Master Project

Risk-Sensitive Randomized Online Algorithms for Selection and Matching

Hossein Nekouyan Jazi
hossein.nekouyan@mail.utoronto.ca
hossein-nekouyan.github.io

Prof. Peyman Mohajerin Esfahani
p.mohajerinesfahani@utoronto.ca
mohajerinesfahani.github.io

Department of Mechanical & Industrial Engineering, University of Toronto

Motivation and Context

Online algorithms make sequential decisions without complete knowledge of future requests. They arise naturally in resource allocation, matching, admissions, hiring, advertising, marketplaces, and other applications in which decisions must be made immediately as opportunities arrive.

Randomization is one of the fundamental tools used in the design of online algorithms. In many online selection and matching problems, randomized algorithms achieve stronger theoretical guarantees than deterministic policies. These guarantees, however, are typically expressed in terms of **expected performance**. Consequently, two executions of the same randomized algorithm on exactly the same arrival sequence may produce very different outcomes.

Despite their strong theoretical guarantees, randomized online algorithms are not always attractive for practical deployment. An actual deployment observes only a single realization of the algorithm's internal randomness, and an algorithm with excellent expected performance may still produce an unfavorable outcome in that realization. This lack of reliability can be especially problematic in applications involving scarce resources, valuable assignments, or irreversible decisions, and may limit the practical appeal of randomized policies even when they outperform deterministic alternatives theoretically.

This creates an important distinction between average performance and reliability. A randomized algorithm may perform very well on average while occasionally producing a substantially worse realization. In applications involving scarce resources, valuable assignments, or irreversible selections, such unfavorable realizations may have significant operational consequences.

The objective of this project is to develop a **risk-sensitive perspective on randomized online algorithms**. Rather than evaluating an algorithm only according to its expected reward or competitive ratio, the project will study the distribution of its realized performance and investigate how downside risk should be measured, understood, and controlled. The initial focus will be on **online selection and online matching**.

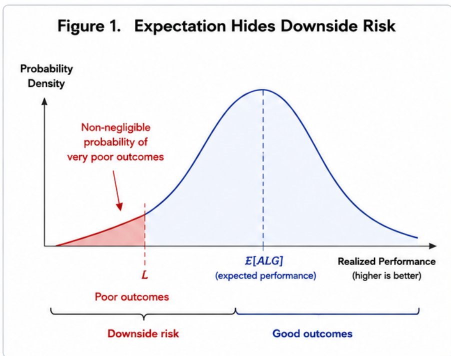

Figure 1: **Expected performance can hide downside risk.**

Source of Risk

The project focuses specifically on uncertainty generated by the **internal randomization of the algorithm**. The underlying sequence of arrivals may be deterministic, adversarial, stochastic, or taken directly from a real-world dataset. Once an arrival sequence is fixed, repeatedly executing a randomized algorithm with different random seeds can produce different decisions and therefore different final outcomes.

The central object of study is this variation in realized performance. Environmental uncertainty may also be present, but the primary objective is to isolate and understand the additional risk introduced by the randomized decision rule itself.

Central research question: Can randomized online algorithms retain their performance advantages while providing meaningful protection against unfavorable realizations of their own internal randomness?

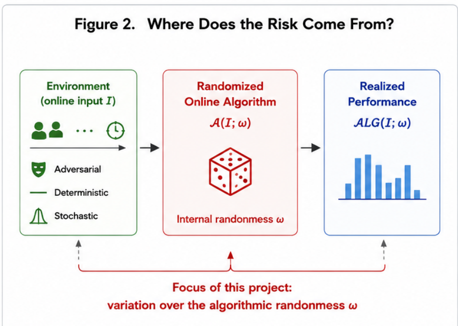

Figure 2: **Source of risk.** We isolate variation caused by the algorithm's internal randomness.

Project Objectives and Research Tasks

The project has two main components. The first component identifies a concrete online matching or selection setting, uses real data to expose the risk created by algorithmic randomization, and determines which risk measure best captures that behavior. The second component uses this understanding to design randomized online algorithms with explicit downside-risk control.

