

Characterizing the Convergence of Game Dynamics via Potentialness

Anonymous authors

Paper under double-blind review

Abstract

Understanding the convergence landscape of multi-agent learning is a fundamental problem of great practical relevance in many applications of artificial intelligence and machine learning. While it is known that learning dynamics converge to Nash equilibrium in potential games, the behavior of dynamics in many important classes of games that do not admit a potential is poorly understood. To measure how “close” a game is to being potential, we consider a distance function, that we call “potentialness”, and which relies on a strategic decomposition of games introduced by Candogan et al. (2011). We introduce a numerical framework enabling the computation of this metric, which we use to calculate the degree of “potentialness” in generic matrix games, as well as (non-generic) games that are important in economic applications, namely auctions and contests. Understanding learning in the latter games has become increasingly important due to the wide-spread automation of bidding and pricing with no-regret learning algorithms. We empirically show that potentialness decreases and concentrates with an increasing number of agents or actions; in addition, potentialness turns out to be a good predictor for the existence of pure Nash equilibria and the convergence of no-regret learning algorithms in matrix games. In particular, we observe that potentialness is very low for complete-information models of the all-pay auction where not pure Nash equilibrium exists, and much higher for Tullock contests, first-, and second-price auctions, explaining the success of learning in the latter. In the incomplete-information version of the all-pay auction, a pure Bayes-Nash equilibrium exists and it can be learned with gradient-based algorithms. Potentialness nicely characterizes these differences to the complete-information version.

1 Introduction

Multi-agent systems and multi-agent learning have drawn considerable attention, owing to their massive deployment in machine learning enabled applications and their increasing economic impact, with agents automatically ordering goods, setting prices, or bidding in auctions (Yang & Wang, 2020). In contrast to other applications of machine learning, the input in multi-agent learning is non-stationary and depends on the strategic behavior and learning of other agents, which leads to challenging computation and learning problems that go well beyond the “business as usual” framework of empirical risk minimization.

The literature on learning in games has a long history and asks what type of equilibrium behavior (if any) may arise in the long run of a process of learning and adaptation, in which agents are trying to (myopically) maximize their payoff while adapting to the actions of other agents through repeated interactions (Fudenberg & Levine, 1998; Hart & Mas-Colell, 2003). To that end, many learning algorithms have been developed ranging from iterative best-response to first-order online optimization algorithms in which agents follow their utility gradient in each step (Mertikopoulos & Zhou, 2019; Bichler et al., 2023).

It is known that the dynamics of learning agents do not always converge to a Nash equilibrium (Daskalakis et al., 2010; Flokas et al., 2020): they may cycle, diverge, or be chaotic, even in zero-sum games, where computing Nash equilibria is tractable (Mertikopoulos et al., 2018; Palaiopanos et al., 2017). In fact, (Hart & Mas-Colell, 2003) showed that in general uncoupled dynamics do not converge to Nash equilibrium in all

games. With this negative result at hand, there lacks a comprehensive characterization of games that are “learnable”; however our understanding is much better for certain classes of games.

On the one hand, no-regret algorithms and dynamics do not converge in games with only mixed Nash equilibria (Flokas et al., 2020; Giannou et al., 2021a). On the other hand, if the underlying game is an exact potential game, then no-regret and other learning algorithms converge to an ε -NE with minimal exploration (Heliou et al., 2017; Mertikopoulos et al., 2024). Importantly, even though many classes of games of interest are not exact potential games (Candogan et al., 2013a;b), experimental evidence shows that even if games are not exact potential games, learning dynamics often converge to Nash equilibrium. Beyond exact potential games there is no good understanding of the behavior of learning dynamics.

Partially motivated by this, Candogan et al. (2011) introduced a game decomposition that allows one to characterize how “close” a game is to being potential by resolving it into a potential and a harmonic component (plus a “non-strategic” part which does not affect the game’s equilibrium structure and unilateral payoff differences). In contrast to potential games, the exponential weight / replicator dynamic – perhaps the most widely studied no-regret dynamics – does not converge in any harmonic game, and instead exhibit a quasi-periodic behavior known as Poincaré recurrence (Legacci et al., 2024). Based on this dynamical dichotomy between potential and harmonic games, we consider a measure of *potentialness* and analyze generic normal-form games, as well as several classes of games stemming in economic applications. We show that potentialness provides a useful indicator for both the existence of pure Nash equilibria and the convergence of no-regret algorithms. While the latter was already part of the motivation of Candogan et al. (2013a) the former is a novel connection emerging from our empirical exploration. In particular, we find that the average potentialness in random games decreases and concentrates on a value with increasing numbers of agents or actions: Games with a potentialness below 0.4 rarely exhibit convergence, while games with values larger than 0.6 mostly do. We also categorize specific games, such as Jordan’s matching pennies game (Jordan, 1993), where learning dynamics are known not to converge (or, more precisely, to converge to a non-terminating cycle of best responses).

Economically motivated games such as auctions and contests have more structure in the payoffs. The analysis of these games is relevant today because pricing and bidding are increasingly being automated via learning agents. Learning agents are used to bid in display ad auctions, but they are also used by automated agents that set prices on online platforms such as Amazon (Chen et al., 2016). Whether we can expect the dynamics of such multi-agent interactions to converge to an equilibrium or exhibit inefficient price cycles or even chaos, is an economically important question.

We find that the potentialness is very low for all-pay auctions and much higher for Tullock contests, first- and second-price auctions. Indeed, our experiments show that learning algorithms do not converge for all-pay auctions, but they do so for the other economic games. The low potentialness of the all-pay auction also connects to the fact that it does not possess a pure Nash equilibrium. Overall, potentialness provides a single indicator for convergence to a Nash equilibrium that is independent of the initializations in individual experiments. This is in stark contrast to a brute-force approach, where initial conditions need to be explored.

2 Related Literature

In this paper, we focus throughout on repeated normal-form games, where players move simultaneously and receive the payoffs as specified by the combination of actions played. The theory of learning in games examines what kind of equilibrium arises as a consequence of a process of learning and adaptation, in which agents are trying to maximize their payoff while learning about the actions of other agents in repeated games (Fudenberg & Levine, 1998). For example, fictitious play is a natural method by which agents iteratively search for a pure NE and play a best response to the empirical frequency of play of other players (Brown, 1951). Several algorithms have been proposed based on best or better response dynamics for finite and simultaneous-move games, ultimately leading to a vast corpus of literature (Abreu & Rubinstein, 1988; Hart & Mas-Colell, 2000; Fudenberg & Levine, 1998; Hart & Mas-Colell, 2003; Young, 2004), while more recent contributions draw on first-order online optimization methods such as online gradient descent or online mirror descent to study the question of convergence (Mertikopoulos & Zhou, 2019; Bichler et al., 2023).

Learning dynamics do not always converge to equilibrium (Daskalakis et al., 2010; Flokas et al., 2020). Learning algorithms can cycle, diverge, or be chaotic; even in zero-sum games, where the NE is tractable (Mertikopoulos et al., 2018; Palaiopoulos et al., 2017). Sanders et al. (2018) argues that chaos is typical behavior for more general matrix games. Recent results have shown that learning dynamics do not converge in games with mixed Nash equilibria (Giannou et al., 2021a;b). On the positive side, Mertikopoulos & Zhou (2019) showed conditions for which no-external-regret learning algorithms result in a NE in finite games if they converge. However, in general, the dynamics of matrix games can be arbitrarily complex and hard to characterize (Andrade et al., 2021).

