

On the Rate of Convergence of Fictitious Play

Felix Brandt · Felix Fischer · Paul
Harrenstein

Abstract Fictitious play is a simple learning algorithm for strategic games that proceeds in rounds. In each round, the players play a best response to a mixed strategy that is given by the empirical frequencies of actions played in previous rounds. There is a close relationship between fictitious play and the Nash equilibria of a game: if the empirical frequencies of fictitious play converge to a strategy profile, this strategy profile is a Nash equilibrium. While fictitious play does not converge in general, it is known to do so for certain restricted classes of games, such as constant-sum games, non-degenerate $2 \times n$ games, and potential games. We study the rate of convergence of fictitious play and show that, in all the classes of games mentioned above, fictitious play may require an exponential number of rounds (in the size of the representation of the game) before *some* equilibrium action is eventually played. In particular, we show the above statement for symmetric constant-sum win-lose-tie games.

Keywords Game Theory · Nash Equilibrium · Fictitious Play · Rate of Convergence

1 Introduction

A common criticism of Nash equilibrium, the most prominent solution concept of the theory of strategic games, is that it fails to capture how players' deliberation processes actually reach a steady state. When considering a set of human or artificial agents engaged in a parlor game or a more austere decision-making

F. Brandt · P. Harrenstein
Institut für Informatik, Technische Universität München, 85748 Garching, Germany
E-mail: {brandtf,harrenst}@in.tum.de

F. Fischer
School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138,
USA
E-mail: fischerf@seas.harvard.edu

situation, it is hard to imagine that they would after some deliberation arrive at a Nash equilibrium, a carefully chosen probability distribution over all possible courses of action. One reason why this behavior is so hard to imagine is that Nash equilibrium rests on rather strong assumptions concerning the rationality of players and the ability to reliably carry out randomizations. Another concern is that in many settings finding a Nash equilibrium is computationally intractable.

A more reasonable scenario would be that agents face a strategic situation by playing a game in their heads, going through several rounds of speculation and counterspeculation as to how their opponents might react and how they would react in turn. This is the idea underlying *fictitious play* (FP). FP proceeds in rounds. In the first round, each player arbitrarily chooses one of his actions. In subsequent rounds, each player looks at the empirical frequency of play of their respective opponents in previous rounds, interprets it as a probability distribution, and myopically plays a pure best response against this distribution. FP can also be seen as a *learning algorithm* for games that are played repeatedly, such that the intermediate best responses are actually played. This interpretation rests on the simplifying assumption that the other players follow a fixed strategy.

FP was originally introduced by Brown [7] as an algorithm to approximate the value of constant-sum games, or equivalently compute approximate solutions to linear programs [10]. Shortly after, it was shown that FP does indeed converge to the desired solution [25]. While convergence does not extend to arbitrary games, as illustrated by Shapley [27], it does so for quite a few interesting classes of games, and much research has focussed—and still focusses—on identifying such classes ([3], and the references therein). Both as a linear program solver and as a learning algorithm, FP is easily outperformed by more sophisticated algorithms. However, FP is of captivating simplicity and is therefore considered one of the most convincing explanations of Nash equilibrium play. As Luce and Raiffa put it: “Brown’s results are not only computationally valuable but also quite illuminating from a substantive point of view. Imagine a pair of players repeating a game over and over again. It is plausible that at every stage a player attempts to exploit his knowledge of his opponent’s past moves. Even though the game may be too complicated or too nebulous to be subjected to an adequate analysis, experience in repeated plays may tend to a statistical equilibrium whose (time) average return is approximately equal to the value of the game” [17, p. 443].

In this paper, we show that in virtually all classes of games where FP is known to converge to a Nash equilibrium, it may take an exponential number of rounds (in the representation of the game) before any equilibrium action is played at all. While it was widely known that FP does not converge rapidly, the strength of our results is still somewhat surprising. They do not depend on the choice of a metric for comparing probability distributions. Rather, we show that the empirical frequency of FP after an exponential number of rounds can be *arbitrarily far* from any Nash equilibrium for *any* reasonable metric. This

casts doubt on the plausibility of FP as an explanation of Nash equilibrium play.

