Exponential Lower Bounds for Fictitious Play in Potential Games

Anonymous Author(s)
Affiliation
Address
email

Abstract

Fictitious Play (FP) is a simple and natural dynamic for repeated play with many applications in game theory and multi-agent reinforcement learning. It was introduced by Brown [3, 4] and its convergence properties for two-player zero-sum games was established later by Robinson [15]. Potential games [12] is another class of games which exhibit the FP property [11], i.e., FP dynamics converges to a Nash equilibrium if all agents follows it. Nevertheless, except for two-player zero-sum games and for specific instances of payoff matrices [1] or for adversarial tie-breaking rules [8], the *convergence rate* of FP is unknown. In this work, we focus on the rate of convergence of FP when applied to potential games and more specifically identical payoff games. We prove that FP can take exponential time (in the number of strategies) to reach a Nash equilibrium, even if the game is restricted to two agents and for arbitrary tie-breaking rules. To prove this, we recursively construct a two-player coordination game with a unique Nash equilibrium. Moreover, every approximate Nash equilibrium in the constructed game must be close to the pure Nash equilibrium in ℓ_1-distance.

1 Introduction

In 1949 Brown proposed an uncoupled dynamics called *fictitious play* so as to capture the behavior of selfish agents engaged in a repeatedly-played game. Fictitious play assumes at round $t = 1$, each agent selects an arbitrary action. At each round $t \geq 2$, each player plays a best response pure action to the opponents' *empirical strategy*; empirical strategy is defined to be the empirical average of the past chosen strategies.

Due to its simplicity and natural behavioral assumptions, fictitious play is one of the most seminal and well-studied game dynamics [6]. Despite the fact that fictitious play does not converge to a Nash equilibrium (NE) in general normal-form games, there are several important classes of games at which the *empirical strategies* always converge to a NE. In her seminal work, Robinson [15] showed that in the case of *two-player zero-sum* games, the empirical strategy profiles converge to a min-max equilibrium of the game; Robinson's proof use a smart inductive argument on the number of strategies of the game. Later, Monderer and Shapley [11] established that in the case of N-player potential games, the empirical strategies also converge to a NE. Summing up, the following theorem is true:

Informal Theorem [15],[11] *In two-player zero-sum and N-player potential games, the empirical strategy profiles of fictitious play converge to a NE for any initialization and tie-breaking rule*[1].

The above convergence results are asymptotic in the sense that they do not provide guarantees on the number of rounds needed by fictitious play to reach an approximate NE. Karlin [10] conjectured

[1] Tie-breaking rule when selecting between two or more different best-response actions.

Submitted to 37th Conference on Neural Information Processing Systems (NeurIPS 2023). Do not distribute.

that in the case of two-player zero-sum games, fictitious play requires $O(1/\epsilon^2)$ rounds to reach an ϵ-approximate NE. Daskalakis et. al. [8] disproved the strong version of Karlin's conjecture by providing an adversarial tie-breaking rule for which fictitious play requires exponential number of rounds (with respect to the number of strategies) in order to converge to an ϵ-approximate NE. However there are no such lower bounds results in the case of potential games. In this paper we investigate the following question.

Q: Does fictitious play in potential games admit convergence to an approximate NE with rates that depend polynomially on the number of actions and the desired accuracy?

Our contributions Our work provides a negative answer on the above question. Specifically, we present a *two-player potential* game for which fictitious play requires super-exponential time with respect to the number of actions to reach an approximate NE.

Theorem 1.1 (Main result, formally stated Theorem 3.1)**.** *There exists a two-player potential game (more specifically both agents have identical payoffs) in which both agents admit n actions and for which fictitious play requires $\Omega\left(4^n\left(\left(\frac{n}{2}-2\right)!\right)^4+\frac{1}{n\sqrt{\epsilon}}\right)$ rounds in order to reach an ϵ-approximate NE. Moreover, the result holds for any tie-breaking rule and uniformly random initialization.*

Remark 1.2. Daskalakis et al. [8] provide an exponential lower bound on the convergence of fictitious play assuming an adversarial tie-breaking rule, meaning that ties are broken in favor of slowing down the convergence rate. To this point, it is not known whether fictitious play with a consistent tie-breaking rule (e.g. lexicographic) converges in polynomial time or not, except for special cases, e.g., diagonal payoff matrices [1]. We would like to note that our lower bound construction holds for any tie-breaking rule. The latter indicates an interesting discrepancy between two-player zero-sum and two-player potential games.

Related Work The work of Daskalakis et al. [8] provides a lower bound on the convergence rate of fictitious play for the case of two-player zero-sum games using an adversarial tie-breaking rule and is the most close to ours. Another work we must highlight is [1] in which the authors show that if the tie-breaking rule is fixed in advance (e.g., lexicographic), then fictitious play converge rate is polynomial in the number of actions/strategies and is $O\left(\frac{1}{\epsilon^2}\right)$. Other works include convergence of fictitious play for near-potential games [5], sufficient conditions games must satisfy so that fictitious play converges to a NE, using decomposition techniques [7]; the aforementioned results do not include rates for fictitious play. Fast convergence rates of *continuous time* fictitious play for *regular* potential games are established [17] (these are games in which the NE are regular according to the definition of Harsanyi). Other works on continuous time fictitious play include [14] (and references there in). Fictitious play dynamics has found various application in Multi-agent Reinforcement Learning as well (see [13, 16, 2] and [18] for a survey and references therein) and extensive form games (using deep RL) [9] to name a few[2].

