

MIT Open Access Articles

A Counter-example to Karlin's Strong Conjecture for Fictitious Play

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Daskalakis, Constantinos, and Qinxuan Pan. "A Counter-Example to Karlin's Strong Conjecture for Fictitious Play." 2014 IEEE 55th Annual Symposium on Foundations of Computer Science (October 2014).

As Published: <http://dx.doi.org/10.1109/FOCS.2014.10>

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <http://hdl.handle.net/1721.1/99979>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



A Counter-Example to Karlin’s Strong Conjecture for Fictitious Play

Constantinos Daskalakis[†]
EECS, MIT
costis@mit.edu

Qinxuan Pan[‡]
EECS, MIT
qinxuan@mit.edu

Abstract—Fictitious play is a natural dynamic for equilibrium play in zero-sum games, proposed by Brown [6], and shown to converge by Robinson [33]. Samuel Karlin conjectured in 1959 that fictitious play converges at rate $O(t^{-\frac{1}{2}})$ with respect to the number of steps t . We disprove this conjecture by showing that, when the payoff matrix of the row player is the $n \times n$ identity matrix, fictitious play may converge (for some tie-breaking) at rate as slow as $\Omega(t^{-\frac{1}{n}})$.

I. INTRODUCTION

Von Neumann’s MinMax theorem for two-person zero-sum games marked the birth of Game Theory [36], and is intimately related to the development of linear programming. Given a payoff matrix A , whose ij -th entry specifies how much the column player playing j pays the row player playing i , the theorem states that

$$\max_x \min_y x^T A y = \min_y \max_x x^T A y,$$

where x, y range over randomized/mixed strategies for the row and column player respectively. In other words, there exists a unique value $z \in \mathbb{R}$ and a pair of mixed strategies \hat{x} and \hat{y} such that:

$$\min_y \hat{x}^T A y = z = \max_x x^T A \hat{y}. \quad (1)$$

Dantzig and von Neumann observed that the MinMax theorem is implied by strong linear programming duality [10], [2]. Dantzig also provided a candidate construction for the opposite implication [10], and this was also established some decades later [1].

Ultimately, the MinMax theorem provides a very sharp prediction in two-player zero-sum games. It shows that there is a unique value z and a pair of strategies \hat{x} and \hat{y} such that, by playing \hat{x} the row player can guarantee himself expected payoff of z regardless of what strategy the column player adopts, and such that, by playing \hat{y} , the column player can guarantee herself expected payoff of $-z$ regardless of what strategy the row player adopts. In particular, (\hat{x}, \hat{y}) comprise a Nash equilibrium of the

game, with expected payoff z for the row player and $-z$ for the column. Moreover, \hat{x}, \hat{y} and z can be computed in polynomial time with linear programming. This type of crisp prediction is rather rare in Game Theory. According to Aumann, zero-sum games are “one of the few areas in game theory, and indeed in the social sciences, where a fairly sharp, unique prediction is made” [3].

Shortly after the proof of the MinMax theorem and the development of linear programming, G. W. Brown proposed *fictitious play* as an iterative procedure for solving a zero-sum game, or equivalently a linear program [6], [7]. The procedure proceeds in steps in which players choose a pure strategy best response to their opponent’s empirical mixed strategy up until that step. Let us describe it a bit more formally (we focus on the *simultaneous* version, but our results also hold for the *asynchronous* version, where the players’ moves alternate): At every step t , the row player chooses some row i_t and the column player chooses some column j_t . At $t = 1$, the choices are arbitrary. At $t + 1 > 1$, the players calculate the *empirical mixed strategies* of their opponents in previous steps, namely¹

$$x(t) = \frac{1}{t} \sum_{\tau \leq t} e_{i_\tau},$$

$$y(t) = \frac{1}{t} \sum_{\tau \leq t} e_{j_\tau}.$$

Then, the row player chooses an arbitrary best response i_{t+1} to $y(t)$ and the column player chooses an arbitrary best response j_{t+1} to $x(t)$, namely

$$i_{t+1} \in \arg \max_i \{e_i^T A y(t)\},$$

$$j_{t+1} \in \arg \min_j \{x(t)^T A e_j\}. \quad (2)$$

The procedure may be viewed as a natural way through which two players could interact in a repeated game with stage game $(A, -A)$. The question is whether the sequence $(x(t), y(t))_t$ converges to something meaningful.

[†]Supported by a Sloan Foundation fellowship, a Microsoft Research faculty fellowship and NSF Award CCF-0953960 (CAREER) and CCF-1101491.

[‡]Supported by ONR grant N00014-12-1-0999.

¹We use e_i to denote the column vector with i -th component 1 and all other components 0. The dimension of e_i is always implied by the context; it is m when describing row player strategies and n when describing column player strategies.

In an elegant paper shortly after Brown’s, Robinson showed that the average payoffs of the players in fictitious play converge to the value of the game [33]. In particular, it was shown that

$$\begin{aligned} f_A(x(t), y(t)) &= \max_i e_i^\top Ay(t) - \min_j x(t)Ae_j \\ &\rightarrow 0, \text{ as } t \rightarrow \infty. \end{aligned}$$

Hence, because $\min_j x(t)Ae_j \leq x(t)^\top Ay(t) \leq \max_i e_i^\top Ay(t)$ and $\min_j x(t)Ae_j \leq z \leq \max_i e_i^\top Ay(t)$, it follows that all three quantities converge to the value of the game z .

