantitatively, in terms of the objective they are trying to optimize – when they delegate to an agent. We develop a methodology for bounding this loss of efficiency, and show that in a very general model of delegation, there is a family of mechanisms achieving a universal bound on the ratio between the quality of the solution obtained through delegation and the quality the principal could obtain in an idealized benchmark where they searched for a solution themself. Moreover, it is possible to achieve such bounds through mechanisms with a natural threshold structure, which are thus structurally simpler than the optimal mechanisms typically considered in the literature on delegation. At the heart of our framework is an unexpected connection between delegation and the analysis of *prophet inequalities*, which we leverage to provide bounds on the behavior of our delegation mechanisms.

CCS Concepts: • **Theory of computation → Design and analysis of algorithms**; • **Applied computing → Economics**;

Additional Key Words and Phrases: Delegated search, prophet inequalities.

1 INTRODUCTION

There are many settings in which a decision-maker is faced with a difficult problem that they cannot solve on their own, and so they instead approach it in two steps: they first *delegate* the search for possible solutions to an *agent* who is able to invest more time in the process, and then they evaluate the solution(s) that the agent proposes. One concrete example arises in organizations or firms, where the management may delegate the search for solutions to a division that reports to them, ultimately making a decision on the solution that is proposed by the division [2, 6, 18]. A second example arises in regulation, where a governmental agency needs to decide whether there is a way to structure a proposed corporate merger in a way that is compatible with regulatory guidelines; the companies seeking to merge study possible ways of structuring the merger, propose one or more to the regulator, and seek the regulator's approval. In this way, the search over structures for the merger has implicitly been delegated from the regulator to the companies, with the regulator

Authors' addresses: Jon Kleinberg, Cornell University, Ithaca, NY, 14853, USA; Robert Kleinberg, Cornell University, Ithaca, NY, 14853, USA.

ACM EC'18, June 18–22, 2018, Ithaca, NY, USA. ACM ISBN 978-1-4503-5829-3/18/06…$15.00
https://doi.org/10.1145/3219166.3219205

retaining decision-making control over the the proposed solution [11, 20]. A similar scenario could be described for regulation in other settings, where a company may be searching over possible solutions that comply with environmental law, employment law, or other guidelines.

The interesting tension in all these situations is that the decision-maker who delegates the task (henceforth referred to as the "principal") has a particular objective function that they are seeking to optimize; but the agent who actually performs the task might have interests that are not directly aligned with the principal's. For example, in the regulatory context, the regulator (acting as principal) may reasonably suspect that a merger proposed by a set of companies (acting as the agent) will be structured in a way that strongly benefits the companies, even if other feasible structures would have been better for the market or for society as a whole. Similarly, a group within an organization tasked with solving a problem may well preferentially search for a solution that benefits them in particular. Given this natural set of incentives, how should the principal structure the delegation to the agent so as to ensure that the solution the agent proposes performs well under the principal's own objective function?

A rich literature has developed in economics around the formalization and analysis of delegation, focusing on this tension between the conflicting objectives of the principal and the agent; see Holmstrom [13, 14] for influential early research, and [3–6] for recent work. A dominant theme in this line of work is that the principal does not offer monetary compensation to the agent as a way of favoring certain proposed solutions over others (though see Krishna and Morgan [17]); this is consistent with the motivating applications, in which for example regulators in many contexts can accept or reject proposals from companies, but cannot selectively offer varying amounts of compensation to these companies based on the content of the proposal. This lack of monetary transfers between the parties imparts a fundamental structure to the problem, in which the principal can simply define a mechanism implicitly specifying the subset of all "eligible" proposals that they are willing to accept, and the agent is then incentivized to search for solutions that are good for them but that also lie in this eligible set. A long line of work has gone into determining the structure of eligible sets that produce optimal mechanisms for the principal, yielding constructions that are often quite intricate [3, 6, 19].

The Present Work. Given how broadly delegation is used across a range of contexts, it is interesting to consider how precarious a process it is — the principal is ceding control of their search problem to an agent whose interests might be completely misaligned with their own, and the only leverage the principal has is to accept or reject the solution that is eventually proposed. How much does the principal give up — quantitatively, measured in terms of the objective they are trying to optimize — when they delegate to an agent? Is there some robust, intrinsic reason why things don't turn out as badly as we might fear? And how do the answers to these questions depend on what the principal is actually able to observe about the agent's solution — including how much effort the agent put into finding the solution, and how good it is not just for the principal but also for the agent?

In their most natural formulation, these are inherently comparative questions, since they seek to relate the solution quality obtained through delegation to the solution quality in an alternate, ideal setting where delegation was not necessary. As such, they address an issue fundamentally distinct from the primary focus of the existing literature on delegation, which as noted above has centered on characterizing mechanisms that produce optimal delegation for the principal, without this type of comparative evaluation.

There is a natural benchmark to use for our comparison: we could measure the quality of the outcome under delegation versus the quality of the solution that the principal could obtain were they to perform the search task themself, investing the same level of effort in the search that the agent does. Now, there are many settings where it may be too costly or otherwise infeasible for the principal to actually conduct this search, but this benchmark nonetheless provides a conceptual reference point to make clear how much payoff to the principal is lost through delegation. In this sense, it plays the role of an optimal point of comparison, much like the role of the intractable optimum in an approximation algorithm or the societal optimum in a price-of-anarchy analysis.

In this paper we develop a methodology to bound the performance of delegated search, relative to the benchmark in which the principal searches for a solution on their own. Our methodology builds on a set of links that we identify between bounds on delegated search and the analysis of some fundamental models of decision-making under uncertainty — in particular, a surprisingly strong connection between delegated search and bounds on *prophet inequalities*. The connections between these formalisms turn out be quite natural and useful, but to our knowledge they have not been previously identified in either of the literatures on delegated search or on prophet inequalities. This connection not only provides bounds on the quality of delegated solutions relative to an ideal benchmark; it also shows that strong bounds can be obtained using eligible sets that are structurally very simple — in a number of cases defined simply by a carefully chosen threshold rule — and hence in contrast with the complex constructions associated with optimal mechanisms.

Overview of Models

A Distributional Model. We begin by describing the models in which we perform our analysis. Our main model, which is essentially the one considered in Armstrong and Vickers [6], has the principal and the agent agree that the agent will consider n candidate solutions and propose one to the principal; n thus represents the level of effort that the agent commits to the problem. The principal will only see the solution that is proposed, not the other $n-1$ that the agent also considers.

What does it means for the agent to consider a candidate solution? We assume that the solutions belong to an abstract space Ω with a probability measure on it, and the agent's search for a solution consists of performing n independent and identically distributed draws from Ω, resulting in a set of candidate solutions $\omega_1, \omega_2, \ldots, \omega_n \in \Omega$.[1] Each solution ω_i drawn by the agent has a quality for the principal, denoted $x(\omega_i)$, and a possibly different quality for the agent, denoted $y(\omega_i)$. The agent selects one of its candidate solutions, say ω_i, to present to the principal. (Below, we will discuss the contrast between the model in which the principal is able to determine both $x(\omega_i)$ and $y(\omega_i)$ — the value to both the agent and themself — for the proposed solution, and the model in which the principal is only able to determine $x(\omega_i)$.)