Technical overview The technical overview of this paper provides a high-level roadmap of the key contributions. In Section 3, we outline the steps towards proving Theorem 3.1 by recursively constructing a payoff matrix of carefully crafted structural properties. In that matrix, starting from the lower-left element, the sequence of successive increments form a spiraling trajectory that converge towards the element of maximum value. We demonstrate that fictitious play has to follow the same trajectory, passing through all non-zero elements, to reach the unique pure Nash equilibrium, as illustrated in Figure 2. A crucial component of the proof is provided by an induction argument that *emulates* the movement of fictitious play and provides a super-exponential lower bound on the number of rounds needed.

2 Preliminaries

2.1 Notation and Definitions

Notation Let $\mathbb{R}$ be the set of real numbers, and $[n]=\{1,2,\ldots,n\}$ be the set of actions. We define Δ_n as the probability simplex, which is the set of n-dimensional probability vectors, i.e.,

[2] Due to the vastness of the literature in fictitious play, it is not possible to include all the works that have either been used or been inspired by this method.

$\Delta_n := \{x \in \mathbb{R}^n : x_i \geq 0, \sum_{i=1}^n x_i = 1\}$. We use e_i to denote the i-th elementary vector, and to refer to a coordinate of a vector, we use either x_i or $[x]_i$. The superscripts are used to indicate the time at which a vector is referring to.

Normal-form Games In a *two-player normal-form* game we are given a pair of payoff matrices $A \in \mathbb{R}^{n\times m}$ and $B \in \mathbb{R}^{n\times m}$ where n and m are the respective pure strategies of the *row* and the *column* player. If the row player selects strategy $i \in [n]$, and the column player selects strategy $j \in [m]$, then their respective payoffs are A_{ij} for the row player and B_{ij} for the column player.

The agents can use randomization. A *mixed strategy* for the row player is a probability distribution $x \in \Delta_n$ over the n rows, and a *mixed strategy* for the column player is a probability distribution $y \in \Delta_m$ over the m columns. After the row player selects mixed strategy $x \in \Delta_n$, and the column player selects mixed strategy $y \in \Delta_m$, their expected payoffs are $x^\top Ay$ and $x^\top By$, respectively.

Potential Games A two-player potential game is a class of games that admit a unique function Φ, referred to as a potential function, which captures the incentive of all players to modify their strategies. In other words, if a player deviates from their strategy, then the difference in payoffs is determined by a potential function Φ evaluated at those two strategy profiles. We can express this formally as follows:

Definition 2.1 (Two-player Potential Game)**.** For any given pair of strategies (x, y) and a pair of unilateral deviations x' by the row player and y' by the column player, the difference in their utility is equivalent to the difference in the potential function.

$$(x')^\top Ay - x^\top Ay = \Phi(x', y) - \Phi(x, y) \quad \text{and} \quad x^\top A(y') - x^\top Ay = \Phi(x, y') - \Phi(x, y)$$

Remark 2.2. We note Φ is a function that characterizes the equilibria of the game as the strategy profiles that maximize the potential function.

In this work, we focus on a specific type of potential games called *identical payoff games*, where both players receive the same payoff.

Definition 2.3 (Identical Payoff Games)**.** A two-player normal-form game (A, B) is called *identical payoff* if and only if $A = B$.

In this scenario, it is apparent that the potential function is given by $\Phi(x, y) = x^\top Ay$. Finally we provide the definition of an approximate NE.

Definition 2.4 (ϵ-Nash Equilibrium)**.** A strategy profile $(x^\star, y^\star) \in \Delta_n \times \Delta_m$ is called an ϵ-approximate NE if and only if

$$(x^\star)^\top Ay^\star \geq x^\top Ay^\star - \epsilon \;\; \forall x \in \Delta_n \quad \text{and} \quad (x^\star)^\top By^\star \geq (x^\star)^\top By - \epsilon \;\; \forall y \in \Delta_m$$

In words, an approximate Nash equilibrium is a strategy profile in which no player can improve their payoff significantly by unilaterally changing their strategy, but the strategy profile may not necessarily satisfy the precise definition of a Nash equilibrium.

Remark 2.5. We highlight two special cases of the ϵ-NE. Firstly, when ϵ is equal to zero, it is referred to as an *exact Nash equilibrium*. Secondly, when the support of strategies is of size 1, it is called a *pure Nash equilibrium*. It is worth noting that potential games always admit a pure NE.

2.2 Fictitious Play

Fictitious play is a natural uncoupled game dynamics at which each agent chooses a *best response* to their opponent's empirical mixed strategy. Since there might be several best response actions at a given round, fictitious play might contain different sequences of play; see Definition 2.6.