Robinson’s proof is an elegant induction argument, which eliminates one row or one column of A at a time. Unraveling the induction, one can also deduce the following bound on the convergence rate of the procedure:

$$f_A(x(t), y(t)) = O(t^{-\frac{1}{m+n-2}}),$$

which appears rather slow, compared to the convergence rate of $O(t^{-\frac{1}{2}})$ that is typically achieved by no-regret learning algorithms [13], [26], [8], and the improved convergence rate of $O(\frac{\log t}{t})$ of some no-regret learning algorithms, obtained recently [11], [32]. Indeed, about ten years after Robinson’s proof and five decades ago, Samuel Karlin conjectured that the convergence rate of fictitious play should be $O(t^{-\frac{1}{2}})$, namely

Conjecture 1 ([23]). *Fictitious play converges at rate $O(t^{-\frac{1}{2}})$ in all games.*

There is some evidence supporting a convergence rate of $O(t^{-\frac{1}{2}})$. As pointed out earlier, a convergence rate of $O(t^{-\frac{1}{2}})$ is quite common with dynamics that are known to converge. Indeed, a close relative of fictitious play, *follow the perturbed leader*, is known to achieve convergence rate of $O(t^{-\frac{1}{2}})$ [8]. Also, a continuous time version of fictitious play has been shown to converge in time $O(t^{-1})$ [19]. Despite this evidence and the apparent simplicity of fictitious play, the convergence rate from Robinson’s proof has remained the state-of-the-art. Our main result is a counter-example, disproving Karlin’s conjecture. If I_n is the $n \times n$ identity matrix, we show the following:

Theorem 1. *For every $n \geq 2$, fictitious play for I_n may converge at rate $\Theta(t^{-\frac{1}{n}})$, if ties are broken arbitrarily.*

Our counter-example, provided in Section III, constructs a valid execution of fictitious play for I_n such that the empirical mixed strategies $x(t), y(t)$ of players satisfy

$$\begin{aligned} f_{I_n}(x(t), y(t)) &= \max_i e_i^\top y(t) - \min_j x(t)e_j \\ &= \Theta(t^{-\frac{1}{n}}). \end{aligned}$$

Remark 1. *It is crucial for our construction that ties in choosing a best response in (2) can be broken arbitrarily at each step. This is allowed in Karlin’s formulation of the conjecture. To distinguish this case from when ties are broken in some consistent way or randomly, we will call Karlin’s conjecture with arbitrary tie-breaking Karlin’s strong conjecture, while that with lexicographic or random tie-breaking Karlin’s weak conjecture. With this terminology, Theorem 1 disproves Karlin’s strong conjecture.*

Interestingly, like Robinson’s upper bound argument, our lower bound also works by induction. We show that slow fictitious play executions for I_2 can be folded inside fictitious play executions for I_3 , etc, leading to an exponentially slow convergence rate for fictitious play in I_n . More intuition about the construction is provided in Section II, and the complete details can be found in Section III.

While outperformed by modern learning algorithms [8], because of its simplicity, fictitious play was thought to provide a convincing explanation of Nash equilibrium play in zero-sum games. According to Luce and Raiffa “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” [27]. In this light, our counterexample sheds doubt on the plausibility of fictitious play in explaining Nash equilibrium behavior. Given our counterexample, it is important to investigate whether fictitious play in random payoff zero-sum games satisfies Karlin’s conjecture, or whether some choice of tie-breaking rule in the definition of fictitious play makes it satisfy Karlin’s conjecture for all zero-sum games. We did perform preliminary simulations of fictitious play with random tie-breaking on our lower bounding instances, as well as on zero-sum games with i.i.d. uniform $[0, 1]$ entries, and they suggest a quadratic rate of convergence. We leave a rigorous study of these important questions for future work.

a) Related Work: Fictitious play is one of the most well-studied dynamics in Game Theory, and we cannot do it justice in a short exposition. We only mention a few highlights here. As we have already mentioned, it was proposed by Brown, in a technical report at RAND corporation [6], and was shown to converge in two-person zero-sum games by Robinson [33]. Miyakawa extended Robinson’s results to two-player

games with two strategies per player assuming a specific tie-breaking rule [29], while Shapley constructed a two-player three-strategy game where fictitious play does not converge [35]. Since then a lot of research has been devoted to understanding classes of games where fictitious play converges (e.g. [28], [31], [21], [17], [34], [4]) or does not converge (e.g. [22], [15], [30], [12], [25]). Surveys can be found in [24], [14], [20]. Other work has studied the approximation performance of fictitious play when used as a heuristic to find approximate Nash equilibria [9], [16].

In two-person zero-sum games, a convergence rate of $O(t^{-\frac{1}{m+n-2}})$ is implied by Robinson's proof, and S. Karlin conjectured that the convergence rate should be $O(t^{-\frac{1}{2}})$, which would match what we know is achievable by no-regret learning algorithms [8]. Indeed, Harris showed that a continuous analog of fictitious play converges in time $O(t^{-1})$ [19]. On the other hand, it is shown in [5] that it may take an exponential number of steps (in the size of the representation of the game) before any Nash equilibrium action is played by the players in fictitious play. However, this is not incompatible with Karlin's conjecture, since the payoffs may nevertheless still converge at rate $O(t^{-\frac{1}{2}})$. In fact, it is not even prohibited by [5] that the empirical strategies converge to Nash equilibrium strategies at rate $O(t^{-\frac{1}{2}})$.

As fictitious play is one of the simplest and most natural dynamics for learning in games it is widely used in applications, and has inspired several algorithms for learning and optimization, including von Neumann's variant of fictitious play for linear programming [37], the regret minimization paradigm [18], and lots of specialized algorithms in AI. See [5] for a survey.

II. PRELIMINARIES

Basic Definitions: A two-player zero-sum game can be represented by an $m \times n$ payoff matrix $A = (a_{ij})$, where m and n are the numbers of *pure strategies* for the row player and the column player, respectively. The game is played when, simultaneously, the row player chooses one of his m strategies, and the column player chooses one of her n strategies. If the row player chooses strategy i and the column player chooses strategy j , then the row player receives a_{ij} from the column player.

The players can randomize their choices of strategies. A *mixed strategy* for the row player is an m -vector x , where $x_i \geq 0$ and $\sum_i x_i = 1$. Similarly, a mixed strategy for the column player is an n -vector y , where $y_j \geq 0$ and $\sum_j y_j = 1$. When the players adopt those mixed strategies, the row player receives $x^T A y = \sum_{ij} a_{ij} x_i y_j$ in expectation from the column player.

A *min-max equilibrium*, or *Nash equilibrium*, of a zero-sum game A is a pair of mixed strategies \hat{x} for

the row player and \hat{y} for the column player such that Eq (1) is satisfied.

Dynamic: We already described fictitious play in Section I. We now introduce the notion of a dynamic as a formal way to describe a valid execution of fictitious play.

