Risk-Sensitive Online Algorithms (Extended Abstract)

Nicolas Christianson NCHRISTIANSON@CALTECH.EDU
California Institute of Technology

Bo Sun BO.SUN@UWATERLOO.CA
University of Waterloo

Steven Low SLOW@CALTECH.EDU
California Institute of Technology

Adam Wierman ADAMW@CALTECH.EDU
California Institute of Technology

Editors: Shipra Agrawal and Aaron Roth

Decision-makers in many domains are sensitive to risks of certain sizes or likelihoods, motivating the incorporation of *risk sensitivity* into decision-making objectives. While such risk sensitivity has been extensively studied in online learning problems (e.g., Even-Dar et al. (2006)), very little work has examined the impact of risk-sensitive objectives on algorithm design for competitive online optimization problems like ski rental, online search, and metrical task systems; closest to this direction is the recent paper of Dinitz et al. (2024) on designing ski rental algorithms with constrained tail risk. In this work, we begin to fill this gap, considering the design of algorithms for competitive online optimization problems with risk-sensitive objectives. In particular, we introduce a variant of the standard competitive ratio metric, the CVaR_δ-competitive ratio (δ-CR), using the conditional value-at-risk of an algorithm's cost (Acerbi and Tasche, 2002), which measures the expectation of the $(1-\delta)$-fraction of worst outcomes against the offline optimal cost. Using this metric, we design algorithms for three prototypical online optimization problems: continuous-time ski rental, discrete-time ski rental, and one-max search. We find that the structure of the optimal δ-CR and algorithm varies significantly across problems:

- The optimal algorithm for continuous-time ski rental arises as the solution to a delay differential equation, and achieves δ-CR $2 - 2^{-\Theta(\frac{1}{1-\delta})}$ as $\delta \uparrow 1$.
- For discrete-time ski rental with buying cost B, the classic randomized algorithm is optimal for $\delta = \mathcal{O}(\frac{1}{B})$, and there is an abrupt phase transition at $\delta = 1 - \Theta(\frac{1}{\log B})$, after which the classic deterministic strategy is optimal.
- One-max search also exhibits a sharp phase transition at $\delta = \frac{1}{2}$, after which the classic deterministic strategy is optimal. Moreover, we propose an algorithm arising as the solution to a delay differential equation that is asymptotically optimal as $\delta \downarrow 0$.

Our results imply limits on the power of randomness to improve performance in certain risk-sensitive settings and suggest avenues for the design of risk-sensitive algorithms in more general online problems.[1]

Acknowledgments

The authors acknowledge support from an NSF Graduate Research Fellowship (DGE-2139433), NSF Grants CNS-2146814, CPS-2136197, CNS-2106403, and NGSDI-2105648, the Resnick Sustainability Institute, and a Caltech S2I Grant.

[1] Extended abstract. Full version appears as [arXiv:2405.09859, v2].

References

Carlo Acerbi and Dirk Tasche. On the coherence of expected shortfall. *Journal of Banking & Finance*, 26 (7):1487–1503, July 2002. ISSN 0378-4266. doi: 10.1016/S0378-4266(02)00283-2.

Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii. Controlling Tail Risk in Online Ski-Rental. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4247–4263, 2024. doi: 10.1137/1.9781611977912.147.

Eyal Even-Dar, Michael Kearns, and Jennifer Wortman. Risk-Sensitive Online Learning. In José L. Balcázar, Philip M. Long, and Frank Stephan, editors, *Algorithmic Learning Theory*, pages 199–213, Berlin, Heidelberg, 2006. Springer. ISBN 978-3-540-46650-5. doi: 10.1007/11894841_18.