A General Class of Adaptive Strategies*

Sergiu Hart[†] Andreu Mas-Colell[‡]

December 26, 1999

Abstract

We exhibit and characterize an entire class of simple adaptive strategies, in the repeated play of a game, having the Hannan-consistency property: In the long-run, the player is guaranteed an average payoff as large as the best-reply payoff to the empirical distribution of play of the other players; i.e., there is no "regret." Smooth fictitious play (Fudenberg and Levine [1995]) and regret-matching (Hart and Mas-Colell [2000]) are particular cases. The motivation and application of the current paper come from the study of procedures whose empirical distribution of play is, in the long-run, (almost) a correlated equilibrium. For the analysis we first develop a generalization of Blackwell's [1956a] approachability strategy for games with vector payoffs.

Keywords: adaptive strategies, approachability, correlated equilibrium, fictitious play, regret.

Journal of Economic Literature Classification: C7, D7, C6.

1. Introduction

Consider a game repeated through time. We are interested in strategies of play which, while simple to implement, generate desirable outcomes. Such strategies,

*Previous version: March 1999. We thank the referees and editors for their useful comments. Research partially supported by grants of: the Israel Academy of Sciences and Humanities, the Spanish Ministry of Education, the Generalitat de Catalunya, CREI, and the EU-TMR Research Network.

[†]Center for Rationality and Interactive Decision Theory; Department of Mathematics; and Department of Economics; The Hebrew University of Jerusalem, 91904 Jerusalem, Israel. *E-mail*: hart@math.huji.ac.il *URL*: http://www.ma.huji.ac.il/~hart

[‡]Department of Economics and Business; and CREI; Universitat Pompeu Fabra, Ramon Trias Fargas, 25-27, 08005 Barcelona, Spain. *E-mail:* mcolell@upf.es

typically consisting of moves in "improving" directions, are usually referred to as adaptive.

In Hart and Mas-Colell [2000] we presented simple adaptive strategies with the property that, if used by all players, the empirical distribution of play is, in the long-run, (almost) a correlated equilibrium of the game (for other procedures leading to correlated equilibria, see[1] Foster and Vohra [1997] and Fudenberg and Levine [1999]). From this work we are led—for reasons we will comment upon shortly—to the study of a concept originally introduced by Hannan [1957]. A strategy of a player is called *Hannan-consistent* if it guarantees that his long-run average payoff is as large as the highest payoff that can be obtained (i.e., the one-shot best-reply payoff) against the empirical distribution of play of the other players. In other words, a strategy is Hannan-consistent if, given the play of the others, there is no regret in the long-run for not having played (constantly) any particular action. As a matter of terminology, the *regret* of player i for an action[2] k at period t is the difference in his average payoff up to t that results from replacing his actual past play by the constant play of action k. Hannan-consistency thus means that all regrets are non-positive, as t goes to infinity.

In this paper we concentrate on the notion of Hannan-consistency, rather than on its stronger conditional version which characterizes convergence to the set of correlated equilibria (see Hart and Mas-Colell [2000]). This is just to focus on essentials. The extension to the conditional setup is straightforward; see Section 5 below.

Hannan-consistent strategies have been obtained by several authors: Hannan [1957], Blackwell [1956b] (see also Luce and Raiffa [1957, pp. 482-483]), Foster and Vohra [1993, 1998], Fudenberg and Levine [1995], Freund and Schapire [1999], and Hart and Mas-Colell [2000, Section 4(c)].[3] The strategy of Fudenberg and Levine [1995] (as well as those of Hannan [1957], Foster and Vohra [1993, 1998] and Freund and Schapire [1999]) is a smoothed out version of fictitious play. (We note

[1]For these and the other topics discussed in this paper, the reader is referred also to the book of Fudenberg and Levine [1998] and the survey of Foster and Vohra [1999] (as well as to the other papers in the special issue of *Games and Economic Behavior* 29 [1999]).

[2]Think of this as the "regret for not having played k in the past."

[3]See also Baños [1968] and Megiddo [1980] and, in the computer science literature, Littlestone and Warmuth [1994], Auer *et al.* [1995] and the book of Borodin and El-Yaniv [1998].

that fictitious play—which may be stated as: at each period play an action with maximal regret—is by itself not Hannan-consistent.) In contrast, the strategy of Hart and Mas-Colell [2000], called "regret-matching," prescribes, at each period, play probabilities that are proportional to the (positive) regrets. That is, if we write $D(k)$ for the regret of i for action k at time t (as defined above) and $D_+(k)$ for the *positive regret* (i.e., $D(k)$ when $D(k) > 0$ and 0 when $D(k) \leq 0$), then the probability of playing action k at period $t+1$ is simply $D_+(k) / \sum_{k'} D_+(k')$.

Clearly, a general examination is called for. Smooth fictitious play and regret-matching should be but particular instances of a whole class of adaptive strategies with the Hannan-consistency property. In this paper we exhibit and characterize such a class. It turns out to contain, in particular, a large variety of new simple adaptive strategies.

In Hart and Mas-Colell [2000] we have introduced Blackwell's [1956a] approachability theory for games with vector payoffs as the appropriate basic tool for the analysis: The vector payoff is simply the vector of regrets. In this paper, therefore, we proceed in two steps. First, in Section 2, we generalize Blackwell's result: Given an approachable set (in vector payoff space), we find the class of ("directional") strategies that guarantee that the set is approached. We defer the specifics to that section. Suffice it to say that Blackwell's strategy emerges as the particular quadratic case of a continuum of strategies where continuity and, interestingly, integrability feature decisively.

Second, in Section 3, we apply the general theory to the regret framework and derive an entire class of Hannan-consistent strategies. A feature common to them all is that, in the spirit of bounded rationality, they aim at "better" rather than "best" play. We elaborate on this aspect, and carry out an explicit discussion of fictitious play in Section 4. Section 5 discusses a number of extensions, including conditional regrets and correlated equilibria.

2. The Approachability Problem

2.1. Model and Main Theorem

In this section we will consider games where a player's payoffs are *vectors* (rather than, as in standard games, scalar real numbers), as introduced by Blackwell [1956a]. This setting may appear unnatural at first. However, it has turned out to be quite useful: The coordinates may represent different commodities; or contingent payoffs in different states of the world (when there is incomplete information); or, as we will see below (in Section 3), regrets in a standard game.

Formally, we are given a game in strategic form played by a player i against an opponent $-i$ (which may be Nature and/or the other players). The *action* sets are the finite[4] sets S^i for player i and S^{-i} for $-i$. The payoffs are *vectors* in some Euclidean space. We denote the payoff function by[5] $A : S \equiv S^i \times S^{-i} \to \mathbb{R}^m$; thus $A(s^i, s^{-i}) \in \mathbb{R}^m$ is the payoff vector when i chooses s^i and $-i$ chooses s^{-i}. As usual, A is extended bi-linearly to mixed actions, thus[6] $A : \Delta(S^i) \times \Delta(S^{-i}) \to \mathbb{R}^m$.

Let time be discrete: $t = 1, 2, \dots$, and denote by $s_t = (s_t^i, s_t^{-i}) \in S^i \times S^{-i}$ the actions chosen by i and $-i$, respectively, at time t. The payoff vector in period t is $a_t := A(s_t)$, and $\overline{a}_t := (1/t)\sum_{\tau \le t} a_\tau$ is the average payoff vector up to t. A *strategy*[7] *for player* i assigns to every history of play $h_{t-1} = (s_\tau)_{\tau \le t-1} \in (S)^{t-1}$ a (randomized) choice of action $\sigma_t^i \equiv \sigma_t^i(h_{t-1}) \in \Delta(S^i)$ at time t, where $[\sigma_t^i(h_{t-1})](s^i)$ is, for each s^i in S^i, the probability that i plays s^i at period t following a history h_{t-1}.

Let $\mathcal{C} \subset \mathbb{R}^m$ be a convex and closed[8] set. The set $\mathcal{C}$ is *approachable* by player i (cf. Blackwell [1956a]; see Remark 3 below) if there is a strategy of i such that,

[4] See however Remark 2 below. Also, we always assume that S^i contains at least two elements.

[5] $\mathbb{R}$ is the real line, and $\mathbb{R}^m$ is the m-dimensional Euclidean space. For $x = (x_k)_{k=1}^m$ and $y = (x_k)_{k=1}^m$ in $\mathbb{R}^m$, we write $x \ge y$ when $x_k \ge y_k$ for all k, and $x >> y$ when $x_k > y_k$ for all k. The non-negative, non-positive, positive and negative orthants of $\mathbb{R}^m$ are, respectively, $\mathbb{R}^m_+ := \{x \in \mathbb{R}^m : x \ge 0\}$, $\mathbb{R}^m_- := \{x \in \mathbb{R}^m : x \le 0\}$, $\mathbb{R}^m_{++} := \{x \in \mathbb{R}^m : x >> 0\}$ and $\mathbb{R}^m_{--} := \{x \in \mathbb{R}^m : x << 0\}$.

[6] We write $\Delta(Z)$ for the set of probability distributions on Z, i.e., the $(|Z|-1)$-dimensional unit simplex $\Delta(Z) := \{p \in \mathbb{R}^Z_+ : \sum_{z \in Z} p(z) = 1\}$.

[7] We use the term "action" for a one-period choice, and the term "strategy" for a multi-period choice.

[8] A set is approachable if and only if its closure is approachable; we thus assume without loss of generality that the set $\mathcal{C}$ is closed.

no matter what $-i$ does,[9] $\text{dist}(\overline{a}_t, \mathcal{C}) \to 0$ almost surely as $t \to \infty$. Blackwell's result can then be stated as follows:

Blackwell's Approachability Theorem.

(1) A convex and closed set $\mathcal{C}$ is approachable if and only if every half-space $\mathcal{H}$ containing $\mathcal{C}$ is approachable.

(2) A half-space $\mathcal{H}$ is approachable if and only if there exists a mixed action of player i such that the expected vector payoff is guaranteed to lie in $\mathcal{H}$; i.e., there is $\sigma^i \in \Delta(S^i)$ such that $A(\sigma^i, s^{-i}) \in \mathcal{H}$ for all $s^{-i} \in S^{-i}$.

The condition for $\mathcal{C}$ to be approachable may be restated as follows (since, clearly, it suffices to consider in (1) only "minimal" half-spaces containing $\mathcal{C}$): For every $\lambda \in \mathbb{R}^m$ there exists $\sigma^i \in \Delta(S^i)$ such that

$$\lambda \cdot A\left(\sigma^i, s^{-i}\right) \leq w(\lambda) := \sup\{\lambda \cdot y : y \in \mathcal{C}\} \text{ for all } s^{-i} \in S^{-i} \tag{2.1}$$

(w is the "support function" of $\mathcal{C}$; note that only those $\lambda \neq 0$ with $w(\lambda) < \infty$ matter for (2.1)). Furthermore, the strategy constructed by Blackwell that yields approachability uses at each step t where the current average payoff $\overline{a}_{t-1}$ is not in $\mathcal{C}$, a mixed choice σ_t^i satisfying (2.1) for that vector $\lambda \equiv \lambda(\overline{a}_{t-1})$ which goes to $\overline{a}_{t-1}$ from that point y in $\mathcal{C}$ that is closest to $\overline{a}_{t-1}$ (see Figure 2.1). To get some intuition, note that the next period expected payoff vector $b := E[a_t | h_{t-1}]$ lies in the half-space $\mathcal{H}$, and thus satisfies $\lambda \cdot b \leq w(\lambda) < \lambda \cdot \overline{a}_{t-1}$, which implies that

$$\lambda \cdot (E[\overline{a}_t | h_{t-1}] - \overline{a}_{t-1}) = \lambda \cdot (\frac{1}{t} b + \frac{t-1}{t} \overline{a}_{t-1} - \overline{a}_{t-1}) = \frac{1}{t}\, \lambda \cdot (b - \overline{a}_{t-1}) < 0.$$

Therefore the expected average payoff $E[\overline{a}_t | h_{t-1}]$ moves from $\overline{a}_{t-1}$ in the "general direction" of $\mathcal{C}$; in fact, it is closer than $\overline{a}_{t-1}$ to $\mathcal{C}$. Hence $E[\overline{a}_t | h_{t-1}]$ converges to $\mathcal{C}$, and so does the average payoff $\overline{a}_t$ (by the Law of Large Numbers).