2 Related Work

As mentioned above, FP does not converge in general. Shapley showed this using a variant of Rock-Paper-Scissors and argued further that “if fictitious play is to fail, the game must contain elements of both coordination and competition” [27, p. 24]. This statement is perfectly consistent with the fact that FP is guaranteed to converge for both constant-sum games [25] and identical interest games, i.e., games that are best-response equivalent (in mixed strategies) to a common payoff game [21]. Other classes of games where FP is known to converge include two-player games solvable by iterated elimination of strictly dominated strategies [22] and non-degenerate 2×2 games [18]. While the proof of Miyasawa [18] was initially thought to apply to the class of all 2×2 games, this was later shown to be false [19]. The result was recently extended to non-degenerate $2 \times n$ games [2]. Since every non-degenerate 2×2 game is best-response equivalent to either a constant-sum game or a common payoff game [21], the result of Miyasawa follows more easily by combining those of Robinson [25] and Monderer and Shapley [21].

Shapiro [26] has shown that in two-player constant-sum games the rate of convergence of fictitious play is at most $O(n^{-1/(r+s-2)})$, where n is the number of rounds and r and s are the numbers of actions available to the two players. Hidden in the asymptotic notation is an additional factor of $a \cdot 2^{r+s}$, where a is the largest absolute payoff value. Karlin [16] conjectured that the correct rate of convergence is $O(n^{-1/2})$, but to our knowledge no progress has been made toward this conjecture. A lower bound of $O(n^{-1/2})$ already holds for the Matching Pennies game [13]. It should be noted that all these bounds concern convergence of payoffs values, whereas we study convergence to equilibrium strategies.

Von Neumann [29] proposed a variant of FP and compared it to Dantzig’s Simplex method. Indeed, there are some interesting similarities between the two. Conitzer [8] recently studied the ability of FP to find approximate Nash equilibria. In addition to worst-case guarantees on the approximation ratio—which are rather weak—Conitzer showed that in random games a good approximation is typically achieved after a relatively small number of rounds. Similarly, the Simplex method is known to work very well in practice. As we show in this paper, FP also shares one of the major shortcomings of the Simplex method—its exponential worst-case running time.

Since FP is one of the earliest and simplest algorithms for learning in games, it inspired many of the algorithms that followed: the variant due to von Neumann, a similar procedure suggested by Bellman [1], improvements like smooth FP [11], the regret minimization paradigm [15], and a large number of specialized algorithms put forward by the artificial intelligence community (e.g., [23, 9]).

Despite its conceptual simplicity and the existence of much more sophisticated learning algorithms, FP continues to be employed successfully in the area of artificial intelligence. Recent examples include equilibrium computation in Poker [12] and in anonymous games with continuous player types [24], and learning in sequential auctions [30].

3 Preliminaries

An accepted way to model situations of strategic interaction is by means of a *normal-form game* (see, e.g., [17]). We will focus on games with two players.

A *two-player game* $\Gamma = (P, Q)$ is given by two matrices $P, Q \in \mathbb{R}^{m \times n}$ for positive integers m and n . Player 1, or the row player, has a set $A = \{1, \dots, m\}$ of actions corresponding to the rows of these matrices, player 2, the column player, a set $B = \{1, \dots, n\}$ of actions corresponding to the columns. To distinguish between them, we usually denote actions of the row player by a^1, \dots, a^m and actions of the column player by b^1, \dots, b^n . Both players are assumed to simultaneously choose one of their actions. For the resulting action profile $(i, j) \in A \times B$, they respectively obtain payoffs p_{ij} and q_{ij} .

A *strategy* of a player is a probability distribution $s \in \Delta(A)$ or $t \in \Delta(B)$ over his actions, i.e., a nonnegative vector $s \in \mathbb{R}^m$ or $t \in \mathbb{R}^n$ such that $\sum_i s_i = 1$ or $\sum_j t_j = 1$, respectively. In a slight abuse of notation, we write p_{st} and q_{st} for the expected payoff of players 1 and 2 given a strategy profile $(s, t) \in \Delta(A) \times \Delta(B)$. A strategy is called *pure* if it selects one action with probability one, and the set of pure strategies can be identified in a natural way with the set of actions.