While there is no comprehensive characterization of games that are “learnable” and one cannot expect that uncoupled dynamics lead to NE in all games (Hart & Mas-Colell, 2003), there are some important results regarding the broad class of no-regret learning algorithms. One can distinguish between internal (or conditional) regret and a weaker version, called ‘external (or unconditional) regret’. External regret compares the performance of an algorithm to the best single action in retrospect; while internal regret allows one to modify the online action sequence by changing every occurrence of a given action by an alternative action. For learning rules that satisfy the stronger no-internal regret condition, the empirical frequency of play converges to the game’s set of correlated equilibria (Foster & Vohra, 1997; Hart & Mas-Colell, 2000). The set of correlated equilibria (CE) is a non-empty, convex polytope that contains the convex hull of the game’s Nash equilibria. The coordination in CE can be implicit via the history of play (Foster & Vohra, 1997). On the other hand, algorithms that are no-external-regret learners converge by definition to the set of coarse correlated equilibria (CCE) in finite games (Foster & Vohra, 1997; Hart & Mas-Colell, 2000). This set, in turn, contains the set of CE such that we get $NE \subset CE \subset CCE$. In contrast to correlated equilibria, in a coarse correlated equilibrium, every player could be playing a strictly dominated strategy for all time (Viossat & Zapechelnyuk, 2013), which makes CCE a fairly weak, non-rationalizable solution concept.

An important class of games in which a variety of learning algorithms converge to an NE is *potential games*. In exact potential games, a pure NE exists, and the change in the utility of any player when moving from one strategy profile to another is exactly equal to the change in a potential function. Any sequence of improvements by players converges to a pure NE (Heliou et al., 2017; Christodoulou et al., 2012; Anagnostides et al., 2022). Congestion games are equivalent to the class of exact potential games (Monderer & Shapley, 1996). However, many games are not exact potential games, yet no-regret algorithms converge.

In their seminal work Candogan et al. (2011), leveraging the so-called *combinatorial Hodge theorem* (Jiang et al., 2011; Friedman, 1996; Munkres, 1984; Dodziuk, 1976), decompose a finite normal form game into three components with distinctive strategical properties. This decomposition can then be used to approximate a given game with a potential game, which can be used to characterize the limiting behavior of dynamics in the original game (Candogan et al., 2013a;b). In particular, Candogan et al. (2013a) examine the convergence of best-response and logit-best-response dynamics – in the sense of Blum & Kalai (1999) – and they show that, if only one player updates per turn, the dynamics remain convergent in slight perturbations of potential games.¹

An important caveat is that the class of dynamics considered by Candogan et al. (2013a) can lead to positive regret and, moreover, in contrast to the setting under study, players cannot move simultaneously, but only one after another. Our focus in this paper is more general, as we seek to understand the behavior of no-regret dynamics over the entire spectrum of potentialness, and to understand where the convergence of regularized no-regret learning breaks down. In so doing, we also provide a first positive answer to the open question stated by Candogan et al. (2013a), who asked whether the replicator dynamic or follow-the-regularized-leader (a staple of no-regret learning) remain convergent in small, near-potential perturbations of potential games.

3 Preliminaries

In this section, we provide some basic definitions. We begin by defining the concept of a normal-form game, a Nash equilibrium, and a potential game.

¹Importantly, the dynamics considered by Candogan et al. (2013a) are not the simultaneous best-reply dynamics of Gilboa & Matsui (1991) or the logit dynamics of Fudenberg & Levine (1999); Hofbauer & Sandholm (2009), but rather turn-by-turn updates where each player observes the play of their opponents and plays a (logit) best-response.

A normal-form game is a representation in game theory that defines the strategies available to each player, their corresponding payoffs, and the resulting outcomes in a simultaneous and strategic interaction.

Definition 1 (Normal-form games). *A finite normal-form game with N players can be described as a tuple $G = (\mathcal{N}, \mathcal{A}, u)$.*

- \mathcal{N} is a finite set of N players indexed by i .
- $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_N$, where \mathcal{A}_i is a finite set of actions available to player i . A vector $a = (a_1, \dots, a_N) \in \mathcal{A}$ is referred to as an action profile. Let $A = \sum_i A_i$ with $A_i = |\mathcal{A}_i|$.
- $u = (u_1, \dots, u_N)$, where $u_i : \mathcal{A} \mapsto \mathbb{R}$ is a payoff or utility function for player $i \in \mathcal{N}$.

One type of strategy available to a player i in a normal-form game is to select a single action a_i and play it. Such a strategy is called a *pure strategy*. But a player is also able to randomize over the set of available actions according to some probability distribution. Such a strategy is called a *mixed strategy*.

Definition 2 (Mixed strategy). *In a normal-form game $G = (\mathcal{N}, \mathcal{A}, u)$ the set of mixed strategies for player i is $S_i = \Delta(\mathcal{A}_i)$, where $\Delta(\mathcal{A}_i)$ is the set of all probability distributions (aka lotteries) over \mathcal{A}_i .*

A Nash equilibrium is a situation in a strategic interaction where each player's strategy is optimal given the strategies chosen by all other players, and no player has an incentive to unilaterally deviate from their chosen strategy.

Definition 3 (Nash equilibrium). *In a normal-form game $G = (\mathcal{N}, \mathcal{A}, u)$, a strategy profile $s^* = (s_1^*, s_2^*, \dots, s_N^*) \in S_1 \times \dots \times S_N$ is a Nash equilibrium if, for every player $i \in \mathcal{N}$, the following condition holds:*

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*), \quad \forall s_i \in S_i, \quad (1)$$

where $s_{-i}^* = (s_1^*, s_2^*, \dots, s_{i-1}^*, s_{i+1}^*, \dots, s_N^*)$.

A potential game is a type of game in which the strategic interaction can be characterized by a potential function, and the individual players' incentives align with the minimization or maximization of this function, facilitating the analysis of equilibrium and strategic dynamics.

Definition 4 (Potential game). *A game $G = (\mathcal{N}, \mathcal{A}, u)$ is a potential game if there exists a potential function $\phi : \mathcal{A} \rightarrow \mathbb{R}$ such that for every player $i \in \mathcal{N}$ and every pair of action profiles $a, a' \in \mathcal{A}$ that differ only in the action of player i , i.e., $a_i \neq a'_i$ and $a_{-i} = a'_{-i}$, the following condition holds:*

$$u_i(a) - u_i(a') = \phi(a) - \phi(a'). \quad (2)$$

4 Potentialness of a game

In this section we present an overview of a combinatorial decomposition technique for finite games in normal form introduced by Candogan et al. (2011), and we use it to define the *potentialness* of a game, a measure of closeness to being a potential game. This measure is closely related to the *maximum pairwise difference* introduced by Candogan et al. (2013a), and can be used as a predictor for the existence of strict pure Nash-equilibria (SPNE) in a game, and of the limiting behavior of learning dynamics thereof.

Deviation map Given a finite game in normal form $G = (\mathcal{N}, \mathcal{A}, u)$, pairs of strategy profiles (a, a') that differ only in the strategy of one player are called *unilateral deviations*, and their space is denoted by \mathcal{E} . Representing a game in terms of *utility differences between unilateral deviations* rather than in terms of utilities themselves captures the strategic structure of a game in an effective way, in the sense that games with different utilities but identical utility differences between unilateral deviations share the same set of Nash equilibria Candogan et al. (2011).

To achieve this representation of a game $G = (\mathcal{N}, \mathcal{A}, u)$ consider its *response graph* $\Gamma(\mathcal{N}, \mathcal{A})$, that is the graph with a node for each of the A pure strategy profile in \mathcal{A} , and an edge for each of the $E := |\mathcal{E}| = \frac{A}{2} \sum_{i \in \mathcal{N}} (A_i - 1)$ unilateral deviations in \mathcal{E} (Biggar & Shames, 2023).

This graph is an instance of a *simplicial complex* K , that is, loosely speaking, a collection of oriented k -dimensional faces (points, segments, triangles, tetrahedrons, ...) with $k \in \{0, 1, \dots\}$ (Jonsson, 2007). Given a simplicial complex K one can build a family of vector spaces $\{C_k\}_{k=0,1,\dots}$, called *chain groups*, where each chain group C_k is the space spanned by the k -dimensional faces of the complex K , i.e., the space of assignments of a real number to each k -dimensional face of the complex (Munkres, 1984).