Definition 2.6 (Fictitious Play)**.** An infinite sequence of pure strategy profiles $(i^{(1)}, j^{(1)}), \ldots, (i^{(t)}, j^{(t)}), \ldots$ is called a fictitious play sequence if and only if at each round $t \geq 2$,

$$i^{(t)} \in \underset{i\in[n]}{\text{argmax}} \sum_{s=1}^{t-1} A_{ij^{(s)}} \quad \text{and} \quad j^{(t)} \in \underset{j\in[m]}{\text{argmax}} \sum_{s=1}^{t-1} B_{i^{(s)}j} \tag{1}$$

The empirical strategy profile of row and column player at time T is defined as $\hat{x}^{(T)} = \left(\frac{1}{T}\sum_{s=1}^{T} e_{i^{(s)}}\right)$ and $\hat{y}^{(T)} = \left(\frac{1}{T}\sum_{s=1}^{T} e_{j^{(s)}}\right)$ where $e_{i^{(t)}}, e_{j^{(t)}}$ are the elementary basis vectors.

Definition 2.7 (Cumulative utility vector)**.** For an infinite sequence of pure strategy profiles $(i^{(1)}, j^{(1)}), \ldots, (i^{(t)}, j^{(t)}), \ldots$, the *cumulative utility* vectors of the row and column player at round $t \geq 1$ are defined as,

$$R^{(t)} = \sum_{s=1}^{t-1} A e_{j^{(s)}} \quad \text{and} \quad C^{(t)} = \sum_{s=1}^{t-1} e_{i^{(s)}}^{\top} B.$$

Remark 2.8. Fictitious play assumes that each agent selects at each round $t \in [T]$ a strategy with *maximum cumulative utility*. The latter decision-making algorithm is also known as *Follow the Leader*. We remark that the latter alternative interpretation provides a direct generalization of fictitious play in N-player games.

In their seminal work, Monderer et al.[11] established that in case of identical payoff games the empirical strategies of any fictitious play sequence converges asymptotically to a NE.

Theorem 2.9 ([11])**.** *Let a fictitious play sequence $(i^{(1)}, j^{(1)}), \ldots, (i^{(t)}, j^{(t)}), \ldots$ for an identical payoff game described with matrix A. Then, there exists a round $T^\star \geq 1$ such that for any $t \geq T^\star$, the empirical strategy profile $(\hat{x}^{(t)}, \hat{y}^{(t)})$ converges to a NE with a rate of $1/t$.*

On the positive side, Theorem 2.9 establishes that *any fictitious play sequence* converges to a Nash equilibrium in the case of potential games [3] On the negative side, Theorem 2.9 does not provide any convergence rates, since the round $T^\star$ depends on the specific fictitious play sequence and its dependence on the number of strategies is rather unclear.

3 Main Result

In this section, we outline the steps towards proving Theorem 3.1, as it is stated below. Firstly, we introduce a carefully constructed payoff matrix A of size $n \times n$ and analyze its structural properties in Section 3.1. Next, in Section 3.2, we investigate the behavior of fictitious play when the game is an two-player identical payoff game with this matrix A. We also present a set of key statements that are necessary for proving the main theorem. Finally, in Section 3.3, we provide a proof for the fundamental Lemma 3.8.

Theorem 3.1. *Let an identical payoff game defined with the matrix A of size $n \times n$ and consider any fictitious play sequence $(i^{(1)}, j^{(1)}), \ldots, (i^{(t)}, j^{(t)}), \ldots$ with $(i^{(1)}, j^{(1)}) = (n, 1)$. In case the empirical strategy profile $(\hat{x}^{(T)}, \hat{y}^{(T)})$ is an ϵ-approximate Nash equilibrium then it holds*

$$T \geq \Omega\left(4^n((n/2-2)!)^4 + \frac{1}{n\sqrt{\epsilon}}\right).$$

Moreover, the lower bound on T is independent of the tie-breaking rule. Finally, if the initialization is chosen uniformly at random, then the expected number of rounds to reach an ϵ-approximate Nash equilibrium is $\Omega\left(\frac{4^n((n/2-2)!)^4+1/n\sqrt{\epsilon}}{n^2}\right)$.

3.1 Construction and Analysis of the Payoff Matrix A

We begin by introducing our recursive construction for the payoff matrix, which we use to establish the formal statement of Theorem 3.1.

Definition 3.2. For any $z > 0$ and n even, consider the following $n \times n$ matrix $K^n(z)$.

1. For $n = 2$, $K^n(z) = \begin{pmatrix} z+2 & z+3 \\ z+1 & 0 \end{pmatrix}$.
2. For $n \geq 4$,
   - $K^n_{n1}(z) = z+1$ and $K^n_{nj}(z) = 0$ for $j \notin \{1\}$. *Row n*
   - $K^n_{n1}(z) = z+1, K^n_{11}(z) = z+2$ and $K^n_{1j}(z) = 0$ for any $j \notin \{1, n\}$. *Column 1*

[3] For the sake of exposition, we have stated Theorem 2.9 only for the case of *identical payoff games*. However, we remark that the same theorem holds for general N-player potential games.

$$K^n(z) = \begin{pmatrix} (z+2) & 0 \cdots 0 & (z+3) \\ 0 & & 0 \\ \vdots & K^{n-2}(z+4) & \vdots \\ & & 0 \\ 0 & & (z+4) \\ (z+1) & 0 \cdots 0 & 0 \end{pmatrix}$$

(a) Recursive construction of A.

$$K^6(0) = \begin{pmatrix} 2 & 0 & 0 & 0 & 0 & 3 \\ 0 & 6 & 0 & 0 & 7 & 0 \\ 0 & 0 & 10 & 11 & 0 & 0 \\ 0 & 0 & 9 & 0 & 8 & 0 \\ 0 & 5 & 0 & 0 & 0 & 4 \\ 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

(b) An example for $z = 0$ and $n = 6$.

