

Learning correlated equilibria in games with compact sets of strategies [★]

Gilles Stoltz ¹

`gilles.stoltz@ens.fr`

*Département de Mathématiques et Applications,
Ecole Normale Supérieure, 75005 Paris, France*

Gábor Lugosi

`lugosi@upf.es`

*Department of Economics, Pompeu Fabra University,
Ramon Trias Fargas 25-27, 08005 Barcelona, Spain*

Abstract

Hart and Schmeidler's extension of correlated equilibrium to games with infinite sets of strategies is studied. General properties of the set of correlated equilibria are described. It is shown that, just like for finite games, if all players play according to an appropriate regret-minimizing strategy then the empirical frequencies of play converge to the set of correlated equilibria whenever the strategy sets are convex and compact.

[★] This article was presented at the Second World Congress of the Game Theory Society (Marseille, July 2004). Part of the work of the first author was done while he was visiting University Pompeu Fabra. The work of the second author was supported by the Spanish Ministry of Science and Technology and FEDER, grant BMF2003-03324. The authors acknowledge support by the PASCAL Network of Excellence under EC grant no. 506778.

¹ Gilles Stoltz is the corresponding author. His phone and fax numbers are respectively (+33) 1 44 32 32 77, and (+33) 1 44 32 20 80.

1 Introduction

Correlated equilibrium, introduced by Aumann [4,5] is arguably one of the most natural notions of equilibrium. A correlated equilibrium is a joint distribution π over the set of strategies of the players that has the property that if, before taking an action, each player receives a recommendation such that the recommendations are drawn randomly according to the joint distribution of π , then no player has an incentive to divert from the recommendation, provided that all other players follow theirs. The distinguishing feature of the notion is that, unlike in the definition of Nash equilibria, the recommendations do not need to be independent. A correlated equilibrium π is a Nash equilibrium if and only if π is a product measure.

A remarkable property of correlated equilibrium, pointed out by Foster and Vohra [12], is that if the game is repeated infinitely many times such that every player plays according to a certain regret-minimization strategy, then the empirical frequencies of play converge to the set of correlated equilibria. (See also Fudenberg and Levine [14], Hart and Mas-Colell [16–18].) No coordination is necessary between the players, and the players do not even need to know the others' payoff functions. Hart and Mas-Colell [19] show that Nash equilibrium does not share this property unless the game has quite special properties.

The purpose of this paper is to study the correlated equilibria of a large class of infinite games. In Section 2 we recall Hart and Schmeidler's extended definition, and propose some equivalent formulations. One of them may be given by discretizing the sets of strategies, considering correlated equilibria of the discretized (finite) games, and taking appropriate limits as the discretization becomes finer (see Theorem 5).

The main result of the paper (Theorems 6 and 7) generalizes the above-mentioned result of Foster and Vohra to the case when the sets of strategies are compact and convex subsets of a normed space, and the payoff function of each player is continuous. It is shown that convergence of the empirical frequencies of play to the set of correlated equilibria can also be achieved in this case, by playing internal regret-minimizing strategies, where Foster and Vohra's notion of internal regret [12] has to be generalized to the case of games with infinite strategy sets.

The proof of the main theorem is given in Section 4 by a sequence of results, by broadening the class of departure functions in each step. We point out in Section 5 that, contrary to the case of finite games, the size of the classes of departure functions plays a role in the rate of convergence to the set of correlated equilibria. We conclude the paper in Section 6 by establishing a connection between the correlated equilibria of a finite game and those of its

mixed extension. We show that, in some sense, these are equivalent.

2 Definition of correlated equilibrium

2.1 Refined definition

The notion of correlated equilibrium was introduced by Aumann [4,5] who assumed that the sets of strategies are finite, and extended later by Hart and Schmeidler [20] to infinite games.

Formally, consider an N -person game in strategic (normal) form

$$\Gamma = \left(\{1, \dots, N\}, (S^k)_{1 \leq k \leq N}, (h^k)_{1 \leq k \leq N} \right) ,$$

where $\{1, \dots, N\}$ is the finite set of players, player k is given a (not necessarily finite) set of strategies S^k and a payoff function $h^k : S \rightarrow \mathbb{R}$. The set of N -tuples of strategies is denoted by $S = S^1 \times S^2 \times \dots \times S^N$. We use the notation $s = (s^{-k}, s^k)$, where

$$s^{-k} = (s^1, \dots, s^{k-1}, s^{k+1}, \dots, s^N)$$

denotes the strategies played by everyone but player k . We write $s^{-k} \in S^{-k}$, where $S^{-k} = \prod_{j \neq k} S^j$.

Some assumptions on the topology of the S^k are required. More precisely, assume that the S^k are topological spaces, equipped with their Borel σ -algebra (that is, the σ -algebra generated by the open sets). Then S is naturally equipped with a (product) topology and a (product) σ -algebra. We can now consider (Borel) probability measures over S .

Hart and Schmeidler's original definition² [20] states that a correlated equilibrium π of the game Γ is a (joint) probability distribution over S such that

² Note that in what follows we only consider games with finitely many players, as opposed to Hart and Schmeidler [20]. We do so for the sake of simplicity. Note however, that the results of Section 3 may be extended to the case of games with countably many players, for which the set of strategies of each player is a compact metric space. This is so because the set of action profiles S is then a compact *metric* space as well, and therefore, Theorem 11, which only relies on Prohorov's theorem (see Proposition 9) via the sequential compactness of the set $\mathcal{P}(S)$ of probability measures on S , still holds. On the other hand, if there are more than countably many players, the set of action profiles S , though compact by Tychonov's theorem, is not necessarily a metric space anymore, and $\mathcal{P}(S)$ may not be (sequentially) compact, as required (see Remark 12).

the extended game Γ' defined below admits the Dirac probability measure concentrated on the N -tuple of identity functions $S^k \rightarrow S^k$ as a Nash equilibrium. The extended game is defined as follows: the strategy set of player k in the game Γ' is given by the set \mathcal{F}_k of all measurable maps $S^k \rightarrow S^k$ (the so-called “departure functions”), and the game is played as follows. Each player k chooses his action $\psi_k \in \mathcal{F}_k$, a signal (sometimes called recommendation) $\mathbf{I} = (I^1, \dots, I^N) \in S$ is drawn randomly according to π , player k is told the k -th component of the signal, I^k , and he finally plays $\psi_k(I^k)$.

The set \mathcal{F}_k of allowed departures for player k may actually be taken as a proper subset of the set $\mathcal{L}^0(S^k)$ of all measurable departures $S^k \rightarrow S^k$, with the only restriction that it should contain the identity map. We then define a $(\mathcal{F}_k)_{1 \leq k \leq N}$ -correlated equilibrium similarly as above except that we consider departure functions ψ_k only from the class \mathcal{F}_k , $k = 1, \dots, N$. In the simplest cases \mathcal{F}_k may be a finite set, but we also consider larger classes \mathcal{F}_k given by the set of all linear functions, all continuous functions, or all measurable functions.

A more formal definition is the following.

Definition 1 *Assume that the payoff functions h^k are measurable and either bounded or nonnegative. An $(\mathcal{F}_k)_{1 \leq k \leq N}$ -correlated equilibrium is a (joint) distribution π over S such that for all players k and all departure functions $\psi_k \in \mathcal{F}_k$, one has*

$$\mathbb{E}_\pi [h^k(I^{-k}, I^k)] \geq \mathbb{E}_\pi [h^k(I^{-k}, \psi_k(I^k))] , \quad (1)$$

where the notation \mathbb{E}_π indicates that the random vector $\mathbf{I} = (I^k)_{1 \leq k \leq N}$, taking values in S , is distributed according to π .

π is an $(\mathcal{F}_k)_{1 \leq k \leq N}$ -correlated ε -equilibrium if for all k and all $\psi_k \in \mathcal{F}_k$,

$$\mathbb{E}_\pi [h^k(I^{-k}, I^k)] \geq \mathbb{E}_\pi [h^k(I^{-k}, \psi_k(I^k))] - \varepsilon .$$

A correlated equilibrium may be interpreted as follows. In an average sense (with respect to the randomization associated with the signal), no player has an incentive to divert from the recommendation, provided that all other players follow theirs. The distinguishing feature of this notion is that, unlike in the definition of Nash equilibria, the random variables I^k do not need to be independent. Indeed, if π is a product measure, it becomes a Nash equilibrium. This also means that correlated equilibria always exist as soon as Nash equilibria do, which is ensured under minimal assumptions (see Remark 2). Their existence may also be seen without underlying fixed point results, see Hart and Schmeidler [20].

Remark 2 In the definition of correlated equilibria we consider an extension

of the original game. But note that under minimal assumptions (e.g., if the sets of strategies S^k are convex compact subsets of topological vector spaces and the payoffs h^k are continuous and concave in the k -th variable) there exists a Nash equilibrium in pure strategies (see, e.g., [1]). Each pure Nash equilibrium corresponds to a $(\mathcal{L}^0(S^k))_{1 \leq k \leq N}$ -correlated equilibrium π given by a Dirac measure over S . Clearly, π is a mixed Nash equilibrium if and only if it is an $(\mathcal{L}^0(S^k))_{1 \leq k \leq N}$ -correlated equilibrium equal to the product of its marginals.

Example 3 Assume that each S^k is a convex and compact subset of a normed vector space and that each payoff function h^k is continuous. In Section 4.5 we show that the set of $(\mathcal{L}^0(S^k))_{1 \leq k \leq N}$ -correlated equilibria coincides with the set of $(\mathcal{C}(S^k))_{1 \leq k \leq N}$ -correlated equilibria, where $\mathcal{C}(S^k)$ is the set of all continuous functions mapping S^k in S^k . This set is convex, compact, and contains the non-empty set of (pure and mixed) Nash equilibria.