Now, if the principal imposed no constraint on the agent's behavior, then the agent would simply choose the solution ω_i that maximizes $y(\omega_i)$, and the principal would receive whatever corresponding $x(\omega_i)$ value resulted from this choice. To improve on this, the principal could specify at the outset that they will only accept ω values that satisfy some predicate on $x(\omega_i)$ and (in the case that they can determine it) $y(\omega_i)$; we will refer to the set of all ω satisfying the principal's predicate as the *eligible set* of solutions. It is thus in the agent's interest to propose a solution belonging to the eligible set; we ask whether one can design eligible sets that provide provable bounds on the expected quality of the solution to the principal, relative to the scenario in which the principal simply were to draw n times from Ω and select the sampled solution ω with maximum $x(\omega)$.

A Binary Model. In our first, distributional model, the agent draws a set of candidate solutions $\omega_1, \ldots, \omega_n$ that the principal cannot observe, and then must choose one to present to the principal. This models a setting in which the agent explores a design space and cannot fully anticipate which

[1]Later we will also consider the case in which different draws by the agent can come from different probability measures on Ω; for example, this can model the case in which the agent is a group of n employees in an organization, and the i^{th} solution is drawn by the i^{th} agent, who may have a different distribution over solutions from the j^{th} agent.

ω_i it will encounter until it begins this exploration. But we can also imagine settings where the principal and agent both know that the set of potential options comes from a large discrete set $\omega_1, \omega_2, \ldots, \omega_m$, and the only question is which of these options is actually feasible to implement. For example, there may be m standard ways of structuring a merger in a given industry, and the only question is which are possible for the companies in question.

We model this version with publicly known binary options as follows. There is a set of options $\omega_1, \omega_2, \ldots, \omega_m$, and for each i, there is a known probability $p_i > 0$ such that option ω_i is *feasible* with probability p_i, and *infeasible* otherwise. If option ω_i turns out to be feasible, then it produces a known payoff of $x(\omega_i)$ for the principal and a known payoff $y(\omega_i)$ for the agent; if it turns out to be infeasible, then it produces a payoff of 0 for both. The only way to evaluate the feasibility of option ω_i is to pay a cost of c_i to investigate it.

The principal delegates to the agent the task of proposing a feasible option, which the principal can either accept or reject. The principal will not be able to see which options the agent decides to pay to evaluate as part of this task, but again the principal can specify a predicate defining the eligible subset of ω_i that they will accept. Subject to this constraint, the agent then must decide how to evaluate options in a way that maximizes its own benefit $y(\omega_i)$ from the option ω_i it proposes, minus the evaluation cost. Here too we evaluate the principal's payoff relative to the scenario in which they performed the evaluation of options themself. We also consider an extension of this model in which there is a budget of $n < m$ on the number of options that the agent can evaluate — a constraint analogous to the bound on the number of samples the agent can evaluate in our first distributional model.

Overview of Results

We begin by showing that for an arbitrary instance of the distributional model, there is a mechanism the principal can specify to the agent so that the principal's expected payoff from delegation is within a factor of 2 of the expected payoff they'd receive were they able to search for the solution by themself. (If the principal were to search by themself, they would examine n candidate solutions and choose the one that was best for them.) This mechanism only requires knowledge of the principal's $x(\omega_i)$ values, not the agent's $y(\omega_i)$ values, and it has a very simple structure: depending on the distribution of values, it can be written as a threshold rule with either a weak threshold, in which the principal only accepts proposals ω for which $x(\omega) \geq \theta$ for some θ, or a strict threshold, in which the principal only accepts proposals ω for which $x(\omega) > \theta$ for some θ. In the case when $x(\omega)$ and $y(\omega)$ are distributed independently with no point masses, the factor of 2 in this bound can be improved to $e/(e-1)$.

There are several things worth remarking on about this result. First, the fact that arbitrary instances of the problem have mechanisms providing provable guarantees of this form suggests a qualitative argument for the robustness of delegation: no matter how misaligned the agent's interests are, the principal can ensure an absolute bound on how much is lost in the quality of the solution. Second, the mechanism that achieves this bound is very simple and detail-free, consisting of just a (weak or strict) threshold on the quality of the solution for the principal. And third, the mechanism requires knowledge of n (the number of samples drawn by the agents) but not the values $y(\omega)$. In this sense, it suggests that it is more important for the principal to know how much effort the agent has spent on the search (via n) than to know how good the proposed solution is for the agent (via $y(\omega)$).

A connection to prophet inequalities. These results on threshold mechanisms and their guarantees follow from a general result at the heart of our analysis — a close connection between bounds for delegated search and *prophet inequalities*. Prophet inequalities are guarantees for the following type of decision under uncertainty: we see a sequence of values $s_1, s_2, \ldots, s_n$ in order, with s_i drawn from a distribution S_i, and when we see the value s_i we must irrevocably decide whether to stop and accept s_i, or continue (in the hope of finding a better value in the future). Research on prophet inequalities has established the non-trivial fact that it is possible to design rules whose expected payoff comes within a constant fraction of the maximum achievable by a decision-maker who could see all the realized values $s_1, s_2, \ldots, s_n$ in advance.

Prophet inequalities tend to be established by designing carefully constructed *threshold rules*, in which the decision-maker accepts s_i if and only if s_i (weakly or strictly) exceeds a specified threshold θ_i that can depend on the position i. The key component of our analysis is to establish a close, though subtle, technical connection between delegated search and prophet inequalities: roughly speaking, the sequence of values $x(\omega_1), x(\omega_2), \ldots, x(\omega_n)$ sampled by the agent from the set of possible solutions Ω plays the role of the process generating $s_1, s_2, \ldots, s_n$; and the principal and the agent jointly — through the principal's specification of the threshold and the agent's incentive to obey it — play the role of the decision-maker who uses a threshold rule for deciding when to stop. Again, the notion of "stopping" here is a bit oblique, since the principal never sees the full sequence $x(\omega_1), x(\omega_2), \ldots, x(\omega_n)$ that the agent generates; this is the sense in which the stopping rule is jointly constructed by the behavior of the principal and the agent together.

Stronger bounds for independent values. Using this connection to threshold rules for prophet inequalities, we can design a much more powerful policy for the principal in the case when the values of $x(\omega)$ and $y(\omega)$ on a draw ω from Ω are distributed independently, and when the principal can see both $x(\omega_i)$ and $y(\omega_i)$ (rather than only $x(\omega_i)$) in the solution ω_i proposed by the agent.

To do this, we begin with a stopping rule from the prophet inequality literature achieving an expected payoff that is at least 0.745 times the optimum when the distributions of the s_i values are independent and identically distributed [1, 9, 12, 15]. This stopping rule uses a sequence of thresholds θ_i that decrease with i, making the decision-maker naturally more prone to stop and accept a value as the end of the sequence nears — effectively following the idea that one should only accept a value early if it's very good.

