

# Algorithmic Delegated Choice: An Annotated Reading List

MOHAMMAD T. HAJIAGHAYI

University of Maryland

and

SUHO SHIN

University of Maryland

---

The problem of delegated choice has been of long interest in economics and recently on computer science. We overview a list of papers on delegated choice problem, from classic works to recent papers with algorithmic perspectives.

Categories and Subject Descriptors: B.6.3 [**Theory of computing**]: Algorithmic mechanism design

General Terms: Mechanism Design,

Additional Key Words and Phrases: Delegation

---

## 1. INTRODUCTION

The *delegated choice* problem is a fundamental model of principal-agent interaction with numerous real-world applications, capturing the tension when a decision maker (principal) delegates the role of decision making to an informed but self-interested agent. The model has its roots in classic economic theory introduced by [Holmström 1978], and has since evolved into a rich interdisciplinary area spanning economics, computer science, and operations research. It considers a scenario where the principal *cannot* commit to a contingent monetary transfer, and thus the principal needs to commit to a mechanism that specifies the characteristics of the agent's proposals that she is willing to accept. Such a model is particularly motivated by practical scenarios such as public regulators who can only accept or reject proposals from private sector, or investors relying on recommendations from financial advisors who are not contractually incentivized by the investor's returns.

In its canonical form, the principal must select an action from a discrete (or often continuous) set  $\Omega$ , where each action  $i \in \Omega$  has a pair of random utility values  $(X_i, Y_i)$ : one for the principal and one for the agent. Only the agent observes these values and proposes an action for selection, while the principal only knows the distribution from which these random values are drawn, introducing an information asymmetry. Once the agent observes the realizations, he sends a signal (*e.g.*, proposing an action) to the principal, who then makes the final decision. The agent seeks to maximize his own utility rather than the principal's, resulting in moral hazard. To mitigate this, the principal commits to a *screening mechanism* that selectively accepts the proposed action based on its values. For instance, an

---

Authors' addresses: hajiagha@umd.edu, suhoshin@umd.edu

eligible set  $E_i \subseteq \mathbb{R}^2$  can be announced to the agent so that the principal will only accept the proposed action  $i$  if  $(X_i, Y_i) \in E_i$ .

Recent work has shifted from classical results on existence and structure of optimal mechanisms to more algorithmic and computational perspectives. This reading list collects key papers in this growing literature, highlighting recent developments in various aspects of the delegated choice problem, including its connections to the prophet inequality [Krengel and Sucheston 1977], Pandora’s box [Weitzman 1978], and broader stochastic optimization problems.

Our goal is to introduce the core problem setup, present both classic and recent contributions, and illustrate how this line of work connects to various research areas in the EC community. As such, this article is neither exhaustive nor comprehensive; we hope it will serve as a useful starting point for readers to grasp the central ideas and emerging directions in delegated choice problem.

- (1) Bengt Rober Holmström. On incentives and control in organizations, *Ph.D. dissertation thesis*, 1978 & Bengt Rober Holmström, On the theory of delegation, *Bayesian Models in Economic Theory*, 1984.

Seminal work by [Holmström 1978; 1984] provides the *foundational framework* for delegation as an optimal screening mechanism. Instead of the discrete choice model described above, the principal delegates an optimization problem by choosing a screening set  $A' \subseteq A$  from which the agent selects an action  $a \in A'$  that yields payoffs  $v(a, \theta)$  and  $u(a, \theta)$  to the principal and agent, respectively, where  $\theta$  is a realized state observed only by the agent. Holmstrom characterizes conditions under which a single interval is optimal, and shows that the discretion given to the agent increases as his preferences align more closely with the principal’s.

- (2) Mark Armstrong, John Vickers. A model of delegated project choice, *Econometrica*, 2010.

[Armstrong and Vickers 2010] provides the first *discrete choice* model, which serves as a foundation for subsequent works. There are  $n$  available actions,<sup>1</sup> each with a type  $(u, v)$  drawn i.i.d. from a distribution, where the agent receives payoff  $v$  and the principal receives  $v + \alpha u$  for some  $\alpha \geq 0$ . Only the proposed action’s type is verifiable, so the principal’s goal is to design a screening mechanism over admissible types  $(u, v)$ , and the authors identify optimal mechanisms under specific payoff assumptions.

- (3) Jon Kleinberg, Robert Kleinberg. Delegated search approximates efficient search, *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, 2018.