Fix an approachable convex and closed set $\mathcal{C}$. We will now consider general strategies of player i which—like Blackwell's strategy above—are defined in terms

[9] We emphasize that the strategies of the opponents $(-i)$ are not in any way restricted; in particular, they may randomize and, furthermore, correlate their actions. All the results in this paper hold against all possible strategies of $-i$, and thus, *a fortiori*, for any specific class of strategies (like independent mixed strategies, and so on). Moreover, requiring independence over $j \neq i$ will not increase the set of strategies of i that guarantee approachability, since the worse $-i$ can do may always be taken to be pure (and thus independent).

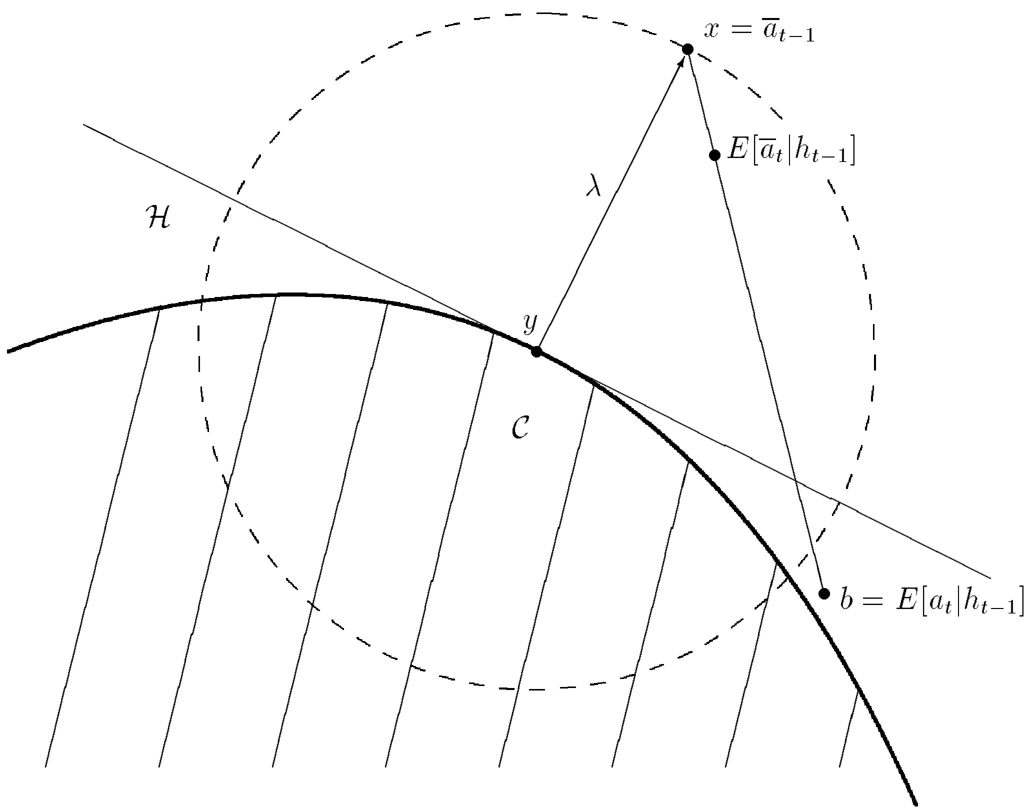

Figure 2.1: Approaching the set $\mathcal{C}$ by Blackwell's strategy

of a *directional mapping,* that is, a function $\Lambda : \mathbb{R}^m \backslash \mathcal{C} \to \mathbb{R}^m$ that associates to every $x \notin \mathcal{C}$ a corresponding "direction" $\Lambda(x)$. Given such a mapping Λ, a strategy of player i is called a Λ-*strategy* if, whenever $\overline{a}_{t-1}$ does not lie in $\mathcal{C}$, it prescribes using at time t a mixed action σ_t^i that satisfies

$$\Lambda(\overline{a}_{t-1}) \cdot A\left(\sigma_t^i, s^{-i}\right) \leq w(\Lambda(\overline{a}_{t-1})) \text{ for all } s^{-i} \in S^{-i} \tag{2.2}$$

(see Figure 2.2: a Λ-strategy guarantees that, when $x = \overline{a}_{t-1} \notin \mathcal{C}$, the next period expected payoff vector $b = E[a_t|h_{t-1}]$ lies in the smallest half-space $\mathcal{H}$ with normal $\Lambda(x)$ that contains $\mathcal{C}$); notice that there is no requirement when $\overline{a}_{t-1} \in \mathcal{C}$. We are interested in finding conditions on the mapping Λ such that, if player i uses a Λ-strategy, then the set $\mathcal{C}$ is guaranteed to be approached, no matter what $-i$ does.

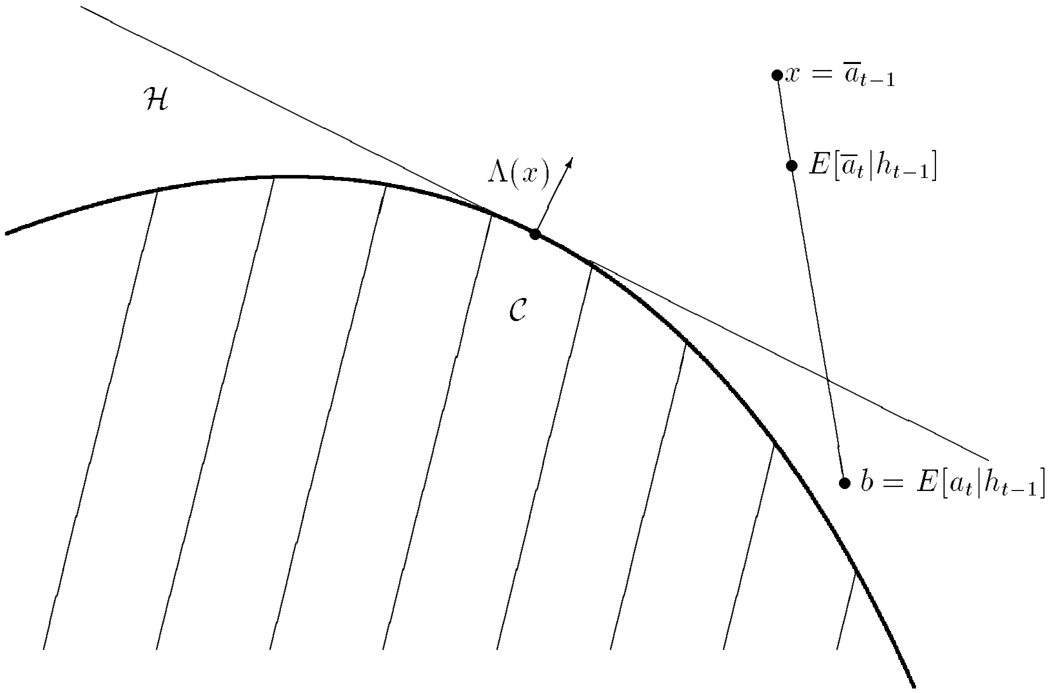

Figure 2.2: A Λ-strategy

We introduce three conditions on a directional mapping Λ, relative to the given set $\mathcal{C}$.

(D1) Λ is continuous.

(D2) Λ is integrable, namely there exists a Lipschitz function[10] $P : \mathbb{R}^m \to \mathbb{R}$ such that $\nabla P(x) = \phi(x)\Lambda(x)$ for almost every $x \notin \mathcal{C}$, where $\phi : \mathbb{R}^m \backslash \mathcal{C} \to \mathbb{R}_{++}$ is a continuous positive function.

(D3) $\Lambda(x) \cdot x > w(\Lambda(x))$ for all $x \notin \mathcal{C}$.

See Figure 2.3. The geometric meaning of (D3) is that the point x is strictly separated from the set $\mathcal{C}$ by $\Lambda(x)$. Note that (D3) implies that all λ with $w(\lambda) = \infty$, as well as $\lambda = 0$, are not allowable directions. Also, observe that the combination

[10]Notice that P is defined on the whole space $\mathbb{R}^m$.

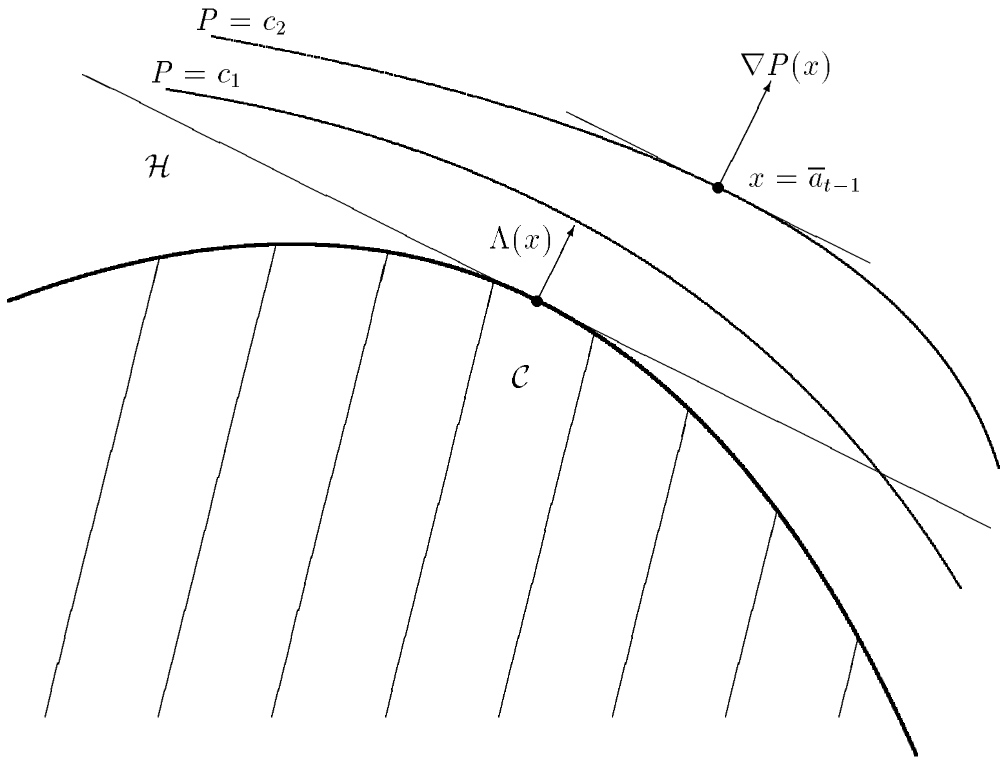

Figure 2.3: The directional mapping Λ and level sets of the potential P

of (D1) and (D2) implies that P is continuously differentiable on $\mathbb{R}^m \backslash \mathcal{C}$ (see Clarke [1983, Corollary to Proposition 2.2.4 and Theorem 2.5.1]). We will refer to the function P as the *potential* of Λ.

The main result of this section is:

Theorem 2.1. *Suppose that player i uses a Λ-strategy, where Λ is a directional mapping satisfying (D1), (D2), and (D3) for the approachable convex and closed set $\mathcal{C}$. Then the average payoff vector is guaranteed to approach the set $\mathcal{C}$; that is,* $\operatorname{dist}(\overline{a}_t, \mathcal{C}) \to 0$ *almost surely as $t \to \infty$, for any strategy of $-i$.*

Before proving the Theorem (in the next subsection), we state a number of comments.

Remarks.

1. The conditions (D1)–(D3) are independent of the game A (they depend on $\mathcal{C}$ only). That is, given a directional mapping Λ satisfying (D1)–(D3), a Λ-strategy is guaranteed to approach $\mathcal{C}$ for *any* game A for which $\mathcal{C}$ is approachable (of course, the specific choice of action depends on A, according to (2.2)). It is in this sense that we refer to the Λ-strategies as "universal."