In the context of delegated search when the principal can observe both $x(\omega_i)$ and $y(\omega_i)$ for a proposed solution ω_i, a related concept is useful for designing mechanisms: the principal should only accept a solution ω_i with $y(\omega_i)$ very large if $x(\omega_i)$ is very large as well. The analogy between requiring strong incentive to accept a value with large $y(\omega_i)$ in delegated search and requiring strong incentive to accept a value early in the sequence in the prophet inequality context can be made precise, and it reveals that the $y(\omega_i)$ values (over the set of candidate solutions considered by the agent) can be used as a kind of "continuous time" parameter for deriving a threshold: if we think of the candidate solution ω_i as arriving at continuous time $y(\omega_i)$, then we can derive a threshold function $\theta(\cdot)$ in which the principal only accepts ω_i if $x(\omega_i)$ (weakly or strictly) exceeds $\theta(y(\omega_i))$.[2] In this sense, the principal and agent again jointly construct the stopping rule, with the agent's payoff providing a type of synthetic temporal ordering that is useful in formulating a threshold policy.

Bounds for the binary model. We also use the connection to prophet inequalities to derive bounds for the binary model, where the agent pays to evaluate the feasibility of pairs from a known list of options $\omega_1, \omega_2, \ldots, \omega_m$. Here too the principal can designate a predefined eligible set of proposals so that the mechanism that accepts any eligible proposed ω_i yields an expected payoff that is within a factor of two of the benchmark in which the principal performs the search on their

[2] We observe that this continuous time defined by the $y(\omega_i)$ runs "in reverse," in the sense that large values of $y(\omega_i)$, like small values of time, place stricter demands on the $x(\omega_i)$ values that can be accepted.

own. However, the eligibility criterion in this case is subtler: it depends not only on the principal's assessment of the proposal's quality, $x(\omega_i)$, but also on the cost c_i and the *a priori* probability of feasibility, p_i.

To establish this bound, we draw on both prophet inequality bounds and on work of Kleinberg et al for the *box problem* [16, 22]; by considering an ordering of the options by the notion of *reservation price* (or, equivalently, *strike price*) defined in those works, we can establish a provable guarantee that correctly handles not only the payoff arising from the $x(\omega_i)$ and $y(\omega_i)$ values but also the cost incurred by the agent in evaluating the feasibility of options.

Finally, we derive similar bounds in the more general case where the agent also has a budget of $n < m$ on the number of options they can evaluate. The approach using reservation prices does not directly extend to this case, but we show that by combining the approach of Kleinberg et al. [16] with bounds for stochastic optimization due to Asadpour and Nazerzadeh [7], we can obtain more general bounds for a budgeted variant of the box problem that contains the case we need for our delegated search guarantee.

We note that it would be a natural open question to consider a variant of the problem combining characteristics of the two main versions we consider here: as in the distributional model, the agent performs independent draws from a space Ω; but as in the binary model, the agent does not have a fixed bound n on the number of allowed draws, instead incurring a cost to perform each draw that must be traded off against the eventual payoff from the sample selected.

Further Discussion of Related Work on Delegation

The theory of delegation in the economics literature is often viewed as beginning with Bengt Holmstrom's Ph.D. thesis [13, 14]; this work articulates the basic tension that we see in these models, between allowing an agent to optimize in a large space and restricting the agent's freedom of action to prevent them from pursuing their own objectives too aggressively. Holmstrom's model considered delegating an optimization problem over an interval, and a sequence of subsequent papers analyzed the case in which the agent optimizes over a continuum [3, 19]. Armstrong and Vickers [6] propose a model that is very close to what we consider here, where the optimization takes place over a discrete set that the agent samples from an underlying distribution. By way of comparison between our work and that of Armstrong and Vickers [6], we noted the key contrast earlier in this section: their paper is largely devoted to identifying cases of the delegated search problem for which the structure of the optimal mechanism can be identified, whereas we focus on bounding the inefficiency of delegated search relative to a benchmark in which the principal performs the task themself. It is through our emphasis on these types of bounds that we develop the connection to the analysis of prophet inequalities.

A distinct line of work in delegation relaxes the constraint that the principal may only allow or forbid each proposed solution, and instead allows the principal to add arbitrary amounts of cost to certain subsets of proposed solutions [4, 5, 8]. One of the key motivations for such a condition is to model the strategic role of bureaucracy within an organization: if management wants to dissuade units within the organization from proposing certain types of solutions, they can use bureaucratic measures (requiring more extensive justifications and processes) that make these solutions selectively more costly without explicitly forbidding them, and without engaging in explicit monetary transfers. Ambrus and Egorov [5] propose a model in which such selective cost increases are in fact part of the optimal delegation scheme.

2 MODEL AND PRELIMINARIES

We begin by making the precise the way in which the principal and the agent interact, resulting in the principal's selection of (at most) one element from a set Ω of potential solutions. There are functions $x : \Omega \to \mathbb{R}$ and $y : \Omega \to \mathbb{R}$ such that if $\omega \in \Omega$ is selected, then the principal's utility is $x(\omega)$ and the agent's utility is $y(\omega)$. To formalize the possibility that the principal selects no solution (i.e., perpetuating the status quo) we identify this possiblity with a special "null outcome", denoted $\perp$, and we extend the utility functions from Ω to $\Omega_+ \stackrel{\Delta}{=} \Omega \cup \{\perp\}$ by setting $x(\perp) = y(\perp) = 0$.

The set Ω is a probability space, with probability measure μ, and the agent has the power to draw independent samples from Ω according to μ. The principal, on the other hand, can neither draw samples from Ω nor directly observe the outcome of the agent's sampling; she must rely on her interaction with the agent to arrive at a selected element of Ω.[3]

Before formalizing our model of interaction, it is useful to first note some of the ways in which our basic model can be generalized or specialized.

- We will initially consider the case of a single probability measure μ on Ω, but it is also useful to consider cases in which there are multiple probability measures $\mu_1, \mu_2, \dots, \mu_m$ on Ω, and the agent has the power to draw independent samples from any of these distributions.
- We will generally assume there is a *sampling budget* of n on the number of samples that the agent can draw. In some of our models, we will also introduce a *sampling cost* $c \geq 0$ for each draw by the agent — or in the case of multiple probability measures, a cost $c_i \geq 0$ for sampling from μ_i.
- We consider both the *full-information case* — in which the principal knows both the functions $x : \Omega \to \mathbb{R}$ and $y : \Omega \to \mathbb{R}$, and hence can evaluate the utility of a solution ω to both herself and to the agent — and the *limited-information case*, in which the principal only knows her own utility function x.
- The functions $x : \Omega \to \mathbb{R}$ and $y : \Omega \to \mathbb{R}$ define random variables on Ω, and we consider both the case in which they can be arbitrary non-negative functions, and the case of *independent utilities*, when they are independent random variables.

In a later section, we will specialize the formalism to the *binary model* discussed in Section 1, in which each distribution μ_i is supported on a two-element set $\{\omega_{0i}, \omega_{1i}\}$ such that $(x(\omega_{0i}), y(\omega_{0i})) = (0, 0)$. In this case we will let (x_i, y_i) denote the pair $(x(\omega_{1i}), y(\omega_{1i}))$. The binary model captures a setting in which the feasibility of the i^{th} solution is unknown until the agent investigates it, but the value of the solution to both parties (if feasible) is known *a priori*.

2.1 A General Definition of Mechanisms for the Principal and Agent

Let us now formalize how the principal and the agent interact, resulting in the principal's selection of a solution. Thus far, our discussion in Section 1 has focused on interactions of a very structured form: the agent draws a set of samples from Ω; the agent selects one of these samples to present to the principal; and the principal accepts or rejects it. But it would be useful to be able to consider more general formulations for their allowed interactions, within which the transmission of a single proposal from the agent to the principal is a particular special case.