[Kleinberg and Kleinberg 2018] studies the approximate efficiency of the optimal delegation mechanism with respect to the *first-best* benchmark, where the principal observes the utilities of all actions in hindsight and can choose what

<sup>1</sup>In fact, they consider a setting where  $n$  is also drawn from a known distribution.

she wants, building on the model of [Armstrong and Vickers 2010].<sup>2</sup> It considers precisely the problem setup described in the introduction, and uncovers a surprising connection to the classical *prophet inequality* [Krengel and Sucheston 1977]. Specifically, they prove an *equivalence* between the delegated choice problem and a version of prophet inequality with oblivious stopping rules. This leads to several delegation gap bounds, including a  $1/2$ -approximation with a threshold mechanism in the general case and a  $(1 - 1/e)$ -approximation in the i.i.d. case.

- (4) Kiarash Banihashem, Mohammad T. Hajiaghayi, Piotr Krysta, Suho Shin. Delegated choice with combinatorial constraints, *Proceedings of the 2025 ACM Conference on Economics and Computation (EC)*, 2025.

[Banihashem et al. 2025] considers a natural follow-up to [Kleinberg and Kleinberg 2018], asking to what extent the connection between delegated choice and prophet inequalities carries over. They study a multi-choice setting with a universe  $U$  of  $n$  actions and a family of feasible sets  $\mathcal{I} \subset 2^U$ , where the principal aims to select a set  $I \in \mathcal{I}$  to maximize  $X_I = \sum_{i \in I} X_i$ . Notably, they provide the first *provable separation* between the two problems by showing that delegated choice under downward-closed constraints allows constant-factor approximation, whereas prophet inequalities does not [Rubinstein 2016]; they further show that the correspondence between the two problems holds *if and only if* the constraint is a matroid.

- (5) Curtis Bechtel, Shaddin Dughmi. Delegated stochastic probing, *12th Innovations in Theoretical Computer Science (ITCS)*, 2021.

[Bechtel and Dughmi 2021] proposes a delegated stochastic probing problem where the agent probes actions under an outer constraint  $\mathcal{I}_{out}$  and proposes a feasible set under an inner constraint  $\mathcal{I}_{in}$ . Without an outer constraint or probing cost, their model coincides with [Banihashem et al. 2025], and they show that a greedy prophet inequality strategy against an almighty adversary, who observes every random bit of the algorithms and environments and decides a worst-case instance, can be implemented in this setting. This yields several immediate corollaries on matroid, matching, and knapsack constraints via greedy online contention resolution schemes [Feldman et al. 2016].<sup>3</sup>

- (6) Ali Kohdabakhsh, Emmanouil Pountourakis, Samuel Taggart. Simple delegated choice, *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2024.

The assumption that the proposed action's utilities are easily verifiable is admissible as misreporting could be verifiable by implementing the action or it

<sup>2</sup>They also consider a search problem where the agent undergoes a costly search process to seek solutions, with connection to Pandora's box problem [Weitzman 1978].

<sup>3</sup>We note here that this does not contradict the necessity of matroid results shown in [Banihashem et al. 2025], since [Bechtel and Dughmi 2021] guarantees only that the principal's utility exceeds the prophet's, not an exact equivalence between the agent's choice and the prophet inequality algorithm's decision as shown by [Banihashem et al. 2025].

may incur a reputation effect. On the other hand, this could often be impractical when the delegation happens as a one-off interaction or if it entails an expensive cost of verification such as the delegated decision of a governmental policy. [Khodabakhsh et al. 2024] considers a mechanism that *does not depend* on the utilities of the proposed values, but the principal completely rules actions in or out based on the distributional knowledge. They show that competing with the first-best benchmark is hopeless, and that the problem of computing the optimal mechanism is NP-hard, which is complemented by their  $1/3$  approximate deterministic mechanism.

- (7) Mohammad T. Hajiaghayi, Piotr Krysta, Mohammad Mahdavi, Suho Shin. Delegation with costly inspection, *Proceedings of the 2025 ACM Conference on Economics and Computation (EC)*, 2025.

[Hajiaghayi et al. 2025] directly addresses the verifiability assumption by allowing the principal to *inspect* the proposed action, and possibly others, at deterministic costs  $c_i$ , to verify utilities. In their extension of [Kleinberg and Kleinberg 2018], the agent may misreport if inspection is unlikely, and delegation itself incurs a fixed cost  $c_{del}$ . This model generalizes the Pandora’s box problem with nonobligatory inspection [Doval 2018], inheriting its NP-hardness [Fu et al. 2023; Beyhaghi and Cai 2023], and they show that while the first-best benchmark cannot be approximated, constant-factor approximate mechanisms exist in both costless and costly delegation settings when the cost of delegation is high or low.

