

On the Approximation Performance of Fictitious Play in Finite Games^{*}

Paul W. Goldberg¹, Rahul Savani¹, Troels Bjerre Sørensen², and Carmine Ventre¹

¹ Department of Computer Science, University of Liverpool, UK

² Department of Computer Science, University of Warwick, UK

Abstract. We study the performance of Fictitious Play, when used as a heuristic for finding an approximate Nash equilibrium of a two-player game. We exhibit a class of two-player games having payoffs in the range $[0, 1]$ that show that Fictitious Play fails to find a solution having an additive approximation guarantee significantly better than $1/2$. Our construction shows that for $n \times n$ games, in the worst case both players may perpetually have mixed strategies whose payoffs fall short of the best response by an additive quantity $1/2 - O(1/n^{1-\delta})$ for arbitrarily small δ . We also show an essentially matching upper bound of $1/2 - O(1/n)$.

1 Introduction

Fictitious Play is a very simple iterative process for computing equilibria of games. A detailed motivation for it is given in [5]. When it converges, it necessarily converges to a Nash equilibrium. For two-player games, it is known to converge for zero-sum games [10], or if one player has just 2 strategies [2]. On the other hand, Shapley exhibited a 3×3 game for which it fails to converge [9, 11].

Fictitious Play (FP) works as follows. Suppose that each player has a number of *actions*, or pure strategies. Initially (at iteration 1) each player starts with a single action. Thereafter, at iteration t , each player has a sequence of $t - 1$ actions which is extended with a t -th action chosen as follows. Each player makes a best response to a distribution consisting of the selection of an opponent's strategy uniformly at random from his sequence. (To make the process precise, a tie-breaking rule should also be specified; however, in the games constructed here, there will be no ties.) Thus the process generates a sequence of mixed-strategy profiles (viewing the sequences as probability distributions), and the hope is that they converge to a limiting distribution, which would necessarily be a Nash equilibrium.

The problem of computing *approximate* equilibria was motivated by the apparent intrinsic hardness of computing exact equilibria [8], even in the two-player case [4]. An ϵ -Nash equilibrium is one where each player's strategy has a payoff of at most ϵ less than the best response. Formally, for 2 players with pure strategy sets M, N and payoff functions $u_i : M \times N \rightarrow \mathbb{R}$ for $i \in \{1, 2\}$, the mixed strategy σ is an ϵ -best-response against the mixed strategy τ , if for any $m \in M$, we have $u_1(\sigma, \tau) \geq u_1(m, \tau) - \epsilon$. A pair of strategies σ, τ is an ϵ -Nash equilibrium if they are ϵ -best responses to each other.

^{*} Supported by EPSRC grants EP/G069239/1 and EP/G069034/1 "Efficient Decentralised Approaches in Algorithmic Game Theory."

Typically one assumes that the payoffs of a game are rescaled to lie in $[0, 1]$, and then a general question is: for what values of ϵ does some proposed algorithm guarantee to find ϵ -Nash equilibria? Previously, the focus has been on various algorithms that run in polynomial time. Our result for FP applies without any limit on the number of iterations; we show that a kind of cyclical behavior persists.

A recent paper of Conitzer [5] shows that FP obtains an approximation guarantee of $\epsilon = (t + 1)/2t$ for two-player games, where t is the number of FP iterations, and furthermore, if both players have access to infinitely many strategies, then FP cannot do better than this. The intuition behind this upper bound is that an action that appears most recently in a player's sequence has an ϵ -value close to 0 (at most $1/t$); generally a strategy that occurs a fraction γ back in the sequence has an ϵ -value of at most slightly more than γ (it is a best response to slightly less than $1 - \gamma$ of the opponent's distribution), and the ϵ -value of a player's mixed strategy is at most the overall average, i.e., $(t + 1)/2t$, which approaches $1/2$ as t increases.

However, as soon as the number of available pure strategies is exceeded by the number of iterations of FP, various pure strategies must get re-used, and this re-usage means, for example, that every previous occurrence of the most recent action all have ϵ -values of $1/t$. This appears to hold out the hope that FP may ultimately guarantee a significantly better additive approximation. We show that unfortunately that is not what results in the worst case. Our hope is that this result may either guide the design of more "intelligent" dynamics having a better approximation performance, or alternatively generalize to a wider class of related algorithms, e.g., the ones discussed in [6].

In Section 2 we give our main result, the lower bound of $1/2 - O(1/n^{1-\delta})$ for any $\delta > 0$, and in Section 3 we give the corresponding upper bound of $1/2 - O(1/n)$. Proofs omitted due to lack of space may be found in the full version of the paper.

2 Lower Bound

We specify a class of games with parameter n , whose general idea is conveyed in Figure 1, which shows the row player's matrix for $n = 5$; the column player's matrix is its transpose. A blank entry indicates the value zero; let $\alpha = 1 + \frac{1}{n^{1-\delta}}$ and $\beta = 1 - \frac{1}{n^{2(1-\delta)}}$ for $\delta > 0$. Both players start at strategy 1 (top left). Generally, let \mathcal{G}_n be a $4n \times 4n$ game in which the column player's payoff matrix is the transpose of the row player's payoff matrix R , which itself is specified as follows. For $i \in [2 : n]$, $R_{i,i-1} = 1$. If $i \in [n + 1 : 4n]$, $R_{i,i} = 1$. If $i \in [n + 1 : 4n]$, $R_{i,i-1} = \alpha$. Also, $R_{2n+1,4n} = \alpha$. Otherwise, if $i > j$ and $j \leq 2n$, $R_{i,j} = \beta$. Otherwise, if $i > j$ and $i - j \leq n$, $R_{i,j} = \beta$. For $j \in [3n + 1 : 4n]$ and $i \in [2n + 1 : j - n]$, $R_{i,j} = \beta$. Otherwise, we have $R_{ij} = 0$. For ease of presentation we analyze FP on \mathcal{G}_n ; the obtained results can be seen to apply to a version of \mathcal{G}_n with payoffs rescaled into $[0, 1]$ (cf. the proof of Theorem 2).