To do this, we begin by defining a *mechanism* as follows. A mechanism $M = (\Sigma, g)$ defines a set of signals, Σ, that the agent may send, and an allocation function $g : \Sigma \to \Omega_+$ that specifies which solution the principal will choose given the agent's signal. In such a mechanism, a *strategy* for the agent is specified by a mapping $\sigma : \Omega^* \to \Sigma$, where Ω^* denotes the set of finite sequences over Ω, such that $\sigma(\omega_1, \dots, \omega_n)$ represents the signal the agent sends if he sampled n solutions and observed the sequence $\omega_1, \dots, \omega_n$.

Suppose the agent observes sequence $\omega_1, \dots, \omega_n$ and sends signal σ, resulting in outcome $\omega = g(\sigma)$. In this case, the principal's and agent's utilities are $x(\omega)$ and $y(\omega)$, respectively, if

[3]Throughout this paper we use feminine pronouns for the principal and masculine pronouns for the agents.

$\omega \in \{\omega_1, \ldots, \omega_n\} \cup \{\perp\}$. Otherwise the principal's utility is 0 and the agent's is -1. In other words, we assume that if the mechanism results in the principal selecting a solution that was never sampled by the agent, that solution cannot be adopted. Instead the status quo is preserved and the agent suffers a penalty. This assumption is consistent with our assumption that the principal lacks the power to directly search for solutions herself; she can only adopt solutions that the agent has discovered.

In models with costly sampling, the specification of an agent's strategy must also include a sequential policy π for deciding which sample (if any) to observe next, given the set of samples already observed. The principal's and agent's utilities are both diminished by the sum of costs c_i for the samples i that the agent observed when running policy π. (We deduct this sum from the principal's utility because we think of the cost incurred by the agent in searching for a solution as a kind of "waste" that the principal views as detracting from the overall utility.)

Under our definition of mechanisms, the sequence of solutions sampled by the agent leads to a signal (via the agent's strategy σ), and this signal leads to a solution in Ω_+ (via the principal's allocation function g). Composing these two functions, we get a mapping from the agent's sampled solutions to a single solution:

Definition 2.1 (interim allocation function). If M is a mechanism and σ is an agent's strategy, the interim allocation function of the pair (M, σ) is the mapping $f_{M,\sigma} : \Omega^* \to \Omega_+$ obtained by composing the strategy σ with the allocation function g. In other words, $f_{M,\sigma}(\omega_1, \ldots, \omega_n)$ is the outcome resulting from mechanism M, when the agent draws sample sequence $(\omega_1, \ldots, \omega_n)$ and plays according to σ.

2.2 Single Proposal Mechanisms

We now show that there is a sense in which it is without loss of generality to focus on interactions in which the agent proposes a single solution from among the ones they sampled, and the principal either accepts or rejects it. To do this, we define a simple type of mechanism called a *single proposal mechanism*, and we show in Lemma 1 below that any other mechanism can be simulated by a single proposal mechanism.

Definition 1. A *single-proposal mechanism* with eligible set $R \subseteq \Omega$ is a mechanism in which the agent proposes one outcome, and the mechanism accepts this proposal if and only if it belongs to R. More formally, $M = (\Sigma, g)$ is a single-proposal mechanism with eligible set R if $\Sigma = \Omega_+$ and g restricts to the identity function on R and the constant function $\omega \mapsto \perp$ on $\Omega_+ \setminus R$.

Lemma 1. *If M is any mechanism and σ is any strategy constituting a best response to M, then there exists a single proposal mechanism M' and a best response σ' to M', such that the interim allocation functions $f_{M,\sigma}$ and $f_{M',\sigma'}$ are identical.*

Proof. Let R be the range of the interim allocation function $f_{M,\sigma}$, i.e. the set of all possible outcomes of M, other than $\perp$, when the agent acts according to σ. Define M' to be the single-proposal mechanism with eligible set R. Let σ' be the strategy in which the agent observes his tuple of samples, $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_n)$, and chooses strategy $\sigma'(\boldsymbol{\omega}) = g(\sigma(\boldsymbol{\omega}))$. By construction the interim allocation functions $f_{M,\sigma}$ and $f_{M',\sigma'}$ are identical. To prove that σ' is a best response to M', consider any $\boldsymbol{\omega} \in \Omega^n$ and any $\nu \neq \sigma'(\boldsymbol{\omega})$. Let $y_0 = y(g'(\sigma'(\boldsymbol{\omega})))$ denote the agent's utility when playing according to σ'; note that $y_0 \geq 0$. We wish to show that the agent cannot benefit by playing ν instead, i.e.

$$y(g'(\nu)) \leq y_0. \tag{1}$$

If $v \notin R$ then $g'(v) = \perp$ which implies (1) since $y_0 \geq 0$. If $v \in R$ then $v = g(\tilde{\sigma})$ for some $\tilde{\sigma} \in \Sigma$. Now (1) follows because strategy σ is a best response for mechanism M, and $y(g'(v))$ denotes the agent's utility when playing strategy $\tilde{\sigma}$ in M whereas y_0 denotes his utility when playing σ. □

3 ANALYZING DELEGATED SEARCH VIA PROPHET INEQUALITIES

In this section we develop a formal link between delegated search mechanisms and prophet inequalities. It turns out that the relevant prophet inequalities involve random variables arriving at discrete points in continuous time, rather than the usual assumption that they arrive at time points $1, 2, \ldots, n$. Accordingly, we will begin by explaining the formal model of continuous-time prophet inequalities in Section 3.1 below. Then in Section 3.2 we explain the reduction from delegated search (in the distributional model) to continuous-time prophet inequalities.

3.1 Continuous-time prophet inequalities

In this section we will be concerned with problems which involve designing a selection rule to choose (at most) one element from a random finite set of pairs $(x_i, t_i) \in \mathbb{R}_+ \times [0, 1]$, with the goal of maximizing the expected x-coordinate of the chosen element. The t-coordinate is thought of as a time coordinate, and we will generally (but not exclusively) be concerned with selection rules that make their choice without looking into the future, as is ordinarily the case in the analysis of prophet inequalities.

Definition 2 (selection rules)**.** A *selection rule* is a function ρ from finite subsets of $\mathbb{R}_+ \times [0, 1]$ to the set $(\mathbb{R}_+ \times [0, 1]) \cup \{\perp\}$, with the property that $\rho(S) \in S \cup \{\perp\}$ for every S.

A *stopping rule* is a selection rule that chooses element (x, t) from set S without looking at the set of elements whose time coordinate is greater than t. Formally, ρ is a stopping rule if it satisfies the following property: for any $(x, t) \in \mathbb{R}_+ \times [0, 1]$ and any two sets S, S' such that $S \cap (\mathbb{R}_+ \times [0, t]) = S' \cap (\mathbb{R}_+ \times [0, t])$, we have

$$\rho(S) = (x, t) \iff \rho(S') = (x, t).$$

An *oblivious stopping rule* with eligible set $Q \subseteq \mathbb{R}_+ \times [0, 1]$ is a stopping rule ρ_Q such that for every S, $\rho_Q(S)$ is an earliest element of $S \cap Q$ (i.e., an element of that set with minimum t coordinate) or $\rho_Q(S) = \perp$ if $S \cap Q$ is empty.