A two-player game is called a *constant-sum game* if $p_{ij} + q_{ij} = p_{i'j} + q_{i'j'}$ for all $i, i' \in A$ and $j, j' \in B$. Since all results in this paper hold invariably under positive affine transformations of the payoffs, such games can conveniently be represented by a single matrix P containing the payoffs of player 1; player 2 is then assumed to minimize the values in P . A constant-sum game is further called *symmetric* if P is a skew-symmetric matrix. In symmetric games, both players have the same set of actions, and we usually denote these actions by a^1, a^2, \dots, a^m . A game is a *common payoff game* if $p_{ij} = q_{ij}$ for all $i \in A$ and $j \in B$. Finally, a game is *non-degenerate* if for each strategy, the number of best responses of the other player is at most the support size of that strategy, i.e., the number of actions played with positive probability.

An action $i \in A$ of player 1 is said to *strictly dominate* another action $i' \in A$ if it provides a higher payoff for every action of player 2, i.e., if for all $j \in B$, $p_{ij} > p_{i'j}$. Dominance among actions of player 2 is defined analogously. A game is then called *solvable via iterated strict dominance* if strictly dominated actions can be removed iteratively such that exactly one action remains for each player.

A pair (s, t) of strategies is called a *Nash equilibrium* if the two strategies are best responses to each other, i.e., if $p_{st} \geq p_{it}$ for every $i \in A$ and $q_{st} \geq q_{sj}$ for every $j \in B$. A Nash equilibrium is *quasi-strict* if actions played

with positive probability yield strictly more payoff than actions played with probability zero. By the minimax theorem [28], every Nash equilibrium (s, t) of a constant-sum game satisfies $\min_j \sum_i p_{ij} s_i = \max_i \sum_j p_{ij} t_j = \omega$ for some $\omega \in \mathbb{R}$, also called the value of the game.

Fictitious play (FP) was originally introduced to approximate the value of constant-sum games, and has subsequently been studied in terms of its convergence to Nash equilibrium in more general classes of games. It proceeds in rounds. In the first round, each player arbitrarily chooses one of his actions. In subsequent rounds, each player looks at the empirical frequency of play of his respective opponents in previous rounds, interprets it as a probability distribution, and myopically plays a pure best response against this distribution. Fix a game $\Gamma = (P, Q)$ with $P, Q \in \mathbb{R}^{m \times n}$. Denote by u_i and v_i the i th unit vector in \mathbb{R}^m and \mathbb{R}^n , respectively. Then, a *learning sequence of Γ* is a sequence $(x^0, y^0), (x^1, y^1), (x^2, y^2), \dots$ of pairs of non-negative vectors $(x^i, y^i) \in \mathbb{R}^m \times \mathbb{R}^n$ such that $x^0 = \mathbf{0}$, $y^0 = \mathbf{0}$, and for all $k \geq 0$,

$$\begin{aligned} x^{k+1} &= x^k + u_i && \text{where } i \text{ is the index of a maximum component of } Py^k \text{ and} \\ y^{k+1} &= y^k + v_j && \text{where } j \text{ is the index of a maximum component of } x^k Q. \end{aligned}$$

A learning sequence $(x^0, y^0), (x^1, y^1), (x^2, y^2), \dots$ of a game Γ is said to converge if for some Nash equilibrium s of Γ ,

$$\lim_{k \rightarrow \infty} \left(\frac{x^k}{k}, \frac{y^k}{k} \right) = s,$$

where both division and limit are to be interpreted component-wise. We then say that FP converges for Γ if every learning sequence of Γ converges to a Nash equilibrium.

An alternative definition of a learning sequence, in which players update their beliefs alternatingly instead of simultaneously, can be obtained by replacing $x^k Q$ by $x^{k+1} Q$ in the last condition above. Berger [3] distinguishes between simultaneous and alternating FP, and points out that Brown actually introduced the latter variant, while almost all subsequent work routinely uses the former. We henceforth concentrate on simultaneous FP, or simply FP, but note that with some additional work all of our results can be shown to hold for alternating FP as well.

4 Results

We now present several results concerning the convergence rate of FP. Taken together, they cover virtually all classes of games for which FP is known to converge.

	a^1	a^2	a^3
a^1	0	-1	$-\epsilon$
a^2	1	0	$-\epsilon$
a^3	ϵ	ϵ	0

Fig. 1 Symmetric constant-sum game used in the proof of Theorem 1. Player 1 chooses rows, player 2 chooses columns. Outcomes are denoted by the payoff of player 1.