In this work, we restrict our attention to the chain groups C_0 and C_1 on the response graph $\Gamma(\mathcal{N}, \mathcal{A})$ of a game. $C_0 \cong \mathbb{R}^A$ is the space of assignments of a real number to each vertex $a \in \mathcal{A}$, and $C_1 \cong \mathbb{R}^E$ is the space of assignments of a real number to each edge $(a, a') \in \mathcal{E}$; an element in C_1 is called *flow* on the graph.

Note that the potential function ϕ of a potential game is the assignment of a number to each pure strategy profile $a \in \mathcal{A}$, and as such, it is an element of C_0 . In a similar way, observe that a utility function u is the assignment of a number $u_i(a) \in \mathbb{R}$ to each pure strategy profile $a \in \mathcal{A}$ for each player $i \in \mathcal{N}$, and as such it can be considered as an element of N copies of C_0 , that is $u \in \mathcal{U} := C_0 \times \dots \times C_0$.

The key observation is that the differences between unilateral deviations of a game $G = (\mathcal{N}, \mathcal{A}, u)$ can be represented as a special flow Du in C_1 on $\Gamma(\mathcal{N}, \mathcal{A})$, called *deviation flow of the game*, by means of the *deviation map*:

Definition 5 (Deviation map). *The deviation map of the game $G = (\mathcal{N}, \mathcal{A}, u)$ is the linear map*

$$\begin{aligned} D : \mathcal{U} &\rightarrow C_1 \\ u &\mapsto Du \end{aligned} \tag{3}$$

such that, for all $u \in \mathcal{U}$ and all $(a, a') \in \mathcal{E}$

$$(Du)_{(a, a')} = u_i(a') - u_i(a) \tag{4}$$

for $i \in \mathcal{N}$ such that $a_i \neq a'_i$.

In words, the deviation flow $Du \in C_1$ assigns to each edge $(a, a') \in \mathcal{E}$ of the response graph, i.e., to every unilateral deviation of the game, the utility difference of the deviating player $i \in \mathcal{N}$. We call $\text{Im } D \in C_1$ the space of *feasible flows* on the response graph of a game.

Potential flows Representing a game u via its deviation flow Du preserves all the strategic information of the game and allows for a concise characterization of potential games.

Definition 6 (Gradient map). *The gradient map² is the linear map*

$$\begin{aligned} d_0 : C_0 &\rightarrow C_1 \\ \phi &\mapsto d_0\phi \end{aligned} \tag{5}$$

such that

$$(d_0\phi)_{(a, a')} = \phi(a') - \phi(a) \tag{6}$$

for all $\phi \in C_0$ and all $(a, a') \in \mathcal{E}$.

It is now immediate to show that

Proposition 1. *A game $G = (\mathcal{N}, \mathcal{A}, u)$ is potential with potential function ϕ if and only if $Du = d_0\phi$ for some $\phi \in C_0$.*

Proof. Let $G = (\mathcal{N}, \mathcal{A}, u)$ be a potential game with potential function ϕ . Then

$$Du_{(a, a')} = u_i(a') - u_i(a) = \phi(a') - \phi(a) = (d_0\phi)_{(a, a')}$$

for all $(a, a') \in \mathcal{E}$, where $i \in \mathcal{N}$ is the actor of the deviation (a, a') . Thus, $Du = d_0\phi$. Conversely, let $G = (\mathcal{N}, \mathcal{A}, u)$ be a game with $Du = d_0\phi$ for some $\phi \in C_0$. Then

$$u_i(a') - u_i(a) = Du_{(a, a')} = (d_0\phi)_{(a, a')} = \phi(a') - \phi(a)$$

for all $i \in \mathcal{N}$ and all $(a, a') \in \mathcal{E}$ acted by i . Thus, the game is potential with potential function ϕ . \square

²The gradient map is an instance of so-called *co-boundary maps*; see Munkres (1984) for details.

The proposition means that the space of potential games is the linear subspace $D^{-1} \text{Im } d_0 \subset \mathcal{U}$; in light of this we call $\text{Im } d_0 \subset C_1$ the space of *potential flows*.

Harmonic flows Endowing C_0 and C_1 with an Euclidean-like inner product $\langle \cdot, \cdot \rangle_k$ one can define the *divergence map* $\partial_1 := d_0^\dagger : C_1 \rightarrow C_0$ as the adjoint operator of the gradient map, namely $\langle X, d_0 \phi \rangle_1 = \langle \partial_1 X, \phi \rangle_0$ for all $X \in C_1$ and all $\phi \in C_0$. The flows in the subspace $\ker \partial_1 \in C_1$ are called *harmonic flows*³, and the games whose flow is harmonic, i.e., the games in $D^{-1} \ker \partial_1 \subset \mathcal{U}$, are called *harmonic games*.

Hodge decomposition of feasible flows Leveraging the *combinatorial Hodge decomposition theorem*⁴ it can be shown that potential flows and harmonic flows completely characterize feasible flows:

Theorem (Candogan et al. (2011) — Combinatorial Hodge decomposition of feasible flows). *The space of feasible flows is the orthogonal direct sum of the subspaces of potential flows and harmonic flows:*

$$\text{Im } D = \text{Im } d_0 \oplus \ker \partial_1 \quad (7)$$

Equivalently, every feasible Du flow can be decomposed in a unique way as $Du = Du_p + Du_h$, where the potential flow $Du_p \in \text{Im } d_0$ is the orthogonal projection of Du onto $\text{Im } d_0$, and the harmonic flow $Du_h \in \ker \partial_1$ is the orthogonal projection of Du onto $\ker \partial_1$. Note that in the space of feasible flows there is no non-strategic component as the latter has vanishing flow Du .

This decomposition in the space $\text{Im } D \subset C_1$ of feasible flows is sufficient to introduce the measure of potentialness used in this work; in the supplementary material we discuss how to obtain a corresponding decomposition in the space \mathcal{U} of payoffs, allowing to decompose in a unique way a game $u \in \mathcal{U}$ as $u = u_p + u_h + u_\kappa$, where u_p is a *normalized potential game*, u_h a *normalized harmonic game*, and u_κ a *non-strategic game*; c.f. Figure 1 for an example. Endowing \mathcal{U} with an inner product structure Candogan et al.

	Payoff	Potential	Harmonic	Non-Strategic
Agent 1	1.00 0.00 0.00	0.17 0.17 -0.33	0.50 -0.50 -0.00	0.33 0.33 0.33
	0.00 1.00 0.00	-0.33 0.17 0.17	-0.00 0.50 -0.50	0.33 0.33 0.33
	0.00 0.00 1.00	0.17 -0.33 0.17	-0.50 0.00 0.50	0.33 0.33 0.33
Agent 2	0.00 1.00 0.00	0.17 0.17 -0.33	-0.50 0.50 -0.00	0.33 0.33 0.33
	0.00 0.00 1.00	-0.33 0.17 0.17	-0.00 -0.50 0.50	0.33 0.33 0.33
	1.00 0.00 0.00	0.17 -0.33 0.17	0.50 -0.00 -0.50	0.33 0.33 0.33

Figure 1: Decomposition of the Shapley game.

(2011) show that u_p is (up to the non-strategic game u_κ) the potential game *closest* to u . This allows Candogan et al. (2013a;b) to characterize the limiting behavior of dynamics in the game u in terms of the properties of the potential game u_p and of the *distance* between u and u_p , a concept that is made precise in the next paragraph.

Potentialness The potential component of the deviation flow Du of a game can be used to build a measure of how close to being a potential game the game is. To compute it Candogan et al. (2011) introduce the orthogonal projection onto the subspace of potential flows; by the properties of the Moore-Penrose pseudo-inverse $\tilde{d}_0 : C_1 \rightarrow C_0$ of the gradient map (Golan, 1992) such projection is $e := d_0 \tilde{d}_0 : C_1 \rightarrow C_1$, so that

$$Du_p = eDu \in \text{Im } d_0 \subset C_1 \quad (8)$$

³The term *harmonic* refers in combinatorial Hodge theory Dodziuk (1976) to flows in the kernel of the *Laplacian operator* $\ker \Delta_1 = \ker \partial_1 \cap \ker d_1$, where d_1 is defined analogously to d_0 . As Candogan et al. (2011) show, each feasible flow belongs to $\ker d_1$, making these two definitions of harmonic flows consistent.