A *threshold stopping rule* with threshold θ is an oblivious stopping rule whose eligible set is of the form $(\theta, \infty) \times [0, 1]$ or $[\theta, \infty) \times [0, 1]$.

Definition 3 (CTSPs and prophet inequalities)**.** A *continuous-time selection problem* (CTSP) is an ordered pair $(\mathscr{D}, R)$ where $\mathscr{D}$ is a set of probability distributions over finite subsets of $\mathbb{R}_+ \times [0, 1]$, and $\mathscr{R}$ is a set of selection rules.

A CTSP $(\mathscr{D}, R)$ satisfies a *prophet inequality with factor* α if it is the case that for every $\mathbf{D} \in \mathscr{D}$ there exists some $\rho \in \mathscr{R}$ such that

$$\mathbb{E}_{S\sim\mathbf{D}}[X_{\rho(S)}] \geq \alpha \cdot \mathbb{E}_{S\sim\mathbf{D}}[X_{*(S)}].$$

Here the random variable $X_{\rho(S)}$ is defined by specifying that if $\rho(S) = (x, t)$ then $X_{\rho(S)} = x$, and if $\rho(S) = \perp$ then $X_{\rho(S)} = 0$. The random variable $X_{*(S)}$ is defined to be $\max\{x | (x, t) \in S\}$.

We now present the prophet inequalities we will use in this work. To state them, we define the following families of stopping rules and distributions on subsets of $\mathbb{R}_+ \times [0, 1]$.

- $\mathscr{R}_{\text{obliv}}$ is the family of oblivious stopping rules.
- $\mathscr{R}_{\text{thresh}}$ is the family of threshold stopping rules.

- $\mathscr{D}_{\text{ind}}$ is the family of random sets whose elements are obtained by sampling independently from n joint distributions. In other words, a distribution $\mathbf{D} \in \mathscr{D}_{\text{ind}}$ is specified by giving a positive number n, a tuple of joint distributions $\mathbf{D}_1, \dots, \mathbf{D}_n$ on $\mathbb{R}_+ \times [0,1]$, and defining $\mathbf{D}$ to be the distribution on n-element sets obtained by drawing one sample independently from each of $\mathbf{D}_1, \dots, \mathbf{D}_n$.
- $\mathscr{D}_{\text{iid},n}$ is the family of random sets whose n elements are i.i.d. samples from an atomless distribution with x and t independent, i.e. a distribution on $\mathbb{R}_+ \times [0,1]$ which is a product of atomless distributions.
- $\mathscr{D}_{\text{iid}}$ is the union of $\mathscr{D}_{\text{iid},n}$ over all $n \geq 1$.

Our first prophet inequality is Samuel-Cahn's famous prophet inequality for threshold stopping rules [21].

Theorem 1. *(Samuel-Cahn [21]) There is a prophet inequality with factor $\frac{1}{2}$ for $(\mathscr{D}_{ind}, \mathscr{R}_{thresh})$.*

The second is an improved prophet inequality for threshold stopping rules when samples are drawn i.i.d. from atomless product distributions; it can be derived as a corollary of either [10, Theorem 19] or [9, Corollary 2.2].

Theorem 2. *(Correa et al. [9], Ehsani et al. [10]) There is a prophet inequality with factor $1 - \frac{1}{e}$ for $(\mathscr{D}_{iid}, \mathscr{R}_{thresh})$.*

Our third prophet inequality again pertains to the case when samples are drawn i.i.d. from atomless product distributions, but it allows for oblivious stopping rules rather than threshold stopping rules. The discrete-time counterpart to this prophet inequality can be found in [9, 12, 15].

Theorem 3. *Let $\alpha = 0.745\dots$ be the solution to $\int_0^1 \left(y - y\ln y + 1 - \frac{1}{\alpha}\right)^{-1} dy = 1$. There is a prophet inequality with factor $\alpha - O\left(\frac{\log n}{n}\right)$ for $(\mathscr{D}_{\text{iid},n}, \mathscr{R}_{obliv})$.*

Since the distincton between discrete time and continuous time is immaterial from the standpoint of analyzing threshold stopping rules, the first two of these theorems are equivalent to the existing results for discrete-time prophet inequalities that we have cited before the theorem statements. On the other hand, because oblivious stopping rules are less powerful than general stopping rules, the third theorem is not an immediate consequence of the corresponding discrete-time prophet inequality.

To complete this section, we will describe the stopping rules which achieve the bounds stated in the three prophet inequalities above.

When the points $\{(x_i, t_i)\}_{i=1}^n$ are independent but not necessarily identically distributed, choose threshold $\theta_{1/2}$ to be the median of the distribution of $\max\{x_i\}$. In other words, $\theta_{1/2}$ is defined such that the events $\max\{x_1, \dots, x_n\} > \theta_{1/2}$ and $\max\{x_1, \dots, x_n\} < \theta_{1/2}$ both have probability at most $\frac{1}{2}$. Consider the threshold stopping rule that selects the first pair (x_i, t_i) with $x_i > \theta_{1/2}$, and consider the one whose selection criterion is $x_i \geq \theta_{1/2}$. The proof of Theorem 1 shows that at least one of these two stopping rules fulfills a prophet inequality with factor 1/2.

When the points $\{(x_i, t_i)\}_{i=1}^n$ are i.i.d. and the distributions of x_i and t_i are atomless and independent, with cumulative distribution functions F and G, respectively, choose threshold $\theta_{1-1/e}$ such that $\Pr(\max\{x_1, \dots, x_n\} > \theta_{1-1/e}) = 1 - \frac{1}{e}$. The proof of Theorem 2 shows that the threshold stopping rule that selects the first pair (x_i, t_i) such that $x_i > \theta_{1-1/e}$ fulfills a prophet inequality with factor $1 - \frac{1}{e}$. Now let β_n be the solution to

$$\int_0^n \left(1 + z + \beta_n e^z\right)^{-1} dz = 1,$$

and let $z_n(s)$ be the solution of the differential equation

$$\frac{dz}{ds} = 1 + z + \beta_n e^z$$

with initial condition $z(0) = 0$. The oblivious stopping rule that accepts the first (x_i, t_i) such that

$$1 - F(x_i) < z(G(t_i))/n$$

fulfills a prophet inequality with factor $\alpha - \frac{\log n}{n}$.

3.2 Reducing delegated search to prophet inequalities

Although delegated search problems and prophet inequalities appear unrelated at first glance, the tight technical connection between them is explained by an observation which is extremely natural in hindsight. Consider a change of variables that maps the agent's utility y to a point $t(y) \in [0, 1]$, where the function t is monotonically decreasing. In a single proposal mechanism with eligible set R, the agent submits the eligible proposal (x, y) with the highest y value. Similarly, an oblivious stopping rule with eligible set Q selects the earliest eligible point (x, t). Since the change of variables $t(y)$ is monotonically decreasing, the two selection criteria are equivalent! Thus, designing single proposal mechanisms that yield high utility for the principal is equivalent to designing oblivious stopping rules that yield a high expected value.