Round i	(a^i, a^i)	Py^i
0	—	$(0, 0, 0)$
1	(a^1, a^1)	$(0, 1, 2^{-k})$
2	(a^2, a^2)	$(-1, 1, 2^{-k}2)$
3	(a^2, a^2)	$(-2, 1, 2^{-k}3)$
	\vdots	\vdots
2^k	(a^2, a^2)	$(-2^k + 1, 1, 1)$

Table 1 A learning sequence of the game depicted in Figure 1, where $\epsilon = 2^{-k}$

4.1 Symmetric Constant-Sum Games and Games Solvable by Iterated Strict Dominance

Let us first consider games with arbitrary payoffs. Our first result concerns two large classes of games where FP is guaranteed to converge: constant-sum games and games solvable by iterated strict dominance.

Theorem 1 *In symmetric two-player constant-sum games, FP may require exponentially many rounds (in the size of the representation of the game) before an equilibrium action is eventually played. This holds even for games solvable via iterated strict dominance.*

Proof Consider the symmetric two-player constant-sum game $\Gamma = (P, Q)$ with payoff matrix P for player 1 as shown in Figure 1, where $0 < \epsilon < 1$. It is readily appreciated that (a^3, a^3) is the only Nash equilibrium of this game, as it is the only action profile that remains after iterated elimination of strictly dominated actions. Consider an arbitrary integer $k > 1$. We show that for $\epsilon = 2^{-k}$, FP may take 2^k rounds before either player plays action a^3 . Since the game can clearly be encoded using $O(k)$ bits in this case, the theorem follows.

Let FP start with both players choosing action a^1 . Since the game is symmetric, we can assume the actions for each step of the learning sequence to be identical for both players. After the first round $Py^1 = (0, 1, 2^{-k})$, and both players will play a^2 in round 2. We claim that they will continue to do so at least until round 2^k . To see this, observe that for all i with $1 \leq i < 2^k$, we have $Py^i = (-i + 1, 1, 2^{-k}i)$. As $2^{-k}i < 1$, both players will choose a^2 round $i + 1$. Table 1 summarizes this development. It follows that the action sequence

$$(a^1, a^1) \underbrace{(a^2, a^2), \dots, (a^2, a^2)}_{2^k - 1 \text{ times}}$$

gives rise to a learning sequence that is exponentially long in k and in which no equilibrium action is played. \square

This result is tight in the sense that FP converges very quickly in symmetric 2×2 games. Up to renaming of actions, every such game can be described by a matrix

$$\begin{array}{cc} & \begin{array}{cc} a^1 & a^2 \end{array} \\ \begin{array}{c} a^1 \\ a^2 \end{array} & \begin{bmatrix} 0 & -\alpha \\ \alpha & 0 \end{bmatrix} \end{array}$$

for some $\alpha \geq 0$. If $\alpha = 0$, every strategy profile is a Nash equilibrium. Otherwise, action a^1 is strictly dominated for both players, and both players will play the equilibrium action a^2 from round 2 onwards.

4.2 Non-Degenerate $2 \times n$ Games and Identical Interest Games

Another class of games where FP is guaranteed to converge are non-degenerate $2 \times n$ games. We again obtain a strong negative result concerning the convergence rate of FP, which also applies to games with identical interests.

Theorem 2 *In non-degenerate 2×3 games, FP may require exponentially many rounds (in the size of the representation of the game) before an equilibrium action is eventually played. This holds even for games with identical interests.*

Proof Consider the 2×3 game $\Gamma = (P, Q)$ shown in Figure 2, where $0 < \epsilon < 1$. It is easily verified that Γ is non-degenerate and that the players have identical interests. The action profile (a^2, b^3) is the only action profile that remains after iterated elimination of strictly dominated actions, and thus the only Nash equilibrium of the game.

Now consider an integer $k > 1$. We show that for $\epsilon = 2^{-k}$, FP may take 2^k rounds before actions a^2 or b^3 are played. Since in this case the game can clearly be encoded using $O(k)$ bits, the theorem follows.