⁴See Jiang et al. (2011) for a concise presentation and proof.

Candogan et al. (2013a;b) use the deviation map to define the *maximum pairwise difference* between two games as $\delta(u, u') = \|Du - Du'\|$.⁵ In particular, since the potential component $u_{\mathcal{P}}$ of a game u is (up to a non-strategic game) the potential game closest to the original game, they use the maximum pairwise difference $\delta(u, u_{\mathcal{P}}) = \|Du - Du_{\mathcal{P}}\| = \|Du_h\|$ between a game u and its potential component $u_{\mathcal{P}}$ as a measure of closeness to being a potential game for the game u . In this spirit we give the following definition:

Definition 7 (Potentialness). *The potentialness of a game $G = (\mathcal{N}, \mathcal{A}, u)$ is the real number*

$$P(u) := \frac{\|Du_{\mathcal{P}}\|}{\|Du_{\mathcal{P}}\| + \|Du_h\|} \quad (9)$$

Proposition 2. *The potentialness of a game fulfills*

1. $P(u) \in [0, 1]$
2. $P(u) = 1 \iff \delta(u, u_{\mathcal{P}}) = 0 \iff u$ is a potential game
3. $P(u) = 0 \iff u$ is a harmonic game

Proof. Consider the game $G = (\mathcal{N}, \mathcal{A}, u)$ and its potentialness

$$P(u) := \frac{\|Du_{\mathcal{P}}\|}{\|Du_{\mathcal{P}}\| + \|Du_h\|}$$

1. It is obvious that $P(u) \in [0, 1]$;
2. Since $\delta(u, u_{\mathcal{P}}) = \|Du_h\|$ it is also obvious that $P(u) = 1 \iff \delta(u, u_{\mathcal{P}}) = 0$. Let this be the case, then $Du = Du_{\mathcal{P}} + Du_h = Du_{\mathcal{P}} \in \text{Im } d_0$, so $Du = d_0\phi$ for some $\phi \in C_0$, i.e. the game is potential by Proposition 1.

Conversely, let $G = (\mathcal{N}, \mathcal{A}, u)$ be a potential game. Then $Du_{\mathcal{P}} = Du$, since $Du_{\mathcal{P}}$ is the orthogonal projection of Du onto $\text{Im } d_0$, and such projection leaves Du invariant since $Du = d_0\phi$ itself belongs to $\text{Im } d_0$. Hence, $\|Du_h\| = 0$ (this can also be seen immediately by the unicity of the decomposition $Du = Du_{\mathcal{P}} + Du_h$.)

3. It is obvious that $P(u) = 0 \iff \|Du_{\mathcal{P}}\| = 0$; if this is the case then $Du = Du_h$, so the game is harmonic by definition.

Conversely, if the game is harmonic then $Du = Du_h$ by an argument analogous to the one in point 2., which implies that $\|Du_{\mathcal{P}}\| = 0$.

□

In light of these properties, the potentialness of a game can be used as a concise measure of how close to being a potential game a given game is. In the next section, we investigate the existence of strict pure Nash-equilibria (SPNE) and the limit behavior of learning dynamics in games in function of their potentialness.

Scalability To compute the potentialness of a game represented by $u \in \mathcal{U}$, one must perform two computationally expensive operations: the calculation of the deviation flow Du and a projection to get the potential flow $Du_{\mathcal{P}}$. As these operations involve linear operators, the calculations essentially boil down to matrix-vector multiplications. The projection $Du_{\mathcal{P}} = eDu$, being the most expensive operation, involves the matrix e of dimensions $\dim C_1 \times \dim C_1$. Therefore, in a game where all N agents have the same number of actions, i.e., $A_i = m$ for all $i \in \mathcal{N}$, the time complexity of the matrix-vector multiplication is of the order $\mathcal{O}(N^2 m^{2N+2})$. The runtime required to compute the potentialness in our experiments (on a standard notebook) is visualized in the subsequent plot (see Figure 2). In this figure, we show the average runtime (over 100 runs) for the computation of the potentialness, assuming that the necessary matrices of the linear operators are already given.

⁵Different choices of norm are possible; in this work we use the 2-norm, whereas Candogan et al. (2013a) use the infinity norm.




Figure 2: Runtime for computation of potentialness.

While this leads to a very fast computation of our metric, computing the necessary matrices can be more time-consuming and memory-intensive. Recall that computing the projection $e : C_1 \rightarrow C_1$ requires a matrix multiplication of the gradient map d_0 with its pseudo-inverse \tilde{d}_0 , which also has to be computed first. Since these operators do not depend on the game $u \in \mathcal{U}$ itself, but only on the number of actions and agents, the matrices need to be computed only once for each number of agents and actions and can be stored and later used for other games of the same size. But even storing the matrix can become an issue as the number of entries grows exponentially in the number N of agents, i.e., $\mathcal{O}(N^2 m^{2N+2})$. For instance, in a game with 5 actions for each agent, the number of entries of the projection matrix is of the order 10^4 for $N = 2$ agents, 10^5 for $N = 3$ agents, and 10^7 for $N = 4$ agents. For the larger settings considered in this example, i.e., 3 agents with 12 actions each or 4 agents with 7 actions each, the computation of all necessary matrices takes 2-3 minutes.

5 Numerical Experiments

We conduct our numerical experiments on randomly generated and economically motivated games. Our analysis focuses mainly on two aspects, namely, the existence of strict pure Nash-equilibria (SPNE) and the convergence of learning dynamics with respect to the potentialness of the games.⁶

To analyze the learning dynamics, we focus on online mirror descent (OMD) (Nemirovskij & Yudin, 1983) with an entropic regularization term, which leads to the update steps described in Algorithm 1. OMD is a natural candidate for no-regret algorithms with good regret properties, which falls in the class of follow-the-regularized-leader (FTRL) algorithms (Shalev-Shwartz, 2012).

Algorithm 1: Online Mirror Descent

Input : initial mixed strategies $s_{i,0}$

for $t = 1$ **to** T **do**

for agent $i = 1$ **to** N **do**

 observe gradient $v_{i,t}$;

 set $s_{i,t} \leftarrow \mathcal{P}_{s_{i,t-1}}(\eta_t v_{i,t})$;

end

end

⁶The code will be published on Github after the review process.

The regularization term induces a prox-mapping $\mathcal{P}_x : \mathbb{R}^d \rightarrow \mathbb{R}^d$ in Algorithm 1, which is given by

$$\mathcal{P}_x(y) = \frac{(x_j \exp(y_j))_{j=1}^d}{\sum_{j=1}^d x_j \exp(y_j)}, \quad (10)$$

where $d \in \mathbb{N}$ is the dimension of the strategy space of the respective player, i.e., $d = A_i$. It is known that convergence properties vary with the step-size. We therefore use a step-size sequence η_t of the form $\eta_t = \eta_0 \cdot t^{-\beta}$ for some $\beta \in (0, 1]$, and we say that the algorithm converges to an (approximate) NE, if the relative utility loss, defined by $\ell_i(s_t) = 1 - u_i(s_t)/u_i(br_i, s_{-i,t})$, is less than some predefined tolerance $\varepsilon = 10^{-8}$ for all agents $i \in \mathcal{I}$ within a fixed number of iterations $T = 2000$ for some step-size in the sequence. Note that br_i denotes the best response of agent i given the opponents' strategy profile $s_{-i,t}$.

Before focusing on randomly generated games, let us briefly review some standard games regarding their potentialness. For the first three games in Table 1, we used payoff matrices as defined in Nisan (2007).

Table 1: Potentialness of some matrix games.