In more detail, let t be any continuous, monotonically decreasing bijection from $[0, \infty)$ to $(0, 1]$, for example $t(y) = e^{-y}$. Under the mapping $H : \Omega \to (\mathbb{R}_+ \times [0, 1])$ defined by $H(\omega) = (x(\omega), t(y(\omega)))$, any distribution on sets of solutions $\{\omega_1, \ldots, \omega_n\}$ induces a distribution $\mathbf{D}$ on sets of pairs $\{(x_1, t_1), \ldots, (x_n, t_n)\}$. In particular, our distributional model in which the agent draws n i.i.d. samples from Ω is mapped, under this correspondence, to a member of the family of distributions $\mathscr{D}_{\text{iid}, n}$.

There is also a reverse correspondence from oblivious stopping rules to single proposal mechanisms and their interim allocation functions. The oblivious stopping rule ρ_Q with eligible set Q corresponds to the single proposal mechanism with eligible set $H^{-1}(Q)$. More precisely, if $R = H^{-1}(Q)$ and σ is a best response to the single proposal mechanism M with eligible set R, then for any sequence of samples $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_n)$, we have

$$\rho_Q(H(\boldsymbol{\omega})) = H(f_{M,\sigma}(\boldsymbol{\omega})).$$

In other words, suppose we run the mechanism M; the agent draws a sequence of samples; and we let the agent choose the best one (for the agent) that belongs to R. This procedure is equivalent to running the oblivious stopping rule ρ_Q on the sequence obtained by transforming all of the agents' samples to points $(x_i, t_i) = (x_i, t(y_i))$, and selecting the earliest such point (ordered by t_i) that belongs to Q. Under this correspondence, threshold stopping rules correspond to single proposal mechanisms in which a solution is deemed eligible if the principal's utility exceeds a specified threshold. Note that this subset of single proposal mechanisms can be implemented even when the agent's utility is unobservable.

Combining these observations with Theorems 1 to 3, we obtain the following theorem.

THEOREM 4. *In the distributional model, suppose the agent draws n i.i.d. samples, and let x^* denote the utility the principal would attain if she could directly choose her favorite among these n samples.*

(1) *There is always a set X of the form (θ, ∞) or $[\theta, \infty)$ such that a single proposal mechanism with eligible set $\{\omega \mid x(\omega) \in X\}$ ensures that the principal's expected utility is at least $\frac{1}{2}\mathbb{E}[x_*]$.*
(2) *If the principal and agent have independent utilities, each drawn from an atomless distribution, then a single proposal mechanism that accepts any proposal satisfying $x(\omega) > \theta$, for a suitable choice of θ, ensures that the principal's expected utility is at least $\left(1 - \frac{1}{e}\right)\mathbb{E}[x_*]$.*

(3) *If the principal and agent have independent utilities, each drawn from an atomless distribution, and the principal can observe the agent's utility, then a single proposal mechanism that accepts any proposal satisfying* $x(\omega) > \theta(y(\omega))$*, for a suitable choice of the function* $\theta(\cdot)$*, ensures that the principal's expected utility is at least* $\left(\alpha - O\left(\frac{\log n}{n}\right)\right)\mathbb{E}[x_*]$*, where* α *is the constant defined in Theorem 3.*

It is possible to show that the bounds in all three parts of the theorem are tight with respect to the assumptions made in their respective statements.

4 BINARY OUTCOMES

Recall the binary model from Section 1: The potential solutions come from a large discrete set $\Omega = \{\omega_1, \omega_2, \ldots, \omega_m\}$ and the agent's role is to explore which of these options are feasible to implement. If ω_i is feasible, it yields utility x_i for the principal and y_i for the agent — where the pair (x_i, y_i) is commonly known to both parties — and if ω_i is infeasible it yields zero utility for both parties. To explore the feasibility of solution ω_i the agent must incur a cost of $c_i \geq 0$, and the probability of success is $p_i > 0$, independently of the success of other solutions. These quantities c_i, p_i are again commonly known to both parties. We will assume that $c_i \leq p_i y_i$ for each solution y_i, since otherwise it is against the agent's self-interest to explore ω_i, even if it were assured that the solution would be adopted if feasible.

4.1 Optimal search policies: Weitzman's box problem

If the principal were conducting the search by herself (without delegation to an agent), this model would correspond to a special case of the *box problem* introduced by Weitzman [22]. The optimal search policy is simple but surprisingly subtle: it assigns to each option a priority z_i satisfying $\mathbb{E}[(x_i - z_i)^+] = c_i$ — which in our case entails setting $z_i = x_i - c_i/p_i$ — and then explores options in decreasing order of priority, selecting the first feasible one in this ordering or stopping when all remaining unexplored options have $z_i < 0$.

Now suppose that the principal instead delegates the search to an agent who bears the cost of exploration, by running a single-proposal mechanism with eligible set R. Then the agent faces a different instance of the box problem, in which the set of options is limited to R, and the costs and success probabilities of the options is the same as before, but the value of option i (if feasible) is y_i rather than x_i. This means the agent prioritizes boxes in decreasing order of $w_i = y_i - c_i/p_i$ rather than $z_i = x_i - c_i/p_i$, and recommends the first box in this ordering that is discovered to be feasible.

To summarize, the delegated search problem in the binary model is analogous to Weitzman's box problem, but with the important distinction that the searcher (the principal) is not allowed to choose the order in which to open the boxes. Instead the problem specifies an exogenous ordering of the boxes — corresponding to the agent's ranking of options by decreasing w_i — and the searcher is only free to decide which boxes in this sequence should be opened and which ones should be skipped, corresponding to the principal's problem of choosing the set R. Since this problem may be of independent interest, we devote Theorem 5 below to presenting a solution that always achieves at least half of the expected value of running the optimal search procedure that is allowed to inspect the boxes in any order it desires. Interestingly, the analysis is based on prophet inequalities, specifically Theorem 1 and its proof. It implies there is an approximately optimal mechanism with the following structure. For any half-infinite interval X of the form $X = (\theta, \infty)$ or $X = [\theta, \infty)$, let $R(X) = \{\omega_i \mid z_i \in X\}$ and define $M(X)$ to be the single-proposal mechanism in which a proposal ω_i is eligible if it is feasible and belongs to $R(X)$.

4.2 The Box Problem with an Exogenous Ordering

In this section we recapitulate some background material about Weitzman's 1979 box problem. In this problem[4] there are m boxes, each containing an independent random prize. The prize in box i is denoted v_i, and the cost of opening the box is c_i. A searcher may open any number of boxes sequentially, or may cease the search at any time and claim a prize from at most one of the open boxes. The problem is to design an optimal sequential search policy. Weitzman proves that if each box is assigned a *priority* z_i defined by the equation $\mathbb{E}[(v_i - z_i)^+] = c_i$, then the optimal sequential search policy opens boxes in decreasing order of priority, stopping at the first time when the highest prize inside an open box exceeds the highest priority of a closed box, or at the first time when the priority of every remaining closed box is negative, whichever comes sooner.

Kleinberg et al. [16] provided a proof of optimality of Weitzman's procedure in which the priority z_i is interpreted as the "strike price" of a real option with fair value c_i. An important quantity in their analysis is the "covered call value", which is simply the random variable $\kappa_i = \min\{v_i, z_i\}$. We restate the following lemma[5] from their work.

