Paradoxes of Game Theoretic Equilibria and Price of Anarchy

Georgios Piliouras* **Ian Gemp** **Siqi Liu** **Luke Marris**
Google DeepMind

Abstract

Static equilibria—Nash, Correlated (CE), and Coarse Correlated Equilibria (CCE)—and the Price of Anarchy (PoA) provide essential, tractable benchmarks for multi-agent systems. However, evaluating learning dynamics purely through C^0 fixed points and discrete empirical distributions abstracts away critical C^1 vector field information. We demonstrate that this reduction diverges significantly from physical learning trajectories. First, the worst-case pure Nash equilibria anchoring canonical PoA bounds mathematically manifest as strict saddles—and in canonical instances, even as global maxima of the exact potential; because they are topologically unstable repellers, natural learning actively bypasses these theoretical inefficiency bounds. Additionally, the PoA metric itself exhibits algebraic sensitivity; relaxing syntactic constraints to accommodate data-driven, strictly positive affine cost models renders the Price of Anarchy unbounded. Furthermore, evaluating algorithms strictly through time-averaged regret minimization to reach CCE (and variants thereof) structurally permits convergence to strictly dominated strategies. Even enforcing optimal $O(1/T)$ swap-regret minimization provably accommodates chaotic limit sets in small normal-form games. Finally, in non-atomic congestion games, discrete-time learning bifurcates into instability and even Li-Yorke chaos, driving time-averaged inefficiency to scale exponentially as 2^p, diverging from static sub-linear bounds. Collectively, these findings highlight the necessity of augmenting classical frameworks with dynamically grounded metrics.

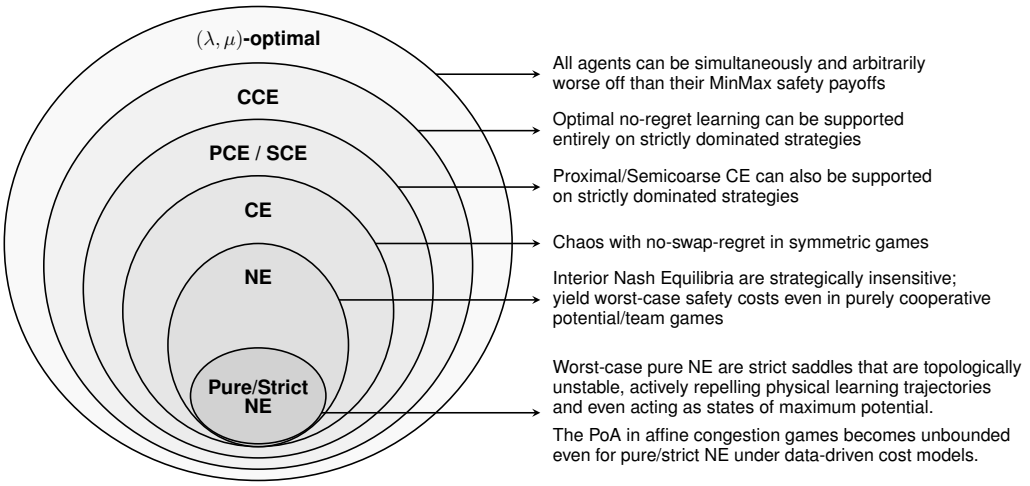

Figure 1: **Paradoxes of Game-Theoretic Solution Concepts.**

*Corresponding author: gpil@google.com

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

1 Introduction

A fundamental challenge in the study of multi-agent systems and economics is the development of target solution concepts that serve as both mathematically rigorous models of agent behavior and reliable predictors of downstream system performance. Algorithmic Game Theory (AGT) has predominantly addressed this challenge through the efficient computation of, and approximations to, classical static equilibria such as Nash Equilibria (NE), Correlated Equilibria (CE), and Coarse Correlated Equilibria (CCE) (Nisan et al., 2007; Roughgarden, 2016). A significant body of research is dedicated to developing fast algorithms for the expedient computation of these concepts and their variants, such as ϵ-NE/CE/CCE or (λ, μ)-optimal states (Syrgkanis et al., 2015; Foster et al., 2016; Daskalakis et al., 2021; Anagnostides et al., 2022c; Piliouras et al., 2022; Farina et al., 2022; Gemp et al., 2024; Soleymani et al., 2025b,a; Li et al., 2024; Peng and Rubinstein, 2024; Anagnostides et al., 2022a,b; Dagan et al., 2024; Daskalakis et al., 2025a; Papadimitriou and Roughgarden, 2008; Jiang and Leyton-Brown, 2015; Farina and Pipis, 2024). Their practical computational tractability is often hailed as a critical test of their real-world applicability, famously captured by Kamal Jain's maxim: "If your laptop can't find it, then neither can the market."

This computational focus has driven influential parallel research aimed at validating the systemic efficiency of these equilibria, most notably through the Price of Anarchy (PoA) literature (Koutsoupias and Papadimitriou, 1999b; Christodoulou and Koutsoupias, 2005b; Vetta, 2002; Roughgarden, 2009a, 2015; Christodoulou and Koutsoupias, 2005a; Roughgarden et al., 2017; Awerbuch et al., 2005; Feldman et al., 2016a; Lykouris et al., 2016). Rooted in worst-case analysis, this framework argues that all equilibria—even worst-case NE/CE/CCE—remain approximately optimal in terms of social welfare. Consequently, enabling proofs of fast convergence to these near optimal states via no-regret learning dynamics has become a foundational paradigm in the analysis of multi-agent systems.

In this paper, we systematically demonstrate that this worst-case AGT framework is misaligned with the natural behavior of game dynamics. We trace this misalignment to a critical information loss: by evaluating multi-agent interactions purely through C^0 fixed-point analysis or discrete empirical distributions, standard solution concepts systematically discard the C^1 continuation of the underlying vector fields and expected game payoffs. When evaluated under typical, gradient-based learning dynamics, this structural and game-theoretic insensitivity routinely accommodates outcomes that are dynamically unstable, highly inefficient, and non-rationalizable.

Our Contributions. The Topological Degeneracy of Worst-Case Pure Nash (Theorems 1,9): To establish efficiency guarantees, worst-case analysis frequently isolates pure NE. We prove that the specific worst-case pure NE dictating the tightness of robust PoA bounds mathematically manifest as strict saddles of the exact potential function. Consequently, under natural learning dynamics, they are topologically unstable and possess a region of attraction of measure zero. In fact, in canonical textbook constructions, these states exhibit such extreme topological degeneracy that they act as global maxima of the potential, supported on almost everywhere strictly dominated strategies, physically repelling all interior learning trajectories. Their inclusion in performance bounds structurally alters efficiency metrics by accounting for dynamically unstable outcomes.

The Origin of the Anomaly: C^0 Fixed Points and Bifurcations (Theorems 2, 3): We trace the origin of these dynamically unstable pure states to topological bifurcations induced by worst-case parameterization in combination with the dynamic instability of generic mixed NE. Evaluated strictly as fixed points, interior NE provide only C^0 information, lacking the local gradient data to distinguish between aligned and strictly opposing incentives. Up to strategic equivalence, an interior NE solves exponentially many (2^{Nm}) maximally distinct games. The worst-case pure NE inherits this exact dynamic insensitivity, distorting and sometimes even outright inverting the optimization landscape.

The Sensitivity of Price of Anarchy (Theorems 4, 10, and 11): Building on these topological properties, we show that classical ratio-based efficiency metrics evaluate algebraic artifacts distinct from the physical learning dynamics. First, the set of (λ, μ)-approximate optimal states accommodates outcomes where agents experience costs strictly higher than their individual min-max safety levels. Second, when the standard syntactic constraint of non-negative intercepts is relaxed to accommodate statistically optimal, data-driven cost models ($x \geq 1$), the worst-case PoA becomes unbounded. We establish this divergence via minimal generic congestion games—comprising exactly two agents, two strategies, and (up to) four resources—where all affine latency functions remain strictly positive and monotonically increasing on the physical domain. Because the learning vector fields are strictly

invariant to the payoff translations that drive these ratio metrics to infinity, the analytical framework evaluates algebraic representations that are strategically irrelevant to the learning agents themselves.

The Combinatorial Shadow of Correlated Play (Theorems 5, 6, and 12): We show that projecting continuous-space learning trajectories onto the discrete simplex of empirical play (CE/CCE) causes critical information loss. Evaluating learning dynamics strictly through the lens of Coarse Correlated Equilibria is structurally insufficient to exclude strictly dominated strategies; we explicitly construct uncoupled no-regret algorithms that converge to (strong) CCEs supported entirely on dominated outcomes. Positive support on strictly dominated strategies persists even under recent proximal refinements (SCE/PCE). Furthermore, we establish that optimal $O(1/T)$ swap-regret minimization (CE) does not preclude macroscopic turbulence, manifesting as chaotic limit sets even in minimal normal-form games.

Phase Transitions to Exponential Inefficiency and Chaos in the Large Population Limit (Theorem 7): Finally, we test the limits of algorithmic stability by examining non-atomic congestion games. Historically, this setting is considered a highly stable environment characterized by tight efficiency guarantees: Wardrop equilibria correspond to global potential minima, static PoA guarantees are exceptionally tight at $\Theta(p/\ln p)$, and no-regret dynamics are mathematically guaranteed to converge to approximate equilibria. However, by analyzing the discrete-time physical reality of these systems, we demonstrate that this macroscopic stability masks underlying microscopic fragility. We prove that generic learning algorithms (including PGD, and FTRL with L_2 regularization) natively induce Li-Yorke chaos. In symmetric polynomial networks of degree p, the system is captured by a global attractor whose empirical time-averaged social cost diverges from the optimum by exactly 2^p, exponentially diverging from the sub-linear static PoA bound. This highlights a critical behavioral dichotomy: the classical convergence guarantees inherently rely on global fine-tuning, normalizing the learning rate by the maximum possible network congestion $O(N^p)$ to force microscopic, stable steps. While mathematically ensuring convergence, this normalization physically requires agents to become highly unresponsive to absolute delays as the population or network degree grows. If agents instead employ standard, responsive learning rates tuned to locally experienced absolute delays, the effective step size strictly exceeds the critical stability threshold. Rather than converging, the dynamics immediately shatter into a highly inefficient, oscillating or even chaotic regime.

2 Background

Let $I = \{1, \dots, N\}$ be the finite set of players of the game G. Each player $i \in I$ has a finite *strategy set* S_i and a cost function $c_i : S_i \times S_{-i} \to [0, 1]$, where $S_{-i} = \prod_{j \neq i} S_j$. A player $i \in I$ selects a strategy from their set of *mixed strategies* $\Delta(S_i)$, defining a probability distribution over S_i. We extend the domain of the cost function to mixed strategies via the standard multilinear extension of expected utility. The aggregate inefficiency of a state is captured by the social cost $C(s) = \sum_i c_i(s)$.

Definition 1 (Nash (1950))**.** *A* Nash equilibrium *(NE) is a vector of independent distributions* $(p_i^*)_{i \in I} \in \prod_{i \in I} \Delta(S_i)$ *such that no agent can strictly decrease their expected cost via a unilateral deviation:* $\forall i \in I, \forall p_i \in \Delta(S_i),\ c_i(p_i^*, p_{-i}^*) \leq c_i(p_i, p_{-i}^*)$

Equilibria supported exclusively on deterministic strategies are termed *pure*; otherwise, they are *mixed*. To accommodate the limitations of distributed learning, algorithmic game theory frequently studies relaxations of the Nash equilibrium that allow for correlated play.

Definition 2 ((Aumann, 1974))**.** *A* correlated equilibrium *(CE) is a joint distribution* π *over the set of action profiles* $S = \prod_i S_i$ *such that for any player* i *and any pair of distinct strategies* $s_i, s_i' \in S_i$*,* $\sum_{s_{-i} \in S_{-i}} c_i(s_i, s_{-i}) \pi(s_i, s_{-i}) \leq \sum_{s_{-i} \in S_{-i}} c_i(s_i', s_{-i}) \pi(s_i, s_{-i})$

Definition 3 ((Moulin and Vial, 1978; Young, 2004))**.** *A* coarse correlated equilibrium *(CCE) is a joint distribution* π *over the set of action profiles* S *such that for any player* i *and any fixed unconditional deviation* $s_i \in S_i$*,* $\sum_{s \in S} c_i(s)\pi(s) \leq \sum_{s_{-i} \in S_{-i}} c_i(s_i, s_{-i})\pi_i(s_{-i})$ *where* $\pi_i(s_{-i}) = \sum_{s_i \in S_i} \pi(s_i, s_{-i})$ *is the marginal distribution of* π *with respect to the opponents of* i*. This solution concept is classically equivalent to Hannan consistency.*

When the inequalities in the CCE definition are strict—meaning any unilateral deviation to a fixed strategy strictly increases expected cost—the distribution constitutes a *strong coarse correlated equilibrium* Anagnostides et al. (2022d). For their connections between (C)CE and online learning, e.g., (swap) regret minimization, see appendix B.

The *Price of Anarchy (PoA)* quantifies the degradation of system efficiency due to decentralized, self-interested behavior. It is defined as the ratio between the worst-case equilibrium cost and the centralized social optimum. Let $\mathcal{E}$ denote the set of valid equilibria in the game (e.g., pure NE, mixed NE, CE, or CCE), and let $s^* \in \operatorname{argmin}_{s \in S} C(s)$ be the social optimum. The Price of Anarchy is formalized as: $\text{PoA}(G) = \frac{\max_{\pi \in \mathcal{E}} C(s)}{\min_{s \in S} C(s)} = \frac{\max_{\pi \in \mathcal{E}} C(s)}{C(s^*)}$

Definition 4 (Smooth Game Roughgarden (2009a, 2015))**.** *A cost-minimization game is (λ, μ)-smooth if for every two outcomes s and s^*: $\sum_i c_i(s_i^*, s_{-i}) \leq \lambda\, C(s^*) + \mu\, C(s)$.*

The smoothness framework abstracts canonical PoA proofs, guaranteeing a *robust PoA* bound of $\frac{\lambda}{1-\mu}$ that applies simultaneously to all aforementioned game-theoretic solution concepts (pure NE, mixed NE, CE, and CCE). (Randomized) states (produced, e.g., by sampling histories of no-regret learning algorithms) satisfying these optimality guarantees are typically called approximately optimal Syrgkanis et al. (2015) or near-optimal Roughgarden (2009b). To resolve any ambiguity, we will call them (λ, μ)-approximate optimal states.

Congestion games: A congestion game is a cost-minimization game defined by a ground set E of resources, a set of N players with strategy sets $S_1, \ldots, S_N \subset 2^E$, and a cost function $c_e : \mathbb{Z}^+ \to \mathbb{R}$ for each resource $e \in E$ (Rosenthal, 1973). We assume the standard conditions that cost functions are non-negative and non-decreasing. Given a strategy profile $s = (s_1, \ldots, s_N)$, the load $x_e(s) = |\{i : e \in S_i\}|$ is the number of agents utilizing resource e. The cost to player i is $c_i(s) = \sum_{e \in s_i} c_e(x_e(s))$, and the social cost is $C(s) = \sum_i c_i(s) = \sum_e c_e(x_e) x_e$. A heavily studied subclass consists of games with affine cost functions $c_e(x) = a_e x + b_e$, where $a_e, b_e \geq 0$. Such games are mathematically guaranteed to be $(5/3, 1/3)$-smooth, resulting in a robust PoA $5/2$ Christodoulou and Koutsoupias (2005a); Roughgarden (2009a).

Valid utility games: A canonical class of smooth payoff-maximization games is the set of *valid utility games* Vetta (2002) (see also *Wonderful Life Utility* in MARL Wolpert and Tumer (1999, 2001)). These environments are defined by a finite ground set E, a non-negative submodular function $V : 2^E \to \mathbb{R}_{\geq 0}$, a strategy set $S_i \subseteq 2^E$ for each agent i, and a private utility function u_i. For any strategy profile $s = (s_1, \ldots, s_N)$, let $U(s) = \bigcup_{i=1}^N s_i$ denote the aggregate set of selected elements. The social objective function is defined as $W(s) = V(U(s))$. By definition, a valid utility game must satisfy two bounding inequalities: (i) $u_i(s) \geq W(s) - V(U(s_{-i}))$, guaranteeing that an agent's private utility is bounded below by their marginal contribution to the social welfare, and (ii) $\sum_{i=1}^N u_i(s) \leq W(s)$, ensuring that the sum of private utilities does not exceed the total social value.

3 The Strategic Insensitivity of Nash Equilibria

We analyze the topological and dynamical properties of the exact Nash equilibria that anchor classical efficiency bounds in algorithmic game theory. We establish that both the pure and interior equilibria utilized by worst-case analysis exhibit critical topological degeneracies.

3.1 The Topological Degeneracy of Worst-Case Pure Nash

To establish deterministic efficiency guarantees, algorithmic game theory frequently isolates pure Nash equilibria. Because generic pure (strict) NE act as topological attractors, they possess C^1 derivative information that strictly bounds local learning trajectories, appearing to resolve the instability of mixed states. However, we demonstrate that the specific worst-case pure NE engineered to prove the tightness of robust Price of Anarchy (PoA) bounds are not standard attractors, but rather topological anomalies.

Unlike generic strict pure equilibria, which are regular and robust to payoff perturbations, worst-case pure NE are mathematically degenerate. To strictly bind the PoA inequalities, worst-case constructions require tuning the game parameters until agents are exactly indifferent between their optimal and suboptimal strategies. Algebraically, this exact indifference forces the marginal cost difference to vanish perfectly at the boundary of the simplex. By smoothly extending the potential function beyond the closed constraints, this tightness condition mathematically mandates that the worst-case boundary state acts as an interior critical point ($\nabla \Phi = \mathbf{0}$) within the unconstrained space.

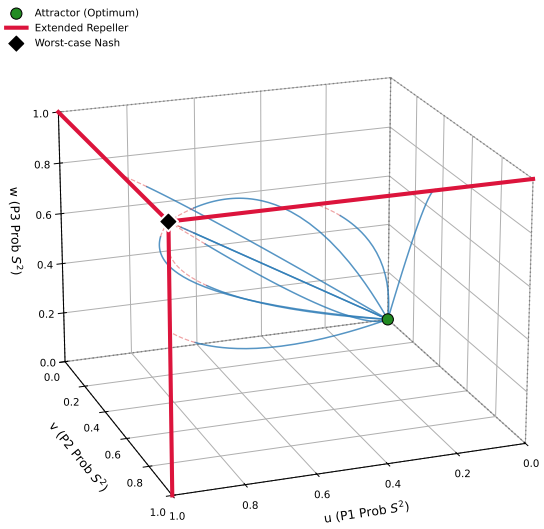

Figure 2: **The Topological Degeneracy of Worst-Case PoA (I).** Replicator dynamics in the linear congestion game producing the tight lower bound of $5/2$ for pure PoA Christodoulou and Koutsoupias (2005a),Roughgarden (2015) (Theorem 5.6, Example 5.7). The worst-case pure NE corresponds to a state of **maximum potential**—the most dynamically unstable state in the system. The strategies supporting the worst case pure Nash are **almost everywhere strictly dominated** in the measure-theoretic sense. Under any gradient-like learning dynamic (where Φ acts as a strict Lyapunov function outside the equilibrium set), the worst-case equilibrium y lies strictly in the **global repeller** of the interior state space. (See Appendix C.2 for a more detailed discussion).

Under the exact algebraic constraints of worst-case parameterization, a topological bifurcation occurs where an interior saddle point is pushed precisely onto the boundary. The resulting degenerate state mathematically satisfies the static inequalities of a pure NE, but its geometry is distorted. These worst-case pure equilibria are not local minima, but rather strict saddles of the extended potential. Because their stable directions can point entirely outside the physical simplex, they can physically manifest as states of maximum potential.

Theorem 1 (Topological Instability of Worst-Case Nash Equilibria)**.** *Let Γ be a canonical exact potential congestion game utilized to establish tight robust Price of Anarchy (PoA) bounds (e.g., the affine double-cycle of Christodoulou and Koutsoupias (2005a); Roughgarden (2015), or the multi-hop network of Awerbuch et al. (2005); Nisan et al. (2007)) where its worst-case pure Nash equilibrium y perfectly binds the smoothness inequalities. Under these exact algebraic constraints, the following properties hold:*

1. *The worst-case pure Nash equilibrium y manifests as a **strict saddle** of the exact continuous potential function Φ.*

2. *Under any smooth learning dynamic defined by an interior-regular Riemannian metric (strictly positive-definite at the boundary), the Lebesgue measure of the basin of attraction for y within the physical probability simplex is exactly zero.*

3. *Under any locally Lipschitz continuous learning dynamic on the strategy space $[0,1]^N$ where the exact potential Φ acts as a strict Lyapunov function outside the equilibrium set, with this equilibrium set contained in the set of rest points of the replicator dynamics (this covers both the case where it coincides with the set of Nash equilibria and the case where it coincides with the full rest-point set of the replicator dynamics), the equilibrium y is **topologically unstable**, actively repelling physically valid learning trajectories away from it.*

3.2 The C^0 Topology of Interior Nash: Deadlocks and Inverse Selection

The topological degeneracy of the worst-case pure equilibria described above is fundamentally inherited from the structural properties of fully mixed (interior) Nash equilibria. Evaluated strictly as fixed points of learning dynamics, interior NE provide only C^0 information about the system. They carry no auxiliary C^1 information regarding the continuation of the vector field infinitesimally close to the fixed point. Because they rely strictly on C^0 evaluation, interior NE lack the local gradient data necessary to distinguish between aligned and strictly opposing incentives. If a game G admits an interior NE, that exact same state is simultaneously an equilibrium of the exact negation $-G$ (where agents act as cost maximizers). Furthermore, up to strategic equivalence, the interior NE remains an equilibrium in any adversarial team game derived by inverting the incentives of an arbitrary subset of agents. We formalize the magnitude of this strategic insensitivity by considering the pre-image of the equilibrium mapping. If an equilibrium concept holds prescriptive power, observing the state p should allow an analyst to largely narrow down the strategic structure of the game being played. Instead, we observe a combinatorial explosion of strategically distinct games.

Theorem 2 (The strategic insensitivity of interior Nash equilibria)**.** *Let G be a generic game with N agents and $m \geq 3$ strategies each, and let p be an interior Nash equilibrium of G. The equivalence class $INV(p)$ of games for which p is a Nash equilibrium includes at least 2^{Nm} distinct games, no two of which are game-theoretically equivalent to each other.*

Interior NE do not solve these games to approximate optimality; they are simply trapped, unable to navigate the exponentially many intersecting games they theoretically inhabit. While prior literature has established the instability of these states under specific dynamics (e.g. replicator, multiplicative weights update) Kleinberg et al. (2009); Panageas et al. (2019), we can distill this instability into a universal geometric law. By directly evaluating the exact potential function, we formally establish that these states are generically strict saddles of the underlying optimization landscape, independent of the choice of learning algorithm.

Theorem 3 (The Universal Zero-Trace Law)**.** *In any exact potential game where players randomize independently, let q^* be a Nash equilibrium. Let $\mathcal{S}^*$ denote the restricted state space defined by the support of q^*. If q^* is partially or fully mixed, and generic within its support (i.e., the restricted Hessian of the potential evaluated at q^* is non-zero), then q^* is strictly a saddle point of the potential function restricted to $\mathcal{S}^*$.*

Remark 1. ***(Structural Robustness of Topological Deadlocks).*** *The strategic insensitivity and topological deadlocks induced by interior Nash equilibria are not fragile artifacts. By Harsanyi's theorem on generic games Harsanyi (1973), generic (interior) equilibria are regular, possessing a non-singular Jacobian of the static payoff mapping. Consequently, by the Implicit Function Theorem, these states are structurally stable. If the payoff entries are subjected to a generic perturbation, the dynamically blind saddle point does not vanish; it merely undergoes a small continuous shift. Thus, the topological traps that actively repel gradient-based learning dynamics persist across a positive Lebesgue measure of the parameter space, rendering local smoothed analysis fundamentally insufficient to resolve the systemic inefficiency.*

4 The Sensitivity of Price of Anarchy

We have established that Nash equilibrium, the predominant game theoretic solution concept, provides certificates of stability to states which are demonstrably and universally unstable from an algorithmic/optimization perspective, the *strict saddles* and even *global maxima* of non-convex minimization landscapes. This introduces some natural structural questions. How can PoA establish approximately optimal social cost when its performance guarantees are anchored to dynamically repelling strict saddles where agents are indifferent between minimization or maximization?

We revisit PoA guarantees from three different perspectives. Critically, we establish the sensitivity of PoA in congestion games, showing that the PoA in affine congestion games is unbounded. Additionally, in Appendix D we explore the sensitivity of PoA guarantees from further perspectives. We introduce a new baseline, where we compare social cost not multiplicatively against the optimal states but additively against min-max safety baselines and show that approximate optimality in the sense of robust PoA is strictly weaker than min-max safety. Finally, we show that even when performing average case performance analysis, PoA ratio-based metrics remain unbounded.

Theorem 4 (The sensitivity of Price of Anarchy)**.** *The Price of Anarchy of affine congestion games is strictly unbounded even when restricting to affine costs with arbitrarily small positive slopes. Moreover, the Pure Price of Anarchy of general affine congestion games is unbounded. This holds even for minimal games with two agents, two strategies, and four congested elements, where all congested elements possess positive, non-decreasing affine latency functions on the active domain.*

These constructions (detailed in Appendix D) refute the possibility of universal positive PoA results for more general classes of empirically derived functions (e.g., polynomials).

5 The Combinatorial Shadow of Learning: Fast Convergence to Strictly Dominated Strategies

A central focus of modern algorithmic game theory is establishing fast, optimal regret-minimization guarantees for learning dynamics in games (Syrgkanis et al., 2015; Foster et al., 2016; Daskalakis et al., 2021; Piliouras et al., 2022; Farina et al., 2022, 2024; Anagnostides et al., 2022b). A common assumption is that optimal algorithmic convergence rates can serve as a proxy for rational, efficient multi-agent behavior. We demonstrate that this reduction introduces critical information loss.

It is an established property of generic, unstructured normal-form games that Coarse Correlated Equilibria (CCE) can support strictly dominated strategies (Viossat and Zapechelnyuk, 2013). Recognizing this vulnerability, recent literature has introduced specialized algorithmic interventions explicitly engineered to filter out unrationalizable states during the learning process (Wang et al., 2023; Wu et al., 2022). However, the mathematical necessity of these bespoke interventions highlights a fundamental limitation in the standard evaluative framework. While specific algorithms can be forcefully engineered to avoid dominated strategies, the fundamental target concept itself—no-regret learning evaluated through the lens of CCE—provides no inherent mathematical guarantee of economic rationality.

When the continuous trajectories of no-regret learning are projected onto the discrete simplex of empirical play via time-averaging, the resulting combinatorial shadow loses critical topological information. We demonstrate that the speed of convergence, even when constrained by strict informational boundaries, is mathematically insufficient to preclude game-theoretic non-rationalizable states, even within the highly structured conditions of smooth games.

Theorem 5 (Fast No-Regret Convergence to Strictly Dominated Strategies in Smooth Games)**.** *There exist smooth games (e.g., linear congestion games and valid utility games) where adversarially selected, strongly uncoupled no-regret algorithms converge at an optimal rate of $O(1/T)$ to a strong coarse correlated equilibrium, whilst simultaneously, on every period, all agents play strictly dominated strategies.*

The proof is detailed in Appendix E. Moreover, the property of assigning positive support to strictly dominated outcomes persists even under recent proximal refinements of CCE, such as Semicoarse/Proximal Correlated Equilibria (SCE/PCE) Cai et al. (2025); Ahunbay and Bichler (2025). See Appendix F for more details.

Remark 2. *Theorem 5 is not an isolated edge-case; it is topologically robust. By leveraging the slackness inherent to the defining inequalities of strong CCE, there exist open families of games exhibiting these properties. This fragility persists both in the native topology of structured games (e.g., perturbing the latency functions of a linear congestion game) and when embedding them into the space of general games, as is inevitable when inferring cost functions from empirical data.*

While the existence of non-rationalizable CCEs has been established in generic games (Viossat and Zapechelnyuk, 2013), Theorem 5 establishes that this problem fundamentally penetrates the highly structured environments of smooth games—the exact environments where CCEs are championed for their robust social welfare guarantees (Anagnostides et al., 2022e; Roughgarden, 2015; Syrgkanis et al., 2015; Foster et al., 2016; Daskalakis et al., 2021). Consequently, standard grey-box algorithmic constraints—such as optimal $O(1/T)$ convergence rates, uncoupledness, and bounded memory—are structurally insufficient to guarantee rational game-theoretic behavior. As these solution concepts are increasingly extended to evaluate complex, high-dimensional domains such as general-sum Markov and extensive-form games (Anagnostides et al., 2022b; Daskalakis et al., 2023; Farina et al., 2020; Jin et al., 2021; Mao and Başar, 2023; Song et al., 2021), recognizing the systemic danger of this

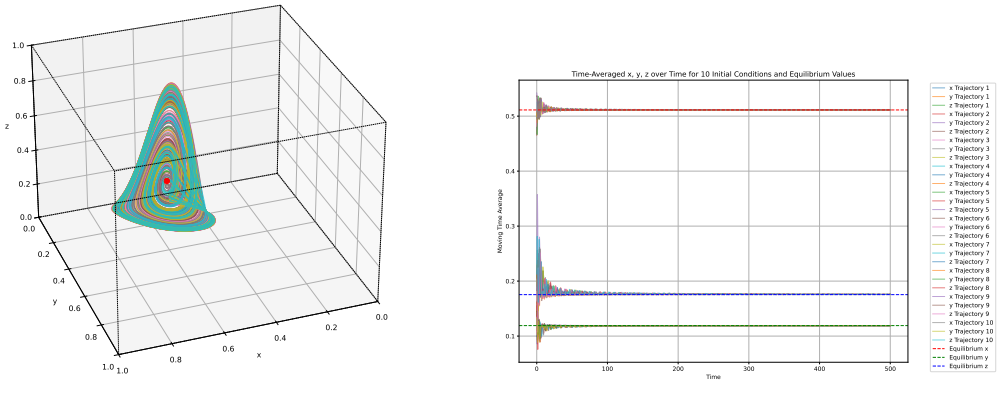

(a) Chaotic limit cycle of replicator dynamics in a normal-form polymatrix game. The red dot is the Nash equilibrium. Colors denote distinct trajectories.

(b) For all interior symmetric initial conditions the time-average of the strategies of all agents converges precisely to the interior Nash equilibrium.

Figure 3: **Chaotic/Strange Correlated Equillibria.** Replicator dynamics simultaneously exhibit macroscopic chaos, $O(1/T)$ convergence to CE, and convergence of time-average marginals to NE.

strategic insensitivity becomes even more important. In conclusion, the issue is merely not that coarse equilibria allow strictly dominated outcomes, but that the standard reduction from learning dynamics to such equilibria erases essential structure: it ignores dynamical accessibility and can certify these outcomes as approximately optimal. As a result, the framework loses the ability to distinguish between stable and unstable, or rational and non-rationalizable behavior, revealing a misalignment between dynamics, equilibrium abstractions, and efficiency guarantees.

6 Beyond the Shadows: The Topological Reality of No Swap Regret

Having established that Coarse Correlated Equilibria and their continuous proximal refinements (SCE/PCE) do not exclude strictly dominated strategies, a natural theoretical recourse is to demand swap-regret minimization. Constrained by the PPAD-hardness of computing Nash equilibria Daskalakis et al. (2006); Chen et al. (2009), swap regret and its sequential extensions have emerged as standard benchmarks for rationalizable learning Brandl (2026); Dagan et al. (2024); Peng and Rubinstein (2024); Celli et al. (2020); Farina et al. (2022); Daskalakis et al. (2025b); Arunachaleswaran et al. (2025). See Appendix G for more discussion.

We prove that evaluating learning algorithms strictly through the discrete combinatorial shadow of swap regret creates an analytical blind spot that obscures complex, non-equilibrating dynamics. By directly analyzing the multi-agent learning trajectories, we prove that even in minimal normal-form games, optimal convergence to CE can perfectly coincide with complex chaotic limit sets.

Theorem 6 (Fast No-Swap-Regret Convergence of Replicator Dynamics to Chaotic Correlated Equilibria)**.** *There exists a one-parameter family of normal-form games with payoff symmetries such that continuous-time replicator dynamics, under symmetric interior initial conditions, exhibits chaos while simultaneously:*

(i) Minimizing swap regret at an optimal rate of $O(1/T)$ (fast convergence to CE).

(ii) Ensuring the time-average strategies of all agents converge precisely to an interior Nash equilibrium.

The games in question are polymatrix games with 6 *agents, each possessing exactly two strategies.*

The proof is detailed in Appendix G.1. See figure 3 for a related visualization. More discussion and visualizations can be found in Appendix G.

7 Phase Transitions, Inefficiency & The Dynamics of Large Populations

In the preceding sections, we established how the combinatorial complexity of finite normal-form games inherently complicates the convergence of independent learning dynamics. One might naturally look to the continuous domain of non-atomic congestion games as a refuge of algorithmic stability. Indeed, this setting typically represents arguably the ideal of game-theoretic and algorithmic optimality. In non-atomic routing, Wardrop equilibria correspond to the global minima of a convex potential function, yielding effectively unique macroscopic flows Beckmann et al. (1956); Nisan et al. (2007). The celebrated Price of Anarchy (PoA) results, based on tight two-link Pigou networks, provide worst-case inefficiency bounds of $\Theta(p/\ln p)$ for polynomial latency functions of degree p Roughgarden and Tardos (2002, 2004). These guarantees are significantly tighter than their atomic counterparts, a favorable property that atomic games are known to formally recover in the limit as the population size increases Feldman et al. (2016b). Furthermore, unlike in atomic games where no-regret algorithms can be consistent with strictly dominated strategies, classic results establish that no-regret learning in non-atomic routing guarantees the day-to-day empirical flow converges to an ϵ-equilibrium flow Blum et al. (2006).

Together, these results form a powerful macroscopic framework. However, a rigorous examination of the physical timescale of these guarantees reveals a fundamental tension in the day-to-day dynamical reality of the system. To achieve a physically meaningful relative error (e.g., bounding regret by a small fraction δ of the optimal equilibrium cost), no-regret algorithms must aggressively throttle their learning rate to account for the maximum possible cost in the network. For a network of m paths with degree p latencies, achieving this relative bound requires an effective step-size scaling of $\mathcal{O}(1/m^p)$, which e.g., given the canonical decreasing step-sizes of $1/\sqrt{t}$, in turn forces the convergence horizon to scale exponentially as $\Omega(m^{2p})$. For standard responsive learning rates, the system operates in a vast meta-stable regime long before asymptotic bounds become binding.