- $K^n_{11}(z) = z + 2$, $K^n_{1n}(z) = z + 3$ and $K^n_{1j}(z) = 0$ for $j \notin \{1, n\}$. *Row* 1
- $K^n_{1n}(z) = z + 3$, $K^n_{n-1n}(z) = z + 4$ and $K^n_{nj}(z) = 0$ for $j \notin \{1, n-1\}$. *Column* n
- For all $i, j \in \{2, \ldots, n-1\} \times \{2, \ldots, n-1\}$, $K^n_{ij}(z) := K^{n-2}_{i-1j-1}(z + 4)$.

In Figures 1a and 1b, we provide a schematic representation of Definition 3.2 and an illustrative example for $n = 6$ and $z = 0$. For the sake of simplicity, we have intentionally omitted the remaining zeros in the outer rows, and columns of the matrix.

The construction of the payoff matrix exhibits an interesting circular pattern, which begins at the lower left corner and extends along the outer layer of the matrix. More specifically, the first increment occurs in the same column as the starting point, i.e., at position $(1, 1)$ on the first row. The pattern then proceeds to the next greater element on the same row but a different column, i.e., at position $(1, n)$, and the last increment before entering the inner sub-matrix is located on the same column but on the $(n-1)$-th row.

$$\text{Sequence of increments: } (n, 1) \to (1, 1) \to (1, n) \to (n-1, n) \to \underbrace{(n-1, 1) \to \cdots}_{K^{n-2}(z+4)}$$

The increments have been carefully selected to ensure that there are alternating changes in row and column when starting from the lower-left corner and following the successive increments, until reaching the sub-matrix in the center. Once inside the sub-matrix, a similar pattern continues. As we will explore later on, the structure of the payoff dictates the behavior of fictitious play. We denote the payoff matrix under consideration as A, which is defined as $A := K_n(0)$. The subsequent statements establish the key properties of A.

Observation 3.3 (Structural Properties of matrix A)**.** Let the matrix $A = K^n(0)$, then for $i \in \{0, \ldots, \frac{n}{2} - 1\}$ the following hold:

- The only elements with non-zero values in column $i + 1$ are located at positions $i + 1$ and $n - i$, and have values $4i + 2$ and $4i + 1$, respectively.
- The only non-zero elements of row $i + 1$ are located at positions $i + 1$ and $n - i$, and have values $4i + 2$ and $4i + 3$, respectively.
- The only non-zero elements of column $n - i$ are located at positions $i + 1$ and $n - i - 1$ and have values $4i + 3$ and $4i + 4$, respectively.
- The only non-zero elements of row $n - i$ are located at positions $i + 1$ and $n - i + 1$, and have values $4i + 1$ and $4i$, respectively.

Observation 3.4. The maximum value of A is $2n - 2$ and is located at the entry $\left(\frac{n}{2}, \frac{n}{2} + 1\right)$.

Proposition 3.5. *For any non-zero element in the matrix A, there is at most one non-zero element that is greater and at most one non-zero element that is smaller in the same column or row.*

Proof. By Observation 3.3, each row and column of matrix A contains at most two non-zero elements that are necessarily different to each other. Thus, if (i, j) is a non-zero element of A, there can be at most two additional non-zero elements in row i and column j combined. Moreover, at most one of these elements can be greater and at most one can be smaller. □

One of the central components of the main theorem is presented below. Lemma 3.6 establishes that in an identical payoff game with matrix A, any approximate Nash equilibrium must distribute the majority of its probability mass to the maximum element in A, which is located at the entry $(\frac{n}{2}, \frac{n}{2}+1)$. The proof of this theorem is based solely on the structural properties presented in Observation 3.3, and is provided in the Appendix.

Lemma 3.6 (Unique ϵ^2-NE)**.** *Let $\epsilon \in O(n^3)$ and consider an ϵ^2-approximate Nash Equilibrium (x^*, y^*). Then the following hold,*

$$x^*_{\frac{n}{2}} \geq 1 - n\epsilon \quad \textit{and} \quad y^*_{\frac{n}{2}+1} \geq 1 - n\epsilon$$

Lemma 3.6 not only establishes that the only Nash equilibrium of the identical-payoff game with matrix A corresponds to the strategies $(\frac{n}{2}, \frac{n}{2}+1)$, but it also implies that this is the only exact Nash equilibrium (i.e., $\epsilon = 0$). This observation follows from Observation 3.4, which states that the entry $(\frac{n}{2}, \frac{n}{2}+1)$ corresponds to a maximum value of A and hence it is a Nash equilibrium as it dominates both the row and the column that it belongs to. We can formally state this observation as follows.

Corollary 3.7 (Unique pure NE)**.** *In an identical-payoff game with payoff matrix A, there exists a unique pure Nash equilibrium at $\left(\frac{n}{2}, \frac{n}{2}+1\right)$.*

3.2 Lower Bound for Fictitious Play in a Game with Matrix A

In this subsection, we present the proof of Theorem 3.1. To achieve this, we first prove that fictitious play requires super-exponential time before placing a positive amount of mass in entry $(\frac{n}{2}, \frac{n}{2}+1)$. This result is established by our main technical contribution of the subsection, which is Lemma 3.8.

