Efficient Multi-Agent Delegated Search

Extended Abstract

Curtis Bechtel
University of Southern California
Los Angeles, United States
bechtel@usc.edu

Shaddin Dughmi
University of Southern California
Los Angeles, United States
shaddin@usc.edu

ABSTRACT

Consider a *principal* who wants to search through a space of stochastic solutions for one maximizing their utility. If the principal cannot conduct this search on their own, they may instead delegate this problem to an *agent* with distinct and potentially misaligned utilities. This is called delegated search, and the principal in such problems faces a mechanism design problem in which they must incentivize the agent to find and propose a solution maximizing the principal's expected utility. Following prior work in this area, we consider mechanisms without payments and aim to achieve a multiplicative approximation of the principal's utility when they solve the problem without delegation.

In the full version [6] of this extended abstract, we investigate a natural and recently studied generalization of this model to multiple agents and find nearly tight bounds on the principal's approximation as the number of agents increases. As one might expect, this approximation approaches 1 with increasing numbers of agents, but, somewhat surprisingly, we show that this is largely not due to direct competition among agents.

KEYWORDS

mechanism design, principal-agent problems, delegation

ACM Reference Format:
Curtis Bechtel and Shaddin Dughmi. 2025. Efficient Multi-Agent Delegated Search: Extended Abstract. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 3 pages.

1 INTRODUCTION

Consider the problem faced by a non-profit organization that funds academic research through grants. As a member of this organization, you are tasked with allocating a fixed amount of resources to fund a single research project, and there are a variety of research groups from which you can receive, evaluate, and approve proposals. You are confident in your ability to evaluate proposals, but recognize that each research group has its own interests that may be misaligned with that of your organization. You must try to design a grant proposal mechanism that motivates these groups to propose research projects that are most beneficial to your non-profit's goals.

This is a practical and simplified example of the kind of problems that we investigate in this paper. More specficially, we consider a model in which the principal faces a stochastic optimization problem where they have to find a solution maximizing the expected value of some objective function. The task of searching for solutions to this problem is then delegated to a fixed group of agents, who each have distinct utility functions. Agents propose solutions to the principal, and the principal picks a single winner who receives utility for their proposal. We focus on models of delegation in which the principal's mechanism can not make outcome-contingent payments, representing situations in which players are confined to a fixed-price contract or are not allowed to make transfers of value for specific outcomes.

Finally, in contrast to designing optimal mechanisms, we build on recent delegation research [5, 7, 8, 13, 15] in which the principal aims for a multiplicative approximation of their first-best expected utility, i.e. their expected utility when the problem is not delegated (alternatively, their utility when they delegate to agents with identical interests as the principal). This approximation factor, which can be called the delegation gap, tells the principal what minimum fraction of their optimal utility they are guaranteed while delegating to arbitrary untrusted agents.

Prior work on approximation guarantees in delegation problems includes that of Kleinberg and Kleinberg [15], which proposed two models of delegation without payments and showed that they could be reduced to certain prophet inequalities and Pandora's box problems. This was expanded on by Bechtel and Dughmi [5], who introduced combinatorial constraints on the principal and agent, and Bechtel et al. [7], who consider different models of delegating Pandora's box problems. This line of work has shown that delegation has close connections to prophet inequalities [10, 11, 16–19] and contention resolution schemes [1, 9, 11], among other related problems. A few other notable papers on related delegation problems include [2–4, 8, 12, 14]. See the full version of this paper [6] or [5, 13, 15] for a more thorough overview.

Most similar is the work of Hajiaghayi et al. [13], who proposed the model of strategic multi-agent delegation that we study. They show that when all agents have the same number of i.i.d. elements, the principal can achieve approximations tending to 1 as αkm increases, where k is the number of agents, m is the number of elements per agent, and α is a parameter of the distributions. In contrast, we achieve an approximation tending to 1 as k increases when agents have symmetric sets of elements (not necessarily i.i.d.), with no conditions on the distributions or number of elements.

2 MODEL

The delegation model of interest in this paper is strategic multi-agent delegation as originally defined by Hajiaghayi et al. [13].