For a vector v , let $\min v$ and $\max v$ denote its minimal and maximal components. A dynamic as defined in the next paragraph is a special case of a vector system as defined in [33] that starts from the zero vectors.

Definition 1. A dynamic (U, V) for A is a sequence of n -dimensional row vectors $U(0), U(1), \dots$ and a sequence of m -dimensional column vectors $V(0), V(1), \dots$ such that²

$$\begin{aligned} U(0) &= [0, 0, \dots, 0]^T, \\ V(0) &= [0, 0, \dots, 0], \end{aligned}$$

and

$$\begin{aligned} U(t+1) &= U(t) + e_i^T A, \\ V(t+1) &= V(t) + A e_j, \end{aligned}$$

where i and j satisfy the conditions

$$\begin{aligned} V_i(t) &= \max V(t), \\ U_j(t) &= \min U(t). \end{aligned}$$

Just like there can be multiple valid executions of fictitious play for a matrix A , due to tie-breakings, there can be multiple possible dynamics for A . In fact, a dynamic for A corresponds uniquely to an execution of fictitious play for A , if we identify $U(t)$ and $V(t)$ with $x(t)^T A$ and $t A y(t)$, respectively. (Recall from Section I that $x(t)$ and $y(t)$ are the empirical mixed strategies of the two players for the first t steps.)

In terms of dynamics, Robinson's argument [33] implies the following: If (U, V) is a dynamic for an m by n matrix A , then

$$\frac{\max V(t) - \min U(t)}{t} = O(t^{-\frac{1}{m+n-2}}).$$

Karlin's conjecture [23] amounts to the following: If (U, V) is a dynamic for a matrix A , then

$$\frac{\max V(t) - \min U(t)}{t} = O(t^{-\frac{1}{2}}).$$

Notice that in both equations above, the constant in $O(\cdot)$ may depend on A . Lastly, our construction implies that there exists a dynamic (U, V) for I_n such that

$$\frac{\max V(t) - \min U(t)}{t} = \Theta(t^{-\frac{1}{n}}),$$

where the constant in $O(\cdot)$ may depend on n .

²Any vector presented using rectangular brackets is a column vector by default, unless it is followed by a transpose sign T .

Outline of our Construction: First notice that, by Definition 1, a dynamic (U, V) for I_n satisfies

$$\begin{aligned} U(0) &= [0, 0, \dots, 0]^T, \\ V(0) &= [0, 0, \dots, 0], \end{aligned}$$

and

$$\begin{aligned} U(t+1) &= U(t) + e_i^T, \\ V(t+1) &= V(t) + e_j, \end{aligned}$$

where i and j satisfy the conditions

$$\begin{aligned} V_i(t) &= \max V(t), \\ U_j(t) &= \min U(t). \end{aligned}$$

A special property of the dynamics for I_n is that permuting the n components of every vector in a dynamic for I_n by a common permutation σ results in another dynamic for I_n , because I_n stays the same when its rows and columns are both permuted by σ . This property allows us to combine many distinct cases in our main proof.

For $n = 2$, we can directly construct a dynamic for I_2 that converges at rate $\Theta(t^{-\frac{1}{2}})$, which we call the *main dynamic for I_2* (Figure 2 and Claim 3). At each step t , ties are simply broken by selecting the strategy that maximizes the ensuing gap $\max V(t) - \min U(t)$.

For $n = 3$, there is no obvious way to directly construct a dynamic for I_3 that converges at rate $\Theta(t^{-\frac{1}{3}})$. But, in the first three steps, it is easy to arrive at

$$\begin{aligned} U(3) &= [1, 1, 1]^T, \\ V(3) &= [0, 1, 2]. \end{aligned}$$

Aiming for an inductive construction, let's in fact assume that, for some P , we can arrive at

$$\begin{aligned} U(3P) &= [P, P, P]^T, \\ V(3P) &= [Q_1, Q_2, Q_3], \end{aligned}$$

where $Q_1 \leq Q_2 \leq Q_3$. For the next few steps, we let U increase only in its third component, and V only in its first two components. We can do this as long as the third component of V , i.e. Q_3 , remains its largest. Thus, we get to

$$\begin{aligned} U(3P+R) &= [P, P, P+R]^T, \\ V(3P+R) &= [Q_3, Q_3, Q_3]. \end{aligned}$$

The crucial component of our construction are the next steps, where we let U and V increase only their first two components, simulating a dynamic for the 2×2 subgame induced by the first two strategies of both players, i.e. I_2 . (We are able to do this as long as the third component of U , i.e. $P+R$, remains its largest.) Since U and V have equal first and second components at step $3P+R$, any initial portion of any dynamic (U', V') for I_2 can be copied, as long as the components of U' remain at most R . Indeed, if we do this, then for all

t the first two components of $U(3P+R+t)$ are P plus, respectively, the two components of $U'(t)$, and the first two components of $V(3P+R+t)$ are Q_3 plus, respectively, the two components of $V'(t)$.

For a dynamic (U', V') for I_2 , suppose that both components of $U'(t)$ are at most R , for all $t \leq t_0$, for some t_0 . It can be easily checked that, if we copy this dynamic in the first two components of our dynamic (U, V) for I_3 for t_0 steps, then the amount by which the gap for (U, V) increases, that is, from

$$\max V(3P+R) - \min U(3P+R)$$

to

$$\max V(3P+R+t_0) - \min U(3P+R+t_0),$$

is exactly the gap $\max V'(t_0) - \min U'(t_0)$ of (U', V') at t_0 .

We have two goals now. The first is to increase the gap for (U, V) as much as possible, and the second is to come back to the pattern we started from (that is, U has three equal components) so that we can apply the process again. To achieve our first goal, we want the gap $\max V'(t_0) - \min U'(t_0)$ to be as large as possible, subject to $\max U'(t_0) \leq R$. Naturally, we want (U', V') to be the main dynamic for I_2 , discussed earlier, as this achieves a rate of convergence of $\Theta(t^{-\frac{1}{2}})$. To achieve our second goal, we wish that $U'(t_0) = [R, R]^T$, so that $U(3P+R+t_0) = [P+R, P+R, P+R]^T$. Clearly, we must have $t_0 = 2R$ in this case. So, is it true that $U'(2R) = [R, R]^T$, if (U', V') is the main dynamic for I_2 ?

