On Multiagent Online Problems with Predictions

Extended Abstract

Gabriel Istrate
University of Bucharest
Bucharest, Romania
gabriel.istrate@unibuc.ro

Cosmin Bonchiş
West University of Timişoara
Timişoara, Romania
cosmin.bonchis@e-uvt.ro

Victor Bogdan
West University of Timişoara
Timişoara, Romania
victor.bogdan@e-uvt.ro

ABSTRACT

We study the power of (competitive) algorithms with predictions in a multiagent setting. We introduce a *two predictor framework*, that assumes that agents use one predictor for their future (self) behavior, and one for the behavior of the other players. The main problem we are concerned with is understanding what are the best competitive ratios that can be achieved by employing such predictors, under various assumptions on predictor quality. As an illustration of our framework, we introduce and analyze a *multiagent version of the ski-rental problem.* In this problem agents can collaborate by pooling resources to get a *group license* for some asset. If the license price is not met then agents have to rent the asset individually for the day at a unit price. Otherwise the license becomes available forever to everyone at no extra cost. We show that blindly following the predictors is not robust to mispredictions of future behavior (even when the others-predictor is perfect). We propose (and benchmark) a more robust meta-algorithm.

KEYWORDS

algorithms with predictions; multiagent ski-rental; metaalgorithms

ACM Reference Format:
Gabriel Istrate, Cosmin Bonchiş, and Victor Bogdan. 2026. On Multiagent Online Problems with Predictions: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/WZRJ7594

1 INTRODUCTION

Recent advances in *algorithms with predictions*, that is (online) algorithms enhanced by (possibly machine-learned) advice, have led to improved performance in many classical single-agent problems such as caching, ski-rental, and scheduling. Yet, virtually all existing work (at least on the competitive analysis of online algorithms) has focused on single decision-makers, leaving open the question of how predictions can be leveraged in *multiagent* environments, where decisions are interdependent and outcomes are shaped by collective behavior. This paper takes a first step in that direction. **Important:** A full version is available [29] on arXiv.org. References included in this extended abstract are those from the full version.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/WZRJ7594

Note: Our present work is decision-theoretic, rather than game-theoretic: Agents do **not** consider (when computing their best response) what may be rational for *other* agents to do. They simply learn about the actions of the other agents through a predictor (which they view as a black box that they may or may not trust). A companion paper [30] takes a game-theoretic perspective, developing a formal notion of *equilibria with predictions.*

2 CONTRIBUTIONS

Our main conceptual and technical contributions are as follows:

- We introduce a framework for *competitive online problems with predictions*, combining ideas from competitive analysis and multi-agent decision-making. Our framework assumes that *each agent is equipped with two predictors:* a *self-predictor* for its own future activity, and a *others-predictor* for the behavior of the other agents, capturing the asymmetric uncertainty agents face about themselves and their peers.
- We use this framework to analyze a **multiagent variant of the classical ski-rental problem**, in which agents can cooperate to buy a group license for a resource, or rent it individually:

Definition 1 (Multiagent Ski Rental)**.** n agents are initially active but may become inactive (once an agent becomes inactive it will be inactive forever.) Active agents need a resource for their daily activity. Each day, active agents have the option to (individually) rent the resource, at a cost of \$1/day. They can also cooperate in order to buy a group license that will cost $B > 1$ dollars. For this, each agent i may pledge some amount $w_i > 0$ or *refrain from pledging* (equivalently, pledge 0. We assume that both pledges w_i and the price B are integers.) If the total sum pledged is at least B then the group license is acquired and the use of the resource becomes free for all remaining active agents from that moment on. Otherwise, the pledges are nullified. Instead, every active agent must (individually) rent the resource for the day. We will assume that when the license is overpledged agents will pay their pledged sums. Agents are strategic, in that they care about their overall costs. They are faced, on the other hand, with deep uncertainty concerning the number of days they will be active. So they choose instead to minimize their competitive ratio. We assume that the others-predictor for an agent i only predicts the total amount pledged by other agents.