Game	Actions	$P(u)$
Matching Pennies	2x2	0.00
Battle of the Sexes	2x2	0.94
Prisoners' Dilemma	2x2	1.00
Shapley Game	3x3	0.36
Jordan Game (α, β)	2x2	$[0.00, 0.50]$

Pure equilibria only exist in *battle of the sexes* and the *prisoners' dilemma*, where we observe the convergence of OMD. All other games are known to have only mixed equilibria, which is why (last iterate) convergence of OMD cannot be observed. Interestingly, the potentialness of the *Jordan game* (Jordan, 1993, Def 2.1) with parameters of the payoff matrices sampled uniformly at random can vary between $[0, 1/2]$.

Random Games Given a number of agents N and actions A_i for each agent i , we create a random game by sampling each entry of the payoff matrices independently from a uniform distribution, i.e., $u_i(a) \sim \mathcal{U}([0, 1])$, $\forall a \in \mathcal{A}$ and $\forall i \in \mathcal{N}$. First, we analyze the potentialness of random games by varying the number of agents and the number of actions. To that end, we sample 10^6 games for each setting and visualize the distribution (Figure 3).

We make the following observations:

Observation 1. *Increasing the size of the games (agents or actions) reduces the variance and the mean of the observed levels of potentialness.*

To analyze the connection between the potentialness and the behavior of learning dynamics, a natural next step is to look at the existence of strict pure Nash equilibria (SPNE). There are two reasons why the existence of SPNE is interesting. First, we know from Vlatakis-Gkaragkounis et al. (2020) that only SPNE can be asymptotically stable under FTRL algorithms (such as Algorithm 1). And second, the result on the existence of pure Nash equilibria in potential games (Monderer & Shapley, 1996) indicates a connection between potentialness and the existence. Note that we focus on SPNE, since weak pure Nash equilibria occur with probability zero in these randomly generated games.

Rinott & Scarsini (2000) observed that the probability that a pure Nash equilibrium exists in randomly drawn games goes to $1 - 1/e$ as the number of players gets large or as the number of actions per player gets large. So both limits converge to the same number. As expected, with increasing numbers of agents and actions, we can see in Figure 4, that the probability of the existence of SPNE converges to $1 - 1/e$.

To analyze the connection between the existence of SPNE and the potentialness of a game of fixed size, we consider the same sampled games from the previous analysis. For each setting, we group games with a similar potentialness, i.e., in the same subinterval $I_k := (\frac{k-1}{20}, \frac{k}{20}]$ for $k = 1, \dots, 20$, and compute the fraction




Figure 3: Distribution of potentialness for random games.

We visualize the empirical distribution of the potentialness for randomly generated games with different numbers of agents and different numbers of actions per agent.




Figure 4: Probability of Existence of SPNE.

We visualize the empirical probability that an SPNE exists on the y-axis and observed potentialness on the x-axis for different sizes of randomly generated games. The lines indicate which levels of potentialness are observed, i.e., the support of the densities visualized in Figure 3. The average potentialness in each settings is represented by the dot on the lines.

of games with at least one SPNE for each group. The results are visualized in Figure 5 and lead to our next observation:

Observation 2. *Within a setting (fixed number of actions and agents), the higher the potentialness, the more likely the game has at least one SPNE.*

Since the existence of SPNE only gives us local convergence of our algorithm (Mertikopoulos & Zhou, 2019), we want to understand if a higher potentialness of a game not only increases the probability of having an SPNE, but also has an influence on the basin of attraction for these equilibria.

Analogous to the previous analyses, for fixed number of agents and actions, we group the games into subgroups of similar potentialness, i.e., potentialness is in I_k , and consider 100 games from each of these groups (group is ignored if it contains less than 100 games). We then apply OMD with the step-size $\eta_t = \eta_0 \cdot t^{-\beta}$




Figure 5: Relationship between potentialness and existence of SPNE.

We visualize the empirical probability of the existence of an SPNE (y-axis) in relation to the potentialness (x-axis) of a game. The analysis includes random games with different numbers of agents and different numbers of actions per agent.

with $\eta_0 = 2^3$ and $\beta = \frac{1}{20}$. We visualize the fraction of games (with SPNE) in each group, where OMD converged with at least one step-size sequence (Figure 6) from a fixed starting point (uniform strategy, i.e., $s_{i,0} = \frac{1}{A_i}\mathbf{1}$ for all agents i).




Figure 6: Convergence of OMD in random games with SPNE.

We visualize the empirical probability of convergence of OMD (cf. Algorithm 1) in relation to the potentialness. The analysis includes random games with at least one SPNE for different numbers of agents and different numbers of actions per agent.

Observation 3. *If a game has a SPNE, the higher the potentialness, the more likely we are to end up in equilibrium using OMD.*

We did the same experiments with several different initial strategies (sampled uniformly from probability simplex) for each game and observed a similar outcome (see supplementary material). The results are also robust w.r.t. different, smaller step-sizes, but yield the best results (higher convergence rates) for larger ones.

Economic Games As compared to random matrix games, many games analyzed in the economic sciences have more structure in the payoffs. Arguably, one of the most important classes of economic games are auctions and contests, which are widely used to describe strategic interaction in markets (Krishna, 2009). In these games, the utility functions of agents have a specific form and the allocation rule is monotonic in the bids. Well-known auctions and contests include:

- First-Price Sealed-Bid (FPSB) Auction

$$u_i(a, v_i) = x_i(a) \cdot (v_i - a_i) \quad (11)$$

- Second-Price Sealed-Bid (SPSB) Auction

$$u_i(a, v_i) = x_i(a) \cdot (v_i - \max_{j \neq i} a_j) \quad (12)$$

- All-Pay Auction

$$u_i(a, v_i) = x_i(a) \cdot v_i - a_i \quad (13)$$

- War of Attrition (WoA)

$$u_i(a, v_i) = x_i(a) \cdot (v_i - \max_{j \neq i} a_j) - (1 - x_i(a)) \cdot a_i \quad (14)$$

- Tullock Contest

$$u_i(a, v_i) = \begin{cases} v \cdot \frac{a_i}{\sum_j a_j} - a_i & \text{if } \sum_j a_j > 0 \\ \frac{v}{n} & \text{else} \end{cases} \quad (15)$$

Note that $a \in \mathcal{A}$ denotes the action profile, v the value, and $x_i(a) : \mathcal{A} \rightarrow [0, 1]$ the allocation function with random tie-breaking rule which is defined by

$$x_i(a) = \begin{cases} \frac{1}{n_{\max}} & \text{if } a_i = \max_{j \neq i} a_j \\ 0 & \text{if } a_i < \max_{j \neq i} a_j \end{cases} \quad (16)$$

where n_{\max} denotes the number of bids that attain the maximum. While the games are generally defined on continuous action spaces, the decomposition only works on finite games. Note that in practice also auctions are discrete as bids can only be submitted up to a certain number of trailing digits. Therefore, we discretize the action space \mathcal{A} to A_i equidistant points for all $i \in \mathcal{I}$, e.g., $\mathcal{A}_i = \{0.0, \dots, 0.95\}$.

First, we compute the potentialness of the game for different discretizations of the action space and different symmetry assumptions on the agents, i.e., different valuations. Especially for higher discretizations we observe that the potentialness of the games does not change much, even though the number of pure NE might change (e.g., the symmetric FPSB has one SPNE for $A_i = 21$ actions, but one WPNE and one SPNE for $A_i = 20$).

Observation 4. *The potentialness of the discretized games is a property of the underlying (continuous) game, and not of the specific discretization.*

Looking closer at the decomposition of the first- and second-price auctions, we observe that their harmonic part coincides; if we vary the payment rule by considering convex combinations of the first- and second-price rule, only the potential and non-strategic parts of the decomposition change. This observation supports the intuition that the allocation rule determines the strategic difficulty posed by the game. In contrast, contests have a smoothed version of this allocation and show a higher level of potentialness.