From (Figure 2/Claim 3), we see that there are indeed infinitely many T 's such that $U'(2T) = [T, T]^T$. However, this is not true for all T . Thus, we can't exactly take (U', V') to be the main dynamic for I_2 , but will need a padded version of it. Hence, we define the *padding dynamic for I_2* as in Figure 1/Claim 2, which reaches

$$\begin{aligned} U''(2k) &= [k, k]^T, \\ V''(2k) &= [k-1, k+1], \end{aligned}$$

for all k . The dynamic (U', V') that we copy into (U, V) first follows the padding dynamic for I_2 , and then the main dynamic for I_2 . By picking the appropriate moment of transition, we can ensure that (U', V') still converges at rate $\Theta(t^{-\frac{1}{2}})$, and $U'(2R) = [R, R]^T$.

Calculation shows that, if we repeat the process successively, the dynamic that will be obtained for I_3 converges at rate $\Theta(t^{-\frac{1}{3}})$. We call the resulting dynamic the *main dynamic for I_3* , and deal with $n = 4$ in similar fashion, etc, leading to our main theorem.

III. THE COUNTEREXAMPLE

In this section, we disprove Karlin's conjecture, by establishing the following.

Theorem 2. For every $n \geq 2$, there exists a dynamic for I_n such that for infinitely many T 's,

$$\max V(nT) - \min U(nT) = \Theta(T^{\frac{n-1}{n}}).$$

Proof of Theorem 2: Theorem 2 follows directly from Part 1 of the following Lemma (Part 2 is useful for showing Part 1 by induction):

Lemma 1. Part 1: For every $n \geq 2$, there exists a dynamic for I_n such that for infinitely many T 's,

$$U(nT) = [T, T, \dots, T]^T,$$

and

$$\max V(nT) - \min U(nT) = \Theta(T^{\frac{n-1}{n}}).$$

Part 2: For every $n \geq 2$ and $T \geq 1$, there exists a dynamic for I_n such that

$$U(nT) = [T, T, \dots, T]^T,$$

and

$$\max V(nT) - \min U(nT) = \Theta(T^{\frac{n-1}{n}}).$$

In either part, the constant hidden by $\Theta(\cdot)$ may depend on n , but not on T .

Proof of Lemma 1: We prove the lemma by induction on n . For each n , we prove Part 1 before Part 2.

Base case $n = 2$: We consider two dynamics for I_2 , which we call the *padding dynamics*. The first steps of the padding dynamics are illustrated on the left and on the right respectively of Figure 1. Notice that the strategy chosen by the row (respectively column) player at each step is exactly the index of the incremented component in U (respectively V).

We claim the following.

Claim 1. The dynamics shown in Figure 1 can be extended so that the dynamic on the left satisfies

$$\begin{aligned} U(2k) &= [k, k]^T, \\ V(2k) &= [k \pm 1, k \mp 1], \end{aligned} \quad (3)$$

for odd $k \geq 1$, while the dynamic on the right satisfies (3) for even $k \geq 2$. The choice of $+$ or $-$ depends on the parity of $\lceil \frac{k}{2} \rceil$.

Proof of Claim 1: To see the claim for the dynamic on the left, compare $U(t), V(t)$ at steps $t = 2$ and $t = 6$. The two components of $U(t)$ are equal, while the two components of $V(t)$ differ by 2. So, after exchanging the strategies $1 \leftrightarrow 2$, we can repeat the players' choices at Steps 3, 4, 5 and 6 in Steps 7, 8, 9 and 10 respectively to arrive at $U(10) = [5, 5]^T$ and $V(10) = [4, 6]$. And, we can continue the same way ad infinitum, which proves the claim for all odd k 's. Similar argument for the dynamic on the right proves for all even k 's. \square

By using either of the padding dynamics for I_2 and exchanging the components as necessary, we see the following:

Claim 2. For any $k \geq 1$, there exists a padding dynamic for I_2 such that

$$\begin{aligned} U(2k) &= [k, k]^T, \\ V(2k) &= [k - 1, k + 1]. \end{aligned}$$

Next, we define the *main dynamic* for I_2 , whose first steps are shown in Figure 2 in the appendix. We claim the following.

Claim 3. The dynamic given in Figure 2 can be extended so that it satisfies the following for all $k \geq 1$:

$$\begin{aligned} U(2k(2k-1)) &= [k(2k-1), k(2k-1)]^T, \\ V(2k(2k-1)) &= [(k \pm 1)(2k-1), (k \mp 1)(2k-1)], \end{aligned} \quad (4)$$

where the choice of $+$ or $-$ depends on the parity of k .

Proof of Claim 3: This can be easily established by induction on k . Indeed, Figure 2 establishes the claim for $k = 1, 2, 3$. In general, suppose that, for some k :

$$\begin{aligned} U(2k(2k-1)) &= [k(2k-1), k(2k-1)]^T, \\ V(2k(2k-1)) &= [(k+1)(2k-1), (k-1)(2k-1)]. \end{aligned}$$

Generalizing what is taking place from Step 13 through Step 30 of Figure 2, the dynamic proceeds with both players playing strategy 1 for one step, the row player playing strategy 1 and the column player playing strategy 2 for the next $4k$ steps, and both players playing strategy 2 for the next $4k+1$ steps, resulting in

$$\begin{aligned} U(2(k+1)(2(k+1)-1)) &= [(k+1)(2(k+1)-1), (k+1)(2(k+1)-1)]^T, \\ V(2(k+1)(2(k+1)-1)) &= [k(2(k+1)-1), (k+2)(2(k+1)-1)]. \end{aligned}$$

This establishes the claim for $k+1$. The derivation is similar, if for k Equation (4) is satisfied with \pm and \mp instantiated by $-$ and $+$ respectively. \square

Notice that Claim 3 proves Part 1 of Lemma 1 for $n = 2$.