We aim to answer the following questions about this problem, from a single agent perspective:

(a). what is the best possible competitive ratio of a predictionless algorithm; how much can one/both predictor(s) improve it?
(b). how do algorithms achieving such optimal improvements look like, and how robust are they to mispredictions?

Table 1: Summary of optimal competitive ratios (and algorithms realizing them).

Self-predictor	Others Predictor	Competitive Ratio	Optimal Algorithm(s)
none/pessimal	none/pessimal	$B+1$	described in Thm. 1 [29]
optimal	none/pessimal	$B+1$	described in Thm. 2 [29]
none/pessimal	optimal	depends on other players' behavior (explicit formula in Thm. 3 [29])	described in Thm. 3 [29]
optimal	optimal	1	Algorithm 1, [29]

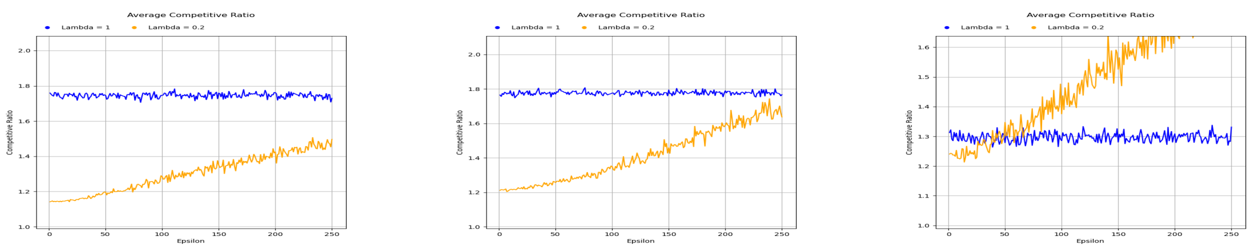

Figure 1: Average competitive ratio for the meta-algorithm. (a). z=0 (b). z=0.5 (c). z=1. In each case 1000 samples were used. Two values of parameter λ are reported: $\lambda = 1$ (blue) and $\lambda = 0.2$ (orange).

- We give complete answers to questions (a) and (b). The results are summarized in Table 1, and are intuitively described as follows:
  - the optimal predictionless algorithm (or with pessimal self/others-predictions) rents on day 1 and pledges the full amount, B, on day 2. It is $(B+1)$-competitive.
  - **when the others-predictor is pessimal, self-predictors don't help:** no algorithm with self-predictions (even perfect ones) can be better than $B+1$-competitive (although perfect self-predictors help more algorithms reach this optimum).
  - for perfect others- predictions but pessimal self-predictions, *the optimal competitive ratio depends on the actions of other players* (Theorem 3 [29]). We characterize all the optimal algorithms.
- We show that **blindly trusting predictions is not robust**: even with perfect predictions of others, errors in self-prediction can lead to arbitrarily poor competitive ratio.
- We design a **meta-algorithm**, parameterized by trust levels $\lambda, \mu \in [0,1]$ in the self and other-predictor, respectively, that "interpolates" between predictor-reliant and predictor-free strategies, achieving simultaneously the optimal competitive ratios in all the four limiting cases from Table 1. We prove analytical bounds on its performance (Theorem 6 in [29]), and benchmark it under random aggregate total bids (see Section 3 below).
- We show (Example 2 in [29]) that *a specialization of our meta-algorithm differs minimally from the meta-algorithm for single-agent ski-rental with predictions from [34], and displays an identical robustness/consistency tradeoff.*

3 EXPERIMENTAL EVALUATION

We experimentally investigated the performance of our metaalgorithm. Since this paper does not assume rational agents, we used random total bids of other players. This made the agent face a varying "price" between 0 and B, representing the amount the agent needs to pledge to ensure the buying of the group license.