Overview. We next give a general overview and intuition on how our main result works, before embarking on the technical details. Number the strategies $1, \dots, 4n$ from top to bottom and left to right, and assume that both players start at strategy 1. Fictitious Play proceeds in a sequence of steps, which we index by positive integer t , so that step t consists of both players adding the t -th element to their sequences of length $t - 1$. It is easy to observe that since the column player's payoff matrix is the transpose of the row

show that in fact the process never converges (it cycles infinitely often) and furthermore, the lengths of the blocks increase in geometric progression. The parameters α and β in \mathcal{G}_n govern the ratio between the lengths of consecutive blocks. We choose a ratio large enough that ensures that the n strategies most recently played, occupy all but an exponentially-small fraction of the sequence. At the same time the ratio is small enough that the corresponding probability distribution does not allocate much probability to any individual strategy.

The proof. We now identify some properties of probabilities assigned to strategies by FP. We let $\ell_t(i)$ be the number of times that strategy i is played by the players until time step t of FP. Let $p_t(i)$ be the corresponding probability assigned by the players to strategy i at step t , also for any subset of actions S we use $p_t(S)$ to denote the total probability of elements of S . So it is immediate to observe that $p_t(i) = \frac{\ell_t(i)}{\sum_{j=1}^{4n} \ell_t(j)} = \frac{\ell_t(i)}{t}$. The next fact easily follows from the FP rule.

Lemma 2. *For all strategies $i \leq n$, $p_t(i) = \frac{1}{t}$ and therefore $\ell_t(i) = 1$ for any $t \geq i$.*

By Lemma 1, each strategy is played a number of consecutive times, in order, until the strategy $4n$ is played; at this point, this same pattern repeats but only for the strategies in $\{2n + 1, \dots, 4n\}$. We let t^* be the length of the longest sequence containing all the strategies in ascending order, that is, t^* is the last step of the first consecutive block of $4n$'s. We also let t_i be the last time step in which i is played during the first t^* steps, i.e., t_i is such that $\ell_{t_i}(i) = \ell_{t_i-1}(i) + 1$ and $\ell_t(i) = \ell_{t_i}(i)$ for $t \in \{t_i, \dots, t^*\}$.

Lemma 3. *For all strategies $n + 1 \leq i \leq 3n$ and all $t \in \{t_i, \dots, t^*\}$, it holds:*

$$\frac{\alpha - \beta}{\alpha - 1} p_t(i - 1) \leq p_t(i) \leq \frac{1}{t} + \frac{\alpha - \beta}{\alpha - 1} p_t(i - 1)$$

and therefore,

$$\frac{\alpha - \beta}{\alpha - 1} \ell_t(i - 1) \leq \ell_t(i) \leq 1 + \frac{\alpha - \beta}{\alpha - 1} \ell_t(i - 1).$$

Proof. By definition of t_i , strategy i is played at step t_i . This means that i is a best response for the players given the probability distributions at step $t_i - 1$. In particular, the expected payoff of i is better than the expected payoff of $i + 1$, that is,

$$\beta \sum_{j=1}^{i-2} p_{t_i-1}(j) + \alpha p_{t_i-1}(i - 1) + p_{t_i-1}(i) \geq \beta \sum_{j=1}^{i-2} p_{t_i-1}(j) + \beta p_{t_i-1}(i - 1) + \alpha p_{t_i-1}(i).$$

Since $\alpha > 1$, the above implies that $p_{t_i-1}(i) \leq \frac{\alpha - \beta}{\alpha - 1} p_{t_i-1}(i - 1)$. By explicitly writing the probabilities, we get

$$\frac{\ell_{t_i-1}(i)}{t_i - 1} \leq \frac{\alpha - \beta}{\alpha - 1} \frac{\ell_{t_i-1}(i - 1)}{t_i - 1} \iff \ell_{t_i}(i) - 1 \leq \frac{\alpha - \beta}{\alpha - 1} \ell_{t_i}(i - 1) \iff \quad (1)$$

$$\frac{\ell_{t_i}(i)}{t_i} \leq \frac{1}{t_i} + \frac{\alpha - \beta}{\alpha - 1} \frac{\ell_{t_i}(i - 1)}{t_i} \iff p_{t_i}(i) \leq \frac{1}{t_i} + \frac{\alpha - \beta}{\alpha - 1} p_{t_i}(i - 1). \quad (2)$$

At step $t_i + 1$ strategy i is not a best response to the opponent's strategy. Then, by Lemma 1, $i + 1$ is the unique best response and so the expected payoff of $i + 1$ is better than the expected payoff of i given the probability distributions at step t_i , that is,

$$\beta \sum_{j=1}^{i-2} p_{t_i}(j) + \alpha p_{t_i}(i-1) + p_{t_i}(i) \leq \beta \sum_{j=1}^{i-2} p_{t_i}(j) + \beta p_{t_i}(i-1) + \alpha p_{t_i}(i).$$

Since $\alpha > 1$, the above implies that

$$p_{t_i}(i) \geq \frac{\alpha - \beta}{\alpha - 1} p_{t_i}(i-1), \text{ and then that } \ell_{t_i}(i) \geq \frac{\alpha - \beta}{\alpha - 1} \ell_{t_i}(i-1). \quad (3)$$

By definition of t_i action i will not be played anymore until time step t^* . Similarly, Lemma 1 shows that $i - 1$ will not be a best response twice in the time interval $[1, t^*]$ and so will not be played until step t^* . So, the claim follows from (1), (2) and (3). \square

Using similar arguments, we can prove a similar result on the actions $3n+1, \dots, 4n-1$.