Now, for any given T , we construct a dynamic (U', V') for I_2 that satisfies the conditions in Part 2 of Lemma 1. Let k be the largest integer such that $k(2k-1) \leq T$, and $l = T - k(2k-1) + 1$. Starting with

$U(0) = [0, 0]^T$, $V(0) = [0, 0]$	$U(0) = [0, 0]^T$, $V(0) = [0, 0]$
Step 1: row chooses 1 column chooses 2	row chooses 1 column chooses 1
$U(1) = [1, 0]^T$, $V(1) = [0, 1]$	$U(1) = [1, 0]^T$, $V(1) = [1, 0]$
Step 2: row chooses 2 column chooses 2	row chooses 1 column chooses 2
$U(2) = [1, 1]^T$, $V(2) = [0, 2]$	$U(2) = [2, 0]^T$, $V(2) = [1, 1]$
Step 3: row chooses 2 column chooses 1	row chooses 2 column chooses 2
$U(3) = [1, 2]^T$, $V(3) = [1, 2]$	$U(3) = [2, 1]^T$, $V(3) = [1, 2]$
Step 4: row chooses 2 column chooses 1	row chooses 2 column chooses 2
$U(4) = [1, 3]^T$, $V(4) = [2, 2]$	$U(4) = [2, 2]^T$, $V(4) = [1, 3]$
Step 5: row chooses 1 column chooses 1	row chooses 2 column chooses 1
$U(5) = [2, 3]^T$, $V(5) = [3, 2]$	$U(5) = [2, 3]^T$, $V(5) = [2, 3]$
Step 6: row chooses 1 column chooses 1	row chooses 2 column chooses 1
$U(6) = [3, 3]^T$, $V(6) = [4, 2]$	$U(6) = [2, 4]^T$, $V(6) = [3, 3]$
...	...

Fig. 1. The padding dynamics for I_2 .

$U'(0) = [0, 0]^T$ and $V'(0) = [0, 0]$, we first evolve the vectors to

$$\begin{aligned} U'(2l) &= [l, l]^T, \\ V'(2l) &= [l-1, l+1], \end{aligned}$$

as enabled by Claim 2. Because the components of $U'(2l)$ and $V'(2l)$ are exactly $l-1$ larger than the corresponding components of $U(2)$ and $V(2)$ of the main dynamic for I_2 , we can further evolve the vectors U' and V' for $2k(2k-1)-2$ steps, mirroring the players' choices from Steps 3 through $2k(2k-1)$ in the main dynamic for I_2 . Using Claim 3, we arrive at

$$\begin{aligned} U'(2T) &= [T, T]^T, \\ V'(2T) &= [T \pm (2k-1), T \mp (2k-1)], \end{aligned}$$

which satisfies

$$\begin{aligned} &\max V'(2T) - \min U'(2T) \\ &= 2k-1 \\ &= \Theta(T^{\frac{1}{2}}). \end{aligned}$$

The constant hidden by $\Theta(\cdot)$ can obviously be chosen uniformly for all T . We have thus proved Part 2 of Lemma 1 for $n=2$.

Induction Step: Assume that Lemma 1 is true for a certain $n \geq 2$. To prove it for $n+1$, we first consider two *padding dynamics* for I_{n+1} , whose first steps are shown in Figure 3 (in the appendix). We suppress the step numbers and strategy choices in the figure, since these can be easily inferred from the vectors. These dynamics generalize the padding dynamics for I_2 appropriately. Similarly to Claim 2, we can show the following:

Claim 4. *For any $k \geq 1$, there exists a padding dynamic for I_{n+1} such that*

$$\begin{aligned} U((n+1)k) &= [k, k, \dots, k]^T, \\ V((n+1)k) &= [k-1, k, \dots, k, k+1]. \end{aligned}$$

Proof of Claim 4: We omit most of the details as the proof is very similar to that of Claim 2. For example, in the top dynamic in Figure 3, we see that U reaches both $[1, 1, \dots, 1]^T$ and $[3, 3, \dots, 3]^T$. Since the corresponding values for V have the same format up to an additive shift and a permutation of the components, we can repeat the pattern ad infinitum to prove the cases for odd k 's. Similarly, the bottom dynamic in Figure 3 deals with even k 's. \square

Next, we define the main dynamic for I_{n+1} , which pieces together parts of various dynamics for I_n obtained from the inductive hypothesis. We describe this dynamic inductively by dividing it into *epochs*:

- 1) *Initial steps leading to 1st epoch:* Starting with $U(0) = [0, \dots, 0]^T$ and $V(0) = [0, \dots, 0]$, we first evolve the vectors to

$$\begin{aligned} U(n+1) &= [1, 1, \dots, 1]^T, \\ V(n+1) &= [0, 1, \dots, 1, 2], \end{aligned} \tag{5}$$

as enabled by Claim 4. We mark those vectors as the beginning of the 1st epoch.

- 2) *Evolution within an epoch:* For $i \geq 1$, suppose that at the beginning of the i -th epoch we satisfy

$$\begin{aligned} U((n+1)P) &= [P, P, \dots, P]^T, \\ V((n+1)P) &= [Q_1, Q_2, \dots, Q_{n+1}]. \end{aligned}$$

Without loss of generality, let us also assume that $Q_1 \leq Q_2 \leq \dots \leq Q_{n+1}$. Because $(n+1)P = \sum_j Q_j$, we have

$$(n+1)(Q_{n+1} - P) = \sum_j (Q_{n+1} - Q_j).$$

For the next $R = (n+1)(Q_{n+1} - P)$ steps, let U increase only in its $(n+1)$ -th component, and V increase $Q_{n+1} - Q_j$ times its j -th component, for all j (the exact order of those increments doesn't matter). The process is compatible with the definition of a dynamic because, in each of those R steps, the $(n+1)$ -th component of V remains maximal in V , and the first n components of U remain minimal in U . At the end of these steps, we arrive at

$$\begin{aligned} U((n+1)Q_{n+1}) &= [P, \dots, P, P+R]^T, \\ V((n+1)Q_{n+1}) &= [Q_{n+1}, \dots, Q_{n+1}]. \end{aligned} \quad (6)$$