Let FP start with both players choosing action a^1 . Then, $Py^1 = (1, 0)$ and $x^1Q = (1, 2, 0)$. Accordingly, in the second round, the row player will choose a^1 , and the column player b^2 . Hence, $Py^2 = (3, 2 + 2^{-k})$ and $x^2Q = (2, 4, 0)$. Hereafter, for at least another $2^k - 1$ rounds, the players will choose the same actions as in round 2, because for all i with $2 \leq i \leq 2^k$, $x^iQ = (i, 2i, 0)$, $Py^i = (2i - 1, 2i - 1 + 2^{-k}(i - 1))$, and $2i - 1 > 2i - 1 + 2^{-k}(i - 1)$. Accordingly, the sequence of pairs of actions

$$(a^1, b^1) \underbrace{(a^1, b^2), \dots, (a^1, b^2)}_{2^k \text{ times}},$$

	b^1	b^2	b^3
a^1	(1, 1)	(2, 2)	(0, 0)
a^2	(0, 0)	$(2 + \epsilon, 2 + \epsilon)$	(3, 3)

Fig. 2 Non-degenerate two-player game with identical interests used in the proof of Theorem 2. Outcomes are denoted by a pair of payoffs for the two players.

Round i	(a^i, b^i)	Py^i	$x^i Q$
0	—	(0, 0)	(0, 0, 0)
1	(a^1, b^1)	(1, 0)	(1, 2, 0)
2	(a^1, b^2)	$(3, 2 + 2^{-k})$	(2, 4, 0)
3	(a^1, b^2)	$(5, 4 + 2^{-k}2)$	(3, 6, 0)
	\vdots	\vdots	\vdots
2^k	(a^1, b^2)	$(2^{k+1} - 1, 2^{k+1} - 1 - 2^{-k})$	$(2^k, 2^{k+1}, 0)$

Table 2 A learning sequence of the game shown in Figure 2, where $\epsilon = 2^{-k}$

which contains no equilibrium actions, gives rise to a learning sequence that is exponentially long in k . Figure 2 illustrates both sequences. \square

This result is again tight: in any 2×2 game, one of the players must always play an equilibrium action almost immediately. Indeed, given that the initial action profile is not itself an equilibrium, one of the players plays his second action in the following round. But what about the other player? By looking at the subgame of the game in Figure 2 induced by actions $\{a^1, a^2\}$ and $\{b^1, b^2\}$, and at the learning sequence used to obtain Theorem 2, we find that it might still take exponentially many rounds for *one* of the two players until he plays an equilibrium action for the first time.

Theorem 2 also applies to potential games [20], which form a superclass of games with identical interests. For the given ordering of its actions, the game of Figure 2 further has strategic complementarities and diminishing returns,¹ which implies results analogous to Theorem 2 for classes of games in which convergence of FP was respectively claimed by Hahn [14]² and shown by Berger [4].

4.3 Games with Constant Payoffs

The proofs of the previous two theorems crucially rely on exponentially small payoffs, so one may wonder if similar results can still be obtained if additional constraints are imposed on the payoffs. While this is certainly not the case for games where both the payoffs and the number of actions are constant, we find

¹ A two-player game with totally ordered sets of actions is said to have strategic complementarities if the advantage of switching to a higher action, according to the ordering, increases when the opponent chooses a higher action, and diminishing returns if the advantage of increasing one's action is decreasing.

² The proof of this claim later turned out to be flawed [5].

	a^1	a^2	a^3	a^4	a^5	a^6	a^7	a^8	a^9
a^1	0	-1	-1	-1	-1	0	0	0	-1
a^2	1	0	0	0	0	-1	1	1	1
a^3	1	0	0	0	0	0	-1	1	1
a^4	1	0	0	0	0	0	0	-1	1
a^5	1	0	0	0	0	0	0	0	-1
a^6	0	1	0	0	0	0	0	0	0
a^7	0	-1	1	0	0	0	0	0	0
a^8	0	-1	-1	1	0	0	0	0	0
a^9	1	-1	-1	-1	1	0	0	0	0

Fig. 3 Symmetric constant-sum game Γ^4 used in the proof of Theorem 3. The game possesses a quasi-strict equilibrium (s^4, s^4) with $s^4 = (0, 0, 0, 0, 0, \frac{8}{15}, \frac{4}{15}, \frac{2}{15}, \frac{1}{15})$.

that a somewhat weaker version of Theorem 1 holds for games with constant payoffs, and in particular for symmetric constant-sum win-lose-tie games, i.e., symmetric constant-sum games with payoffs in $\{-1, 0, 1\}$.