As potentialness describes the relative weight of the potential compared to the harmonic component in a game, We can now increase or decrease the weight of the potential part by building a game $u_\alpha := \alpha u_{\mathcal{P}} + (1 - \alpha)u_{\mathcal{H}}$. In Figure 8, we can see the convergence behavior of online mirror descent in games constructed by convex combinations of the harmonic and potential parts.




Figure 7: Potentialness of economic games.

We consider economic games with two agents, different discretizations, i.e., number of actions A_i , and different valuations. The solid lines show the potentialness in the symmetric setting, where both agents have values $v_1 = v_2 = 1$, while the dashed line shows the asymmetric settings with $v_1 = \frac{3}{4}$, $v_2 = 1$.




Figure 8: Convergence of OMD in economic games.

Each square shows a convex combination of the potential and harmonic part of the respective economic game with a level of potentialness that can be seen on the x-axis. The color of the square shows the empirical frequency of convergence to some equilibrium using OMD ($\eta = 2^8, \beta = \frac{1}{20}$) for 10^2 random initial strategies. The stars indicate the potentialness of the original economic game.

We observe that for each game, there is a threshold for potentialness at which the game starts to converge. This threshold coincides with the existence of SPNE. If the potentialness exceeds a certain value, an SPNE exists, and the learning algorithm converges. Below that value, as we can see for the all-pay auction, no pure equilibrium exists and online mirror descent does not converge.

Interestingly, while the all-pay auction in the complete-information setting only has a mixed equilibrium, its incomplete-information counterpart can admit pure Bayes-Nash equilibria (BNE). There is also literature that suggests that these BNE are learnable in discrete games using first-order methods such as online mirror descent (Bichler et al., 2023). To better understand these differences, we want to extend our analysis to the more complex Bayesian framework and compare these frameworks through the lens of potentialness.

Bayesian Economic Games Bayesian (or incomplete-information) games are the standard approach in auction theory to model the economic games we considered (Krishna, 2009). Compared to normal-form games, Bayesian games $B = (\mathcal{N}, \mathcal{V}, \mathcal{A}, F, u)$ are characterized by an additional type space \mathcal{V} and a known prior distribution F over this type space. Each player $i \in \mathcal{N}$ observes a private type, which is drawn from \mathcal{V}_i

according to the prior distribution F . In our examples, the private type is the valuation v_i . After observing their values, the agents play an action, i.e., submit a bid $a \in \mathcal{A}_i$, and receive their utility $u_i(a_i, a_{-i}, v_i)$, which depends on their type. A pure strategy in such a Bayesian game is a function $\beta_i : \mathcal{V}_i \rightarrow \mathcal{A}_i$ that maps players' types to actions. Given a strategy profile, one can compute the expected (ex-ante) utility $\tilde{u}_i = \mathbb{E}_{v \sim F}[u_i(\beta_i(v_i), \beta_{-i}(v_{-i}), v_i)]$. A Bayes-Nash equilibrium (BNE) is a strategy profile β , where no player can increase the expected utility \tilde{u} by deviating.

Similar to the complete-information economic games, we only consider finite action spaces and also assume finite type spaces. This allows us to reformulate the Bayesian game as a finite normal-form game and to analyze the potentialness. The constructed normal-form game consists of actions that correspond to all possible strategies. If an agent has V_i types and A_i actions in the Bayesian game, we have $A_i^{V_i}$ possible strategies. To reduce the size of the constructed normal-form game, we restrict ourselves to non-decreasing strategies, as is usually done in auction theory. This gives us $\binom{V_i + A_i - 1}{A_i - 1}$ strategies, which allows us to analyze Bayesian games with two players and up to four actions and types (corresponds to 35 non-decreasing strategies). The payoffs of the normal-form game correspond to the expected utility in the Bayesian game. The pure NE in this normal-form game corresponds to BNE in the Bayesian game.

Using this procedure, we constructed the normal-form games from Bayesian economic games with four actions, up to four types, and uniform prior distribution. In Figure 9, we can see that the Bayesian version of each game has a higher potentialness than the complete-information setting (number of types = 1).




Figure 9: Potentialness of Bayesian economic games.

We consider the Bayesian version of economic games with two agents, $A_i = 4$ actions, $\mathcal{A}_i = \{0.0, 0.3, 0.6, 0.9\}$ and different number of types. The types are given by $\mathcal{V}_i = \{\frac{i}{V_i} : i = 1, \dots, V_i\}$ where V_i is the number of types. The types are independent and uniformly distributed.

In particular, for the all-pay auction, the increase in potentialness is substantial and for two or four types, it even has a pure NE, i.e., a pure strategy BNE. This is interesting because learning algorithms do converge to equilibrium in these Bayesian games, but not in the complete-information version of the all-pay auction that only exhibits mixed Nash equilibria.

Observation 5. *Mirror ascent does not converge to the mixed Nash equilibrium in the all-pay auction, but it does converge to pure Bayes-Nash equilibrium in the Bayesian version of the game. Potentialness only increases, but never decreases with additional types in the Bayesian game versions.*

6 Conclusions

Understanding when learning algorithms converge to a Nash equilibrium has long been analyzed in game theory. While it is well-known that exact potential games allow for convergence, we show that learning may also converge in games that are not exact potential games. Characterizing which games converge and which

do not, is an open problem in the literature on learning in games. Game decompositions allow us to measure how “potential” a specific normal-form game is. For random matrix games, we find that with a larger number of actions or players the potentialness decreases and concentrates. Importantly, potentialness proves to be a good predictor for the existence of a strict Nash equilibrium and for convergence of online mirror descent. For example, we found that random games with a potentialness greater than 0.8 have a very high probability of having an SPNE to which OMD converges to. Economically motivated games provide more structure in the payoffs, and potentialness is also a very good predictor for convergence in these games. In contrast to individual simulation runs where convergence depends on the initialization of the algorithm, potentialness serves as a predictor for average convergence. So, rather than running many simulations from any possible initial conditions to study convergence, potentialness provides a useful predictor for convergence, addressing a long-standing question in the literature on learning in games. It will be crucial to further solidify this finding by considering other no-regret learning dynamics and larger classes of games.

References

- Dilip Abreu and Ariel Rubinstein. The structure of nash equilibrium in repeated games with finite automata. *Econometrica*, 56(6):1259–1281, 1988.
- Ioannis Anagnostides, Ioannis Panageas, Gabriele Farina, and Tuomas Sandholm. On last-iterate convergence beyond zero-sum games. In *International Conference on Machine Learning*, pp. 536–581. PMLR, 2022.
- Gabriel P Andrade, Rafael Frongillo, and Georgios Piliouras. Learning in matrix games can be arbitrarily complex. In *Conference on Learning Theory*, pp. 159–185. PMLR, 2021.
- Martin Bichler, Maximilian Fichtl, and Matthias Oberlechner. Computing Bayes-Nash equilibrium in auction games via gradient dynamics. *Operations Research*, 2023. doi: 10.1287/opre.2022.0287.
- Oliver Biggar and Iman Shames. The graph structure of two-player games. *Scientific Reports*, 13(1):1833, February 2023. ISSN 2045-2322. doi: 10.1038/s41598-023-28627-8.
- Avrim Blum and Adam Tauman Kalai. Universal portfolios with and without transaction costs. *Machine Learning*, 35(3):193–205, 1999.
- George W Brown. Iterative solution of games by fictitious play. *Activity Analysis of Production and Allocation*, 13(1):374–376, 1951.
- Ozan Candogan, Ishai Menache, Asuman Ozdaglar, and Pablo A Parrilo. Flows and decompositions of games: Harmonic and potential games. *Mathematics of Operations Research*, 36(3):474–503, 2011.
- Ozan Candogan, Asuman Ozdaglar, and Pablo A Parrilo. Dynamics in near-potential games. *Games and Economic Behavior*, 82:66–90, 2013a.
- Ozan Candogan, Asuman Ozdaglar, and Pablo A Parrilo. Near-potential games: Geometry and dynamics. *ACM Transactions on Economics and Computation (TEAC)*, 1(2):1–32, 2013b.
- Le Chen, Alan Mislove, and Christo Wilson. An empirical analysis of algorithmic pricing on amazon marketplace. In *Proceedings of the 25th international conference on World Wide Web*, pp. 1339–1349, 2016.
- George Christodoulou, Vahab S Mirrokni, and Anastasios Sidiropoulos. Convergence and approximation in potential games. *Theoretical Computer Science*, 438:13–27, 2012.
- Constantinos Daskalakis, Rafael Frongillo, Christos H Papadimitriou, George Pierrakos, and Gregory Valiant. On learning algorithms for Nash equilibria. In *International Symposium on Algorithmic Game Theory*, pp. 114–125. Springer, 2010.
- Jozef Dodziuk. Finite-Difference Approach to the Hodge Theory of Harmonic Forms. *American Journal of Mathematics*, 98(1):79–104, 1976. ISSN 0002-9327. doi: 10.2307/2373615.
- Beno Eckmann. Harmonische Funktionen und Randwertaufgaben in einem Komplex. *Commentarii Mathematici Helvetici*, 17(1):240–255, December 1944. ISSN 1420-8946. doi: 10.1007/BF02566245.