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

Definition 2.1 (Strategic Multi-Agent Delegation). An instance of *strategic multi-agent delegation* consists of k agents labeled $1, \ldots, k$, each of which has a collection of distinct elements. Each element e is associated with an a-priori unknown *outcome* $\omega_e \in \Omega$ drawn from a known distribution μ_e, where Ω is an abstract set of outcomes. We assume the distributions μ_e for different elements e have disjoint support. Every outcome ω has a fixed utility $x(\omega)$ for the principal and $y(\omega)$ for the agent.

The principal starts by committing to a mechanism, which consists of sets of valid signals $\Sigma_1, \ldots, \Sigma_k$ for each agent and a *conditional acceptance function* $g : \Sigma_1 \times \cdots \times \Sigma_k \to \Omega \cup \{\perp\}$ that maps the signals from all agents into a single winning outcome ω or a rejection of all outcomes $\perp$. After learning the mechanism, each agent observes their own type, the full set of outcomes sampled from their elements, and must select a signal to send to the principal. The principal receives these signals and uses g to transform them into a winning outcome or a rejection. If the winning outcome ω from agent i is verified by the principal, then the principal receives utility $x(\omega)$ and agent i receives utility $y(\omega)$, while all other agents receive nothing. Otherwise, everyone receives nothing.

Recall the example from the introduction of a non-profit that funds academic research through grants. Clearly, the non-profit is represented by the principal and the agents by the research groups. Each element represents a rough idea for a proposal with a-priori unknown random value, and sampling the element represents the process of investigating and evaluating this idea to determine its true utility for the principal and agent. The mechanism set by the principal consists of a description of what constitutes a valid proposal, as well as the specific details of the process they will use to determine the winner.

One challenge posed by this model is the complexity of analyzing agents' equilibrium strategies, how these equilibria are affected by the choice of mechanism, and how they affect the principal's utility. To this end, we also study a simplified model in which agents are assumed to act adversarially against the principal. More specifically, adversarial multi-agent delegation is the same as the strategic model above, except that we do not define agents' utilities and all agents instead aim to minimize the principal's expected utility subject to maintaining a positive probability of winning.

3 RESULTS

We start by showing that bounds on the delegation gap in the strategic case can be reduced to identical bounds in the adversarial case. This comes in two parts: the simple observation that delegating to strategic agents is at least as easy as delegating to adversarial agents (Lemma 3.2), and, less obviously, that for every adversarial instance within a central class of problems, there is an analogous strategic instance with identical behavior from the principal's perspective (Lemma 3.3). This may be somewhat surprising, since it implies that any increase in utility from delegating to multiple agents is not, in general, attributable to strategic competition between those agents. Rather, the principal's utility seems to increase simply as a consequence of the larger pool of acceptable options afforded by a larger pool of agents.

Turning our focus toward adversarial delegation, we find a harsh 1/2-approximation upper bound for any number of agents that carries over from the related single-agent model. This is due to the fact that the general form of the model allows for one agent to hold all elements that contribute non-zero expected utility to the principal, so, in essence, the principal must delegate to just that one agent.

However, moving beyond this impossibility, we show that when all agents have identical sets of elements, it is possible to achieve a competitive delegation gap of $1 - O\left(\frac{\ln k}{k}\right)$. This is done in two parts: first achieving this approximation for instances with only atomless distributions (Proposition 4.2), and then showing how to modify the strategy to deal with atoms (Theorem 4.3). Notably, this approximation uses only a very simple threshold mechanism, so it is easy to implement.

In the interest of demonstrating that other forms of symmetry also give competitive approximations, we consider also a randomized version of the adversarial model in which each element is given to a uniform random agent, and the principal's utility is measured in expectation over this randomness. For this model of *shuffled multi-agent delegation*, we show that the delegation gap has the same $1-O\left(\frac{\ln k}{k}\right)$ lower-bound (Proposition 4.5 and Theorem 4.6). Noting that a different symmetry assumption gives the same result, we conjecture that this is an instance of a more general phenomenon.