2. The action sets S^i and S^{-i} need not be finite; as we will see in the proof, it suffices for the range of A to be bounded.

3. As in Blackwell's result, our proof below yields *uniform* approachability: For every ε there is $t_0 \equiv t_0(\varepsilon)$ such that $E\left[\operatorname{dist}(\overline{a}_t, \mathcal{C})\right] < \varepsilon$ for all $t > t_0$ and all strategies of $-i$ (i.e., t_0 is independent of the strategy of $-i$).

4. The conditions on P are invariant under strictly increasing monotone transformations (with positive derivative); that is, only the level sets of P matter.

5. If the potential P is a convex function and $\mathcal{C} = \{y : P(y) \leq c\}$ for some constant c, then (D3) is automatically satisfied: $P(x) > P(y)$ implies $\nabla P(x) \cdot x > \nabla P(x) \cdot y$.

6. Given a norm $\|\cdot\|$ on $\mathbb{R}^m$, consider the resulting "distance from $\mathcal{C}$" function $P(x) := \min_{y \in \mathcal{C}} \|x - y\|$. If P is a smooth function (which is always the case when either the norm is smooth—i.e., the corresponding unit ball has smooth boundary—or when the boundary of $\mathcal{C}$ is smooth), then the mapping $\Lambda = \nabla P$ satisfies (D1)–(D3) (the latter by the previous Remark 5). In particular, the l_2 Euclidean norm yields precisely the Blackwell strategy, since then $\nabla P(x)$ is proportional to $x - y(x)$, where $y(x) \in \mathcal{C}$ is the point in $\mathcal{C}$ closest to x. The l_p norm is smooth for $1 < p < \infty$; therefore it yields strategies that guarantee approachability for *any* approachable set $\mathcal{C}$. However, if the boundary of $\mathcal{C}$ is not smooth—for instance, when $\mathcal{C}$ is an orthant, an important case in applications—then (D1) is *not* satisfied in the extreme cases $p = 1$ and $p = \infty$ (see Figure 2.4; more on these two cases below).

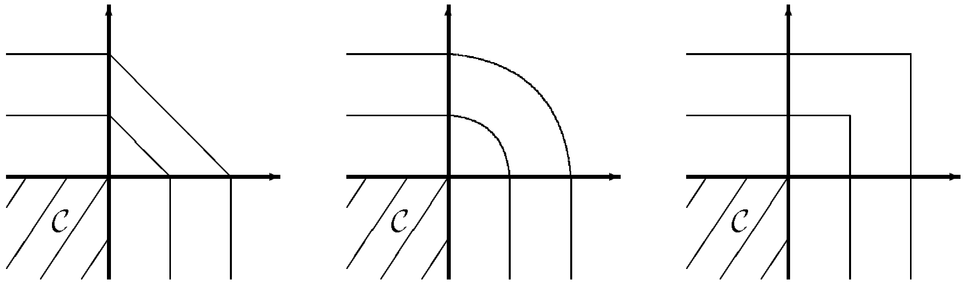

$p = 1$ $\quad$ $1 < p < \infty$ $\quad$ $p = \infty$

Figure 2.4: The l_p potential for an orthant $\mathcal{C}$

7. When $\mathcal{C} = \mathbb{R}^m_-$ and P is given by (D2), condition (D3) becomes $\nabla P(x) \cdot x > 0$ for every $x \notin \mathcal{C}$, which means that P is increasing along any ray from the origin that goes outside the negative orthant.

2.2. Proof of Theorem 2.1

We begin by proving two auxiliary results. The first applies to functions Q that satisfy conditions similar to but stronger than (D1)–(D3); the second allows us to reduce the general case to such a Q. The set $\mathcal{C}$, the mappings Λ and P, and the strategy of i (which is a Λ-strategy) are fixed throughout. Also, let K be a convex and compact set containing in its interior the range of A (recall that S is finite).

Lemma 2.2. *Let $Q : \mathbb{R}^m \to \mathbb{R}$ be a continuously differentiable function that satisfies:*

(i) $Q(x) \geq 0$ *for all* x;

(ii) $Q(x) = 0$ *for all* $x \in \mathcal{C}$;

(iii) $\nabla Q(x) \cdot x - w(\nabla Q(x)) \geq Q(x)$ *for all* $x \in K \backslash \mathcal{C}$; *and*

(iv) $\nabla Q(x)$ *is non-negatively proportional to* $\Lambda(x)$ *(i.e.,* $\nabla Q(x) = \phi(x)\Lambda(x)$ *where* $\phi(x) \geq 0$*) for all* $x \notin \mathcal{C}$.

Then $\lim_{t\to\infty} Q(\overline{a}_t) = 0$ *a.s. for any strategy of* $-i$.

Proof. We have $\overline{a}_t - \overline{a}_{t-1} = (1/t)(a_t - \overline{a}_{t-1})$; thus, writing x for $\overline{a}_{t-1}$,

$$Q(\overline{a}_t) = Q(x) + \nabla Q(x) \cdot \frac{1}{t}(a_t - x) + o\left(\frac{1}{t}\right), \tag{2.3}$$

since Q is (continuously) differentiable. Moreover, the remainder $o(1/t)$ is uniform, since all relevant points lie in the compact set K. If $x \notin \mathcal{C}$ then player i plays at time t so that

$$\nabla Q(x) \cdot E[a_t | h_{t-1}] \leq w(\nabla Q(x)) \tag{2.4}$$

(by (2.2) and (iv)); if $x \in \mathcal{C}$ then $\nabla Q(x) = 0$ (by (i) and (ii)), and (2.4) holds too. Taking conditional expectation in (2.3) and then substituting (2.4) yields

$$\begin{aligned} E[Q(\overline{a}_t)|h_{t-1}] &\leq Q(x) + \frac{1}{t}\left(w(\nabla Q(x)) - \nabla Q(x) \cdot x\right) + o\left(\frac{1}{t}\right) \\ &\leq Q(x) - \frac{1}{t} Q(x) + o\left(\frac{1}{t}\right), \end{aligned}$$

where we have used (iii) when $x \notin \mathcal{C}$ and (i)–(ii) when $x \in \mathcal{C}$. Thus

$$E[Q(\overline{a}_t)|h_{t-1}] \leq \frac{t-1}{t} Q(\overline{a}_{t-1}) + o\left(\frac{1}{t}\right).$$

This may be rewritten as[11]

$$E[\zeta_t | h_{t-1}] \leq o(1), \tag{2.5}$$

where $\zeta_t := tQ(\overline{a}_t) - (t-1)Q(\overline{a}_{t-1})$. Hence $\limsup_{t\to\infty}(1/t)\sum_{\tau\leq t} E[\zeta_\tau | h_{\tau-1}] \leq 0$. The Strong Law of Large Numbers for Dependent Random Variables (see Loève [1978, Theorem 32.1.E]) implies that $(1/t)\sum_{\tau\leq t}(\zeta_\tau - E[\zeta_\tau | h_{\tau-1}]) \to 0$ a.s. as $t \to \infty$ (note that the ζ_t's are uniformly bounded, as can be immediately seen from equation (2.3): $\zeta_t = Q(\overline{a}_{t-1}) + \nabla Q(\overline{a}_{t-1}) \cdot (a_t - \overline{a}_{t-1}) + o(1)$, and from the fact that everything happens in the compact set K). Therefore $\limsup_{t\to\infty}(1/t)\sum_{\tau\leq t}\zeta_\tau \leq 0$. But $0 \leq Q(\overline{a}_t) = (1/t)\sum_{\tau\leq t}\zeta_\tau$, so $\lim_{t\to\infty} Q(\overline{a}_t) = 0$. ■

[11] Recall that the remainder term $o(1/t)$ was uniform; that is, for every $\varepsilon > 0$ there is $t_0(\varepsilon)$ such that $o(1) < \varepsilon$ is guaranteed for all $t > t_0(\varepsilon)$.

Lemma 2.3. *The function P satisfies:*

(c1) *If the boundary of $\mathcal{C}$ is connected, then there exists a constant c such that*

$$\begin{cases} P(x) = c, & \text{if } x \in \operatorname{bd} \mathcal{C}; \\ P(x) > c, & \text{if } x \notin \mathcal{C}. \end{cases}$$

(c2) *If the boundary of $\mathcal{C}$ is not connected, then there exists a $\lambda \in \mathbb{R}^m \backslash \{0\}$ such that*[12] *$\mathcal{C} = \{x \in \mathbb{R}^m : -w(-\lambda) \leq \lambda \cdot x \leq w(\lambda)\}$ (where $w(\lambda) < \infty$ and $w(-\lambda) < \infty$), and there are constants c_1 and c_2 such that*

$$\begin{cases} P(x) = c_1, & \text{if } x \in \operatorname{bd} \mathcal{C} \text{ and } \lambda \cdot x = w(\lambda); \\ P(x) = c_2, & \text{if } x \in \operatorname{bd} \mathcal{C} \text{ and } (-\lambda) \cdot x = w(-\lambda); \\ P(x) > c_1, & \text{if } x \notin \mathcal{C} \text{ and } \lambda \cdot x > w(\lambda); \\ P(x) > c_2, & \text{if } x \notin \mathcal{C} \text{ and } (-\lambda) \cdot x > w(-\lambda). \end{cases}$$

Proof. Let $x_0, x_1 \in \operatorname{bd} \mathcal{C}$, and denote by λ_j, for $j = 0, 1$, an outward unit normal to $\mathcal{C}$ at x_j; thus $\|\lambda_j\| = 1$ and $\lambda_j \cdot x_j = w(\lambda_j)$.

Step 1: If $\lambda_1 \neq -\lambda_0$, we claim that there is a path on bd $\mathcal{C}$ connecting x_0 and x_1, and moreover $P(x_0) = P(x_1)$. Indeed, there exists a vector[13] $z \in \mathbb{R}^m$ such that $\lambda_0 \cdot z > 0$ and $\lambda_1 \cdot z > 0$. The straight line segment connecting x_0 and x_1 lies in $\mathcal{C}$; we move it in the direction z until it reaches the boundary of $\mathcal{C}$. That is, for each $\eta \in [0,1]$, let $y(\eta) := \eta x_1 + (1 - \eta) x_0 + \alpha(\eta) z$, where $\alpha(\eta) := \max\{\beta : \eta x_1 + (1 - \eta) x_0 + \beta z \in \mathcal{C}\}$; this maximum exists by the choice of z. Note that $y(\cdot)$ is a path on bd $\mathcal{C}$ connecting x_0 and x_1.

It is easy to verify that $\alpha(0) = \alpha(1) = 0$ and that $\alpha : [0,1] \to \mathbb{R}_+$ is a concave function—thus differentiable a.e. For each $k = 1, 2, \ldots$, define $y_k(\eta) := y(\eta) + (1/k) z$; then $y_k(\cdot)$ is a path in $\mathbb{R}^m \backslash \mathcal{C}$, the region where P is continuously differentiable. Let $\overline{\eta} \in (0,1)$ be a point of differentiability of $\alpha(\cdot)$, thus also of $y(\cdot)$, $y_k(\cdot)$ and $P(y_k(\cdot))$; we have $dP(y_k(\overline{\eta}))/d\eta = \nabla P(y_k(\overline{\eta})) \cdot y_k'(\overline{\eta}) = \nabla P(y_k(\overline{\eta})) \cdot y'(\overline{\eta})$. By (D3), $\nabla P(y_k(\overline{\eta})) \cdot y_k(\overline{\eta}) > w(\nabla P(y_k(\overline{\eta}))) \geq \nabla P(y_k(\overline{\eta})) \cdot y(\eta)$ for any $\eta \in [0,1]$ (the second inequality since $y(\eta) \in \mathcal{C}$). Thus, for any accumulation point q of the

[12] This is a general fact about convex sets: The only case where the boundary of a convex closed set $\mathcal{C} \subset \mathbb{R}^m$ is not path-connected is when $\mathcal{C}$ is the set of points lying between two parallel hyperplanes. We prove this in Steps 1–3 below, independently of the function P.