Lemma 3.8. *Let an identical-payoff game with payoff matrix A and a fictitious play sequence $(i^{(1)}, j^{(1)}), \ldots, (i^{(t)}, j^{(t)}), \ldots$ with $(i^{(1)}, j^{(1)}) = (n, 1)$. Then, for all $\ell = \{0, \ldots, \frac{n}{2}-1\}$ there exists a round $T_\ell \geq 1$ such that:*

1. *the agents play the strategies $(n-\ell, \ell+1)$ for the first time,*
2. *all rows $r \in [\ell+1, n-\ell-1]$ admit 0 cumulative utility, $R_r^{(T_\ell)} = 0$,*
3. *all columns $c \in [\ell+2, n-\ell]$ admit 0 cumulative utility, $C_c^{(T_\ell)} = 0$.*

Moreover for $\ell \geq 2$, the cumulative utility of row $n-\ell$ at round T_ℓ is greater than

$$R_{n-\ell}^{(T_\ell)} \geq 4\ell(4\ell-1)(4\ell-2)(4\ell-3) \cdot R_{n-\ell}^{(T_{\ell-1})} \;\; \textit{while} \;\; R_{n-1}^{(T_1)} \geq 4. \tag{2}$$

Using Lemma 3.8 we are able to establish that for a very long period of time the row player has never played row $\frac{n}{2}$ and that the column player has never played column $\frac{n}{2}+1$.

Lemma 3.9 (Exponential Lower Bound)**.** *Let an identical-payoff game with matrix A and a fictitious play sequence $(i^{(1)}, j^{(1)}), \ldots, (i^{(t)}, j^{(t)}), \ldots$ with $(i^{(1)}, j^{(1)}) = (n, 1)$. In case $(i^{(T)}, j^{(T)}) = (\frac{n}{2}, \frac{n}{2}+1)$ then $T \geq \Omega(4^n((n/2-2)!)^4)$.*

Proof. Based on Lemma 3.8, we can guarantee the existence of a round $T^\star := T_{n/2-1}$ when the players choose the strategy profile $(\frac{n}{2}+1, \frac{n}{2})$ for the first time. In addition, at round $T^\star$, it holds that $R_{n/2-1}^{(T^\star)} > 0$ and $C_{n/2+1}^{(T^\star)} = R_{n/2}^{(T^\star)} = 0$. The latter condition ensures that the strategy profile $(\frac{n}{2}, \frac{n}{2}+1)$ has not been played up to time $T^\star$.

As indicated by Observation 3.3, row $\frac{n}{2}$ has non-zero entries at columns $\frac{n}{2}$ and $\frac{n}{2}+1$. Therefore, if the cumulative utilities $R_{n/2}^{(T^\star)}$ at time $T^\star$ is zero, this implies that neither of these columns has been chosen up to that point. By the same reasoning, column $\frac{n}{2}+1$ has a non-zero entry only at row $\frac{n}{2}$, indicating that this row has not been chosen as well.

In order to continue, we require an estimate of the duration during which the strategy profile $(\frac{n}{2}+1, \frac{n}{2})$ will be played. Observation 3.3 guarantees that the utility vector of the row player is the following.

$$Ae_{\frac{n}{2}} = (0, \ldots, 0, \underbrace{2n-2}_{\frac{n}{2}}, \underbrace{2n-3}_{\frac{n}{2}+1}, 0, \ldots, 0) \tag{3}$$

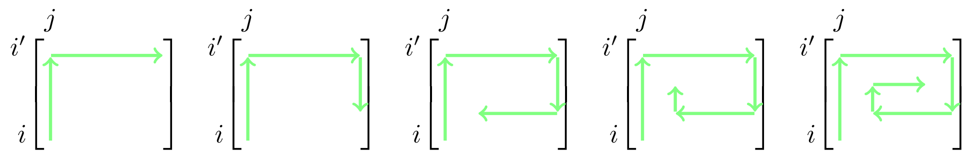

Figure 2: This figure shows the spiral trajectory generated by the fictitious play dynamic in a game with matrix A. The starting point is the lower-left element, and as the dynamic progresses, it visits all non-zero elements in ascending order of value.

We can combine this information with the fact that $R^{(T^\star)}_{n/2+1} \neq 0$, $R^{(T^\star)}_{n/2} = 0$, and that the respective entries in the payoff vector (3) of the row player differ by exactly one. This allows us to conclude that it will take at least $R^{(T^\star)}_{\frac{n}{2}+1}$ iterations for the cumulative utilities to become equal, i.e $R_{\frac{n}{2}+1} = R_{\frac{n}{2}}$. Therefore, the lower bound on the number of iterations holds regardless of the tie-breaking rule.

Now, if at any time T the agents play the strategy profile $(i^{(T)}, j^{(T)}) = (\frac{n}{2}, \frac{n}{2}+1)$, we can conclude that $T \geq R^{(T^\star)}_{\frac{n}{2}+1}$. Using Equation (2) of Lemma 3.8 we obtain:

$$T \geq R^{(T^\star)}_{\frac{n}{2}-1} \geq 16^{\frac{n}{2}-1} \left(\prod_{\ell=2}^{\frac{n}{2}-1} (\ell - 1) \right)^4 R^{(T_1)}_{n-1} \geq \Omega \left(4^n \left(\left(\frac{n}{2} - 2 \right)! \right)^4 \right)$$

□

Now that we have established the necessary technical result in Lemma 3.9, we are ready to present the proof of Theorem 3.1.