1. Data + Risk Measurement	→	2. Risk-Sensitive Algorithm Design
Online matching/selection, real data, and risk measures		Randomized policies with controlled downside risk

Task 1: Study the risk-sensitivity of randomized online selection and matching algorithms in practical applications. Online selection and matching capture a broad range of sequential resource-allocation applications, including marketplaces, advertising, hiring and admissions, and task-to-resource assignment. These problems also have a rich literature of randomized algorithms whose performance guarantees can outperform those achievable by deterministic policies [4, 5, 6, 7, 8, 9, 10]. Since this advantage is obtained through randomization, understanding the associated downside risk is important for assessing their practical reliability.

The student will identify a concrete selection or matching application together with an appropriate real-world dataset and use it to study the risk-sensitive behavior of randomized online algorithms. Based on the observed behavior, the student will then investigate which notion of risk—such as **CVaR, VaR, lower-tail quantiles, threshold-violation probability, or entropic risk**—provides the most meaningful characterization of undesirable realizations [2, 3]. The outcome of this task will be a practically motivated problem setting and risk measure that provide the foundation for the algorithm-design stage.

Task 2: Design risk-sensitive randomized online algorithms. Building on the problem setting and risk measure identified in Task 1, the second objective is to design randomized online algorithms that reduce downside risk while preserving the performance advantages of randomization. The project will investigate structured or correlated randomization, risk-aware decision rules, or other problem-specific mechanisms. The resulting algorithms will be evaluated theoretically and empirically to characterize the trade-off between classical performance guarantees and downside protection.

Expected Outcomes

A successful project may produce a concrete real-world dataset and experimental framework for evaluating the risk of randomized online matching or selection algorithms; an empirical demonstration that strong expected performance can conceal significant downside risk; a principled comparison of risk measures for randomized online algorithms; and a new risk-sensitive randomized algorithm with theoretical and empirical guarantees describing the trade-off between classical performance and downside protection.

Student Background

A strong background in algorithms, probability, optimization, or operations research is desirable, together with mathematical maturity and programming skills in Python or a similar language. Prior exposure to online algorithms, competitive analysis, online matching, prophet inequalities, or risk measures is useful but not required.

Starting References

References

[1] H. Nekouyan Jazi et al., *Risk-Sensitive Online Selection with Bounded Adaptivity*, arXiv:2512.02427.

[2] N. Christianson, B. Sun, S. Low, and A. Wierman, *Risk-Sensitive Online Algorithms*, Conference on Learning Theory (COLT), 2024.

[3] T. Soma and Y. Yoshida, *Online Risk-Averse Submodular Maximization*, Annals of Operations Research, 320:393–414, 2023.

[4] S. R. Balseiro, H. Lu, and V. Mirrokni, *The Best of Many Worlds: Dual Mirror Descent for Online Allocation Problems*, Operations Research, 71(1):101–119, 2023.

[5] C. Papadimitriou, T. Pollner, A. Saberi, and D. Wajc, *Online Stochastic Max-Weight Bipartite Matching: Beyond Prophet Inequalities*, Mathematics of Operations Research, 49(3):1607–1628, 2024.

[6] S. Delong, A. Farhadi, R. Niazadeh, B. Sivan, and R. Udwani, *Online Bipartite Matching with Reusable Resources*, Mathematics of Operations Research, 49(3):1825–1854, 2024.

[7] B. Kalyanasundaram, K. Pruhs, and C. Stein, *A Randomized Algorithm for Online Metric b-Matching*, Operations Research Letters, 2023.

[8] B. Peng and Z. G. Tang, *Order Selection Prophet Inequality: From Threshold Optimization to Arrival Time Design*, Proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS), 2022.

[9] M. Hoefer, P. Krysta, and B. Sina, *Threshold Testing and Semi-Online Prophet Inequalities*, Proceedings of the European Symposium on Algorithms (ESA), 2023.

[10] K. Banihashem, A. Eden, M. Feldman, A. Fiat, and K. Goldner, *Power of Posted-Price Mechanisms for Prophet Inequalities*, 2024.