Lemma 2. *(Kleinberg et al. [16]) For any sequential search procedure and any box i, let A_i, B_i be the indicator random variables of the event that the procedure selects box i and the event that it opens box i, respectively. The inequality*

$$\mathbb{E}[A_i v_i - B_i c_i] \leq \mathbb{E}[A_i \kappa_i] \tag{2}$$

is satisfied by every search procedure, and equality holds if and only if the search procedure is non-exposed*, meaning that $A_i = B_i$ at every sample point where $v_i > z_i$.*

Corollary 1. *For any sequential search procedure, the expected net value of running the procedure (i.e., the value of the selected box minus the combined cost of opening boxes) is bounded above by the expectation of the maximum covered call value, i.e.*

$$\mathbb{E}\left[\sum_{i=1}^{m} A_i v_i - \sum_{i=1}^{m} B_i c_i\right] \leq \mathbb{E}\left[\sum_{i=1}^{m} A_i \kappa_i\right]. \tag{3}$$

The corollary is immediate, by summing inequality (2) over boxes $i = 1, \ldots, m$.

Now consider the box problem with an exogenous ordering of boxes, where the searcher is limited to considering the boxes one by one in the specified order, and once she decides to leave a box closed or to leave the prize within unclaimed, she cannot later return to the box and open it or claim its prize. We define a type of policy that we call a *κ-thresholding policy*; the reason for the name will become apparent in the subsequent Lemma 3, which shows that these policies correspond to a threshold rule applied to the sequence of covered call values κ_i.

Definition 4. A *κ-thresholding policy* for the box problem with exogenous ordering is a policy that operates as follows. There is a half-infinite interval $X = (\theta, \infty)$ or $X = [\theta, \infty)$ called the *target interval*. The policy declines to open any box i with $z_i \notin X$. Otherwise, if $z_i \in X$, the policy opens the box and claims the prize inside if and only if $v_i \in X$.

Lemma 3. *Every κ-thresholding policy is non-exposed. The expected net value of running a κ-thresholding policy with target interval X is exactly the same as the expected value selected by the threshold stopping rule that observes the random sequence $\kappa_1, \kappa_2, \ldots, \kappa_n$ and selects the first element of this sequence that belongs to X.*

[4]The following description constitutes a special case of Weitzman's problem. The general case incorporates geometric time discounting and time delays.

[5]Lemma 1 of the full version of their paper, http://dx.doi.org/10.2139/ssrn.2753858.

Proof. The policy is non-exposed because $z_i \notin X$ implies $A_i = B_i = 0$, while $z_i \in X$ and $v_i > z_i$ imply $A_i = B_i = 1$. Hence the left and right sides of (2) are equal for every box, and the net value of running the policy is $\mathbb{E}\left[\sum_{i=1}^{m} A_i \kappa_i\right]$, i.e. the expected covered call value of the box the policy selects. By design, the policy perfectly simulates the threshold stopping rule that chooses the first element of the sequence $\kappa_1, \ldots, \kappa_n$ that belongs to X; this is because it selects the first box such that z_i and v_i both belong to X, which is also the first box such that κ_i belongs to X. □

Theorem 5. *For every instance of the box problem with exogenous ordering, there is a κ-thresholding policy whose expected net value is at least half that of Weitzman's optimal search procedure (which endogenously selects the ordering of the boxes).*

Proof. Lemma 3 reduces the analysis of κ-thresholding policies to a question about prophet inequalities. In particular, the expected net value of running a κ-thresholding policy is equal to the expected covered call value of the random element selected from the sequence $\kappa_1, \ldots, \kappa_n$ by a particular threshold stopping rule. Since Samuel-Cahn's 1984 prophet inequality (Theorem 1 above) implies that threshold stopping rules can always attain at least half the expectation of the maximum random variable in the sequence, it follows that there is a κ-thresholding policy whose expected net value is at least half the expectation of the maximum covered call value. Corollary 1 ensures that the latter is an upper bound on the expected net value of Weitzman's optimal search procedure. □

4.3 An approximately optimal mechanism

Recall that for a half-infinite interval $X = (\theta, \infty)$ or $X = [\theta, \infty)$, the mechanism $M(X)$ is defined to be the single proposal mechanism whose eligible set consists of solutions ω_i that are feasible and satisfy $z_i \in X$.

Theorem 6. *There exists a choice of X such that the expected net value of mechanism $M(X)$ — i.e., the principal's value for adopting the agent's proposal, if adopted, minus combined cost of all the alternatives explored — is at least half of the expected net value the principal could achieve by performing the optimal search herself (without delegation).*

Proof. Convert the delegated search problem into a box problem with exogenous order, where the order is defined by sorting the solutions $\omega_1, \ldots, \omega_m$ in non-increasing order of the agent's priority value $w_i = y_i - c_i/p_i$, and the value v_i inside box i is defined to be x_i if ω_i turns out to be feasible, 0 otherwise.

According to Theorem 5 there exists a choice of X such that the κ-thresholding policy with target set X attains at least half the expected net value of the optimal search procedure. This thresholding policy goes through boxes in the given order, i.e. descending w_i, and opens only those with $z_i \in X$, selecting the first one such that $v_i \in X$. Note that among the boxes which the policy opens, the first one with $v_i \in X$ is also the first one corresponding to a feasible ω_i. This is because an infeasible ω_i has $v_i = 0$ hence $v_i \notin X$, whereas a feasible ω_i has $v_i = x_i \geq z_i$, hence $v_i \in X$.

Recall from Section 4.1 that the agent's best response to mechanism $M(X)$ is to go through the elements of $R(X)$ in decreasing order of w_i, stopping and proposing the first one that is discovered to be feasible. This is exactly the behavior of the κ-thresholding policy with target set X, as derived in the preceding paragraph. Hence the mechanism $M(X)$ coupled with the agent's best response behavior emulates the κ-thresholding policy which attains at least half the expected net value of the optimal search procedure. □

4.4 Limiting the number of samples

In some cases the number of distinct potential solutions, m, may be prohibitively large, and the agent may only have the power to explore the feasibility of a limited number of them, $n < m$. In this case, if the principal were to conduct the search autonomously without delegation — subject to the same costs c_i and the same upper bound, n, on the total number of solutions that can be tested for feasibility — it may require a very complex procedure. Nevertheless, we will provide in this section a simple delegated search mechanism such that it is easy for the agent to compute a search procedure that is a best response to the mechanism, and the outcome of running the mechanism with this best response attains at least $\frac{1}{2}\left(1-\frac{1}{e}\right) \approx 0.316$ of the net expected value of the (potentially complex) optimal procedure.

The key observation is the following lemma, which provides a useful upper bound on the value of running the optimal search procedure.