For each integer k we define a symmetric constant-sum game Γ^k with a unique (mixed) Nash equilibrium and show that FP may take a number of rounds exponential in k before an equilibrium action is played. In contrast to the previous result, however, this result not only assumes a worst-case initial action profile, but also a worst-case learning sequence.

Theorem 3 *In symmetric constant-sum win-lose-tie games, FP may require exponentially many rounds (in the size of the game) before an equilibrium action is eventually played.*

Proof Fix an integer $k > 1$. We construct a symmetric constant-sum win-lose-tie game $\Gamma^k = (P^k, Q^k)$ with a $(2k+1) \times (2k+1)$ payoff matrix $P^k = (p_{ij}^k)$ for player 1 such that for all i, j with $1 \leq j \leq i \leq 2k+1$,

$$p_{ij}^k = \begin{cases} 1 & \text{if } j = 1 \text{ and } 2 \leq i \leq k+1, \text{ or} \\ & \text{if } j = 1 \text{ and } i = 2k+1, \text{ or} \\ & \text{if } j \neq 1 \text{ and } i = j+k, \\ -1 & \text{if } j \neq 1 \text{ and } i > j+k, \\ 0 & \text{otherwise.} \end{cases}$$

For $i < j$, let $p_{ij}^k = -p_{ji}^k$. Thus Γ^k clearly is a symmetric constant-sum game. To illustrate the definition, Γ^4 is shown in Figure 3.

Further define, for each k , a strategy profile (s^k, s^k) of Γ^k such that for all i with $1 \leq i \leq 2k+1$,

$$s_i^k = \begin{cases} 2^{2k+1-i}/(2^k-1) & \text{if } i > k+1, \\ 0 & \text{otherwise.} \end{cases}$$

Round i	(a^j, a^j)	$P^4 y^i$
0	—	(0, 0, 0, 0, 0, 0, 0, 0)
1	(a^1, a^1)	(0, 1, 1, 1, 1, 0, 0, 1)
2	(a^2, a^2)	(-1, 1, 1, 1, 1, 1, -1, 0)
3	(a^3, a^3)	(-2, 1, 1, 1, 1, 1, 0, -2, -1)
4	(a^3, a^3)	(-3, 1, 1, 1, 1, 1, 1, -3, -2)
5	(a^4, a^4)	(-4, 1, 1, 1, 1, 1, 1, -2, -3)
6	(a^4, a^4)	(-5, 1, 1, 1, 1, 1, 1, -1, -4)
7	(a^4, a^4)	(-6, 1, 1, 1, 1, 1, 1, 0, -5)
8	(a^4, a^4)	(-7, 1, 1, 1, 1, 1, 1, 1, -6)
9	(a^5, a^5)	(-8, 1, 1, 1, 1, 1, 1, 1, -5)
10	(a^5, a^5)	(-9, 1, 1, 1, 1, 1, 1, 1, -4)
11	(a^5, a^5)	(-10, 1, 1, 1, 1, 1, 1, 1, -3)
12	(a^5, a^5)	(-11, 1, 1, 1, 1, 1, 1, 1, -2)
13	(a^5, a^5)	(-12, 1, 1, 1, 1, 1, 1, 1, -1)
14	(a^5, a^5)	(-13, 1, 1, 1, 1, 1, 1, 1, 0)
15	(a^5, a^5)	(-14, 1, 1, 1, 1, 1, 1, 1, 1)
16	(a^5, a^5)	(-15, 1, 1, 1, 1, 1, 1, 1, 2)

Table 3 A learning sequence of the game Γ^4 shown in Figure 3

It is not hard to see that (s^k, s^k) is a quasi-strict equilibrium of Γ^k . Moreover, since Γ^k is both a symmetric and a constant-sum game, the support of any equilibrium strategy of Γ^k is contained in that of s^k (cf. [6]). We will now show that, when starting with (a^1, a^1) , FP in Γ^k may take at least 2^k rounds before an equilibrium action is played for the first time.