For the sake of completeness, we give an analog of the conditional definition, usually proposed as a definition for correlated equilibria in the case of finite games (see, e.g., Aumann [5], Hart and Mas-Colell [16]). Provided that the S^k are finite sets, a *correlated equilibrium* is a (joint) distribution π over S such that for all players k and all functions $\psi_k : S^k \rightarrow S^k$, one has

$$\sum_{s \in S} \pi(s^{-k} | s^k) \left(h^k(s^{-k}, s^k) - h^k(s^{-k}, \psi_k(s^k)) \right) \geq 0 ,$$

where $\pi(\cdot | s^k)$ is the conditional distribution of S^{-k} given that player k is advised to play s^k . Recalling that we denote by $\mathcal{L}^0(S^k)$ the set of all measurable functions over S^k , we have the following conditional definition in the general case where the game may be finite or infinite. The proof is immediate.

Proposition 4 *Under the same measurability and boundedness assumptions as in Definition 1, a distribution π over S is an $(\mathcal{L}^0(S^k))_{1 \leq k \leq N}$ -correlated equilibrium if and only if for all players k and all measurable departure maps $\psi_k : S^k \rightarrow S^k$,*

$$\mathbb{E}_\pi \left[h^k(I^{-k}, I^k) | I^k \right] \geq \mathbb{E}_\pi \left[h^k(I^{-k}, \psi_k(I^k)) | I^k \right] ,$$

where \mathbb{E}_π indicates that the random vector $\mathbf{I} = (I^k)_{1 \leq k \leq N}$ is distributed according to π .

2.2 Discretized games

An alternative natural definition of correlated equilibrium in games with infinite strategy spaces is obtained by discretization. The idea is to “discretize”

the sets of strategies and consider the set of correlated equilibria of the obtained finite game. Appropriate “limits” may be taken as the discretization gets finer. In this section we make this definition precise and show that the obtained definition coincides, under general conditions, with the definition given above when one allows all measurable departure functions.

A $(\mathcal{P}, \mathcal{D})$ *discretization* of the game $\Gamma = (\{1, \dots, N\}, (S^i)_{1 \leq i \leq N}, (h^i)_{1 \leq i \leq N})$ is given by a product partition \mathcal{P} , a grid \mathcal{D} and induced payoffs h_d^k , $1 \leq k \leq N$. More precisely, a product partition is an N -tuple $(\mathcal{P}^1, \dots, \mathcal{P}^N)$, where each \mathcal{P}^k is a finite measurable partition of the corresponding strategy set S^k , which we denote by $\mathcal{P}^k = \{V_1^k, \dots, V_{N_k}^k\}$. In every set V_i^k , $1 \leq k \leq N$, $1 \leq i \leq N_k$, we pick an arbitrary element $t_i^k \in V_i^k$. These points form a *grid* $\mathcal{D}^k = \{t_1^k, \dots, t_{N_k}^k\}$. We write $\mathcal{D} = \mathcal{D}^1 \times \dots \times \mathcal{D}^N$. The induced payoffs h_d^k are obtained by restricting the original payoff functions to the grid \mathcal{D} .

For a given discretization $(\mathcal{P}, \mathcal{D})$, a distribution π over S induces a discrete distribution π_d over the grid \mathcal{D} by

$$\pi_d(t_{i_1}^1 \times \dots \times t_{i_N}^N) = \pi(V_{i_1}^1 \times \dots \times V_{i_N}^N) .$$

The size r of a discretization $(\mathcal{P}, \mathcal{D})$ is the maximal diameter of the sets V_i^k , $1 \leq k \leq N$, $1 \leq i \leq N_k$. If each S^k is compact, then every discretization has a finite size. Then we have the following characterization of correlated equilibria with respect to all measurable departures. (The fairly straightforward proof is given in the Appendix.)

Theorem 5 *Assume that all strategy spaces S^k are convex and compact subsets of a normed space and that the h^k are continuous functions over S . Then a probability distribution over S is an $(\mathcal{L}^0(S^k))_{1 \leq k \leq N}$ -correlated equilibrium of the game Γ if and only if there exists a function ε with $\lim_{r \rightarrow 0} \varepsilon(r) = 0$ such that for all discretizations $(\mathcal{P}, \mathcal{D})$ of size r , π induces an $\varepsilon(r)$ -correlated equilibrium.*

Note that, in the case of a finite number of players, the above result is more precise than the general results contained in the proofs of Theorems 2 and 3 of Hart and Schmeidler [20], where correlated equilibria of a given game with infinite strategy sets (and with an infinite number of players) were shown to be cluster points of the set of correlated equilibria of the discretized games.

3 Regret minimization and convergence in repeated games

One of the remarkable properties of correlated equilibrium in finite games is that if the game is played repeatedly many times such that every player

plays according to a certain regret-minimization strategy then the empirical frequencies of play converge to the set of correlated equilibria. No coordination is necessary between the players, the players do not even need to know the others' payoff functions. This property was first proved by Foster and Vohra [12], see also Fudenberg and Levine [14], Hart and Mas-Colell [16–18], Lehrer [24,25].

The purpose of this section is to investigate to what extent the above-mentioned convergence result can be extended to games with possibly infinite strategy spaces. We consider a situation in which the game Γ is played repeatedly at time instances $t = 1, 2, \dots$. The players are assumed to know their own payoff function and the sequence of strategies played by all players up to time $t - 1$.

3.1 Internal regret

The notion of correlated equilibrium is intimately tied to that of internal (or conditional) regret. Intuitively, internal regret is concerned with the increase of a player's payoff gained by simple modifications of the played strategy. If a simple modification results in a substantial improvement then a large internal regret is suffered.

The formal definition of internal regret (see, e.g., [13]) may be extended to general games in a straightforward manner as follows. Let \mathcal{F}_k be a class of functions $\psi_k : S^k \rightarrow S^k$. As the game Γ is repeated, at each round t , player k could play consistently $\psi_k(s_t^k)$ whenever his strategy³ prescribes to play $s_t^k \in S^k$. This results in a different strategy, called the ψ_k -modified strategy. The maximal cumulative difference in the obtained payoffs for player k , for n rounds of play, equals

$$\mathcal{R}_{\psi_k, n}^k = \max_{s_1^{-k}, \dots, s_n^{-k}} \left(\sum_{t=1}^n h^k(s_t^{-k}, \psi_k(s_t^k)) - \sum_{t=1}^n h^k(s_t^{-k}, s_t^k) \right),$$

where the maximum is taken over all possible sequences of opponent players' actions. We call $\mathcal{R}_{\psi_k, n}^k$ the *internal regret of player k with respect to the departure ψ_k at round n* . The intuition is that if $\mathcal{R}_{\psi_k, n}^k$ is not too large, then the original strategy cannot be improved significantly in a simple way.

We say that a strategy for player k *suffers no internal regret (or minimizes his internal regret) with respect to a class \mathcal{F}_k of departures* whenever

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \mathcal{R}_{\psi_k, n}^k \leq 0,$$

³ Even though it is suppressed in the notation, it is important to keep it in mind that the choice of s_t^k depends on the past action profiles s_1, \dots, s_{t-1} .

for all $\psi_k \in \mathcal{F}_k$. The departure functions play a similar role as in the general definition of correlated equilibria.

3.2 Main convergence result

The main result of the paper, summarized in the following theorem, shows that under general conditions, if all players follow a certain regret-minimizing strategy, the empirical frequencies of play converge to the set of correlated equilibria. Thus, on the average, a correlated equilibrium is achieved without requiring any cooperation among the players.

Theorem 6 (Main result) *Assume that all the strategy spaces S^k are convex and compact subsets of a normed space and all payoff functions h^k are continuous over S and concave in the k -th strategy. Then there exists an internal-regret minimizing strategy such that, if every player follows such a strategy, then joint convergence of the sequence of empirical plays to the set of $(\mathcal{L}^0(S^k))_{1 \leq k \leq N}$ -correlated equilibria is achieved.*

Thus, the convergence result extends, under quite general assumptions, even if all possible measurable departure functions are allowed in the definition of correlated equilibrium. The only restrictive assumption is the concavity of the payoffs. This condition may be removed by allowing the players to use randomized strategies. The next theorem asserts that almost sure convergence of the empirical frequencies of play to the set of correlated equilibria is achieved under the only assumption that the payoff functions are continuous.

Theorem 7 (Main result, randomized version) *Assume that all strategy spaces S^k are convex and compact subsets of a normed space and all payoff functions h^k are continuous over S . If the players are allowed to randomize, then there exists an internal-regret minimizing strategy such that, if every player follows such a strategy, then joint convergence of the sequence of empirical plays to the set of $(\mathcal{L}^0(S^k))_{1 \leq k \leq N}$ -correlated equilibria is achieved almost surely.*

Theorems 6 and 7 are proved below by a series of results, some of which may be of independent interest. In particular, we give upper bounds for the internal regret in some cases, see Theorems 15 and 17, as well as Section 5.

Remark 8 The regret minimizing strategies considered in Theorem 6 are deterministic in the sense that players do not need to randomize. This is made possible because of the concavity assumption on the payoffs. An example is the mixed extension of a finite game, which may be seen to satisfy the assumptions of Theorem 6. This means that if the game is played in the mixed extension (i.e., in each round the players output a probability distribution over the set of

actions), then joint convergence to the set of correlated equilibria (with respect to all measurable departures or just linear departures) may be achieved in the mixed extension, in a deterministic way. It is easy to see that any of these sets of correlated equilibria of the mixed extension induces, in a natural way, the set of correlated equilibria of the underlying finite game. Thus, our algorithm generalizes the (randomized) algorithms designed for the case of finite games. See Section 6 for more details.

3.3 Some topological properties

Theorems 6 and 7 involve convergence properties in the set of probability measures over S . We assume that each S^k is a compact metric space. Then the product S is also a compact metric space, for instance, under the supremum metric. Denote by $\mathcal{C}(S, \mathbb{R})$ the set of (bounded) continuous real-valued functions over S . The set of Borel probability measures over S , denoted by $\mathcal{P}(S)$, is equipped with the weak-* topology. This is the weakest topology such that, for each $f \in \mathcal{C}(S, \mathbb{R})$, the linear map $\mu \rightarrow \mu[f]$ defined for $\mu \in \mathcal{P}(S)$ is continuous, where $\mu[f] = \int_S f d\mu$. That is, the open sets of this topology are generated by the sets

$$\{\mu \in \mathcal{P}(S) : \mu[f] < \alpha\} ,$$

where f is any element of $\mathcal{C}(S, \mathbb{R})$ and α is any real number.