Recent important contributions have illuminated this meta-stable regime, demonstrating the emergence of Li-Yorke chaos and qualitative inefficiencies in non-atomic congestion games under Multiplicative Weights Update (MWU) and Follow-the-Regularized-Leader (FTRL) with steep regularizers Chotibut et al. (2020); Bielawski et al. (2021, 2025). Building upon this work, we provide a generalized and exact quantification of this inefficiency for polynomial latencies, contrasting the dynamical reality directly against the static PoA guarantees.

Theorem 7 (Exponential Degradation of the Global Attractor)**.** *Consider a symmetric non-atomic congestion game with two parallel paths and identical monomial latency functions $c(x) = x^p$ for $p \geq 1$. Under MWU with a fixed effective step-size parameter $a > 0$[2]:*

1. *The unique Wardrop equilibrium $x^* = 1/2$ is locally asymptotically stable if $a < \frac{2^{p+1}}{p}$.*

2. *If $a > \frac{2^{p+1}}{p}$, the equilibrium strictly destabilizes and the interior system is captured by a global, period-2 attracting orbit $\{\sigma_a, 1-\sigma_a\}$. Accounting for the population size N, the empirical time-averaged social cost of this chaotic orbit is exactly $SC_{emp}(a) = N^{p+1}[\sigma_a^{p+1} + (1-\sigma_a)^{p+1}]$, which strictly exceeds the optimal social cost of $N^{p+1}(1/2)^p$.*

3. *As the effective load $a \to \infty$ ($\sigma_a \to 0$), the ratio of this empirical time-averaged social cost to the optimal social cost converges exactly to 2^p.*

The proof is detailed in Appendix I. Theorem 7 establishes a severe divergence between static guarantees and discrete-time dynamics. While the classical Price of Anarchy ensures that the worst-case equilibrium is bounded by $\Theta(p/\ln p)$, the physical, day-to-day behavior of the learning dynamics falls into an attractor that is 2^p strictly worse than the optimum. For further discussion including the emergence of Li-Yorke chaos as well as its connection to Gradient Equilibrium Angelopoulos et al. (2025) see Appendix H.

[2]The parameter a formally couples the MWU learning rate ϵ, the population size N, and the polynomial degree p. It is defined as $a = N^p \ln\left(\frac{1}{1-\epsilon}\right) \approx \epsilon N^p$. This relationship explicitly demonstrates that to maintain a stable effective step-size a, the raw learning rate ϵ must be aggressively throttled inversely to the maximum network cost $\mathcal{O}(N^p)$.

8 Conclusion

Algorithmic Game Theory has traditionally derived its prescriptive guarantees from the efficient computation of static equilibria and the formulation of worst-case performance bounds. Concurrently, the reduction of multi-agent learning to black-box regret minimization has provided a standard mathematical framework for bounding system inefficiency. In this paper, we demonstrated that this analytical paradigm is structurally decoupled from the topological and measure-theoretic realities of game dynamics. By evaluating multi-agent systems purely through C^0 fixed-point analysis or discrete empirical distributions, standard solution concepts inherently abstract away the C^1 continuation of the underlying learning maps and expected game payoffs. Through a series of constructive results, we established that this restriction introduces fundamental analytical limitations:

1. **The Strategic Insensitivity of (Interior) Equilibria:** (Interior) Nash equilibria lack the local gradient information necessary to distinguish between aligned and opposing incentives. Up to strategic equivalence, they function as adversarial "each-vs.-all" games, locking agents into minimal, min-max safety payoffs even in perfectly cooperative potential games.
2. **The Topological Degeneracy of Worst-Case Pure Nash:** The worst-case pure Nash equilibria that dictate the tightness of robust Price of Anarchy bounds mathematically manifest as strict saddles of the extended potential. Consequently, under natural learning dynamics, they are topologically unstable and possess a region of attraction of measure zero. In canonical instances, they are dynamically inaccessible and physically manifest as states of maximum potential, actively repelling valid trajectories and supported on *almost everywhere strictly dominated* strategies—actions that are strictly inferior against all but a measure-zero set of mixed opponent profiles.
3. **The Sensitivity of Efficiency Metrics:** Classical PoA guarantees in affine congestion games are highly sensitive to specific algebraic representations. When standard non-negativity constraints are relaxed to accommodate strictly increasing, data-driven cost models ($x \geq 1$), both the worst-case PoA and the Average PoA become strictly unbounded.
4. **The Combinatorial Shadow of Correlated Play:** Projecting trajectories from the continuous state space of learning algorithms onto the discrete simplex of correlated equilibria is structurally insufficient to preclude non-rationalizable behavior. Uncoupled no-regret learning—and continuous proximal refinements (SCE/PCE)—can be supported (sometimes entirely) on strictly dominated strategies. Furthermore, optimal $O(1/T)$ swap-regret minimization does not preclude complex, non-equilibrating dynamics, manifesting as chaotic limit sets even in small normal-form games.
5. **Exponential Inefficiency and Chaos in the Non-Atomic Limit:** In non-atomic congestion games, standard discrete-time learning dynamics natively induce Li-Yorke chaos. Rather than stabilizing at the Wardrop equilibrium, the system can be captured by global chaotic attractors or other limitsets (e.g., periodic attractors) whose empirical time-averaged social cost degrades exponentially as 2^p, severely diverging from static sub-linear bounds.

These divergences stem directly from a worst-case methodology that assigns equal analytical weight to all algebraically feasible states, regardless of their dynamic stability. By constructing efficiency bounds that must mathematically account for measure-zero topological deadlocks, dynamically inaccessible states, and chaotic limit sets, the current framework risks structurally decoupling theoretical guarantees from the typical trajectories of learning agents. The paradigms of worst-case equilibrium analysis and black-box regret minimization have undeniably served as indispensable foundations for algorithmic game theory. However, as the field matures toward modeling data-driven, empirically grounded environments, our analytical tools must evolve to capture higher-fidelity dynamic realities.

The path forward requires directly interrogating the continuous state-space geometry of the learning dynamics—irrespective of whether the algorithmic updates occur in discrete or continuous time—using tools from ergodic theory and the study of dynamical invariant measures (Bielawski et al., 2026). By shifting focus to physical and natural invariant measures—limit sets possessing basins of attraction with strictly positive Lebesgue measure—future research can systematically isolate the specific, dynamically stable outcomes within the broader worst-case sets. Integrating the topological and measure-theoretic mechanics of dynamical systems into Algorithmic Game Theory provides the necessary foundation for formulating robust, predictive, and empirically grounded guarantees for complex multi-agent systems.

References

Mete Şeref Ahunbay. 2024. First-order (coarse) correlated equilibria in non-concave games. *arXiv preprint arXiv:2403.18174* (2024).

Mete Şeref Ahunbay and Martin Bichler. 2025. Semicoarse Correlated Equilibria and LP-Based Guarantees for Gradient Dynamics in Normal-Form Games. *arXiv preprint arXiv:2502.20466* (2025).

Sebastian Aland, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann. 2011. Exact price of anarchy for polynomial congestion games. *SIAM J. Comput.* 40, 5 (2011), 1211–1233.

Ioannis Anagnostides, Constantinos Daskalakis, Gabriele Farina, Maxwell Fishelson, Noah Golowich, and Tuomas Sandholm. 2022a. Near-optimal no-regret learning for correlated equilibria in multi-player general-sum games. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. 736–749.

Ioannis Anagnostides, Gabriele Farina, Christian Kroer, Andrea Celli, and Tuomas Sandholm. 2022b. Faster no-regret learning dynamics for extensive-form correlated and coarse correlated equilibria. *arXiv preprint arXiv:2202.05446* (2022).

Ioannis Anagnostides, Gabriele Farina, Christian Kroer, Chung-Wei Lee, Haipeng Luo, and Tuomas Sandholm. 2022c. Uncoupled learning dynamics with $\mathcal{O}(\log T)$ swap regret in multiplayer games. *Advances in Neural Information Processing Systems* 35 (2022), 3292–3304.

Ioannis Anagnostides, Gabriele Farina, Ioannis Panageas, and Tuomas Sandholm. 2022d. Optimistic mirror descent either converges to Nash or to strong coarse correlated equilibria in bimatrix games. *Advances in Neural Information Processing Systems* 35 (2022), 16439–16454.

Ioannis Anagnostides, Ioannis Panageas, Gabriele Farina, and Tuomas Sandholm. 2022e. On last-iterate convergence beyond zero-sum games. In *International Conference on Machine Learning*. PMLR, 536–581.

Ioannis Anagnostides, Ioannis Panageas, Gabriele Farina, and Tuomas Sandholm. 2023. On the convergence of no-regret learning dynamics in time-varying games. *Advances in Neural Information Processing Systems* 36 (2023), 16367–16405.

Anastasios N. Angelopoulos, Michael I. Jordan, and Ryan J. Tibshirani. 2025. Gradient Equilibrium in Online Learning: Theory and Applications. *Journal of Machine Learning Research* 26, 305 (2025), 1–68. `http://jmlr.org/papers/v26/25-0356.html`

Sanjeev Arora, Elad Hazan, and Satyen Kale. 2012. The Multiplicative Weights Update Method: a Meta-Algorithm and Applications. *Theory of Computing* 8, 1 (2012), 121–164.

Eshwar Ram Arunachaleswaran, Natalie Collina, Yishay Mansour, Mehryar Mohri, Jon Schneider, and Balasubramanian Sivan. 2025. Swap regret and correlated equilibria beyond normal-form games. In *Proceedings of the 26th ACM Conference on Economics and Computation*. 130–157.

Robert J Aumann. 1974. Subjectivity and correlation in randomized strategies. *Journal of mathematical Economics* 1, 1 (1974), 67–96.

Baruch Awerbuch, Yossi Azar, and Amir Epstein. 2005. The price of routing unsplittable flow. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*. 57–66.

David Balduzzi, Sébastien Racanière, James Martens, Jakob N. Foerster, Karl Tuyls, and Thore Graepel. 2018. The Mechanics of n-Player Differentiable Games. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018 (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer G. Dy and Andreas Krause (Eds.). PMLR, 363–372. `http://proceedings.mlr.press/v80/balduzzi18a.html`

Martin Beckmann, Charles B McGuire, and Christopher B Winsten. 1956. *Studies in the Economics of Transportation*. Technical Report.

Jakub Bielawski, Thiparat Chotibut, Fryderyk Falniowski, Grzegorz Kosiorowski, Michał Misiurewicz, and Georgios Piliouras. 2021. Follow-the-regularized-leader routes to chaos in routing games. In *International Conference on Machine Learning*. PMLR, 925–935.

Jakub Bielawski, Thiparat Chotibut, Fryderyk Falniowski, Michał Misiurewicz, and Georgios Piliouras. 2024. Memory loss can prevent chaos in games dynamics. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 34, 1 (2024).

Jakub Bielawski, Thiparat Chotibut, Fryderyk Falniowski, Michał Misiurewicz, and Georgios Piliouras. 2025. Heterogeneity, reinforcement learning, and chaos in population games. *Proceedings of the National Academy of Sciences* 122, 25 (2025), e2319929121.

Jakub Bielawski, Thiparat Chotibut, Fryderyk Falniowski, MichaĹ Misiurewicz, and Georgios Piliouras. 2026. Natural Invariant Measures for Chaotic Game Dynamics: Finding Order in Chaos. *arXiv preprint arXiv:2607.21805* (2026).

Louis S Block and William A Coppel. 2006. *Dynamics in one dimension*. Springer.

Avrim Blum, Eyal Even-Dar, and Katrina Ligett. 2006. Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*. ACM, 45–52.

Avrim Blum and Yishay Mansour. 2007. From external to internal regret. *Journal of Machine Learning Research* 8, 6 (2007).

Florian Brandl. 2026. *Axioms for Correlated Equilibrium*. Working paper. University of Bonn. `https://brandlf.com/docs/correlated.pdf`

Yang Cai, Constantinos Daskalakis, Haipeng Luo, Chen-Yu Wei, and Weiqiang Zheng. 2025. Proximal regret and proximal correlated equilibria: A new tractable solution concept for online learning and games. *arXiv preprint arXiv:2511.01852* (2025).

Yang Cai, Haipeng Luo, Chen-Yu Wei, and Weiqiang Zheng. 2024. Near-optimal policy optimization for correlated equilibrium in general-sum markov games. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 3889–3897.

Andrea Celli, Alberto Marchesi, Gabriele Farina, and Nicola Gatti. 2020. No-regret learning dynamics for extensive-form correlated equilibrium. *Advances in Neural Information Processing Systems* 33 (2020), 7722–7732.

Nikolo Cesa-Bianchi and Gabor Lugoisi. 2006. *Prediction, Learning, and Games*. Cambridge University Press.

Xi Chen, Xiaotie Deng, and Shang-Hua Teng. 2009. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM (JACM)* 56, 3 (2009), 14.

Thiparat Chotibut, Fryderyk Falniowski, Michał Misiurewicz, and Georgios Piliouras. 2020. The route to chaos in routing games: When is Price of Anarchy too optimistic? *Advances in Neural Information Processing Systems* 33 (2020).

George Christodoulou and Elias Koutsoupias. 2005a. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Algorithms–ESA 2005*. Springer, 59–70.

George Christodoulou and Elias Koutsoupias. 2005b. The Price of Anarchy of Finite Congestion Games. *STOC* (2005), 67–73.

Yuval Dagan, Constantinos Daskalakis, Maxwell Fishelson, and Noah Golowich. 2024. From external to swap regret 2.0: An efficient reduction for large action spaces. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 1216–1222.

Constantinos Daskalakis, Alan Deckelbaum, and Anthony Kim. 2011. Near-optimal no-regret algorithms for zero-sum games. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*. SIAM, 235–254.

Constantinos Daskalakis, Gabriele Farina, Maxwell Fishelson, Charilaos Pipis, and Jon Schneider. 2025a. Efficient Learning and Computation of Linear Correlated Equilibrium in General Convex Games. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing* (Prague, Czechia) *(STOC '25)*. Association for Computing Machinery, New York, NY, USA, 542–553. https://doi.org/10.1145/3717823.3718307

Constantinos Daskalakis, Gabriele Farina, Maxwell Fishelson, Charilaos Pipis, and Jon Schneider. 2025b. Efficient learning and computation of linear correlated equilibrium in general convex games. *STOC* (2025).

Constantinos Daskalakis, Maxwell Fishelson, and Noah Golowich. 2021. Near-optimal no-regret learning in general games. *Advances in Neural Information Processing Systems* 34 (2021), 27604–27616.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. 2006. The Complexity of Computing a Nash Equilibrium. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing* (Seattle, WA, USA) *(STOC '06)*. Association for Computing Machinery, New York, NY, USA, 71–78. https://doi.org/10.1145/1132516.1132527

Constantinos Daskalakis, Noah Golowich, and Kaiqing Zhang. 2023. The complexity of Markov equilibrium in stochastic games. In *The Thirty Sixth Annual Conference on Learning Theory*. PMLR, 4180–4234.

Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. 2018. Training GANs with optimism. In *ICLR*.

Yuan Deng, Jon Schneider, and Balasubramanian Sivan. 2019. Strategizing against no-regret learners. *Advances in neural information processing systems* 32 (2019).

Paul Dupuis and Anna Nagurney. 1993. Dynamical Systems and Variational Inequalities. *Annals of Operations Research* 44, 1 (1993), 7–42.

Gabriele Farina, Ioannis Anagnostides, Haipeng Luo, Chung-Wei Lee, Christian Kroer, and Tuomas Sandholm. 2022. Near-optimal no-regret learning dynamics for general convex games. *Advances in Neural Information Processing Systems* 35 (2022), 39076–39089.

Gabriele Farina, Tommaso Bianchi, and Tuomas Sandholm. 2020. Coarse correlation in extensive-form games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 1934–1941.

Gabriele Farina, Julien Grand-Clément, Christian Kroer, Chung-Wei Lee, and Haipeng Luo. 2024. Regret matching+:(in) stability and fast convergence in games. *Advances in Neural Information Processing Systems* 36 (2024).

Gabriele Farina, Chun Kai Ling, Fei Fang, and Tuomas Sandholm. 2019. Efficient regret minimization algorithm for extensive-form correlated equilibrium. *Advances in Neural Information Processing Systems* 32 (2019).

Gabriele Farina and Charilaos Pipis. 2024. Polynomial-Time Computation of Exact Phi-Equilibria in Polyhedral Games. *arXiv preprint arXiv:2402.16316* (2024).

Michal Feldman, Nicole Immorlica, Brendan Lucier, Tim Roughgarden, and Vasilis Syrgkanis. 2016a. The price of anarchy in large games. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. 963–976.

Michal Feldman, Nicole Immorlica, Brendan Lucier, Tim Roughgarden, and Vasilis Syrgkanis. 2016b. The Price of Anarchy in Large Games. In *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing* (Cambridge, MA, USA) *(STOC '16)*. ACM, New York, NY, USA, 963–976.

Maxwell Fishelson, Robert Kleinberg, Princewill Okoroafor, Renato Paes Leme, Jon Schneider, and Yifeng Teng. 2025. Full swap regret and discretized calibration. *arXiv preprint arXiv:2502.09332* (2025).

Lampros Flokas, Emmanouil-Vasileios Vlatakis-Gkaragkounis, Thanasis Lianeas, Panayotis Mertikopoulos, and Georgios Piliouras. 2020. No-regret learning and mixed Nash equilibria: They do not mix. In *NeurIPS*.

Dylan J Foster, Thodoris Lykouris, Karthik Sridharan, and Eva Tardos. 2016. Learning in games: Robustness of fast convergence. In *Advances in Neural Information Processing Systems*. 4727–4735.

Dean P Foster and Rakesh V Vohra. 1997. Calibrated learning and correlated equilibrium. *Games and Economic Behavior* 21, 1 (1997), 40–55.

Yoav Freund and Robert E Schapire. 1999. Adaptive game playing using multiplicative weights. *Games and Economic Behavior* 29, 1-2 (1999), 79–103.

Ian Gemp, Andreas Alexander Haupt, Luke Marris, Siqi Liu, and Georgios Piliouras. 2025. Convex Markov Games: A New Frontier for Multi-Agent Reinforcement Learning. In *Forty-second International Conference on Machine Learning*.

Ian Gemp, Luke Marris, and Georgios Piliouras. 2024. Approximating nash equilibria in normal-form games via stochastic optimization. *ICLR* (2024).

John C Harsanyi. 1973. Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International journal of game theory* 2, 1 (1973), 1–23.

Sergiu Hart and Andreu Mas-Colell. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68, 5 (2000), 1127–1150.

Sergiu Hart and Andreu Mas-Colell. 2003. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review* 93, 5 (2003), 1830–1836.

Sergiu Hart and Andreu Mas-Colell. 2006. Stochastic uncoupled dynamics and Nash equilibrium. *Games and economic behavior* 57, 2 (2006), 286–303.

Elad Hazan, Karan Singh, and Cyril Zhang. 2017. Efficient regret minimization in non-convex games. In *International Conference on Machine Learning*. PMLR, 1433–1441.

Albert Xin Jiang and Kevin Leyton-Brown. 2015. Polynomial-time computation of exact correlated equilibrium in compact games. *Games and Economic Behavior* 91 (2015), 347–359.

Chi Jin, Qinghua Liu, Yuanhao Wang, and Tiancheng Yu. 2021. V-Learning–A Simple, Efficient, Decentralized Algorithm for Multiagent RL. *arXiv preprint arXiv:2110.14555* (2021).

J.S. Jordan. 1993. Three Problems in Learning Mixed-Strategy Nash Equilibria. *Games and Economic Behavior* 5, 3 (1993), 368 – 386. `https://doi.org/10.1006/game.1993.1022`

Hassan K. Khalil. 2002. *Nonlinear Systems* (3rd ed.). Prentice Hall, Upper Saddle River, NJ.

Robert Kleinberg, Georgios Piliouras, and Éva Tardos. 2009. Multiplicative Updates Outperform Generic No-Regret Learning in Congestion Games. In *ACM Symposium on Theory of Computing (STOC)*.

E. Koutsoupias and C. Papadimitriou. 1999a. Worst-case equilibria. In *(STACS)*. Springer-Verlag, 404–413.

Elias Koutsoupias and Christos H. Papadimitriou. 1999b. Worst-case Equilibria. In *STACS*. 404–413.

Jason D Lee, Ioannis Panageas, Georgios Piliouras, Max Simchowitz, Michael I Jordan, and Benjamin Recht. 2019. First-order methods almost always avoid strict saddle points. *Mathematical programming* 176, 1-2 (2019), 311–337.

Renato P Leme, Georgios Piliouras, and Jon Schneider. 2024. Convergence of no-swap-regret dynamics in self-play. *Advances in Neural Information Processing Systems* 37 (2024), 76210–76231.

Hanyu Li, Wenhan Huang, Zhijian Duan, David Henry Mguni, Kun Shao, Jun Wang, and Xiaotie Deng. 2024. A survey on algorithms for Nash equilibria in finite normal-form games. *Computer Science Review* 51 (2024), 100613.

Tien-Yien Li and James A. Yorke. 1975. Period Three Implies Chaos. *The American Mathematical Monthly* 82, 10 (1975), 985–992.

Nick Littlestone and Manfred K Warmuth. 1994. The weighted majority algorithm. *Information and computation* 108, 2 (1994), 212–261.

V. Losert and E. Akin. 1983. Dynamics of Games and Genes: Discrete Versus Continuous Time. *Journal of Mathematical Biology* (1983).

Thodoris Lykouris, Vasilis Syrgkanis, and Éva Tardos. 2016. Learning and efficiency in games with dynamic population. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 120–129.

Yishay Mansour, Mehryar Mohri, Jon Schneider, and Balasubramanian Sivan. 2022. Strategizing against learners in bayesian games. In *Conference on Learning Theory*. PMLR, 5221–5252.

Weichao Mao and Tamer Başar. 2023. Provably efficient reinforcement learning in decentralized general-sum Markov games. *Dynamic Games and Applications* 13, 1 (2023), 165–186.

Eric Mazumdar, Lillian J. Ratliff, and S. Shankar Sastry. 2020. On Gradient-Based Learning in Continuous Games. *SIAM Journal on Mathematics of Data Science* 2, 1 (2020), 103–131.

Panayotis Mertikopoulos, Bruno Lecouat, Houssam Zenati, Chuan-Sheng Foo, Vijay Chandrasekhar, and Georgios Piliouras. 2019. Optimistic mirror descent in saddle-point problems: Going the extra(-gradient) mile. In *ICLR*. `https://openreview.net/forum?id=Bkg8jjC9KQ`

Panayotis Mertikopoulos, Christos Papadimitriou, and Georgios Piliouras. 2018. Cycles in Adversarial Regularized Learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms* (New Orleans, Louisiana) *(SODA '18)*. Society for Industrial and Applied Mathematics, USA, 2703–2717.

Jason Milionis, Christos Papadimitriou, Georgios Piliouras, and Kelly Spendlove. 2023. An impossibility theorem in game dynamics. *Proceedings of the National Academy of Sciences* 120, 41 (2023), e2305349120.

Hervé Moulin and J P Vial. 1978. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory* 7 (1978), 201–221.

Anna Nagurney and Ding Zhang. 1996. *Projected Dynamical Systems and Variational Inequalities with Applications*. Kluwer Academic Publishers, Boston.

J. Nash. 1950. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences* (1950), 48–49.

Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. 2007. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA.

Shayegan Omidshafiei, Christos Papadimitriou, Georgios Piliouras, Karl Tuyls, Mark Rowland, Jean-Baptiste Lespiau, Wojciech M Czarnecki, Marc Lanctot, Julien Perolat, and Remi Munos. 2019. α-Rank: Multi-Agent Evaluation by Evolution. *arXiv preprint arXiv:1903.01373* (2019).

Ioannis Panageas and Georgios Piliouras. 2016a. Average case performance of replicator dynamics in potential games via computing regions of attraction. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. 703–720.

Ioannis Panageas and Georgios Piliouras. 2016b. Gradient descent only converges to minimizers: Non-isolated critical points and invariant regions. *arXiv preprint arXiv:1605.00405* (2016).

Ioannis Panageas, Georgios Piliouras, and Xiao Wang. 2019. Multiplicative weights updates as a distributed constrained optimization algorithm: Convergence to second-order stationary points almost always. In *International Conference on Machine Learning*. PMLR, 4961–4969.

Christos H Papadimitriou and Tim Roughgarden. 2008. Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM)* 55, 3 (2008), 1–29.

Telmo Peixe and Alexandre Rodrigues. 2022. Persistent strange attractors in 3D polymatrix replicators. *Physica D: Nonlinear Phenomena* 438 (2022), 133346.

Binghui Peng and Aviad Rubinstein. 2024. Fast swap regret minimization and applications to approximate correlated equilibria. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 1223–1234.

Georgios Piliouras, Ryann Sim, and Stratis Skoulakis. 2022. Beyond time-average convergence: Near-optimal uncoupled online learning via clairvoyant multiplicative weights update. *Advances in Neural Information Processing Systems* 35 (2022), 22258–22269.

R. Tyrrell Rockafellar and Roger J.-B. Wets. 1998. *Variational Analysis*. Springer.

R.W. Rosenthal. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 1 (1973), 65–67.

Tim Roughgarden. 2009a. Intrinsic robustness of the price of anarchy. In *Proc. of STOC*. 513–522.

T. Roughgarden. 2009b. Intrinsic robustness of the price of anarchy. In *ACM Symposium on Theory of Computing (STOC)*. ACM, 513–522.

Tim Roughgarden. 2015. Intrinsic Robustness of the Price of Anarchy. *J. ACM* 62, 5, Article 32 (Nov. 2015), 42 pages. `https://doi.org/10.1145/2806883`

Tim Roughgarden. 2016. *Twenty lectures on algorithmic game theory*. Cambridge University Press.

Tim Roughgarden, Vasilis Syrgkanis, and Eva Tardos. 2017. The price of anarchy in auctions. *Journal of Artificial Intelligence Research* 59 (2017), 59–101.

Tim Roughgarden and Éva Tardos. 2002. How bad is selfish routing? *Journal of the ACM (JACM)* 49, 2 (2002), 236–259.

Tim Roughgarden and Éva Tardos. 2004. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and economic behavior* 47, 2 (2004), 389–403.

Joseph Sakos, Stefanos Leonardos, Stelios Stavroulakis, Will Overman, Ioannis Panageas, and Georgios Piliouras. 2024. Beating Price of Anarchy and Gradient Descent without Regret in Potential Games. In *The Twelfth International Conference on Learning Representations (ICLR 2024)*.

Peter Schuster and Karl Sigmund. 1983. Replicator dynamics. *Journal of Theoretical Biology* 100, 3 (1983), 533 – 538. `https://doi.org/10.1016/0022-5193(83)90445-9`

L.S. Shapley. 1964. Some topics in two-person games. *Advances in game theory, edited by M. Dresher, R. J. Aumann, L. S. Shapley, A. W. Tucker* (1964).

AN Sharkovsky, SF Kolyada, AG Sivak, and VV Fedorenko. 2013. *Dynamics of one-dimensional maps*. Vol. 407. Springer Science & Business Media.

Ashkan Soleymani, Georgios Piliouras, and Gabriele Farina. 2025a. Cautious Optimism: A Meta-Algorithm for Near-Constant Regret in General Games. In *Proceedings of the 26th ACM Conference on Economics and Computation* (Stanford University, Stanford, CA, USA) *(EC '25)*. Association for Computing Machinery, New York, NY, USA, 870. `https://doi.org/10.1145/3736252.3742639`

Ashkan Soleymani, Georgios Piliouras, and Gabriele Farina. 2025b. Faster Rates for No-Regret Learning in General Games via Cautious Optimism. *STOC* (2025).

Ziang Song, Song Mei, and Yu Bai. 2021. When can we learn general-sum Markov games with a large number of players sample-efficiently? *arXiv preprint arXiv:2110.04184* (2021).

Sylvain Sorin. 2009. Exponential weight algorithm in continuous time. *Mathematical Programming* 116, 1 (2009), 513–528.

Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E. Schapire. 2015. Fast Convergence of Regularized Learning in Games. In *Proceedings of the 28th International Conference on Neural Information Processing Systems* (Montreal, Canada) *(NIPS'15)*. MIT Press, Cambridge, MA, USA, 2989–2997. http://dl.acm.org/citation.cfm?id=2969442.2969573

Peter D. Taylor and Leo B. Jonker. 1978. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences* 40, 1–2 (1978), 145 – 156. https://doi.org/10.1016/0025-5564(78)90077-9

Adrian Vetta. 2002. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.* IEEE, 416–425.

Yannick Viossat and Andriy Zapechelnyuk. 2013. No-regret dynamics and fictitious play. *Journal of Economic Theory* 148, 2 (2013), 825–842.

Yuanhao Wang, Dingwen Kong, Yu Bai, and Chi Jin. 2023. Learning Rationalizable Equilibria in Multiplayer Games. In *International Conference on Learning Representations (ICLR)*.

David H Wolpert and Kagan Tumer. 1999. An introduction to collective intelligence. *arXiv preprint cs/9908014* (1999).

David H Wolpert and Kagan Tumer. 2001. Optimal payoff functions for members of collectives. *Advances in Complex Systems* 4, 02n03 (2001), 265–279.

Jibang Wu, Haifeng Xu, and Fan Yao. 2022. Multi-agent learning for iterative dominance elimination: Formal barriers and new algorithms. In *Conference on Learning Theory*. PMLR, 543–543.

H. P. Young. 2004. *Strategic Learning and Its Limits*. Oxford University Press, USA.

Brian Hu Zhang, Ioannis Anagnostides, Gabriele Farina, and Tuomas Sandholm. 2024. Efficient Phi-Regret Minimization with Low-Degree Swap Deviations in Extensive-Form Games. *arXiv preprint arXiv:2402.09670* (2024).

A Other Related Work

This paper bridges classical algorithmic game theory with the continuous state-space analysis of dynamical systems, prompting a formal re-evaluation of prescriptive solution concepts. Our work extends beyond adjusting existing performance bounds to systematically analyze the geometric and topological validity of the target states anchoring three major pillars of the classical literature.

A.1 Equilibrium Relaxations and the Topology of Potential Games

The observation that natural learning dynamics fail to converge to Nash equilibria is well documented, tracing back to classical non-convergence results in cyclic environments by Shapley and Jordan Shapley (1964); Jordan (1993). Historically, the algorithmic game theory and multi-agent learning communities responded to these impossibility results—such as the non-convergence of uncoupled dynamics Hart and Mas-Colell (2003, 2006); Milionis et al. (2023)—by relaxing the target solution concepts (e.g., to Correlated and Coarse Correlated Equilibria) or by augmenting the learning algorithms with stochasticity and memory. Under these relaxations, the equilibrium-centric paradigm has remained foundational, scaling to extensive-form games Farina et al. (2019); Anagnostides et al. (2022b); Zhang et al. (2024), Markov and stochastic games Daskalakis et al. (2023); Jin et al. (2021); Mao and Başar (2023); Song et al. (2021); Cai et al. (2024); Gemp et al. (2025), polyhedral games Farina and Pipis (2024), convex and non-convex games Daskalakis et al. (2025b); Hazan et al. (2017), time-varying environments Anagnostides et al. (2023), and continuous/differentiable

games Balduzzi et al. (2018); Mazumdar et al. (2020); Mertikopoulos et al. (2019); Daskalakis et al. (2018).

Rather than evaluating these equilibrium concepts in adversarial or cyclic environments where non-convergence is topologically expected, we focus our analysis on exact potential and congestion games. These environments form the canonical baseline for Algorithmic Game Theory as a whole Nisan et al. (2007); Roughgarden (2016); they serve as the foundation for seminal Price of Anarchy (PoA) bounds Koutsoupias and Papadimitriou (1999a); Christodoulou and Koutsoupias (2005a); Awerbuch et al. (2005); Roughgarden (2015) and are conventionally characterized by guaranteed convergence and macroscopic stability. By examining these constructions through the formalism of continuous optimization theory, we show that the pure Nash equilibria engineered to tightly bind the classical PoA inequalities are mathematically degenerate. Because their construction enforces exact indifference—embedding the algebraic structure of a mixed equilibrium onto the boundary of the simplex—they manifest locally as strict saddles of the extended potential. Furthermore, in the 3-player construction of Christodoulou and Koutsoupias (2005a), the stable manifold of the worst-case equilibrium does not meet the interior of the valid physical state space, and the equilibrium is moreover a constrained global maximum of the potential function. Consequently, the canonical tight bounds are governed by implicit geometric constraints, anchoring efficiency metrics to topological repellers that diverge under gradient-like learning trajectories.

A.2 Correlated Play and the Limitations of C^0 Analysis

Given the topological instability of these deterministic anchors, their probabilistic generalizations require analogous geometric scrutiny. The reduction of multi-agent learning to regret minimization Littlestone and Warmuth (1994); Freund and Schapire (1999); Hart and Mas-Colell (2000) and its extension into robust Price of Anarchy Roughgarden (2015) provided a tractable algebraic framework. Supported by results establishing near-optimal convergence rates for time-averaged (social) regret and robust PoA social welfare guarantees Syrgkanis et al. (2015); Daskalakis et al. (2021), this paradigm has motivated numerous follow-ups Foster et al. (2016); Anagnostides et al. (2022b); Daskalakis et al. (2023); Farina et al. (2020); Jin et al. (2021); Mao and Başar (2023); Song et al. (2021); Piliouras et al. (2022); Soleymani et al. (2025b); Peng and Rubinstein (2024); Celli et al. (2020) including proximal refinements as well as gradient-based variations Angelopoulos et al. (2025); Cai et al. (2025); Ahunbay (2024); Ahunbay and Bichler (2025).

We establish that reliance on discrete empirical histories (time-averaging) introduces severe structural slackness. By evaluating trajectories exclusively through time-averaged distributions, optimal algebraic regret minimization can obscure underlying physical instability. We prove that uncoupled learning and continuous proximal refinements can yield empirical convergence while maintaining support on strictly dominated strategies. Furthermore, optimal zero swap-regret dynamics inherently admit chaos and positive Lyapunov exponents even in minimal normal-form games. Thus, while the algorithms satisfy standard algebraic regret criteria, the underlying physical system remains in a non-equilibrating chaotic regime.