Proof of Theorem 3.1. Let $T^\star \geq 1$ denote the first time at which $\left(i^{(t)}, j^{(t)}\right) = (\frac{n}{2}, \frac{n}{2}+1)$. For any $t \leq T^\star - 1$, it holds that $\hat{x}^{(t)}_{\frac{n}{2}} = \hat{y}^{(t)}_{\frac{n}{2}+1} = 0$. Thus, Lemma 3.6 implies that $(\hat{x}^{(t)}, \hat{y}^{(t)})$ is not an approximate NE. On the other hand, at each round $t \geq T^\star$ we are ensured that $\hat{x}^{(t)}_{\frac{n}{2}}, \hat{y}^{(t)}_{\frac{n}{2}+1}$ converges to 1 with rate $1/t$. Applying Lemma 3.6 for $\epsilon := \sqrt{\epsilon}$ we get that if $(\hat{x}^{(t)}, \hat{y}^{(t)})$ is an ϵ-NE then $t \geq T^\star + \frac{1}{n\sqrt{\epsilon}} \geq \Omega\left(4^n((n/2-2)!)^4\right) + \frac{1}{n\sqrt{\epsilon}}$. It is evident that, even with a uniformly random initialization, the probability of selecting $(n, 1)$ as the starting point for fictitious play is $1/n^2$ and so the claim of the theorem follows. □

3.3 Strategy Switches and Proof of Lemma 3.8

This subsection is dedicated to presenting the proof of Lemma 3.8. However, before delving into the proof, we first provide additional statements that serve to shed light on the sequence of strategy profiles generated by fictitious play.

According to Proposition 3.5, if the strategy (i, j) is chosen by fictitious play, then either in row i or in column j, there is at most one element greater than A_{ij}. Therefore, in a subsequent round, fictitious play will necessarily choose this element as its strategy. Let i' be the row where the element of greater value is located. Intuitively, we can imagine that as fictitious play continues to play the same strategy (i, j), the payoff vector Ae_j is repeatedly added to the cumulative vector of the row player. Since this vector has a greater value at coordinate i' than at coordinate i, a strategy switch will eventually occur. We state this observation in Proposition 3.10 and defer the proof to the Appendix.

Proposition 3.10 (Strategy Switch)**.** *Let $(i^{(t)}, j^{(t)})$ be a strategy selected by fictitious play at round t, and $(i^{(t)}, j^{(t)}) \neq (\frac{n}{2}, \frac{n}{2})$. Then, in a subsequent round, fictitious play will choose the strategy of greater value that is either on row $i^{(t)}$ or column $j^{(t)}$.*

As previously hinted, the structure of the payoff matrix dictates the sequence of strategy profiles chosen by fictitious play. The sequence of strategy switches alternates between rows and columns.

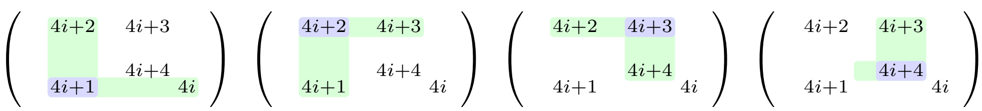

(a) $t \in [T_i^0,\, T_i^1)$ (b) $t \in [T_i^1,\, T_i^2)$ (c) $t \in [T_i^2,\, T_i^3)$ (d) $t \in [T_i^3,\, T_i^4)$

Figure 3: The figure illustrates the active row and column player for each time period, with the played strategy highlighted in purple and the corresponding payoff vectors of the row and column player highlighted in green. Additionally, the time period of each played strategy is indicated for clarity.

Continuing with the example from the previous paragraph, let us assume that the row player made the most recent strategy switch from (i, j) to (i', j). This implies that the element $A_{i'j}$ is greater than A_{ij}, as otherwise a strategy switch would not have taken place, as established in Proposition 3.10. Moreover, by Proposition 3.5, we know that there must be an element of greater value in either row i' or column j. Since $A_{i'j}$ is greater than A_{ij}, any element of greater value must be located in row i'. The same reasoning applies for the case where the column player was the last to switch strategies. We can summarize this observation by stating that if one player is the last to switch, then the other player must switch next. We formally state this in Corollary 3.11 and defer the proof to the Appendix.

Corollary 3.11 (Successive Strategy Switches)**.** *Let t be a round in which a player changes their strategy. Then exactly one of the following statements is true:*

1. *If the row player changes their strategy at round t, i.e.* $i^{(t)} \neq i^{(t-1)}$*, then the column player can only make the next strategy switch.*
2. *If the column player changes their strategy at round t, i.e.* $j^{(t)} \neq j^{(t-1)}$*, then the row player can only make the next strategy switch.*

Applying the same concept, we can observe that starting from the lower-left corner, fictitious play follows a spiral trajectory. The resulting spiral is illustrated in Figure 2. We now proceed with the main result of this section, Lemma 3.8.

Proof. Since $(i^{(1)}, j^{(1)}) = (n, 1)$ all the above claims trivially for $T_0 = 1$. We assume that the claim holds for i and will now establish it inductively for $i + 1$.