Lemma 4. For $i \in \{3n+1, \dots, 4n-1\}$ and $t \in \{t_i, \dots, t^*\}$, it holds:

$$\frac{\alpha - \beta}{\alpha - 1} p_t(i-1) \leq p_t(i) \leq \frac{1}{t} + \frac{\alpha - \beta}{\alpha - 1} p_t(i-1) + \frac{\beta}{\alpha - 1} p_t(i-n)$$

and therefore,

$$\frac{\alpha - \beta}{\alpha - 1} \ell_t(i-1) \leq \ell_t(i) \leq 1 + \frac{\alpha - \beta}{\alpha - 1} \ell_t(i-1) + \frac{\beta}{\alpha - 1} \ell_t(i-n).$$

Exploiting the properties given by Lemmata 2, 3 and 4, the next lemma shows that we can “forget” about the first $2n$ actions at the cost of paying an exponentially small addend in the payoff function.

Lemma 5. For any $\delta > 0$, $\alpha = 1 + \frac{1}{n^{1-\delta}}$ and $\beta = 1 - \frac{1}{n^{2(1-\delta)}}$, $\sum_{j=1}^{2n} p_{t^*}(j) \leq 2^{-n^\delta}$.

The theorem below generalizes the above arguments to the cycles that FP visits in the last block of the game, i.e., the block which comprises strategies $S = \{2n+1, \dots, 4n\}$. Since we focus on this part of the game, to ease the presentation, our notation uses circular arithmetic on the elements of S . For example, the action $j+2$ will denote action $2n+2$ for $j=4n$ and the action $j-n$ will be the strategy $3n+1$ for $j=2n+1$. Note that under this notation $j-2n = j+2n = j$ for each $j \in S$.

Theorem 1. For any $\delta > 0$, $\alpha = 1 + \frac{1}{n^{1-\delta}}$ and $\beta = 1 - \frac{1}{n^{2(1-\delta)}}$, n sufficiently large, any $t \geq t^*$ we have

$$\frac{p_t(i)}{p_t(i-1)} \geq 1 + \frac{1}{n^{1-\delta}} \quad \forall i \in S \text{ with } i \neq s_t, s_t + 1, \text{ and } \frac{p_t(i)}{p_t(i-1)} \leq 1 + \frac{3}{n^{1-\delta}} \quad \forall i \in S.$$

Proof. The proof is by induction on t .

Base. For the base of the induction, consider $t = t^*$ and note that at that point $s_{t^*} = 4n$ and $s_{t^*} + 1 = 2n + 1$. Therefore we need to show the lower bound for any strategy $i \in \{2n + 2, \dots, 4n - 1\}$. From Lemmata 3 and 4 we note that for $i \neq 4n, 2n + 1$,

$$\frac{p_{t^*}(i)}{p_{t^*}(i-1)} \geq \frac{\alpha - \beta}{\alpha - 1} = 1 + \frac{1}{n^{1-\delta}}.$$

As for the upper bound, we first consider the case of $i \neq 4n, 2n + 1$. Lemma 3 implies that for $i = 2n + 2, \dots, 3n$,

$$\frac{p_{t^*}(i)}{p_{t^*}(i-1)} \leq \frac{1}{t^*} + \frac{\alpha - \beta}{\alpha - 1},$$

while Lemma 4 implies that for $i = 3n + 1, \dots, 4n - 1$,

$$\frac{p_{t^*}(i)}{p_{t^*}(i-1)} \leq \frac{1}{t^*} + \frac{\alpha - \beta}{\alpha - 1} + \frac{\beta}{\alpha - 1} \frac{p_{t^*}(i-n)}{p_{t^*}(i-1)} = \frac{1}{t^*} + \frac{\alpha - \beta}{\alpha - 1} + \frac{\beta}{\alpha - 1} \frac{\ell_{t^*}(i-n)}{\ell_{t^*}(i-1)}.$$

To give a unique upper bound for both cases, we only focus on the above (weaker) upper bound and next are going to focus on the ratio $\frac{\ell_{t^*}(i-n)}{\ell_{t^*}(i-1)}$. We use Lemmata 3 and 4 and get

$$\ell_{t^*}(i-1) \geq \frac{\alpha - \beta}{\alpha - 1} \ell_{t^*}(i-2) \geq \left(\frac{\alpha - \beta}{\alpha - 1}\right)^2 \ell_{t^*}(i-3) \geq \left(\frac{\alpha - \beta}{\alpha - 1}\right)^{n-1} \ell_{t^*}(i-n).$$

By setting α and β as in the hypothesis and noticing that $t^* \geq n \geq n^{1-\delta}$ we then obtain that

$$\frac{p_{t^*}(i)}{p_{t^*}(i-1)} \leq 1 + \frac{2}{n^{1-\delta}} + \left(1 + \frac{1}{n^{1-\delta}}\right)^{1-n} \frac{n^{2(1-\delta)} - 1}{n^{1-\delta}}.$$

