

Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory ••• (••••) ••••••••

JOURNAL OF
**Economic
Theory**www.elsevier.com/locate/jet

When is the lowest equilibrium payoff in a repeated game equal to the min max payoff? [☆]

Olivier Gossner ^{a,b,*}, Johannes Hörner ^c^a *Paris School of Economics, 48 Boulevard Jourdan, 75014 Paris, France*^b *Department of Mathematics, London School of Economics, Houghton Street, WC2A 2AE London, United Kingdom*^c *Department of Economics, Yale University, 30 Hillhouse Avenue, New Haven, CT 06520-8281, United States*

Received 30 March 2007; final version received 8 July 2009; accepted 8 July 2009

Abstract

We study the relationship between a player's lowest equilibrium payoff in a repeated game with imperfect monitoring and this player's min max payoff in the corresponding one-shot game. We characterize the signal structures under which these two payoffs coincide for any payoff matrix. Under an identifiability assumption, we further show that, if the monitoring structure of an infinitely repeated game "nearly" satisfies this condition, then these two payoffs are approximately equal, independently of the discount factor. This provides conditions under which existing folk theorems exactly characterize the limiting payoff set.

© 2009 Elsevier Inc. All rights reserved.

JEL classification: C72; C73; D82

Keywords: Folk theorem; Repeated game; Individually rational payoff; Min max payoff; Signals; Entropy; Conditional independence

[☆] The authors thank Tristan Tomala and Jonathan Weinstein for very useful discussions, as well as audiences at Brown University, the University of Pennsylvania, Yale University, the University of Montréal and McGill University. Suggestions from the editor, Christian Hellwig, significantly helped improve the exposition of the paper. They are gratefully acknowledged.

* Corresponding author.

E-mail address: ogossner@gmail.com (O. Gossner).

1. Introduction

Folk theorems aim at characterizing the entire set of payoff vectors that can be attained at equilibrium in repeated games. That is, the purpose of a folk theorem is to determine which payoff vectors are, and which are not, equilibrium payoffs when players are sufficiently patient. While the hallmark of this literature lies in the wealth of payoffs that can be supported, it might not be as well known that, in notable cases, these results only provide lower bounds on the limit set of equilibrium payoffs, rather than actual characterizations.

For instance, the folk theorem under imperfect public monitoring [7] asserts that, under some statistical conditions, every feasible and individually rational payoff vector is an equilibrium payoff vector under low discounting. Individual rationality refers to the (stage-game, mixed) min max payoff, defined as

$$\min_{\alpha_{-i} \in \prod_{j \neq i} \Delta A_j} \max_{a_i \in A_i} g_i(a_i, \alpha_{-i}),$$

where $a_i \in A_i$ is player i 's pure action, $\alpha_j \in \Delta A_j$ is player j 's mixed action and g_i is player i 's payoff function. That is, the min max payoff is the lowest payoff player i 's opponents can hold him to in the stage game by any choice α_{-i} of independent actions, provided that player i correctly foresees α_{-i} and plays a best-response to it.

Yet as Fudenberg, Levine and Maskin [7] acknowledge, in some games, equilibria can be constructed in which a player's equilibrium payoff is strictly lower than his min max payoff. (See Exercise 5.10 in [9] for an illuminating example.) This is because actions are unobserved, so that, if the stage game is such that player i 's correlated min max payoff is strictly below his min max payoff, players $-i$ might be able to use their private histories to correlate their actions.¹

This paper provides conditions under which the min max payoff provides a tight bound to the equilibrium set, in repeated games with imperfect public, or private monitoring. Doing so does not merely provide a converse for some of these folk theorems, but also helps understand in which situations there is scope for punishments that are harsher than those usually assumed. This is important because, to compute the greatest equilibrium payoff for a fixed discount factor, one must typically also compute the lowest such payoff.

To understand the statistical requirement under which the min max payoff provides the lower bound on the equilibrium payoff set, we start our analysis with static Bayesian games: each player receives a payoff-irrelevant signal before choosing his action. In this framework, we characterize the correlation devices that do not lead to equilibrium payoffs that are strictly worse than the uncorrelated minmax payoffs.²

It is rather immediate to see that a player i can always assure himself of no less than his uncorrelated minmax payoff if the signal structure is such that either the other player's signals are independently distributed conditional on player i 's signal, or there exists a garbling of player i 's signal, conditional on which the other player's signals are independently distributed. By conditioning his actions on such an "independence-inducing" garbling, player i ensures himself against the possibility that the other players use their signals to correlate their actions. We prove

¹ The definition of correlated min max payoff is obtained by replacing $\prod_{j \neq i} \Delta A_j$ by $\Delta \prod_{j \neq i} A_j$ as the domain of the minimization in the definition of the min max payoff.

² A complementary problem is to determine the payoff matrices for which no signal structure allows some player to be held down below his min max payoff. The answer is rather immediate, as it amounts to comparing the min max and the correlated min max payoff of the payoff matrix. Our question is motivated by the folk theorems, in which conditions are identified on the signal structure that are sufficient for all games.

that the existence of an independence-inducing garbling is not only a sufficient condition, but is also necessary: the existence of an independence-inducing garbling for player i characterizes the correlation structures for which player i cannot be held below his min max payoff.