- Lampros Flokas, Emmanouil Vasileios Vlatakis-Gkaragkounis, Thanasis Lianas, Panayotis Mertikopoulos, and Georgios Piliouras. No-regret learning and mixed Nash equilibria: They do not mix. In *NeurIPS '20: Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020.
- Dean P Foster and Rakesh V Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1-2):40, 1997.
- Joel Friedman. Computing Betti numbers via combinatorial Laplacians. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing - STOC '96*, pp. 386–391, Philadelphia, Pennsylvania, United States, 1996. ACM Press. ISBN 978-0-89791-785-8. doi: 10.1145/237814.237985.
- Drew Fudenberg and David K. Levine. *The Theory of Learning in Games*, volume 2 of *Economic learning and social evolution*. MIT Press, Cambridge, MA, 1998.
- Drew Fudenberg and David K. Levine. Conditional universal consistency. *Games and Economic Behavior*, 29(1):104–130, 1999.
- Angeliki Giannou, Emmanouil Vasileios Vlatakis-Gkaragkounis, and Panayotis Mertikopoulos. Survival of the strictest: Stable and unstable equilibria under regularized learning with partial information. In *COLT '21: Proceedings of the 34th Annual Conference on Learning Theory*, 2021a.
- Angeliki Giannou, Emmanouil Vasileios Vlatakis-Gkaragkounis, and Panayotis Mertikopoulos. The convergence rate of regularized learning in games: From bandits and uncertainty to optimism and beyond. In *NeurIPS '21: Proceedings of the 35th International Conference on Neural Information Processing Systems*, 2021b.
- Itzhak Gilboa and Akihiko Matsui. Social stability and equilibrium. *Econometrica*, 59(3):859–867, May 1991.
- Jonathan S. Golan. *Foundations of Linear Algebra*. Springer Science & Business Media, 1992. ISBN 978-94-015-8502-6.
- Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- Sergiu Hart and Andreu Mas-Colell. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93(5):1830–1836, 2003.
- Amélie Heliou, Johanne Cohen, and Panayotis Mertikopoulos. Learning with bandit feedback in potential games. *Advances in Neural Information Processing Systems*, 30, 2017.
- Josef Hofbauer and William H. Sandholm. Stable games and their dynamics. *Journal of Economic Theory*, 144(4):1665–1693, July 2009.
- Xiaoye Jiang, Lek-Heng Lim, Yuan Yao, and Yinyu Ye. Statistical ranking and combinatorial Hodge theory. *Mathematical Programming*, 127(1):203–244, 2011.
- Jakob Jonsson. *Simplicial Complexes of Graphs*. Springer Science & Business Media, 2007.
- James S. Jordan. Three problems in learning mixed strategy Nash equilibria. *Games and Economic Behavior*, 5(3):368–386, July 1993.
- Vijay Krishna. *Auction theory*. Academic press, 2009.
- Davide Legacci, Panayotis Mertikopoulos, and Bary S. R. Pradelski. A geometric decomposition of finite games: Convergence vs. recurrence under exponential weights. In *ICML '24: Proceedings of the 41st International Conference on Machine Learning*, 2024.
- Panayotis Mertikopoulos and Zhengyuan Zhou. Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming*, 173(1-2):465–507, January 2019.

- Panayotis Mertikopoulos, Christos H. Papadimitriou, and Georgios Piliouras. Cycles in adversarial regularized learning. In *SODA '18: Proceedings of the 29th annual ACM-SIAM Symposium on Discrete Algorithms*, 2018.
- Panayotis Mertikopoulos, Ya-Ping Hsieh, and Volkan Cevher. A unified stochastic approximation framework for learning in games. *Mathematical Programming*, 203:559–609, January 2024.
- Dov Monderer and Lloyd S Shapley. Potential games. *Games and economic behavior*, 14(1):124–143, 1996.
- James R Munkres. *Elements of Algebraic Topology*. Perseus Books, 1984.
- Arkadij Semenovič Nemirovskij and David Borisovich Yudin. Problem complexity and method efficiency in optimization. 1983.
- Noam Nisan (ed.). *Algorithmic Game Theory*. Cambridge University Press, Cambridge ; New York, 2007. ISBN 978-0-521-87282-9. OCLC: ocn122526907.
- Gerasimos Palaiopanos, Ioannis Panageas, and Georgios Piliouras. Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. In *NIPS '17: Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017.
- Yosef Rinott and Marco Scarsini. On the number of pure strategy nash equilibria in random games. *Games and Economic Behavior*, 33(2):274–293, 2000.
- James BT Sanders, J Doyne Farmer, and Tobias Galla. The prevalence of chaotic dynamics in games with many players. *Scientific Reports*, 8(1):1–13, 2018.
- Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012. ISSN 1935-8237. doi: 10.1561/22000000018.
- Yannick Viossat and Andriy Zapechelnyuk. No-regret dynamics and fictitious play. *Journal of Economic Theory*, 148(2):825–842, 2013.
- Emmanouil-Vasileios Vlatakis-Gkaragkounis, Lampros Flokas, Thanasis Lianas, Panayotis Mertikopoulos, and Georgios Piliouras. No-regret learning and mixed nash equilibria: They do not mix. *Advances in Neural Information Processing Systems*, 33:1380–1391, 2020.
- Yaodong Yang and Jun Wang. An overview of multi-agent reinforcement learning from game theoretical perspective. *arXiv:2011.00583*, 2020.
- H Peyton Young. *Strategic learning and its limits*. OUP Oxford, 2004.

A Appendix

A.1 Decomposition in the space of payoffs

The reason two work in the space C_1 of flows rather than in the space \mathcal{U} of payoffs to achieve a decomposition of games is twofold: first, the combinatorial Hodge theorem holds true on the space C_1 of flows (and in general on every chain group C_k of a simplicial complex), while there is no such a theorem for the space \mathcal{U} of payoffs; second, two games $u \neq u'$ such that $Du = Du'$, albeit having different payoffs, display the same strategical properties and are effectively the “same” game⁷, so looking at the deviation flow rather than at the payoff one gets rid of a redundancy that is intrinsic in the formulation of a game.

Non-strategic games To make this last point more precise Candogan et al. (2011) introduce the notion of *non-strategic games*.

Definition 8. A finite game in normal form $G = (\mathcal{N}, \mathcal{A}, u)$ is called *non-strategic* if it has zero deviation flow:

$$Du = 0 \quad (17)$$

The space of non-strategic games is denote by

$$\mathcal{K} := \ker D \quad (18)$$

In a non-strategic game, all players are indifferent among all of their choices:

Proposition 3 (Candogan et al. 2011). The game $G = (\mathcal{N}, \mathcal{A}, u)$ is non-strategic if and only if

$$u_i(a'_i, a_{-i}) = u_i(a''_i, a_{-i}) \quad (19)$$

for all $i \in \mathcal{N}$, all $a_{-i} \in \mathcal{A}_{-i}$, and all $a'_i, a''_i \in \mathcal{A}_i$.