We end this part of the proof by showing that the last addend on the right-hand side of the above expression is upper bounded by $\frac{1}{4n^{1-\delta}}$. To do so we need to prove

$$\left(1 + \frac{1}{n^{1-\delta}}\right)^{1-n} \leq \frac{1}{4n^{1-\delta}} \frac{n^{1-\delta}}{n^{2(1-\delta)} - 1}, \quad (4)$$

which is equivalent to $\left(1 + \frac{1}{n^{1-\delta}}\right)^{n^{1-\delta}} \geq 4(n^{2(1-\delta)} - 1)$. We now lower

bound the left-hand side of the latter inequality: $\left(1 + \frac{1}{n^{1-\delta}}\right)^{n^{1-\delta}} > \frac{2^{n^\delta}}{2^{\frac{1}{n^{1-\delta}}}} >$

$\frac{2^{n^\delta}}{2}$, where the first inequality follows from the fact that the function $(1 + 1/x)^x$ is greater than 2 for $x > 2$ and the second one follows from the fact that $2^{\frac{1}{n^{1-\delta}}} < 2$ for $n^{1-\delta} > 1$. Then, since for $n \geq \sqrt[2(1-\delta)]{4}$, $5n^{2(1-\delta)} \geq 4(n^{2(1-\delta)} - 1)$, to prove (4) is enough to show $2^{n^\delta} \geq 2(5n^{2(1-\delta)}) \iff n^\delta \geq 2(1-\delta) \log_2(10n)$. To prove the latter, since $\delta > 0$, it is enough to observe that the function n^δ is certainly bigger than the function $2 \log_2(10n) > 2(1-\delta) \log_2(10n)$ for n large enough (e.g., for $\delta = 1/2$, this

is true for $n > 639$). Similarly to the proof of Lemma 4, we can prove that the upper bound holds at time step t^* for $i = 4n, 2n + 1$. This concludes the proof of the base of the induction.

Inductive step. Now we assume the claim is true until time step $t - 1$ and we show it for time step t . By inductive hypothesis, the following is true, with $j \neq s_{t-1}, s_{t-1} + 1$

$$1 + \frac{1}{n^{1-\delta}} \leq \frac{p_{t-1}(j)}{p_{t-1}(j-1)} \leq 1 + \frac{3}{n^{1-\delta}}, \quad (5)$$

$$\frac{p_{t-1}(s_{t-1}+1)}{p_{t-1}(s_{t-1})} \leq 1 + \frac{3}{n^{1-\delta}}. \quad (6)$$

We first consider the case in which $s_t \neq s_{t-1}$. By Lemma 1, the strategy played at time t is $s_{t-1} + 1$, i.e., $s_t = s_{t-1} + 1$. Let $s_{t-1} = i$ and then we have $s_t = i + 1$. By inductive hypothesis, for all the actions $j \neq i, i + 1, i + 2$ we have

$$\frac{\alpha - \beta}{\alpha - 1} = 1 + \frac{1}{n^{1-\delta}} \leq \frac{p_t(j)}{p_t(j-1)} \leq 1 + \frac{3}{n^{1-\delta}}. \quad (7)$$

Indeed, for these actions j , $\ell_{t-1}(j) = \ell_t(j)$ and $\ell_{t-1}(j-1) = \ell_t(j-1)$. Therefore the probabilities of j and $j-1$ at time t are simply those at time $t-1$ rescaled by the same amount and the claim follows from (5). The upper bound on the ratio $\frac{p_t(i+2)}{p_t(i+1)}$ easily follows from the upper bound in (5) as $\ell_{t-1}(i+2) = \ell_t(i+2)$ and $\ell_{t-1}(i+1) < \ell_t(i+1) = \ell_{t-1}(i+1) + 1$. However, as $s_t = i + 1$ here we need to prove lower and upper bound also for the ratio $\frac{p_t(i)}{p_t(i-1)}$ and the upper bound for the ratio $\frac{p_t(i+1)}{p_t(i)}$ (this proof can be found in the full version of the paper).

Claim. It holds $1 + \frac{1}{n^{1-\delta}} \leq \frac{p_t(i)}{p_t(i-1)} \leq 1 + \frac{3}{n^{1-\delta}}$.

Proof. To prove the claim we first focus on the last block of the game, i.e., the block in which players have strategies in $\{2n+1, \dots, 4n\}$. Recall that our notation uses circular arithmetic on the number of actions of the block.

The fact that action $i + 1$ is better than action i after $t - 1$ time steps implies that

$$p_{t-1}(i) + \alpha p_{t-1}(i-1) + \beta p_{t-1}(i-n) \leq \alpha p_{t-1}(i) + p_{t-1}(i+1) + \beta p_{t-1}(i-1)$$

and then since $\alpha > 1$

$$\frac{p_{t-1}(i)}{p_{t-1}(i-1)} \geq \frac{\alpha - \beta}{\alpha - 1} + \frac{\beta}{\alpha - 1} \frac{p_{t-1}(i-n)}{p_{t-1}(i-1)} - \frac{1}{\alpha - 1} \frac{p_{t-1}(i-2n+1)}{p_{t-1}(i-1)}. \quad (8)$$

