Learning and Computation of Φ-Equilibria at the Frontier of Tractability

BRIAN HU ZHANG*, Carnegie Mellon University, USA
IOANNIS ANAGNOSTIDES*, Carnegie Mellon University, USA
EMANUEL TEWOLDE, Carnegie Mellon University, USA
RATIP EMIN BERKER, Carnegie Mellon University, USA
GABRIELE FARINA, Massachusetts Institute of Technology, USA
VINCENT CONITZER, Carnegie Mellon University, USA and University of Oxford, UK
TUOMAS SANDHOLM†, Carnegie Mellon University, USA

Φ-*equilibria*—and the associated notion of Φ-*regret*—are a powerful and flexible framework at the heart of online learning and game theory, whereby enriching the set of deviations Φ begets stronger notions of rationality. Recently, Daskalakis, Farina, Fishelson, Pipis, and Schneider (STOC '24) settled the existence of efficient algorithms when Φ contains only linear maps under a general, d-dimensional convex constraint set $\mathcal{X}$. In this paper, we significantly extend their work by resolving the case where Φ is k-dimensional; degree-ℓ polynomials constitute a canonical such example with $k = d^{O(\ell)}$. In particular, we obtain two main positive results: i) a $\text{poly}(n, d, k, \log(1/\epsilon))$-time algorithm for computing ϵ-approximate Φ-equilibria in n-player multilinear games, and ii) an efficient online algorithm that incurs average Φ-regret at most ϵ using $\text{poly}(d, k)/\epsilon^2$ rounds.

We also show nearly matching—up to constant factors in the exponents—lower bounds parameterized by k in the online learning setting. We thus obtain for the first time a family of deviations that captures the learnability of Φ-regret. At the heart of our approach is a polynomial-time algorithm for computing an *expected fixed point* of any $\phi : \mathcal{X} \to \mathcal{X}$—that is, a distribution $\mu \in \Delta(\mathcal{X})$ such that $\mathbb{E}_{\boldsymbol{x} \sim \mu}[\phi(\boldsymbol{x}) - \boldsymbol{x}] \approx 0$—based on the seminal *ellipsoid against hope (EAH)* algorithm of Papadimitriou and Roughgarden (JACM '08). In particular, our algorithm for computing Φ-equilibria is based on executing EAH in a nested fashion—each step of EAH itself being implemented by invoking a separate EAH call.

The full version of the paper is available at arxiv.org/pdf/2502.18582.

CCS Concepts: • **Theory of computation → Convergence and learning in games**.

ACM Reference Format:
Brian Hu Zhang, Ioannis Anagnostides, Emanuel Tewolde, Ratip Emin Berker, Gabriele Farina, Vincent Conitzer, and Tuomas Sandholm. 2025. Learning and Computation of Φ-Equilibria at the Frontier of Tractability. In *The 26th ACM Conference on Economics and Computation (EC '25), July 7–10, 2025, Stanford, CA, USA.* ACM, New York, NY, USA, 1 page. https://doi.org/10.1145/3736252.3742670

*Equal contribution.
†Also with Strategy Robot, Inc., Strategic Machine, Inc., Optimized Markets, Inc.

Authors' Contact Information: Brian Hu Zhang, bhzhang@cs.cmu.edu, Carnegie Mellon University, USA; Ioannis Anagnostides, ianagnos@cs.cmu.edu, Carnegie Mellon University, USA; Emanuel Tewolde, etewolde@cs.cmu.edu, Carnegie Mellon University, USA; Ratip Emin Berker, rberker@cs.cmu.edu, Carnegie Mellon University, USA; Gabriele Farina, gfarina@mit.edu, Massachusetts Institute of Technology, USA; Vincent Conitzer, conitzer@cs.cmu.edu, Carnegie Mellon University, USA and University of Oxford, UK; Tuomas Sandholm, sandholm@cs.cmu.edu, Carnegie Mellon University, USA.

EC '25, July 7–10, 2025, Stanford, CA, USA

ACM ISBN 979-8-4007-1943-1/25/07
https://doi.org/10.1145/3736252.3742670