The setting mimics the one in [34] as much as possible; in particular, we also measure *average* rather than worst-case competitive ratio. As with the experiments in [34], we assume that total price $B = 100$, the number of days T of the given agent is randomly chosen in the interval $[1, 4B]$, the predicted active time $\widehat{T} = T + \epsilon$, where ϵ is normally distributed with average 0 and standard deviation σ. "Prices" fluctuate randomly in $[B - \lfloor zB \rfloor, B]$ for a parameter $0 \leq z \leq 1$. We display three values for z: $z = 0$ (classical ski rental), $z = 0.5$, and $z = 0.1$. The algorithm performs reasonably well: Figures 1 (a,b,c) shows this in three scenarios. By Theorem 1 in [29], the baseline competitive ratio to compare the plots against is the competitive ratio for the $\lambda = \mu = 0$ case, which is $B + 1 = 101$. The conclusion is that **results are qualitatively similar with those from [34] for all values of z, and acceptable in all cases.**

4 FURTHER DIRECTIONS

Our work leaves many issues open (see the full paper for details). For instance, questions of interaction between prediction and optimization similar to those in Newcomb's problem have to be clarified: our results implicitly used *causal decision theory* [51], but this is not the only option [2, 23]. There exists, on the other hand, a huge, multi-discipline literature that deals with learning game-theoretic equilibria through repeated interactions (e.g. [33, 53**?**]). Applying such results to the setting with predictions and, generally, allowing predictions to be *adaptive* is an open problem. Note also that [21] considers a model in which programs P are benchmarked *against themselves*, i.e. *everyone uses* P. Optimal programs P are examples of *Kantian equilibria* [27, 44]. How to do define such equilibria *with predictions* is interesting. Finally, a promising direction is related to the Minority game [17], a model from the Complex System literature with a similar inteplay between collective decision and individual prediction.

ACKNOWLEDGMENTS

Partially supported by the "Romanian Hub for Artificial Intelligence - HRIA", Smart Growth, Digitization and Financial Instruments Program, 2021-2027, MySMIS no. 334906, and a grant of the Romanian Ministry of Education and Research, CCCDI - UEFISCDI, number PN-IV-P6-6.1-CoEx-2024-0214, within PNCDI IV.