For the analysis we need to establish a topological property of $\mathcal{P}(S)$, namely its (sequential) compactness. It is ensured by the following simple statement of Prohorov's theorem, see, for example, [11].

Proposition 9 (Prohorov's theorem) *If S is a compact metric space, then the space $\mathcal{P}(S)$ is compact. Its topology is equivalent to the topology of the so-called Prohorov metric. In particular, $\mathcal{P}(S)$ is sequentially compact, that is, every sequence of elements from $\mathcal{P}(S)$ contains a convergent subsequence.*

Topological properties of the set of correlated equilibria

This paragraph is given for the sake of completeness and the results presented below are not needed in the rest of the paper.

Fix the set of allowed departures \mathcal{F}_k , $1 \leq k \leq N$, and denote by Π the set of all $(\mathcal{F}_k)_{1 \leq k \leq N}$ -correlated equilibria. It is immediate from the definition that Π is a convex set, and that it contains the set of Nash equilibria (which is known to be non-empty under minimal assumptions, see Remark 2 above).

The next result summarizes some of the basic properties of the set Π . Recall that by Example 3, under some mild conditions, the set of correlated equilibria

with respect to all measurable departures equals the set of correlated equilibria with respect to all continuous departures. Thus, the assumption in the following theorem that departure functions are continuous may be weakened in some important cases.

Theorem 10 *Assume that the strategy spaces S^k are compact metric spaces. The set Π of $(\mathcal{F}_k)_{1 \leq k \leq N}$ -correlated equilibria is non-empty whenever the payoff functions h^k are continuous over S . Moreover, Π is a convex set, which contains the convex hull of Nash equilibria. If, in addition, for all k , $\mathcal{F}_k \subset \mathcal{C}(S^k)$, where $\mathcal{C}(S^k)$ is the set of all continuous functions mapping S^k into S^k , then Π is compact.*

Proof The non-emptiness of Π under the assumption of continuity of the payoff functions follows either from Theorem 3 of Hart and Schmeidler [20] or, alternatively, from the existence of a mixed Nash equilibrium. (The latter may be shown by checking the hypotheses of a version of Nash's theorem stated in Remark 2, which follows easily by Prohorov's and Stone-Weierstrass theorems.)

It remains to prove the compactness of Π under the given assumptions. By Prohorov's theorem (see Proposition 9), Π is included in a compact set, therefore it is enough to prove that Π is a closed set. To this end, consider the continuous real-valued function over S defined by

$$f_{k,\psi_k}(s) = h^k(s^{-k}, s^k) - h^k(s^{-k}, \psi_k(s^k)) ,$$

where $1 \leq k \leq N$ and $\psi_k \in \mathcal{F}_k \subset \mathcal{C}(S^k)$. Each f_{k,ψ_k} is a continuous real-valued function over S and Π is the intersection of the closed half-spaces

$$\{\mu \in \mathcal{P}(S) : \mu[f_{k,\psi_k}] \geq 0\} . \quad \square$$

3.4 Joint minimization of internal regret

We are now ready to prove that, similarly to the case of finite games, if all players minimize their internal regrets, then joint convergence of the sequence of empirical distributions of plays to the set of correlated equilibria is guaranteed.

Denote by s_1, \dots, s_n the played strategies up to time n . We denote the empirical distribution of plays up to time n by

$$\pi_n = \frac{1}{n} \sum_{t=1}^n \delta_{s_t} ,$$

where δ_s is the Dirac mass on $s \in S$. We have the following convergence result generalizing the corresponding statement of Foster and Vohra [12] for finite games.

Theorem 11 *If each player k minimizes his internal regret with respect to a departure class \mathcal{F}_k , then, provided that the S^k are compact metric spaces, the h^k are continuous, and $\mathcal{F}_k \subset \mathcal{C}(S^k)$ for all k , the sequence of empirical distributions of plays $(\pi_n)_{n \in \mathbb{N}}$ converges to the set of $(\mathcal{F}_k)_{1 \leq k \leq N}$ -correlated equilibria.*

Recall that under the assumptions of the theorem, S is a compact metric space, and so is $\mathcal{P}(S)$ by Prohorov's theorem (see Proposition 9). We denote by d a metric of $\mathcal{P}(S)$. The convergence of the sequence of empirical distributions of plays to a subset C of the set $\mathcal{P}(S)$ of all probability distributions over S means that

$$\lim_{n \rightarrow \infty} \inf_{\mu \in C} d(\pi_n, \mu) = 0 .$$

Proof The assumption on the internal regrets may be rewritten as

$$\limsup_{n \rightarrow \infty} \mathbb{E}_{\pi_n} [h^k(I^{-k}, \psi_k(I^k))] - \mathbb{E}_{\pi_n} [h^k(I^{-k}, I^k)] \leq 0 \quad (2)$$

for all k and all $\psi_k \in \mathcal{F}_k$, where $\mathbf{I} = (I^k)_{1 \leq k \leq N}$ is the identity map over S , defined on the probability space (S, π_n) . (\mathbb{E}_{π_n} denotes expectation with respect to this probability measure π_n over S .) By Prohorov's theorem (see Proposition 9), the sequence $(\pi_n)_{n \in \mathbb{N}}$ lies in the compact metric space $\mathcal{P}(S)$. Thus, if the whole sequence did not converge to the set of $(\mathcal{F}_k)_{1 \leq k \leq N}$ -correlated equilibria, we could extract from it a subsequence $(\pi_{\varphi(n)})_{n \in \mathbb{N}}$, where φ is an increasing function $\mathbb{N} \rightarrow \mathbb{N}$, such that $(\pi_{\varphi(n)})_{n \in \mathbb{N}}$ converges to a probability measure π which is not a $(\mathcal{F}_k)_{1 \leq k \leq N}$ -correlated equilibrium. That is, there exists a player k , $1 \leq k \leq N$, and a departure $\psi_k \in \mathcal{F}_k$ such that

$$\mathbb{E}_{\pi} [h^k(I^{-k}, I^k)] < \mathbb{E}_{\pi} [h^k(I^{-k}, \psi_k(I^k))] . \quad (3)$$

But (2) ensures that

$$\limsup_{n \rightarrow \infty} \mathbb{E}_{\pi_{\varphi(n)}} [h^k(I^{-k}, \psi_k(I^k))] - \mathbb{E}_{\pi_{\varphi(n)}} [h^k(I^{-k}, I^k)] \leq 0 .$$

By continuity of the function f_{k, ψ_k} defined by

$$f_{k, \psi_k}(s) = h^k(s^{-k}, s^k) - h^k(s^{-k}, \psi_k(s^k)) , \quad s \in S$$

and by the definition of weak-* topology over $\mathcal{P}(S)$, we have

$$\begin{aligned} & \lim_{n \rightarrow \infty} \mathbb{E}_{\pi_\varphi(n)} \left[h^k(I^{-k}, \psi_k(I^k)) \right] - \mathbb{E}_{\pi_\varphi(n)} \left[h^k(I^{-k}, I^k) \right] \\ &= \mathbb{E}_\pi \left[h^k(I^{-k}, \psi_k(I^k)) \right] - \mathbb{E}_\pi \left[h^k(I^{-k}, I^k) \right] \leq 0, \end{aligned}$$

which contradicts (3), thus proving the desired convergence. \square

Remark 12 The assumption of compactness in Theorem 11 is crucial. Consider, for instance, the one-player game with strategy set \mathbb{N} and payoff function h defined by $h(0) = 0$ and $h(n) = -1/n$ for all $n \geq 1$. The strategy that consists in playing the index n at the n -th round of the repeated game has an average payoff converging to zero, and hence, asymptotically minimizes its internal regret. However, the empirical frequencies of play do not converge to the unique correlated equilibrium of the game, which is the probability distribution that puts probability mass 1 on 0.

Theorem 11 shows that in order to guarantee convergence of the empirical frequencies of play to the set of correlated equilibria, it suffices that all players use a strategy that minimizes their internal regret.

4 Internal-regret minimizing algorithms

The main issues in designing internal-regret minimizing strategies concern the size of the set of allowed departures \mathcal{F}_k . For finite games, the measurable departures $S^k \rightarrow S^k$ are given by all functions $S^k \rightarrow S^k$, whose number is $m_k^{m_k}$, where m_k is the cardinality of S^k . If S^k is infinite (countably or continuously infinite), there is *a priori* an infinite number of departures. In particular, a simple procedure allocating a weight per each departure function, as was proposed in the finite case in Foster and Vohra [13] and Hart and Mas-Colell [17], would be impossible if the set of allowed departures was too large. Thus, learning algorithms designed for finite games cannot be generalized as easily as the definition could be carried over to the infinite case. Constructing new learning algorithms for some general classes of infinite games is the point of the present section.

4.1 Blackwell's condition

Regret-minimization strategies have been often derived from Blackwell's approachability theorem [7]. Here however, we do not need the full power of Blackwell's theory, only a few simple inequalities derived in Cesa-Bianchi and Lugosi [10] which we briefly recall. Consider a sequential decision problem parameterized by a *decision space* \mathcal{X} , by an *outcome space* \mathcal{Y} . At each step

$t = 1, 2, \dots$, the decision maker selects an element \hat{x}_t from the decision space \mathcal{X} . In return, an outcome $y_t \in \mathcal{Y}$ is received, and the decision maker suffers a vector $\mathbf{r}_t = \mathbf{r}_t(\hat{x}_t, y_t) \in \mathbb{R}^N$ of regret. The cumulative regret after t rounds of play is $\mathbf{R}_t = \sum_{s=1}^t \mathbf{r}_s$. The goal of the decision maker is to minimize $\max_{i=1, \dots, N} R_{i,n}$, that is, the largest component of the cumulative regret vector after n rounds of play.