Now, from our inductive hypothesis, there exists a dynamic (\hat{U}, \hat{V}) for I_n such that

$$\begin{aligned} \hat{U}(nR) &= [R, R, \dots, R]^T, \\ \hat{V}(nR) &= [S_1, S_2, \dots, S_n], \end{aligned}$$

and

$$\max \hat{V}(nR) - \min \hat{U}(nR) = \Theta(R^{\frac{n-1}{n}}),$$

where the constant hidden by $\Theta(\cdot)$ is independent of R . Starting from (6), for the next nR steps, we increment only the first n components of U and V , in a way that mirrors the strategy choices of the players in the evolution of \hat{U} and \hat{V} , starting from $\tilde{U}(0) = [0, \dots, 0]^T$ and $\tilde{V}(0) = [0, \dots, 0]$, until $\tilde{U}(nR)$ and $\tilde{V}(nR)$. Because the $(n+1)$ -th component of V remains minimal in V , we see that, in each of those nR steps, a maximal component among the first n components of V is also a maximal component of the entire vector V . Similarly, a minimal component among the first n components of U is also a minimal component of the entire vector U . Therefore, the process is compatible with the definition of a dynamic. At the end of the nR steps, we have

$$\begin{aligned} U((n+1)(P+R)) &= [P+R, \dots, P+R, P+R]^T, \\ V((n+1)(P+R)) &= [Q_{n+1} + S_1, \dots, Q_{n+1} + S_n, Q_{n+1}], \end{aligned}$$

which we mark as the beginning of the $(i+1)$ -th epoch. Notice that the vectors have a format that allows the induction to continue.

We analyze the convergence rate of the main dynamic for I_{n+1} . For each i , let $(n+1)T_i$ be the step number at the beginning of the i -th epoch, and G_i the gap

$$G_i = \max V((n+1)T_i) - \min U((n+1)T_i).$$

Using the P , Q , R , and S notation above, we have the following relations:

$$\begin{aligned} T_i &= P, \\ T_{i+1} &= P + R, \\ G_i &= Q_{n+1} - P, \\ G_{i+1} &= \max_j (S_j + Q_{n+1}) - (P + R) \\ &= (Q_{n+1} - P) + (\max \hat{V}(nR) - \min \hat{U}(nR)) \\ &= (Q_{n+1} - P) + \Theta(R^{\frac{n-1}{n}}), \\ R &= (n+1)(Q_{n+1} - P). \end{aligned}$$

From the above, along with the initial values from (5), we obtain the following recursive relations:

$$\begin{aligned} G_1 &= 1, \\ T_1 &= 1, \\ G_{i+1} &= G_i + \Theta([(n+1)G_i]^{\frac{n-1}{n}}), \\ T_{i+1} &= T_i + (n+1)G_i, \end{aligned}$$

where the constants hidden by the $\Theta(\cdot)$'s depend only on $n+1$. A simple calculation based on those relations yields

$$\begin{aligned} G_i &= \Theta(i^n), \\ T_i &= \Theta(i^{n+1}), \end{aligned}$$

and so

$$G_i = \Theta(T_i^{\frac{n}{n+1}}),$$

where the constants hidden by the $\Theta(\cdot)$'s depend only on $n+1$. Consequently, by considering the beginning of each of the infinitely many epoches, the main dynamic for I_{n+1} satisfies Part 1 of Lemma 1 for $n+1$.

We are now ready to construct, for any given T , a dynamic (U', V') for I_{n+1} satisfying the conditions in Part 2. Let k be the largest integer so that $T_k \leq T$, and $l = T - T_k + 1$. Starting from $U'(0) = [0, \dots, 0]^T$ and $V'(0) = [0, \dots, 0]$, we first evolve the vectors to

$$\begin{aligned} U'((n+1)l) &= [l, l, \dots, l]^T, \\ V'((n+1)l) &= [l-1, l, \dots, l, l+1], \end{aligned}$$

as enabled by Claim 4. Because the components of $U'((n+1)l)$ and $V'((n+1)l)$ are exactly $l-1$ larger than the corresponding components of $U(n+1)$ and $V(n+1)$ in the main dynamic for I_{n+1} (i.e. the vectors marking the beginning of the 1st epoch), we can further evolve the vectors U' and V' for $(n+1)T_k - (n+1)$ steps, mirroring the players' choices in Steps $n+2$ through $(n+1)T_k$ (i.e. up until the beginning of the k -th epoch) in the main

dynamic for I_{n+1} . The components of $U'((n+1)T)$ and $V'((n+1)T)$ at the end of this process are $l-1$ plus the corresponding components of $U((n+1)T_k)$ and $V((n+1)T_k)$ in the main dynamic for I_{n+1} . Thus, we have

$$U'((n+1)T) = [T, T, \dots, T]^T,$$

and

$$\begin{aligned} & \max V'((n+1)T) - \min U'((n+1)T) \\ &= G_k \\ &= \Theta(T_k^{\frac{n}{n+1}}) \\ &= \Theta(T^{\frac{n}{n+1}}). \end{aligned}$$

The constant hidden by the $\Theta(\cdot)$'s can obviously be chosen uniformly for all T . We have thus proved Part 2 of Lemma 1 for $n+1$. By induction, the proof of Lemma 1 is completed. \square

\square

Remark 2. Notice that, even though we do not explicitly state it in Theorem 2, our proof implies something stronger, namely that for every $n \geq 2$, there exists a dynamic for I_n such that for all t (as opposed to just infinitely many t 's):