REFERENCES

[1] Priyank Agrawal, Eric Balkanski, Vasilis Gkatzelis, Tingting Ou, and Xizhi Tan. 2022. Learning-Augmented Mechanism Design: Leveraging Predictions for Facility Location. *Proceedings of EC'22* (2022), 497–528.
[2] Arif Ahmed. 2021. *Evidential decision theory.* Cambridge University Press.
[3] Kareem Amin, Travis Dick, Mikhail Khodak, and Sergei Vassilvitskii. 2022. Private Algorithms with Private Predictions. *arXiv preprint arXiv:2210.11222* (2022).
[4] Keerti Anand, Rong Ge, Amit Kumar, and Debmalya Panigrahi. 2022. Online Algorithms with Multiple Predictions. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, 582–598. https://proceedings.mlr.press/v162/anand22a.html
[5] Keerti Anand, Rong Ge, and Debmalya Panigrahi. 2020. Customizing ML predictions for online algorithms. In *International Conference on Machine Learning*. PMLR, 303–313.
[6] Spyros Angelopoulos, Christoph Dürr, Shendan Jin, Shahin Kamali, and Marc Renault. 2020. Online Computation with Untrusted Advice. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*.
[7] Itai Ashlagi, Brendan Lucier, and Moshe Tennenholtz. 2013. Equilibria of online scheduling algorithms. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 27. 67–73.
[8] Eric Balkanski, Vasilis Gkatzelis, Xizhi Tan, and Cherlin Zhu. 2023. *Online Mechanism Design with Predictions.* Technical Report. arXiv.org/cs.GT/2310.02879. arXiv:2310.02879 [cs.GT]
[9] Soumya Banerjee. 2020. Improving online rent-or-buy algorithms with sequential decision making and ML predictions. *Advances in Neural Information Processing Systems* 33 (2020), 21072–21080.
[10] Magnus Berg, Joan Boyar, Lene M Favrholdt, and Kim S Larsen. 2024. Complexity Classes for Online Problems with and without Predictions. *arXiv preprint arXiv:2406.18265* (2024).
[11] Allan Borodin and Ran El-Yaniv. 2005. *Online computation and competitive analysis.* Cambridge University Press.
[12] Richmond Campbell and Lanning Sowden. 1985. *Paradoxes of rationality and cooperation: Prisoner's dilemma and Newcomb's problem.* U.B.C. Press.
[13] Clément L. Canonne, Kenny Chen, and Julián Mestre. 2025. *With a Little Help From My Friends: Exploiting Probability Distribution Advice in Algorithm Design.* Technical Report 2505.04949. arXiv.org.
[14] Ioannis Caragiannis and Georgios Kalantzis. 2024. Randomized learning-augmented auctions with revenue guarantees. *arXiv preprint arXiv:2401.13384*.
[15] Nicolo Cesa-Bianchi and Gábor Lugosi. 2006. *Prediction, learning, and games.* Cambridge University Press.
[16] Justin Y Chen, Talya Eden, Piotr Indyk, Honghao Lin, Shyam Narayanan, Ronitt Rubinfeld, Sandeep Silwal, Tal Wagner, David Woodruff, and Michael Zhang. 2022. Triangle and Four Cycle Counting with Predictions in Graph Streams. In *International Conference on Learning Representations.* https://openreview.net/forum?id=8in_5gN9I0
[17] M. Marsili D. Challet and Y-C. Zhang. 2004. *Minority Games: Interacting Agents in Financial Markets.* Oxford University Press.
[18] Ilias Diakonikolas, Vasilis Kontonis, Christos Tzamos, Ali Vakilian, and Nikos Zarifis. 2021. Learning Online Algorithms with Distributional Advice. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 2687–2696. https://proceedings.mlr.press/v139/diakonikolas21a.html
[19] Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii. 2022. Algorithms with Prediction Portfolios. In *Advances in Neural Information Processing Systems*, Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (Eds.). https://openreview.net/forum?id=wwWCZ7sER_C
[20] Marina Drygala, Sai Ganesh Nagarajan, and Ola Svensson. 2023. Online Algorithms with Costly Predictions. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 206)*, Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent (Eds.). PMLR, 8078–8101. https://proceedings.mlr.press/v206/drygala23a.html
[21] Roee Engelberg and Joseph Seffi Naor. 2016. Equilibria in online games. *SIAM J. Comput.* 45, 2 (2016), 232–267.
[22] Paolo Ferragina and Giorgio Vinciguerra. 2020. Learned data structures. In *Recent Trends in Learning From Data*. Springer, 5–41.
[23] Ghislain Fourny. 2020. Perfect Prediction in normal form: Superrational thinking extended to non-symmetric games. *Journal of Math. Psychology* 96 (2020), 102332.
[24] Vasilis Gkatzelis, Kostas Kollias, Alkmini Sgouritsa, and Xizhi Tan. 2022. Improved Price of Anarchy via Predictions. In *Proceedings of EC'22*. 529–557.
[25] Sreenivas Gollapudi and Debmalya Panigrahi. 2019. Online algorithms for rent-or-buy with expert advice. In *International Conference on Machine Learning*. PMLR, 2319–2327.