[13] Take for instance $z = \lambda_0 + \lambda_1$.

bounded[14] sequence $(\nabla P(y_k(\overline{\eta})))_{k=1}^{\infty}$, we get $q \cdot y(\overline{\eta}) \geq q \cdot y(\eta)$ for all $\eta \in [0,1]$. Therefore $q \cdot y(\eta)$ is maximized at $\eta = \overline{\eta}$, which implies that $q \cdot y'(\overline{\eta}) = 0$. This holds for *any* accumulation point q, hence $\lim_{k\to\infty} dP(y_k(\overline{\eta}))/d\eta = 0$ for almost every $\overline{\eta}$. Therefore

$$\begin{aligned}
P(x_1) - P(x_0) &= P(y(1)) - P(y(0)) = \lim_k [P(y_k(1)) - P(y_k(0))] \\
&= \lim_k \int_0^1 (dP(y_k(\eta))/d\eta)\, d\eta = \int_0^1 \left(\lim_k dP(y_k(\eta))/d\eta\right) d\eta = 0,
\end{aligned}$$

(again, P is Lipschitz, so $dP(y_k(\eta))/d\eta$ are uniformly bounded).

Step 2: If $\lambda_1 = -\lambda_0$ and there is another boundary point x_2 with outward unit normal λ_2 different from both $-\lambda_0$ and $-\lambda_1$, then we get paths on $\operatorname{bd} \mathcal{C}$ connecting x_0 to x_2 and x_1 to x_2, and also $P(x_0) = P(x_2)$ and $P(x_1) = P(x_2)$ —thus we get the same conclusion as in Step 1.

Step 3: If $\lambda_1 = -\lambda_0$ and no x_2 and λ_2 as in Step 2 exist, it follows that the unit normal to every point on the boundary of $\mathcal{C}$ is either λ_0 or $-\lambda_0$; thus $\mathcal{C}$ is the set bounded between the two parallel hyperplanes $\lambda_0 \cdot x = w(\lambda_0)$ and $-\lambda_0 \cdot x = w(-\lambda_0)$. In particular, the boundary of $\mathcal{C}$ is not connected, and we are in case (c2). Note that in this case when x_0 and x_1 lie on the same hyperplane then $P(x_0) = P(x_1)$ by Step 1 (since $\lambda_1 = \lambda_0 \neq -\lambda_0$).

Step 4: If it is case (c1)—thus not (c2)—then the situation of Step 3 is not possible; thus for any two boundary points x_0 and x_1 we get $P(x_0) = P(x_1)$ by either Step 1 or Step 2.

Step 5: Given $x \notin \mathcal{C}$, let $x_0 \in \operatorname{bd} \mathcal{C}$ be the point in $\mathcal{C}$ that is closest to x. Then the line segment from x to x_0 lies outside $\mathcal{C}$, i.e., $y(\eta) := \eta x + (1-\eta)x_0 \notin \mathcal{C}$ for all $\eta \in (0,1]$. By (D3) and $x_0 \in \mathcal{C}$, it follows that $\nabla P(y(\eta)) \cdot y(\eta) > w(\nabla P(y(\eta))) \geq \nabla P(y(\eta)) \cdot x_0$, or, after dividing by $\eta > 0$, that $\nabla P(y(\eta)) \cdot (x - x_0) > 0$, for all $\eta \in (0,1]$. Hence $P(x) - P(x_0) = \int_0^1 \nabla P(y(\eta)) \cdot y'(\eta)\, d\eta = \int_0^1 \nabla P(y(\eta)) \cdot (x - x_0)\, d\eta > 0$, showing that $P(x) > c$ in case (c1) and $P(x) > c_1$ or $P(x) > c_2$ in case (c2). ■

We can now prove the main result of this section.

Proof of Theorem 2.1. First, use Lemma 2.3 to replace P by P_1 as follows: When the boundary of $\mathcal{C}$ is connected (case (c1)), define $P_1(x) := (P(x) - c)^2$ for

[14] Recall that P is Lipschitz.

$x \notin \mathcal{C}$ and $P_1(x) := 0$ for $x \in \mathcal{C}$; when the boundary of $\mathcal{C}$ is not connected (case (c2)), define $P_1(x) := (P(x)-c_1)^2$ for $x \notin \mathcal{C}$ with $\lambda \cdot x > w(\lambda)$, $P_1(x) := (P(x)-c_2)^2$ for $x \notin \mathcal{C}$ with $(-\lambda)\cdot x > w(-\lambda)$, and $P_1(x) := 0$ for $x \in \mathcal{C}$. It is easy to verify that: P_1 is continuously differentiable; $\nabla P_1(x)$ is positively proportional to $\nabla P(x)$ and thus to $\Lambda(x)$ for $x \notin \mathcal{C}$; $P_1(x) \geq 0$ for all x; and $P_1(x) = 0$ if and only if $x \in \mathcal{C}$.

Given $\varepsilon > 0$, let $k \geq 2$ be a large enough integer such that

$$\frac{\nabla P_1(x) \cdot x - w(\nabla P_1(x))}{P_1(x)} \geq \frac{1}{k} \tag{2.6}$$

for all x in the compact set $K \cap \{x : P_1(x) \geq \varepsilon\}$ (the minimum of the above ratio is attained and it is positive by (D3)). Put[15] $Q(x) := ([P_1(x) - \varepsilon]_+)^k$. Then Q is continuously differentiable (since $k \geq 2$) and it satisfies all the conditions of Lemma 2.2. To check (iii): When $Q(x) = 0$ we have $\nabla Q(x) = 0$, and when $Q(x) > 0$ we have

$$\begin{aligned}
\nabla Q(x) \cdot x - w(\nabla Q(x)) &= k(P_1(x) - \varepsilon)^{k-1} \left[\nabla P_1(x) \cdot x - w(\nabla P_1(x))\right] \\
&\geq (P_1(x) - \varepsilon)^{k-1} P_1(x) \\
&\geq Q(x).
\end{aligned}$$

(the first inequality follows from (2.6)).

By Lemma 2.2, it follows that the Λ-strategy guarantees a.s. $\lim_{t\to\infty} Q(\overline{a}_t) = 0$, or $\limsup_{t\to\infty} P_1(\overline{a}_t) \leq \varepsilon$. Since $\varepsilon > 0$ is arbitrary, this yields a.s. $\lim_{t\to\infty} P_1(\overline{a}_t) = 0$, or $\overline{a}_t \to \mathcal{C}$. ∎

Remark. P may be viewed (up to a constant, as in the definition of P_1 above) as a generalized distance to the set $\mathcal{C}$ (compare with Remark 6 in Subsection 2.1).

2.3. Counterexamples

In this subsection, we provide counterexamples showing the indispensability of the conditions (D1)–(D3) for the validity of Theorem 2.1. The first two examples refer to (D1), the third to (D2), and the last one to (D3).

[15] We write $[z]_+$ for the positive part of z, i.e., $[z]_+ := \max\{z, 0\}$.

Example 2.4. *The role of (D1).*

Consider the following 2-dimensional vector payoff matrix A

	C1	C2
R1	$(0,-1)$	$(0,1)$
R2	$(1,0)$	$(-1,0)$

.

Let i be the Row player and $-i$ the Column player. The set $\mathcal{C} := \mathbb{R}^2_-$ is approachable by the Row player since $w(\lambda) < \infty$ whenever $\lambda \geq 0$, and then the mixed action $\sigma^{Row}(\lambda) := (\lambda_1/(\lambda_1+\lambda_2), \lambda_2/(\lambda_1+\lambda_2))$ of the Row player yields $\lambda \cdot A(\sigma^{Row}(\lambda), \gamma) = 0 = w(\lambda)$ for any action γ of the Column player.

We define a directional mapping Λ_∞ on $\mathbb{R}^2 \backslash \mathbb{R}^2_-$:

$$\Lambda_\infty(x) := \begin{cases} (1,0), & \text{if } x_1 > x_2; \\ (0,1), & \text{if } x_1 \leq x_2. \end{cases}$$

Clearly Λ_∞ is *not continuous*, i.e., it does not satisfy (D1); it does however satisfy (D3) and (D2) (with $P(x) = \max\{x_1, x_2\}$, the l_∞ potential; see Remark 6 in Subsection 2.1). Consider a Λ_∞-strategy for the Row player that, when $x := \overline{a}_{t-1} \notin \mathcal{C}$, plays $\sigma^{Row}(\Lambda_\infty(x))$ at time t; that is, he plays R1 when $x_1 > x_2$, and R2 when $x_1 \leq x_2$. Assume that the Column player plays[16] C2 when $x_1 > x_2$, and C1 when $x_1 \leq x_2$. Then, starting with, say, $a_1 = (0,1) \notin \mathcal{C}$, the vector payoff a_t will always be either $(0,1)$ or $(1,0)$, thus on the line $x_1 + x_2 = 1$, so the average $\overline{a}_t$ does not converge to $\mathcal{C} = \mathbb{R}^2_-$.

Example 2.5. *The role of (D1), again.*

The same as in Example 2.4, but now the directional mapping is Λ_1, defined on $\mathbb{R}^2 \backslash \mathbb{R}^2_-$ by

$$\Lambda_1(x) := \begin{cases} (1,1), & \text{if } x_1 > 0 \text{ and } x_2 > 0; \\ (1,0), & \text{if } x_1 > 0 \text{ and } x_2 \leq 0; \\ (0,1), & \text{if } x_1 \leq 0 \text{ and } x_2 > 0. \end{cases}$$

Again, the mapping Λ_1 is *not continuous*—it does not satisfy (D1)—but it satisfies (D3) and (D2) (with $P(x) := [x_1]_+ + [x_2]_+$, the l_1 potential). Consider a Λ_1-strategy for the Row player where at time t he plays $\sigma^{Row}(\Lambda_1(x))$ when $x :=$

[16] In order to show that the strategy of the Row player does not guarantee approachability to $\mathcal{C}$, we exhibit one strategy of the Column player for which $\overline{a}_t$ does not converge to $\mathcal{C}$.

$\overline{a}_{t-1} \notin \mathcal{C}$, and assume that the Column player plays C1 when $x_1 \leq 0$ and $x_2 > 0$, and plays C2 otherwise. Thus, if $x \notin \mathcal{C}$ then a_t is:

- $(0,1)$ or $(-1,0)$ with equal probabilities, when $x_1 > 0$ and $x_2 > 0$;
- $(0,1)$, when $x_1 > 0$ and $x_2 \leq 0$;
- $(1,0)$, when $x_1 \leq 0$ and $x_2 > 0$.

In all cases the second coordinate of a_t is non-negative; therefore, if we start with, say, $a_1 = (0,1) \notin \mathcal{C}$, then, inductively, the second coordinate of $\overline{a}_{t-1}$ will be strictly positive, so that $\overline{a}_{t-1} \notin \mathcal{C}$ for all t. But then $E[a_t|h_{t-1}] \in \mathcal{D} := \text{conv}\{(-1/2, 1/2), (0,1), (1,0)\}$, and $\mathcal{D}$ is disjoint from $\mathcal{C}$ and at a positive distance from it. Therefore $(1/t)\sum_{\tau \leq t} E[a_\tau|h_{\tau-1}] \in \mathcal{D}$ and so, by the Strong Law of Large Numbers, $\lim \overline{a}_t = \lim(1/t)\sum_{\tau \leq t} a_\tau \in \mathcal{D}$ too (a.s.), so $\overline{a}_t$ does not approach[17] $\mathcal{C}$.

To get some intuition, consider the deterministic system where a_t is replaced by $E[a_t|h_{t-1}]$. Then the point $(0, 1/3)$ is a stationary point for this dynamic. Specifically (see Figure 2.5), if $\overline{a}_{t-1}$ is on the line segment joining $(-1/2, 1/2)$ with $(1,0)$, then $E[\overline{a}_t|h_{t-1}]$ will be there too, moving towards $(-1/2, 1/2)$ when $\overline{a}_{t-1}$ is in the positive orthant and towards $(1,0)$ when it is in the second orthant.