A.3 Non-Atomic Games and Macroscopic Instability

Finally, we re-evaluate the stability of the non-atomic limit, an environment historically characterized by tight sub-linear efficiency bounds Roughgarden and Tardos (2002, 2004); Blum et al. (2006). Recent literature has identified instability and Li-Yorke chaos in non-atomic routing games for numerous dynamics including Multiplicative Weights, Follow-the-Regularized-Leader dynamics with steep-regularizers, as well as canonical behavioral models such as Experienced Weighted Attraction learning Chotibut et al. (2020); Bielawski et al. (2021, 2025, 2024). Despite the prevalence of such results, the impact of such unstable behavior to system performance is much less understood. The only known prior result established globally attracting period-2 limit cycles only in the case of linear congestion games ($p = 1$), which correspond only to a mild degradation of social welfare over the canonical PoA bounds (from $4/3$ to 2 times worse than optimal).

Our analysis extends the geometric scope of this macroscopic instability. First, we establish the first chaotic limit sets for non-steep regularizers (such as L_2), demonstrating that instability is a robust topological feature of standard discrete-time optimization algorithms such as (Projected) Gradient Descent. Second, we show that addressing polynomial cost functions of degree $p \geq 3$

necessitates more intricate arguments compared to typically analyzed linear variants. We prove that in these higher-degree environments, the system destabilizes into a global period-two attractor whose the time-averaged inefficiency scales exponentially as 2^p, diverging fundamentally from the sub-linear $\Theta(p/\ln p)$ static bounds. This formally decouples the classical algebraic guarantees from the discrete-time dynamical reality, demonstrating that without strictly bounding the learning rate by factors scaling with network degree and population, the non-atomic limit natively exhibits exponential inefficiency.

B Background on learning in games

We provide some expanded background on online learning, regret minimization and specific examples of learning dynamics (i.e., replicator dynamics).

Online Learning and Regret Minimization

We conceptualize learning through the standard framework of online optimization. For notational brevity, we adopt the perspective of a single arbitrary agent, omitting player indices to focus on the temporal evolution of their mixed strategy $p^t \in \Delta(S)$ over rounds $t = 1, \dots, T$.

In each round t:

- The player selects a probability distribution p^t over their strategies.
- The environment (or opposing agents) determines a cost vector c^t, specifying a cost $c^t(s)$ for every available strategy s.
- The player incurs an expected cost of $c^t(p^t) = p^t \cdot c^t$.

The distribution p^t chosen at time t depends strictly on the historical sequence of observations up to $t-1$. Upon committing to p^t, the agent receives full-information feedback, observing the entire realized cost vector c^t.

Given a learning algorithm A generating the sequence $\{p^1, p^2, \dots\}$, the cumulative expected cost by time T is $\sum_{t=1}^{T} c^t(p^t)$. Regret quantifies the performance of an algorithm relative to optimal static baselines in hindsight.

Definition 5 (Regret and Swap Regret)**.** *The* external regret *and* swap regret *of algorithm A at time T are defined respectively as:*

$$\textit{Regret}(T) = \sum_{t=1}^{T} c^t(p^t) - \min_{p \in \Delta(S)} \sum_{t=1}^{T} c^t(p) \tag{1}$$

$$\textit{SwapRegret}(T) = \sum_{t=1}^{T} c^t(p^t) - \min_{\rho: S \to S} \sum_{t=1}^{T} c^t(\rho(p^t)). \tag{2}$$

In Equation (2), the maximum is over all strategy modification functions ρ mapping the action space to itself, extended to act multilinearly on mixed strategies. An algorithm is Hannan consistent (or exhibits *no-regret*) if its external regret grows strictly sublinearly, $o(T)$. It is a standard result that efficient algorithms exist bounding both external and swap regret at an optimal rate of $O(\sqrt{T})$ (Cesa-Bianchi and Lugoisi, 2006; Blum and Mansour, 2007; Young, 2004).

Continuous-Time Dynamics and Strategic Equivalence

Shifting to continuous time, the *replicator dynamic* Taylor and Jonker (1978); Schuster and Sigmund (1983) serves as a foundational model of evolutionary game theory and acts as the exact fluid limit of the discrete-time Multiplicative Weights Update (MWU) algorithm Kleinberg et al. (2009); Arora et al. (2012); Losert and Akin (1983). It is governed by the differential equation:

$$\frac{dp_i(t)}{dt} = \dot{p}_i = -p_i[c_i(p) - \hat{c}(p)], \quad \hat{c}(p) = \sum_{i=1}^{N} p_i c_i(p) \tag{3}$$

Benefiting from continuous-time smoothing, the replicator dynamic achieves strictly bounded $O(1)$ regret in general adversarial environments Mertikopoulos et al. (2018); Sorin (2009). Standard AGT analysis routinely projects the continuous trajectories of such no-regret algorithms onto discrete empirical distributions via time-averaging, invoking convergence to CCE (or CE for swap-regret).

However, this linear projection fundamentally discards the topological reality of the vector field. To rigorously analyze the exact limits of these dynamics, it is necessary to identify when distinct games generate identical dynamical trajectories.

Definition 6 (Strategic and Dynamic Equivalence)**.** *Two games G and G' over the same strategy space are strategically equivalent if for each agent i, there exists a non-strategic dummy function $d_i(s_{-i})$ such that $c_i(s) = c'_i(s) + d_i(s_{-i})$ for all $s \in S$. If we further permit a shared positive scalar multiplier $a > 0$ such that $c_i(s) = ac'_i(s) + d_i(s_{-i})$, the games are dynamically equivalent. Finally, if each agent possesses an independent multiplier $a_i > 0$ such that $c_i(s) = a_i c'_i(s) + d_i(s_{-i})$, the games are game-theoretically equivalent.*

Two games are strategically equivalent if and only if the marginal payoff differences between any two unilateral deviations are identical. Consequently, an external observer cannot mathematically distinguish which game within an equivalence class is driving the observed behavior; it represents an innate gauge degree of freedom in the game's definition. Dynamically equivalent games share the exact same set of equilibria and generate identical continuous-time orbits (e.g., under replicator dynamics). Game-theoretical equivalence relaxes this further, preserving all static equilibrium sets but allowing dynamic trajectories to diverge.

C Appendix of Section 3

C.1 Proof of Theorem 1

Before formally proving theorem 1, we introduce the primary mathematical tool utilized to establish topological instability in cases where the boundary metric becomes degenerate. When the Jacobian evaluates identically to the zero matrix, standard linear approximation fails, and we must rely entirely on the exact geometric properties of the non-linear vector field. For this, we employ Chetaev's Instability Theorem, a foundational result in non-linear dynamical systems theory Khalil (2002).

Theorem 8 (Chetaev's Instability Theorem Khalil (2002), Theorem 4.3)**.** *Let $\mathbf{0}$ be an equilibrium point of the continuous dynamical system $\dot{x} = f(x)$, where f is continuously differentiable on a local neighborhood $D \subset \mathbb{R}^N$ around the origin. Suppose there exists a continuously differentiable function $W(x) : D \to \mathbb{R}$ and an open sub-region $\Omega \subset D$ such that:*

1. *The origin $\mathbf{0}$ lies on the boundary of the region Ω.*
2. *$W(x) > 0$ for all $x \in \Omega$.*
3. *$W(x) = 0$ for all $x \in \partial\Omega \cap D$ (the boundary of Ω that lies strictly inside D).*
4. *$\dot{W}(x) = \nabla W(x) \cdot f(x) > 0$ for all $x \in \Omega$ (meaning W strictly increases along trajectories).*

Then, the equilibrium point $\mathbf{0}$ is topologically unstable.

We will use the following relative version. Let $\mathcal{K} \subset \mathbb{R}^N$ be closed with $\mathbf{0} \in \mathcal{K}$, let f be locally Lipschitz on $\mathcal{K}$, and assume $\mathcal{K}$ is forward-invariant (every solution of $\dot{x} = f(x)$ starting in $\mathcal{K}$ exists and remains in $\mathcal{K}$ for all $t \geq 0$); f need not be defined outside $\mathcal{K}$. Let W be continuously differentiable on a neighborhood of $\mathcal{K}$ with $W(\mathbf{0}) = 0$, and suppose that for some $r > 0$ the set $U = \{x \in \mathcal{K} : \|x\| < r,\ W(x) > 0\}$ has $\mathbf{0}$ in its closure, and $\dot{W}(x) > 0$ for every $x \in \mathcal{K}$ with $\|x\| \leq r$ and $W(x) > 0$. Then $\mathbf{0}$ is unstable relative to $\mathcal{K}$. Indeed, Khalil's proof applies verbatim. A trajectory starting at $x_0 \in U$ with $W(x_0) = a > 0$ remains in $\mathcal{K}$, and W increases along it, so it cannot reach the set $\{W = 0\}$. Hence it can leave U only through the sphere $\|x\| = r$. While it remains in U, it lies in the compact set $\{x \in \mathcal{K} : \|x\| \leq r,\ W(x) \geq a\}$, on which $\dot{W}$ is bounded below by a positive constant. Since W is bounded on this set, the trajectory must leave the ball $\|x\| < r$ in finite time. In this relative version, the part of the boundary of U lying on the boundary of $\mathcal{K}$ imposes no condition.

This theorem provides the rigorous mathematical foundation for Property (iii) of our main result. By defining a localized escape region inside the physically valid probability space and demonstrating that an inverted potential function strictly grows within it, Chetaev's Theorem allows us to bypass the zero-Jacobian boundary degeneracy entirely and guarantee that the equilibrium actively repels interior trajectories.

To analyze the local geometry of the worst-case equilibria, we must evaluate the landscape of the potential function. Because the worst-case state lies on the boundary of the probability simplex, we formalize the notion of a strict saddle via smooth analytic extension into an open domain.

Definition 7 (Strict Saddle via Analytic Extension)**.** *Let Φ be a twice-continuously differentiable function defined on a closed manifold with boundaries $\mathcal{X}$. Suppose Φ admits a smooth extension to an open neighborhood $U \subset \mathbb{R}^N$ encompassing $\mathcal{X}$. A point $q^* \in \mathcal{X}$ is a **strict saddle** if, evaluated within the extended interior space U, its gradient vanishes ($\nabla\Phi(q^*) = \mathbf{0}$) and its Hessian matrix, $H = \nabla^2\Phi(q^*)$, admits at least one strictly negative eigenvalue ($\lambda_{\min}(H) < 0$).*

Theorem 1 (Topological Instability of Worst-Case Nash Equilibria)**.** *Let Γ be a canonical exact potential congestion game utilized to establish tight robust Price of Anarchy (PoA) bounds (e.g., the affine double-cycle of Christodoulou and Koutsoupias (2005a); Roughgarden (2015), or the multi-hop network of Awerbuch et al. (2005); Nisan et al. (2007)) where its worst-case pure Nash equilibrium y perfectly binds the smoothness inequalities. Under these exact algebraic constraints, the following properties hold:*

1. *The worst-case pure Nash equilibrium y manifests as a **strict saddle** of the exact continuous potential function Φ.*
2. *Under any smooth learning dynamic defined by an interior-regular Riemannian metric (strictly positive-definite at the boundary), the Lebesgue measure of the basin of attraction for y within the physical probability simplex is exactly zero.*
3. *Under any locally Lipschitz continuous learning dynamic on the strategy space $[0,1]^N$ where the exact potential Φ acts as a strict Lyapunov function outside the equilibrium set, with this equilibrium set contained in the set of rest points of the replicator dynamics (this covers both the case where it coincides with the set of Nash equilibria and the case where it coincides with the full rest-point set of the replicator dynamics), the equilibrium y is **topologically unstable**, actively repelling physically valid learning trajectories away from it.*

Proof. The proof proceeds in four structural parts. First, we analytically extend the potential function to circumvent the boundary topology and link algebraic tightness to an interior critical point. Second, we calculate the exact Hessians to establish the strict saddle. Third, we compute explicit eigenvectors to show that directions of negative curvature point into the physical interior. Finally, we map this geometry to continuous learning dynamics, proving physical repulsion and topological instability.

Part I: Analytic Extension and the Critical Point The Smoothness Framework algebraically bounds the inefficiency of a game via parameters (λ, μ) such that for any states y and y^*:

$$\sum_{i=1}^{N} C_i(y_i^*, y_{-i}) \leq \lambda C(y^*) + \mu C(y) \tag{4}$$

To prove that this bound is perfectly tight, the framework requires a game parameterization where the inequality becomes a strict equality at the worst-case pure Nash equilibrium y. This algebraically mandates that every player i is exactly indifferent between their worst-case strategy y_i and a unilateral deviation to their socially optimal strategy y_i^*.

Let the independent mixed strategy profile be represented by $q \in [0,1]^N$, where q_i represents the probability that player i plays y_i^*, and $1 - q_i$ is the probability they play y_i. Let the worst-case state y be mapped to the origin, $q = \mathbf{0}$.

Because the latency functions of canonical bounding games are polynomial, the exact potential function $\Phi(q)$ is a continuous polynomial over the closed simplex $[0,1]^N$. Because polynomials are infinitely differentiable (C^∞) and analytic everywhere, we smoothly extend Φ from the closed simplex to the open, unconstrained space $\mathbb{R}^N$. Within $\mathbb{R}^N$, the boundary origin $q = \mathbf{0}$ formally acts as an interior point of a manifold without boundary.

Let $\Delta_i(q) = \mathbb{E}[C_i(y_i, q_{-i})] - \mathbb{E}[C_i(y_i^*, q_{-i})]$ denote the expected cost difference. The strict indifference condition mathematically mandates:

$$\Delta_i(\mathbf{0}) = 0 \quad \forall i \in \{1, \ldots, N\} \tag{5}$$

In an exact potential game, the expected cost difference corresponds to the negative gradient of the potential function: $\Delta_i(q) = -\frac{\partial \Phi}{\partial q_i}$. Because $\Delta_i(\mathbf{0}) = 0$ for all i, the gradient vanishes locally at the origin in the extended space:

$$\nabla \Phi(\mathbf{0}) = \mathbf{0} \tag{6}$$

Therefore, the algebraic requirement of tightness formally forces the worst-case Nash equilibrium to act as an interior critical point of the extended potential function.

Part II: Explicit Calculation of the Hessian and the Strict Saddle We evaluate the local geometry of the potential function at the interior critical point by computing its Hessian matrix, H, where $H_{ij} = \frac{\partial^2 \Phi}{\partial q_i \partial q_j} = -\frac{\partial \Delta_i}{\partial q_j}$.

First, we examine the main diagonal elements $H_{ii} = -\frac{\partial \Delta_i}{\partial q_i}$. In any congestion game utilizing independent mixed strategies, the expected latency of a path depends exclusively on the background traffic generated by the opposing players (q_{-i}). Player i's own mixed probability q_i dictates the weighting of their choice, but it does not alter the physical distribution of the opposing traffic. Therefore, the partial derivative of player i's marginal cost with respect to their own probability is identically zero:

$$\frac{\partial \Delta_i}{\partial q_i} = 0 \quad \forall i \in \{1, \dots, N\} \tag{7}$$

Consequently, the main diagonal of the Hessian is entirely zero, guaranteeing that the trace of the Hessian vanishes universally:

$$\text{Trace}(H) = \sum_{i=1}^{N} H_{ii} = 0 \tag{8}$$

By Clairaut's Theorem, H is a real symmetric matrix. To definitively establish that $q = \mathbf{0}$ is a strict saddle, we must prove that $H \neq \mathbf{0}$. We achieve this by explicitly calculating the cross-derivatives for the canonical tightness construction.

Case 1: The 4-Player Awerbuch et al. (2005) Network
We recall the standard 4-player construction yielding the $5/2$ bound (the "AAE example" Awerbuch et al. (2005); see also Nisan et al. (2007), Chapter 18). The network has vertices u, v, w. The directed edges $(u, v), (u, w), (v, w), (w, v)$ have cost $c(x) = x$, and the directed edges $(v, u), (w, u)$ have cost $c(x) = 0$. Four players each route one unit of traffic, with source–sink pairs (u, v), (u, w), (v, w) and (w, v) respectively. Each player chooses between the direct one-hop path y_i^* and a two-hop path y_i:

- Player 1: $y_1^* = u \to v$ and $y_1 = u \to w \to v$;
- Player 2: $y_2^* = u \to w$ and $y_2 = u \to v \to w$;
- Player 3: $y_3^* = v \to w$ and $y_3 = v \to u \to w$;
- Player 4: $y_4^* = w \to v$ and $y_4 = w \to u \to v$.

The profile $y = (y_1, \dots, y_4)$ is a pure Nash equilibrium with social cost 10, while the optimum $y^* = (y_1^*, \dots, y_4^*)$ has social cost 4, giving the ratio $5/2$. A direct computation of the expected loads gives the expected cost differences $\Delta_i(q)$ at the worst-case pure Nash equilibrium:

$$\begin{aligned}
\Delta_1(q) &= 2q_2 - q_3 + 2q_4 \\
\Delta_2(q) &= 2q_1 + 2q_3 - q_4 \\
\Delta_3(q) &= -q_1 + 2q_2 \\
\Delta_4(q) &= 2q_1 - q_2
\end{aligned}$$

Taking the partial derivatives $H_{ij} = -\frac{\partial \Delta_i}{\partial q_j}$ yields the exact, explicit Hessian matrix at the origin:

$$H = \begin{pmatrix} 0 & -2 & 1 & -2 \\ -2 & 0 & -2 & 1 \\ 1 & -2 & 0 & 0 \\ -2 & 1 & 0 & 0 \end{pmatrix} \tag{9}$$

This matrix is explicitly non-zero and symmetric. Because its trace is 0, its eigenvalues must cross the origin. Calculations verify the eigenvalues are $-1+\sqrt{2} \approx 0.414$, $-1-\sqrt{2} \approx -2.414$, $1+\sqrt{10} \approx 4.162$, and $1-\sqrt{10} \approx -2.162$. The existence of strictly negative eigenvalues formally guarantees the state is a strict saddle.

Case 2: The Worst-Case Topologies of Christodoulou and Koutsoupias (2005a),Roughgarden (2015)

We evaluate the exact topological construction that tightly binds the affine Price of Anarchy at $5/2$ (Christodoulou and Koutsoupias (2005a), Roughgarden (2015) Theorem 5.6, Example 5.7). The game consists of $k = 3$ players $\{0, 1, 2\}$ and two disjoint cycles of resources u and v, totaling 6 resources: $\{u_0, u_1, u_2, v_0, v_1, v_2\}$. The cost functions are the identity $c(x) = x$.

Each player i chooses between an optimal strategy $Q_i = \{u_i, v_i\}$ and a worst-case strategy $P_i = \{u_{i+1}, v_{i+1}, v_{i+2}\}$ (indices modulo 3). Let q_i represent the probability that player i plays Q_i.

To determine if the four-way interaction between P_i, Q_i and P_{i-1}, Q_{i-1} cancels out, we explicitly calculate the expected cost difference $\Delta_0(q)$ for Player 0 as a function of the probabilities of the other players (q_1, q_2).

First, we calculate the expected cost of Player 0's worst-case strategy, $\mathbb{E}[C_0(P_0)]$, evaluating the load on its specific edges $P_0 = \{u_1, v_1, v_2\}$. Assuming Player 0 plays P_0 deterministically, they add a load of 1 to these edges:

- Load on u_1: Player 0 plus Player 1 (if playing Q_1). Expected load $= 1 + q_1$.
- Load on v_1: Player 0 plus Player 1 (if playing Q_1) plus Player 2 (if playing P_2). Expected load $= 1 + q_1 + (1 - q_2) = 2 + q_1 - q_2$.
- Load on v_2: Player 0 plus Player 1 (if playing P_1) plus Player 2 (if playing Q_2). Expected load $= 1 + (1 - q_1) + q_2 = 2 - q_1 + q_2$.

Summing these yields the expected cost of the worst-case path:

$$\mathbb{E}[C_0(P_0)] = (1 + q_1) + (2 + q_1 - q_2) + (2 - q_1 + q_2) = 5 + q_1 \tag{10}$$

Next, we calculate the expected cost of Player 0's optimal strategy, $\mathbb{E}[C_0(Q_0)]$, evaluating the load on $Q_0 = \{u_0, v_0\}$ assuming Player 0 plays Q_0 deterministically:

- Load on u_0: Player 0 plus Player 2 (if playing P_2). Expected load $= 1 + (1 - q_2) = 2 - q_2$.
- Load on v_0: Player 0 plus Player 1 (if playing P_1) plus Player 2 (if playing P_2). Expected load $= 1 + (1 - q_1) + (1 - q_2) = 3 - q_1 - q_2$.

Summing these yields the expected cost of the optimal path:

$$\mathbb{E}[C_0(Q_0)] = (2 - q_2) + (3 - q_1 - q_2) = 5 - q_1 - 2q_2 \tag{11}$$

The expected marginal cost difference for Player 0 is:

$$\Delta_0(q) = \mathbb{E}[C_0(P_0)] - \mathbb{E}[C_0(Q_0)] = (5 + q_1) - (5 - q_1 - 2q_2) = 2q_1 + 2q_2 \tag{12}$$

At the worst-case Nash equilibrium ($q = \mathbf{0}$), the indifference condition holds perfectly: $\Delta_0(\mathbf{0}) = 0$. However, the four-way interaction structurally does not cancel out. Taking the partial derivatives yields strictly non-zero values:

$$\frac{\partial \Delta_0}{\partial q_1} = 2 \quad \text{and} \quad \frac{\partial \Delta_0}{\partial q_2} = 2 \tag{13}$$

By the symmetric construction of the cycles, this holds for all players: $\Delta_1(q) = 2q_0 + 2q_2$ and $\Delta_2(q) = 2q_0 + 2q_1$. Thus, the Hessian matrix evaluated at the origin is explicitly:

$$H = \begin{pmatrix} 0 & -2 & -2 \\ -2 & 0 & -2 \\ -2 & -2 & 0 \end{pmatrix} \tag{14}$$

Because this matrix is generated by the structure of the universal worst-case topology, $H \neq \mathbf{0}$ is a physical mandate of theoretical tightness.

The characteristic equation of H yields the explicit eigenvalues $\lambda \in \{2, 2, -4\}$. Because the spectrum contains at least one strictly negative eigenvalue ($\lambda_{\min} = -4$), the critical point $q = \mathbf{0}$ satisfies the definition of a **strict saddle**, establishing Property (i).

Part III: Physical Intersection via Explicit Local Geometry The standard invariant manifold theorems apply to vector fields defined on open subsets. Because we smoothly extended the potential function from the closed simplex to an open neighborhood $U \subset \mathbb{R}^N$, the state $q = \mathbf{0}$ formally acts as an interior fixed point within U.

We must guarantee that the local stable manifold W^s of any continuous learning dynamic does not unilaterally consume the physically valid probability simplex $(0, 1)^N$. First, we note that the explicit Hessians derived in Part II are non-singular. For any dynamic whose linearization at $\mathbf{0}$ is non-degenerate (Case 1 of Part IV, where $J = -M(\mathbf{0})H$ with $M(\mathbf{0}) \succ 0$), all eigenvalues of the Jacobian are therefore non-zero. Consequently, the boundary state $\mathbf{0}$ is a strictly hyperbolic fixed point, and the center manifold is empty. (For degenerate dynamics such as the replicator dynamics, $J = \mathbf{0}$ and linearization is inconclusive; this case is handled in Case 2 of Part IV.)

The repelling directions of the potential landscape are the directions of negative curvature of the Hessian H. We must establish that these directions intersect the physical interior, i.e., that there exists a strictly positive vector v with $v^T H v < 0$. Rather than relying on aggregate quadratic forms, we exhibit such vectors directly as eigenvectors of H.

For the 3-player worst-case topology (Equation 14), evaluating the vector $v = (1, 1, 1)^T$ yields:

$$Hv = \begin{pmatrix} 0 & -2 & -2 \\ -2 & 0 & -2 \\ -2 & -2 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} -4 \\ -4 \\ -4 \end{pmatrix} = -4v$$

Thus, $v = (1, 1, 1)^T$ is an exact eigenvector corresponding to the strictly negative eigenvalue $\lambda = -4$.

Similarly, for the 4-player topology (Equation 9), the characteristic polynomial yields a strictly negative eigenvalue $\lambda = -1 - \sqrt{2}$. Solving the system $(H - \lambda I)v = 0$ yields the corresponding exact eigenvector:

$$v = \begin{pmatrix} 1 \\ 1 \\ \sqrt{2} - 1 \\ \sqrt{2} - 1 \end{pmatrix}$$

Because $\sqrt{2} > 1$, every component of this eigenvector is strictly positive. (The 4-player Hessian has a second negative eigenvalue, $1 - \sqrt{10}$, whose eigenvector has components of mixed sign. One strictly positive negative-curvature direction suffices for our purposes.) In both canonical topologies, the negative-curvature subspace of H thus contains a strictly positive vector v. Since $\nabla\Phi(\mathbf{0}) = \mathbf{0}$, we have $\Phi(tv) = \Phi(\mathbf{0}) + \frac{t^2}{2} v^T H v + o(t^2) < \Phi(\mathbf{0})$ for all sufficiently small $t > 0$. The potential therefore strictly decreases along a direction pointing into the valid probability simplex $(0, 1)^N$. This establishes the geometric foundation for physical repulsion.

Part IV: Continuous Learning Dynamics and Topological Instability We now link the local geometry of the potential function to the vector field of the continuous-time learning dynamic. We evaluate the local stability of the critical point $q = \mathbf{0}$ by bifurcating our analysis into two broad classes of gradient-like dynamics.

First, we consider smooth Riemannian gradient flows of the form:

$$V(q) = M(q)\Delta(q) = -M(q)\nabla\Phi(q)$$

where $M(q)$ is a smooth, positive semi-definite matrix representing the inverse Riemannian metric tensor. Second, we generalize our analysis to encompass any locally Lipschitz continuous learning dynamic on $[0, 1]^N$ for which the exact potential function Φ acts as a strict Lyapunov function outside the equilibrium set. That is, $\frac{d}{dt}\Phi(q(t)) < 0$ at every state outside a set $\mathcal{E}$, where $\mathcal{E}$ is contained

in the set of rest points of the replicator dynamics; e.g., $\mathcal{E}$ may be the set of Nash equilibria. The mathematical machinery required to establish instability depends strictly on whether the dynamic's boundary behavior permits linear approximation.

Case 1: Interior-Regular Metrics. Assume the continuous dynamic is governed by an interior-regular Riemannian metric, where $M(q)$ remains symmetric and strictly positive-definite at the boundary ($M(\mathbf{0}) \succ 0$). We evaluate local stability via the Jacobian matrix, $J = DV(\mathbf{0})$. Because $q = \mathbf{0}$ is a critical point ($\nabla\Phi(\mathbf{0}) = \mathbf{0}$), the product rule yields $J = -M(\mathbf{0})H$.

Because $M(\mathbf{0})$ is positive-definite, it admits a unique positive-definite square root $M^{1/2}$, allowing us to express the Jacobian as similar to a symmetric matrix: $J = -M^{1/2}(M^{1/2}HM^{1/2})M^{-1/2}$. By Sylvester's Law of Inertia, the congruent matrix $M^{1/2}HM^{1/2}$ shares the exact same inertia as the Hessian H. As established in Part II, H is non-singular and possesses at least one strictly negative eigenvalue, meaning $-H$ possesses at least one strictly positive eigenvalue.

Consequently, J is similar to the symmetric matrix $-M^{1/2}HM^{1/2}$, which is non-singular and, by Sylvester's Law of Inertia, has as many strictly negative eigenvalues as H has strictly positive ones. Hence, for any interior-regular metric, all eigenvalues of J are real and non-zero, and the state is a strictly hyperbolic fixed point. Moreover, J possesses at least one strictly positive eigenvalue. By the Stable Manifold Theorem, the local stable manifold W^s_{loc} of the (extended) flow φ_t at $\mathbf{0}$ is a C^1 embedded manifold whose dimension equals the number of positive eigenvalues of H. This number is at most $N-1$; it equals 2 in both canonical games. The basin of attraction of $\mathbf{0}$ is the global stable set $W^s(\mathbf{0}) = \{q : \varphi_t(q) \to \mathbf{0}\}$. Every trajectory converging to $\mathbf{0}$ eventually enters W^s_{loc}, so $W^s(\mathbf{0}) = \bigcup_{n\in\mathbb{N}} \varphi_n^{-1}(W^s_{\text{loc}})$, where φ_n is the time-n map, a C^1 diffeomorphism from its (open) domain onto an open set. Because any smooth manifold of dimension strictly less than N has a Lebesgue measure of exactly zero in $\mathbb{R}^N$, and C^1 diffeomorphisms map null sets to null sets, $W^s(\mathbf{0})$ is a countable union of null sets. Its intersection with the open N-dimensional physical state space $(0,1)^N$ therefore has a Lebesgue measure of exactly zero. This globally establishes Property (ii) without requiring the dynamic's specific unstable eigenvectors to perfectly align with those of the potential landscape.

Case 2: Degenerate Metrics and General Gradient-Like Dynamics. For proportional evolutionary dynamics (e.g., Replicator Dynamics) or other highly non-linear learning rules, the metric tensor may evaluate to exactly zero at pure strategy profiles ($M(\mathbf{0}) = \mathbf{0}$). In such cases, the Jacobian evaluates identically to the zero matrix ($J = \mathbf{0}$), rendering standard linear stability and manifold dimension theorems inapplicable.

However, we can establish Property (iii) for any gradient-like learning dynamic on $[0,1]^N$ where the potential Φ acts as a strict Lyapunov function outside an equilibrium set $\mathcal{E}$ contained in the set of rest points of the replicator dynamics. We do this directly via Chetaev's Instability Theorem, in the relative form stated after Theorem 8, with $\mathcal{K} = [0,1]^N$, which is forward-invariant by assumption (the dynamic is a flow on the strategy space). We define a continuously differentiable Chetaev function by inverting the potential function relative to the boundary equilibrium:

$$W(q) = \Phi(\mathbf{0}) - \Phi(q)$$

We define an escape region Ω as the set of physical states within a local neighborhood $B_\epsilon(\mathbf{0})$, with $\epsilon < 1$, where the potential is strictly less than the boundary equilibrium:

$$\Omega = \{q \in [0,1]^N \cap B_\epsilon(\mathbf{0}) \mid \Phi(q) < \Phi(\mathbf{0})\}$$

Let v be the strictly positive negative-curvature vector exhibited in Part III. Then $tv \in \Omega$ for all sufficiently small $t > 0$. Thus, Ω is non-empty, and the origin $\mathbf{0}$ lies on its boundary. By definition, $W(q) > 0$ on Ω, and $W(q) = 0$ on the relative boundary of Ω in $[0,1]^N$ inside the open ball $B_\epsilon(\mathbf{0})$. Faces of the cube contained in Ω impose no condition in the relative version of Theorem 8, because trajectories cannot leave the strategy space.

We next show that Ω contains no point of $\mathcal{E}$. For affine costs, the Rosenthal potential is quadratic in the resource loads. Each load is a sum of indicators of independent player choices, so $\Phi(q)$ is a polynomial of degree at most two. Since $\nabla\Phi(\mathbf{0}) = \mathbf{0}$, its Taylor expansion is exact: $\Phi(q) = \Phi(\mathbf{0}) + \frac{1}{2}q^THq$, and $\Delta(q) = -Hq$. With two strategies per player, the replicator dynamics read $\dot{q}_i = q_i(1-q_i)\Delta_i(q)$. Its rest points are the states where, for each i, either $q_i \in \{0,1\}$ or $\Delta_i(q) = 0$; every Nash equilibrium is such a state. Let $q \in B_\epsilon(\mathbf{0}) \cap [0,1]^N$ be such a rest point. Since $\epsilon < 1$, no coordinate equals 1, so

$(Hq)_i = 0$ for every i with $q_i > 0$. Hence $q^T Hq = \sum_{i:q_i>0} q_i(Hq)_i = 0$, i.e., $\Phi(q) = \Phi(\mathbf{0})$. Every rest point of the replicator dynamics near y therefore lies on the level set $\{\Phi = \Phi(\mathbf{0})\}$, and none lies in Ω. (In the 3-player game these rest points form the three edges of the cube through y; in the 4-player network, the edges together with the face $\{q_1 = q_2 = 0\}$.)

We evaluate the time derivative of $W(q)$ along the learning trajectories inside Ω. Because $\Phi(q)$ is a strict Lyapunov function outside $\mathcal{E}$, and Ω (indeed every state of $[0,1]^N \cap \overline{B_\epsilon(\mathbf{0})}$ with $\Phi(q) < \Phi(\mathbf{0})$) contains no point of $\mathcal{E}$, the potential must strictly decrease along any physical trajectory within Ω. Therefore:

$$\frac{d}{dt}\Phi(q) < 0 \implies \dot{W}(q) = -\frac{d}{dt}\Phi(q) > 0$$

Because $W(q)$ strictly increases along trajectories inside Ω, the vector field actively forces physically valid states away from the origin toward higher values of $W(q)$ (lower values of $\Phi(q)$). By Chetaev's Theorem (relative to $[0,1]^N$), this mathematically guarantees that the worst-case Nash equilibrium y is topologically unstable. The dynamic actively repels interior trajectories away from it, establishing Property (iii). □

Remark 3 (Physical Geometry and Manifold Intersections)**.** *While the universal zero-trace law guarantees that the worst-case pure Nash equilibrium $q = \mathbf{0}$ is mathematically a strict saddle in the unconstrained space $\mathbb{R}^N$, its physical manifestation depends entirely on how its stable manifold W^s intersects the valid probability simplex $[0,1]^N$.*

In the 3-player affine bounding game of Christodoulou and Koutsoupias (2005a), the quadratic form of the Hessian dictates that all directions of positive curvature point strictly outside the non-negative orthant. Consequently, the stable manifold W^s does not intersect the interior of the physical probability space. Within the confines of the game, every physically valid trajectory flows strictly downhill, causing the saddle to act as a pure topological repeller (and a constrained global maximum of the potential function). See Figure 2 and Appendix C.2 for a more detailed discussion.