Similarly to Hart and Mas-Colell [17], we consider potential-based decision-making strategies, based on a convex and twice differentiable *potential function* $\Phi : \mathbb{R}^N \rightarrow \mathbb{R}^+$. Even though the results below hold for a general class of potential functions, for concreteness and to get the best bounds, we restrict our attention to the special case of the exponential potential given by

$$\Phi(\mathbf{u}) = \sum_{i=1}^N \exp(\eta u_i) ,$$

where the parameter $\eta > 0$ will be tuned by the analysis below.

We recall the following bound, proved in [10].

Proposition 13 *Assume that the decision-maker plays such that in each round t of play, the regret vector \mathbf{r}_t satisfies the so-called “Blackwell condition”*

$$\nabla \Phi(\mathbf{R}_{t-1}) \cdot \mathbf{r}_t \leq 0 . \tag{4}$$

If $\|\mathbf{r}_t\|_\infty \leq M$ for all t then in the case of an exponential potential and for the choice $\eta = 1/M\sqrt{2 \ln N/n}$,

$$\max_{1 \leq i \leq N} R_{i,n} \leq M\sqrt{2n \ln N} .$$

Observe that the value of the parameter η requires the knowledge of the number of rounds n . We remark here that similar bounds hold if, instead of the exponential potential function, the polynomial potential

$$\Phi(\mathbf{u}) = \sum_{i=1}^N (u_i)_+^p$$

is used with $p = 2 \ln N$, see [10]. (Here previous knowledge of the horizon n is unnecessary.)

4.2 Finite classes of departure functions

As a first step assume that the set \mathcal{F}_k of allowed departures for player k is finite, with cardinality m_k . For any $s \in S$ and departure $\psi_k \in \mathcal{F}_k$, denote by

$$r_{\psi_k}^k(s) = h^k(s^{-k}, \psi_k(s^k)) - h^k(s^{-k}, s^k)$$

the associated instantaneous internal regret, and by

$$\mathbf{r}^k(s) = \left(r_{\psi_k}^k(s) \right)_{\psi_k \in \mathcal{F}_k}$$

the regret vector formed by considering all departures. For a given sequence $s_1, \dots, s_n \in S$ of plays, the cumulative internal regrets are given by the vector

$$\mathbf{R}^k(s_1^n) = \sum_{t=1}^n \mathbf{r}^k(s_t),$$

where s_1^n denotes the sequence (s_1, \dots, s_n) . $\mathbf{R}^k(s_1^n)$ will be referred to as $(R_{\psi_k}^k(s_1^n))_{\psi_k \in \mathcal{F}_k}$.

Consider the following algorithm for player k . For $t = 1, 2, \dots$, at round t , player k chooses any $s_t^k \in S^k$ such that

$$s_t^k = \sum_{\psi_k \in \mathcal{F}_k} \Delta_{\psi_k, t-1}^k \psi_k(s_t^k), \quad (5)$$

with

$$\Delta_{\psi_k, t-1}^k = \frac{\varphi \left(R_{\psi_k}^k(s_1^{t-1}) \right)}{\sum_{g \in \mathcal{F}_k} \varphi \left(R_g^k(s_1^{t-1}) \right)}, \quad t \geq 2$$

where $\varphi(x) = \exp(\eta x)$. For $t = 1$ we set $\Delta_{\psi_k, 0}^k = 1/m_k$. (The parameter η will be tuned by the analysis below.)

Thus, each player is assumed to choose his action by solving⁴ the fixed-point equation (5). The existence of such a fixed point (under the assumptions of Theorem 15) follows easily by the Schauder-Cauty fixed-point theorem [8], which we recall below. The Schauder-Cauty theorem is applicable since the right-hand side of (5) is a fixed convex combination of continuous mappings $S^k \rightarrow S^k$.

⁴ Note that an approximate solution of (5) is sufficient for our purposes. Provided that S^k is included in a normed vector space and h^k is a Lipschitz function, a simple modification of the proof of Cesa-Bianchi and Lugosi [10, Theorem 1] shows that the internal regret would still be $o(n)$ had we used a strategy s_t^k such that $\|s_t^k - \sum_{\psi_k \in \mathcal{F}_k} \Delta_{\psi_k, t-1}^k \psi_k(s_t^k)\| \leq \varepsilon_n$, where ε_n decreases quickly enough to 0. In particular, when the S^k are included in finite-dimensional vector spaces, an algorithm partitioning S^k into a thin grid is able to find a suitable approximate fixed-point.

Theorem 14 (Schauder-Cauty fixed-point theorem) *Let C be a non-empty convex and compact subset of a topological Hausdorff vector space. Then each continuous map $T : C \rightarrow C$ has a fixed point.*

Note that if several fixed points of (5) exist, then the player is free to choose any of them.

Theorem 15 *Assume that S^k is a convex and compact subset of a topological Hausdorff vector space and that the payoff function h^k is bounded over S by $M_k \in \mathbb{R}$ and is concave in the k -th strategy. Then, whenever \mathcal{F}_k is a finite subset of $\mathcal{C}(S^k)$ with cardinality m_k , the above algorithm guarantees that the cumulative internal regret satisfies*

$$\max_{\psi_k \in \mathcal{F}_k} \mathcal{R}_{\psi_k, n}^k \leq M_k \sqrt{2n \ln m_k} ,$$

if the exponential potential is used with $\eta = 1/M_k \sqrt{2 \ln m_k/n}$.

Remark 16 (Rates of convergence) The theorem implies that, for a given horizon n , if all \mathcal{F}_k are finite, and all players play according to the above procedure, then, at round n , the empirical distribution is an $(\mathcal{F}_k)_{1 \leq k \leq N}$ -correlated ε_n -equilibrium, with ε_n of the order $1/\sqrt{n}$.

Proof The statement follows easily by Theorem 13. It suffices to prove that our choice of s_t^k satisfies the Blackwell condition

$$\nabla \Phi(\mathbf{R}^k(s_1^{t-1})) \cdot \mathbf{r}^k(s_t) \leq 0$$

or equivalently

$$\sum_{\psi_k \in \mathcal{F}_k} \Delta_{\psi_k, t-1}^k h^k(s_t^{-k}, \psi_k(s_t^k)) \leq h^k(s_t^{-k}, s_t^k) ,$$

which is implied by the equality

$$h^k \left(s_t^{-k}, \sum_{\psi_k \in \mathcal{F}_k} \Delta_{\psi_k, t-1}^k \psi_k(s_t^k) \right) = h^k(s_t^{-k}, s_t^k)$$

and by the concavity of h^k in its k -th argument. This equality ensured by the choice (5). \square

4.3 Countably infinite classes of departure functions

The next step is to extend the result of the previous section to countably infinite classes of departure functions. In this section we design an internal-

regret minimizing procedure in the case when the set of allowed departures for player k is countably infinite. Denote by

$$\mathcal{F}_k = \{\psi_{k,q}, q \in \mathbb{N}\}$$

the set of departure functions of player k .

Theorem 17 *Assume that S^k is a convex and compact subset of a topological Hausdorff vector space and that the payoff function h^k is bounded by M_k and is concave in the k -th strategy. If \mathcal{F}_k is a countable subset of $\mathcal{C}(S^k)$, there exists a procedure such that for all $q \in \mathbb{N}$ and n ,*

$$\mathcal{R}_{\psi_{k,q},n}^k \leq M_k \left(2(\ln q)^2 + 4.2n^{3/4} \right) .$$

Consequently, this procedure suffers no internal regret.

Proof We use a standard tool, often called the “doubling trick” (see, e.g., [9]) to extend the procedure of Theorem 15. Time is divided into blocks of increasing lengths such that the t -th block is $\llbracket 2^{t-1}, 2^t - 1 \rrbracket$. At the beginning of the t -th block, the algorithm for player k takes a fresh start and uses the method presented in Section 4.2, with the departures indexed by the integers between 1 and m_t and with $\eta = \eta_t$ tuned as

$$\eta_t = \frac{1}{M_k} \sqrt{2 \frac{\ln m_t}{2^{t-1}}} .$$

We take, for instance, $m_t = \lfloor \exp \sqrt{2^t} \rfloor$, where $\lfloor x \rfloor$ denotes the greatest integer smaller than x .

Denote $\bar{n} = 2^{\lfloor \log_2 n \rfloor + 1}$. Define

$$H^k(s_j^n) = \sum_{t=j}^n h^k(s_t^{-k}, s_t^k) , \quad \text{and} \quad H_{\psi_{k,q}}^k(s_j^n) = \sum_{t=j}^n h^k(s_t^{-k}, \psi_{k,q}(s_t^k)) .$$

Theorem 15 ensures that

$$\begin{aligned} H^k(s_1^n) &= \sum_{t=1}^{\lfloor \log_2 n \rfloor} H^k(s_{2^{t-1}}^{2^t-1}) + H^k(s_{\bar{n}/2}^{\bar{n}}) \\ &\geq \sum_{t=1}^{\lfloor \log_2 n \rfloor} \left(\max_{1 \leq q \leq m_t} H_{\psi_{k,q}}^k(s_{2^{t-1}}^{2^t-1}) - M_k \sqrt{2^t \ln m_t} \right) \\ &\quad + \left(\max_{1 \leq q \leq m_{\lfloor \log_2 n \rfloor + 1}} H_{\psi_{k,q}}^k(s_{\bar{n}/2}^{\bar{n}}) - M_k \sqrt{\bar{n} \ln m_{\lfloor \log_2 n \rfloor + 1}} \right) . \end{aligned}$$

The departure function $\psi_{k,q}$ is considered from the time segment indexed by t_q , where t_q is the smallest integer such that $q \leq m_{t_q}$, that is, $2^{t_q-1} < (\ln q)^2 \leq 2^{t_q}$.