Example 2.6. *The role of (D2).*

Consider the following 2-dimensional vector payoff matrix A:

	C1	C2	C3	C4
R1	$(0,1)$	$(0,0)$	$(0,-1)$	$(0,0)$
R2	$(-1,0)$	$(0,0)$	$(1,0)$	$(0,0)$
R3	$(0,0)$	$(0,-1)$	$(0,0)$	$(0,1)$
R4	$(0,0)$	$(-1,0)$	$(0,0)$	$(1,0)$

.

Again, the Row player is i and the Column player is $-i$. Let $\mathcal{C} := \{(0,0)\}$. For every $\lambda \in \mathbb{R}^2 \backslash \{(0,0)\}$, put $\mu_1 := |\lambda_1| / (|\lambda_1| + |\lambda_2|)$ and $\mu_2 := |\lambda_2| / (|\lambda_1| + |\lambda_2|)$, and define a mixed action $\sigma^{Row}(\lambda)$ for the Row player and a pure action $c(\lambda)$ for the Column player, as follows:

[17] One way to see this formally is by a separation argument: Let $f(x) := x_1 + 3x_2$; then $E[f(a_t)|h_{t-1}] \geq 1$, so $\liminf f(\overline{a}_t) = \liminf(1/t)\sum_{\tau \leq t} f(a_t) \geq 1$, whereas $f(x) \leq 0$ for all $x \in \mathcal{C}$.

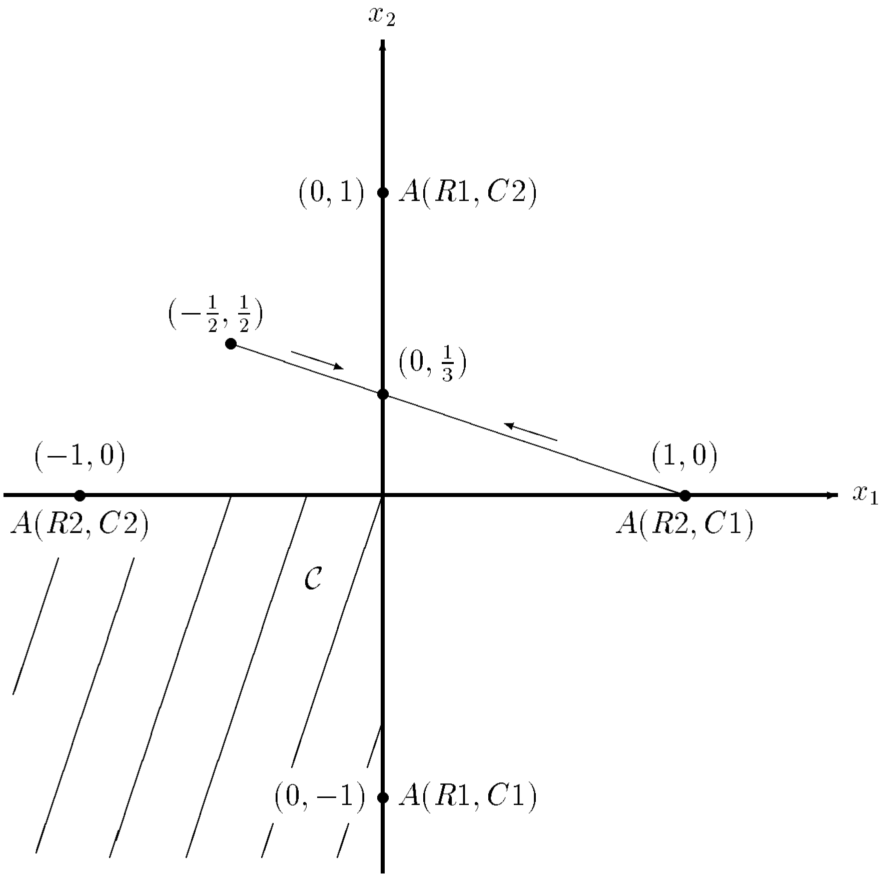

Figure 2.5: The deterministic dynamic in Example 2.4

- If $\lambda_1 \geq 0$ and $\lambda_2 \geq 0$ then $\sigma^{Row}(\lambda) := (\mu_1, \mu_2, 0, 0)$ and $c(\lambda) :=$ C1;
- If $\lambda_1 < 0$ and $\lambda_2 \geq 0$ then $\sigma^{Row}(\lambda) := (0, 0, \mu_1, \mu_2)$ and $c(\lambda) :=$ C2;
- If $\lambda_1 < 0$ and $\lambda_2 < 0$ then $\sigma^{Row}(\lambda) := (\mu_1, \mu_2, 0, 0)$ and $c(\lambda) :=$ C3; and
- If $\lambda_1 \geq 0$ and $\lambda_2 < 0$ then $\sigma^{Row}(\lambda) := (0, 0, \mu_1, \mu_2)$ and $c(\lambda) :=$ C4.

It is easy to verify that in all four cases:

(1) $\lambda \cdot A(\sigma^{Row}(\lambda), \gamma) \leq 0 = w(\lambda)$ for any action γ of the Column player; and

(2) $A(\sigma^{Row}(\lambda), c(\lambda)) = (|\lambda_1| + |\lambda_2|)^{-1} \widehat{\lambda}$, where $\widehat{\lambda} := (-\lambda_2, \lambda_1)$.

Condition (1) implies, by (2.1), that $\mathcal{C}$ is approachable by the Row player; and condition (2) means that $A(\sigma^{Row}(\lambda), c(\lambda))$ is $90°$ counterclockwise from λ.

Consider now the directional mapping Λ given by $\Lambda(x) := (x_1 + \alpha x_2, x_2 - \alpha x_1)$, where $\alpha > 0$ is a fixed constant.[18] Then (D1) and (D3) hold (for the latter, we have $x \cdot \Lambda(x) = (x_1)^2 + (x_2)^2 > 0 = w(\Lambda(x))$ for all $x \notin \mathcal{C}$), but *integrability* (D2) *is not satisfied.* To see this, assume that $\nabla P(x) = \phi(x)\Lambda(x)$ for all $x \notin \mathcal{C}$, where $\phi(x) > 0$ is a continuous function. Consider the path $y(\eta) := (\sin\eta, \cos\eta)$ for $\eta \in [0, 2\pi]$. We have $0 = P(y(2\pi)) - P(y(0)) = \int_0^{2\pi}(dP(y(\eta))/d\eta)d\eta = \int_0^{2\pi} \nabla P(y(\eta)) \cdot y'(\eta)d\eta$. But the integrand equals $(\phi(y(\eta)))(\sin\eta + \alpha\cos\eta, \cos\eta - \alpha\sin\eta) \cdot (\cos\eta, -\sin\eta) = \alpha\phi(y(\eta))$ which is everywhere positive, a contradiction.

We claim that if the Row player uses the Λ-strategy where at time t he plays[19] $\sigma^{Row}(\Lambda(\overline{a}_{t-1}))$, and if the Column player chooses $c(\Lambda(\overline{a}_{t-1}))$, then the distance to the set $\mathcal{C} = \{(0,0)\}$ does not approach 0. Indeed, let $b_t := E[a_t | h_{t-1}]$; then the strategies played imply that the vector $b_t = A(\sigma^{Row}(\Lambda(\overline{a}_{t-1})), c(\Lambda(\overline{a}_{t-1})))$ is perpendicular to $\Lambda(\overline{a}_{t-1})$ and makes an acute angle with the vector $\overline{a}_{t-1}$. Specifically,[20] $b_t \cdot \overline{a}_{t-1} = \beta \|b_t\| \|\overline{a}_{t-1}\|$, where $\beta := \alpha/\sqrt{1+\alpha^2} \leq 1$. Therefore

$$\begin{aligned}
\left(E[t\,\|\overline{a}_t\|\,|h_{t-1}]\right)^2 &\geq \|E[t\overline{a}_t|h_{t-1}]\|^2 \\
&= (t-1)^2\|\overline{a}_{t-1}\|^2 + 2(t-1)b_t \cdot \overline{a}_{t-1} + \|b_t\|^2 \\
&\geq (t-1)^2\|\overline{a}_{t-1}\|^2 + 2(t-1)\beta\|b_t\|\|\overline{a}_{t-1}\| + \beta^2\|b_t\|^2 \\
&= \left((t-1)\|\overline{a}_{t-1}\| + \beta\|b_t\|\right)^2.
\end{aligned}$$

Now $\|b_t\| \geq 1/\sqrt{2}$ (for instance, see (2)), and we have thus obtained

$$E[t\,\|\overline{a}_t\| - (t-1)\,\|\overline{a}_{t-1}\|\,|h_{t-1}] \geq \beta/\sqrt{2} > 0,$$

[18] This will provide a counterexample for any value of α, showing that it is not an isolated phenomenon.

[19] For all λ except those on the two axes (i.e., $\lambda_1 = 0$ or $\lambda_2 = 0$), the mixed action $\sigma^{Row}(\lambda)$ is uniquely determined by (2.2). If the Row player were to play in these exceptional cases another mixed action satisfying (2.2), it is easy to verify that the Column player could respond appropriately so that $\mathcal{C}$ is not approached. Thus, no Λ-strategy guarantees approachability.

[20] Writing b for b_t; x for $\overline{a}_{t-1}$; and λ for $\Lambda(\overline{a}_{t-1}) = \Lambda(x)$, we have: $x = (1+\alpha^2)^{-1}\left(\lambda + \alpha\widehat{\lambda}\right)$ (recall the definition of Λ and invert), thus $b \cdot x = (|\lambda_1| + |\lambda_2|)^{-1}\widehat{\lambda} \cdot (1+\alpha^2)^{-1}\left(\lambda + \alpha\widehat{\lambda}\right) = \alpha\left(1+\alpha^2\right)^{-1}(|\lambda_1| + |\lambda_2|)^{-1}\|\lambda\|^2$ (we have used (2), $\lambda \cdot \widehat{\lambda} = 0$ and $\|\lambda\| = \left\|\widehat{\lambda}\right\|$). Now $\|b\| = (|\lambda_1| + |\lambda_2|)^{-1}\|\lambda\|$ and $\|x\| = (1+\alpha^2)^{-1/2}\|\lambda\|$, thus completing the proof.

from which it follows that $\liminf \|\overline{a}_t\| = \liminf (1/t) \sum_{\tau \leq t} [\tau \|\overline{a}_\tau\| - (\tau - 1) \|\overline{a}_{\tau-1}\|] \geq \beta/\sqrt{2} > 0$ a.s., again by the Strong Law of Large Numbers. Thus the distance of the average payoff vector $\overline{a}_t$ from the set $\mathcal{C} = \{(0,0)\}$ is, with probability one, bounded away from 0 from some time on.

To get some intuition for this result, note that direction of movement from $\overline{a}_{t-1}$ to $E[\overline{a}_t | h_{t-1}]$ is at a fixed angle $\theta \in (0, \pi/2)$ from $\overline{a}_{t-1}$, which, if the dynamic was deterministic, would generate a counterclockwise spiral that goes away from $(0,0)$.

Example 2.7. *The role of (D3).*

Consider the 2-dimensional vector payoff matrix A

$$\begin{array}{c|c|} \cline{2-2} \text{R1} & (0,1) \\ \cline{2-2} \text{R2} & (-1,0) \\ \cline{2-2} \end{array} \quad ,$$

(where i is the Row player, and $-i$ has one action). The set $\mathcal{C} := \mathbb{R}^2_-$ is approachable by the Row player (by playing "R2 forever"). Consider the directional mapping Λ defined on $\mathbb{R}^2 \backslash \mathbb{R}^2_-$ by $\Lambda(x) := (1,0)$. Then (D1) and (D2) are satisfied (with $P(x) := x_1$), but (D3) is not: $\Lambda(0,1) \cdot (0,1) = 0 = w(\Lambda(0,1))$. Playing "R1 forever" is a Λ-strategy, but the payoff is $(0,1) \notin \mathcal{C}$.

3. Regrets

3.1. Model and Preliminary Results

In this section we consider standard N-person games in strategic form (with *scalar* payoffs for each player). The set of players is a finite set N, the action set of each player i is a finite set S^i, and the payoff function of i is $u^i : S \to \mathbb{R}$, where $S := \Pi_{j \in N} S^j$; we will denote this game $< N, (S^i)_i, (u^i)_i >$ by Γ.