Lemma 4. *In the box problem with $m > n$ boxes, if the searcher is limited to open at most n boxes before claiming a prize, then the expected net value of any search procedure is bounded above by $\mathbb{E}[\max_{i \in S} \kappa_i]$ where S is the random set of boxes that the procedure opens.*

Proof. Sum up the inequality (2) over all boxes and note that $A_i = 0$ for $i \notin S$, to derive

$$\mathbb{E}\left[\sum_{i=1}^{m} A_i v_i - \sum_{i=1}^{m} B_i c_i\right] \leq \mathbb{E}\left[\sum_{i \in S} A_i \kappa_i\right].$$

The lemma follows by noting that $\sum_{i \in S} A_i \kappa_i \leq \max_{i \in S} \kappa_i$ because $\sum_{i \in S} A_i \leq 1$. □

Lemma 5. *There exists a (non-random) set T of cardinality n, such that $\mathbb{E}[\max_{i \in T} \kappa_i] \geq (1 - \frac{1}{e})\mathbb{E}[\max_{i \in S} \kappa_i]$, where S is the random set of solutions explored by the optimal search procedure subject to a contraint of exploring at most n solutions.*

Proof. The problem of adaptively exploring a random set S of at most n solutions to maximize $\mathbb{E}[\max_{i \in S} \kappa_i]$ is a special case of the stochastic monotone submodular function maximization problem studied by Asadpour and Nazerzadeh [7], in which the role of the monotone submodular function $f : \mathbb{R}_+^n \to \mathbb{R}_+$ is played by the function $f(\lambda_1, \ldots, \lambda_n) = \max\{\lambda_i\}$, and role of the matroid constraint is played by the cardinality constraint that at most n elements may be probed. Theorem 1 of [7], which asserts that the adaptivity gap of stochastic monotone submodular maximization is $\frac{e}{e-1}$, specializes in the present case to the assertion stated in the lemma. □

Theorem 7. *Consider delegated search in the binary model with a constraint that no more than n solutions can be examined for feasibility. There exists a mechanism that attains at least $\frac{1}{2}\left(1-\frac{1}{e}\right)$ fraction of the expected net value of the optimal search procedure subject to the same limitation of examining at most n solutions.*

Proof. According to Lemmas 4 and 5, there is an n-element set $T \subseteq \Omega$ such that the optimal search procedure that is limited to explore only solutions in T is able to attain at least $1 - \frac{1}{e}$ fraction of the expected net value of the optimal search procedure that is limited to examine at most n solutions but can (adaptively) choose any n elements of Ω during its search. When the set of solutions is restricted to T, the constraint that at most n solutions can be examined becomes irrelevant since T only has n elements. Thus, Theorem 6 guarantees the existence of a delegated search mechanism that is at least half as good as the optimal search procedure limited to T, and is consequently at least $\frac{1}{2}\left(1-\frac{1}{e}\right)$ times as good as the optimal search procedure limited to examine at most n solutions. Moreover, by applying the algorithm in Asadpour and Nazerzadeh [7] used to prove Lemma 5, we can implement this policy in polynomial time with a loss of a further additive ε in the approximation ratio, thus obtaining a bound of $\frac{1}{2}\left(1 - \frac{1}{e} - \varepsilon\right)$ efficiently. □

Acknowledgements. This work was supported in part by NSF grants CCF-1512964, CCF-1740822, and SES-1741441, a grant from the MacArthur Foundation, and a Simons Investigator Award. The authors would like to thank Brendan Lucier, Jens Ludwig, Sendhil Mullainathan, Rad Niazadeh, and Glen Weyl for helpful discussions, and The Nines in Ithaca, NY, for the many deep-dish pizzas that were consumed during the course of this research.

REFERENCES

[1] Melika Abolhassani, Soheil Ehsani, Hossein Esfandiari, MohammadTaghi Hajiaghayi, Robert D. Kleinberg, and Brendan Lucier. 2017. Beating 1-1/e for ordered prophets. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017.* 61–71. DOI: http://dx.doi.org/10.1145/3055399.3055479
[2] Philippe Aghion and Jean Tirole. 1997. Formal and Real Authority in Organizations. *Journal of Political Economy* 105, 1 (Feb. 1997), 1–29.
[3] Ricardo Alonso and Niko Matouschek. 2008. Optimal Delegation. *Review of Economic Studies* 75, 1 (Jan. 2008), 259–293.
[4] Manuel Amador and Kyle Bagwell. 2013. The Theory of Optimal Delegation With an Application to Tariff Caps. *Econometrica* 81, 4 (July 2013), 1541–1599.
[5] Attila Ambrus and Georgy Egorov. 2017. Delegation and nonmonetary incentives. *Journal of Economic Theory* 171 (2017), 101–135.
[6] Mark Armstrong and John Vickers. 2010. A model of delegated project choice. *Econometrica* 78, 1 (Jan. 2010), 213–244.
[7] Arash Asadpour and Hamid Nazerzadeh. 2016. Maximizing Stochastic Monotone Submodular Functions. *Management Science* 62, 8 (2016), 2374–2391. DOI: http://dx.doi.org/10.1287/mnsc.2015.2254
[8] Susan Athey, Kyle Bagwell, and Chris Sanchirico. 2004. Collusion and Price Rigidity. *Review of Economic Studies* 71, 2 (April 2004), 317–349.
[9] José R. Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. 2017. Posted Price Mechanisms for a Random Stream of Customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017.* 169–186. DOI: http://dx.doi.org/10.1145/3033274.3085137
[10] Soheil Ehsani, MohammadTaghi Hajiaghayi, Thomas Kesselheim, and Sahil Singla. 2018. Prophet Secretary for Combinatorial Auctions and Matroids. In *Proc. 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2018).* 700–714.
[11] Joseph Farrell and Michael L. Katz. 2006. The economics of welfare standards in antitrust. *Competition Policy International* 2 (2006), 3–28.
[12] Theodore P Hill and Robert P Kertz. 1982. Comparisons of stop rule and supremum expectations of iid random variables. *The Annals of Probability* 10, 2 (1982), 336–345.
[13] Bengt Holmstrom. 1977. *On Incentives and Control in Organizations.* Ph.D. Dissertation. Stanford University.
[14] Bengt Holmstrom. 1984. On the Theory of Delegation. In *Bayesian Models in Economic Theory*, M. Boyer and R. Kihlstrom (Eds.). Elsevier, 115–141.
[15] Robert P Kertz. 1986. Stop rule and supremum expectations of iid random variables: a complete comparison by conjugate duality. *Journal of Multivariate Analysis* 19, 1 (1986), 88–112.
[16] Robert D. Kleinberg, Bo Waggoner, and E. Glen Weyl. 2016. Descending Price Optimally Coordinates Search. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016.* 23–24. DOI: http://dx.doi.org/10.1145/2940716.2940760
[17] Vijay Krishna and John Morgan. 2008. Contracting for information under imperfect commitment. *RAND Journal of Economics* 39, 4 (2008), 905–925.
[18] Jin Li, Niko Matouschek, and Michael Powell. 2017. Power Dynamics in Organizations. *American Economic Journal: Microeconomics* 9, 1 (2017), 217–241.
[19] Nahum D. Melumad and Toshiyuki Shibano. 1991. Communication in Settings with No Transfers. *RAND Journal of Economics* 22, 2 (1991), 173–198.
[20] Volker Nocke and Michael D. Whinston. 2010. Dynamic Merger Review. *Journal of Political Economy* 118, 6 (2010), 1200–1251.
[21] Ester Samuel-Cahn. 1984. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Annals of Probability* 12, 4 (1984), 1213–1216.
[22] Martin L. Weitzman. 1979. Optimal Search for the Best Alternative. *Econometrica* 47 (1979), 641–654.