$$\max V(t) - \min U(t) = \Theta(t^{\frac{n-1}{n}}).$$

REFERENCES

- [1] I. Adler. The equivalence of linear programs and zero-sum games. *International Journal of Game Theory*, 42(1):165–177, 2013. **I**
- [2] D. J. Albers and C. Reid. An interview with George B. Dantzig: the father of linear programming. *The College Mathematics Journal*, 17(4):293–314, 1986. **I**
- [3] R. Aumann. Game theory. *The new palgrave—a dictionary of Economics*. Edited by John Eatwell, Murray Milgate and Peter Newman. The Macmillan Press Limited, 1987. **I**
- [4] U. Berger. Fictitious play in $2 \times n$ games. *Journal of Economic Theory*, 120(2):139–154, 2005. **I-0a**
- [5] F. Brandt, F. Fischer, and P. Harrenstein. On the rate of convergence of fictitious play. *Theory of Computing Systems*, 53:41–52, 2013. **I-0a**
- [6] G. W. Brown. Some notes on computation of games solutions. Technical report, P-78, The Rand Corporation, 1949. **(document), I, I-0a**
- [7] G. W. Brown. Iterative solution of games by fictitious play. *Activity Analysis of Production and Allocation*, 13(1):374–376, 1951. **I**
- [8] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006. **I, I, I, I-0a**
- [9] V. Conitzer. Approximation guarantees for fictitious play. In *the 47th Annual Allerton Conference on Communication, Control, and Computing*, pages 636–643. IEEE, 2009. **I-0a**
- [10] G. B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, 1963. **I**
- [11] C. Daskalakis, A. Deckelbaum, and A. Kim. Near-optimal no-regret algorithms for zero-sum games. *Games and Economic Behavior*, 2014. **I**
- [12] D. P. Foster and H. P. Young. On the nonconvergence of fictitious play in coordination games. *Games and Economic Behavior*, 25(1):79–96, 1998. **I-0a**
- [13] Y. Freund and R. E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1):79–103, 1999. **I**
- [14] D. Fudenberg. *The theory of learning in games*, volume 2. MIT press, 1998. **I-0a**
- [15] A. Gaunersdorfer and J. Hofbauer. Fictitious play, Shapley polygons, and the replicator equation. *Games and Economic Behavior*, 11(2):279–303, 1995. **I-0a**
- [16] P. W. Goldberg, R. Savani, T. B. Sørensen, and C. Ventre. On the approximation performance of fictitious play in finite games. *International Journal of Game Theory*, 42(4):1059–1083, 2013. **I-0a**
- [17] S. Hahn. The convergence of fictitious play in 3×3 games with strategic complementarities. *Economics Letters*, 64(1):57–60, 1999. **I-0a**
- [18] J. Hannan. Approximation to Bayes risk in repeated play. *Contributions to the Theory of Games*, 3:97–139, 1957. **I-0a**
- [19] C. Harris. On the rate of convergence of continuous-time fictitious play. *Games and Economic Behavior*, 22(2):238–259, 1998. **I, I-0a**
- [20] J. Hofbauer and K. Sigmund. Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 40(4):479–519, 2003. **I-0a**
- [21] S. Hon-Snir, D. Monderer, and A. Sela. A learning approach to auctions. *Journal of Economic Theory*, 82(1):65–88, 1998. **I-0a**
- [22] J. S. Jordan. Three problems in learning mixed-strategy Nash equilibria. *Games and Economic Behavior*, 5(3):368–386, 1993. **I-0a**
- [23] S. Karlin. *Mathematical Methods and Theory in Games, Programming, and Economics*. Addison-Wesley, 1959. **I, II**
- [24] V. Krishna and T. Sjöström. Learning in games: Fictitious play dynamics. In *Cooperation: Game-Theoretic Approaches*, pages 257–273. Springer, 1997. **I-0a**
- [25] V. Krishna and T. Sjöström. On the convergence of fictitious play. *Mathematics of Operations Research*, 23(2):479–511, 1998. **I-0a**
- [26] N. Littlestone and M. K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, 1994. **I**
- [27] R. D. Luce and H. Raiffa. *Games and Decisions: Introduction and Critical Survey*. Courier Dover Publications, 1957. **I**
- [28] P. Milgrom and J. Roberts. Adaptive and sophisticated learning in normal form games. *Games and Economic Behavior*, 3(1):82–100, 1991. **I-0a**
- [29] K. Miyasawa. On the convergence of the learning process in a 2×2 non-zero-sum two-person game. Technical report, DTIC Document, 1961. **I-0a**
- [30] D. Monderer and A. Sela. A 2×2 game without the fictitious play property. *Games and Economic Behavior*, 14(1):144–148, 1996. **I-0a**
- [31] D. Monderer and L. S. Shapley. Fictitious play property for games with identical interests. *Journal of Economic Theory*, 68(1):258–265, 1996. **I-0a**
- [32] A. Rakhlin and K. Sridharan. Optimization, learning, and games with predictable sequences. In *the 27th Annual Conference on Neural Information Processing Systems (NIPS)*, 2013. **I**
- [33] J. Robinson. An iterative method of solving a game. *Annals of Mathematics*, pages 296–301, 1951. **(document), I, I-0a, II, II**
- [34] A. Sela. Fictitious play in 2×3 games. *Games and Economic Behavior*, 31(1):152–162, 2000. **I-0a**
- [35] L. S. Shapley. Some topics in two-person games. *Advances in Game Theory*, 52:1–29, 1964. **I-0a**
- [36] J. von Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1928. **I**
- [37] J. von Neumann. A numerical method to determine optimum strategy. *Naval Research Logistics Quarterly*, 1(2):109–115, 1954. **I-0a**

APPENDIX

	$U(0) = [0, 0]^T,$	$V(0) = [0, 0]$
Step 1:	row chooses 1	column chooses 2
	$U(1) = [1, 0]^T,$	$V(1) = [0, 1]$
Step 2:	row chooses 2	column chooses 2
	$U(2) = [1, 1]^T,$	$V(2) = [0, 2]$
Step 3:	row chooses 2	column chooses 2
	$U(3) = [1, 2]^T,$	$V(3) = [0, 3]$
Step 4:	row chooses 2	column chooses 1
	$U(4) = [1, 3]^T,$	$V(4) = [1, 3]$
Step 5:	row chooses 2	column chooses 1
	$U(5) = [1, 4]^T,$	$V(5) = [2, 3]$
Step 6:	row chooses 2	column chooses 1
	$U(6) = [1, 5]^T,$	$V(6) = [3, 3]$
Step 7:	row chooses 2	column chooses 1
	$U(7) = [1, 6]^T,$	$V(7) = [4, 3]$
Step 8:	row chooses 1	column chooses 1
	$U(8) = [2, 6]^T,$	$V(8) = [5, 3]$

Step 12:	row chooses 1	column chooses 1
	$U(12) = [6, 6]^T,$	$V(12) = [9, 3]$
Step 13:	row chooses 1	column chooses 1
	$U(13) = [7, 6]^T,$	$V(13) = [10, 3]$
Step 14:	row chooses 1	column chooses 2
	$U(14) = [8, 6]^T,$	$V(14) = [10, 4]$