Finally, we show that the optimal delegation gap achievable with k agents in the agent-symmetric case is upper bounded by $1 - \Omega\left(\frac{1}{k}\right)$ (Theorem 4.7). We leave open for future work whether the gap between these upper and lower bounds can be closed.

4 OPEN QUESTIONS

We conclude by listing some open questions and interesting directions for future work:

(1) Are either of our bounds on the delegation gap tight?
(2) Could Myerson-type mechanisms (which we define analogously to Myerson's optimal auctions) be optimal for strategic multi-agent delegation?
(3) Does the delegation gap improve when considering the larger class of randomized mechanisms?
(4) Previous work [5, 7, 15] has considered single-agent delegation in the presence of "probing constraints" such as probing costs and combinatorial constraints on sets of probed elements. How do these affect multi-agent delegation?
(5) Previous work has also considered single-agent delegation in which the principal can accept sets of outcomes subject to a hard constraint. Can this be extended to multiple agents?
(6) Are there weaker or alternative forms of symmetry under which the principal can still achieve strictly better than a 1/2-approximation in general?

ACKNOWLEDGMENTS

This work is supported by the Air Force Office of Scientific Research under award number FA9550-24-1-0261, and the National Science Foundation under award number CCF-2009060. Any opinions, findings, and conclusions or recommendations expressed in this document are those of the authors and do not necessarily reflect the views of the United States Air Force or the National Science Foundation.

REFERENCES

[1] Marek Adamczyk and Michał Włodarczyk. 2018. Random order contention resolution schemes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 790–801.

[2] Ricardo Alonso, Isabelle Brocas, and Juan D Carrillo. 2014. Resource allocation in the brain. *Review of Economic Studies* 81, 2 (2014), 501–534.

[3] Ricardo Alonso and Niko Matouschek. 2008. Optimal delegation. *The Review of Economic Studies* 75, 1 (2008), 259–293.

[4] Mark Armstrong and John Vickers. 2010. A model of delegated project choice. *Econometrica* 78, 1 (2010), 213–244.

[5] Curtis Bechtel and Shaddin Dughmi. 2021. Delegated Stochastic Probing. In *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.

[6] Curtis Bechtel and Shaddin Dughmi. 2024. Efficient Multi-Agent Delegated Search. arXiv:2411.00181 [cs.GT] https://arxiv.org/abs/2411.00181

[7] Curtis Bechtel, Shaddin Dughmi, and Neel Patel. 2022. Delegated Pandora's box. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 666–693.

[8] Pirmin Braun, Niklas Hahn, Martin Hoefer, and Conrad Schecker. 2022. Delegated Online Search. *arXiv preprint arXiv:2203.01084* (2022).

[9] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. 2014. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM J. Comput.* 43, 6 (2014), 1831–1879.

[10] Paul Dütting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. 2017. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 540–551.

[11] Moran Feldman, Ola Svensson, and Rico Zenklusen. 2016. Online contention resolution schemes. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 1014–1033.

[12] Tan Gan, Ju Hu, and Xi Weng. 2023. Optimal contingent delegation. *Journal of Economic Theory* 208 (2023), 105597.

[13] Mohammadtaghi Hajiaghayi, Keivan Rezaei, and Suho Shin. 2023. Delegating to Multiple Agents. In *Proceedings of the 24th ACM Conference on Economics and Computation*. 1081–1126.

[14] Ali Khodabakhsh, Emmanouil Pountourakis, and Samuel Taggart. 2022. Simple Delegated Choice. (2022).

[15] Jon Kleinberg and Robert Kleinberg. 2018. Delegated search approximates efficient search. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. 287–302.

[16] Robert Kleinberg and Seth Matthew Weinberg. 2012. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM Symposium on Theory of Computing*. ACM, 123–136.

[17] Ulrich Krengel and Louis Sucheston. 1977. Semiamarts and finite values. *Bull. Amer. Math. Soc.* 83, 4 (1977), 745–747.

[18] Ulrich Krengel and Louis Sucheston. 1978. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces* 4 (1978), 197–266.

[19] Brendan Lucier. 2017. An economic view of prophet inequalities. *ACM SIGecom Exchanges* 16, 1 (2017), 24–47.