*In the 4-player bounding network of Awerbuch et al. (2005), the worst-case state acts as a geometric "mountain pass/bottleneck" suspended between the socially optimal Nash equilibrium (the global minimum of the potential) and the highly congested "anti-equilibria" (the global maxima of the potential). Trajectories nearby the bad Nash equilibrium exhibit **slowing down**—crawling asymptotically toward the worst-case Nash before the local geometry of the potential landscape inevitably captures and repels the flow toward the optimum. Consequently, the worst-case state is **topologically unstable**, actively repelling physically valid learning trajectories and rendering it dynamically inaccessible. When trajectories are simulated back in time, they converge to the two repellers.*

Theorem 1 establishes that the Nash equilibria defining the tightness of classical PoA bounds are dynamically inaccessible, and sometimes even strictly anti-optimal (acting as states of maximum potential). By anchoring efficiency bounds to these dynamically unstable states, worst-case analysis yields performance metrics that are inherently more pessimistic than the empirical trajectories observed under natural learning.

Even if these worst-case games are slightly perturbed to break the degeneracy—undoing the bifurcation so the suboptimal pure state becomes a weakly strict NE (ϵ-attractor)—the fundamental issue remains. The region of attraction for the suboptimal NE becomes a set of negligible measure. Thus, the inefficiencies captured by worst-case bounds are structurally decoupled from the typical behavior of natural learning dynamics.

Remark 4 (Genericity in Arbitrary Cost Functions)**.** *We have explicitly established topological instability for the canonical bounding topologies of Awerbuch et al. (2005); Christodoulou and Koutsoupias (2005a); Roughgarden (2015). For extensions to arbitrary non-decreasing cost functions, the analytical recipe established in this proof fundamentally persists. First, the algebraic tightness of the efficiency bound universally forces the worst-case equilibrium to act as an interior critical point. Second, the independence of mixed strategies universally guarantees a zero-trace Hessian. Therefore, the sole remaining requirement to establish a strict saddle is that the symmetric Hessian does not perfectly vanish ($H \neq \mathbf{0}$). Because exact tightness requires highly coupled, overlapping resource cycles—such as the generalized families provided in Roughgarden (2015)—the presence of non-zero cross-derivatives is a generic algebraic condition of the network design. Consequently, we expect any network topology constructed to perfectly bind the Price of Anarchy to structurally inherit this*

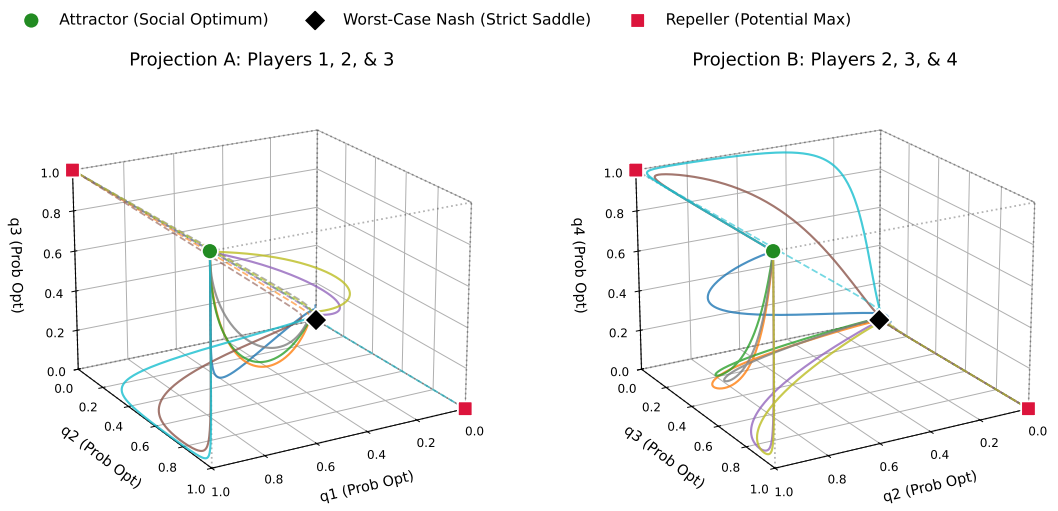

Figure 4: **The Topological Degeneracy of Worst-Case PoA (II)**. Replicator dynamics for the congestion game establishing the $5/2$ pure PoA lower bound Awerbuch et al. (2005), Nisan et al. (2007) (Chapter 18). Shown are two 3D projections of the 4D state space capturing the player trajectories. The worst-case pure Nash equilibrium (black diamond) manifests as a geometric "mountain pass" (**strict saddle**), perfectly suspended between the socially optimum Nash equilibrium (green circle, global minimum of the potential) and the heavily congested "anti-equilibria" (red squares, global maxima). Trajectories nearby the bad Nash equilibrium exhibit **slowing down**—crawling asymptotically toward the worst-case Nash before the local geometry of the potential landscape inevitably captures and repels the flow toward the socially optimum Nash equilibrium. Consequently, the worst-case bound is **topologically unstable**, actively repelling learning trajectories. When trajectories are simulated back in time they converge to the two repellers.

exact saddle geometry, rendering the zero-measure physical basin of attraction a generic feature of tight bounding games.

This topological instability follows a line of reasoning first explored in learning in congestion games in Kleinberg et al. (2009). In that context, it was established that under the continuous-time limit of the Multiplicative Weights Update (replicator dynamics), players inherently avoid unstable mixed Nash equilibria, converging almost everywhere to weakly stable states, typically pure/strict Nash. By proving that the exact worst-case configurations of canonical PoA bounds are structurally mandated to be strict saddles (algebraically behave as mixed Nash), we effectively bridge these two narratives: the worst-case static bounds correspond precisely to the unstable fixed points that natural evolutionary dynamics are mathematically guaranteed to escape. Furthermore, this evasion of strict saddles is not merely a continuous-time artifact. More recent measure-theoretic advances has generalized this phenomenon to a vast family of discrete-time algorithms. It is now well-established that standard first-order methods—including gradient descent even in the presence of non-isolated critical points, as well as mirror descent and block coordinate descent—almost always avoid strict saddles Panageas and Piliouras (2016b); Lee et al. (2019). Consequently, the algebraic necessity of the strict saddle ensures that practically any natural, decentralized learning algorithm will dynamically bypass the theoretical worst-case Price of Anarchy.

C.2 Worst case pure Nash equilibria are global maxima of the potential in Christodoulou and Koutsoupias (2005a); Roughgarden (2015)

Theorem 9. *Let G be the congestion game constructed via the canonical smoothness arguments to yield the worst-case Price of Anarchy for affine cost functions in Christodoulou and Koutsoupias (2005a); Roughgarden (2015) for $N = 3$ agents. Let y^* be the socially optimal pure Nash equilibrium and y be the worst-case pure Nash equilibrium. Then:*

(i) The strategies supporting y are almost everywhere strictly dominated in the measure-theoretic sense.

(ii) The state y, alongside the manifold of states formed by unilateral deviations from it, constitutes the global maximum of the exact potential function Φ.

(iii) Under any gradient-like learning dynamic (where Φ acts as a strict Lyapunov function outside the equilibrium set), the worst-case equilibrium y lies strictly in the global repeller of the interior state space.

Proof. The proof proceeds by explicitly unpacking the canonical lower-bound construction that establishes the tightness of robust PoA bounds for affine latencies (formalized in Roughgarden (2015), Theorem 5.6, and Christodoulou and Koutsoupias (2005b), Theorem 2). We proceed in three stages: explicitly defining the worst-case game geometry, establishing the measure-theoretic dominance via partial derivatives, and deriving the topological consequences.

Step 1: The Geometry of the Worst-Case Double-Cycle Game
To achieve the exact worst-case Price of Anarchy ratio of $5/2$ for affine cost functions $c(x) = ax + b$, the canonical construction requires a highly specific topological structure: the $N = 3$ "double-cycle" congestion game.

The game consists of exactly 3 active players and a ground set of 6 resources partitioned into two disjoint sets, $E_1 = \{h_1, h_2, h_3\}$ and $E_2 = \{g_1, g_2, g_3\}$. The resources in each set are arranged in a cyclic topology. Following the formal construction (e.g., Theorem 5.6 in Roughgarden [2015]), the resources in E_1 and E_2 are assigned the proportionally weighted cost functions $\eta \cdot c(x)$ and $(1 - \eta) \cdot c(x)$, respectively.

Each player $i \in \{1, 2, 3\}$ possesses exactly two strategies: a suboptimal strategy P_i and an optimal strategy Q_i. The optimal block sizes required to strictly bind the $5/2$ affine inequalities dictate the lengths of these contiguous segments:

- Strategy P_i consists of precisely $x_1 = 2$ consecutive resources in E_1 and $x_2 = 1$ resource in E_2, constructed cyclically: $P_i = \{g_{i+1}, h_{i-1}, h_{i+1}\}$ (evaluating indices modulo 3).
- Strategy Q_i consists of precisely $x_1^* = 1$ resource in E_1 and $x_2^* = 1$ resource in E_2, constructed as $Q_i = \{h_i, g_i\}$.
- By construction, for every player i, the strategies P_i and Q_i are strictly disjoint ($P_i \cap Q_i = \emptyset$), while $Q_i \cap Q_j = \emptyset$ for all $i \neq j$.

Let y denote the worst-case pure Nash equilibrium where every player i selects P_i, and let y^* denote the socially optimal state where every player selects Q_i.

In order to maximize the PoA ratio, the parameter η must be algebraically tuned such that at the state y, every player is exactly indifferent between their two strategies ($C_i(y) = C_i(Q_i, y_{-i})$). Equating the expected costs under the base latency $c(x) = ax$ yields the exact balancing condition:

$$\eta a(x_1^2) + (1 - \eta)a(x_2^2) = \eta a(x_1^*)(x_1 + 1) + (1 - \eta)a(x_2^*)(x_2 + 1)$$

Substituting the exact block sizes ($x_1 = 2, x_2 = 1, x_1^* = 1, x_2^* = 1$) yields:

$$\eta a(4) + (1 - \eta)a(1) = \eta a(3) + (1 - \eta)a(2)$$

$$3\eta + 1 = \eta + 2 \implies 2\eta = 1 \implies \eta = 1/2$$

The indifference constraint mathematically mandates that $\eta = 1/2$. (see also Example 5.7 in Roughgarden (2015)). Consequently, the affine cost functions for both cycles possess identical slopes ($a/2$), ensuring uniform marginal congestion costs across the entire topology. To simplify notation moving forward with absorb the multiplicative constant $\eta = 1/2$ into the slope and intercept variables $a \leftarrow a/2, b \leftarrow b/2$. The tight construction uses zero intercepts. With a common intercept $b > 0$, the strategy P_i (three resources) would cost exactly b more than Q_i (two resources) at y, so y would cease to be a Nash equilibrium. This cannot be repaired by adding fixed non-strategic loads, which

cannot decrease the total cost gap $\sum_i [C_i(y) - C_i(Q_i, y_{-i})]$. We therefore work with $b = 0$, for which the indifference condition reads:

$$C_i(y) = C_i(Q_i, y_{-i})$$

Step 2: Almost Everywhere Strict Dominance

Let $q \in [0,1]^3$ represent the mixed strategy profile of the population, where q_j denotes the probability that agent j plays their optimal strategy Q_j (and $1 - q_j$ is the probability they play P_j). For any agent i, let $\Delta_i(q)$ denote the expected cost difference between their two strategies:

$$\Delta_i(q) = \mathbb{E}[C_i(P_i, q_{-i})] - \mathbb{E}[C_i(Q_i, q_{-i})]$$

By the strict indifference condition established in Step 1, when all opponents play the suboptimal strategy ($q_{-i} = \mathbf{0}$), we have $\Delta_i(\mathbf{0}) = 0$.

We evaluate the stability of this indifference by calculating the partial derivative of the cost gap with respect to an opponent's shifting probability q_j. By the linearity of expectation for affine costs, any constant intercept terms vanish under differentiation. The rate of change depends strictly on the slope $a > 0$ and the collision combinatorics between the sets:

$$\frac{\partial \mathbb{E}[C_i(P_i)]}{\partial q_j} = a\left(|P_i \cap Q_j| - |P_i \cap P_j|\right)$$

$$\frac{\partial \mathbb{E}[C_i(Q_i)]}{\partial q_j} = a\left(|Q_i \cap Q_j| - |Q_i \cap P_j|\right)$$

Evaluating Player 1 against Opponent 2 ($j = 2$):

- $P_1 \cap Q_2 = \{h_2, g_2\} \implies$ Size 2. $P_1 \cap P_2 = \{h_3\} \implies$ Size 1. Thus, $\frac{\partial \mathbb{E}[C_1(P_1)]}{\partial q_2} = a(2-1) = a$.
- $Q_1 \cap Q_2 = \emptyset \implies$ Size 0. $Q_1 \cap P_2 = \{h_1\} \implies$ Size 1. Thus, $\frac{\partial \mathbb{E}[C_1(Q_1)]}{\partial q_2} = a(0-1) = -a$.
- The gradient of the gap is: $\frac{\partial \Delta_1}{\partial q_2} = a - (-a) = 2a > 0$.

Evaluating Player 1 against Opponent 3 ($j = 3$):

- $P_1 \cap Q_3 = \{h_3\} \implies$ Size 1. $P_1 \cap P_3 = \{h_2\} \implies$ Size 1. Thus, $\frac{\partial \mathbb{E}[C_1(P_1)]}{\partial q_3} = a(1-1) = 0$.
- $Q_1 \cap Q_3 = \emptyset \implies$ Size 0. $Q_1 \cap P_3 = \{h_1, g_1\} \implies$ Size 2. Thus, $\frac{\partial \mathbb{E}[C_1(Q_1)]}{\partial q_3} = a(0-2) = -2a$.
- The gradient of the gap is: $\frac{\partial \Delta_1}{\partial q_3} = 0 - (-2a) = 2a > 0$.

Because both directional derivatives are strictly positive ($2a > 0$), any infinitesimal shift by any opponent away from $q = \mathbf{0}$ strictly increases the cost gap $\Delta_i(q)$, making P_i strictly more expensive than Q_i. Consequently, for any profile $q_{-i} \neq \mathbf{0}$ (which constitutes the entire Lebesgue measure of the opponents' mixed strategy space), Q_i strictly outperforms P_i.[3] This mathematically guarantees that the strategy P_i supporting the worst-case pure NE is strictly dominated everywhere except at the exact, measure-zero coordinate of the equilibrium itself.

Step 3: Maximum Potential

Every congestion game admits an exact potential function $\Phi(s)$ that precisely mirrors the change in cost for any unilateral deviator:

$$\Phi(s) - \Phi(s'_i, s_{-i}) = C_i(s) - C_i(s'_i, s_{-i})$$

Because y is constructed as a weakly strict Nash equilibrium ($\Delta_i(\mathbf{0}) = 0$), a unilateral deviation by agent i to Q_i yields zero marginal change in cost. Substituting this into the potential function identity yields $\Phi(y) - \Phi(Q_i, y_{-i}) = 0$. This establishes that the potential is completely flat along the k edges of the hypercube connecting y to the states (Q_i, y_{-i}).

[3] As a reminder, the cost functions are linear in the load probabilities jointly so the gradient analysis is sufficient to judge the cost gaps for joint changes in q_j's, i.e., all $q_{-i} \neq 0$.

Crucially, by the strict dominance established in Step 2, for any state s not residing on these specific edges, unilaterally moving toward the suboptimal strategy P_i strictly increases the agent's expected cost. Because the potential function tracks the agent's marginal cost perfectly, moving toward y strictly increases the global potential Φ. Therefore, the set containing y and its immediately adjacent degenerate edges constitutes the strict global maximum of the potential landscape.

Step 4: The Global Repeller
Consider any continuous-time, gradient-like learning dynamic (e.g., replicator dynamics) operating on the state space. In exact potential games, the potential function Φ serves as a strict Lyapunov function for these dynamics, satisfying $\frac{d}{dt}\Phi(p(t)) \leq 0$, with equality holding strictly if and only if the state $p(t)$ is a stationary fixed point.

Let $p(0) \in (0,1)^k$ be an arbitrary interior initial condition. Since the state space $[0,1]^k$ is compact and Φ is bounded below, the sequence $\Phi(p(t))$ must converge monotonically to some limit Φ^*. Let $\Omega(p(0))$ denote the ω-limit set of the trajectory (the set of all accumulation points). By standard topological arguments, $\Omega(p(0))$ is a non-empty, compact, and forward-invariant set. For any accumulation point $q \in \Omega(p(0))$, the continuity of Φ guarantees that $\Phi(q) = \lim_{t\to\infty} \Phi(p(t)) = \Phi^*$.

Because $\Omega(p(0))$ is an invariant set, any trajectory initialized at a point $q \in \Omega(p(0))$ must remain entirely within $\Omega(p(0))$ for all future times. This implies that Φ must remain perfectly constant along any such trajectory. However, the strict Lyapunov property dictates that Φ strictly decreases everywhere except at equilibria. Therefore, $\Omega(p(0))$ must consist exclusively of equilibrium points.

We now evaluate the dynamical viability of the worst-case pure Nash equilibrium y. Because y strictly maximizes the potential function across the interior (as proven in Step 3), we have $\Phi(p(0)) < \Phi(y)$ for any interior initialization. The strict monotonic decrease of the potential function mathematically guarantees that $\Phi(p(t)) \leq \Phi(p(0)) < \Phi(y)$ for all $t \geq 0$. Consequently, the entire trajectory $p(t)$ is permanently trapped in a lower contour set of Φ, strictly bounded away from y. The state y, therefore, cannot be an element of the limit set $\Omega(p(0))$ for any interior trajectory. It acts as a global topological repeller, mathematically forbidding the learning dynamics from ever approaching the worst-case efficiency bound. □

C.3 Proof of Theorem 2

Prior to the formal proof, we rigorously define the genericity condition necessary to preclude degenerate collinearity in the agents' payoff spaces. By fixing the opponents' strategy profile $s_{-i} \in S_{-i}$, we can express the cost function of agent i for a pure strategy j as a vector $\mathbf{c}_{i,j} \in \mathbb{R}^{|S_{-i}|}$.

Definition 8 (Non-Collinear Genericity)**.** *Let G be an N-player game where each agent possesses $m \geq 3$ pure strategies. Let $p \in \prod_i \Delta(S_i)$ be an interior Nash equilibrium. For each agent i and pure strategy j, define the normalized cost vector $\mathbf{c}'_{i,j} \in \mathbb{R}^{|S_{-i}|}$ such that its coordinate for opponent profile s_{-i} is given by $c'_i(j, s_{-i}) = c_i(j, s_{-i}) - c_i(p_i, s_{-i})$. The game G satisfies non-collinear genericity if, for every agent i:*

(i) No normalized cost vector is identically zero ($\mathbf{c}'_{i,j} \neq \mathbf{0}$ for all j).

(ii) The set of normalized cost vectors $\{\mathbf{c}'_{i,1}, \ldots, \mathbf{c}'_{i,m}\}$ spans a subspace of $\mathbb{R}^{|S_{-i}|}$ of dimension strictly greater than 1.

Theorem 2 (The strategic insensitivity of interior Nash equilibria)**.** *Let G be a generic game with N agents and $m \geq 3$ strategies each, and let p be an interior Nash equilibrium of G. The equivalence class $INV(p)$ of games for which p is a Nash equilibrium includes at least 2^{Nm} distinct games, no two of which are game-theoretically equivalent to each other.*

Proof. The proof proceeds constructively in four stages: (1) extracting the normalized base game, (2) generating a combinatorial explosion of 2^{Nm} derivative games, (3) verifying the strict preservation of the equilibrium, and (4) proving the mutually exclusive game-theoretic non-equivalence of the entire set.

Step 1: Normalization of the Base Game
Given the base game G with cost functions c_i, we construct a strategically equivalent normalized

game G' by applying a *strategy-independent cost offset* corresponding to the expected cost of the equilibrium strategy p_i against the opponents' profile:

$$c'_i(j, s_{-i}) = c_i(j, s_{-i}) - c_i(p_i, s_{-i})$$

where $j \in \{1, \dots, m\}$ is a pure strategy and $c_i(p_i, s_{-i}) = \sum_{k=1}^{m} p_i(k) c_i(k, s_{-i})$. Because the subtracted term $c_i(p_i, s_{-i})$ is strictly independent of agent i's chosen action j, G' is strategically equivalent to G.

By construction, the expected cost of playing the mixed strategy p_i against any s_{-i} in G' is identically zero:

$$\sum_{j=1}^{m} p_i(j) c'_i(j, s_{-i}) = 0$$

Since p is an interior Nash equilibrium in G, it remains an interior equilibrium in G'. A fundamental property of interior equilibria is that all pure strategies in the support must yield the exact same expected cost. Therefore, the expected cost of any pure strategy j evaluated against the equilibrium profile p_{-i} must also be exactly zero:

$$\mathbb{E}_{s_{-i} \sim p_{-i}}[c'_i(j, s_{-i})] = 0 \quad \text{for all } i, j$$

Step 2: Constructing the 2^{Nm} Derivative Games
We establish a family of combinatorial sign assignments. For each agent i and each pure strategy j, we assign a scalar sign $\sigma_{i,j} \in \{-1, 1\}$. Across N agents and m strategies, this yields 2^{Nm} distinct sign matrices $\sigma \in \{-1, 1\}^{N \times m}$.

For each sign matrix σ, we define a unique game G^σ with cost functions:

$$c^\sigma_i(j, s_{-i}) = \sigma_{i,j} c'_i(j, s_{-i})$$

Step 3: Verification of Equilibrium Preservation
We now verify that p remains an exact interior Nash equilibrium in every derived game G^σ. We calculate the expected cost of playing pure strategy j against p_{-i} in G^σ:

$$\mathbb{E}_{s_{-i} \sim p_{-i}}[c^\sigma_i(j, s_{-i})] = \mathbb{E}_{s_{-i} \sim p_{-i}}[\sigma_{i,j} c'_i(j, s_{-i})] = \sigma_{i,j} \mathbb{E}_{s_{-i} \sim p_{-i}}[c'_i(j, s_{-i})]$$

Because we established in Step 1 that $\mathbb{E}_{s_{-i} \sim p_{-i}}[c'_i(j, s_{-i})] = 0$, it immediately follows that:

$$\sigma_{i,j} \cdot 0 = 0$$

Because every pure strategy yields an expected cost of exactly zero against p_{-i}, no agent has a strictly profitable deviation. Thus, p is a valid interior Nash equilibrium across all 2^{Nm} games.

Step 4: Proof of Strict Non-Equivalence
We demonstrate that the constructed set of 2^{Nm} games consists entirely of pairwise game-theoretically non-equivalent instances, thereby establishing a strict exponential lower bound on the size of the equivalence class $INV_G(p)$. Assume, for the sake of contradiction, that there exist two distinct sign matrices, $\sigma \neq \tau$, such that G^σ and G^τ are game-theoretically equivalent.

By definition of game-theoretic equivalence, this implies that for every agent i, there exists a positive scalar $a_i > 0$ and a *strategy-independent offset* $d_i(s_{-i})$ such that for all j and all s_{-i}:

$$c^\sigma_i(j, s_{-i}) = a_i c^\tau_i(j, s_{-i}) + d_i(s_{-i})$$

Substituting the definition of our derivative games from Step 2 yields:

$$\sigma_{i,j} c'_i(j, s_{-i}) = a_i \tau_{i,j} c'_i(j, s_{-i}) + d_i(s_{-i})$$

Let $\gamma_{i,j} = \sigma_{i,j} - a_i \tau_{i,j}$. We can rewrite the equation as:

$$\gamma_{i,j} c'_i(j, s_{-i}) = d_i(s_{-i})$$

Crucially, the right-hand side of this equation, $d_i(s_{-i})$, is entirely independent of the chosen strategy j. Formulated in $\mathbb{R}^{m^{N-1}}$, this requires that the scaled vector $\gamma_{i,j} \mathbf{c}'_{i,j}$ must be identical for all pure strategies j. Let us denote this common vector as $\mathbf{D}_i \in \mathbb{R}^{m^{N-1}}$.

We now face two mutually exclusive cases:

- **Case 1: $\mathbf{D}_i \neq \mathbf{0}$ ($\mathbf{D}_i \neq [0, 0, \ldots, 0]$).** If $\mathbf{D}_i$ is a strictly non-zero vector, then the scalar multiplier $\gamma_{i,j}$ must be non-zero for all j. If there existed even a single strategy j where $\gamma_{i,j} = 0$, the scalar multiplication would yield $\gamma_{i,j}\mathbf{c}'_{i,j} = \mathbf{0}$, forcing the common vector $\mathbf{D}_i = \mathbf{0}$ and contradicting the premise of this case. Because $\gamma_{i,j} \neq 0$ for all j, we can divide by it to obtain $\mathbf{c}'_{i,j} = \frac{1}{\gamma_{i,j}}\mathbf{D}_i$. This implies that the normalized payoff vectors for all pure strategies j are merely scalar multiples of the single vector $\mathbf{D}_i$. Consequently, the set $\{\mathbf{c}'_{i,1}, \ldots, \mathbf{c}'_{i,m}\}$ spans a 1-dimensional subspace. However, by condition (ii) of Non-Collinear Genericity for $m \geq 3$, these vectors must span a subspace of dimension strictly greater than 1. This yields a strict contradiction.
- **Case 2: $\mathbf{D}_i = \mathbf{0}$.** In this case, we have $\gamma_{i,j}\mathbf{c}'_{i,j} = \mathbf{0}$ for all j. By condition (i) of Non-Collinear Genericity, the normalized cost vector $\mathbf{c}'_{i,j}$ is strictly non-zero. Therefore, the scalar multiplier must be zero: $\gamma_{i,j} = 0$ for all j.

Since Case 1 yields a contradiction, Case 2 must hold. Thus, for all j:

$$\gamma_{i,j} = 0 \implies \sigma_{i,j} = a_i\tau_{i,j}$$

We are given that $\sigma_{i,j}, \tau_{i,j} \in \{-1, 1\}$ and $a_i > 0$. The only way for a strictly positive scalar to map one sign element identically to another is if $a_i = 1$. Consequently, $\sigma_{i,j} = \tau_{i,j}$ for all i, j.

This implies the matrices are identical ($\sigma = \tau$), which directly contradicts our initial assumption that the games were derived from distinct sign assignments. Therefore, all 2^{Nm} games in $INV_G(p)$ are strictly distinct and game-theoretically non-equivalent. □

C.4 Proof of Theorem 3

Theorem 3 (The Universal Zero-Trace Law)**.** *In any exact potential game where players randomize independently, let q^* be a Nash equilibrium. Let $\mathcal{S}^*$ denote the restricted state space defined by the support of q^*. If q^* is partially or fully mixed, and generic within its support (i.e., the restricted Hessian of the potential evaluated at q^* is non-zero), then q^* is strictly a saddle point of the potential function restricted to $\mathcal{S}^*$.*

Proof. Let $\Phi(q)$ be the exact expected potential function of the game. Because the agents randomize independently, $\Phi(q)$ is formed via the standard multilinear extension over the players' mixed strategy spaces.

Geometrically, if q^* is partially mixed, it resides in the relative interior of a specific face of the global probability polytope. We restrict our analysis to this face, denoted $\mathcal{S}^*$, which corresponds to the subspace spanned by the strategies in the support of q^*. To rigorously evaluate the local geometry within $\mathcal{S}^*$, we formulate the restricted gradient $\nabla_{\mathcal{S}^*}\Phi$ and the restricted Hessian matrix $H_{\mathcal{S}^*} = \nabla^2_{\mathcal{S}^*}\Phi(q^*)$ using strictly independent variables.

For an arbitrary player i with $m \geq 2$ strategies in their support, let $x_1, \ldots, x_{m-1}$ be their independent probability variables within $\mathcal{S}^*$, such that the final support strategy is constrained as $q_{i,m} = 1 - \sum_{k=1}^{m-1} x_k$. The expected potential restricted to $\mathcal{S}^*$ can be expressed as a linear combination of the conditional potentials when player i plays a pure strategy within their support:

$$\Phi(x, q_{-i}) = \sum_{k=1}^{m-1} x_k\Phi_{i,k}(q_{-i}) + \left(1 - \sum_{k=1}^{m-1} x_k\right)\Phi_{i,m}(q_{-i})$$

where $\Phi_{i,k}(q_{-i})$ is the expected potential conditioned on player i deterministically playing strategy k. Regrouping the terms yields:

$$\Phi(x, q_{-i}) = \Phi_{i,m}(q_{-i}) + \sum_{k=1}^{m-1} x_k\left[\Phi_{i,k}(q_{-i}) - \Phi_{i,m}(q_{-i})\right]$$

First, we establish that q^* is a critical point within the restricted subspace $\mathcal{S}^*$. Taking the first derivative with respect to an arbitrary independent variable x_k yields the exact marginal expected potential difference:

$$\frac{\partial\Phi}{\partial x_k} = \Phi_{i,k}(q_{-i}) - \Phi_{i,m}(q_{-i})$$

Because q^* is a Nash equilibrium, player i must be exactly indifferent between all strategies actively played in their support. Consequently, the marginal potentials evaluate to exactly the same value: $\Phi_{i,k}(q^*_{-i}) = \Phi_{i,m}(q^*_{-i})$. Therefore, $\frac{\partial \Phi}{\partial x_k} = 0$ for all independent variables defining $\mathcal{S}^*$, mathematically guaranteeing that $\nabla_{\mathcal{S}^*}\Phi(q^*) = \mathbf{0}$. Thus, q^* is a critical point of the restricted potential.

Next, we evaluate the diagonal entries of the restricted Hessian $H_{\mathcal{S}^*}$ to classify this critical point. Taking the second partial derivative with respect to x_k requires differentiating the marginal difference:

$$[H_{\mathcal{S}^*}]_{x_k x_k} = \frac{\partial^2 \Phi}{\partial x_k^2} = \frac{\partial}{\partial x_k}\left[\Phi_{i,k}(q_{-i}) - \Phi_{i,m}(q_{-i})\right]$$

Crucially, $\Phi_{i,k}(q_{-i})$ and $\Phi_{i,m}(q_{-i})$ depend exclusively on the opposing players' distributions. They possess strictly zero dependence on player i's own mixing probabilities. Consequently, the second derivative vanishes universally:

$$[H_{\mathcal{S}^*}]_{x_k x_k} = 0$$

Because this property holds for every independent probability variable defining the support for every player, the main diagonal of the restricted Hessian $H_{\mathcal{S}^*}$ is entirely populated by zeros. This mandates that the trace of the restricted Hessian is equal to zero:

$$\text{Trace}(H_{\mathcal{S}^*}) = \sum [H_{\mathcal{S}^*}]_{x_k x_k} = 0$$

By the spectral theorem for symmetric matrices, the trace is equal to the sum of the eigenvalues. Since q^* is generic within its support ($H_{\mathcal{S}^*} \neq \mathbf{0}$), it is impossible for the non-zero eigenvalues to sum to zero unless the restricted spectrum contains at least one strictly negative eigenvalue and at least one strictly positive eigenvalue. Because q^* is a restricted critical point possessing directions of both strict positive and negative curvature within its own support manifold, it satisfies the definition of a strict saddle within $\mathcal{S}^*$. □

D Appendix of Section 4

We revisit PoA guarantees from three different perspectives. We establish the sensitivity of PoA in congestion games, showing that the PoA in affine congestion games is unbounded. Furthermore, we introduce a new baseline, where we compare social cost not multiplicatively against the optimal states but additively against min-max safety baselines. We show that approximate optimality in the sense of robust PoA is strictly weaker than min-max safety. Finally, we show that even when performing average case performance analysis, PoA ratio-based metrics remain unbounded.

D.1 The Algebraic Sensitivity of Price of Anarchy Bounds

The standard robust Price of Anarchy (PoA) analysis for affine congestion games rests on a specific algebraic axiom: all resource cost functions must take the form $c_e(x) = a_e x + b_e$, where both $a_e, b_e \geq 0$. While the condition $a_e \geq 0$ correctly models the negative externalities of congestion, enforcing a non-negative intercept $b_e \geq 0$ introduces strict analytical constraints that may diverge from data-driven physical models.

Traditionally, the $b_e \geq 0$ constraint is defended via continuous-domain intuition: a hypothetical "empty" resource ($x = 0$) cannot incur negative latency. However, atomic congestion games inherently operate over discrete, strictly positive integer loads ($x \geq 1$). In real-world systems, it is trivial to encounter strictly positive, monotonically increasing latency functions that mathematically necessitate a negative intercept. Consider a shared machine requiring a fixed "cool-down" or setup period between successive uses. If the first task requires 1 second, and every subsequent task requires 2 seconds (1 for setup, 1 for execution), the exact affine representation of the physical latency is $c(x) = 2x - 1$. The costs are strictly positive for all physically realizable loads ($x \geq 1$), yet the intercept is definitively negative.

The exclusion of such valid physical models highlights a structural limitation within worst-case equilibrium analysis. It is a fundamental property of the Price of Anarchy that it is highly sensitive to affine transformations and cost translations. Adding a fixed positive constant to all costs artificially

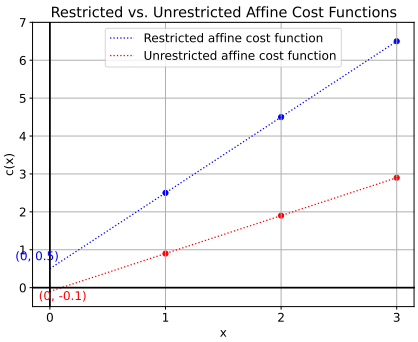

Figure 5: **Algebraic Constraints vs. Physical Validity.** A comparison between a restricted and an unrestricted affine function in robust Price of Anarchy analysis. While both cost functions map strictly positive integer loads to positive costs, the critical parameter dictating their inclusion in PoA bounds relies entirely on whether the extrapolated, physically inaccessible continuous line crosses the y-axis above or below zero.

	A	B
A	$(w+1, w+1)$	$(\mathbf{w}, \mathbf{w})$
B	$(\mathbf{1}, \mathbf{1})$	$(w+1, w+1)$

	A	B
A	(ϵ, ϵ)	$(\epsilon^{\mathbf{2}}, \epsilon^{\mathbf{2}})$
B	$(\epsilon^{\mathbf{2}}, \epsilon^{\mathbf{2}})$	(ϵ, ϵ)

Figure 6: **Unbounded Price of Anarchy in Affine Congestion Games.** Two parametric families of affine congestion games where the PoA is strictly unbounded. Pure Nash equilibrium states are bolded. **Left:** For games with $w > 1$, the pure PoA is unbounded when accommodating affine cost functions of arbitrarily large slopes. **Right:** The mixed PoA is unbounded even for families of affine congestion games with an arbitrarily small positive Lipschitz constant (slope).

inflates the optimal social cost (the denominator), thereby artificially deflating the PoA ratio. Because the PoA metric is inherently *not* invariant to these strategic equivalences, the framework operates under the strict premise that a game's cost functions represent exact, un-normalized empirical measurements. If the applicability of the PoA ratio relies on costs reflecting exact physical measurements, excluding cost functions with negative intercepts structurally limits the application of these bounds, particularly when those functions accurately represent ground-truth measurements of everywhere-positive, strictly increasing physical costs.