By the induction hypothesis, agents play strategies $(n - i, i + 1)$ at round $T_i^0 := T_i$. Furthermore, row $n - i$ admits cumulative utility of $R_{n-i}^{(T_i^0)}$ while row $i + 1$ admits cumulative utility of $R_{i+1}^{(T_i^0)} = 0$. According to Observation 3.3, the payoff vectors of the row and column agent are highlighted in Figure 3a. By combining these facts, we can establish the following.

Proposition 3.12 (Abridged; Full Version in Proposition B.3)**.** *There exists a round* $T_i^1 > T_i^0$ *at which the strategy profile is* $(i + 1, i + 1)$ *for the first time, column* $i + 1$ *admits cumulative utility* $C_{i+1}^{(T_i^1)} \geq (4i + 1) \cdot (R_{i+1}^{(T_i^0)} + 1)$*, and* $C_{n-i}^{(T_i^1)} = 0$.

By Proposition 3.12, at round T_i^1, the agents play strategies $(i+1, i+1)$. Furthermore, the cumulative utility of column $n - i$ equals to $C_{n-i}^{(T_i^1)} = 0$. According to Observation 3.3, the payoff vectors of the row and column agent are highlighted Figure 3b. Combining these facts, we get the following:

Proposition 3.13 (Abridged; Full Version in Proposition B.4)**.** *There exists a round* $T_i^2 > T_i^1$ *at which the strategy profile is* $(i + 1, n - i)$ *for the first time, row* $i + 1$ *admits cumulative utility* $R_{i+1}^{(T_i^2)} \geq (4i + 2) \cdot C_{i+1}^{(T_i^1)}$*, and* $R_{n-i-1}^{(T_i^2)} = 0$.

By Proposition 3.13, at round T_i^2, the agents play strategies $(i+1, n-i)$. Furthermore, the cumulative utility of row $n - i - 1$ equals to $R_{n-i-1}^{(T_i^2)} = 0$. According to Observation 3.3, the payoff vectors of the row and column agent are highlighted Figure 3c. By combining these facts, we get the following:

Proposition 3.14 (Abridged; Full Version in Proposition B.5)**.** *There exists a round* $T_i^3 > T_i^2$ *at which the strategy profile is* $(n - (i + 1), n - i)$ *for the first time,* $R_{i+2}^{(T_i^3)} = 0$.

t	*player*	*from*	*to*
2	i	4	1
4	j	1	4
11	i	1	3
38	j	4	2
174	i	3	2
990	j	2	3

(a) Table of transitions

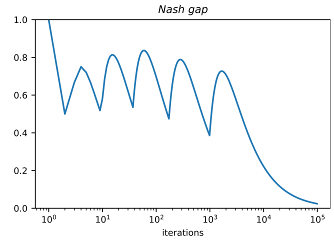

(b) Nash Gap

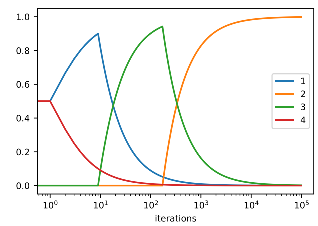

(c) Empirical strategy of x player

Figure 4: The figure displays the three aspects that were analyzed in the experiments. To preserve the crucial qualitative features in both diagrams, the x-axis was set to a logarithmic scale.

By Proposition 3.14, at round T_i^3, the agents play strategies $(n-i-1, n-i)$. the cumulative of column $i+2$ equals to $R_{i+2}^{(T_i^3)} = 0$. According to Observation 3.3, the payoff vectors of the row and column agent are highlighted Figure 3d. By combining these facts, we can establish the following.

Proposition 3.15 (Abridged; Full Version in Proposition B.6)**.** *There exists a round $T_i^4 > T_i^3$ at which the strategy profile is $(n-(i+1), (i+1)+1)$ for the first time, row $n-i-1$ admits cumulative utility $R_{n-(i+1)}^{(T_i^4)} \geq (4i+4) \cdot C_{n-i}^{(T_i^3)}$, all rows $k \in [(i+1)+1, n-(i+1)-1]$ admit $R_k^{(T_i^4)} = 0$ and all columns $k \in [(i+1)+2, n-(i+1)]$ admit $C_k^{(T_i^4)} = 0$.*

Proposition 3.15 establishes that there exits a round $T_{i+1} := T_i^4$ at which the strategy profile $(n-(i+1), (i+1)+1)$ is played for the first time. Furthermore, Proposition 3.15 confirms that all rows $k \in \{(i+1)+1, n-(i+1)-1\}$ admit $R_k^{(T_{i+1})} = 0$ and all columns $k \in \{(i+1)+2, n-(i+1)\}$ admit $C_k^{(T_{i+1})} = 0$. We still need to verify the the recursive relation Equation (2). By combining Proposition 3.12, 3.13, 3.14 and 3.15, we can deduce

$$R_{n-i-1}^{(T_{i+1})} \geq (4i+4)(4i+3)(4i+2)(4i+1)R_{n-i}^{(T_i)}.$$

□

4 Experiments

In this section, we aim to experimentally validate our findings on a 4×4 payoff matrix. Our analysis focuses on three key aspects: the round in which a new strategy switch occurs, the Nash gap throughout the game, and the empirical strategy employed by the x player. We present the plot from the row player's perspective, which is identical to that of the column player.