Observe that the total length of the previous time segments is $2^{t_q} - 1 \leq 2(\ln q)^2$. Thus, we obtain, for any $q \in \mathbb{N}$,

$$\begin{aligned} H^k(s_1^n) &\geq H_{\psi_{k,q}}^k(s_1^n) - M_k \left(2(\ln q)^2 + \sum_{t=1}^{\lfloor \log_2 n \rfloor + 1} \sqrt{2^t \ln(\exp \sqrt{2^t})} \right) \\ &\geq H_{\psi_{k,q}}^k(s_1^n) - M_k \left(2(\ln q)^2 + \sum_{t=1}^{\lfloor \log_2 n \rfloor + 1} (2^{3/4})^t \right) \\ &\geq H_{\psi_{k,q}}^k(s_1^n) - M_k \left(2(\ln q)^2 + \frac{2^{3/2}}{2^{3/4} - 1} n^{3/4} \right), \end{aligned}$$

which concludes the proof. \square

Remark 18 Theorem 17 does not provide any uniform bound for the internal regrets (where uniformity is understood with respect to the elements of the class of allowed departures \mathcal{F}_k), contrary to the case of finitely many departure functions of Theorem 15 (see Remark 16). In fact, in general, no non-trivial rate can be given for the convergence of the empirical distribution of plays to the set of $(\mathcal{F}_k)_{1 \leq k \leq n}$ -correlated equilibria. However, in the special case of totally bounded classes of departures, rates of convergence may be established, and the rates depend on the size of the classes, see Section 5. This means that the choice of the departure classes may be an important issue in practice.

4.4 Separable sets of departure functions

The extension to separable sets of departure functions is now quite straightforward. Recall that compact or totally bounded spaces are special cases of separable spaces so the next result covers quite general situations.

Theorem 19 *Assume that all strategy spaces S^k are convex and compact subsets of normed vector spaces. Let the payoff functions h^k be continuous over S and concave in the k -th strategy and assume that the \mathcal{F}_k are separable subsets of $\mathcal{C}(S^k)$ (equipped with the supremum norm). Then there exist regret minimizing strategies such that, if every player follows such a strategy, then joint convergence of the sequence of empirical plays to the set of $(\mathcal{F}_k)_{1 \leq k \leq N}$ -correlated equilibria is achieved.*

The proof is based on the following lemma that can be shown by a simple dominated-convergence argument.

Lemma 20 *Assume that the h^k are continuous, and let (\mathcal{G}_k) , $1 \leq k \leq N$, be classes of departure functions. Let π be a $(\mathcal{G}_k)_{1 \leq k \leq N}$ -correlated equilibrium. If*

for every k , \mathcal{F}_k denotes the set of functions that may be obtained as π -almost sure limits of elements from \mathcal{G}_k , then π is an $(\mathcal{F}_k)_{1 \leq k \leq N}$ -correlated equilibrium.

Proof of Theorem 19 For each player k , consider a countable dense subset \mathcal{G}_k of \mathcal{F}_k and apply the algorithm given in the proof of Theorem 17. Then Theorems 17 and 11 show that the empirical distribution of plays converges to the set of $(\mathcal{G}_k)_{1 \leq k \leq N}$ -correlated equilibria. By Lemma 20 the set of $(\mathcal{G}_k)_{1 \leq k \leq N}$ -correlated equilibria coincides with the set of $(\mathcal{F}_k)_{1 \leq k \leq N}$ -correlated equilibria. \square

4.5 Proof of Theorem 6

To prove Theorem 6, we need two intermediate results. The first establishes separability needed to apply Theorem 19.

Lemma 21 *If X is a convex and compact subset of a normed vector space, then the set $\mathcal{C}(X)$ of continuous functions $X \rightarrow X$ is separable (in the sense of the supremum norm).*

The proof is an extension of Hirsch and Lacombe [21, Proposition 1.1]. Second we need a characterization of correlated equilibria with respect to all measurable departures. The proofs of both results are given in the Appendix.

Lemma 22 *Assume that the strategy spaces S^k are convex and compact subsets of a normed vector space and that the h^k are continuous functions over S . Then the set of correlated equilibria with respect to all continuous departures $(\mathcal{C}(S^k))_{1 \leq k \leq N}$ equals the set of correlated equilibria with respect to all measurable departures $(\mathcal{L}^0(S^k))_{1 \leq k \leq N}$.*

Proof of the Main theorem By the separability property stated in Lemma 21, Theorem 19 applies and gives an algorithm leading to convergence to the set of $(\mathcal{C}(S^k))_{1 \leq k \leq N}$ -correlated equilibria. In view of Lemma 22, this is equivalent to convergence to the set of $(\mathcal{L}^0(S^k))_{1 \leq k \leq N}$ -correlated equilibria, thus concluding the proof. \square

4.6 Proof of Theorem 7

Sections 4.3, 4.4, and 4.5 only rely on the results of Section 4.2, and therefore it suffices to extend the results of Section 4.2 to the case of non-concave payoffs.

We assume that the strategy sets S^k are convex and compact subsets of normed vector spaces, and the payoff functions h^k are continuous over S^k . The players are allowed to randomize (which they do independently of each other). More, precisely, player k chooses his action s_t^k at round t according to the probability distribution $\mu_t^k \in \mathcal{P}(S^k)$, where $\mathcal{P}(S^k)$ denotes the set of probability distributions over S^k . We also assume that the departure class \mathcal{F}_k is a finite subset of $\mathcal{C}(S^k)$, with cardinality m_k .

For any $\mu^k \in \mathcal{P}(S^k)$, $s^{-k} \in S^{-k}$ and any departure $\psi_k \in \mathcal{F}_k$, we denote

$$h^k(s^{-k}, \mu^k) = \int_{S^k} h^k(s^{-k}, s^k) d\mu^k(s^k) ,$$

and by $(\mu^k)^{\psi_k}$ the image measure of μ^k by ψ_k , which means, in particular, that

$$h^k(s^{-k}, (\mu^k)^{\psi_k}) = \int_{S^k} h^k(s^{-k}, \psi_k(s^k)) d\mu^k(s^k) .$$

Below we design a procedure for player k such that for all possible sequences of opponents' plays, $s_1^{-k}, s_2^{-k}, \dots$,

$$\sum_{t=1}^n \left(h^k(s_t^{-k}, (\mu_t^k)^{\psi_k}) - h^k(s_t^{-k}, \mu_t^k) \right) = o(n) . \quad (6)$$

Then, thanks to the boundedness of the payoff function h^k , we may use a simple martingale convergence result such as the Hoeffding-Azuma inequality [6,22], as well as the Borel-Cantelli lemma, to show that (6) implies

$$\sum_{t=1}^n \left(h^k(s_t^{-k}, \psi_k(s_t^k)) - h^k(s_t^{-k}, s_t^k) \right) = o(n) \quad \text{a.s. .}$$

The latter is enough to apply Theorem 11, and prove the desired almost sure convergence.

It thus only remains to see how to design a procedure for player k guaranteeing (6). The techniques of Section 4.2 extend easily to this case. For any $\mu^k \in \mathcal{P}(S^k)$, $s^{-k} \in S^{-k}$ and any departure $\psi_k \in \mathcal{F}_k$, denote by

$$r_{\psi_k}^k(s^{-k}, \mu^k) = h^k(s^{-k}, (\mu^k)^{\psi_k}) - h^k(s^{-k}, \mu^k)$$

the associated instantaneous internal regret, and by

$$\mathbf{r}^k(s^{-k}, \mu^k) = \left(r_{\psi_k}^k(s^{-k}, \mu^k) \right)_{\psi_k \in \mathcal{F}_k}$$

the regret vector formed by considering all departures. For a given sequence $s_1^{-k}, \dots, s_n^{-k}$ of opponents' plays, and the sequence of probability distributions

μ_1^k, \dots, μ_n^k , the cumulative internal regrets are given by the vector

$$\mathbf{R}^k \left((s^{-k})_1^n, (\mu^k)_1^n \right) = \sum_{t=1}^n \mathbf{r}^k(s_t^{-k}, \mu_t^k),$$

where $(s^{-k})_1^n$ denotes the sequence $(s_1^{-k}, \dots, s_n^{-k})$, and $(\mu^k)_1^n$ is $(\mu_1^k, \dots, \mu_n^k)$. Now, assume that player k can select his distribution μ_t^k at time t as a solution $\mu \in \mathcal{P}(S^k)$ of the equation

$$\mu = \sum_{\psi_k \in \mathcal{F}_k} \Delta_{\psi_k, t-1}^k \mu^{\psi_k}, \quad (7)$$

where

$$\Delta_{\psi_k, t-1}^k = \frac{\varphi \left(R_{\psi_k}^k \left((s^{-k})_1^{t-1}, (\mu^k)_1^{t-1} \right) \right)}{\sum_{g \in \mathcal{F}_k} \varphi \left(R_g^k \left((s^{-k})_1^{t-1}, (\mu^k)_1^{t-1} \right) \right)}, \quad t \geq 2,$$

with $\varphi(x) = \exp(\eta x)$, $\Delta_{\psi_k, 0}^k = 1/m_k$, and the parameter η is tuned as in Section 4.2. If μ_t^k is proved to exist for all t , then we may obtain an upper bound of the order of \sqrt{n} on the right-hand side of (6), by mimicking the argument of the proof of Theorem 15.

But the existence of such a distribution μ_t^k follows by the Schauder-Cauty fixed-point theorem. Recall that the weak-* topology induced on $\mathcal{P}(S^k)$ is such that, for all $\psi \in \mathcal{C}(S^k)$, the map that assigns the element μ^ψ to $\mu \in \mathcal{P}(S^k)$ is continuous. Thus, on the right-hand side of (7), we have a continuous function of μ . The existence of μ_t^k follows by the application of the claimed fixed-point theorem to the convex and compact subset $\mathcal{P}(S^k)$ of the vector space of all Borel, finite, real-valued and regular measures over S^k , equipped with its weak-* topology. \square

5 A note on rates of convergence

Up to this point we have only focused on asymptotic statements and have not paid attention to rates of convergence. In particular, in Sections 4.3 and 4.4, we did not consider the way the elements of the countable classes were ordered, and we set up some parameters quite arbitrarily. However, under some assumptions, precise non-asymptotic bounds may be derived for the internal regret.