This formulation dependence becomes particularly relevant in modern, data-driven applications. In practice, latency functions are rarely provided as exact theoretical primitives; they are typically inferred via regression over historical data. When modeling cost dynamics using standard parametric approaches, statistically optimal best-fit curves frequently yield negative coefficients to accurately capture the curvature over the active physical domain ($x \geq 1$). Imposing strict algebraic non-negativity constraints on these learned parameters restricts the model's expressivity, potentially reducing its fidelity to the empirical observations.[4]

The subtlety of this algebraic boundary highlights a deep structural dichotomy in the classical literature (e.g., Nisan et al. (2007)). While the non-negativity requirement ($b_e \geq 0$) is structurally necessary for the PoA bounds to hold—often formalized within supporting lemmas—this parameter remains dynamically and strategically invariant to the agents themselves. Consequently, the sole algebraic constraint bounding the efficiency metric operates entirely independently from the learning trajectories of the agents.

Consequently, because classical worst-case PoA bounds diverge when valid, everywhere-positive latency functions possess negative intercepts, applying these static bounds to data-driven, learned cost models requires careful structural re-evaluation.

Theorem 4 (The sensitivity of Price of Anarchy)**.** *The Price of Anarchy of affine congestion games is strictly unbounded even when restricting to affine costs with arbitrarily small positive slopes.*

[4]In scientific computing, fitting data using the standard monomial basis is generally avoided due to numerical instability. Regression libraries project data onto orthogonal polynomial bases (e.g., Chebyshev polynomials $T_2(x) = 2x^2 - 1$). Expanding any stably learned cost function back into the standard monomial form can thus naturally yield negative coefficients.

Moreover, the Pure Price of Anarchy of general affine congestion games is unbounded. This holds even for minimal games with two agents, two strategies, and four congested elements, where all congested elements possess positive, non-decreasing affine latency functions on the active domain.

D.2 Divergence from Min-Max Safety Baselines

Since its inception, the Price of Anarchy has evaluated multi-agent systems by benchmarking equilibrium outcomes against a centralized social optimum. While this approach provides a metric for social efficiency, it inherently bypasses a more fundamental game-theoretic criterion: baseline individual rationality. A rigorous test of any solution concept is whether agents can guarantee the performance of their strictly adversarial safety strategies.

We formalize this via min-max costs: the minimum cost an agent can guarantee themselves even if all opponents play a coordinated, adversarial strategy designed to maximize that agent's cost. Any rational agent can independently secure this safety level; thus, any solution concept that systematically yields costs higher than this min-max threshold admits non-rationalizable behavior.

Definition 9 (MinMax-Centered Performance)**.** *Given any game G and any state s, the guaranteed safety gap for agent i is:*

$$\begin{aligned} u_i^{\textit{MinMaxCenter}}(s) &= \min \max c_i - c_i(s) \\ &= \min_{p_i \in \Delta(S_i)} \max_{\pi_{-i} \in \Delta(S_{-i})} \sum_{s_{-i} \in S_{-i}} \sum_{s_i} c_i(s_i, s_{-i}) p_i(s_i) \pi_{-i}(s_{-i}) - c_i(s) \end{aligned}$$

Definition 10 (MinMax-Regret Performance)**.**

$$u_i^{\textit{MinMaxRegret}}(s) = c_i(\textit{minmax}_i^*, s_{-i}) - c_i(s)$$

where $\textit{minmax}_i^*$ is the max-entropy min-max strategy of agent i against a meta-agent controlling the opponents via correlated randomness.

A score of zero serves as the absolute baseline of economic rationality. Any rationalizable outcome must satisfy $u_i^{\text{MinMaxCenter}}(s) \geq u_i^{\text{MinMaxRegret}}(s) \geq 0$. However, the broad class of (λ, μ)-approximate optimal states structurally fails to secure this safety baseline, even in the most canonical, highly symmetric settings.

Theorem 10 (Divergence from Baseline Rationality in Balls and Bins Games)**.** *Consider the class of canonical balls and bins games (symmetric load balancing games with N agents, $M \geq 2$ identical resources, and monomial cost functions $c(x) = x^d$ for $d \geq 1$). The set of (λ, μ)-approximate optimal states dictated by the robust Price of Anarchy framework strictly includes states where every agent simultaneously incurs a cost arbitrarily worse—scaling as $-\Theta(N^d)$—than their absolute (correlated) min-max safety guarantee.*

The proof (detailed in Appendix D.5) is constructive. By explicitly computing the optimal social cost, the exact min-max safety threshold, and the theoretical (λ, μ) upper bound, we demonstrate that the standard robust PoA inequalities are structurally permissive enough to include states of perfect miscoordination within the 'approximately optimal' set.

To understand the severity of this failure, we must analyze the exact probabilistic nature of the load induced by the adversaries. In these canonical congestion games, the agents' incentives are perfectly aligned to minimize collisions. Yet, the fully mixed interior Nash equilibrium extracts exactly zero cooperative surplus. As formalized in the following lemma, the interior equilibrium perfectly mimics the optimal adversarial attack.

Lemma 1 (Independent Min-Max Equivalence)**.** *In symmetric balls and bins games with strictly convex monomial costs $c(x) = x^d$ ($d \geq 2$), the expected cost evaluated at the fully mixed interior Nash equilibrium is mathematically identical to the exact independent min-max safety cost.*

Its proof can be found in Appendix D.6. Modulo the constraint of independent play, the fully mixed interior Nash equilibrium mathematically collapses to the adversarial min-max safety strategy. (For linear congestion games where $d = 1$, linearity of expectation forces this independent safety cost to perfectly collapse to the absolute correlated min-max cost).

This mathematical equivalence between a cooperative equilibrium and a zero-sum adversarial safety strategy is the exact quantitative manifestation of the C^0 topological insensitivity established in

Section 3. Because the worst-case framework must analytically accommodate these zero-surplus topological deadlocks, it systematically generates efficiency bounds that fall arbitrarily below the threshold of economic rationality.

D.3 Average Price of Anarchy is Similarly Unbounded

To salvage efficiency guarantees in the presence of unstable worst-case equilibria, a natural analytical recourse is the Average Price of Anarchy (APoA) Panageas and Piliouras (2016a); Sakos et al. (2024). Rather than evaluating the single worst-case equilibrium, APoA evaluates the expected social cost, weighting each Nash equilibrium proportional to the Lebesgue measure of its basin of attraction under natural learning dynamics. We show that while this methodology recovers optimal performance for certain game families, it inherits the critical limitations of classical PoA.

Theorem 11 (The Dynamical Disconnect of Average-Case Ratios)**.** *For the first parametric family of affine congestion games yielding unbounded PoA (Figure 6), under a uniform Lebesgue measure prior, the expected limit social cost is at most* 2 *times the optimal social cost, with instances reaching an average-case inefficiency of at least* $1 + \frac{1}{27}(18 - 2\sqrt{3}\pi) \approx 1.2636$*. For the second family, given any prior absolutely continuous with respect to the Lebesgue measure, the average-case performance strictly equals the optimal social cost.*

Conversely, by linearly translating the costs of the first family of games, we can construct affine congestion games with two agents and two strategies where, under a uniform Lebesgue prior, the average-case system performance becomes arbitrarily worse than the optimal social welfare.

There are two interpretations of this result. The first is an econometric perspective, which assumes that the observable costs are strictly accurate and exogenously given. Under this lens, Theorem 11 establishes that even when deploying the sophisticated tools of dynamical systems and measure theory to account for the exact regions of attraction of pure/strict Nash equilibria, there is no universal ratio-based efficiency guarantee; performance evaluation naturally reduces to case-by-case empiricism.

The second, more fundamental interpretation is that the Price of Anarchy (and its average-case variants) is neither dynamically nor game-theoretically well-posed. Both static equilibria and the continuous orbits of game dynamics remain invariant to strategically equivalent transformations of the game. However, the social cost—and its corresponding performance ratio—does not. Consequently, multi-agent behavior is evaluated against an efficiency metric that is sensitive to global cost translations, even though these translations are strategically irrelevant to the learning agents and the resulting equilibria.

The performance ratio is sensitive to the specific affine representation chosen from the strategic equivalence class. The convention of enforcing non-negative coefficients, while sufficient to bound this ratio, relies on a specific algebraic representation that does not generalize to all valid physical translations. Even when finite Average PoA bounds can be derived under these restricted representations, they possess no inherent game-theoretic or dynamical meaning. The metric evaluates the system's empirical limit sets based on sensitive coordinate choices, fundamentally decoupling theoretical efficiency guarantees from the topological and measure-theoretic reality of the game.

Proofs of Claims in Section 4 and Appendix D

This part of the appendix provides the formal proofs establishing the structural sensitivity of both individual rationality baselines and macroscopic efficiency metrics (Price of Anarchy and Average Price of Anarchy) in congestion games.

D.4 Proof of Theorem 4

Proof. We establish the theorem constructively by mapping the two minimal normal-form matrices presented in Section D.1 to exact, physically realizable affine congestion games.

Step 1: Unbounded Pure Price of Anarchy
To realize the first payoff matrix, we define a congestion game with $N = 2$ agents and a ground set of four resources: $E = \{e_{AA}, e_{AB}, e_{BA}, e_{BB}\}$. The strategy set for the row agent (Agent 1)

is $S_1 = \{\{e_{AA}, e_{AB}\}, \{e_{BA}, e_{BB}\}\}$. The strategy set for the column agent (Agent 2) is $S_2 = \{\{e_{AA}, e_{BA}\}, \{e_{AB}, e_{BB}\}\}$.

We define the affine latency functions $c(x) = ax + b$ such that the costs at loads $x \in \{1, 2\}$ are given by:

$$\begin{aligned} c_{AA}(1) &= c_{AB}(1) = c_{BA}(1) = c_{BB}(1) = \delta \\ c_{AA}(2) &= w + 1 - \delta \\ c_{AB}(2) &= w - \delta \\ c_{BA}(2) &= 1 - \delta \\ c_{BB}(2) &= w + 1 - \delta \end{aligned}$$

To ensure all latency functions are strictly monotonically increasing on the physical domain ($x \geq 1$), we require $c_e(2) > c_e(1) > 0$ for all e. Solving this system of inequalities yields the single strict parameter condition: $0 < 2\delta < 1 < w$. Provided this condition holds, all elements exhibit positive, strictly increasing affine latency functions.

By mapping the agents' action profiles to the resulting congestion on these elements, we exactly recover the first normal-form matrix. The optimal pure Nash equilibrium resides at (B, A) with a social cost of exactly 2. The worst-case pure Nash equilibrium resides at (A, B) with a social cost of exactly $2w$. Consequently, the Pure Price of Anarchy is exactly w. Since w can be set arbitrarily large, the Pure PoA is strictly unbounded, irrespective of the non-decreasing physical properties of the latencies.

Step 2: Unbounded Mixed Price of Anarchy (Small Slopes)
To realize the second payoff matrix, we define a completely symmetric congestion game with $N = 2$ agents and exactly two resources: $E = \{e_1, e_2\}$. Both agents share the identical strategy set $S_1 = S_2 = \{\{e_1\}, \{e_2\}\}$.

Both resources share an identical affine cost function defined by:

$$c(x) = (\epsilon - \epsilon^2)x - \epsilon + 2\epsilon^2$$

for some parameter $0 < \epsilon < 1/2$. Because $\epsilon < 1/2$, the slope $(\epsilon - \epsilon^2)$ is strictly positive, ensuring the latency function is strictly monotonically increasing. Evaluated at integer loads, the costs are $c(1) = \epsilon^2$ and $c(2) = \epsilon$. Because $\epsilon > \epsilon^2 > 0$, the costs are strictly positive everywhere on the active domain. Notably, the intercept $(-\epsilon + 2\epsilon^2)$ is strictly negative.

This maps perfectly to the second normal-form matrix. The pure Nash equilibria are the coordinated states (A, B) and (B, A), yielding the optimal social cost of $2\epsilon^2$. However, the fully mixed interior Nash equilibrium (where both agents randomize uniformly) yields an expected social cost of $\epsilon + \epsilon^2$. The Mixed Price of Anarchy is therefore exactly $\frac{\epsilon+\epsilon^2}{2\epsilon^2} = \frac{1+\epsilon}{2\epsilon}$.

By taking the limit as $\epsilon \to 0^+$, the positive slope $(\epsilon - \epsilon^2)$ becomes arbitrarily small (approaching zero), yet the Mixed Price of Anarchy diverges to $+\infty$. This mathematically proves that classical PoA bounds collapse entirely when negative intercepts are permitted, even for arbitrarily flat, everywhere-positive latency functions. □

D.5 Proof of Theorem 10

Proof. We establish the theorem by explicitly constructing a smooth affine congestion game where the bounds guaranteed by robust Price of Anarchy are permissive, mathematically classifying strictly dominated outcomes as "approximately optimal."

Step 1: The Centralized Optimum and Min-Max Baseline
Consider a symmetric congestion game with N agents (assume N is even for simplicity) and $M = 2$ identical resources. Let the cost function for each resource be the standard identity affine function, $c_e(x) = x$.

The centralized social optimum s^* is achieved when the agents distribute themselves equally across both resources, inducing a load of $N/2$ on each. The corresponding optimal social cost is $C(s^*) = 2(N/2)^2 = N^2/2$.

We compute the baseline individual rationality threshold (the min-max safety cost) for an arbitrary agent i. The adversarial coalition of $N-1$ opponents maximizes agent i's minimum expected cost by uniformly distributing their mass across both resources, inducing an expected load of $(N-1)/2$ on each. Agent i's best response to this symmetric environment yields an expected cost of exactly $1+(N-1)/2=(N+1)/2$. Thus, any rationalizable state must guarantee agent i an expected cost no greater than $(N+1)/2$.

Step 2: Evaluation of the Robust PoA Bound
We evaluate the system under the robust Price of Anarchy framework. For affine congestion games, the canonical smoothness parameters are $(\lambda=5/3,\mu=1/3)$, yielding a robust PoA bound of $\frac{\lambda}{1-\mu}=5/2$. Consequently, the worst-case analysis framework guarantees, and inherently accepts as approximately optimal, any state distribution σ whose expected social cost satisfies $\mathbb{E}[C(\sigma)]\leq \frac{5}{2}C(s^*)=\frac{5}{4}N^2$.

Step 3: The Inclusion of Perfect Miscoordination
We identify a problematic state within this sanctioned set. Let s_{crowd} be the state of perfect miscoordination where all N agents select the exact same resource. The social cost of this state is $C(s_{\text{crowd}})=N^2$. Because $N^2\leq\frac{5}{4}N^2$, this state explicitly satisfies the smoothness inequality and is mathematically classified as a (λ,μ)-approximate optimal state.

However, at s_{crowd}, every single agent experiences a realized cost of N. Evaluating this against the individual rationality baseline yields a strict violation of the safety guarantee:

$$u_i^{\text{MinMaxCenter}}(s_{\text{crowd}})=\frac{N+1}{2}-N=-\frac{N-1}{2}=-\Theta(N)$$

Thus, the set of approximately optimal states defined by robust PoA guarantees structurally includes outcomes where every single agent suffers an efficiency loss arbitrarily worse than their absolute worst-case adversarial safety guarantee. □

The structural slackness identified in Theorem 10 is not an isolated artifact of affine costs; rather, it amplifies exponentially with the non-linearity of the game. We formalize this divergence below.

Proposition 1 (Polynomial Amplification of Structural Slackness)**.** *In smooth congestion games with polynomial cost functions of degree $d\geq 1$, the set of (λ,μ)-approximate optimal states includes outcomes where the violation of the min-max individual rationality baseline scales as $-\Theta(N^d)$.*

Proof. We extend the construction of Theorem 10 to a symmetric congestion game with N agents and $M=2$ identical resources, where the latency function is given by the monomial $c(x)=x^d$.

The centralized social optimum s^* again equidistributes the agents, yielding a social cost of $C(s^*)=2(N/2)^{d+1}=N^{d+1}/2^d$.

Next, we compute the exact min-max safety cost for an arbitrary agent i. The coalition of $N-1$ opponents seeks to choose a joint distribution over the resources to maximize the minimum expected cost of agent i. By the strict convexity of x^d for $d>1$, the variance of the load strictly increases the expected cost. Therefore, the absolute worst-case adversarial strategy is to perfectly correlate and assign all $N-1$ opponents to a single resource, chosen uniformly at random. Agent i's best response to this adversarial environment is to select either resource, encountering $N-1$ opponents with probability $1/2$ and 0 opponents with probability $1/2$. This establishes the exact min-max safety cost:

$$c_i(\text{minmax}_i^*,s_{-i})=\frac{1}{2}(N^d)+\frac{1}{2}(1^d)=\frac{N^d+1}{2}$$

We now evaluate the state of perfect miscoordination, s_{crowd}, where all N agents select the exact same resource. The social cost of this state is $C(s_{\text{crowd}})=N^{d+1}$. The ratio of this cost to the social optimum is exactly $C(s_{\text{crowd}})/C(s^*)=2^d$.

It is a standard result in the Price of Anarchy literature that the robust PoA bound Φ_d for atomic congestion games with polynomial costs of degree d scales as $\Phi_d=\Theta\left(\frac{d}{\ln d}\right)^d$ Christodoulou and Koutsoupias (2005a); Aland et al. (2011). Because Φ_d strictly dominates 2^d for all $d\geq 1$ (e.g.,

$\Phi_1 = 2.5 > 2$, $\Phi_2 \approx 9.6 > 4$), the state s_{crowd} strictly satisfies the robust PoA inequality:

$$\mathbb{E}[C(s_{\text{crowd}})] \leq \Phi_d C(s^*)$$

Consequently, s_{crowd} is mathematically validated by the worst-case framework as a (λ, μ)-approximate optimal state. Yet, evaluating the individual performance of the agents at this state yields:

$$u_i^{\text{MinMaxCenter}}(s_{\text{crowd}}) = \frac{N^d + 1}{2} - N^d = -\frac{N^d - 1}{2} = -\Theta(N^d)$$

Thus, as the degree of the cost functions increases, the analytical bounds provided by robust PoA become exponentially decoupled from the threshold of baseline individual rationality. □

D.6 Proof of Lemma 1

To rigorously evaluate the min-max safety cost subject to independent randomization (as claimed in Section D.2), we must first formalize the exact probabilistic nature of the load induced by an adversarial coalition.

Lemma 2 (Poisson Binomial Opponent Load)**.** *Let $\pi_{-i} \in \prod_{j \neq i} \Delta(M)$ be an arbitrary mixed strategy profile for the $N-1$ opponents of agent i in an M-bin congestion game. For any fixed bin $e \in [M]$, let X_e denote the total load induced on bin e by the opponents. Then X_e follows a Poisson Binomial distribution parameterized by $\mathbf{p}_e = (p_{j,e})_{j \neq i}$, where $p_{j,e}$ is the probability that agent j selects bin e.*

Proof. By the definition of a mixed strategy profile in a normal-form game, the agents randomize their actions independently. Let $s_j \in [M]$ denote the pure strategy realized by agent j. The joint probability measure over the opponents' pure action profiles $s_{-i} \in [M]^{N-1}$ is the exact product measure: $\mathbb{P}[s_{-i}] = \prod_{j \neq i} p_{j,s_j}$.

For each opponent $j \neq i$, define the indicator random variable $B_{j,e} = \mathbb{I}\{s_j = e\}$. Under the product measure, the sequence of variables $\{B_{j,e}\}_{j \neq i}$ is mutually independent. The expected value of each indicator is $\mathbb{E}[B_{j,e}] = \mathbb{P}[s_j = e] = p_{j,e}$. Thus, $B_{j,e} \sim \text{Bernoulli}(p_{j,e})$.

The total opponent load on bin e is the algebraic sum of these indicators: $X_e = \sum_{j \neq i} B_{j,e}$. Because X_e is the sum of $N-1$ mutually independent Bernoulli random variables, it strictly satisfies the definition of a Poisson Binomial random variable. Thus, $X_e \sim \text{PB}(\mathbf{p}_e)$. □

We are now ready to prove Lemma 1 that we remind for clarity below:

Lemma 1 (Independent Min-Max Equivalence)**.** *In symmetric balls and bins games with strictly convex monomial costs $c(x) = x^d$ ($d \geq 2$), the expected cost evaluated at the fully mixed interior Nash equilibrium is mathematically identical to the exact independent min-max safety cost.*

Proof. Consider an agent i facing $N-1$ opponents. The adversarial coalition selects independent mixed strategies $\mathbf{p}_j \in \Delta(M)$ for each opponent $j \neq i$ to maximize agent i's minimum expected cost across all bins. The adversary's objective is to maximize:

$$\min_{e \in [M]} C_e(\mathbf{p}) \quad \text{where} \quad C_e(\mathbf{p}) = \mathbb{E}\left[(1 + X_e)^d\right]$$

For any fixed bin e, let $\mu_e = \mathbb{E}[X_e] = \sum_{j \neq i} p_{j,e}$ denote the expected opponent load. Because each of the $N-1$ opponents must allocate exactly 1 unit of probability mass across the M bins, the sum of the expected loads is strictly conserved: $\sum_{e=1}^{M} \mu_e = N - 1$.

By Lemma 2, the load $X_e \sim \text{PB}(\mathbf{p}_e)$. By Hoeffding's Theorem (1956) on the extrema of expected values of convex functions of independent Bernoulli sums, the expectation $C_e(\mathbf{p})$ is strictly Schur-concave with respect to the probability vector $\mathbf{p}_e$. Consequently, for any fixed mean μ_e, the expectation $C_e(\mathbf{p})$ is strictly upper-bounded by evaluating the expectation under a homogeneous Binomial distribution with the identical mean:

$$C_e(\mathbf{p}) \leq \mathbb{E}\left[(1 + Y(\mu_e))^d\right] := g(\mu_e)$$

where $Y(\mu_e) \sim \text{Binomial}\left(N-1, \frac{\mu_e}{N-1}\right)$.

To evaluate the adversary's objective, we establish the monotonic behavior of the expected cost. For any two probabilities $p < p'$, the distribution Binomial$(N-1, p')$ strictly first-order stochastically dominates (FOSD) the distribution Binomial$(N-1, p)$. Because the latency function $(1+x)^d$ is strictly monotonically increasing for $x \geq 0$, its expectation under the Binomial distribution strictly increases with the success probability. Consequently, the composite function $g(\mu_e)$ is strictly monotonically increasing with respect to the mean μ_e.

We bound the adversary's objective using a direct chain of inequalities. Because $\sum_{e=1}^{M} \mu_e = N-1$, the minimum expected load across the bins is strictly bounded: $\min_{e\in[M]} \mu_e \leq \frac{N-1}{M}$. Thus:

$$\min_{e\in[M]} C_e(\mathbf{p}) \leq \min_{e\in[M]} g(\mu_e) = g\left(\min_{e\in[M]} \mu_e\right) \leq g\left(\frac{N-1}{M}\right)$$

The upper bound $g\left(\frac{N-1}{M}\right)$ represents the expected cost when the opponent load on the target bin is distributed according to Binomial$(N-1, 1/M)$. By the conditions of equality in Hoeffding's Theorem, this exact maximum is achieved if and only if the adversary sets the success probabilities to be perfectly homogeneous: $p_{j,e} = 1/M$ for all j, e.

This homogeneous strategy profile is identically the fully mixed interior Nash equilibrium. Therefore, the adversary's optimal independent min-max strategy strictly collapses to the mixed Nash equilibrium, proving the exact mathematical equivalence. □

D.7 Proof of Theorem 11

Proof. We establish the exact Average Price of Anarchy by computing the Lebesgue measure of the regions of attraction for the pure Nash equilibria. We proceed by establishing a universal geometric upper bound for all w, followed by an exact measure-theoretic integration for $w = 2$.

Step 1: Volume-Preserving Diffeomorphism
Let x and y denote the probability that the row and column agents, respectively, play their first strategy in the 2×2 affine congestion game (Family 1 from Figure 6). The system admits an interior Nash equilibrium at $(x, y) = (\frac{w}{w+1}, \frac{1}{w+1})$.

To symmetrically align the phase space, we apply the change of variables $x' = x$ and $y' = 1 - y$. The Jacobian determinant of this transformation is identically -1, ensuring the mapping is a volume-preserving diffeomorphism. The Lebesgue measure of any region of attraction in the transformed space is strictly identical to its measure in the original game. This transformation maps the congestion game into a classic w-coordination game, possessing two strict pure Nash equilibria at $(0, 0)$ and $(1, 1)$, and a symmetric interior saddle at $(\frac{w}{w+1}, \frac{w}{w+1})$.

Step 2: Universal Upper Bound via Forward-Invariant Triangles
To establish the universal upper bound of 2 on the Average Price of Anarchy for this family, we reproduce a geometric bounding argument originally developed by Panageas and Piliouras Panageas and Piliouras (2016a). We include this analysis to ensure our measure-theoretic evaluation remains rigorously self-contained.

In the transformed coordinates, the socially optimal pure Nash equilibrium resides at $(0, 0)$ with a normalized social cost of 1, while the suboptimal pure Nash equilibrium resides at $(1, 1)$ with a cost of w. Consider the two geometric triangles formed by the vertices $\{(0, 0), (1, 0), (\frac{w}{w+1}, \frac{w}{w+1})\}$ and $\{(0, 0), (0, 1), (\frac{w}{w+1}, \frac{w}{w+1})\}$. Evaluation of the replicator vector field along the boundaries of these triangles confirms they are strictly forward-invariant. Because the interior Nash equilibrium is a strict saddle (and thus its stable manifold has Lebesgue measure zero), the interior of these triangles resides entirely within the basin of attraction of the optimal equilibrium $(0, 0)$.

The combined Lebesgue measure (area) of these two triangles is exactly $\frac{w}{w+1}$. By pessimistically assigning the maximum possible cost w to the entire remaining state space $1 - \frac{w}{w+1} = \frac{1}{w+1}$, we

obtain a strict upper bound on the expected social cost under a uniform prior:

$$\text{APoA} \leq \left(\frac{w}{w+1} \times 1\right) + \left(\frac{1}{w+1} \times w\right) = \frac{2w}{w+1}$$

For all $w \geq 1$, we have $\frac{2w}{w+1} < 2$. This formally guarantees that the Average PoA is strictly bounded by 2 under the standard algebraic constraints.

Step 3: The Exact Invariant Manifold for $w = 2$
To prove the lower bound on inefficiency, we evaluate the specific instance where $w = 2$. The continuous-time replicator equations in the transformed coordinates are given by $\dot{x} = x(1-x)(3y-2)$ and $\dot{y} = y(1-y)(3x-2)$. By dividing the differential equations, the time dependency is strictly eliminated, yielding the separable ordinary differential equation:

$$\frac{3x-2}{x(1-x)}dx = \frac{3y-2}{y(1-y)}dy$$

Integrating both sides by partial fractions explicitly constructs the time-invariant function $I(x, y)$ that remains strictly constant along any trajectory:

$$I(x, y) = -2\ln(x) - \ln(1-x) + 2\ln(y) + \ln(1-y) = C$$

The stable manifold (the separatrix) of the interior saddle at $(2/3, 2/3)$ must satisfy $I(2/3, 2/3) = 0$. Consequently, the exact algebraic equation defining the separatrix that partitions the state space is $x^2(1-x) = y^2(1-y)$. Solving this algebraic curve for y yields the closed-form boundary of the basin of attraction:

$$y(x) = \frac{1}{2}\left(1 - x + \sqrt{1 + 2x - 3x^2}\right)$$

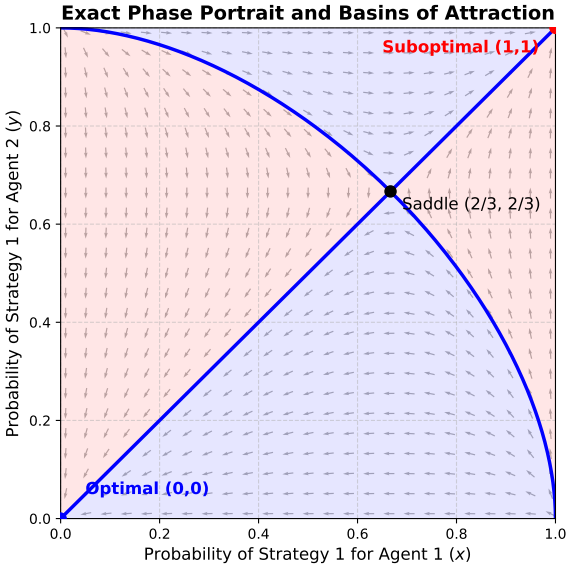

Figure 7: **Exact Measure-Theoretic Basins of Attraction.** The simulated phase portrait of the replicator vector field after the volume-preserving transformation. The blue curve represents the mathematically exact invariant stable manifold $x^2(1-x) = y^2(1-y)$ defining the separatrix. Because the saddle point is positioned at $(2/3, 2/3)$, it severely constricts the basin of the high-potential (worst-case) equilibrium at $(1, 1)$. Trajectories originating below the separatrix converge to the socially optimal equilibrium at $(0, 0)$. The forward-invariant triangles described in Step 2 are strictly contained within the optimal basin.

Step 4: Measure-Theoretic Integration
Under a uniform Lebesgue measure prior over initial conditions, the probability of converging to the

socially optimal equilibrium $(0,0)$ is strictly equal to the area under the separatrix $y(x)$. Integrating this closed-form expression over the unit interval evaluates exactly to:

$$\int_0^1 \frac{1}{2}\left(1 - x + \sqrt{1 + 2x - 3x^2}\right) dx = \frac{1}{27}\left(9 + 2\sqrt{3}\pi\right) \approx 0.7364$$

Consequently, the region of attraction for the worst-case, high-potential Nash equilibrium at $(1,1)$ has a Lebesgue measure of $1 - 0.7364 = 0.2636$.

Step 5: Average Price of Anarchy Evaluation
Let the optimal social cost at $(0,0)$ be normalized to 1. In this configuration, the social cost of the worst-case Nash equilibrium at $(1,1)$ evaluates to $w = 2$. The Average Price of Anarchy is the expected social cost evaluated over the Lebesgue measure of the respective attractors:

$$\text{APoA} = (0.7364 \times 1) + (0.2636 \times 2) = 1.2636$$

The strict analytical evaluation of the integral guarantees this average-case inefficiency bound is robust and exact.

For the second parametric family in Figure 6, an identical analysis reveals that the mixed Nash equilibrium possesses a region of attraction of exactly Lebesgue measure zero. Consequently, almost all initial conditions converge to the optimal pure Nash equilibrium, yielding an APoA of strictly 1.

Step 6: The Collapse of APoA under Data-Driven Relaxations
Finally, we establish that Average PoA is structurally unbounded when the non-negativity constraint on affine intercepts is relaxed. Consider the first family of games, and subtract a uniform constant $1 - \epsilon$ from all payoff entries.

To formally execute the uniform translation of all payoff entries by $1 - \epsilon$, we apply a uniform constant shift to the underlying affine latency functions. Because every strategy in the constructed congestion game utilizes exactly two resources, subtracting a constant $C = \frac{1-\epsilon}{2}$ from each resource cost exactly subtracts $1 - \epsilon$ from the total cost of any strategy profile. We define the translated resource costs as $c'_e(x) = c_e(x) - \frac{1-\epsilon}{2}$.

To guarantee that the resulting latency functions remain strictly positive and strictly monotonically increasing on the physical domain ($x \geq 1$), we constrain the free parameter δ. By choosing $\delta = \frac{1}{2} - \frac{\epsilon}{4}$, we satisfy both requirements simultaneously: the base cost evaluates to $c'_e(1) = \delta - C = \frac{\epsilon}{4} > 0$, ensuring strict positivity, and the minimum marginal cost increment evaluates to $c'_{BA}(2) - c'_{BA}(1) = 1 - 2\delta = \frac{\epsilon}{2} > 0$, preserving strict monotonicity. Consequently, this transformation yields a mathematically valid, everywhere-positive affine congestion game.

Crucially, the continuous-time replicator vector field is invariant to uniform translations of the payoff matrix. Because the vector field is identical, the exact invariant manifold $I(x,y)$, the separatrix $y(x)$, and the resulting Lebesgue measures of the basins of attraction (0.7364 and 0.2636) remain perfectly unchanged.

However, the social cost metric is strictly non-invariant to this translation. After the translation, the cost of the socially optimal state approaches ϵ, while the cost of the suboptimal state approaches $1 + \epsilon$. As $\epsilon \to 0^+$, the ratio of the suboptimal cost to the optimal cost diverges to $+\infty$. Because the suboptimal state retains a strictly positive basin of attraction (0.2636), the expected social cost evaluated over the Lebesgue prior diverges to $+\infty$. □

E Proof of Theorem 5

To formally establish Theorem 5, we prove two fundamental structural lemmas. The first establishes the existence of non-rationalizable Coarse Correlated Equilibria (CCEs) fully supported on strictly dominated strategies in smooth games. The second establishes that strongly uncoupled learning dynamics cannot prevent convergence to these arbitrary target CCEs.

Lemma 3 (Unnatural CCEs in Smooth Games)**.** *There exist smooth games (e.g., linear congestion games and valid utility games) possessing strong coarse correlated equilibria (CCE) where all agents assign strictly positive probability exclusively to strictly dominated strategies. Furthermore, there exist extreme CCEs (vertices of the feasible CCE polytope) exhibiting this exact property.*

Proof. We prove this lemma constructively for both canonical classes of smooth games.

Step 1: Construction in Linear Congestion Games
We define a symmetric congestion game with two agents and four strategies per agent. The game comprises three distinct resources. Resources e_1 and e_2 have identical affine latency functions $c(x) = x$. The third resource, e_d, acts as a fixed-cost penalty with $c(x) = \epsilon$, where $0 < \epsilon < 1/2$. The strategy set for both agents is $S = \{\{e_1\}, \{e_1, e_d\}, \{e_2\}, \{e_2, e_d\}\}$.