We next show that $\frac{\beta p_{t-1}(i-n) - p_{t-1}(i-2n+1)}{(\alpha-1)p_{t-1}(i-1)} \geq -\frac{1}{4n^{1-\delta}}$ or equivalently that $\frac{p_{t-1}(i-n)}{p_{t-1}(i-2n+1)} \geq \frac{1}{\beta} - \frac{(\alpha-1)p_{t-1}(i-1)}{4\beta n^{1-\delta} p_{t-1}(i-2n+1)}$. To prove this it is enough to show that $\frac{p_{t-1}(i-n)}{p_{t-1}(i-2n+1)} \geq \frac{1}{\beta}$. We observe that

$$\frac{p_{t-1}(i-n)}{p_{t-1}(i-2n+1)} = \frac{p_{t-1}(i-n)}{p_{t-1}(i-n-1)} \cdots \frac{p_{t-1}(i-2n+2)}{p_{t-1}(i-2n+1)} \geq \left(\frac{\alpha - \beta}{\alpha - 1}\right)^{n-1} \geq \frac{1}{\beta},$$

where the first inequality follows from inductive hypothesis (we can use the inductive hypothesis as all the actions involved above are different from i and $i + 1$) and the second inequality follows from the aforementioned observation that, for sufficiently large n , $n^\delta \geq 2 \log_2(2n)$. Then to summarize, for α and β as in the hypothesis, (8) implies that

$$\frac{p_t(i)}{p_t(i-1)} = \frac{p_{t-1}(i)}{p_{t-1}(i-1)} \geq 1 + \frac{1}{n^{1-\delta}} - \frac{1}{4n^{1-\delta}},$$

where the first equality follows from $\ell_{t-1}(i) = \ell_t(i)$ and $\ell_{t-1}(i-1) = \ell_t(i-1)$, which are true because $s_t = i + 1$.

Since action $i + 1$ is worse than strategy i at time step $t - 1$ we have that

$$p_{t-1}(i) + \alpha p_{t-1}(i-1) + \beta p_{t-1}(i-n) \geq \alpha p_{t-1}(i) + p_{t-1}(i+1) + \beta p_{t-1}(i-1)$$

and then since $\alpha > 1$, $\frac{p_{t-1}(i)}{p_{t-1}(i-1)} \leq \frac{\alpha-\beta}{\alpha-1} + \frac{\beta}{\alpha-1} \frac{p_{t-1}(i-n)}{p_{t-1}(i-1)} - \frac{1}{\alpha-1} \frac{p_{t-1}(i-2n+1)}{p_{t-1}(i-1)}$. Similarly to the proof of Lemma 3 above this can be shown to imply

$$\begin{aligned} \frac{p_t(i)}{p_t(i-1)} &\leq \frac{1}{t} + \frac{\alpha-\beta}{\alpha-1} + \frac{\beta}{\alpha-1} \frac{p_t(i-n)}{p_t(i-1)} - \frac{1}{\alpha-1} \frac{p_t(i-2n+1)}{p_t(i-1)} \\ &\leq \frac{1}{t} + \frac{\alpha-\beta}{\alpha-1} + \frac{\beta}{\alpha-1} \frac{p_t(i-n)}{p_t(i-1)}. \end{aligned} \quad (9)$$

We now upper bound the ratio $\frac{\beta}{\alpha-1} \frac{p_t(i-n)}{p_t(i-1)}$. By repeatedly using the inductive hypothesis (7) we have that

$$p_t(i-1) \geq \frac{\alpha-\beta}{\alpha-1} p_t(i-2) \geq \left(\frac{\alpha-\beta}{\alpha-1}\right)^2 p_t(i-3) \geq \left(\frac{\alpha-\beta}{\alpha-1}\right)^{n-1} p_t(i-n).$$

(Note again that we can use the inductive hypothesis as none of the actions above is i or $i + 1$.) This yields

$$\frac{\beta}{\alpha-1} \frac{p_t(i-n)}{p_t(i-1)} \leq \frac{\beta}{\alpha-1} \left(\frac{\alpha-1}{\alpha-\beta}\right)^{n-1} \leq \frac{1}{4n^{1-\delta}},$$

where the last inequality is proved above (see (4)). Therefore, since $t \geq n^{1-\delta}$, (9) implies the following

$$\frac{p_t(i)}{p_t(i-1)} \leq 1 + \frac{2}{n^{1-\delta}} + \frac{1}{4n^{1-\delta}}.$$

To conclude the proof we must now consider the contribution to the payoffs of the actions $1, \dots, 2n$ that are not in the last block. However, Lemma 5 shows that all those actions are played with probability $1/2^{n^\delta}$ at time t^* . Since we prove above (see Lemma 1) that these actions are not played anymore after time step t^* this implies that $\sum_{j=1}^{2n} p_t(j) \leq \sum_{j=1}^{2n} p_{t^*}(j) \leq 2^{-n^\delta}$. Thus the overall contribution of these strategies is upper bounded by $\frac{1}{2^{n^\delta}}(\alpha-\beta) \leq \frac{1}{2^{n^\delta}} \leq \frac{1}{4n^{1-\delta}}$ where the last bound follows from the aforementioned fact that, for n sufficiently large, $n^\delta \geq (1-\delta) \log_2(4n)$. This concludes the proof of this claim. \square

Finally, we consider the case in which $s_{t-1} = s_t$. In this case, for the actions $j \neq s_t, s_t + 1$ it holds $\ell_{t-1}(j) = \ell_t(j)$ and $\ell_{t-1}(j-1) = \ell_t(j-1)$. Therefore, similarly to the above, for these actions j the claim follows from (5). The upper bound for the ratio $\frac{p_t(s_t+1)}{p_t(s_t)}$ easily follows from (6) as $\ell_{t-1}(s_t+1) = \ell_t(s_t+1)$ and $\ell_{t-1}(s_t) < \ell_t(s_t) = \ell_{t-1}(s_t) + 1$. The remaining case to analyze is the upper bound on the ratio $\frac{p_t(s_t)}{p_t(s_t-1)}$. To prove this we can use *mutatis mutandis* the proof of the upper bound on the ratio $\frac{p_t(i)}{p_t(i-1)}$ for the case $s_{t-1} \neq s_t$ (in that proof simply set with $s_t = i$). \square

The claimed performance of Fictitious Play, in terms of the approximation to the best response that it computes, follows directly from this theorem.