In Figure 4a, we provide the time steps of all strategy switches. As it is expected from the analysis, fictitious play *visits* all strategies, specifically in increasing order of their utility, to reach the pure Nash equilibrium. Moreover, in Figure 4b we observe a recurring pattern in the Nash gap diagram, where the gap increases after the selection of a new strategy with a higher utility and decreases until the next strategy switch. However, this pattern stops after the pure Nash equilibrium is reached, which is the unique approximate Nash equilibrium in accordance with Lemma 3.8.

5 Conclusion

In summary, this paper has provided a thorough examination of the convergence rate of fictitious play within a specific subset of potential games. Our research has yielded a recursive rule for constructing payoff matrices, demonstrating that fictitious play, regardless of the tie-breaking rule employed, may require super exponential time to reach a Nash equilibrium even in two-player identical payoff games. This contribution to the literature differs from previous studies and sheds new light on the limitations of fictitious play in the particular class of potential games.

Limitations and Broader Impacts: Our work is of theoretical nature and we do not see any limitations or negative ethical, societal implication.

References

[1] J. D. Abernethy, K. A. Lai, and A. Wibisono. Fast convergence of fictitious play for diagonal payoff matrices. In D. Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 1387–1404. SIAM, 2021. doi: 10.1137/1.9781611976465.84. URL `https://doi.org/10.1137/1.9781611976465.84`.

[2] L. Baudin and R. Laraki. Fictitious play and best-response dynamics in identical interest and zero-sum stochastic games. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvári, G. Niu, and S. Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 1664–1690. PMLR, 2022. URL `https://proceedings.mlr.press/v162/baudin22a.html`.

[3] G. W. Brown. Some notes on computation of games solutions. Technical report, RAND CORP SANTA MONICA CA, 1949.

[4] G. W. Brown. Iterative solution of games by fictitious play. *Act. Anal. Prod Allocation*, 13(1): 374, 1951.

[5] O. Candogan, A. E. Ozdaglar, and P. A. Parrilo. Dynamics in near-potential games. *Games Econ. Behav.*, 82:66–90, 2013. doi: 10.1016/j.geb.2013.07.001. URL `https://doi.org/10.1016/j.geb.2013.07.001`.

[6] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006. ISBN 978-0-521-84108-5. doi: 10.1017/CBO9780511546921. URL `https://doi.org/10.1017/CBO9780511546921`.

[7] Y. Chen, X. Deng, C. Li, D. Mguni, J. Wang, X. Yan, and Y. Yang. On the convergence of fictitious play: A decomposition approach. In L. D. Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 179–185. ijcai.org, 2022. doi: 10.24963/ijcai.2022/26. URL `https://doi.org/10.24963/ijcai.2022/26`.

[8] C. Daskalakis and Q. Pan. A counter-example to karlin's strong conjecture for fictitious play. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 11–20. IEEE Computer Society, 2014. doi: 10.1109/FOCS.2014.10. URL `https://doi.org/10.1109/FOCS.2014.10`.

[9] J. Heinrich, M. Lanctot, and D. Silver. Fictitious self-play in extensive-form games. In F. R. Bach and D. M. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 805–813. JMLR.org, 2015. URL `http://proceedings.mlr.press/v37/heinrich15.html`.

[10] S. Karlin. *Mathematical Methods and Theory in Games, Programming & Economics*. Addison-Wesley Professional, 1959.

[11] D. Monderer and L. S. Shapley. Fictitious play property for games with identical interests. *Journal of economic theory*, 68(1):258–265, 1996.

[12] D. Monderer and L. S. Shapley. Potential games. *Games and economic behavior*, 14(1): 124–143, 1996.

[13] P. Muller, S. Omidshafiei, M. Rowland, K. Tuyls, J. Pérolat, S. Liu, D. Hennes, L. Marris, M. Lanctot, E. Hughes, Z. Wang, G. Lever, N. Heess, T. Graepel, and R. Munos. A generalized training approach for multiagent learning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL `https://openreview.net/forum?id=Bkl5kxrKDr`.

[14] G. Ostrovski and S. van Strien. Payoff performance of fictitious play. *CoRR*, abs/1308.4049, 2013. URL `http://arxiv.org/abs/1308.4049`.

[15] J. Robinson. An iterative method of solving a game. *Annals of Mathematics*, 54(2):296–301, 1951. ISSN 0003486X. URL `http://www.jstor.org/stable/1969530`.

[16] M. O. Sayin, F. Parise, and A. E. Ozdaglar. Fictitious play in zero-sum stochastic games. *SIAM J. Control. Optim.*, 60(4):2095–2114, 2022. doi: 10.1137/21m1426675. URL `https://doi.org/10.1137/21m1426675`.

[17] B. Swenson and S. Kar. On the exponential rate of convergence of fictitious play in potential games. In *55th Annual Allerton Conference on Communication, Control, and Computing, Allerton 2017, Monticello, IL, USA, October 3-6, 2017*, pages 275–279. IEEE, 2017. doi: 10.1109/ALLERTON.2017.8262748. URL `https://doi.org/10.1109/ALLERTON.2017.8262748`.

[18] Y. Yang and J. Wang. An overview of multi-agent reinforcement learning from game theoretical perspective. *CoRR*, abs/2011.00583, 2020. URL `https://arxiv.org/abs/2011.00583`.