Last-iterate Convergence for Symmetric, General-sum, 2×2 Games Under the Exponential Weights Dynamic

Guanghui Wang[1], Krishna Acharya[2], Lokranjan Lakshmikanthan[2], Juba Ziani[2], Vidya Muthukumar[3,2]
[1]College of Computing, [2]School of Industrial and Systems Engineering, [3]School of Electrical and Computer Engineering, Georgia Institute of Technology {gwang369,krishna.acharya,llakshmi3,jziani3,vmuthukumar8}@gatech.edu

Full paper: https://arxiv.org/abs/2502.08063 (a preliminary version is accepted at ALT 2026)

Motivation. A central question in learning in games is whether natural, uncoupled learning dynamics converge to a Nash equilibrium (NE) in the *last iterate*, rather than merely in *time-average* to the (typically larger) polytope of coarse correlated equilibria. For the exponential weights (EW) dynamic, known last-iterate results are predominantly negative: with a constant step size, EW oscillates in zero-sum games whose only NE are mixed [1], and is chaotic in broad classes including general-sum 2×2 and congestion games [2, 3].

Setting and main result. We give a surprisingly positive and complete answer for a pervasive class overlooked by prior work: general-sum, symmetric, two-player, two-action games, with payoff $(A, A^\top)$ for $A = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$. This captures agents with identically defined utilities (Bertrand competition, atomic congestion, multi-agent performative prediction). This problem has a rich equilibrium landscape: pure NEs always exist, whereas strictly mixed NE arise when $\varepsilon_1 := a - b$ and $\varepsilon_2 := c - d$ differ in sign. Moreover, outside the same-sign cases the correlated-equilibrium polytope is non-singleton, so black-box regret arguments cannot imply even average-iterate convergence to a NE. Our main theorem, summarized in Table 1, fully characterizes the dynamics of discrete-time EW with constant step size η on this class of games: under a mild non-degeneracy assumption, *for any bounded step size* $\eta < 8/(|\varepsilon_1| + |\varepsilon_2|)$*, the last iterate converges to a pure or strictly mixed Nash equilibrium from every initialization, in every symmetric* 2×2 *game*. In most regimes no step-size bound is needed. In the cases where it is needed the bound is mild ($\eta < 2$ for payoffs in $[-1, 1]$) and qualitatively necessary, as above it there exist games on which EW provably cycles. Interestingly, with a constant step size the dynamic is not even no-regret, and yet it equilibrates. As a corollary, we recover the last-iterate convergence of multiplicative weights for 2×2 congestion games ($\varepsilon_1 = -\varepsilon_2$) [4].

Utility condition	Initialization	Nature of result	Requirement on η
r1. $\varepsilon_1 < 0,\ \varepsilon_2 < 0$	Any	Exponential convergence to pure NE	None
r2. $\varepsilon_1 > 0,\ \varepsilon_2 > 0$	Any	Exponential convergence to pure NE	None
r3. $\varepsilon_1 < 0,\ \varepsilon_2 > 0$	Opposite signs	Exponential convergence to one of the pure NEs	None
r4. $\varepsilon_1 < 0,\ \varepsilon_2 > 0$	Same sign	Asymptotic convergence to one of the pure NEs	None
r5. $\varepsilon_1 < 0,\ \varepsilon_2 > 0$	Identical	Asymptotic convergence to strictly mixed NE	$< 8/(\lvert\varepsilon_1\rvert + \lvert\varepsilon_2\rvert)$
r6. $\varepsilon_1 > 0,\ \varepsilon_2 < 0$	Same sign	Exponential convergence to one of the pure NEs	None
r7. $\varepsilon_1 > 0,\ \varepsilon_2 < 0$	Opposite signs	Asymptotic convergence to one of the NEs	$< 8/(\lvert\varepsilon_1\rvert + \lvert\varepsilon_2\rvert)$
r8. $\varepsilon_1 = 0,\ \varepsilon_2 < 0$	Any	Exponential convergence to pure NE	None
r9. $\varepsilon_1 = 0,\ \varepsilon_2 > 0$	Identical	Asymptotic convergence to symmetric pure NE	None
r10. $\varepsilon_1 = 0,\ \varepsilon_2 > 0$	Non-identical	Asymptotic convergence to set of mixed NE	None

Table 1: Summary of our convergence results (Theorem 1). "Initialization" refers to the signs of $\Delta_i^{(1)} = \varepsilon_1 p_{i,1}^{(1)} + \varepsilon_2 p_{i,2}^{(1)}$, where $p_{i,1}^t$ is the weight assigned to action 1 of player 1 at round t.

Proof technique. Our analysis is first-principles in nature. The ratios $r_i = p_{i,1}/p_{i,2}$ turn EW into the coupled scalar system $r_i^{(t+1)} = r_i^{(t)} \exp\big(\eta f(r_j^{(t)})\big)$, where $f(r) := \frac{\varepsilon_1 r + \varepsilon_2}{1+r}$ is a monotone function whose fixed points (0, ∞, and $r^* = |\varepsilon_2|/|\varepsilon_1|$ when the signs differ) are exactly the candidate equilibria, and $f(r^*) = 0$. The crux of our analysis is a novel study of the number of times r_1 and r_2 cross the fixed point r^* during the learning process, under different conditions on the utilities and initialization.

Application and outlook. For multi-agent performative prediction, we formulate a mortgage-competition game in which two banks repeatedly choose between (low credit-score threshold, high interest rate) and (high threshold, low rate), and borrowers select the qualifying lender with the lowest rate — so each bank's customer distribution, hence its utility, depends on *both* deployed policies. The induced game is symmetric and 2×2, and our theorem implies global last-iterate convergence of EW-learning banks to genuine NE. Sometimes the convergence is to *asymmetric* pure NE, i.e., endogenous specialization between *ex-ante* identical lenders. Open directions include symmetric games with more than two actions, heterogeneous step sizes, and non-asymptotic convergence rates. Notably, we believe our first-principles approach has the potential to extend beyond the symmetric setting to some *asymmetric* 2×2 games.

References

[1] James P Bailey and Georgios Piliouras. Multiplicative weights update in zero-sum games. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 321–338, 2018.

[2] Yun Kuen Cheung and Georgios Piliouras. Vortices instead of equilibria in minmax optimization: Chaos and butterfly effects of online learning in zero-sum games. In *Conference on Learning Theory*, pages 807–834. PMLR, 2019.

[3] Gerasimos Palaiopanos, Ioannis Panageas, and Georgios Piliouras. Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. *Advances in Neural Information Processing Systems*, 30, 2017.

[4] Thiparat Chotibut, Fryderyk Falniowski, Michał Misiurewicz, and Georgios Piliouras. Family of chaotic maps from game theory. *Dynamical Systems*, 36(1):48–63, 2021.