Theorem 2. *For any value of $\delta > 0$ and any time step t , Fictitious Play returns an ϵ -NE with $\epsilon \geq \frac{1}{2} - O\left(\frac{1}{n^{1-\delta}}\right)$.*

Proof. For $t \leq n$ the result follows since the game is similar to [5]. In details, for $t \leq n$ the payoff associated to the best response, which in this case is $s_t + 1$, is upper bounded by 1. On the other hand, the payoff associated to the current strategy prescribed by FP is lower bounded by $\frac{\beta}{i^2} \sum_{j=0}^{i-1} j$ where $i = s_t$. Therefore, the regret of either player normalized to the $[0, 1]$ interval satisfies: $\epsilon \geq \frac{1}{\alpha} - \frac{\beta}{\alpha} \frac{i-1}{2i}$. Since $\frac{i-1}{2i} < 1/2$, the fact that $1 - \frac{\beta}{2} - \frac{\alpha}{2} + \frac{\alpha}{n^{1-\delta}} \geq 0$ (which is true given the values of α and β) yields the claim. For $t \leq t^*$ the result follows from Lemmata 3 and 4; while the current strategy s_t (for $t \leq t^*$) has payoff approximately 1, the players' mixed strategies have nearly all their probability on the recently played strategies, but with no pure strategy having very high probability, so that some player is likely to receive zero payoff; by symmetry each player has payoff approximately $\frac{1}{2}$. This is made precise below, where it is applied in more detail to the case of $t > t^*$.

We now focus on the case $t > t^*$. Recall that for a set of strategies S , $p_t(S) = \sum_{i \in S} p_t(i)$. Let S_t be the set $\{2n+1, \dots, s_t\} \cup \{s_t+n, \dots, 4n\}$ if $s_t \leq 3n$, or the set $\{s_t-n, \dots, s_t\}$ in the case that $s_t > 3n$. Let $S'_t = \{2n+1, \dots, 4n\} \setminus S_t$. Also, let $s_t^{\max} = \arg \max_{i \in \{2n+1, \dots, 4n\}} (p_t(i))$; note that by Theorem 1, s_t^{\max} is equal to either s_t or s_t^- , where $s_t^- = s_t - 1$ if $s_t > 2n$, or $4n$ if $s_t = 2n$. We start by establishing the following claim:

Claim. For sufficiently large n , $p_t(S_t) \geq 1 - \frac{2n-1}{2n^\delta}$.

Proof. To see this, note that for all $x \in S'_t$, by $p_t(s_t^{\max}) \geq p_t(s_t^{\max} - 1)$ and Theorem 1 we have

$$\frac{p_t(s_t^{\max})}{p_t(x)} = \frac{p_t(s_t^{\max})}{p_t(s_t^{\max} - 1)} \frac{p_t(s_t^{\max} - 1)}{p_t(s_t^{\max} - 2)} \cdots \frac{p_t(x+1)}{p_t(x)} \geq \left(1 + \frac{1}{n^{1-\delta}}\right)^{k-1},$$

where k is the number of factors on the right-hand side of the equality above, i.e., the number of strategies between x and s_t^{\max} . Thus, as $k \geq n$,

$$p_t(x) \leq \frac{p_t(s_t^{\max})}{\left(1 + \frac{1}{n^{1-\delta}}\right)^{k-1}} \leq \left(1 + \frac{1}{n^{1-\delta}}\right)^{1-k} \leq \left(1 + \frac{1}{n^{1-\delta}}\right)^{1-n} \leq 4^{(1-n)/(n^{1-\delta})}.$$

Hence $p_t(S'_t) \leq (2n)4^{(1-n)/(n^{1-\delta})} = \frac{n4^{1/n^{1-\delta}}}{2n^\delta} < \frac{2n}{2n^\delta}$, where the last inequality follows from the fact that, for large n , $4^{1/n^{1-\delta}} < 2$. Then $p_t(S_t) \geq 1 - p_t(S'_t) - p_t(\{1, \dots, 2n\})$, which establishes the claim, since Lemma 5 establishes a strong enough upper bound on $p_t(\{1, \dots, 2n\})$. \square

Claim. The current best response at time t , s_t , has payoff at least $\beta \left(1 - \frac{2n-1}{2n^\delta}\right)$.

Proof. The action s_t receive a payoff of at least β when the opponent plays any strategy from S_t ; the claim follows using Claim 2. \square

Let E_t denote the expected payoff to either player that would result if they both select a strategy from the mixed distribution that allocates to each strategy x , the probability $p_t(x)$. The result will follow from the following claim:

Claim. For sufficiently large n , $E_t \leq \frac{\alpha}{2} + \frac{6}{n^{1-\delta}} + \alpha \frac{2n}{2n^\delta}$.