[26] Moritz Hardt and Celestine Mendler-Dünner. 2025. Performative prediction: Past and future. *Statist. Sci.* 40, 3 (2025), 417–436.
[27] Gabriel Istrate. 2021. Game-Theoretic Models of Moral and Other-Regarding Agents. *Proceedings of the Eighteenth Conference on Theoretical Aspects of Rationality and Knowledge (TARK'2021), Electronic Proceedings in Theoretical Computer Science 335* (2021), 213–228.
[28] Gabriel Istrate and Cosmin Bonchiş. 2022. Mechanism Design With Predictions for Obnoxious Facility Location. *arXiv preprint arXiv:2212.09521* (2022).
[29] Gabriel Istrate, Cosmin Bonchiş, and Victor Bogdan. 2025. *On Multiagent Problems with Prediction.* Technical Report arXiv:2507.12486. arXiv.org.
[30] Gabriel Istrate, Cosmin Bonchiş, and Victor Bogdan. 2024. Equilibria in multiagent online problems with predictions. *arXiv preprint arXiv:2405.11873v2* (2024).
[31] Tanqiu Jiang, Yi Li, Honghao Lin, Yisong Ruan, and David P. Woodruff. 2020. Learning-Augmented Data Stream Algorithms. In *International Conference on Learning Representations.* https://openreview.net/forum?id=HyxJ1xBYDH
[32] Neil F Johnson, Paul Jefferies, and Pak Ming Hui. 2003. *Financial market complexity.* Oxford University Press.
[33] Ehud Kalai and Ehud Lehrer. 1993. Rational learning leads to Nash equilibrium. *Econometrica: Journal of the Econometric Society* (1993), 1019–1045.
[34] Ravi Kumar, Manish Purohit, and Zoya Svitkina. 2018. Improving online algorithms via ML predictions. *Advances in Neural Information Processing Systems* 31.
[35] Honghao Lin, Tian Luo, and David Woodruff. 2022. Learning Augmented Binary Search Trees. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, 13431–13440. https://proceedings.mlr.press/v162/lin22f.html
[36] Alexander Lindermayr and Nicole Megow. 2022. Algorithms With Predictions Webpage. https://algorithms-with-predictions.github.io accessed July 2025.
[37] Jessica Maghakian, Russell Lee, Mohammad Hajiesmaili, Jian Li, Ramesh K. Sitaraman, and Zhenhua Liu. 2023. Applied Online Algorithms with Heterogeneous Predictors. In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA (Proceedings of Machine Learning Research, Vol. 202)*, Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (Eds.). PMLR, 23484–23497. https://proceedings.mlr.press/v202/maghakian23a.html
[38] Li-Yuan Meng, Chang-Jun Wang, and Qing-Jie Ye. 2025. Group Rent-or-Buy: the Benefits of Having Grouped Consumers: L.-Y. Meng et al. *Journal of the Operations Research Society of China* (2025), 1–18.
[39] Michael Mitzenmacher. 2018. A model for learned Bloom filters and optimizing by sandwiching. *Advances in Neural Information Processing Systems* 31 (2018).
[40] M. Mitzenmacher and S. Vassilvitskii. 2021. Algorithms with Predictions. Chapter 29 in [45].
[41] Adhyyan Narang, Evan Faulkner, Dmitriy Drusvyatskiy, Maryam Fazel, and Lillian J Ratliff. 2023. Multiplayer performative prediction: Learning in decision-dependent games. *Journal of Machine Learning Research* 24, 202 (2023), 1–56.
[42] Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. 2020. Performative prediction. In *International Conference on Machine Learning*. PMLR, 7599–7609.
[43] Georgios Piliouras and Fang-Yi Yu. 2023. Multi-agent performative prediction: From global stability and optimality to chaos. In *Proceedings of the 24th ACM Conference on Economics and Computation.* 1047–1074.
[44] John E Roemer. 2019. *How We Cooperate: A Theory of Kantian Optimization.* Yale University Press.
[45] Tim Roughgarden. 2021. *Beyond the worst-case analysis of algorithms.* Cambridge University Press.
[46] Judy Hanwen Shen, Ellen Vitercik, and Anders Wikum. 2025. *Algorithms with Calibrated Machine Learning Predictions.* Technical Report 2502.02861. arXiv.org.
[47] Bo Sun, Jerry Huang, Nicolas Christianson, Mohammad Hajiesmaili, and Adam Wierman. 2023. Online Algorithms with Uncertainty-Quantified Predictions. *arXiv preprint arXiv:2310.11558* (2023).
[48] Rahul Vaze. 2023. *Online Algorithms.* Cambridge University Press.
[49] Xuchuang Wang, Bo Sun, Hedyeh Beyhaghi, John Lui, Mohammad Hajiesmaili, and Adam Wierman. 2025. Competitive Algorithms for Multi-Agent Ski-Rental Problems. *arXiv preprint arXiv:2507.15727* (2025).
[50] Alexander Wei and Fred Zhang. 2020. Optimal robustness-consistency trade-offs for learning-augmented online algorithms. *Advances in Neural Information Processing Systems* 33 (2020), 8042–8053.
[51] Paul Weirich. 2008. Causal decision theory. *Cambridge University Press* (2008).
[52] Chenyang Xu and Pinyan Lu. 2022. Mechanism Design with Predictions. *Proceedings of IJCAI'22* (2022), 571–577.
[53] H Peyton Young. 2004. *Strategic learning and its limits.* Oxford University Press.