By construction, strategy $\{e_1, e_d\}$ is strictly dominated by $\{e_1\}$, and $\{e_2, e_d\}$ is strictly dominated by $\{e_2\}$, as the inclusion of e_d incurs a strictly positive additional cost regardless of the opponent's action. We propose the correlated distribution π that assigns $1/2$ probability to the outcome $(\{e_1, e_d\}, \{e_2, e_d\})$ and $1/2$ probability to $(\{e_2, e_d\}, \{e_1, e_d\})$. Under π, the expected cost for each agent is exactly $1 + \epsilon$.

We evaluate the CCE inequalities. The optimal unconditional deviation for either agent is to permanently play an undominated strategy, e.g., $\{e_1\}$. Against the marginal distribution of the opponent (which places $1/2$ probability on utilizing e_1), the expected cost of this deviation is $\frac{1}{2}(1)+\frac{1}{2}(2) = 1.5$. Since we constrained $\epsilon < 1/2$, it holds that $1 + \epsilon < 1.5$. Thus, any unilateral deviation strictly increases the expected cost, mathematically confirming that π is a strong CCE supported entirely on strictly dominated strategies.

Step 2: Construction in Valid Utility Games
To eliminate the ambiguity inherent in the inequality-based definition of valid utility games, we construct an exact *basic utility system* Vetta (2002). We define a ground set of six independent elements $E = \{a_1, a_2, b_1, b_2, c_1, c_2\}$. We define a modular (and thus submodular) social welfare function $W(s) = \sum_{e \in U(s)} v(e)$, with values $v(a_1) = v(a_2) = v(b_1) = v(b_2) = \epsilon$, and $v(c_1) = v(c_2) = 1$, where $0 < \epsilon < 1/2$.

The strategy set for agent 1 is $S_1 = \{\{a_1, c_1\}, \{c_1\}, \{a_2, c_2\}, \{c_2\}\}$, and for agent 2 is $S_2 = \{\{b_1, c_1\}, \{c_1\}, \{b_2, c_2\}, \{c_2\}\}$. In a basic utility system, the private utility of agent i is defined exactly as their marginal contribution to the social welfare: $u_i(s) = W(s) - W(\emptyset, s_{-i})$. Because W is modular, this evaluates exactly to the sum of the values of the elements agent i *exclusively* secures. This exact specification natively satisfies both defining inequalities of a valid utility game.

By construction, the singleton strategy $\{c_1\}$ is strictly dominated by $\{a_1, c_1\}$ for agent 1. Because the element a_1 is uniquely available to agent 1, agent 1 secures a_1 exclusively whenever it is played. Thus, playing $\{a_1, c_1\}$ guarantees a payoff exactly ϵ strictly greater than playing $\{c_1\}$, regardless of agent 2's strategy. By identical logic, $\{c_2\}$ is strictly dominated by $\{a_2, c_2\}$ for agent 1, and the corresponding singleton strategies are strictly dominated for agent 2.

Consider the distribution π assigning $1/2$ probability to the outcome $(\{c_1\}, \{c_2\})$ and $1/2$ probability to $(\{c_2\}, \{c_1\})$. Under π, both agents exclusively secure an element of value 1 with probability 1. The expected utility for each agent is exactly 1.

We evaluate the CCE condition. The optimal unconditional deviation for agent 1 is to permanently play a strictly dominating strategy, e.g., $\{a_1, c_1\}$. Against the marginal distribution of π, agent 2 plays $\{c_2\}$ with probability $1/2$, allowing agent 1 to exclusively secure $\{a_1, c_1\}$ for a payoff of $1 + \epsilon$. With probability $1/2$, agent 2 plays $\{c_1\}$, causing agent 1 to exclusively secure only a_1 for a payoff of ϵ. The expected utility of this optimal deviation is $\frac{1}{2}(1 + \epsilon) + \frac{1}{2}(\epsilon) = 0.5 + \epsilon$.

Because we constrained $\epsilon < 1/2$, we have $0.5 + \epsilon < 1$. Thus, any unilateral deviation strictly decreases expected utility, confirming the distribution is a strong CCE.

Step 3: Extremality via Polytope Geometry
Finally, we establish that these properties hold for extreme CCEs. The set of valid CCEs constitutes a bounded convex polytope defined by linear inequalities. We formulate a linear program over this polytope with the objective of maximizing the sum of probabilities assigned to the outcomes where all agents play their strictly dominated strategies. We have demonstrated in Steps 1 and 2 that a feasible solution exists with an objective value of exactly 1. Because the maximum of a linear objective over a bounded convex polytope is guaranteed to be achieved at an extreme point (vertex), there must exist an extreme CCE that places probability 1 exclusively on strictly dominated profiles. □

Lemma 4 (Strongly Uncoupled Convergence to Target CCEs)**.** *Given any coarse correlated equilibrium (CCE) defined by rational probabilities in a normal-form game, there exists a set of N strongly uncoupled no-regret learning algorithms such that, in self-play, they converge to the target CCE at an optimal time-average rate of $O(1/T)$.*

Proof. We prove the lemma by explicitly constructing a compliant algorithm.

Step 1: Cyclic Representation of the Target Equilibrium
Let C be a target CCE. Because the probabilities defining C are rational, there exists an integer K such that C can be represented as a uniform distribution over a deterministic, cyclic sequence of pure action profiles $V = (v^0, v^1, \dots, v^{K-1})$.

Step 2: Construction of the Grim-Trigger Dynamic
We construct a learning dynamic for each agent i operating in two strict phases:

- **Cooperative Phase:** At time t, agent i plays their assigned coordinate from the profile $v^{t \pmod K}$. The agent then observes their realized cost vector c^t. If this observed cost vector exactly matches a privately stored, pre-computed target cost vector $C_i^{t \pmod K}$, agent i remains in the cooperative phase.
- **Penalty Phase:** If the observed cost vector deviates from the expected target cost vector $C_i^{t \pmod K}$, agent i immediately and permanently switches to a standard, optimal no-regret algorithm (e.g., Optimistic Mirror Descent), initializing it as if time t were round 1.

Step 3: Verification of the Optimal Regret Bound
In self-play, all agents adhere to the sequence V, and the penalty phase is never triggered. The empirical distribution of play cycles through V, converging to the target CCE exactly. The total external regret evaluated at any time T is strictly bounded by $O(K)$ (the length of the cycle), yielding an optimal time-average regret rate of $O(1/T)$.

Under arbitrary adversarial play, the total regret is partitioned into the rounds prior to the first observed deviation, and the rounds following it. Because C is a valid CCE, the regret accumulated during the cooperative phase is strictly bounded by $O(K)$. Once a deviation is detected, the agent assumes worst-case adversarial play and transitions to a standard no-regret algorithm (e.g., Optimistic Mirror Descent). While such algorithms can achieve $O(\text{polylog}(T))$ regret under mutual predictable play, against an arbitrary adversary they guarantee the optimal worst-case bound of $O(\sqrt{T})$. The total regret is therefore bounded by $O(K)+O(\sqrt{T}) = O(\sqrt{T})$, mathematically establishing the algorithm as Hannan consistent under all conditions.

Step 4: Satisfaction of Strongly Uncoupled Constraints
Crucially, this dynamic satisfies the strict informational constraints of *strongly uncoupled learning* Daskalakis et al. (2011). Agent i possesses zero knowledge of the opponents' utility functions or the full game matrix. The private storage requirement is strictly bounded by $O(1)$, as the agent only needs to store their private K-step sequence of scheduled actions and expected target cost vectors. Verification requires only a direct numerical comparison of the observed vector against the stored target, demanding absolutely no structural knowledge of the opponents' matrices or the global sequence V. The computations performed in each round are highly efficient, running in polynomial time relative to the action space. □

Proof of Theorem 5

Proof. The theorem follows immediately from the synthesis of Lemma 3 and Lemma 4. Lemma 3 guarantees the existence of smooth games possessing extreme strong CCEs supported entirely on strictly dominated strategies, which by construction are representable by cyclic sequences of length $K = 2$. Lemma 4 demonstrates that strongly uncoupled, optimal $O(1/T)$ no-regret algorithms can be explicitly constructed to converge to these exact cyclic sequences. Consequently, optimal no-regret learning in smooth games provides no structural guarantee against convergence to strictly dominated, non-rationalizable outcomes. □

F Proximal Refinements and Strictly Dominated Strategies

Recognizing the permissive nature of Coarse Correlated Equilibria, recent literature has attempted to bridge the gap toward Correlated Equilibria (CE) by introducing intermediate solution concepts rooted in continuous optimization. Chief among these are Semicoarse Correlated Equilibria (SCE) Ahunbay and Bichler (2025) and Proximal Correlated Equilibria (PCE) Cai et al. (2025), together with their Bregman generalization (BPCE), induced by the Bregman proximal regret of Cai et al. (2025). These refinements demand that agents exhibit no regret against counterfactual strategy modifications generated by symmetric affine maps (SCE) or by (Bregman) proximal operators (PCE, BPCE).

From the perspective of continuous online optimization, these frameworks are mathematically elegant: standard, uncoupled learning algorithms (such as Projected Gradient Descent or Mirror Descent) drive the average proximal regret to zero at $O(1/\sqrt{T})$ rates Cai et al. (2025). However, the constraints certified by these guarantees are evaluated on the players' mixed iterates (or, for SCE, on affine modifications of correlated pure actions), rather than on the best-response structure of the underlying game. We show that this continuous-to-discrete translation structurally permits non-rationalizable outcomes.

Definitions. Let $X_i = \Delta(S_i)$ and $X = \prod_i X_i$. For $\rho \in [0, 1)$, let $\mathcal{F}_\rho(X_i)$ denote the class of proper, lower semi-continuous, ρ-weakly convex functions $f : X_i \to \mathbb{R} \cup \{+\infty\}$, i.e., functions for which $f + \frac{\rho}{2}\|\cdot\|^2$ is convex. Let $\phi : X_i \to \mathbb{R}$ be a distance-generating function that is 1-strongly convex with respect to the same norm $\|\cdot\|$, with Bregman divergence $D_\phi(y\|x)$. For $f \in \mathcal{F}_\rho(X_i)$, the Bregman proximal operator Cai et al. (2025) is

$$\operatorname{prox}_f^\phi(x) = \operatorname*{argmin}_{y \in X_i} \{f(y) + D_\phi(y\|x)\},$$

which is well defined since the objective is $(1-\rho)$-strongly convex. Following the Φ-equilibrium framework of Cai et al. (2025), we call a distribution $\mu \in \Delta(X)$ over *mixed* strategy profiles a Bregman Proximal Correlated Equilibrium if, for every player i, every $\rho \in [0, 1)$ and every $f \in \mathcal{F}_\rho(X_i)$,

$$\mathbb{E}_{x\sim\mu}\Big[u_i\big(\operatorname{prox}_f^\phi(x_i), x_{-i}\big) - u_i(x)\Big] \le 0. \tag{15}$$

For $\phi = \frac{1}{2}\|\cdot\|_2^2$ this is the (Euclidean) Proximal Correlated Equilibrium of Cai et al. (2025).[5] A distribution $\sigma \in \Delta(S)$ over pure profiles is identified with $\mu_\sigma = \sum_s \sigma(s)\,\delta_{(e_{s_1},\dots,e_{s_n})}$, and every $\mu \in \Delta(X)$ induces the outcome distribution $\bar{\mu}(s) = \mathbb{E}_{x\sim\mu}\big[\prod_j x_j(s_j)\big] \in \Delta(S)$, which describes how play is actually realized. Finally, $\sigma \in \Delta(S)$ is a Semicoarse Correlated Equilibrium if no player benefits from applying to her recommended action any symmetric affine strategy modification $x_i \mapsto Mx_i + b$ with M symmetric Ahunbay and Bichler (2025). Since proximal operators generate all such modifications, PCE $\subseteq$ SCE in normal-form games Cai et al. (2025); the inclusion is strict (Remark 5).

We establish that both refinements nonetheless permit strictly dominated play, through two distinct mechanisms.

Theorem 12 (SCE/PCE Support Strictly Dominated Strategies)**.** *Consider the Euclidean distance-generating function $\phi = \frac{1}{2}\|\cdot\|_2^2$.*

(i) SCE: *There exist a symmetric two-player game and a Semicoarse Correlated Equilibrium $\sigma \in \Delta(S)$ that assigns probability 0.2 to a strictly dominated strategy.*

(ii) PCE: *There exist a three-player game and a Proximal Correlated Equilibrium $\mu \in \Delta(X)$, satisfying* (15) *for every $\rho \in [0, 1)$ and every $f \in \mathcal{F}_\rho$ simultaneously, whose induced outcome distribution $\bar{\mu}$ assigns probability greater than 0.04 to a strictly dominated strategy. Moreover, μ is the time-average distribution of a periodic orbit of continuous-time projected gradient play.*

The proofs are given in Appendix F.1. Part (i) exhibits an *affine shielding effect*: an admissible modification can shift mass from the dominated action A^- to the dominating action A only by also shifting mass from A back to A^- (its symmetric part) or by also involving the third action C (its

[5] Variants of the definition fix ρ or restrict the test class. Requiring (15) for all $\rho \in [0, 1)$ only strengthens the concept, so the existence result below applies to every such variant.

affine part), and under the constructed distribution each of these side effects costs more than the shift gains. Part (ii) exhibits a *projection shielding effect*: under Euclidean projected gradient play, the projection subtracts a common multiplier from all eligible coordinates, so the dominated action becomes active as soon as this multiplier falls below its payoff, even though a strictly better action is active. Meanwhile, every proximal deviation gain along the resulting periodic orbit is bounded above by the negative time derivative of a Moreau envelope, and therefore integrates to at most zero over a period. The two parts cannot be merged into a single construction: restricted to distributions over pure profiles, PCE does exclude strictly dominated strategies (Remark 5 in Appendix F.1). The phenomenon in part (ii) is driven precisely by correlation over mixed strategies—the object that no-regret dynamics produce.

Because the Euclidean divergence $D_\phi(y\|x) = \frac{1}{2}\|y - x\|_2^2$ is an admissible Bregman divergence, PCE is the special case of BPCE with $\phi = \frac{1}{2}\|\cdot\|_2^2$. The Bregman framework therefore inherits part (ii).

Corollary 1. *The Bregman Proximal Correlated Equilibrium framework does not, in general, exclude strictly dominated strategies: for the Euclidean distance-generating function, there exist BPCE whose induced outcome distribution assigns strictly positive probability to a strictly dominated strategy.*[6]

F.1 Proofs of Theorem 12

F.2 Part (i): Semicoarse Correlated Equilibria

Proof of part (i). We define a symmetric two-player normal-form game and verify that a specific distribution, which places strictly positive mass on a strictly dominated strategy, satisfies all conditions of a Semicoarse Correlated Equilibrium (SCE).

Step 1: Game Definition and Dominated Strategy Support
Let the action set of both players be $S_i = \{A, A^-, C\}$. The utility function of Player 1 is

$$\begin{aligned} u_1(A,A) &= 1, & u_1(A,A^-) &= 1, & u_1(A,C) &= 0 \\ u_1(A^-,A) &= 0.8, & u_1(A^-,A^-) &= 0.8, & u_1(A^-,C) &= -0.2 \\ u_1(C,A) &= 0, & u_1(C,A^-) &= 0, & u_1(C,C) &= 1 \end{aligned}$$

and $u_2(a_1,a_2) = u_1(a_2,a_1)$. Action A strictly dominates A^- by a margin of $\epsilon = 0.2$ against every opponent strategy. Consider the correlated distribution

$$\sigma(A,A) = 0.4, \quad \sigma(A^-,A^-) = 0.2, \quad \sigma(C,C) = 0.4,$$

which assigns probability 0.2 to the profile (A^-,A^-), so that each player plays the strictly dominated action A^- with probability 0.2. By symmetry it suffices to verify the conditions for Player 1.

Step 2: Primal Verification via the Generating Modifications
Define the marginal expected gain of deviating from a to a', weighted by the probability of a:

$$G(a',a) = \sum_{a_2 \in S_2} \sigma(a,a_2)\left[u_1(a',a_2) - u_1(a,a_2)\right].$$

Then

$$\begin{aligned} G(A,A^-) &= 0.04 & G(A^-,A) &= -0.08 & G(C,A) &= -0.40 \\ G(A,C) &= -0.40 & G(C,A^-) &= -0.16 & G(A^-,C) &= -0.48 \end{aligned}$$

By Theorem 3.3 of Ahunbay and Bichler (2025), σ is an SCE if and only if the expected gain is non-positive for every modification in the finite generating family of that theorem, namely the cycle deviations and the subset deviations of proper subsets. For a modification that maps each action a to a distribution $P(\cdot \mid a)$, the expected gain is $\sum_a \sum_{a' \neq a} P(a' \mid a)\, G(a',a)$.

Cycle Deviations:

[6] Whether the same holds for steep distance-generating functions such as negative entropy is left open. For such ϕ, with the convention $D_\phi(y\|x) = +\infty$ whenever $\mathrm{supp}(y) \not\subseteq \mathrm{supp}(x)$, we have $D_\phi(y\|e_a) = +\infty$ for every $y \neq e_a$, so the Bregman proximal operator fixes every vertex at which f is finite, and the BPCE constraints are vacuous on pure-profile distributions.

- 2-cycle $\{A, A^-\}$: $G(A^-, A) + G(A, A^-) = -0.08 + 0.04 = -0.04 \leq 0$
- 2-cycle $\{A, C\}$: $G(C, A) + G(A, C) = -0.40 - 0.40 = -0.80 \leq 0$
- 2-cycle $\{A^-, C\}$: $G(C, A^-) + G(A^-, C) = -0.16 - 0.48 = -0.64 \leq 0$
- 3-cycle $\{A, A^-, C\}$ (each action mapped to the uniform distribution over the other two): $\frac{1}{2}\left[G(A^-, A) + G(C, A) + G(A, A^-) + G(C, A^-) + G(A, C) + G(A^-, C)\right] = -0.74 \leq 0$

Subset Deviations (Sizes 1 and 2):

- $\{A^-\} \to \text{Unif}(\{A, C\})$: $0.5\, G(A, A^-) + 0.5\, G(C, A^-) = 0.02 - 0.08 = -0.06 \leq 0$
- $\{A\} \to \text{Unif}(\{A^-, C\})$: $0.5\, G(A^-, A) + 0.5\, G(C, A) = -0.04 - 0.20 = -0.24 \leq 0$
- $\{C\} \to \text{Unif}(\{A, A^-\})$: $0.5\, G(A, C) + 0.5\, G(A^-, C) = -0.20 - 0.24 = -0.44 \leq 0$
- $\{A^-, C\} \to A$: $G(A, A^-) + G(A, C) = 0.04 - 0.40 = -0.36 \leq 0$
- $\{A, C\} \to A^-$: $G(A^-, A) + G(A^-, C) = -0.08 - 0.48 = -0.56 \leq 0$
- $\{A, A^-\} \to C$: $G(C, A) + G(C, A^-) = -0.40 - 0.16 = -0.56 \leq 0$

Every generating modification yields a non-positive expected gain, so σ is an SCE.

Step 3: Verification via the Extended Formulation
For completeness we also give an exact certificate for the polynomial-size extended formulation of the SCE constraints (Theorem 3.4 in Ahunbay and Bichler (2025)): σ is an SCE if there exist $\gamma_1(a', a) \geq 0$ and $\kappa_1(a, a')$ satisfying

$$\gamma_1(a', a) + \kappa_1(a, a') - \kappa_1(a', a) + G(a', a) = 0 \quad \forall a \neq a', \tag{16}$$

$$\sum_{a \neq a'} \gamma_1(a', a) + \bar{G}(a') = 0 \quad \forall a', \tag{17}$$

where $\bar{G}(a') = \sum_{a \in S} \sigma(a)\big[u_1(a', a_2) - u_1(a)\big]$ is the unconditional gain of always playing a'. The expected utility under σ is $U(\sigma) = 0.4(1) + 0.2(0.8) + 0.4(1) = 0.96$, and

$$\begin{aligned}
\bar{G}(A) &= (0.4(1) + 0.2(1) + 0.4(0)) - 0.96 = -0.36 \\
\bar{G}(A^-) &= (0.4(0.8) + 0.2(0.8) + 0.4(-0.2)) - 0.96 = -0.56 \\
\bar{G}(C) &= (0.4(0) + 0.2(0) + 0.4(1)) - 0.96 = -0.56
\end{aligned}$$

The assignment

$$\begin{aligned}
\gamma_1(A, A^-) &= 0.04, & \gamma_1(A^-, A) &= 0, \\
\gamma_1(A, C) &= 0.32, & \gamma_1(C, A) &= 0.48, \\
\gamma_1(A^-, C) &= 0.56, & \gamma_1(C, A^-) &= 0.08
\end{aligned}$$

is non-negative and satisfies (17). Choosing κ_1 antisymmetric ($\kappa_1(x, y) = -\kappa_1(y, x)$), constraint (16) then determines

$$\kappa_1(A, A^-) = 0.04, \qquad \kappa_1(A, C) = -0.04, \qquad \kappa_1(C, A^-) = -0.04,$$

and all six constraints (16) hold with equality. This confirms that σ is an SCE assigning probability 0.2 to a strictly dominated strategy. □

F.3 Part (ii): Proximal Correlated Equilibria

Notation. For the simplex $\Delta_d \subset \mathbb{R}^d$ and $x \in \Delta_d$, let $T_\Delta(x)$ and $N_\Delta(x)$ denote the tangent and normal cones, and Π_K the Euclidean projection onto a closed convex set K. For $g \in \mathbb{R}^d$, the Karush–Kuhn–Tucker conditions give

$$\big(\Pi_{T_\Delta(x)}[g]\big)_k = \begin{cases} g_k - \lambda & \text{if } x_k > 0, \\ (g_k - \lambda)_+ & \text{if } x_k = 0, \end{cases} \tag{18}$$

where λ is the unique scalar for which the entries sum to zero; equivalently, λ is the mean of g over the active set $I = \{k : x_k > 0\} \cup \{k : x_k = 0,\ g_k > \lambda\}$. For a candidate active set J, we write λ_J for the mean of g over J. In *continuous-time projected gradient play*, each player i follows $\dot{x}_i(t) = \Pi_{T_\Delta(x_i(t))}\big[g_i(x_{-i}(t))\big]$, where $g_i(x_{-i}) = \big(u_i(s_i, x_{-i})\big)_{s_i \in S_i}$ is her payoff vector. Throughout, $\mathrm{prox}_f := \mathrm{prox}_f^\phi$ for $\phi = \frac{1}{2}\|\cdot\|_2^2$.

Lemma 5 (Proximal regret of projected gradient flow)**.** *Let $g : [t_0, t_1] \to \mathbb{R}^d$ be integrable and let $x : [t_0, t_1] \to \Delta_d$ be absolutely continuous with $\dot{x}(t) = \Pi_{T_\Delta(x(t))}[g(t)]$ for almost every t. For $\rho \in [0, 1)$ and $f \in \mathcal{F}_\rho(\Delta_d)$, let $M_f(x) = \min_{y \in \Delta_d}\{f(y) + \frac{1}{2}\|y - x\|_2^2\}$ be the Moreau envelope. Then*

$$\int_{t_0}^{t_1} \big\langle g(t),\, \mathrm{prox}_f(x(t)) - x(t)\big\rangle\, dt \;\le\; M_f(x(t_0)) - M_f(x(t_1)).$$

Proof. The function $h = f + \iota_{\Delta_d} + \frac{1}{2}\|\cdot\|_2^2$ is proper, lower semi-continuous and $(1-\rho)$-strongly convex. Hence the minimization defining $M_f(x)$ has a unique solution $\mathrm{prox}_f(x)$ for every $x \in \mathbb{R}^d$, and $M_f(x) = \frac{1}{2}\|x\|_2^2 - h^*(x)$, where the conjugate h^* is differentiable with $(1-\rho)^{-1}$-Lipschitz gradient $\nabla h^*(x) = \mathrm{prox}_f(x)$ (Rockafellar and Wets, 1998, Ch. 12). Thus $M_f \in C^1$ with $\nabla M_f(x) = x - \mathrm{prox}_f(x)$, and by the chain rule $\frac{d}{dt}M_f(x(t)) = \langle x(t) - \mathrm{prox}_f(x(t)), \dot{x}(t)\rangle$ for almost every t. By Moreau's decomposition, $g(t) - \dot{x}(t) = \Pi_{N_\Delta(x(t))}[g(t)] \in N_\Delta(x(t))$; since $\mathrm{prox}_f(x(t)) \in \Delta_d$, we have $\langle g(t) - \dot{x}(t), \mathrm{prox}_f(x(t)) - x(t)\rangle \le 0$. Therefore

$$\big\langle g, \mathrm{prox}_f(x) - x\big\rangle \le \big\langle \dot{x}, \mathrm{prox}_f(x) - x\big\rangle = -\tfrac{d}{dt}M_f(x(t)),$$

and integrating over $[t_0, t_1]$ completes the proof. □

Proof of part (ii). Step 1: Game definition. Players 2 and 3 have actions $\{\mathrm{H}, \mathrm{T}\}$ and play matching pennies, independently of Player 1: $u_2(a) = 1$ if $a_2 = a_3$ and $u_2(a) = -1$ otherwise, and $u_3 = -u_2$. Player 1 has actions $\{A, A^-, B\}$, and her payoff depends only on a_2:

$$u_1(A, \cdot) = 1, \qquad u_1(A^-, \cdot) = 0, \qquad u_1(B, a_2, a_3) = \begin{cases} 20 & a_2 = \mathrm{H}, \\ -20 & a_2 = \mathrm{T}. \end{cases}$$

Action A strictly dominates A^- by a margin of 1.

Step 2: The clock. Write $y = x_2(\mathrm{H})$ and $z = x_3(\mathrm{H})$, so that $g_2 = (2z - 1,\, 1 - 2z)$ and $g_3 = (1 - 2y,\, 2y - 1)$. With $y(0) = \frac{1}{2}$ and $z(0) = \frac{3}{4}$, the functions

$$y(t) = \tfrac{1}{2} + \tfrac{1}{4}\sin 2t, \qquad z(t) = \tfrac{1}{2} + \tfrac{1}{4}\cos 2t$$

solve projected gradient play: they remain in $[\frac{1}{4}, \frac{3}{4}]$, so no projection is active and the dynamics reduce to $\dot{y} = \frac{1}{2}(g_2(\mathrm{H}) - g_2(\mathrm{T})) = 2z - 1$ and $\dot{z} = \frac{1}{2}(g_3(\mathrm{H}) - g_3(\mathrm{T})) = 1 - 2y$, which both functions satisfy. The clock is π-periodic, and Player 1's payoff vector is

$$g_1(t) = \big(1,\ 0,\ 10\, s(t)\big), \qquad s(t) := \sin 2t.$$

Write $a = x_1(A)$, $m = x_1(A^-)$ and $b = x_1(B)$, and let $\alpha = \arcsin(1/10) \approx 0.1002$, $\tau_1 = (\pi - \alpha)/2$ and $\tau_2 = (\pi + \alpha)/2$. Fix any initial condition $x_1(0) \in X_1$ (for instance $x_1(0) = e_B$). The joint dynamics $\dot{x} = \Pi_{T_X(x)}\big[F(x)\big]$, with $F(x) = \big(g_1(x_{-1}), g_2(x_{-2}), g_3(x_{-3})\big)$, form a projected dynamical system on the polytope X with a Lipschitz vector field, which has a unique absolutely continuous solution from every initial condition Dupuis and Nagurney (1993); Nagurney and Zhang (1996). As Step 3 shows, the choice of $x_1(0)$ does not affect the periodic orbit.

Step 3: Reset and periodicity. Suppose $s(t) \ge \frac{1}{2}$ and $b(t) < 1$. Then $g_B \ge 5 > 1 \ge \max\{g_A, g_{A^-}\}$. In (18), B belongs to the active set I: otherwise $g_B \le \lambda \le 1$. Moreover $|I| \ge 2$, since some other coordinate is positive. Hence

$$\dot{b} = g_B - \lambda = \tfrac{1}{|I|}\textstyle\sum_{k \in I}(g_B - g_k) \ge \frac{|I|-1}{|I|}(g_B - 1) \ge 2.$$

Since $s \ge \frac{1}{2}$ on $[k\pi + \frac{\pi}{12},\, k\pi + \frac{5\pi}{12}]$, an interval of length $\frac{\pi}{3} > \frac{1}{2}$, Player 1 reaches the vertex e_B during this interval. At e_B, (18) gives $\dot{x}_1 = 0$ as long as $10s \ge 1$, which holds on $[k\pi + \frac{\alpha}{2},\, k\pi + \tau_1]$.

Hence $x_1(\tau_1 + k\pi) = e_B$ for every $k \geq 0$. Player 1's dynamics satisfy $\dot{x}_1 \in g_1(t) - N_\Delta(x_1)$, whose solutions are forward-unique: for two solutions, $\frac{d}{dt}\frac{1}{2}\|x_1 - x_1'\|^2 = -\langle x_1 - x_1', n - n'\rangle \leq 0$ by monotonicity of N_Δ. Since g_1 is π-periodic, $x_1(t+\pi) = x_1(t)$ for all $t \geq \tau_1$. Together with Step 2, the joint trajectory $x(t)$ is π-periodic on $[\tau_1, \infty)$.

Step 4: μ is an exact PCE. Define $\mu = \frac{1}{\pi}\int_{\tau_1}^{\tau_1+\pi} \delta_{x(t)}\, dt \in \Delta(X)$. Fix a player i, $\rho \in [0,1)$ and $f \in \mathcal{F}_\rho(X_i)$. By multilinearity, $u_i(\mathrm{prox}_f(x_i), x_{-i}) - u_i(x) = \langle g_i(x_{-i}), \mathrm{prox}_f(x_i) - x_i\rangle$. Each player's trajectory satisfies the hypotheses of Lemma 5 on $[t_0, t_1] = [\tau_1, \tau_1 + \pi]$ with $g = g_i(x_{-i}(t))$, so

$$\mathbb{E}_{x\sim\mu}\big[u_i(\mathrm{prox}_f(x_i), x_{-i}) - u_i(x)\big] \leq \tfrac{1}{\pi}\big(M_f(x_i(\tau_1)) - M_f(x_i(\tau_1+\pi))\big) = 0$$

by periodicity. This holds for every i, every $\rho \in [0,1)$ and every $f \in \mathcal{F}_\rho$, so μ satisfies (15).

Step 5: Mass on the dominated strategy. The induced outcome distribution satisfies $\bar{\mu}(a_1 = A^-) = \mathbb{E}_\mu[m] = \frac{1}{\pi}\int_{\tau_1}^{\tau_1+\pi} m(t)\, dt$. We follow one period.

(a) Edge A–B on $[\tau_1, \tau_2]$. Here $|s| \leq 0.1$, since $\sin 2t$ decreases from $\frac{1}{10}$ to $-\frac{1}{10}$ on this interval. With $I = \{A, B\}$, $\lambda = \lambda_{\{A,B\}} = \frac{1+10s}{2} \geq 0 = g_{A^-}$, so A^- stays inactive, and $\dot{a} = -\dot{b} = \frac{1}{2}(1 - 10s) \in [0,1]$; in particular $b \geq 1 - \alpha > 0$, so B stays positive. The interval is symmetric about $\pi/2$ ($\tau_1 + \tau_2 = \pi$) and $\sin 2t$ is odd about $t = \pi/2$, so $\int_{\tau_1}^{\tau_2} \sin 2t\, dt = 0$. Hence $m(\tau_2) = 0$, $a(\tau_2) = \frac{1}{2}(\tau_2 - \tau_1) = \alpha/2$, and $b(\tau_2) = 1 - \alpha/2$. Since $0.1 < \alpha < 0.1002$, this gives $0.9499 < b(\tau_2) < 0.95$.