Proof. The contribution to E_t from strategies in $\{1, \dots, n\}$, together with strategies in S'_t , may be upper-bounded by α times the probability that any of that strategies get played. This probability is by Lemma 5 and the claim above exponentially small, namely $2n/2^{n^\delta}$.

Suppose instead that both players play from S_t . If they play different strategies, their total payoff will be at most α , since one player receives payoff 0. If they play the same strategy, they both receive payoff 1. We continue by upper-bounding the probability that they both play the same strategy. This is upper-bounded by the largest probability assigned to any single strategy, namely $p_t(s_t^{\max})$.

Suppose for contradiction that $p_t(s_t^{\max}) > 6/n^{1-\delta}$. At this point, note that by Theorem 1, for any strategy $s \in S_t$, we have $\frac{p_t(s_t^{\max})}{p_t(s)} \leq \left(1 + \frac{3}{n^{1-\delta}}\right)^k$, where k is the distance between s and s_t^{\max} . Therefore, denoting $r = \left(1 + \frac{3}{n^{1-\delta}}\right)^{-1}$, we obtain

$$p_t(S_t) = \sum_{s \in S_t} p_t(s) = p_t(s_t) + \sum_{i=s_t-n}^{s_t-1} p_t(i) \geq p_t(s_t^{\max}) \sum_{k=0}^{n-1} r^k.$$

Applying the standard formula for the partial sum of a geometric series we have $p_t(S_t) \geq \frac{6}{n^{1-\delta}} \left(\frac{1-r^n}{1-r}\right)$. Noting that $1 - r^n > \frac{1}{2}$ we have $p_t(S_t) > \frac{6}{n^{1-\delta}} \cdot \left(\frac{1}{2}\right) \cdot \left(\frac{n^{1-\delta}}{3}\right)$ which is greater than 1, a contradiction.

The expected payoff E_t to either player, is, by symmetry, half the expected total payoff: we have $E_t \leq \left(1 - \frac{2n}{2n^\delta} - \frac{6}{n^{1-\delta}}\right) \frac{\alpha}{2} + \frac{6}{n^{1-\delta}} + \frac{2n}{2n^\delta} \alpha$ which yields the claim. \square

We now show that Fictitious Play never achieves an ϵ -value better than $\frac{1}{2} - O\left(\frac{1}{n^{1-\delta}}\right)$. From the last two claims the regret of either player normalized to $[0, 1]$ satisfies:

$$\epsilon \geq \frac{\beta}{\alpha} \left(1 - \frac{2n-1}{2n^\delta}\right) - \frac{1}{2} - \frac{6}{\alpha n^{1-\delta}} - \frac{2n}{2n^\delta} = \frac{1}{2} - O\left(\frac{1}{n^{1-\delta}}\right).$$

This concludes the proof. \square

3 Upper Bound

In this section, n denotes the number of pure strategies of both players. Let a, b denote the FP sequences of pure strategies of length t , for players 1 and 2 respectively. Let $a_{[k:\ell]}$ denote the subsequence a_k, \dots, a_ℓ . We overload notation and use a to also denote the mixed strategy that is uniform on the corresponding sequence.

Let m^* be a best response against b , and let ϵ denote the smallest ϵ' for which a is an ϵ' -best-response against b . To derive a bound on ϵ , we use the most recent occurrence of pure strategy in a . For $k \in \{1, \dots, t\}$, let $f(k)$ denote the last occurrence of a_k in the sequence a , that is, $f(k) := \max_{\ell \in \{1, \dots, t\}, a_\ell = a_k} \ell$. It is not hard to show that

$$\epsilon \leq 1 + \frac{1}{t} - \frac{1}{t^2} \sum_{i=1}^t f(i). \quad (10)$$

To provide a guarantee on the performance of FP, we find the sequence a that maximizes the right hand-side of (10), i.e., that minimizes $\sum_{i=1}^t f(i)$.

Definition 1. For a FP sequence a , let $S(a) := \sum_{i=1}^t f(a_i)$ and let $\hat{a} = \arg \min_a S(a)$.

The following three lemmata allow to characterize \hat{a} , the sequence that minimizes $S(a)$.

Lemma 6. The entries of \hat{a} take on exactly n distinct values.

We now define a transformation of an FP sequence a into a new sequence a' so that $S(a') < S(a)$ if $a \neq a'$.

Definition 2. Suppose the entries of a take on d distinct values. We define x_1, \dots, x_d to be the last occurrences, $\{f(a_i) \mid i \in [t]\}$, in ascending order. Formally, let $x_d := a_t$ and for $k < d$ let $x_k := a_i$ be such that $i := \arg \max_{j=1, \dots, t} a_j \notin \{x_{k+1}, \dots, x_d\}$. For $i = 1, \dots, d$, let $\#(x_i) := |\{a_j \mid j \in [t], a_j = x_i\}|$, which is the number of occurrences of x_i in a . Define a' as $a' := \underbrace{x_1, \dots, x_1}_{\#(x_1)}, \underbrace{x_2, \dots, x_2}_{\#(x_2)}, \dots, \underbrace{x_d, \dots, x_d}_{\#(x_d)}$.

Lemma 7. For any FP sequence a , let a' be as in Definition 2. If $a' \neq a$ then $S(a') < S(a)$.

Lemma 8. Let $n, t \in \mathbb{N}$ be such that $n \mid t$. Let a be a sequence of length t of the form $a = 1, \dots, 1, 2, \dots, 2, \dots, n, \dots, n$, where the blocks of consecutive actions have length c_1, c_2, \dots, c_n , respectively. Then $S(a)$ is minimized if and only if $c_1 = \dots = c_n = t/n$.

