

Survival of the Strictest: Stable and Unstable Equilibria under Regularized Learning with Partial Information

Angeliki Giannou

National Technical University of Athens

Emmanouil-Vasileios Vlatakis-Gkaragkounis

Columbia University

Panayotis Mertikopoulos

Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, 38000 Grenoble, France & Criteo AI Lab

GIANNOUANGELIKI@GMAIL.COM

EMVLATAKIS@CS.COLUMBIA.COM

PANAYOTIS.MERTIKOPOULOS@IMAG.FR

Editors: Mikhail Belkin and Samory Kpotufe

Abstract

In this paper, we examine the Nash equilibrium convergence properties of no-regret learning in general N -player games. For concreteness, we focus on the archetypal “follow the regularized leader” (FTRL) family of algorithms, and we consider the full spectrum of uncertainty that the players may encounter – from noisy, oracle-based feedback, to bandit, payoff-based information. In this general context, we establish a comprehensive equivalence between the stability of a Nash equilibrium and its support: *a Nash equilibrium is stable and attracting with arbitrarily high probability if and only if it is strict* (i.e., each equilibrium strategy has a unique best response). This equivalence extends existing continuous-time versions of the “folk theorem” of evolutionary game theory to a bona fide algorithmic learning setting, and it provides a clear refinement criterion for the prediction of the day-to-day behavior of no-regret learning in games.

In more detail, we address the following questions: *Is there a class of Nash equilibria that consistently attract no-regret processes?* Conversely, *are all Nash equilibria equally likely to emerge as outcomes of a no-regret learning process?* To address them in a general setting, we focus on the “follow the regularized leader” (FTRL) algorithm and we prove the following result

x^* is a strict Nash equilibrium $\iff x^*$ is stochastically asymptotically stable under FTRL

Formally, we get the following precise statements for a range of specific feedback models:

Theorem 1. *Let $x^* \in \mathcal{X}$ be a strict Nash equilibrium of the game under study. If FTRL is run with inexact payoff vector estimates with vanishing bias and moderately increasing variance, x^* is stochastically asymptotically stable.*

Theorem 2. *Let x^* be a mixed Nash equilibrium of a generic game. If FTRL is run with inexact payoff vector estimates with vanishing bias and moderately increasing variance, x^* is not stochastically asymptotically stable.*

These results – and, in particular, the implications for the bandit case – provide a learning justification to the abundance of arguments that have been made in the refinement literature against selecting mixed Nash equilibria [3, 6] and strengthen existing results on continuous-time game dynamics [1, 2, 5], sometimes referred to as the “folk theorem” of evolutionary game theory [4].

Keywords: No-regret learning, Nash Equilibrium, follow the regularized leader, asymptotic stability.

*Extended abstract. Full version appears as [arxiv:2101.04667, v3].

Acknowledgments

This research was partially supported by the COST Action CA16228 “European Network for Game Theory” (GAMENET) and the Onassis Foundation undr Scholarship ID: F ZN 010-1/2017-2018. P. Mertikopoulos is also grateful for financial support by the French National Research Agency (ANR) in the framework of the “Investissements d’avenir” program (ANR-15-IDEX-02), the LabEx PERSYVAL (ANR-11-LABX-0025-01), MIAI@Grenoble Alpes (ANR-19-P3IA-0003), and the grants ORACLESS (ANR-16-CE33-0004) and ALIAS (ANR-19-CE48-0018-01). E.V. Vlatakis-Gkaragkounis is grateful to be supported by NSF grants CCF-1703925, CCF1763970, CCF-1814873, CCF-1563155, and by the Simons Collaboration on Algorithms and Geometry.

References

- [1] Pierre Coucheney, Bruno Gaujal, and Panayotis Mertikopoulos. Penalty-regulated dynamics and robust learning procedures in games. *Mathematics of Operations Research*, 40(3):611–633, August 2015.
- [2] Lampros Flokas, Emmanouil Vasileios Vlatakis-Gkaragkounis, Thanasis Lianeas, Panayotis Mertikopoulos, and Georgios Piliouras. No-regret learning and mixed Nash equilibria: They do not mix. In *NeurIPS ’20: Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020.
- [3] Drew Fudenberg and Jean Tirole. *Game Theory*. The MIT Press, 1991.
- [4] Josef Hofbauer and Karl Sigmund. Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 40(4):479–519, July 2003.
- [5] Panayotis Mertikopoulos and William H. Sandholm. Learning in games via reinforcement and regularization. *Mathematics of Operations Research*, 41(4):1297–1324, November 2016.
- [6] Eric van Damme. *Stability and perfection of Nash equilibria*. Springer-Verlag, Berlin, 1987.