Since $D : \mathcal{U} \rightarrow C_1$ is a linear map between vector spaces, the space of non-strategic games is a linear subspace $\mathcal{K} \subset \mathcal{U}$. It follows by the definition of non-strategic games that two games have the same deviation flow if and only if their difference is a non-strategic game, and in this case we say that the two games are *strategically equivalent*.

Normalized games Being strategically equivalent is an equivalence relation on the space \mathcal{U} of payoffs; one can select a representative element in each equivalence class $[u]$ by choosing a complement of \mathcal{K} in \mathcal{U} and projecting $u \in \mathcal{U}$ onto such complement along \mathcal{K} . A natural choice is that of using the *orthogonal* complement \mathcal{K}^\perp of the space of non-strategic games with respect to the Euclidean inner-product in \mathcal{U} ; we refer to such procedure as *normalization*, and following Candogan et al. (2011) we give the following definition:

Definition 9. A finite game in normal form $G = (\mathcal{N}, \mathcal{A}, u)$ is called *normalized* if

$$u \in \mathcal{K}^\perp \quad (20)$$

Normalized games enjoy the “no-leftover” property: the sum of any player’s payoffs over their choices is zero for any fixed choice by the other players.

Proposition 4 (Candogan et al. 2011). The game $G = (\mathcal{N}, \mathcal{A}, u)$ is normalized if and only if

$$\sum_{a'_i \in \mathcal{A}_i} u(a'_i, a_{-i}) = 0 \quad (21)$$

for all $i \in \mathcal{N}$ and all $a_{-i} \in \mathcal{A}_{-i}$.

⁷Quoting Candogan et al. (2013b), if [two games have the same deviation flow], then the equilibrium sets of these games are identical. However, payoffs at equilibria may differ, and hence they may be different in terms of their efficiency (such as Pareto efficiency) properties (see Candogan et al. (2011)).

Decomposition in the space of payoffs After normalizing the space \mathcal{U} of payoffs⁸ one can translate the decomposition of feasible flows (Theorem in the main text) from $\text{Im } D \subset C_1$ to \mathcal{U} by means of the Moore-Penrose pseudo-inverse $\tilde{D} : C_1 \rightarrow \mathcal{U}$ of the deviation map.

Recall from Section 4 that the space of potential games is $D^{-1} \text{Im } d_0 \subset \mathcal{U}$, and that the space of harmonic games is $D^{-1} \ker \partial_1 \subset \mathcal{U}$. Their intersections with the space $\mathcal{K}^\perp \subset \mathcal{U}$ of normalized games give the spaces of *normalized potential games* and of *normalized harmonic games*:

Definition 10. *The space of normalized potential games is the linear subspace*

$$\mathcal{P} := (\text{potential games}) \cap \mathcal{K}^\perp \subset \mathcal{U} \quad (22)$$

The space of normalized harmonic games is the linear subspace

$$\mathcal{H} := (\text{harmonic games}) \cap \mathcal{K}^\perp \subset \mathcal{U} \quad (23)$$

Theorem (Candogan et al. (2011) — Combinatorial Hodge decomposition of finite normal form games). *The space \mathcal{U} is the direct sum of the subspaces of normalized potential games, normalized harmonic games, and non-strategic games:*

$$\mathcal{U} = \mathcal{P} \oplus \mathcal{H} \oplus \mathcal{K} \quad (24)$$

Equivalently, given a finite normal form game $G = (\mathcal{N}, \mathcal{A}, u)$ the payoff function u can be uniquely decomposed as the sum $u = u_{\mathcal{P}} + u_{\mathcal{H}} + u_{\mathcal{K}}$ of a normalized potential game $u_{\mathcal{P}}$, a normalized harmonic game $u_{\mathcal{H}}$, and a non-strategic game $u_{\mathcal{K}}$.

Decomposition components Recall that the deviation map is a linear map $D : \mathcal{U} \rightarrow C_1$ from the space of payoffs to the space of flows. By the properties of the Moore-Penrose pseudo-inverse $\tilde{D} : C_1 \rightarrow \mathcal{U}$ of the deviation map (Golan, 1992), the operator $\Pi := \tilde{D}D : \mathcal{U} \rightarrow \mathcal{U}$ is the orthogonal projection onto $\mathcal{K}^\perp = (\ker D)^\perp$. Recall furthermore that $e : C_1 \rightarrow C_1$ is the orthogonal projection onto the subspace of potential flows.

These operator can be used to obtained explicit expressions for the components of the decomposition in the space of payoffs:

Proposition 5 (Candogan et al. (2011)). *Given the finite normal form game $G = (\mathcal{N}, \mathcal{A}, u)$ the components $u_{\mathcal{P}}, u_{\mathcal{H}}$ and $u_{\mathcal{K}}$ of the combinatorial Hodge decomposition of finite normal form games are given by*

- $u_{\mathcal{K}} = u - \Pi u \in \mathcal{K}$
- $u_{\mathcal{P}} = \tilde{D}eDu \in \mathcal{P}$
- $u_{\mathcal{H}} = u - u_{\mathcal{K}} - u_{\mathcal{P}} \in \mathcal{H}$

A.2 Additional Numerical Experiments

Similar to the experiment visualized in Figure 5, we considered games with a similar potentialness and applied OMD with $(\eta_0 = 2^8, \beta = \frac{1}{2})$. But this time, we randomly sampled 25 initial strategies for each game. The first plot in Figure 10 shows the estimated convergence probability over different games with a similar potentialness and different initial points. On the second plot, we consider the same setting (only restricted to 2 actions), and compare the convergence probability on the level of the different games. To this end, we visualize the mean \pm std (colored area) of the convergence probability (for different initialization points) over similar games.

We observe that even though games have a similar potentialness, the probability of convergence, i.e., the basin of attraction of the SPNE, can differ a lot from game to game. On the other hand, we can still observe a clear connection between the potentialness and convergence in expectation.

⁸That is, after quotienting away the kernel of the deviation map.

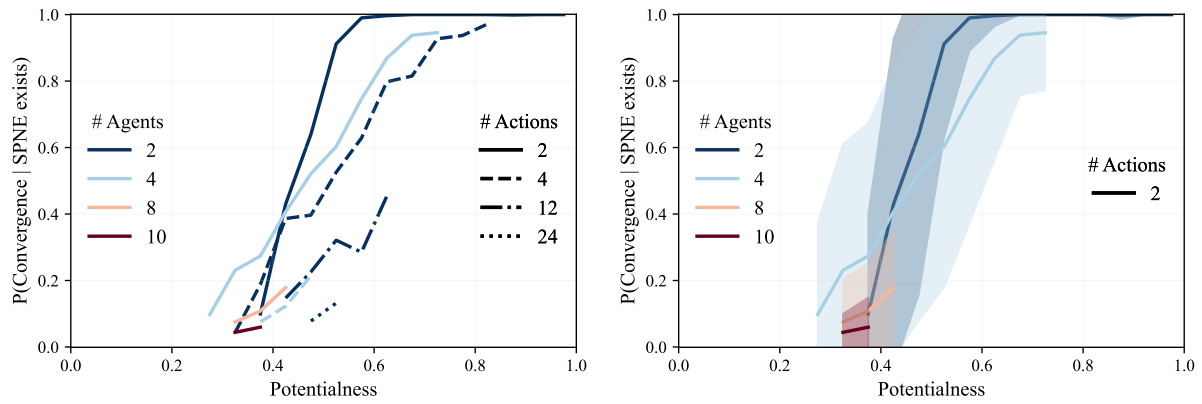


Figure 10: Convergence of OMD in random games with SPNE and with different levels of potentialness and different initialization points.