Consider the sequence a_1, \dots, a_{2^k} with $a_j = a^{1+\lceil \log_2 j \rceil}$ for all j with $1 \leq j \leq 2^k$, i.e., the sequence

$$a^1, a^2, a^3, a^3, \dots, \underbrace{a^i, \dots, a^i}_{2^{i-2} \text{ times}}, \dots, \underbrace{a^{k+1}, \dots, a^{k+1}}_{2^{k-1} \text{ times}}.$$

The length of this sequence is clearly exponential in k . Further define vectors x^0, \dots, x^{2^k} of dimension $2k+1$ such that $x^0 = \mathbf{0}$, and for i with $1 \leq j \leq 2k+1$, $x^{j+1} = x^j + u_i$ when $a_{j+1} = i$.

We now claim that $(x^0, x^0), \dots, (x^{2^k}, x^{2^k})$ is a learning sequence of Γ^k , i.e., that $j+1$ is the index of a maximal component of both $P^k y^j$ and $x^j Q^k$. Table 3 shows the development of this sequence for $k=4$.

By symmetry of Γ^k it suffices to prove the claim for $P^k y^j$. After the first round, we have for all i with $1 \leq i \leq 2k+1$,

$$(P^k y^1)_i = \begin{cases} 1 & \text{if } 1 < i \leq k+1, \\ 0 & \text{otherwise.} \end{cases}$$

Furthermore, since $\{a_2, \dots, a_{2^k}\} \subseteq \{a^2, \dots, a^{k+1}\}$, we have that $(P^k y^j)_i = 1$ for all i with $1 < i \leq k+1$ and all j with $1 < j \leq 2^k$. It, therefore, suffices to show that $(P^k y^j)_i$ for all i with $i = 1$ or $k+1 < i < 2k+1$ and all j

with $1 < j \leq 2^k$. Since, $p_{1i} = -1$ for all i with $1 < i \leq k+1$, the former is obvious. For the latter, it can be shown by a straightforward if somewhat tedious induction on j that for all i with $1 \leq i < k$ and all j with $1 < j \leq 2^k$,

$$(P^k y^j)_{i+k+1} = \begin{cases} 1-j & \text{if } j \leq 2^{i-1}, \\ 1+j-2^i & \text{if } 2^{i-1} < j \leq 2^i, \\ 1 & \text{otherwise, and} \end{cases}$$

$$(P^k y^j)_{2k+1} = \begin{cases} 2-j & \text{if } j \leq 2^{k-1}, \\ 2+j-2^k & \text{otherwise.} \end{cases}$$

It follows that $(P^k y^j)_i \leq 1$ for all i with $1 \leq i \leq 2k+1$ and all j with $1 \leq j < 2^k$, thus proving the claim. \square

5 Conclusion

We have studied the rate of convergence of fictitious play, and obtained mostly negative results: for almost all of the classes of games where FP is known to converge, it may take an exponential number of rounds before some equilibrium action is eventually played. These results hold already for games with very few actions, given that one of the payoffs is exponentially small compared to the others. Slightly weaker results can still be salvaged for symmetric constant-sum games and games solvable by iterated strict dominance, even if payoffs are in the set $\{-1, 0, 1\}$. It is an open question whether this result can be strengthened to match that for games with arbitrary payoffs, and whether a similar result can be obtained for the classes of games covered by Theorem 2, i.e., for potential games and identical interest games.

While it was known that fictitious play does not converge rapidly, the strength of our results is still somewhat surprising. They do not depend on the choice of a metric for comparing probability distributions. Rather, the empirical frequency of FP after an exponential number of rounds can be *arbitrarily far* from any Nash equilibrium for *any* reasonable metric. This casts doubt on the plausibility of fictitious play as an explanation of Nash equilibrium play.

Acknowledgements This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/3-2, BR 2312/3-3, BR 2312/7-1 and FI 1664/1-1. We would like to thank Vincent Conitzer, Paul Goldberg, Peter Bro Miltersen, and Troels Bjerre Sørensen for valuable discussions.

References

1. R. Bellman. On a new iterative algorithm for finding the solutions of games and linear programming problems. Research Memorandum P-473, The RAND Corporation, 1953.