(b) Interior phase. For $t \in (\tau_2, \tau_2 + 0.32]$ we have $2t \in (\pi + \alpha, \pi + \alpha + 0.64] \subset (\pi + \alpha, 2\pi - \alpha)$, so $s < -0.1$. Then $\lambda_{\{A,B\}} = \frac{1+10s}{2} < 0 = g_{A^-}$, so A^- enters, and while all coordinates are positive, $\lambda = \lambda_{\{A,A^-,B\}} = \frac{1+10s}{3}$ and

$$\dot{a} = \tfrac{2-10s}{3} > 0, \qquad \dot{m} = -\tfrac{1+10s}{3} > 0, \qquad \dot{b} = \tfrac{20s-1}{3} < 0. \qquad (*)$$

Integrating from τ_2, where $m(\tau_2) = 0$, we obtain, as long as $b > 0$,

$$b(t) = b(\tau_2) + \tfrac{10}{3}\big(\cos 2\tau_2 - \cos 2t\big) - \tfrac{t-\tau_2}{3}, \qquad m(t) = \tfrac{5}{3}\big(\cos 2t - \cos 2\tau_2\big) - \tfrac{t-\tau_2}{3}.$$

Using $\cos 2\tau_2 = -\cos\alpha$, we get $\cos 2\tau_2 - \cos 2t = \cos(\alpha + 2(t - \tau_2)) - \cos\alpha$. At $t = \tau_2 + 0.32$ this is less than -0.256, so the right-hand side of the formula for b is less than $0.95 - \frac{2.56}{3} - \frac{0.32}{3} = -0.01 < 0$ there; hence b must vanish before $\tau_2 + 0.32$. At $t = \tau_2 + 0.31$ it is greater than -0.244, so $b(\tau_2 + 0.31) > 0.9499 - \frac{2.44+0.31}{3} > 0.03 > 0$. Since b is decreasing, it hits zero at a unique $t^* \in (\tau_2 + 0.31,\ \tau_2 + 0.32)$. From $(*)$, $\dot{m} = -\frac{1}{2}\dot{b} - \frac{1}{2}$, so

$$m^* := m(t^*) = \tfrac{1}{2}b(\tau_2) - \tfrac{1}{2}(t^* - \tau_2) > \tfrac{1}{2}(0.9499) - 0.16 > 0.3149.$$

Moreover, m is nonnegative and increasing on $[\tau_2, t^*]$, and $t^* > \tau_2 + 0.31$. Using $\sin 2\tau_2 = -\sin\alpha$, this gives

$$\int_{\tau_2}^{t^*} m(t)\, dt > \int_{\tau_2}^{\tau_2+0.31} m(t)\, dt = \tfrac{5}{6}\big(\sin\alpha - \sin(\alpha + 0.62) + 0.62\cos\alpha\big) - \tfrac{0.31^2}{6} > 0.0317.$$

(c) Edge A–A^-. For $t \in [t^*, t^* + 1]$ we have $2t \in (\pi + \alpha, \pi + \alpha + 2.64] \subset (\pi, 2\pi)$, so $s < 0$. With $I = \{A, A^-\}$, $\lambda = \lambda_{\{A,A^-\}} = \frac{1}{2} > 10s = g_B$, so B stays inactive, $\dot{a} = \frac{1}{2}$ and $\dot{m} = -\frac{1}{2}$. Since $m^* \leq \frac{1}{2}b(\tau_2) < \frac{1}{2}$, m decreases linearly to 0 at $t^* + 2m^* < t^* + 1$, and

$$\int_{t^*}^{t^*+2m^*} m(t)\, dt = (m^*)^2 > 0.3149^2 > 0.0991.$$

Both windows $[\tau_2, t^*]$ and $[t^*, t^* + 2m^*]$ lie in $[\tau_1, \tau_1 + \pi]$, since $t^* + 2m^* < \tau_2 + 1.32 < \tau_1 + \pi$. Since $m \geq 0$ throughout the period,

$$\bar{\mu}(a_1 = A^-) = \mathbb{E}_\mu[m] \ \geq\ \frac{1}{\pi}\Big(\int_{\tau_2}^{t^*} m\, dt + \int_{t^*}^{t^*+2m^*} m\, dt\Big) > \frac{0.0317 + 0.0991}{\pi} > 0.04,$$

although A^- is strictly dominated by a margin of 1. $\square$

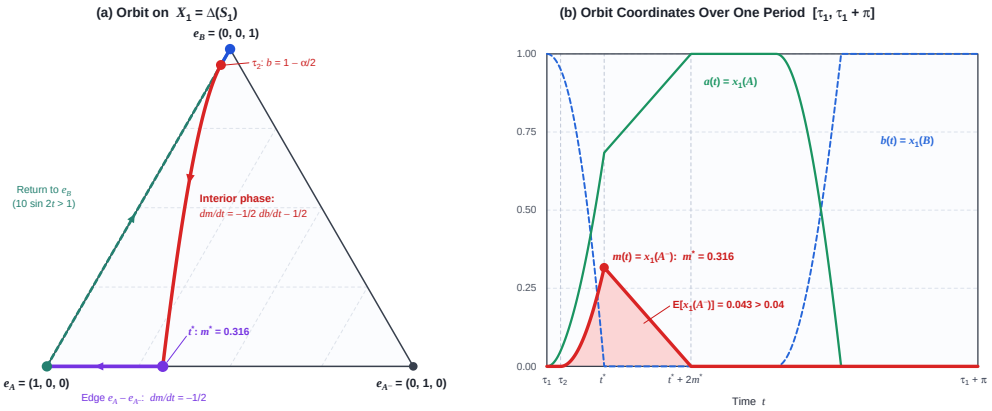

Figure 8: **Projection shielding: a periodic projected-gradient orbit supporting a strictly dominated strategy (Theorem 12(ii)).** Players 2 and 3 play matching pennies, which acts as a π-periodic clock and gives Player 1 the payoff vector $g_1(t) = (1, 0, 10 \sin 2t)$ over actions (A, A^-, B); A strictly dominates A^- by a margin of 1. **(a)** Orbit of Player 1's mixed strategy $(a, m, b) = \big(x_1(A), x_1(A^-), x_1(B)\big) \in X_1 = \Delta(S_1)$. At $\tau_1 = (\pi - \alpha)/2$, $\alpha = \arcsin \frac{1}{10}$, the orbit leaves e_B along the edge e_B–e_A with A^- inactive, reaching $b(\tau_2) = 1 - \alpha/2 \approx 0.950$ at $\tau_2 = (\pi + \alpha)/2$. From there the Euclidean projection lets A^- enter alongside A and B: the orbit crosses the interior until b vanishes at t^*, where the dominated mass peaks at $m^* \approx 0.316$, then slides along e_A–e_{A^-} to e_A, arriving at $t^* + 2m^*$ (where $2t = \pi + 2$), and rests there. Once $10 \sin 2t > 1$ it returns along e_A–e_B to e_B and rests there until $\tau_1 + \pi$. **(b)** Coordinates $a(t), m(t), b(t)$ over one period $[\tau_1, \tau_1 + \pi]$; the shaded area is $\int_{\tau_1}^{\tau_1+\pi} m\, dt$. By Lemma 5 and periodicity, the time average $\mu = \frac{1}{\pi} \int_{\tau_1}^{\tau_1+\pi} \delta_{(x_1,x_2,x_3)(t)}\, dt$ of the joint profile is an exact Euclidean PCE for every $\rho \in [0, 1)$, yet $\mathbb{E}_\mu[x_1(A^-)] \approx 0.043$ (Theorem 12 guarantees > 0.04).

Remark 5 (Pure-profile PCE excludes strictly dominated strategies)**.** *Let Γ be any finite game in which $A^- \in S_i$ is strictly dominated by $A \in S_i$ with margin $\epsilon > 0$, and let $\sigma \in \Delta(S)$ satisfy $\sigma(s_i = A^-) > 0$. Then μ_σ violates* (15) *already for $\rho = 0$. Indeed, let $K = \{y \in X_i : y(A^-) \le y(A)\}$ and $f = \iota_K \in \mathcal{F}_0(X_i)$, so that $\mathrm{prox}_f = \Pi_K$. This map fixes e_a for every $a \neq A^-$ and maps e_{A^-} to $\frac{1}{2}(e_A + e_{A^-})$, since $e_{A^-} - \frac{1}{2}(e_A + e_{A^-}) = \frac{1}{2}(e_{A^-} - e_A) \in N_K\big(\frac{1}{2}(e_A + e_{A^-})\big)$. Hence the left-hand side of* (15) *equals*

$$\tfrac{1}{2} \textstyle\sum_{s_{-i}} \sigma(A^-, s_{-i}) \big[u_i(A, s_{-i}) - u_i(A^-, s_{-i})\big] \;\ge\; \tfrac{\epsilon}{2}\, \sigma(s_i = A^-) > 0.$$

If test functions are required to be finite-valued, the smooth convex function $f_c(y) = \frac{c}{2} \max\{0, y(A^-) - y(A)\}^2$ gives $\mathrm{prox}_{f_c}(e_{A^-}) = (1 - \theta) e_{A^-} + \theta e_A$ with $\theta = \frac{c}{1+2c}$, and the same argument applies. In particular, the SCE of part (i) is not a PCE, so PCE $\subsetneq$ SCE, which is why part (ii) requires correlation over mixed strategies.

G Appendix of Section 6

Having established that Coarse Correlated Equilibria and their continuous proximal refinements (SCE/PCE) do not exclude strictly dominated strategies, a natural theoretical recourse is to demand swap-regret minimization. Constrained by the PPAD-hardness of computing exact Nash equilibria Daskalakis et al. (2006); Chen et al. (2009), swap regret and its sequential extensions have emerged as standard benchmarks for rationalizable learning.

Recent work has expanded the application of swap regret far beyond classical normal-form Correlated Equilibria (CE).[7] It provides the theoretical foundation for online calibrated forecasting Foster and

[7] Beyond its traditional algorithmic justification as the limit of swap-regret minimization, CE has recently been established as the unique solution concept satisfying a fundamental set of normative logical axioms, including consistency under payoff uncertainty and invariance to strategy cloning Brandl (2026).

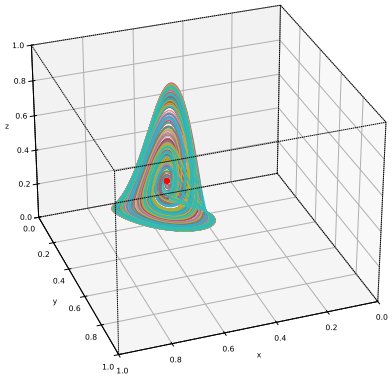

(a) Chaotic limit cycle of replicator dynamics in a normal-form polymatrix game. The red dot is the Nash equilibrium. Colors denote distinct trajectories.

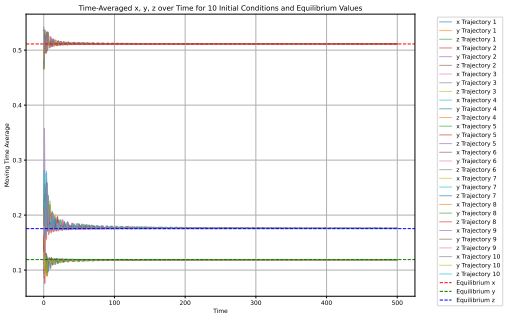

(b) For all interior symmetric initial conditions the time-average of the strategies of all agents converges precisely to the interior Nash equilibrium.

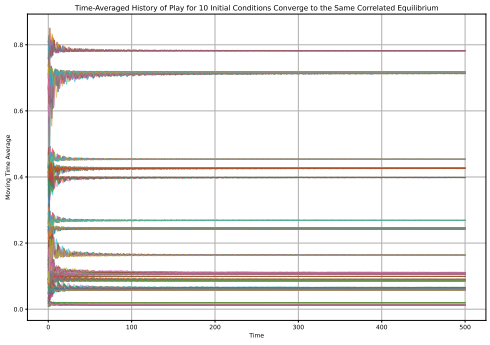

(c) The time-average of the empirical pairwise probabilities converges pointwise to their value at a correlated equilibrium (CE).

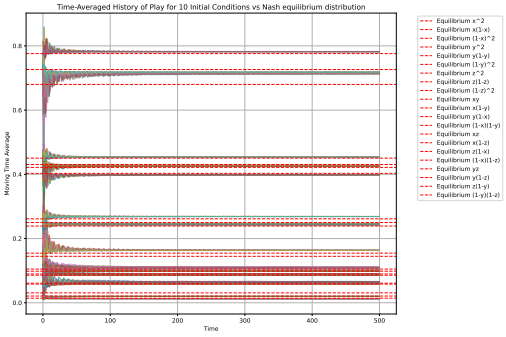

(d) The time-averages computed in (c) strictly do not converge to their independent Nash equilibrium values (encoded by the dashed red lines).

Figure 9: **The Combinatorial Shadow of Chaos (Expanded).** Replicator dynamics in normal-form games simultaneously exhibit macroscopic chaos, optimal $O(1/T)$ convergence to CE, and optimal state-average convergence to NE.

Vohra (1997); Fishelson et al. (2025), and has been extended to imperfect-information, Bayesian, and general convex games via Extensive-Form, Linear, and Profile swap regret Celli et al. (2020); Farina et al. (2022); Daskalakis et al. (2025b); Arunachaleswaran et al. (2025). Concurrently, novel algorithmic reductions have circumvented the classical polynomial dependence on the action space; however, these theoretical improvements mathematically necessitate an exponential dependence on the precision parameter, requiring $\exp(\Omega(1/\epsilon))$ rounds to converge Dagan et al. (2024); Peng and Rubinstein (2024).

Beyond assigning zero probability mass to strictly dominated strategies, swap-regret minimization imposes strict utility-theoretic constraints. It provably guarantees non-manipulability against strategic adversaries Deng et al. (2019); Mansour et al. (2022), ensures the Pareto-optimality of the learning algorithm's asymptotic menu Arunachaleswaran et al. (2025), and forces frequent-iterate convergence to Nash equilibria in symmetric zero-sum settings Leme et al. (2024). Given these theoretical guarantees, it is a standard premise that optimal $O(1/T)$ convergence to the CE polytope imposes sufficient mathematical structure to ensure predictable, stable, and rationalizable trajectories.

We demonstrate that this assumption does not hold dynamically. Evaluating learning algorithms strictly through the discrete combinatorial shadow of swap regret creates an analytical blind spot that obscures complex, non-equilibrating dynamics. By directly analyzing the multi-agent learning trajectories, we prove that even in minimal normal-form games, optimal convergence to CE can perfectly coincide with complex chaotic limit sets.

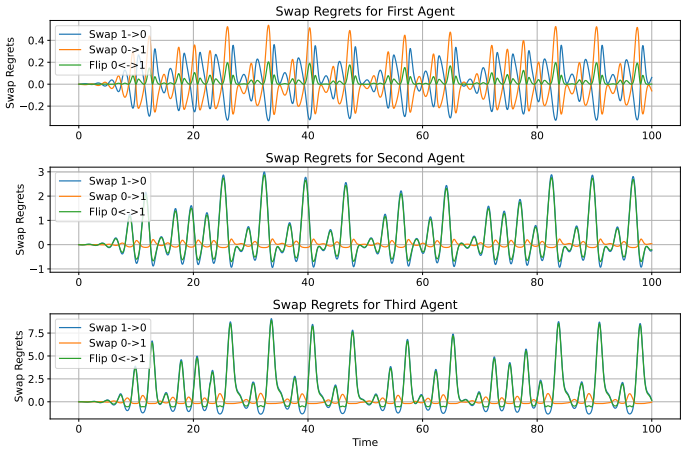

Figure 10: **Bounded Swap Regret amidst Turbulence.** The instantaneous swap regret of all agents remains strictly bounded, ensuring the time-average of the replicator dynamics converges provably at an optimal $O(1/T)$ rate to a chaotic correlated equilibrium.

The implications of this result are clarified by contrasting it with canonical 3×3 zero-sum games, such as Rock-Paper-Scissors. In standard 3-strategy games, (even in symmetric RPS games under symmetric initializations) replicator dynamics minimize external regret but provably fail to minimize swap regret, thereby converging only to the CCE polytope Leme et al. (2024). However, because our construction utilizes exactly two strategies per agent, testing against fixed deviations is sufficient to check against all swap deviations. In this specific environment, replicator dynamics inherently functions as an optimal, no-swap-regret algorithm. When combined with provable time-average convergence of agents' marginal distributions to NE, it satisfies arguably the most stringent game-theoretic desiderata (outside pointwise convergence to NE) while simultaneously sustaining chaotic trajectories.

The emergence of this chaotic limit set relies on a bipartite mirroring technique that translates non-concave evolutionary game dynamics into multilinear normal-form optimization. By coupling a 3-agent chaotic polymatrix replicator Peixe and Rodrigues (2022) with an exact clone population, the resulting 6-agent system preserves the chaotic vector field along a symmetric invariant manifold.

Crucially, standard topological constraints dictate that generic normal-form games cannot possess attracting limit sets strictly within the interior of the state space under continuous-time no-regret learning Omidshafiei et al. (2019); Flokas et al. (2020). Consequently, this chaotic regime mathematically manifests not as an open attractor, but as an invariant limit set strictly bound to the symmetric initialization manifold.

From a measure-theoretic perspective, under a prior absolutely continuous with respect to the Lebesgue measure (e.g., a uniform prior over the entire state space), this exact chaotic limit set constitutes a measure-zero event. However, this invariant set acts as a topological mechanism for disequilibrium. While the strange attractor exhibits strictly positive internal Lyapunov exponents (verified via perturbation analysis in Appendix E.1) generating chaos within the manifold, the invariant manifold itself is also transversely unstable. As demonstrated numerically in Appendix E.2, introducing microscopic asymmetric noise to break the initial synchronization does not cause the system to relax to a static Nash equilibrium. Instead, the transverse instability exponentially amplifies the asymmetric perturbation, permanently breaking the mirror symmetry. Depending on the exact microscopic noise vector, arbitrarily close initial conditions are repelled toward entirely distinct, lower-dimensional limit cycles residing on different boundary faces of the state space.

This extreme sensitivity to initial conditions ensures that even when the measure-zero chaos is transient, the ultimate macroscopic support of the agents' strategies remains fundamentally unpredictable. Whether the system remains trapped in the chaotic manifold under symmetric priors, or is ejected into distinct boundary cycles under generic perturbations, the game-theoretic conclusion is identical. The projection of these continuous orbits onto the discrete CE simplex systematically obscures

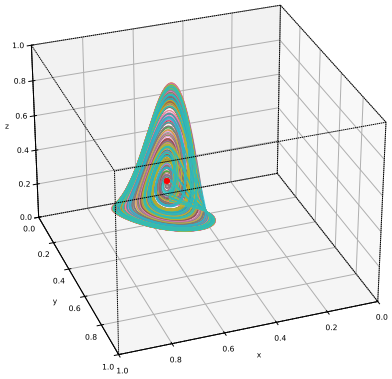

(a) Chaotic limit cycle of replicator dynamics in a normal-form polymatrix game. The red dot is the Nash equilibrium. Colors denote distinct trajectories.

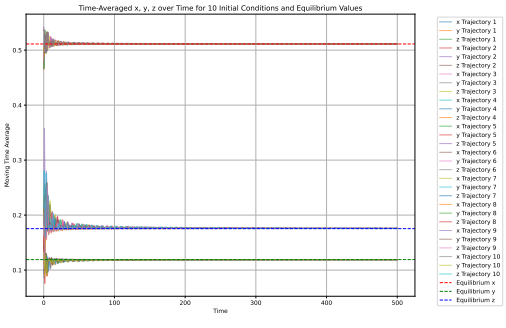

(b) For all interior symmetric initial conditions the time-average of the strategies of all agents converges precisely to the interior Nash equilibrium.

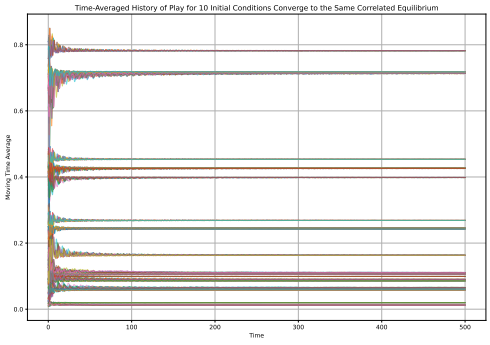

(c) The time-average of the empirical pairwise probabilities converges pointwise to their value at a correlated equilibrium (CE).

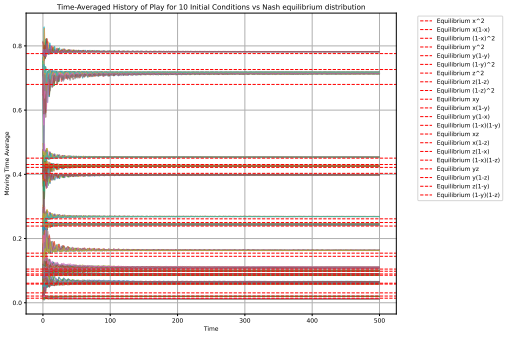

(d) The time-averages computed in (c) strictly do not converge to their independent Nash equilibrium values (encoded by the dashed red lines).

Figure 11: **The Combinatorial Shadow of Chaos.** Replicator dynamics in normal-form games simultaneously exhibit macroscopic chaos and time-average convergence to NE.

the underlying dynamics. A proof of fast convergence to Correlated Equilibria guarantees neither behavioral predictability nor dynamic equilibration.

G.1 Proof of Theorem 6

To visualize the algebraic transformation required for the proof of Theorem 6, we provide a structural diagram of the bipartite mirroring technique in Figure 13. This transformation is strictly necessary to convert non-concave evolutionary dynamics into multilinear normal-form optimization while preserving the target invariant manifold.

thm:chaos There exists a one-parameter family of normal-form games with payoff symmetries such that continuous-time replicator dynamics, given symmetric, interior initial conditions, exhibits chaos while simultaneously:

(i) Minimizing swap regret at an optimal rate of $O(1/T)$ (fast convergence to CE).

(ii) Ensuring the time-average strategies of all agents converge precisely to an interior Nash equilibrium.

The games in question are polymatrix games with 6 agents, each possessing exactly two strategies.

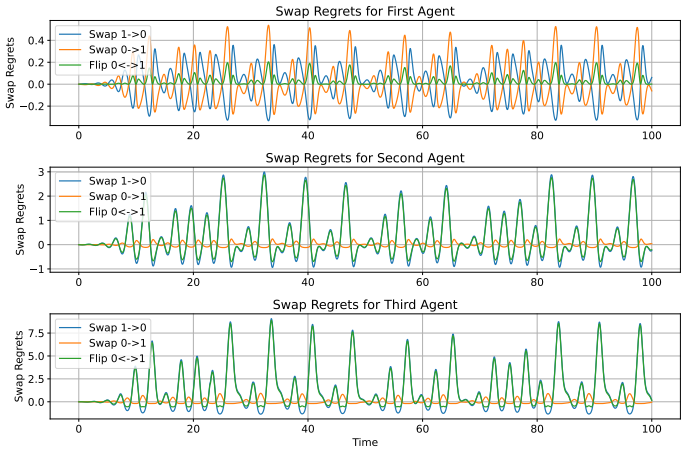

Figure 12: **Bounded Swap Regret amidst Turbulence.** The instantaneous swap regret of all agents remains strictly bounded, ensuring the time-average of the replicator dynamics converges provably at an optimal $O(1/T)$ rate to a chaotic correlated equilibrium.

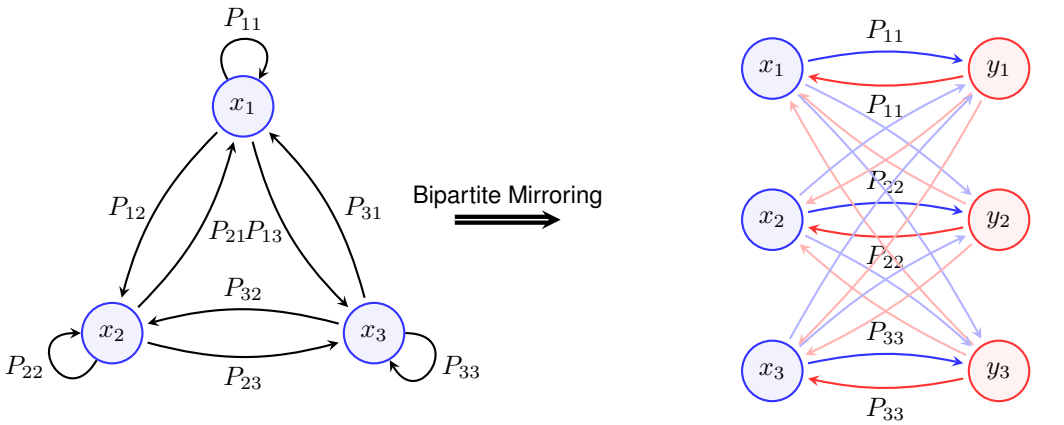

Figure 13: **Restoring Multilinearity via Bipartite Mirroring.** The base 3-agent polymatrix game (left) contains self-loops (P_{ii}), forcing an agent's utility to depend quadratically on their own mixed strategy, rendering the game non-concave and violating standard regret frameworks. By cloning the population into sets X (blue) and Y (red), self-loops are unfolded into bidirectional bipartite interactions between identical agents. The resulting 6-agent game is strictly multilinear, enabling optimal regret minimization while preserving the original chaotic vector field along the symmetric invariant manifold $x_i = y_i$.

Proof. We prove the theorem by formally constructing the requisite multilinear state space, verifying the bounded regret properties without relying on asymptotic approximations, and characterizing the exact topological nature of the strange attractor.

Step 1: Self-Contained $O(1/T)$ Swap Regret
We first establish that continuous-time replicator dynamics inherently minimize swap regret at an optimal $O(1/T)$ rate for any two-strategy game. Let $p(t) \in \Delta_2$ denote the mixed strategy of an agent over two actions, and let $u(t) \in \mathbb{R}^2$ be the instantaneous utility vector. The continuous-time replicator equation is $\dot{p}_j(t) = p_j(t)(u_j(t) - p(t)^\top u(t))$.

To bound the external regret against any fixed action in hindsight, we evaluate the Kullback-Leibler (KL) divergence between a fixed target strategy q and the state $p(t)$. Taking the time derivative yields

$\frac{d}{dt} D_{KL}(q \parallel p(t)) = -\sum_j q_j \frac{\dot{p}_j(t)}{p_j(t)}$. Substituting the replicator equation, this simplifies exactly to $p(t)^\top u(t) - q^\top u(t)$, which is the instantaneous negative regret. Integrating over time T yields the total external regret bound:

$$\text{ExtReg}(T, q) = \int_0^T (q^\top u(t) - p(t)^\top u(t))dt = D_{KL}(q \parallel p(0)) - D_{KL}(q \parallel p(T)) \leq D_{KL}(q \parallel p(0))$$

Because the initial condition $p(0)$ is strictly interior, $D_{KL}(q \parallel p(0))$ is a strictly bounded $O(1)$ constant. Thus, the total external regret for both pure strategies 1 and 2 is strictly bounded by $O(1)$.

Because the game possesses exactly two strategies, bounding external regret completely bounds swap regret. There are exactly four swap functions mapping $\{1, 2\}$ to itself: the identity map (zero regret), the map to constant 1 (external regret for 1), the map to constant 2 (external regret for 2), and the flip map ($1 \leftrightarrow 2$). By the linearity of expectation, the regret of the flip map is strictly upper-bounded by the sum of the external regrets for 1 and 2. Consequently, the total swap regret for any modification rule is bounded by $O(1)$, yielding an optimal time-average swap regret rate of $O(1/T)$ and guaranteeing fast convergence to the Correlated Equilibrium polytope.

Step 2: The Non-Concave Strange Attractor
We leverage a continuous-time 3-player, 2-strategy evolutionary polymatrix game introduced by Peixe and Rodrigues Peixe and Rodrigues (2022). This work rigorously establishes a one-parameter family of vector fields that undergo a supercritical Hopf bifurcation, leading to a Shilnikov homoclinic cycle to a saddle-focus. By the Shilnikov-Ovsyannikov theorem, the unfolding of this cycle generates suspended topological horseshoes, strictly positive Lyapunov exponents, and persistent Hénon-type strange attractors that survive over parameter sets of positive Lebesgue measure.

However, this base system is an evolutionary game that includes "self-loops," meaning an agent's sub-population interacts with itself. Algebraically, this forces the utility function of agent i to depend on the quadratic form $x_i^\top P_{ii} x_i$. Because the utility functions are non-concave, the standard potential functions used in Step 1 fail, and standard regret guarantees do not apply.

Step 3: Bipartite Mirroring and the Invariant Manifold
To resolve the non-concavity while strictly preserving the Shilnikov cycle, we apply a bipartite mirroring transformation (visualized in Figure 13). We duplicate the 3 agents to create two distinct mirror populations: $X = (x_1, x_2, x_3)$ and $Y = (y_1, y_2, y_3)$. We configure the interactions such that X agents only play against Y agents, replacing the self-loop matrices P_{ii} with bipartite matrices $P_{i,i'}$ mapping x_i against y_i.

The resulting 6-agent normal-form game is strictly multilinear, rendering the $O(1)$ swap-regret bounds from Step 1 fully applicable to all 6 agents. Crucially, the replicator vector field preserves symmetric initial conditions. If initialized such that $x_i(0) = y_i(0)$ for all $i \in \{1, 2, 3\}$, the dynamics guarantee $\dot{x}_i(t) = \dot{y}_i(t)$ for all $t > 0$. The system is completely confined to a 3-dimensional symmetric invariant manifold $\mathcal{M} = \{(X, Y) \mid X = Y\}$ within the 6-dimensional state space. Restricted to $\mathcal{M}$, the multilinear vector field perfectly collapses back into the exact non-concave vector field of the Peixe-Rodrigues system, fully inheriting its Hénon-type strange attractor and positive Lyapunov exponents.

Step 4: Time-Average Equilibration
Finally, we prove that despite the macroscopic chaos on $\mathcal{M}$, the time-average of the trajectory converges exactly to an interior Nash equilibrium.

The replicator equation for agent i's probability of playing strategy 1 is $\dot{x}_i = x_i(1 - x_i)(u_{i,1} - u_{i,2})$. Dividing by $x_i(1 - x_i)$ and integrating over $[0, T]$ yields:

$$\frac{1}{T}\left(\ln\left(\frac{x_i(T)}{1 - x_i(T)}\right) - \ln\left(\frac{x_i(0)}{1 - x_i(0)}\right)\right) = \frac{1}{T}\int_0^T (u_{i,1}(t) - u_{i,2}(t))dt$$

Because the strange attractor is strictly contained within the interior of the simplex, $x_i(t)$ is strictly bounded away from the boundaries 0 and 1 for all t. Therefore, the logarithmic terms on the left-hand side are bounded by an absolute constant $O(1)$. As $T \to \infty$, the left-hand side rigorously evaluates to 0.

This implies the time-average difference in expected utilities evaluates to zero. Because the unfolded game is a multilinear polymatrix game, the expected utility differences are strict linear functions of the state variables (X, Y). By the linearity of the integral, evaluating the utility differences at any accumulation point $(\hat{X}, \hat{Y})$ of the time-average trajectory yields exactly zero. Thus, all agents are rendered perfectly indifferent between their strategies. Because the attractor is strictly interior, any such accumulation point must be the unique interior Nash equilibrium of the game. □

G.2 Numerical Verification of Structural Stability and Persistence of Chaos

To further validate the structural stability of the chaotic dynamics observed in the one-parameter family of polymatrix replicators, we conducted a rigorous perturbation analysis on the generalized 12-coefficient system derived from Equation (6) in Peixe and Rodrigues (2022). Setting the base parameter to $\mu = 3.6$—which lies strictly within the chaotic regime—we calculated the maximal Lyapunov exponent (LE). The unperturbed base game yielded a positive LE of approximately 0.0489, mathematically confirming the presence of a strange attractor.

To test the persistence of this attractor under structural deformations (approximating a C^2 neighborhood of vector fields), we applied single ($k = 1$) and pairwise ($k = 2$) perturbations of $\pm\rho$ (with $\rho = 0.01$) to the polynomial coefficients of the vector field. Following the methodology of the main text, we adopted a conservative threshold of 5×10^{-3} to filter out numerical noise, classifying any LE above this value as strictly positive and indicative of chaotic dynamics.

The perturbation results provide robust computational support for the persistence of the chaotic regime identified in Theorem 6. Out of 24 single-parameter perturbations, 23 (95.8%) retained a positive LE. For pairwise perturbations, 251 out of 264 (95.1%) remained chaotic. The distribution of the maximum Lyapunov exponents across all 289 simulated games (Figure 14) illustrates that the vast majority of perturbed systems cluster strictly above the chaotic threshold. These findings provide strong computational evidence that the Hénon-type strange attractors in this system are persistently observable and highly robust to small, generalized structural perturbations in the space of symmetric games.

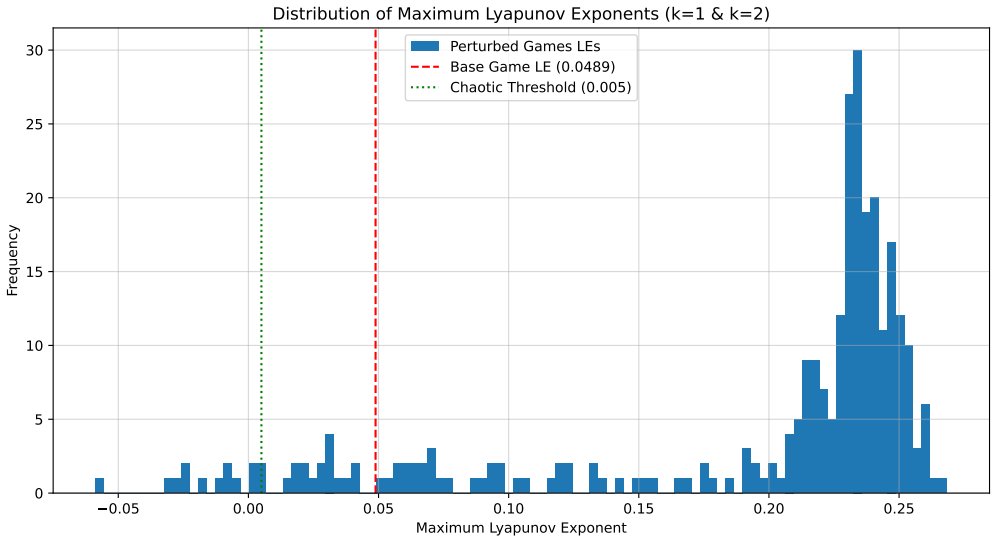

Figure 14: **Structural Robustness of Chaos.** The distribution of the maximum Lyapunov exponents across all 289 simulated games illustrates that the vast majority of perturbed systems $(> 95\%)$ cluster well above the positive chaotic threshold.

G.3 Numerical Demonstration of Symmetry Breaking and Distinct Limit Cycles

In Theorem 6, we established that continuous-time replicator dynamics can exhibit chaotic limit sets in a 6-player, 2-strategy normal-form game, provided the system is initialized with perfect symmetry between the two mirror populations (X and Y). To demonstrate the topological brittleness of this

chaotic limit behavior under generic conditions, we numerically simulate the full 6D replicator vector field under asymmetric initializations.

Experimental Setup: We initialize the system near the unstable interior equilibrium, O_μ. To break the invariant 3D manifold, we introduce a microscopic asymmetric perturbation. Specifically, the X agents are initialized at $X_{init} = O_\mu + \xi$, and the mirror Y agents are initialized at $Y_{init} = X_{init} + \delta$, where $\delta \sim 10^{-5}$ represents a random, asymmetric noise vector between the two populations. We simulate 10 such trajectories for $T = 300$ time units.

Observations: As shown in Figure 15, the trajectories initially shadow the chaotic limit set, remaining tightly coupled. However, because the chaotic regime is not an open attractor in the full 6D space, the positive Lyapunov exponents rapidly amplify the microscopic noise δ. The Euclidean distance between the X and Y populations grows exponentially, permanently breaking the mirror symmetry.

Crucially, the destruction of the chaotic limit set does not lead the system to a global, static point-mass equilibrium. As the dynamics are repelled toward the boundary of the simplex, four of the state variables asymptotically converge to pure strategies (0 or 1), while the remaining two variables are captured by a periodic orbit. Depending on the exact vector of the initial microscopic perturbation δ, the trajectories are expelled toward entirely distinct, lower-dimensional limit cycles residing on different 2D boundary faces of the state space. This reinforces the central thesis of Section 6: even when fragile, problematic behaviors like measure-zero chaos are destroyed by generic noise, the resulting statistically typical behavior systematically evades the predictive power of static equilibria.