Proof. We refer to the maximal length subsequence of entries with value $u \in \{1, \dots, n\}$ as *block* u . Consider two adjacent blocks u and $u + 1$, where block u starts at i and block $u + 1$ starts at j and finishes at k . The contribution of these two blocks to $S(a)$ is $\sum_{i}^{j-1} (j-1) + \sum_j^k k = j^2 - (k+i)j + (i+k)$. If $k+i$ is even, this contribution is minimized when $j = \frac{k+i}{2}$. If $k+i$ is odd, this contribution is minimized for both values $j = \lfloor \frac{k+i}{2} \rfloor$ and $j = \lceil \frac{k+i}{2} \rceil$.

Now suppose for the sake of contradiction that $S(a)$ is minimized when $c_1 = \dots = c_n = t/n$ does not hold. There are two possibilities. Either there are two adjacent

blocks whose lengths differ by more than one, in which case we immediately have a contradiction. If not, then it must be the case that all pairs of adjacent blocks differ in length by at most one. In particular, there must be a block of length $t/n + 1$ and another of length $t/n - 1$ with all blocks in between of length t/n . Flipping the leftmost of these blocks with its right neighbor will not change the sum $S(a)$. Repeatedly doing this until the blocks of lengths $t/n + 1$ and $t/n - 1$ are adjacent, does not change $S(a)$. Then we have two adjacent blocks that differ in length by more than one, which contradicts the fact that $S(a)$ was minimized. \square

Theorem 3. *If $n|t$, the FP strategies (a, b) are a $(\frac{1}{2} + \frac{1}{t} - \frac{1}{2n})$ -equilibrium.*

Proof. By symmetry, it suffices to show that \hat{a} is a $(\frac{1}{2} + \frac{1}{t} - \frac{1}{2n})$ -best-response against b . Applying Lemma 6, Lemma 7 and Lemma 8, we have that $\hat{a} = m_1, \dots, m_1, m_2, \dots, m_2, \dots, m_n, \dots, m_n$, where m_1, \dots, m_n is an arbitrary labeling of player 1's pure strategies and where each block of actions has length t/n . Using (10), we have that

$$\epsilon \leq 1 + \frac{1}{t} - \frac{1}{t^2} \sum_{i=1}^t f(\hat{a}_i) = 1 + \frac{1}{t} - \frac{1}{t^2} \frac{t}{n} \sum_{i=1}^n \left(\frac{i \cdot t}{n}\right) = 1 + \frac{1}{t} - \frac{n+1}{2n} = \frac{1}{2} + \frac{1}{t} - \frac{1}{2n}.$$

This concludes the proof. \square

4 Discussion

Daskalakis et al. [7] gave a very simple algorithm that achieves an approximation guarantee of $\frac{1}{2}$; subsequent algorithms e.g. [3,12] improved on this, but at the expense of being more complex and centralized, commonly solving one or more derived LPs from the game. Our result suggests that further work on the topic might address the question of whether $\frac{1}{2}$ is a fundamental limit to the approximation performance obtainable by certain types of algorithms that are in some sense simple or decentralized. The question of specifying appropriate classes of algorithms is itself challenging, and is also considered in [6] in the context of algorithms that provably fail to find Nash equilibria without computational complexity theoretic assumptions.

References

1. Brandt, F., Fischer, F., Harrenstein, P.: On the Rate of Convergence of Fictitious Play. In: 3rd Symposium on Algorithmic Game Theory, pp. 102–113 (2010)
2. Berger, U.: Fictitious play in $2 \times n$ games. *Journal of Economic Theory* 120, 139–154 (2005)
3. Bosse, H., Byrka, J., Markakis, E.: New Algorithms for Approximate Nash Equilibria in Bimatrix Games. In: Proceedings of the 3rd International Workshop on Internet and Network Economics, pp. 17–29 (2007)
4. Chen, X., Deng, X., Teng, S.-H.: Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM* 56(3), 1–57 (2009)
5. Conitzer, V.: Approximation Guarantees for Fictitious Play. In: Proc of 47th Annual Allerton Conference on Communication, Control, and Computing, pp. 636–643 (2009)

6. Daskalakis, C., Frongillo, R., Papadimitriou, C.H., Pierrakos, G., Valiant, G.: On Learning Algorithms for Nash Equilibria. In: Kontogiannis, S., Koutsoupias, E., Spirakis, P.G. (eds.) AGT 2010. LNCS, vol. 6386, pp. 114–125. Springer, Heidelberg (2010)
7. Daskalakis, C., Mehta, A., Papadimitriou, C.H.: A Note on Approximate Nash Equilibria. *Theoretical Computer Science* 410(17), 1581–1588 (2009)
8. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The Complexity of Computing a Nash Equilibrium. *SIAM Journal on Computing* 39(1), 195–259 (2009)
9. Fudenberg, D., Levine, D.K.: *The Theory of Learning in Games*. MIT Press, Cambridge (1998)
10. Robinson, J.: An Iterative Method of Solving a Game. *Annals of Mathematics* 54(2), 296–301 (1951)
11. Shapley, L.: Some topics in two-person games. In: *Advances in Game Theory*. *Annals of Mathematics Studies*, vol. 52, Princeton University Press, Princeton (1964)
12. Tsaknakis, H., Spirakis, P.G.: An Optimization Approach for Approximate Nash Equilibria. *Internet Mathematics* 5(4), 365–382 (2008)