Step 20:	row chooses 1	column chooses 2
	$U(20) = [14, 6]^T,$	$V(20) = [10, 10]$
Step 21:	row chooses 1	column chooses 2
	$U(21) = [15, 6]^T,$	$V(21) = [10, 11]$
Step 22:	row chooses 2	column chooses 2
	$U(22) = [15, 7]^T,$	$V(22) = [10, 12]$

Step 30:	row chooses 2	column chooses 2
	$U(30) = [15, 15]^T,$	$V(30) = [10, 20]$

Fig. 2. The main dynamic for I_2 .

$$\begin{array}{ll}
U(0) = [0, 0, 0, \dots, 0]^T, & V(0) = [0, 0, 0, \dots, 0] \\
U(1) = [1, 0, 0, \dots, 0]^T, & V(1) = [0, 1, 0, \dots, 0] \\
U(2) = [1, 1, 0, \dots, 0]^T, & V(2) = [0, 1, 1, \dots, 0] \\
& \dots \\
U(n) = [1, 1, 1, \dots, 1, 0]^T, & V(n) = [0, 1, 1, \dots, 1, 1] \\
U(n+1) = [1, 1, 1, \dots, 1, 1]^T, & V(n+1) = [0, 1, 1, \dots, 1, 2] \\
U(n+2) = [1, 1, 1, \dots, 1, 2]^T, & V(n+2) = [1, 1, 1, \dots, 1, 2] \\
U(n+3) = [1, 1, 1, \dots, 1, 3]^T, & V(n+3) = [2, 1, 1, \dots, 1, 2] \\
U(n+4) = [2, 1, 1, \dots, 1, 3]^T, & V(n+4) = [2, 2, 1, \dots, 1, 2] \\
& \dots \\
U(2n+2) = [2, 2, 2, \dots, 2, 1, 3]^T, & V(2n+2) = [2, 2, 2, \dots, 2, 2, 2] \\
U(2n+3) = [2, 2, 2, \dots, 2, 2, 3]^T, & V(2n+3) = [2, 2, 2, \dots, 2, 3, 2] \\
U(2n+4) = [2, 2, 2, \dots, 2, 3, 3]^T, & V(2n+4) = [3, 2, 2, \dots, 2, 3, 2] \\
U(2n+5) = [3, 2, 2, \dots, 2, 3, 3]^T, & V(2n+5) = [3, 3, 2, \dots, 2, 3, 2] \\
& \dots \\
U(3n+2) = [3, 3, 3, \dots, 3, 2, 3, 3]^T, & V(3n+2) = [3, 3, 3, \dots, 3, 3, 3, 2] \\
U(3n+3) = [3, 3, 3, \dots, 3, 3, 3, 3]^T, & V(3n+3) = [3, 3, 3, \dots, 3, 4, 3, 2] \\
& \dots, & \dots,
\end{array}$$

and

$$\begin{array}{ll}
U(0) = [0, 0, 0, \dots, 0]^T, & V(0) = [0, 0, 0, \dots, 0] \\
U(1) = [1, 0, 0, \dots, 0]^T, & V(1) = [1, 0, 0, \dots, 0] \\
U(2) = [2, 0, 0, \dots, 0]^T, & V(2) = [1, 1, 0, \dots, 0] \\
U(3) = [2, 1, 0, \dots, 0]^T, & V(3) = [1, 1, 1, \dots, 0] \\
& \dots \\
U(n+1) = [2, 1, 1, \dots, 1, 0]^T, & V(n+1) = [1, 1, 1, \dots, 1, 1] \\
U(n+2) = [2, 1, 1, \dots, 1, 1]^T, & V(n+2) = [1, 1, 1, \dots, 1, 2] \\
U(n+3) = [2, 1, 1, \dots, 1, 2]^T, & V(n+3) = [1, 2, 1, \dots, 1, 2] \\
U(n+4) = [2, 2, 1, \dots, 1, 2]^T, & V(n+4) = [1, 2, 2, \dots, 1, 2] \\
& \dots \\
U(2n+1) = [2, 2, 2, \dots, 2, 1, 2]^T, & V(2n+1) = [1, 2, 2, \dots, 2, 2] \\
U(2n+2) = [2, 2, 2, \dots, 2, 2, 2]^T, & V(2n+2) = [1, 2, 2, \dots, 2, 3, 2] \\
U(2n+3) = [2, 2, 2, \dots, 2, 3, 2]^T, & V(2n+3) = [2, 2, 2, \dots, 2, 3, 2] \\
U(2n+4) = [2, 2, 2, \dots, 2, 4, 2]^T, & V(2n+4) = [2, 2, 2, \dots, 2, 3, 3] \\
U(2n+5) = [2, 2, 2, \dots, 2, 4, 3]^T, & V(2n+5) = [3, 2, 2, \dots, 2, 3, 3] \\
& \dots \\
U(3n+3) = [3, 3, 3, \dots, 3, 2, 4, 3]^T, & V(3n+3) = [3, 3, 3, \dots, 3, 3, 3, 3] \\
U(3n+4) = [3, 3, 3, \dots, 3, 3, 4, 3]^T, & V(3n+4) = [3, 3, 3, \dots, 3, 4, 3, 3] \\
U(3n+5) = [3, 3, 3, \dots, 3, 4, 4, 3]^T, & V(3n+5) = [3, 3, 3, \dots, 3, 4, 3, 4] \\
U(3n+6) = [3, 3, 3, \dots, 3, 4, 4, 4]^T, & V(3n+6) = [4, 3, 3, \dots, 3, 4, 3, 4] \\
& \dots \\
U(4n+3) = [4, 4, 4, \dots, 4, 3, 4, 4, 4]^T, & V(4n+3) = [4, 4, 4, \dots, 4, 4, 4, 3, 4] \\
U(4n+4) = [4, 4, 4, \dots, 4, 4, 4, 4, 4]^T, & V(4n+4) = [4, 4, 4, \dots, 4, 5, 4, 3, 4] \\
& \dots, & \dots
\end{array}$$

Fig. 3. The padding dynamics for I_{n+1} .