As in the previous section, the game is played repeatedly in discrete time $t = 1, 2, ...$; denote by $s^i_t \in S^i$ the choice of player i at time t, and put $s_t = (s^i_t)_{i \in N} \in S$. The payoff of i in period t is $U^i_t := u^i(s_t)$, and $\overline{U}^i_t := (1/t) \sum_{\tau \leq t} U^i_\tau$ is his average payoff up to t.

Fix a player $i \in N$. Following Hannan [1957], we consider the *regrets* of player i, namely, for each one of his actions $k \in S^i$, the change in his average payoff if he were always to choose k (while no one else makes any change in his realized actions):

$$D_t^i(k) := \frac{1}{t}\sum_{\tau=1}^{t} u^i(k, s_\tau^{-i}) - \overline{U}_t^i = u^i(k, z_t^{-i}) - \overline{U}_t^i,$$

where $z_t^{-i} \in \Delta(S^{-i})$ is the empirical distribution of the actions chosen by the other players in the past.[21] A strategy of player i is called *Hannan-consistent* if, as t increases, all regrets are guaranteed—no matter what the other players do—to become almost surely non-positive in the limit; that is, with probability one, $\limsup_{t\to\infty} D_t^i(k) \leq 0$ for all $k \in S^i$.

Following Hart and Mas-Colell [2000], it is useful to view the regrets of i as an m-dimensional vector payoff, where $m := |S^i|$. We thus define $A \equiv A^i : S \to \mathbb{R}^m$, *the i-regret vector-payoff game associated to* Γ, by

$$\begin{aligned} A_k(s^i, s^{-i}) &:= u^i(k, s^{-i}) - u^i(s^i, s^{-i}) \text{ for all } k \in S^i, \text{ and} \\ A(s^i, s^{-i}) &:= (A_k(s^i, s^{-i}))_{k\in S^i}, \end{aligned}$$

for all $s = (s^i, s^{-i}) \in S^i \times S^{-i} = S$. Rewriting the regret as

$$D_t^i(k) = \frac{1}{t}\sum_{\tau\leq t}\left[u^i(k, s_\tau^{-i}) - u^i(s_\tau^i, s_\tau^{-i})\right]$$

shows that the vector of regrets at time t is just the average of the A vector payoffs in the first t periods: $D_t^i = (1/t)\sum_{\tau\leq t} A(s_\tau)$. The existence of a Hannan-consistent strategy in Γ is thus equivalent to the approachability by player i of the non-positive orthant $\mathbb{R}_-^{S^i}$ in the vector-payoff game A, and a strategy is Hannan-consistent if and only if it guarantees that $\mathbb{R}_-^{S^i}$ is approached.

We now present two important results that apply in all generality to the regret setup.

Proposition 3.1. *For any (finite) N-person game Γ, the non-positive orthant $\mathbb{R}_-^{S^i}$ is approachable by player i in the i-regret vector-payoff associated game.*

This Proposition follows immediately from the next one. Observe that the approachability of $\mathbb{R}_-^{S^i}$ is equivalent, by the Blackwell condition (2.1), to: For

[21] I.e., $z_t^{-i}(s^{-i}) := \left|\{\tau \leq t : s_\tau^{-i} = s^{-i}\}\right|/t$ for each $s^{-i} \in S^{-i}$, where we write $|\Omega|$ for the cardinality of the set Ω.

every $\lambda \in \Delta(S^i)$ there exists $\sigma^i(\lambda) \in \Delta(S^i)$, a mixed action of player i, such that

$$\lambda \cdot A(\sigma^i(\lambda), s^{-i}) \leq 0 \text{ for all } s^{-i} \in S^{-i} \tag{3.1}$$

(indeed, $w(\lambda)$ equals 0 for $\lambda \geq 0$ and it is infinite otherwise). That is, the expected regret obtained by playing $\sigma^i(\lambda)$ lies in the half-space (through the origin) with normal λ. In this regret setup, the mixture $\sigma^i(\lambda)$ may be actually chosen in a simple manner:

Proposition 3.2. *For any (finite) N-person game Γ and every $\lambda \in \Delta(S^i)$, condition (3.1) is satisfied by $\sigma^i(\lambda) = \lambda$.*

Proof. Given $\lambda \in \Delta(S^i)$, a $\sigma^i \equiv (\sigma^i_k)_{k \in S^i} \in \Delta(S^i)$ satisfies (3.1) if and only if

$$\sum_{k \in S^i} \lambda_k \sum_{j \in S^i} \sigma^i_j [u^i(k, s^{-i}) - u^i(j, s^{-i})] \leq 0 \tag{3.2}$$

for all $s^{-i} \in S^{-i}$. This may be rewritten as

$$\sum_{k \in S^i} u^i(k, s^{-i})[\lambda_k \sum_{j \in S^i} \sigma^i_j - \sigma^i_k \sum_{j \in S^i} \lambda_j] = \sum_{k \in S^i} u^i(k, s^{-i})[\lambda_k - \sigma^i_k] \leq 0.$$

Therefore, by choosing σ^i so that all coefficients in the square brackets vanish—that is, by choosing $\sigma^i_k = \lambda_k$—we guarantee (3.2) and thus (3.1) for all s^{-i}. ∎

3.2. Regret-Based Strategies

The general theory of Section 2 is now applied to the regret situation. A *stationary regret-based* strategy for player i is a strategy of i such that the choices depend only on i's regret vector; that is, for every history h_{t-1}, the mixed action of i at time t is a function[22] of D^i_{t-1} only: $\sigma^i_t = \sigma^i(D^i_{t-1}) \in \Delta(S^i)$. The main result of this section is:

Theorem 3.3. *Consider a stationary regret-based strategy of player i given by a mapping $\sigma^i : \mathbb{R}^{S^i} \to \Delta(S^i)$ that satisfies:*

[22]Notice that the time t does not matter: The strategy is "stationary."

(R1) *There exists a continuously differentiable function* $P : \mathbb{R}^{S^i} \to \mathbb{R}$ *such that* $\sigma^i(x)$ *is positively proportional to* $\nabla P(x)$ *for every* $x \notin \mathbb{R}_-^{S^i}$*; and*

(R2) $\sigma^i(x) \cdot x > 0$ *for every* $x \notin \mathbb{R}_-^{S^i}$.

Then this strategy is Hannan-consistent for any (finite) N-person game.

Proof. Apply Theorem 2.1 for $\mathcal{C} = \mathbb{R}_-^{S^i}$ together with Propositions 3.1 and 3.2: (D1) and (D2) yield (R1), and (D3) yields (R2). ■

We have thus obtained a wide class of strategies that are Hannan-consistent. It is noteworthy that these are "universal" strategies: The mapping σ^i is independent of the game (see also the "variable game" case in Section 5).

Condition (R2) says that when $D_{t-1}^i \notin \mathbb{R}_-^{S^i}$—i.e., when some regret is positive—the mixed choice σ_t^i of i satisfies $\sigma_t^i \cdot D_{t-1}^i > 0$. This is equivalent to

$$u^i(\sigma_t^i, z_{t-1}^{-i}) > \overline{U}_{t-1}^i. \tag{3.3}$$

That is, the expected payoff of i from playing σ_t^i against the empirical distribution z_{t-1}^{-i} of the actions chosen by the other players in the past, is higher than his realized average payoff. Thus σ_t^i is a *better-reply*, where "better" is relative to the obtained payoff. By comparison, fictitious play always chooses an action that is a *best-reply* to the empirical distribution z_{t-1}^{-i}. For more on this "better vs. best" issue, see Subsection 4.2 below and Hart and Mas-Colell [2000, Section 4(e)].

We now describe a number of interesting special cases, in order of increasing generality.

1. l_2 *potential*: $P(x) = \left(\sum_{k \in S^i} ([x_k]_+)^2\right)^{1/2}$. This yields (after normalization) $\Lambda(x) = (1/\left\|[x]_+\right\|_1)[x]_+$ for $x \notin \mathbb{R}_-^{S^i}$, and the resulting strategy is $\sigma_t^i(k) = [D_{t-1}^i(k)]_+ / \sum_{k' \in S^i} [D_{t-1}^i(k')]_+$ when $D_{t-1}^i \notin \mathbb{R}_-^{S^i}$. This is the Hannan-consistent strategy introduced in Hart and Mas-Colell [2000, Theorem C], where the play probabilities are proportional to the positive regrets.

2. l_p *potential:* $P(x) = \left(\sum_{k \in S^i} ([x_k]_+)^p\right)^{1/p}$ for some $1 < p < \infty$, which yields $\sigma_t^i(k) = ([D_{t-1}^i(k)]_+)^{p-1} / \sum_{k' \in S^i} ([D_{t-1}^i(k')]_+)^{p-1}$; i.e., play probabilities that are proportional to a fixed positive power $(p-1 > 0)$ of the positive regrets.

3. *Separable potential:* A *separable* strategy is one where σ_t^i is proportional to a vector whose k-th coordinate depends *only* on the k-th regret; i.e., σ_t^i is proportional to a vector of the form $(\psi_k(D_{t-1}^i(k)))_{k\in S^i}$. Conditions (R1) and (R2) result in the following requirements[23]: For each k in S^i, the function $\psi_k : \mathbb{R} \to \mathbb{R}$ is continuous; $\psi_k(x_k) > 0$ for $x_k > 0$; and $\psi_k(x_k) = 0$ for $x_k \leq 0$. The corresponding potential is $P(x) = \sum_{k\in S^i} \Psi_k(x_k)$, where $\Psi_k(x) := \int_{-\infty}^{x} \psi_k(y)dy$. Note that, unlike the previous two cases, the functions ψ_k may differ for different k, and they need not be monotonic (thus a higher regret may not lead to a higher probability).

Finally, observe that in all the above cases, actions with negative or zero regret are never chosen. This need no longer be true in the general (non-separable) case; see Subsection 4.2 below.

3.3. Counterexamples

The counterexamples of Subsection 2.3 translate easily into the regret setup.

- *The role of "better" (R2).* Consider the 1-person game

R1	0
R2	1

.

  The resulting regret game is given in Example 2.7. The strategy of playing "R1 forever" satisfies condition (R1) but not condition (R2) (or (3.3)), and it is indeed not Hannan-consistent.

- *The role of continuity in (R1).* Consider the simplest 2-person coordination game (a well-known stumbling block for many strategies).

	C1	C2
R1	$1, 1$	$0, 0$
R2	$0, 0$	$1, 1$

.

  The resulting regret game for the Row player is precisely the vector-payoff game of Examples 2.4 and 2.5, where we looked at the approachability question for the non-positive orthant. The two strategies we considered there—which we have shown not to be Hannan-consistent—are not continuous.

[23]Consider points x with $x_j = \pm\varepsilon$ for all $j \neq k$.

They correspond to the l_∞ and the l_1 potentials, respectively, which are not differentiable. (Note in particular that the l_∞ case yields "fictitious play," which is further discussed in Subsection 4.1 below.)

- *The role of integrability in (R1).* The vector-payoff game of our Example 2.6 can easily be seen to be a regret game. However, the approachable set there was not the non-positive orthant. In order to get a counterexample to the result of Theorem 3.3 when integrability is not satisfied, one would need to resort to additional dimensions, that is, more than 2 strategies; we do not do it here, although it is plain that such examples are easy—though painful—to construct.

4. Fictitious Play and Better Play

4.1. Fictitious Play and Smooth Fictitious Play

As we have already pointed out, fictitious play may be viewed as a stationary regret-based strategy, corresponding to the l_∞ mapping (the directional mapping generated by the l_∞ potential). It does not guarantee Hannan-consistency (see Example 2.4 and Subsection 3.3); the culprit for this is the lack of continuity (i.e., (D1)).