2. U. Berger. Fictitious play in $2 \times n$ games. *Journal of Economic Theory*, 120:139–154, 2005.
3. U. Berger. Brown’s original fictitious play. *Journal of Economic Theory*, 135:572–578, 2007.
4. U. Berger. Learning in games with strategic complementarities revisited. *Journal of Economic Theory*, 143:292–301, 2008.
5. U. Berger. The convergence of fictitious play in games with strategic complementarities: A comment. MPRA Paper No. 20241, Munich Personal RePEc Archive, 2009.
6. F. Brandt and F. Fischer. On the hardness and existence of quasi-strict equilibria. In B. Monien and U.-P. Schroeder, editors, *Proceedings of the 1st International Symposium on Algorithmic Game Theory (SAGT)*, volume 4997 of *Lecture Notes in Computer Science (LNCS)*, pages 291–302. Springer-Verlag, 2008.
7. G. W. Brown. Iterative solutions of games by fictitious play. In T. C. Koopmans, editor, *Activity Analysis of Production and Allocation*, pages 374–376. John Wiley and Sons, Inc., 1951.
8. V. Conitzer. Approximation guarantees for fictitious play. In *Proceedings of the 47th Annual Allerton Conference on Communication, Control, and Computing*, pages 636–643, 2009.
9. V. Conitzer and T. Sandholm. AWESOME: A general multiagent learning algorithm that converges in self-play and learns a best response against stationary opponents. *Machine Learning*, 67(1–2):23–43, 2007.
10. G. B. Dantzig. A proof of the equivalence of the programming problem and the game problem. In T. C. Koopmans, editor, *Activity Analysis of Production and Allocation*, pages 330–335. John Wiley & Sons Inc., 1951.
11. D. Fudenberg and D. Levine. Consistency and cautious fictitious play. *Journal of Economic Dynamics and Control*, 19:1065–1089, 1995.
12. S. Ganzfried and T. Sandholm. Computing an approximate jam/fold equilibrium for 3-player no-limit texas hold’em tournaments. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 919–925, 2008.
13. S. Gjerstad. The rate of convergence of continuous fictitious play. *Journal of Economic Theory*, 7:161–178, 1996.
14. S. Hahn. The convergence of fictitious play in 3×3 games with strategic complementarities. *Economic Letters*, 64:57–60, 1999.
15. J. Hannan. Approximation to Bayes risk in repeated plays. In M. Dresher, A. W. Tucker, and P. Wolfe, editors, *Contributions to the Theory of Games*, volume 3, pages 97–139. Princeton University Press, 1957.
16. S. Karlin. *Mathematical Methods and Theory in Games*, volume 1–2. Addison-Wesley, 1959.
17. R. D. Luce and H. Raiffa. *Games and Decisions: Introduction and Critical Survey*. John Wiley & Sons Inc., 1957.
18. K. Miyasawa. On the convergence of the learning process in a 2×2 nonzero sum two-person game. Research Memorandum 33, Econometric Research Program, Princeton University, 1961.

19. D. Monderer and A. Sela. A 2×2 game without the fictitious play property. *Games and Economic Behavior*, 14:144–148, 1996.
20. D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
21. D. Monderer and L. S. Shapley. Fictitious play property for games with identical interests. *Journal of Economic Theory*, 68:258–265, 1996.
22. J. H. Nachbar. “Evolutionary” selection dynamics in games: Convergence and limit properties. *International Journal of Game Theory*, 19:59–89, 1990.
23. R. Powers and Y. Shoham. New criteria and a new algorithm for learning in multi-agent systems. In *Advances in Neural Information Processing Systems 17*, pages 1089–1096. MIT Press, 2004.
24. Z. Rabinovich, E. Gerding, M. Polukarov, and N. R. Jennings. Generalised fictitious play for a continuum of anonymous players. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 245–250, 2009.
25. J. Robinson. An iterative method of solving a game. *Annals of Mathematics*, 54(2):296–301, 1951.
26. H. Shapiro. Note on a computation model in the theory of games. *Communications on Pure and Applied Mathematics*, 11:587–593, 1958.
27. L. Shapley. Some topics in two-person games. In M. Dresher, L. S. Shapley, and A. W. Tucker, editors, *Advances in Game Theory*, volume 52 of *Annals of Mathematics Studies*, pages 1–29. Princeton University Press, 1964.
28. J. von Neumann. Zur Theorie der Gesellschaftspiele. *Mathematische Annalen*, 100:295–320, 1928.
29. J. von Neumann. A numerical method to determine optimum strategy. *Naval Research Logistics Quarterly*, 1(2):109–115, 1954.
30. W. Zhu and P. R. Wurman. Structural leverage and fictitious play in sequential auctions. In *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI)*, pages 385–390, 2002.