H Appendix of Section 7

In the preceding sections, we observed how the combinatorial complexity of finite normal-form games inherently complicates the convergence of independent learning dynamics. One might naturally look to the continuous domain of non-atomic congestion games as a refuge of algorithmic stability. Indeed, this setting typically represents arguably the ideal of game-theoretic and algorithmic optimality. In non-atomic routing, Wardrop equilibria correspond to the global minima of a convex potential function, yielding effectively unique macroscopic flows Beckmann et al. (1956); Nisan et al. (2007). The celebrated Price of Anarchy (PoA) results, based on tight two-link Pigou networks, provide worst-case inefficiency bounds of $\Theta(p/\ln p)$ for polynomial latency functions of degree p Roughgarden and Tardos (2002, 2004). These guarantees are significantly tighter than their atomic counterparts, a favorable property that atomic games are known to formally recover in the limit as the population size increases Feldman et al. (2016b). Furthermore, unlike in atomic games where no-regret algorithms can be consistent with strictly dominated strategies, classic results establish that no-regret learning in non-atomic routing guarantees the day-to-day empirical flow converges to an ϵ-equilibrium flow Blum et al. (2006).

Together, these results form a powerful macroscopic framework. However, a rigorous examination of the physical timescale of these guarantees reveals a fundamental tension in the day-to-day dynamical reality of the system. To achieve a physically meaningful relative error (e.g., bounding regret by a small fraction δ of the optimal equilibrium cost), no-regret algorithms must aggressively throttle their learning rate to account for the maximum possible cost in the network. For a network of m paths with degree p latencies, achieving this relative bound requires an effective step-size scaling of $\mathcal{O}(1/m^p)$, which e.g., given the canonical decreasing step-sizes of $\theta(1/\sqrt{t})$, in turn forces the convergence horizon to scale exponentially as $\Omega(m^{2p})$. For standard responsive learning rates, the system operates in a vast meta-stable regime long before asymptotic bounds become binding.

Recent important contributions have illuminated this meta-stable regime, demonstrating the emergence of Li-Yorke chaos and qualitative inefficiencies in non-atomic congestion games under Multiplicative Weights Update (MWU) and Follow-the-Regularized-Leader (FTRL) with steep regularizers Chotibut et al. (2020); Bielawski et al. (2021, 2025). Building upon this work, we provide a generalized and exact quantification of this inefficiency for polynomial latencies, contrasting the dynamical reality directly against the static PoA guarantees.

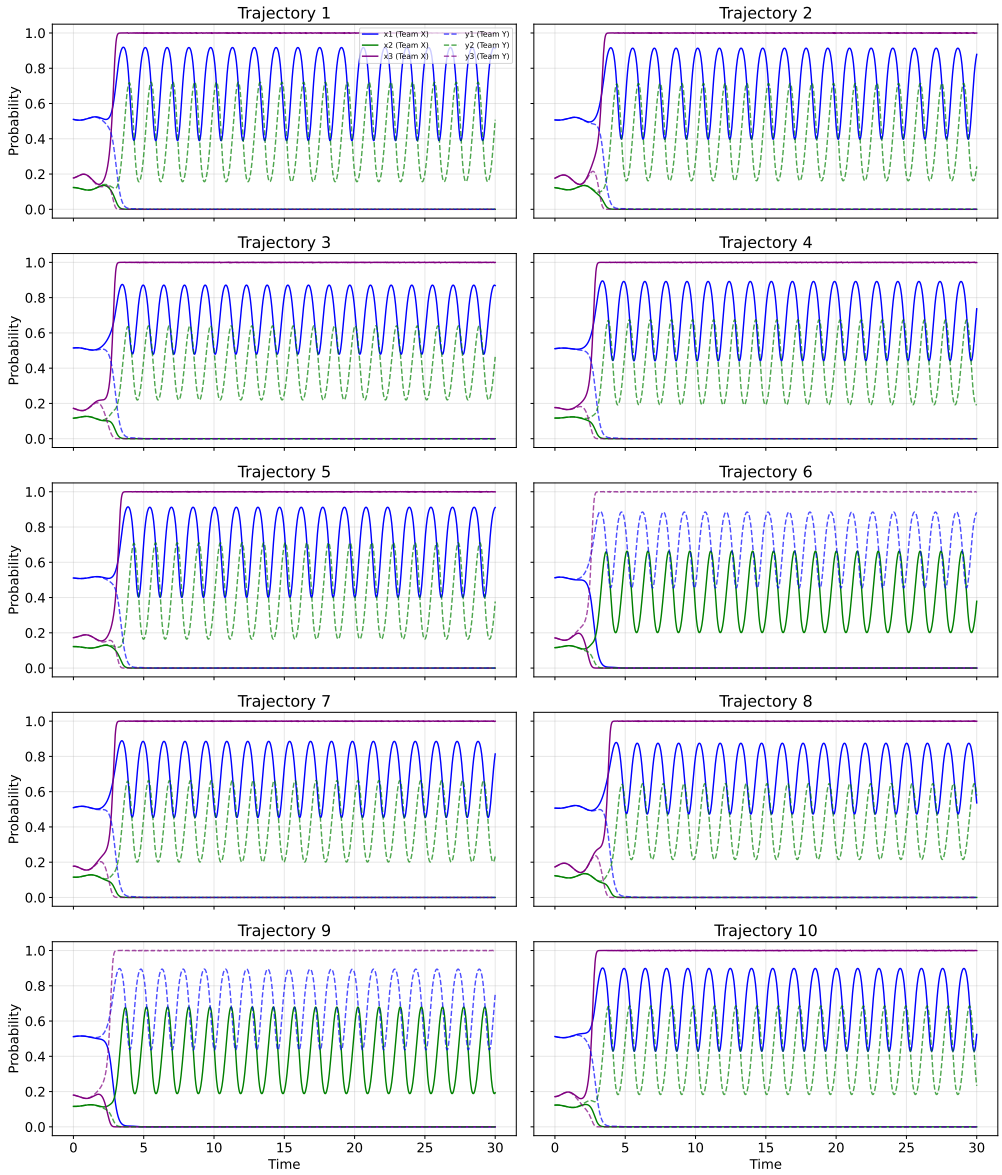

Figure 15: **Symmetry Breaking and Emergence of Distinct Limit Cycles.** Numerical simulation of the 6-agent mirrored polymatrix game under asymmetric initializations. Ten trajectories are initialized near the interior equilibrium with a microscopic asymmetric perturbation (10^{-5}) between the X and Y populations. While the trajectories initially shadow the chaotic limit set, the positive Lyapunov exponents rapidly amplify the noise, leading to an exponential divergence that permanently breaks the mirror symmetry. The system subsequently escapes to the boundary of the simplex, where the agents are captured by entirely distinct, lower-dimensional limit cycles residing on different boundary faces of the state space.

Theorem 7 (Exponential Degradation of the Global Attractor)**.** *Consider a symmetric non-atomic congestion game with two parallel paths and identical monomial latency functions $c(x) = x^p$ for $p \geq 1$. Under MWU with a fixed effective step-size parameter $a > 0$*[8]*:*

[8]The parameter a formally couples the MWU learning rate ϵ, the population size N, and the polynomial degree p. It is defined as $a = N^p \ln\left(\frac{1}{1-\epsilon}\right) \approx \epsilon N^p$. This relationship explicitly demonstrates that to maintain a stable effective step-size a, the raw learning rate ϵ must be aggressively throttled inversely to the maximum network cost $\mathcal{O}(N^p)$.

1. *The unique Wardrop equilibrium $x^* = 1/2$ is locally asymptotically stable if $a < \frac{2^{p+1}}{p}$.*

2. *If $a > \frac{2^{p+1}}{p}$, the equilibrium strictly destabilizes and the interior system is captured by a global, period-2 attracting orbit $\{\sigma_a, 1-\sigma_a\}$. Accounting for the population size N, the empirical time-averaged social cost of this chaotic orbit is exactly $SC_{emp}(a) = N^{p+1}[\sigma_a^{p+1} + (1-\sigma_a)^{p+1}]$, which strictly exceeds the optimal social cost of $N^{p+1}(1/2)^p$.*

3. *As the effective load $a \to \infty$ ($\sigma_a \to 0$), the ratio of this empirical time-averaged social cost to the optimal social cost converges exactly to 2^p.*

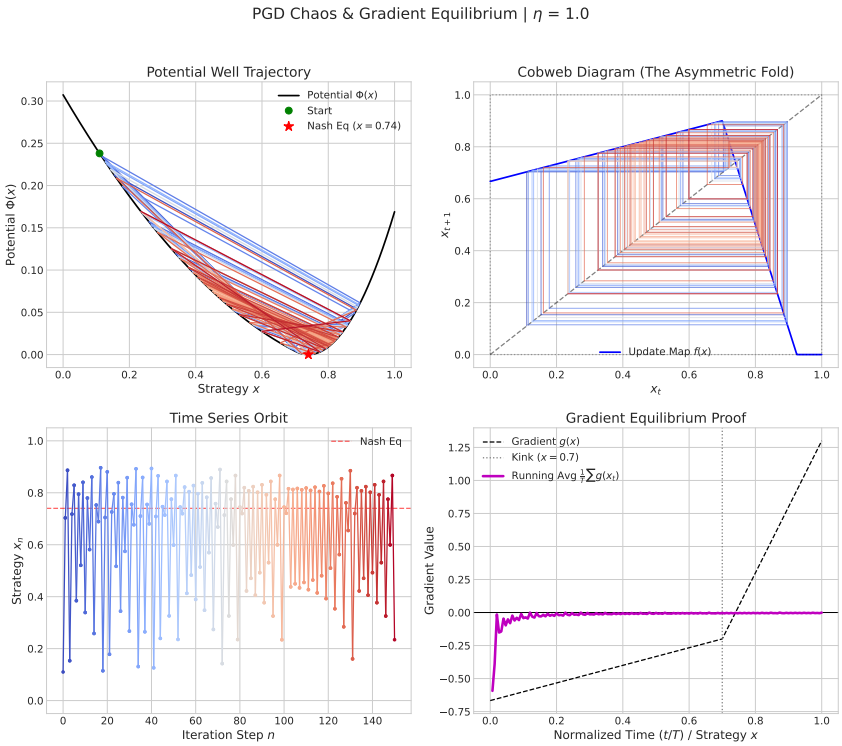

Figure 16: The emergence of Li-Yorke chaos and the concurrent satisfaction of Gradient Equilibrium (GEQ) under discrete-time Projected Gradient Descent (PGD) with L_2 regularization ($\eta = 1.0$) in an asymmetric two-link congestion game. **(Top-Left)** The state trajectory plotted on the exact convex potential well $\Phi(x)$. Rather than descending to the global minimum (Nash Equilibrium at $x = 0.74$), the discrete step-size forces the system to perpetually overshoot, oscillating across the well. **(Top-Right)** The cobweb diagram of the unprojected update map $f(x)$. The asymmetric bottleneck induces a steep "fold" in the map. Because the resulting chaotic attractor is strictly bounded away from the boundaries (0 and 1), the Euclidean projection operator is never activated, demonstrating that the chaos is native to the continuous potential landscape itself. **(Bottom-Left)** The time series of the strategy x_n, showcasing persistent, non-equilibrating chaotic fluctuations around the equilibrium. **(Bottom-Right)** The empirical validation of the GEQ condition Angelopoulos et al. (2025). Despite the severe chaotic variance of the primal state x_n, the time-averaged gradient strictly converges to zero over time. This illustrates the fundamental paradox: time-averaged convergence of the dual/gradient variables perfectly masks the volatile, inefficient reality of the primal day-to-day flows. See Appendix I.1 for the formal construction.

Remark 6 (The Fragility of Global Finetuning)**.** *From a strictly centralized optimization perspective, one might object that in this highly stylized, symmetric two-path game, an omniscient designer could algebraically reverse-engineer a "sweet spot" learning rate ϵ that yields rapid convergence. Fastest local convergence occurs when the derivative of the update map at the equilibrium is zero ($f_a'(1/2) = 0$), which requires an optimal effective load of $a_{opt} = \frac{2^p}{p}$.*

Given the parameter coupling $a \approx \epsilon N^p$, the corresponding optimal raw learning rate is exactly:

$$\epsilon_{opt} \approx \frac{2^p}{pN^p} = \mathcal{O}(N^{-p}) \tag{19}$$

While this sweet spot algebraically exists, realizing it is challenging in practice and theoretically misaligned with decentralized learning.

The fundamental issue is not merely the magnitude of the step-size (this itself is not scale invariant and is dependent on the choice of units), but the fragility of the required finetuning. First, matching this optimal rate requires an independent agent to perfectly calibrate their learning against the global maximum network cost $\mathcal{O}(N^p)$. This demands oracle-like knowledge of global system parameters—the exact total population N and the precise latency degree p—whereas real-world demand is inherently unpredictable and cost functions are only approximate.

Crucially, rational, selfish agents do not coordinate such fragile global finetuning; they reactively adapt to localized historical cost signals. As established recently by Bielawski et al. (2025), even when populations consist of heterogeneous subpopulations of agents adapting from general initial conditions with generic, uncoordinated learning rates, the macroscopic system natively bifurcates into robust Li-Yorke chaos. Consequently, while an exact $\mathcal{O}(N^{-p})$ step-size guarantees rapid convergence in principle, its realization relies on a degree of systemic coordination and precise global knowledge that is fundamentally at odds with decentralized learning.

Furthermore, this dynamical instability is not merely an artifact of the boundary-avoiding nature of steep regularizers (like Shannon entropy in MWU). It is a robust topological consequence of discrete-time updates on continuous potential landscapes. As illustrated in Figure 16, applying standard L_2 regularization (i.e., Projected Gradient Descent or Lazy Mirror Descent) in an asymmetric network natively induces Li-Yorke chaos within the interior of the simplex.

Remark 7 (Gradient Equilibria, Jensen's Gap, and Physical Attractors)**.** *An elegant phenomenon observed in these unstable regimes is that the time-averaged cost difference between the two paths strictly vanishes over time Chotibut et al. (2020). From the perspective of recent literature, this formally implies that the system satisfies the conditions of a Gradient Equilibrium (GEQ) Angelopoulos et al. (2025). This establishes a profound parallel: just as Correlated Equilibria accommodate complex chaotic attractors in continuous-time dynamics, Gradient Equilibria formally accommodate discrete-time Li-Yorke chaos, even in the most benign convex optimization settings.*

The GEQ framework explicitly acknowledges that such states do not enforce standard optimality and can incur linear regret. Theorem 7 provides the exact, worst-case quantification of this inefficiency via Jensen's gap: the system locks into a maximal-variance period-2 orbit where the gradients perfectly average to zero, yet the strict convexity of the macroscopic cost function guarantees a severe 2^p inefficiency.

However, the most critical limitation of GEQ lies not in its inefficiency, but in its predictive looseness. Because GEQ is fundamentally a kinematic condition on the time-averaged gradients, it is entirely agnostic to dynamical stability. By the definition of Li-Yorke chaos, the GEQ condition is simultaneously satisfied by infinitely many unstable periodic orbits, as well as an uncountable scrambled set of aperiodic trajectories. Yet, in well-behaved physical reality, almost all initial conditions are captured by just one or a few structurally attracting limit sets (such as the period-2 attractor). The GEQ solution concept lacks the mathematical machinery to distinguish between these uncountably many unobservable "ghost" orbits and the true physical attractor. Consequently, to generate more accurate predictions within the GEQ framework, one must deploy the tools of ergodic theory and dynamical systems to isolate the specific invariant measures that physically manifest.

I Proof of Theorem 7

We analyze the discrete-time map for a symmetric congestion game with latency $c_1(x) = c_2(x) = x^p$ under Multiplicative Weights Update. Let $x_t \in (0, 1)$ be the flow on the first path. The update map is defined as:

$$x_{t+1} = f_a(x_t) = \frac{x_t}{x_t + (1 - x_t) \exp\left(a\left[x_t^p - (1 - x_t)^p\right]\right)} \tag{20}$$

where $a > 0$ acts as the effective system load parameter. By symmetry, $x^* = 1/2$ is a fixed point. The proof will proceed in several steps.

Proof.

Part 1: Local Asymptotic Stability. Let $h_a(x) = a(x^p - (1-x)^p)$. Applying the quotient rule to $f_a(x)$ and evaluating at the equilibrium $x = 1/2$, we note $h_a(1/2) = 0$ and $h'_a(1/2) = ap2^{2-p}$. The derivative of the map at the fixed point simplifies exactly to:

$$f'_a(1/2) = 1 - \frac{1}{4}h'_a(1/2) = 1 - ap2^{-p} \tag{21}$$

For local asymptotic stability, we require $|f'_a(1/2)| < 1$. Because $a, p > 0$, the upper bound holds trivially. The lower bound requires $1 - ap2^{-p} > -1$, which rigorously isolates the critical threshold $a < \frac{2^{p+1}}{p}$. If a exceeds this threshold, $f'_a(1/2) < -1$ and the equilibrium strictly destabilizes.

Part 2: Global Convergence to the Period-2 Attractor. Before establishing the global attractor, we state a foundational result from one-dimensional topological dynamics regarding continuous maps lacking period-2 points.

Lemma 6 (Convergence in Maps without Period-2 Points Block and Coppel (2006); Sharkovsky et al. (2013))**.** *Let $I \subset \mathbb{R}$ be a closed interval and $g : I \to I$ a continuous map. If g possesses no periodic points of period 2, then the sequence of iterates $g^n(x)$ converges to a fixed point of g for every $x \in I$.*

We define the involution $\varphi(x) = 1 - x$. Because the game is symmetric, f_a commutes with φ. We introduce the auxiliary map $g_a = \varphi \circ f_a = 1 - f_a(x)$. Because φ is an involution, evaluating g_a twice yields $g_a^2(x) = 1 - f_a(1 - f_a(x)) = f_a^2(x)$. Thus, the period-2 points of f_a are mathematically identical to the fixed points of g_a^2.

The fixed points of g_a satisfy $g_a(x) = x \iff f_a(x) = 1 - x$. Substituting the exact definition of $f_a(x)$ and taking the natural logarithm, this condition is equivalent to finding the roots of:

$$D(x) = \ln\left(\frac{x}{1-x}\right) - \frac{a}{2}\left(x^p - (1-x)^p\right) = 0 \tag{22}$$

To rigorously determine the number of roots of $D(x)$ on the active interior $(0, 1/2)$, we must evaluate its critical points by taking the first derivative:

$$D'(x) = \frac{1}{x(1-x)} - \frac{ap}{2}\left(x^{p-1} + (1-x)^{p-1}\right) \tag{23}$$

Setting $D'(x) = 0$ is algebraically equivalent to finding the intersections of an auxiliary polynomial $H_p(x)$ with a horizontal target line $y = \frac{2}{ap}$, where:

$$H_p(x) = x(1-x)\left(x^{p-1} + (1-x)^{p-1}\right) \tag{24}$$

While $H_p(x)$ is trivially strictly monotonic on $(0, 1/2)$ for $p \leq 3$ (as $H_1(x) = 2x(1-x)$ and $H_2(x) = x(1-x)$ are downward-opening parabolas, and $H'_3(x) = (1-2x)^3 > 0$),[9] it loses monotonicity for higher degrees. To rigorously map its geometric profile for $p \geq 4$, let $v = x$ and $w = 1 - x$. Because $x \in (0, 1/2)$, we strictly have $w > v > 0$. The critical points of $H_p(x)$ satisfy $\frac{d}{dx}(v^p w + v w^p) = 0$, which yields:

$$w^p - v^p = pvw(w^{p-2} - v^{p-2}) \tag{25}$$

Dividing both sides by v^p and substituting the ratio $z = w/v$ (noting that $z > 1$ because $w > v$), this requirement transforms into finding the roots of the polynomial $K(z)$:

$$K(z) = z^p - pz^{p-1} + pz - 1 = 0 \tag{26}$$

[9] In all three cases H_p is strictly monotone increasing on $(0, 1/2)$, so the equation $H_p(x) = \frac{2}{ap}$ has exactly one solution on $(0, 1/2)$ whenever $\frac{2}{ap} < 2^{-p}$, i.e., whenever $a > \frac{2^{p+1}}{p}$. Thus $D'(x) = 0$ exactly once on $(0, 1/2)$, D achieves a unique interior maximum, and since $D(0^+) = -\infty$, $D(1/2) = 0$, and $D'(1/2) = 4 - ap \cdot 2^{1-p} < 0$ in the unstable regime, D crosses zero from below exactly once, yielding a unique root $\sigma_a \in (0, 1/2)$.

We evaluate $K(z)$ strictly for the domain $z > 1$. At the boundary, $K(1) = 0$, and the initial slope evaluates to $K'(1) = p(3-p)$. Because $p \geq 4$, the initial slope is strictly negative, forcing the polynomial to initially dip below zero.

To prove $K(z)$ subsequently crosses zero exactly once, we analyze its concavity via the second derivative:

$$K''(z) = p(p-1)z^{p-3}(z-(p-2)) \tag{27}$$

Because $p(p-1)z^{p-3} > 0$, the sign of $K''(z)$ is dictated entirely by $(z-(p-2))$. Thus, $K''(z) < 0$ for $z \in (1, p-2)$ and $K''(z) > 0$ for $z > p-2$. This implies the slope $K'(z)$ strictly decreases to a unique global minimum at $z = p-2$ before strictly and monotonically increasing to $+\infty$. Consequently, $K'(z)$ transitions from negative to positive exactly once.

Because $K'(z)$ has exactly one root, $K(z)$ starts at 0, decreases to a unique global minimum, and then strictly increases to $+\infty$, guaranteeing it crosses zero exactly once for $z > 1$. Translated back to our original variables, this mathematically guarantees that $H_p(x)$ possesses exactly one local extremum—a global maximum—on the interval $(0, 1/2)$.

Thus, for $p \geq 4$, $H_p(x)$ ascends from $H_p(0) = 0$ to its unique peak, then strictly descends to $H_p(1/2) = 2^{-p}$ (for $p \leq 3$, it simply strictly ascends to 2^{-p}). Under the unstable regime, our threshold parameter is $a > \frac{2^{p+1}}{p}$, which dictates that the horizontal target line satisfies $\frac{2}{ap} < 2^{-p}$. Because this target line is mathematically trapped below 2^{-p}, it is forced to intersect $H_p(x)$ exactly once on its ascending branch for all $p \geq 1$.

Therefore, $D'(x) = 0$ exactly once. As x increases, $H_p(x)$ initially increases from 0, causing $D'(x)$ to transition from strictly positive to strictly negative. Because $D(x) \to -\infty$ as $x \to 0^+$ and $D(1/2) = 0$, the transition of $D'(x)$ dictates that $D(x)$ increases from $-\infty$ to a positive maximum before decreasing to 0 at $x = 1/2$. By the Intermediate Value Theorem, $D(x)$ must cross zero exactly once on $(0, 1/2)$. We denote this unique root as σ_a.

By symmetry, $1 - \sigma_a$ is the unique root in $(1/2, 1)$. Thus, g_a has exactly three fixed points: σ_a, $1/2$, and $1 - \sigma_a$.

Because $g_a^2 = f_a^2$, any period-2 point of g_a must be a fixed point of f_a^2. We now rigorously determine the complete set of solutions to $f_a^2(x) = x$ on the interior $(0, 1)$. Let $h_p(x) = x^p - (1-x)^p$. From the definition of $f_a(x)$, we have the algebraic identity $\frac{1-f_a(x)}{f_a(x)} = \frac{1-x}{x}\exp(ah_p(x))$. Imposing the period-2 condition yields:

$$f_a(f_a(x)) = \frac{f_a(x)}{f_a(x) + (1-f_a(x))\exp(ah_p(f_a(x)))} = x \tag{28}$$

Dividing the numerator and denominator by $f_a(x)$ and substituting our identity gives:

$$\frac{1}{1 + \frac{1-x}{x}\exp\left(a[h_p(x) + h_p(f_a(x))]\right)} = x \tag{29}$$

Multiplying by the denominator and simplifying yields $(1-x)\exp(a[h_p(x) + h_p(f_a(x))]) = 1 - x$. Because $x \in (0, 1)$, we divide by $1 - x$ and take the natural logarithm. Since $a > 0$, the condition simplifies exactly to $h_p(f_a(x)) = -h_p(x)$.

By definition, $-h_p(x) = (1-x)^p - x^p = h_p(1-x)$. Therefore, the period-2 condition is entirely equivalent to $h_p(f_a(x)) = h_p(1-x)$. We note that $h_p'(x) = px^{p-1} + p(1-x)^{p-1} > 0$ for all $x \in (0, 1)$ and $p \geq 1$. Because h_p is strictly monotonic, it is injective, which strictly mandates that $f_a(x) = 1 - x$.

Thus, the fixed points of f_a^2 are mathematically identical to the solutions of $f_a(x) = 1 - x$, which are exactly the fixed points of g_a. Therefore, g_a has absolutely no points of period 2 strictly on the open interior $(0, 1)$.

To rigorously apply Lemma 6, which requires a closed interval, we must analyze the boundaries. Extending $f_a(x)$ continuously to $[0, 1]$ yields $f_a(0) = 0$ and $f_a(1) = 1$, which corresponds to $g_a(0) = 1$ and $g_a(1) = 0$. Thus, $\{0, 1\}$ forms a period-2 orbit for g_a on the boundary. However, evaluating the derivative of the original map at the boundaries yields $f_a'(0) = f_a'(1) = e^a$. Because $a > 0$, we strictly have $e^a > 1$, making 0 and 1 strictly repelling fixed points of f_a.

Because the boundaries are strictly repelling ($f_a'(0) > 1$), the map locally expands away from the boundaries, satisfying $f_a(x) > x$ for sufficiently small $x > 0$. Therefore, there exists a small $\epsilon > 0$ such that the closed, compact subinterval $I = [\epsilon, 1-\epsilon]$ is forward-invariant under f_a (and consequently under g_a), eventually absorbing all asymptotic trajectories originating in $(0, 1)$. Because I is a closed interval and $g_a|_I$ has no period-2 points, we may validly apply Lemma 6. The g_a-trajectory of every $x \in I$ must converge to one of the fixed points of g_a. Consequently, for any initial condition $x \in (0, 1)$, the sequence $f_a^n(x)$ must asymptotically converge either to the equilibrium $1/2$ or to the interior 2-cycle $\{\sigma_a, 1-\sigma_a\}$.

It remains to formally prove that the basin of attraction of the fixed point $1/2$ has Lebesgue measure zero, ensuring that almost every initial condition is captured by the 2-cycle.

Let $W^s(1/2) = \{x \in (0,1) \mid \lim_{n\to\infty} f_a^n(x) = 1/2\}$ denote the stable set of the equilibrium. In the unstable regime $a > \frac{2^{p+1}}{p}$, we established in Part 1 that the derivative at the fixed point satisfies $|f_a'(1/2)| > 1$. Because f_a is continuously differentiable, there exists a local neighborhood $U = (1/2 - \epsilon, 1/2 + \epsilon)$ and a constant $\lambda > 1$ such that $|f_a'(x)| \geq \lambda$ for all $x \in U$.

Suppose there exists a trajectory that asymptotically converges to $1/2$ without ever landing exactly on it. By definition of convergence, there exists an integer N such that for all $n \geq N$, $f_a^n(x) \in U$. Applying the Mean Value Theorem, for any $n \geq N$, there exists a $c \in U$ such that:

$$|f_a^{n+1}(x) - 1/2| = |f_a(f_a^n(x)) - f_a(1/2)| = |f_a'(c)||f_a^n(x) - 1/2| \geq \lambda |f_a^n(x) - 1/2| \tag{30}$$

By induction, for any $k > 0$, we strictly have:

$$|f_a^{N+k}(x) - 1/2| \geq \lambda^k |f_a^N(x) - 1/2| \tag{31}$$

Because we assumed the trajectory never lands exactly on the fixed point, $|f_a^N(x) - 1/2| > 0$. Since $\lambda > 1$, the term $\lambda^k \to \infty$ as $k \to \infty$, which strictly forces the trajectory to leave the bounded neighborhood U, contradicting the assumption of asymptotic convergence.

Therefore, asymptotic convergence to $1/2$ is strictly impossible. The only way a trajectory can converge to $1/2$ is if it maps exactly onto $1/2$ in finite time. Mathematically, this implies the stable set is exactly the countable union of all pre-images of $1/2$:

$$W^s(1/2) = \bigcup_{k=0}^{\infty} f_a^{-k}(1/2) \tag{32}$$

Because the latency functions x^p are polynomials, the update map $f_a(x)$ is real-analytic on the open domain $(0, 1)$. The pre-images of $1/2$ correspond to the roots of the analytic equations $f_a^k(x) - 1/2 = 0$. Because $f_a^k(x)$ is not identically $1/2$, the identity theorem for analytic functions guarantees that these roots are strictly isolated within $(0, 1)$.

While an infinite number of roots could theoretically accumulate at the boundaries 0 or 1, any set of strictly isolated points in $\mathbb{R}$ is at most countable. Therefore, for any finite iterate k, the set of pre-images $f_a^{-k}(1/2)$ is a countable set.

Consequently, the entire stable set $W^s(1/2) = \bigcup_{k=0}^{\infty} f_a^{-k}(1/2)$ is a countable union of countable sets, making $W^s(1/2)$ strictly countable. Because any countable subset of $\mathbb{R}$ has a Lebesgue measure of exactly zero, the set of initial conditions asymptotically converging to $1/2$ is a zero-measure set. The trajectory of almost every $x \in (0, 1)$ is thus captured by the global period-2 attractor.

Part 3: The 2^p Asymptotic Inefficiency. Accounting for the total population mass N, the true latency on a path with fractional flow x is $c(Nx) = (Nx)^p = N^p x^p$. The optimal social cost at the Wardrop equilibrium ($x^* = 1/2$) is therefore $SC_{opt} = N(1/2) \cdot N^p(1/2)^p + N(1/2) \cdot N^p(1/2)^p = N^{p+1}(1/2)^p$.

Because the 2-cycle satisfies $g_a(\sigma_a) = \sigma_a$, we can solve exactly for the system parameter a as a function of the amplitude σ_a:

$$a = \frac{2 \ln\left(\frac{1-\sigma_a}{\sigma_a}\right)}{(1-\sigma_a)^p - \sigma_a^p} \tag{33}$$

Because $a(\sigma_a)$ is continuous and finite on the open interval $(0, 1/2)$, any sequence of states yielding $a \to \infty$ must topologically accumulate at the boundaries of this domain. As $\sigma_a \to 1/2$, the

denominator approaches zero, but the numerator concurrently approaches zero; applying L'Hôpital's rule yields a finite limit of $\frac{2^{p+1}}{p}$[10]. Because divergence is strictly impossible at any interior point and at the boundary $\sigma_a \to 1/2$, the condition $a \to \infty$ leaves only one possibility, forcing the limit $\lim_{a\to\infty} \sigma_a = 0$.

The system flow strictly alternates between the states σ_a and $1 - \sigma_a$. Due to the symmetric identical latency functions, the absolute social cost evaluates to exactly the same value at both states of the cycle, yielding a constant empirical time-averaged social cost of $SC_{emp} = N^{p+1}[\sigma_a^{p+1} + (1 - \sigma_a)^{p+1}]$. Evaluating the limit as $a \to \infty$ ($\sigma_a \to 0$) yields $\lim_{a\to\infty} SC_{emp} = N^{p+1}$. The ultimate ratio of empirical performance to optimal performance perfectly cancels the population scaling, yielding exactly $SC_{emp}/SC_{opt} = N^{p+1}/[N^{p+1}(1/2)^p] = 2^p$. □

I.1 Construction of L_2 Chaos (Reference for Figure 16)

We construct a non-atomic game to prove that Li-Yorke chaos emerges under L_2 regularization (PGD/FTRL) without activating boundary projections. Let Link 2 have constant latency $c_2(1 - x) = 2/3$. Let Link 1 have piecewise linear latency $c_1(x) = \frac{2}{3}x$ for $x \leq 0.7$, and $c_1(x) = 5x - 3.033\ldots$ for $x > 0.7$. The potential gradient is $g(x) = c_1(x) - c_2(1 - x)$.

Under PGD with learning rate $\eta = 1$, the unprojected map $f(x) = x - g(x)$ traces a forward orbit from $x_0 = 0.1$:

1. $x_0 = 0.1 \implies g(0.1) = -0.6 \implies x_1 = 0.7$
2. $x_1 = 0.7 \implies g(0.7) = -0.2 \implies x_2 = 0.9$
3. $x_2 = 0.9 \implies g(0.9) = 0.8 \implies x_3 = 0.1$

The sequence $0.1 \to 0.7 \to 0.9 \to 0.1$ is a period-3 cycle. Because $\max(x_t) = 0.9 < 1$ and $\min(x_t) = 0.1 > 0$, the Euclidean projection operator $\Pi_{[0,1]}$ remains the identity map. Because the projection is inactive, the unconstrained dual variable in FTRL is exactly equal to the primal variable in PGD. By the Li-Yorke Theorem Li and Yorke (1975), the existence of a period-3 interior orbit guarantees the presence of Li-Yorke chaos for both L_2 learning dynamics.

[10]To rigorously evaluate the $0/0$ indeterminate form, we apply L'Hôpital's rule: $\lim_{\sigma_a\to 1/2} a(\sigma_a) = \lim_{\sigma_a\to 1/2} \frac{\frac{d}{d\sigma_a}[2\ln(1-\sigma_a)-2\ln(\sigma_a)]}{\frac{d}{d\sigma_a}[(1-\sigma_a)^p-\sigma_a^p]} = \frac{-8}{-p(1/2)^{p-1}-p(1/2)^{p-1}} = \frac{-8}{-p2^{2-p}} = \frac{2^{p+1}}{p}$. This mathematically recovers the exact local stability threshold from Part 1, dynamically verifying that the period-2 orbit bifurcates continuously from the fixed point $1/2$ at this parameter value.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: All proofs of theorems of are included.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: All theorems set exactly the parameters of their applicability.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All proofs and assumptions are included.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [N/A]

Justification: This is a theory paper. Some simulations are provided to help with the understanding of the theorems. Exact descriptions of the game dynamics are included.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [N/A]

   Justification: This is a theory paper.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [N/A]

   Justification: This is a theory paper.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [N/A]

   Justification: This is a theory paper.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [N/A]

   Justification: This is a theory paper.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: This is a theory paper.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A]

    Justification: This is a theory paper.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: This is a theory paper.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [N/A]

Justification: This is a theory paper.

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: This is a theory paper.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: This is a theory paper.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [N/A]

Justification: This is a theory paper.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: This is a theory paper.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.