Recall that in the case of finite classes of departure functions, the internal regret can be made of the order of $n^{1/2}$. For richer classes of departure functions this may become larger, depending on the richness of the class. In this short remark we point out this phenomenon by considering totally bounded classes of departures.

Here we assume that the strategy set S^k of player k is a convex and compact subset of a normed vector space, that his payoff functions h^k is a Lipschitz function concave in the k -th strategy, and that his class of departures \mathcal{F}_k is a totally bounded set under the corresponding supremum norm.

Recall that a metric space X is said to be *totally bounded* if for all $\varepsilon > 0$, there exists a finite cover of X by balls of radius ε . For a given ε , the minimal number of such balls is called the ε -covering number of X , and is denoted by $N(\varepsilon)$. Any cover of X of size $N(\varepsilon)$ will be referred to as an ε -cover of X .

Denote by $N_k(\varepsilon)$ the ε -covering number of \mathcal{F}_k , let δ_k be a Lipschitz constant of h^k , and let M_k be an upper bound for $|h^k|$. For any $\alpha > 0$, introduce

$$\varepsilon_k(\alpha) = \inf \left\{ \varepsilon : \alpha \delta_k^2 \varepsilon^2 \geq 4M_k^2 \ln N_k(\varepsilon) \right\} . \quad (8)$$

Clearly, $\varepsilon_k(\alpha)$ is decreasing. Moreover, $\varepsilon_k(\alpha)$ tends to 0 as $\alpha \rightarrow \infty$.

To obtain a bound on the cumulative regret with respect to a totally bounded class of departure functions, we use the doubling trick similarly to Section 4.3. Time is divided again in segments such that the r -th segment ($r \geq 1$) corresponds to the time instances t between 2^{r-1} and $2^r - 1$. In the r -th segment, the procedure for player k is the one of Section 4.2, with a departure class given by the centers of the balls which form an $(\varepsilon_k(2^r) + 2^{-r})$ -cover of \mathcal{F}_k . Denoting $\varepsilon'_k(2^r) = \varepsilon_k(2^r) + 2^{-r}$, this implies, using the uniform continuity of h^k , that for all sequences of opponents' plays, $s_1^{-k}, s_2^{-k}, \dots$, and for all departure functions $\psi_k \in \mathcal{F}_k$,

$$H^k(s_1^n) \geq H_{\psi_k}^k(s_1^n) - \sum_{r=1}^{\lfloor \log_2 n \rfloor + 1} \left(M_k \sqrt{2^r \ln N_k(\varepsilon'_k(2^r))} + 2^{r-1} \delta_k \varepsilon'_k(2^r) \right) .$$

Noting that $\varepsilon_k(2^r)$ is defined as the infimum of an interval, so that $\varepsilon = \varepsilon'_k(2^r)$ satisfies $2^r \delta_k^2 \varepsilon^2 \geq 4M_k^2 \ln N_k(\varepsilon)$, leads to the following proposition.

Proposition 23 *Assume that the strategy set S^k of player k is a convex and compact subset of a normed vector space, that his payoff functions h^k is a Lipschitz function concave in the k -th strategy, bounded by M_k and with Lipschitz constant δ_k , and that his class of departures \mathcal{F}_k is a totally bounded set under the corresponding supremum norm. Then the internal regret of the above procedure is bounded as*

$$\max_{\psi_k \in \mathcal{F}_k} \mathcal{R}_{\psi_k, n}^k \leq \delta_k \sum_{r=1}^{\lfloor \log_2 n \rfloor + 1} 2^r \varepsilon'_k(2^r) = \delta_k \left(1 + \log_2 n + \sum_{r=1}^{\lfloor \log_2 n \rfloor + 1} 2^r \varepsilon_k(2^r) \right) , \quad (9)$$

where the $\varepsilon_k(2^r)$ are defined by (8).

Observe that the quantity on the right-hand side of (9) is always $o(n)$ by an

application of Cesaro's lemma and the fact that $\varepsilon_k(2^r) \rightarrow 0$ as $r \rightarrow \infty$.

Example 24 As a concrete example, consider the case when the strategy set of player k is the d -dimensional cube $S^k = [0, 1]^d$ and the class \mathcal{F}_k of departures is the class of all Lipschitz functions $[0, 1]^d \rightarrow [0, 1]^d$ with Lipschitz norm bounded by L_k . It is equipped with the metric associated to the supremum norm. Kolmogorov and Tihomirov [23, Theorem XIV] show that the metric entropy $\log N_k(\varepsilon)$ of this class of functions is of the order of ε^{-d} , that is⁵, $\log N_k(\varepsilon) = \Theta(\varepsilon^{-d})$. It follows that $\varepsilon_k(\alpha) = \Theta(\alpha^{-1/(d+2)})$, and (9) implies that

$$\max_{\psi_k \in \mathcal{F}_k} \mathcal{R}_{\psi_k, n}^k \leq c n^{\frac{d+1}{d+2}}$$

for a constant c (depending only on δ_k , M_k , and L_k).

Many other examples of metric entropies $\log N_k(\varepsilon)$, of totally bounded classes may be found in Kolmogorov and Tihomirov [23, Sections 5–9].

6 A link with correlated equilibrium of finite games

In this final section we assume that Γ is a finite game, with strategy sets given by finite sets S^k . Assume that the players play in the mixed extension, that is, at round t , each player k chooses privately a probability distribution p_t^k over S^k , all probability distributions $p_t = (p_t^1, \dots, p_t^N)$ are made public, and player k gets the payoff $h^k(p_t)$, where we linearly extend the definition of h^k by

$$h^k(p_t) = \sum_{s \in S} \left(\prod_{j=1}^N p_t^j(s^j) \right) h^k(s) .$$

The results of the previous sections show that the players can ensure that the empirical frequencies of play in the mixed extension,

$$\mu_n = \frac{1}{n} \sum_{t=1}^n \delta_{(p_t^1, \dots, p_t^N)} ,$$

converge to some set of correlated equilibria of the mixed extension of Γ , for instance, the set $E_{\mathcal{L}^0}$ of correlated equilibria with respect to all measurable departures, or the set E_L of correlated equilibria with respect to all linear departures. The convergence to $E_{\mathcal{L}^0}$ may be seen by Theorem 6, whereas the convergence to E_L is given by Theorem 15, since the set of all linear mappings $\mathcal{P}(S^k) \rightarrow \mathcal{P}(S^k)$ is the convex hull of a finite number of mappings.

⁵ The notation $x_\varepsilon = \Theta(y_\varepsilon)$ means that the ratio $x_\varepsilon/y_\varepsilon$ is bounded above and below by positive numbers as ε tends to 0.

Recall that this is done by minimizing the internal regrets, that is, by ensuring that for all players k and all $\varphi_k \in \mathcal{F}_k$,

$$\sum_{t=1}^n h^k(p_t^{-k}, \varphi_k(p_t^k)) - \sum_{t=1}^n h^k(p_t^{-k}, p_t^k) = o(n) , \quad (10)$$

where \mathcal{F}_k is either a countable dense subset of the set of all continuous functions $\mathcal{P}(S^k) \rightarrow \mathcal{P}(S^k)$, or the finite set of the mappings generating all linear functions $\mathcal{P}(S^k) \rightarrow \mathcal{P}(S^k)$.

We are actually interested in playing in the original finite game, and to do so⁶, we assume that at each round $t = 1, 2, \dots$, each player k draws an action $s_t^k \in S^k$ according to p_t^k . We denote by $(\hat{\pi}_n)_{n \in \mathbb{N}^*}$ the sequence of joint empirical frequencies of play,

$$\hat{\pi}_n = \frac{1}{n} \sum_{t=1}^n \delta_{(s_t^1, \dots, s_t^N)} ,$$

and study its convergence properties. One may wonder whether it converges almost surely to a set strictly smaller than the set E_Γ of correlated equilibria of the finite game Γ for some game Γ .

Here we point out that the results of this paper do not imply convergence of the empirical frequencies to a set smaller than the set of correlated equilibria of the finite game. More precisely, we show that the set of correlated equilibria of the mixed extension and that of the original finite game are the same in a natural sense.

For any distribution μ over $\mathcal{P}(S^1) \times \dots \times \mathcal{P}(S^N)$, denote by $\psi(\mu)$ the distribution $\psi(\mu) = \pi$ over $S^1 \times \dots \times S^N$ defined, for all $i^k \in S^k$, by

$$\pi(i^1, \dots, i^N) = \int_{\mathcal{P}(S^1) \times \dots \times \mathcal{P}(S^N)} p^1(i^1) \dots p^N(i^N) d\mu(p^1, \dots, p^N) . \quad (11)$$

By this definition and by considering the linear extension of h^k , we have that $\mathbb{E}_{\psi(\mu)} h^k = \mathbb{E}_\mu h^k$ for all k and μ .

Denote $\pi_n = \psi(\mu_n)$ and note that $\|\pi_n - \hat{\pi}_n\| \rightarrow 0$ by martingale convergence. $(\hat{\pi}_n)_{n \in \mathbb{N}^*}$ and $(\pi_n)_{n \in \mathbb{N}^*}$ have therefore the same convergence properties. But since ψ is continuous, the $\pi_n = \psi(\mu_n)$ converge to the set $\psi(E_{\mathcal{M}})$, and therefore, so do the $\hat{\pi}_n$.

By Remark 8, $\psi(E_{\mathcal{L}^0}) \subseteq E_\Gamma$. Below we show that, in fact, $\psi(E_{\mathcal{L}^0}) = \psi(E_L) =$

⁶ Note that by martingale convergence, (10) is ensured almost surely whenever for all players k and all $\varphi_k \in \mathcal{F}_k$, $\sum_{t=1}^n h^k(s_t^{-k}, \varphi_k(p_t^k)) - \sum_{t=1}^n h^k(s_t^{-k}, p_t^k) = o(n)$. This can be done in the finite game by using the fixed-point techniques of Section 4.3, in the sense that it can be achieved in the game where only the chosen action profiles (s_t^1, \dots, s_t^N) (and not the probability distributions (p_t^1, \dots, p_t^N)) are made public.