- (8) Suho Shin, Keivan Rezaei, Mohammad T. Hajiaghayi. Delegating to multiple agents, *Proceedings of the 2023 ACM Conference on Economics and Computation (EC)*, 2023.

While the preceding works focus on Bayesian mechanisms, a few have explored prior-independent mechanisms in relaxed settings.<sup>4</sup> [Shin et al. 2023] studies a multi-agent delegated choice problem where each agent proposes an action, but only the selected agent receives nonzero utility, introducing competition that benefits the principal. They consider both Bayesian and prior-independent mechanisms and show that a constant-factor prior-independent mechanism exists in the complete information setting with symmetric agents, whereas the benefit of having multiple agents depends heavily on the agents’ information and symmetry.

- (9) Curtis Bechtel, Shaddin Dughmi. Efficient multi-agent delegated search, *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2025.

[Bechtel and Dughmi 2025] improves the results for Bayesian mechanisms under the incomplete information setting introduced by [Shin et al. 2023]. They achieve an approximation factor tending to 1 as the number of agents increases,

<sup>4</sup>In the standard setup with a single agent, one can easily see that no prior-independent mechanism can approximate the first-best benchmark.

when the agents have symmetric sets of actions that are not necessarily i.i.d. This strengthens the approximation factor and relaxes the constraints introduced by [Shin et al. 2023].

- (10) Mohammad T. Hajiaghayi, Mohammad Mahdavi, Keivan Rezaei, Suho Shin. Regret analysis of repeated delegated choice, *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI)*, 2024.

Distributional knowledge, in practice, is usually constructed from historical data. [Hajiaghayi et al. 2024] considers a variant with prior-independent mechanisms where the principal can repeatedly interact with the agent to construct estimates of the distributions. They frame their setup as a stochastic multi-armed bandit problem, and propose no-regret learning algorithms for myopic and farsighted agents under the Lipschitzness assumption of the utilities, using tools from bandits with perturbed outputs.

## REFERENCES

- ARMSTRONG, M. AND VICKERS, J. 2010. A model of delegated project choice. *Econometrica* 78, 1, 213–244.
- BANIHASHEM, K., HAJIAGHAYI, M. T., KRYSTA, P., AND SHIN, S. 2025. Delegation with costly inspection. *arXiv preprint arXiv:2506.07162*.
- BECHTEL, C. AND DUGHMI, S. 2021. Delegated stochastic probing. *Innovations in Theoretical Computer Science (ITCS)*.
- BECHTEL, C. AND DUGHMI, S. 2025. Efficient multi-agent delegated search. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems*. 2431–2433.
- BEYHAGHI, H. AND CAI, L. 2023. Pandora’s problem with nonobligatory inspection: Optimal structure and a ptas. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 803–816.
- DOVAL, L. 2018. Whether or not to open pandora’s box. *Journal of Economic Theory* 175, 127–158.
- FELDMAN, M., SVENSSON, O., AND ZENKLUSEN, R. 2016. Online contention resolution schemes. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 1014–1033.
- FU, H., LI, J., AND LIU, D. 2023. Pandora box problem with nonobligatory inspection: Hardness and approximation scheme. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 789–802.
- HAJIAGHAYI, M., MAHDAVI, M., REZAEI, K., AND SHIN, S. 2024. Regret analysis of repeated delegated choice. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 38. 9757–9764.
- HAJIAGHAYI, M. T., KRYSTA, P., MAHDAVI, M., AND SHIN, S. 2025. Delegation with costly inspection. *Available at SSRN 5284400*.
- HOLMSTRÖM, B. R. 1978. *On Incentives and Control in Organizations*. Stanford University.
- HOLMSTRÖM, B. R. 1984. On the theory of delegation. *Bayesian Models in Economic Theory*, 115–141.
- KHODABAKHSH, A., POUNTOURAKIS, E., AND TAGGART, S. 2024. Simple delegated choice. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 569–590.
- KLEINBERG, J. AND KLEINBERG, R. 2018. Delegated search approximates efficient search. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. 287–302.
- KRENGEL, U. AND SUCHESTON, L. 1977. Semiamarts and finite values.

- RUBINSTEIN, A. 2016. Beyond matroids: Secretary problem and prophet inequality with general constraints. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. 324–332.
- SHIN, S., REZAEI, K., AND HAJIAGHAYI, M. 2023. Delegating to multiple agents. In *Proceedings of the 24th ACM Conference on Economics and Computation*. 1081–1126.
- WEITZMAN, M. 1978. *Optimal search for the best alternative*. Vol. 78. Department of Energy.