Before continuing the discussion it is useful to note a property of fictitious play: *The play at time t does not depend on the realized average payoff $\overline{U}^i_{t-1}$.* Indeed, $\max_k D^i_{t-1}(k) = \max_k u^i(k, z^{-i}_{t-1}) - \overline{U}^i_{t-1}$, so an action $k \in S^i$ maximizes regret if and only if it maximizes the payoff against the empirical distribution z^{-i}_{t-1} of the actions of $-i$. In the general approachability setup of Section 2 (with $\mathcal{C} = \mathbb{R}^{S^i}_-$), this observation translates into the requirement that the directional mapping Λ be invariant to adding the same constant to all coordinates. That is, writing $e := (1, 1, ..., 1) \in \mathbb{R}^{S^i}$,

$$\Lambda(x) = \Lambda(y) \text{ for any } x, y \notin \mathbb{R}^{S^i}_- \text{ with } x - y = \alpha e \text{ for some scalar } \alpha. \tag{4.1}$$

Note that, as it should be, the l_∞ mapping satisfies this property (4.1).

Proposition 4.1. *A directional mapping Λ satisfies (D2), (D3) and (4.1) for*

$\mathcal{C} = \mathbb{R}^m_-$ *if and only if it is equivalent to the* l_∞ *mapping, i.e., its potential* P *satisfies* $P(x) = \phi(\max_k x_k)$ *for some strictly increasing function* ϕ.

Proof. Since $\mathcal{C} = \mathbb{R}^m_-$, the allowable directions are $\lambda \geq 0, \lambda \neq 0$. Thus $\nabla P(x) \geq 0$ for a.e. $x \notin \mathbb{R}^m_-$ by (D2), implying that the limit of $\nabla P(x) \cdot x$ is ≤ 0 as x approaches the boundary of $\mathbb{R}^m_-$. But $\nabla P(x) \cdot x > 0$ for a.e. $x \notin \mathbb{R}^m_-$ by (D3), implying that the limit of $\nabla P(x) \cdot x$ is in fact 0 as x approaches bd $\mathbb{R}^m_-$. Because P is Lipschitz, it follows that P is constant on bd $\mathbb{R}^m_-$, i.e., $P(x) = P(0)$ for every $x \in$ bd $\mathbb{R}^m_-$. By (D3) again we have $P(x) > P(0)$ for all $x \notin \mathbb{R}^m_-$. Adding to this the invariance condition (4.1) implies that the level sets of P are all translates by multiples of e of bd $\mathbb{R}^m_- = \{x \in \mathbb{R}^m : \max_k x_k = 0\}$. ■

Since the l_∞ mapping does not guarantee that $\mathcal{C} = \mathbb{R}^m_-$ is approached (again, see Example 2.4 and Subsection 3.3), we have:

Corollary 4.2. *There is no stationary regret-based strategy that satisfies (R1) and (R2) and is independent of realized average payoff.*

The import of the Corollary (together with the indispensability of conditions (R1) and (R2), as shown by the counterexamples in Subsection 3.3) is that one cannot simultaneously have independence of realized payoffs and guarantee Hannan-consistency in every game.

We must weaken one of the two properties. One possibility is to weaken the Hannan consistency requirement to ε-consistency: $\limsup_t D^i_t(k) \leq \varepsilon$ for all k. Fudenberg and Levine [1995] propose a smoothing of fictitious play that—like fictitious play itself—is independent of realized payoffs. In essence, their function P is convex, smooth, and satisfies the property that its level sets are obtained from each other by translations along the $e = (1, ..., 1)$ direction (see Figure 4.1).[24] The level set of P through 0 is therefore smooth; it is very close to the boundary of the negative orthant but unavoidably distinct from it. The resulting strategy approaches $\mathcal{C} = \{x : P(x) \leq P(0)\}$ (recall Remark 5 in Subsection 2.1: a set of the form $\{x : P(x) \leq c\}$, for constant c, is approachable when $c \geq P(0)$—since it

[24]Specifically, $P(x) = \max_{\sigma^i \in \Delta(S^i)} \{\sigma^i \cdot x + \varepsilon v(\sigma^i)\}$, where $\varepsilon > 0$ is small and v is a strictly differentiably concave function, with gradient vector approaching infinite length as one approaches the boundary of $\Delta(S^i)$.

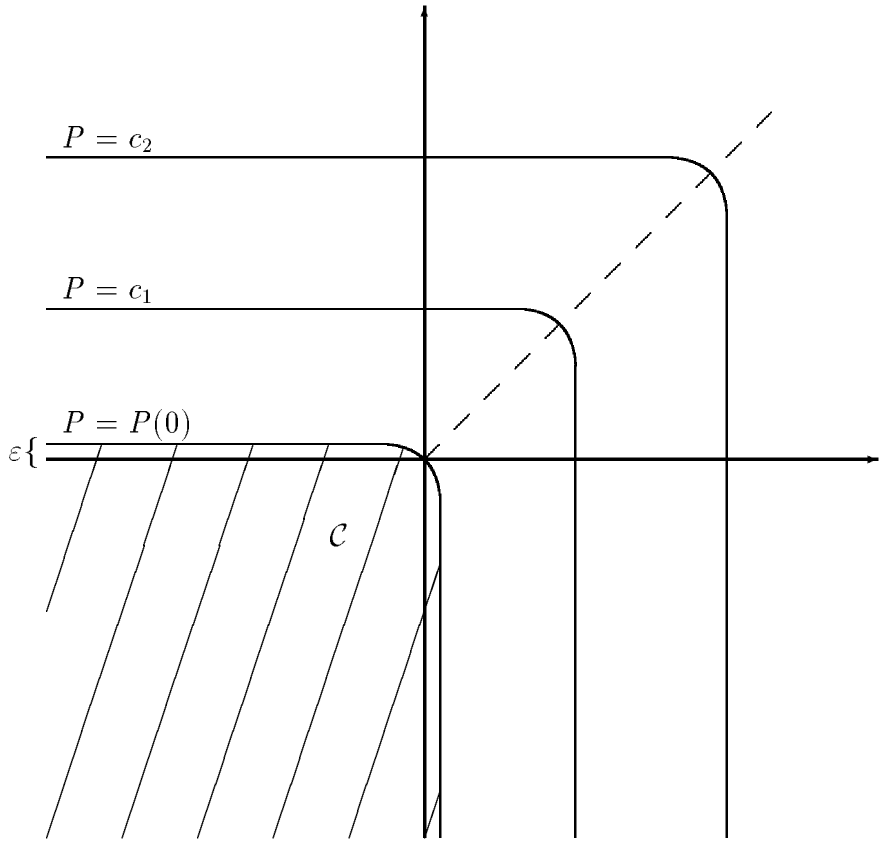

Figure 4.1: The level sets of the potential of smooth fictitious play

contains $\mathbb{R}_-^{S^i}$—and is not approachable when $c < P(0)$—since it does not contain 0). The set $\mathcal{C}$ is strictly larger than $\mathbb{R}_-^{S^i}$; it is an ε-neighborhood of the negative orthant $\mathbb{R}_{--}^{S^i}$. Thus one obtains only ε-consistency.[25,26]

The other possibility is to allow the strategy to depend also on the realized

[25]Other smoothings have been proposed, including Hannan [1957], Foster and Vohra [1993, 1998] and Freund and Schapire [1999] (in the later—which corresponds to exponential smoothing—the strategy is non-stationary, i.e., it depends not only on the point in regret space but also on the time t; of course, non-stationary strategies where ε decreases with t may yield exact consistency).

[26]Smooth fictitious play may be equivalently viewed, in our framework, as first taking a set $\mathcal{C}$ that is close to the negative orthant and has smooth boundary, and then using the l_∞ distance from $\mathcal{C}$ as a potential (recall Remark 6 in Subsection 2.1).

payoffs. Then there are strategies that are close to fictitious play and guarantee Hannan-consistency in any game. Take, for instance, the l_p potential strategy for large enough p (see Subsection 3.2).[27]

4.2. Better Play

All the examples presented up to now satisfy an additional natural requirement, namely, that only actions with positive regret are played (provided, of course, that there are such actions). Formally, consider a stationary regret-based strategy of player i that is given by a mapping $\sigma^i : \mathbb{R}^{S^i} \to \Delta(S^i)$ (see Theorem 3.3); we add to (R1) and (R2) the following condition:[28]

(R3) For every $x \notin \mathbb{R}^{S^i}_{--}$, if $x_k < 0$ then $[\sigma^i(x)]_k = 0$.

Since x is the i-regret vector, (R3) means that σ^i gives probability 1 to the set of actions with non-negative regret (unless all regrets are negative, in which case there is no requirement[29]). This may be rewritten as

$$[\sigma^i_t]_k > 0 \text{ only if } u^i(k, z^{-i}_{t-1}) \geq \overline{U}^i_{t-1}. \tag{4.2}$$

That is, only those actions k are played whose payoff against the empirical distribution z^{-i}_{t-1} of the opponents' actions is at least as large as the actual realized average payoff $\overline{U}^i_{t-1}$; in short, only the "better actions."[30] For an example where (R3) is *not* satisfied, see Figure 4.2.

The l_p potential strategies, for $1 < p < \infty$, and in fact all separable strategies (see Subsection 3.2) essentially[31] satisfy (R3). Fictitious play (with the l_∞ potential) also satisfies (R3): The action chosen is a "best" one (rather than just "better"). At the other extreme, the l_1 potential strategy gives equal probability

[27]This amounts to smoothing the norm and keeping $\mathcal{C}$ equal to the negative orthant, whereas the previous construction smoothed the boundary of $\mathcal{C}$ and kept the l_∞ norm. These are the two "dual" ways of generating a smooth potential (again, see Remark 6 in Subsection 2.1).

[28]Note that the condition needs to be satisfied not only for $x \notin \mathbb{R}^{S^i}_-$, but also for $x \in \text{bd}\, \mathbb{R}^{S^i}_-$.

[29]See Footnote 33 below.

[30]Observe that (R2) (or (3.3)) is a requirement on the average over all played actions k, whereas (R3) (or (4.2)) applies to each such k separately.

[31]The condition in (R3) is automatically satisfied in these cases for $x \notin \mathbb{R}^{S^i}$; one needs to impose it explicitly for $x \in \text{bd}\, \mathbb{R}^{S^i}_-$ (where, up to now, we had no requirements).

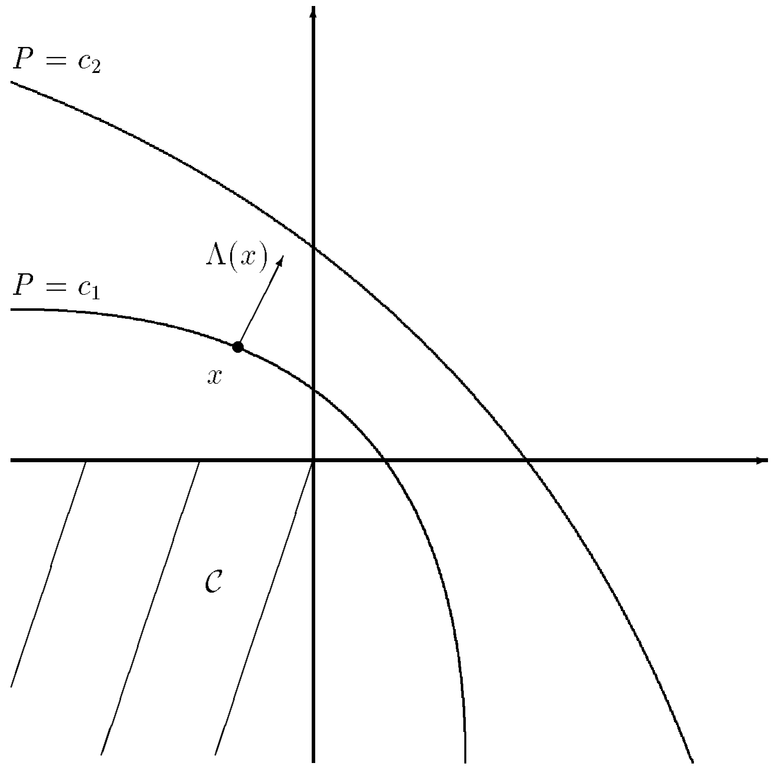

Figure 4.2: (R3) is not satisfied

to *all* better actions, so it also satisfies (R3). (However, these last two do not satisfy (R1)).