E_Γ . In this sense, the sets of correlated equilibria of the mixed extension and of the original finite game are the same. Thus, one cannot hope tighter convergence results by minimizing the internal regret in the mixed extension of the game.

Lemma 25 $\psi(E_{\mathcal{L}^0}) = \psi(E_L) = E_\Gamma$.

Proof The equality between E_Γ and $\psi(E_L)$ is immediate, by linearity and in view of (11).

We now prove that each correlated equilibrium π of Γ may be written as $\psi(\mu)$, where $\mu \in E_{\mathcal{L}^0}$, that is, μ is a probability distribution over $\mathcal{P}(S^1) \times \dots \times \mathcal{P}(S^N)$ that is a correlated equilibrium with respect to all measurable departures.

For a given correlated equilibrium $\pi \in E_\Gamma$, we choose

$$\mu = \sum_{i^1, \dots, i^N} \pi(i^1, \dots, i^N) \delta_{(\delta_{i^1}, \dots, \delta_{i^N})},$$

where δ_{i^j} is the probability distribution over S^j that puts probability mass 1 on i^j . We have to prove that for all players k , for all measurable departures φ_k ,

$$\begin{aligned} \int_{\mathcal{P}(S^1) \times \dots \times \mathcal{P}(S^N)} h^k(p^{-k}, p^k) d\mu(p^1, \dots, p^N) \\ \geq \int_{\mathcal{P}(S^1) \times \dots \times \mathcal{P}(S^N)} h^k(p^{-k}, \varphi_k(p^k)) d\mu(p^1, \dots, p^N). \end{aligned}$$

In view of the form of μ , only the values p^k of the form δ_{i_k} where $i_k \in S^k$ matter in the above integrals. Define a linear mapping L_k from $\mathcal{P}(S^k)$ to $\mathcal{P}(S^k)$ by $L_k(\delta_{i^k}) = \varphi_k(\delta_{i^k})$, for all $i^k \in S^k$. Then,

$$\begin{aligned} \int_{\mathcal{P}(S^1) \times \dots \times \mathcal{P}(S^N)} h^k(p^{-k}, \varphi_k(p^k)) d\mu(p^1, \dots, p^N) \\ = \int_{\mathcal{P}(S^1) \times \dots \times \mathcal{P}(S^N)} h^k(p^{-k}, L_k(p^k)) d\mu(p^1, \dots, p^N). \end{aligned}$$

This concludes the proof in view of the first equality noted above.

We conclude this section by pointing out that the minimization proposed by (10) is, using the terminology of Greenwald and Jafari [15], a matter of Φ -no regret, with Φ including all (extremal) linear functions as well as many other continuous maps. This solves the first half of the question posed in the conclusion of [15]. The second part of the question is to determine if, by performing the regret minimization (10), one could achieve convergence to tighter solution concepts than simply the set of all correlated equilibria. We showed strong evidence that this is not so.

Appendix

Proof of Theorem 5

For $s = (s^1, \dots, s^N) \in S$, we write $\|s\|_\infty = \max_{i=1, \dots, N} \|s^i\|$ where $\|s^i\|$ is the norm of s^i in S^i .

First, we prove the direct implication. Fix π , a $(\mathcal{L}^0(S^k))_{1 \leq k \leq N}$ -correlated equilibrium of the continuous game Γ . Choose any number $\varepsilon > 0$. It suffices to show that there exists a r_0 such that for every discretization of size $r < r_0$, π induces a 2ε -correlated equilibrium.

Each h^k is uniformly continuous, so we may choose r_0 such that for all $k \leq N$, $s, t \in S$, $\|s - t\|_\infty \leq r_0$ implies $|h^k(s) - h^k(t)| \leq \varepsilon$. Fix a discretization $(\mathcal{P}, \mathcal{D})$ of size r less than r_0 .

Fix a player k and a departure $g_k : \mathcal{D}^k \rightarrow \mathcal{D}^k$. We need to prove that

$$\sum_{t \in \mathcal{D}} \pi_d(t) h^k(t) + 2\varepsilon \geq \sum_{t \in \mathcal{D}} \pi_d(t) h^k(t^{-k}, g_k(t^k)) . \quad (12)$$

Define $\psi_k : S \rightarrow S$ by $\psi_k(s^k) = g_k(t_j^k)$ for all $s^k \in V_j^k$. ψ_k is a measurable function. Now, for all $s \in V_{i_1}^1 \times \dots \times V_{i_N}^N$,

$$\left\| (s^{-k}, \psi_k(s^k)) - (t_{i_1}^1, \dots, g_k(t_{i_k}^k), \dots, t_{i_N}^N) \right\|_\infty \leq r .$$

Therefore, due to the uniform continuity of the h^k , we have for all k ,

$$\left| \int_S h^k(s^{-k}, \psi_k(s^k)) d\pi(s) - \sum_{t \in \mathcal{D}} \pi_d(t) h^k(t^{-k}, g_k(t^k)) \right| \leq \varepsilon .$$

It is even easier to see that

$$\left| \int_S h^k(s^{-k}, s^k) d\pi(s) - \sum_{t \in \mathcal{D}} \pi_d(t) h^k(t^{-k}, t^k) \right| \leq \varepsilon .$$

Now, as π is a correlated equilibrium of the original game,

$$\int_S h^k(s^{-k}, s^k) d\pi(s) \geq \int_S h^k(s^{-k}, \psi_k(s^k)) d\pi(s) .$$

Combining the last three inequalities leads to (12), thus concluding the direct part.

The converse implication is proved in a similar way. First, note that thanks to Lemma 22, we can restrict our attention to continuous departures. Assume

that there exists a function ε with $\lim_{r \rightarrow 0} \varepsilon(r) = 0$ such that for all discretizations $(\mathcal{P}, \mathcal{D})$ of size r , π induces an $\varepsilon(r)$ -correlated equilibrium.

Fix an arbitrary $\eta > 0$. We show that for all k and all continuous functions $\psi_k : S^k \rightarrow S^k$,

$$\int_S h^k(s^{-k}, s^k) d\pi(s) + \eta \geq \int_S h^k(s^{-k}, \psi_k(s^k)) d\pi(s). \quad (13)$$

(The conclusion will follow by letting η decrease to 0.)

Fix a player k and a continuous departure ψ_k . As h^k is uniformly continuous, we can choose $\delta > 0$ such that for all $s, t \in S$, $\|s - t\|_\infty \leq \delta$ implies $|h^k(s) - h^k(t)| \leq \eta/3$. Now, ψ_k is also uniformly continuous, so that there exists $\delta' > 0$ such that for all $s, t \in S$, $\|s - t\|_\infty \leq \delta'$ implies $|\psi_k(s) - \psi_k(t)| \leq \delta/2$. Finally, take $r_0 > 0$ sufficiently small so that for all $r \leq r_0$, $\varepsilon(r) \leq \eta/3$. We consider $r = \min(r_0, \delta, \delta')$.

There exists a finite cover of each S^j by open balls of radius r , denoted by $B(x_i^j, r)$, $1 \leq j \leq N$, $1 \leq i \leq N_j$. Each open cover is converted into a measurable partition in the following way. For $1 \leq j \leq N$, $1 \leq i \leq N_j$,

$$V_i^j = B(x_i^j, r) \setminus \left(\bigcup_{m=1}^{i-1} B(x_m^j, r) \right).$$

We take the grid given by the centers, that is, with the above notation, $t_i^j = x_i^j$, $1 \leq j \leq N$, $1 \leq i \leq N_j$. We thus have obtained a discretization of size less than r , and denote by π_d the probability measure induced by π .

We define $g_k : \mathcal{D}^k \rightarrow \mathcal{D}^k$ as follows. For $1 \leq j \leq N_k$, $g_k(x_j^k) = x_m^k$ where $1 \leq m \leq N_k$ is the index such that $\psi_k(x_j^k) \in V_m^k$. Note that in particular, $\|g_k(x_j^k) - \psi_k(x_j^k)\| \leq r \leq \delta/2$.

But if $s^k \in V_{i_k}^k$, $\|s^k - x_{i_k}^k\| \leq r \leq \delta'$, so $\|\psi_k(s^k) - \psi_k(x_{i_k}^k)\| \leq \delta/2$. Finally, $\|\psi_k(s^k) - g_k(x_{i_k}^k)\| \leq \delta$. Thus, if $s \in V_{i_1}^1 \times \dots \times V_{i_N}^N$,

$$\|(s^{-k}, \psi_k(s^k)) - (x_{i_1}^1, \dots, g_k(x_{i_k}^k), \dots, x_{i_N}^N)\|_\infty \leq \delta.$$

Therefore, by uniform continuity of h^k ,

$$\left| \int_S h^k(s^{-k}, \psi_k(s^k)) d\pi(s) - \sum_{x \in \mathcal{D}} \pi_d(x) h^k(x^{-k}, g_k(x^k)) \right| \leq \frac{\eta}{3}.$$

Again, it is even easier to see that

$$\left| \int_S h^k(s^{-k}, s^k) d\pi(s) - \sum_{x \in \mathcal{D}} \pi_d(x) h^k(x^{-k}, x^k) \right| \leq \frac{\eta}{3}.$$

Since π is an $\varepsilon(r)$ -correlated equilibrium (and since $\varepsilon(r) \leq \eta/3$), it is true that

$$\sum_{x \in \mathcal{D}} \pi_d(x) h^k(x) + \frac{\eta}{3} \geq \sum_{x \in \mathcal{D}} \pi_d(x) h^k(x^{-k}, g_k(x^k)) .$$

Combining these last three inequalities, we get (13), concluding the proof. \square

Proof of Lemma 21

Hirsch and Lacombe [21] consider the set of continuous functions defined on a compact metric set X into \mathbb{R} and show that this set is separable (Proposition 1.1). But it turns out that this proof easily extends to the case of Lemma 21, giving, in addition, an example of a dense countable subset of $\mathcal{C}(X)$. We simply need the following well-known lemma, see for instance Rudin [26].

Lemma 26 (Partition of unity) *If X is a locally compact Hausdorff space, then, given a finite number of open sets V_1, \dots, V_N and a compact $K \subset \cup_{i=1, \dots, N} V_i$, there exist N non-negative continuous functions h_1, \dots, h_N summing to 1 over K , such that h_i vanishes outside V_i .*

Proof of Lemma 21 As X is a compact set, for a given $n \in \mathbb{N}^*$, there exist finitely many x_n^j , $j = 1, \dots, N_n$, such that the collection of open balls of common radius $1/n$ and centered in these x_n^j forms a finite cover of X ,

$$X = \cup_{j=1}^{N_n} B(x_n^j, 1/n) .$$

We denote the set formed by these x_n^j by X_n . By Lemma 26 (with $K = X$), denote by φ_j^n , $j = 1, \dots, N_n$, a partition of unity constructed over this open cover of X . We denote by A_n the set formed by

$$A_n = \left\{ \sum_{j=1}^{N_n} y_n^j \varphi_j^n, (y_n^j)_{j=1, \dots, N_n} \in (X_n)^{N_n} \right\} .$$

A_n is a finite set. By convexity of X , each element of A_n maps X into X . By continuity of the φ_j^n , A_n is finally seen as a subset of $\mathcal{C}(X)$.

We consider the countable subset A formed by the union of the A_n , $A = \cup_{n \in \mathbb{N}^*} A_n$, and claim that A is dense in $\mathcal{C}(X)$. To see this, fix a continuous function $f \in \mathcal{C}(X)$. As f maps the compact metric space X into itself, f is uniformly continuous over X . Fix $\varepsilon > 0$ and choose $\delta > 0$ small enough to ensure that $\|x - y\| < \delta$ implies $\|f(x) - f(y)\| < \varepsilon$, where $\|\cdot\|$ denotes the norm of the underlying normed space that contains X . Now, fix a sufficiently large integer n such that $1/n < \min(\delta, \varepsilon)$. For every $j = 1, \dots, N_n$, choose y_n^j

such that $\|y_n^j - f(x_n^j)\| \leq \varepsilon$. Introduce the functions

$$g = \sum_{j=1}^{N_n} f(x_n^j) \varphi_j^n, \quad h = \sum_{j=1}^{N_n} y_n^j \varphi_j^n.$$

It is clear that $h \in A$, and we prove that $\|f - h\|_\infty \leq 2\varepsilon$.

For a given $x \in X$,

$$\|f(x) - g(x)\| = \left\| \sum_{j=1}^{N_n} (f(x) - f(x_n^j)) \varphi_n^j(x) \right\| \leq \sum_{j=1}^{N_n} \|f(x) - f(x_n^j)\| \varphi_n^j(x).$$

Now, $\|f(x) - f(x_n^j)\| \varphi_n^j(x) \leq \varepsilon \varphi_n^j(x)$, simply because φ_n^j vanishes outside $B(x_n^j, 1/n)$ (which is included in $B(x_n^j, \delta)$), whereas, thanks to uniform continuity, the norm of the difference $f(x) - f(x_n^j)$ is less than ε over this ball. Finally, recalling that the φ_n^j sum to 1, we get $\|f - g\|_\infty \leq \varepsilon$.

A similar argument, using the fact that for every j , $\|y_n^j - f(x_n^j)\| \leq \varepsilon$, shows that $\|g - h\|_\infty \leq \varepsilon$, thus concluding the proof. \square

Proof of Lemma 22

The proof is a combination of Lemma 20 and Corollary 28, which is derived from the following version of Lusin's theorem tailored for our needs.

Proposition 27 *If X is a convex and compact subset of a normed space, equipped with a probability measure μ (defined over the Borel σ -algebra), then for every measurable function $f : X \rightarrow X$ and for every $\delta, \varepsilon > 0$, there exists a continuous function $g : X \rightarrow X$ such that*

$$\mu \{ \|f - g\| \geq \delta \} \leq \varepsilon.$$

Proof We use the notation (and the techniques) of the proof of Lemma 21. First note that μ is regular, since it is a finite measure over the Borel σ -algebra of a Polish space (compact metric spaces are Polish).

Fix n large enough such that $1/n < \delta$. Consider the N_n measurable sets

$$M_j^n = f^{-1} \left(B(x_j^n, 1/n) \right).$$

By regularity of μ , one can find compact sets K_j^n and open sets V_j^n such that, for all j ,

$$K_j^n \subset M_j^n \subset V_j^n, \quad \mu(V_j^n \setminus K_j^n) \leq \frac{\varepsilon}{N_n}.$$

By construction, the M_j^n form a cover of X . Therefore, the V_j^n form an open cover of X . By Lemma 26 (with $K = X$), fix a partition of unity based on this open cover, which we denote by $\xi_1^n, \dots, \xi_{N_n}^n$. Consider the continuous function g given by

$$g = \sum_{j=1}^{N_n} x_j^n \xi_j^n .$$

By convexity of X , g maps X into X . Now, as above, for all $x \in X$,

$$\|f(x) - g(x)\| \leq \sum_{j=1}^{N_n} \|f(x) - x_j^n\| \xi_j^n(x) .$$

By construction, $\|f(x) - x_j^n\| \xi_j^n(x) \leq \xi_j^n(x)/n$ provided that $x \in M_j^n \cup (V_j^n)^c$. So, $\|f(x) - g(x)\| \leq 1/n < \delta$, except, possibly, on the measurable subset Δ defined by

$$\Delta = \cup_{j=1}^{N_n} V_j^n \setminus M_j^n ,$$

whose μ -measure is seen to be less than ε by subadditivity of the measure. \square

Now, setting $\delta_n = \varepsilon_n = 1/2^n$, and using Borel-Cantelli lemma, one easily gets the following corollary.

Corollary 28 *If X is a convex and compact subset of a normed space, equipped with a probability measure μ (over the Borel σ -algebra), then every measurable function $f : X \rightarrow X$ may be obtained as a μ -almost sure limit of continuous functions $(g_n)_{n \in \mathbb{N}^*}$ mapping X into X .*

Acknowledgements

We thank Andreu Mas-Colell for many helpful suggestions, one of which led us to Theorem 7. We are grateful to the two anonymous referees for their helpful and profound comments and suggestions. We also thank Thomas Duquesne for interesting discussions and relevant pointers to the literature of functional analysis.

References

- [1] J. Aubin. *Mathematical Methods of Game and Economic Theory*. North Holland Publishing, Amsterdam, 1979.
- [2] P. Auer, N. Cesa-Bianchi, Y. Freund, and R.E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32:48–77, 2002.

- [3] P. Auer, N. Cesa-Bianchi, and C. Gentile. Adaptive and self-confident on-line learning algorithms. *Journal of Computer and System Sciences*, 64(1), 2002.
- [4] R.J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.
- [5] R.J. Aumann. Correlated equilibrium as an expression of bayesian rationality. *Econometrica*, 55:1–18, 1987.
- [6] K. Azuma. Weighted sums of certain dependent random variables. *Tohoku Mathematical Journal*, 68:357–367, 1967.
- [7] D. Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6:1–8, 1956.
- [8] R. Cauty. Solution du problème de point fixe de Schauder. *Fundamenta Mathematicæ*, 170:231–246, 2001.
- [9] N. Cesa-Bianchi, Y. Freund, D.P. Helmbold, D. Haussler, R. Schapire, and M.K. Warmuth. How to use expert advice. *Journal of the ACM*, 44(3):427–485, 1997.
- [10] N. Cesa-Bianchi and G. Lugosi. Potential-based algorithms in on-line prediction and game theory. *Machine Learning*, 51, 2003.
- [11] S. Ethier and T. Kurtz. *Markov processes*. Wiley, New York, 1986.
- [12] D. Foster and R. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behaviour*, 21:40–55, 1997.
- [13] D. Foster and R. Vohra. Regret in the on–line decision problem. *Games and Economic Behavior*, 29:7–36, 1999.
- [14] D. Fudenberg and D. Levine. Universal conditional consistency. *Games and Economic Behavior*, 29:104–130, 1999.
- [15] A. Greenwald and A. Jafari. A general class of no-regret learning algorithms and game-theoretic equilibria. *Proceedings of the 16th Annual Conference on Learning Theory and 7th Kernel Workshop*, 2–12, 2003.
- [16] S. Hart and A. Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68:1127–1150, 2000.
- [17] S. Hart and A. Mas-Colell. A general class of adaptive strategies. *Journal of Economic Theory*, 98:26–54, 2001.
- [18] S. Hart and A. Mas-Colell. A reinforcement procedure leading to correlated equilibrium. In G. Debreu, W. Neuefeind, and W. Trockel, editors, *Economic Essays: A Festschrift for Werner Hildenbrand*, pages 181–200. Springer, New York, 2002.
- [19] S. Hart and A. Mas-Colell. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93:1830–1836, 2003.
- [20] S. Hart and D. Schmeidler. Existence of correlated equilibria. *Mathematics of Operations Research*, 14:18–25, 1989.

- [21] F. Hirsch and G. Lacombe. *Eléments d'Analyse Fonctionnelle*. Masson, Paris, 1997.
- [22] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.
- [23] A.N. Kolmogorov and V.M. Tikhomirov. ε -entropy and ε -capacity of sets in function spaces. *American Mathematical Society Translations*, 17:277–364, 1961.
- [24] E. Lehrer. Approachability in infinite dimensional spaces. *International Journal of Game Theory*, 31:255–270, 1997.
- [25] E. Lehrer. A wide range no-regret theorem. *Games and Economic Behavior*, 42:101–115, 2003.
- [26] W. Rudin. *Real and Complex Analysis*. Mcgraw-Hill, New York. 1974.