Using condition (R3) yields the following result, which is a generalization of Theorem A of Monderer, Samet and Sela [1997] for fictitious play:

Proposition 4.3. *Consider a stationary regret-based strategy of player i given by a mapping $\sigma^i : \mathbb{R}^{S^i} \to \Delta(S^i)$ that satisfies*[32] *(R3). Then, in any (finite) N-person game, the maximal regret of i is always non-negative:*

$$\max_{k \in S^i} D_t^i(k) \geq 0 \text{ for all } t.$$

[32]Note that (R1) and (R2) are not assumed.

Proof. By induction, starting with $D_0^i(k) = 0$ for all k. Assume that $\max_k D_{t-1}^i(k) \geq 0$ (or, $D_{t-1}^i \notin \mathbb{R}_{--}^{S^i}$). By (R3), $D_{t-1}^i(k) \geq 0$ for any k chosen at time t; since $A_k(k, s_t^{-i}) = 0$ it follows that $D_t^i(k) = (1/t)((t-1)D_{t-1}^i(k) + t\, 0) \geq 0$. ■

Thus, the vector of regrets never enters the negative orthant.[33] Recall that the result of Theorem 3.3 is that the vector of regrets approaches the non-positive orthant. To combine the two, we define *better play* as any stationary regret-based strategy of player i that is given by a mapping σ^i satisfying (R1)–(R3). We thus have:

Corollary 4.4. *In any (finite) N-person game, if player i uses a better play strategy, then his maximal regret converges to 0 a.s.*

$$\lim_{t\to\infty} \max_{k\in S^i} D_t^i(k) = \lim_{t\to\infty}\left(\max_{k\in S^i} u^i(k, z_t^{-i}) - \overline{U}_t^i\right) = 0 \text{ a.s.}$$

That is, the average payoff $\overline{U}_t^i$ of player i up to time t is close, as $t \to \infty$, to i's best-reply payoff against the empirical distribution of the other players' actions. In particular, in a two-person zero-sum game we obtain the following:

Corollary 4.5. *In any (finite) two-person zero-sum game, if both players use better play strategies, then:*

(i) *For each player, the empirical distribution of play converges to the set of optimal actions.*[34]

(ii) *The average payoff converges to the value of the game.*

Proof. Let 1 be the maximizer and 2 the minimizer, and denote by v the minimax value of the game. Then $\max_{k\in S^1} u^1(k, z_t^2) \geq v$, so by Corollary 4.4 we have $\liminf_t \overline{U}_t^1 \geq v$. The same argument for player 2 yields the opposite inequality, thus $\lim_t \overline{U}_t^1 = v$. Therefore $\lim_t \max_{k\in S^1} u^1(k, z_t^2) = v$ (apply the Corollary again), hence any limit point of the sequence z_t^2 must be an optimal action of player 2; similarly for player 1. ■

[33] Therefore at every period there are always actions with non-negative regret—out of which the next action is chosen (and so the condition $x \notin \mathbb{R}_{--}^{S^i}$ in (R3) always holds).

[34] That is, the set of mixed actions that guarantee the value.

Thus, better play enjoys the same properties as fictitious play in two-person zero-sum games (for fictitious play, see Robinson [1951] for the convergence to the set of optimal strategies, and see Monderer, Samet and Sela [1997, Theorem B] and Rivière [1997] for the convergence of the average payoff).

5. Discussion and Extensions

In this section we discuss a number of extensions of our results.

1. *Conditional regrets*

As stated in the Introduction, we have been led to the "no regret" Hannan-consistency property from considerations of "no conditional regret" that correspond to correlated equilibria (see Hart and Mas-Colell [2000]). Given two actions k and j of player i, the *conditional regret* from j to k is the change that would have occurred in the average payoff of i if he had played action k in all those periods where he did play j (and everything else is left unchanged). That is,

$$DC_t^i(j,k) := \frac{1}{t} \sum_{\tau \leq t : s_\tau^i = j} \left[u^i(k, s_\tau^{-i}) - u^i(s_\tau) \right]$$

for every[35] $j, k \in S^i$. The vector of regrets DC_t^i is now in $\mathbb{R}^L$, where $L := S^i \times S^i$, and the empirical distribution of actions up to time t constitutes a correlated equilibrium if and only if $DC_t^i \leq 0$. Thus, the set to be approached is the non-positive orthant $\mathbb{R}^L_-$. The corresponding game with vector payoffs A is defined as follows: The (j,k) coordinate of the vector payoff $A(s^i, s^{-i}) \in \mathbb{R}^L$ is $u^i(k, s^{-i}) - u^i(j, s^{-i})$ when $s^i = j$, and it is 0 otherwise; hence $DC_t^i = (1/t) \sum_{\tau \leq t} A(s_\tau)$.

As in Propositions 3.1 and 3.2 (see Hart and Mas-Colell [2000, Section 3]), it can easily be verified that:

- $\mathcal{C} = \mathbb{R}^L_-$ is always approachable.

[35] Note that each Hannan regret $D_t^i(k)$ is the sum of the conditional regrets $DC_t^i(j,k)$ over $j \neq k$. Thus the set of distributions of N-tuples of actions that satisfy the Hannan no-regret conditions includes the set of correlated equilibrium distributions. The inclusion is, in general, strict (the two sets coincide when every player has only two actions).

- For every $\lambda \in \mathbb{R}^L_+$, the Blackwell approachability condition for $\mathcal{C} = \mathbb{R}^L_-$ ((2.1) or (3.1)) holds for any mixed action $\sigma^i = (\sigma^i_k)_{k \in S^i} \in \Delta(S^i)$ that satisfies

$$\sum_{j \in S^i} \sigma^i_j \lambda(j,k) = \sigma^i_k \sum_{j \in S^i} \lambda(k,j) \text{ for all } k \in S^i. \tag{5.1}$$

  Viewing λ as an $S^i \times S^i$ matrix, condition (5.1) says that σ^i is an invariant vector for the (non-negative) matrix λ.

- For every $\lambda \in \mathbb{R}^L_+$, there exists a $\sigma^i \in \Delta(S^i)$ satisfying (5.1).

Applying Theorem 2.1 yields a large class of strategies. For example (as in Subsection 3.2), if P is the l_p potential for some $1 < p < \infty$, then σ^i is an invariant vector of the matrix of the $p-1$ powers of the non-negative regrets.[36] In the more general separable case, σ^i is an invariant vector of the matrix whose (j,k) coordinate is $\psi_{(j,k)}(DC^i_{t-1}(j,k))$, where $\psi_{(j,k)}$ is any real continuous function which vanishes for $x \leq 0$ and is positive for $x > 0$. As in Hart and Mas-Colell [2000, Theorem A], if every player uses a strategy in this class (of course, different players may use different types of strategies), then the empirical distribution of play converges to the set of correlated equilibria of Γ.

Since finding invariant vectors is by no means a simple matter, in Hart and Mas-Colell [2000] much effort is devoted to obtaining simple adaptive procedures, which use the matrix of regrets as a one-step transition matrix. To do the same here, one would use instead the matrix $\Lambda\left(DC^i_{t-1}\right)$.

2. *Variable game*

We noted in Subsection 3.2 that our strategies are game-independent. This allows us to consider the case where, at each period, a different game is being played (for example, a stochastic game). The strategy set of player i is the same set S^i in all games, but he does not know which game is currently being played. All our results—in particular, Theorem 3.3—continue to hold[37] provided player i is told, *after* each period t, which game was played at time t and what were the

[36] For $p = 2$ we get the matrix of regrets—which yields precisely Theorem A of Hart and Mas-Colell [2000].

[37] Assuming the payoffs of i are uniformly bounded.

chosen actions s_t^{-i} of the other players. Indeed, as in Section 3, i can then compute the vector $a := A(s_t^i, s_t^{-i}) \in \mathbb{R}^{S^i}$, update his regret vector: $D_t^i = (1/t)((t-1)D_{t-1}^i + a)$, and then play $\sigma^i(D_t^i)$ in the next period, where σ^i is any mapping satisfying (R1) and (R2).

3. *Unknown game*

When the player does not know the (fixed) game Γ that is played and is told, at each stage, only his own realized payoff (but not the choices of the other players; this may be referred to as a "stimulus-response" model), Hannan-consistency may nonetheless be obtained (see Foster and Vohra [1993, 1998], Auer *et al.* [1995], Fudenberg and Levine [1998, Section 4.8], Hart and Mas-Colell [2000, Section 4(j)], and also Baños [1968] and Megiddo [1980] for related work). For instance, one can replace the regrets—which cannot be computed here—by appropriate estimates.[38]

References

Auer P., N. Cesa-Bianchi, Y. Freund and R. E. Schapire [1995], Gambling in a Rigged Casino: The Adversarial Multi-Armed Bandit Problem, *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, 322–331.

Baños, A. [1968], On Pseudo-Games, *The Annals of Mathematical Statistics* 39, 1932–1945.

Blackwell, D. [1956a], An Analog of the Minmax Theorem for Vector Payoffs, *Pacific Journal of Mathematics* 6, 1–8.

Blackwell, D. [1956b], Controlled Random Walks, *Proceedings of the International Congress of Mathematicians* 1954, Vol. III, North-Holland, 335–338.

Borodin, A. and R. El-Yaniv [1998], *Online Computation and Competitve Analysis*, Cambridge University Press.

Clarke, F. [1983], *Optimization and Nonsmooth Analysis*, Wiley.

[38]Specifically, use $\widehat{D}_t^i(k) := (1/t)\sum_{\tau \le t: s_\tau^i = k}(1/[\sigma_\tau^i]_k)u^i(s_\tau) - \overline{U}_t^i$ instead of $D_t^i(k)$ (see Hart and Mas-Colell [2000]).

Foster, D. and R. V. Vohra [1993], A Randomized Rule for Selecting Forecasts, *Operations Research* 41, 704–709.

Foster, D. and R. V. Vohra [1997], Calibrated Learning and Correlated Equilibrium, *Games and Economic Behavior* 21, 40–55.

Foster, D. and R. V. Vohra [1998], Asymptotic Calibration, *Biometrika* 85, 379–390.

Foster, D. and R. V. Vohra [1999], Regret in the On-Line Decision Problem, *Games and Economic Behavior* 29, 7–35.

Freund, Y. and R. E. Schapire [1999], Adaptive Game Playing Using Multiplicative Weights, *Games and Economic Behavior* 29, 79–103.

Fudenberg, D. and D. K. Levine [1995], Universal Consistency and Cautious Fictitious Play, *Journal of Economic Dynamics and Control* 19, 1065–1090.

Fudenberg, D. and D. K. Levine [1998], *Theory of Learning in Games*, MIT Press.

Fudenberg, D. and D. K. Levine [1999], Conditional Universal Consistency, *Games and Economic Behavior* 29, 104–130.

Hannan, J. [1957], Approximation to Bayes Risk in Repeated Play, in *Contributions to the Theory of Games*, Vol. III (*Annals of Mathematics Studies* 39), M. Dresher, A. W. Tucker and P. Wolfe (eds.), Princeton University Press, 97–139.

Hart, S. and A. Mas-Colell [2000], A Simple Adaptive Procedure Leading to Correlated Equilibrium, *Econometrica.*

Littlestone, N. and M. K. Warmuth [1994], The Weighted Majority Algorithm, *Information and Computation* 108, 212–261.

Loève, M. [1978], *Probability Theory*, Vol. II, 4th Edition, Springer-Verlag.

Luce, R. D. and H. Raiffa [1957], *Games and Decisions*, Wiley.

Megiddo, N. [1980], On Repeated Games with Incomplete Information Played by Non-Bayesian Players, *International Journal of Game Theory* 9, 157–167.

Monderer, D., D. Samet and A. Sela [1997], Belief Affirming in Learning Processes, *Journal of Economic Theory* 73, 438–452.

Rivière, P. [1997], Quelques Modèles de Jeux d'Evolution, Ph.D. thesis, Université Paris 6.

Robinson, J. [1951], An Iterative Method of Solving a Game, *Annals of Mathematics* 54, 296–301.