In repeated games, studied next, signals are influenced by actions. Thus, signal and action sets can no longer be treated independently. The condition must be modified: a player's "signal" includes now both his own action and the actual signal he observed. The same conditional independence (for some garbling) requirement guarantees that the (stage-game) min max payoff and the repeated game min max payoff—the lowest payoff player i 's opponents can hold him to by any choice of strategies in the repeated game—coincide. This is not only a feasibility statement, but also an equilibrium statement. Indeed, by a result of von Stengel and Koller [27], the repeated game min max payoff in a given game is an equilibrium payoff in the game obtained by setting the payoff of all players but i equal to the opposite of player i 's payoff in the original game. In this sense, the result is tight: if the condition is violated, there exists a payoff matrix for which the lowest equilibrium payoff is below the stage-game min max payoff; if it is satisfied, then the lowest equilibrium payoff is greater than, and for some payoff matrices equal to, the min max payoff.

Because a growing literature examines the robustness of folk theorems with respect to small perturbations in the monitoring structures, starting from either perfect monitoring (see [25,6,2,23,14]), or imperfect public monitoring [18,19], we actually prove a stronger result: as the distance of the monitoring structure to any monitoring structure satisfying the aforementioned condition converges to zero, so does the distance between the stage-game and repeated-game min max payoffs. This convergence is uniform in the discount factor, provided that the monitoring structure satisfies some standard identifiability assumptions.

The condition that is identified is by no means mild: as mentioned, there are known examples in games with public monitoring, where the stage-game and repeated-game min max payoffs fail to coincide. In fact, we provide simple examples to show that this is possible even when:

- monitoring is almost-perfect;
- the punished player perfectly monitors his opponents.

But our result implies, for instance, that the two min max payoffs are arbitrarily close if monitoring is almost-perfect monitoring *and* attention is restricted to the canonical signal structure, in which players' signals are (not necessarily correct) action profiles of their opponents. This provides a converse to Theorem 1 of [14]. Our condition also generalizes the various special cases for which it is well known that these two payoffs coincide, namely:

- if there are only two players, as correlated min max and min max payoffs then coincide;
- if monitoring is perfect, as all players then hold the same information at any point in time, so that the probability distribution over player i 's opponents' actions given his information corresponds to independent randomizations by his opponents;³
- if monitoring is public, but information is semi-standard (as in [16]);
- if monitoring is public, but attention is restricted to public strategies, as in this case as well the information relevant to forecasting future play is commonly known.

Section 2 presents examples that motivate the characterization. Section 3 considers static Bayesian games. Section 4 extends the analysis to the case of infinitely repeated games. Section 5 concludes.

³ See, among others, [1,24,8].




Fig. 1. The duenna game.

2. The duenna game

Many of the ideas can be conveyed through a simple example, which we call the duenna game.⁴ Two lovers (players 1 and 2) attempt to coordinate on a place of secret rendezvous. They can either meet on the landscape garden bridge (B) or at the woodcutter’s cottage (C). Unfortunately, the incorruptible duenna (player 3) prevents any communication between them, and wishes to disrupt their meeting. Therefore, the rendezvous only succeeds if both lovers choose the same place and the duenna picks the other place. In all other cases, the duenna exults.

The common payoff to the lovers is the probability of a successful meeting, and the duenna’s payoff is the opposite of this probability. Fig. 1 displays this probability, as a function of the players’ actions (lovers choose row and column; the duenna chooses the matrix).

In the absence of any correlation device, players 1 and 2 (the “team”) can secure a common payoff of 1/4 by randomizing evenly and independently. This payoff of 1/4 is also the best equilibrium payoff for the team, as it is an equilibrium that all three players randomize evenly. Yet if the team could secretly coordinate, they could guarantee a probability of 1/2, by randomizing evenly between (B, B) and (C, C).

Now, suppose that this game is repeated infinitely often, and that monitoring is imperfect. Let Ω_i denote player i ’s (finite) set of signal, with generic element ω_i . The distribution of $\omega := (\omega_1, \omega_2, \omega_3) \in \Omega := \prod_i \Omega_i$ under action profile $a \in A$ is denoted q^a , with marginal distribution on player i ’s signal given by q_i^a . A monitoring structure is denoted (Ω, q) , where $q := \{q^a : a \in A\}$. We examine, for different examples of monitoring structures, whether signals allow the team to generate some amount of secret correlation or not.

Example 1 (Almost-perfect monitoring). Recall from [18] that the monitoring structure (Ω, q) is ε -perfect if there exist signaling functions $f_i : \Omega_i \rightarrow A_{-i}$ for all i such that, for all $a \in A, i = 1, 2, 3$:

$$q^a(\{\omega : f_i(\omega_i) = a_{-i}\}) \geq 1 - \varepsilon.$$

That is, a monitoring structure is ε -perfect if the probability that the action profile suggested by the signal is incorrect does not exceed $\varepsilon > 0$, for all possible action profiles.

Let $\Omega_1 = \{\omega_1^a, \omega_1'^a : a \in A\}, \Omega_2 = \{\omega_2^a, \omega_2'^a : a \in A\}, \Omega_3 = \{\omega_3^a : a \in A\}$. Consider

$$q^a((\omega_1^a, \omega_2^a, \omega_3^a)) = q^a((\omega_1'^a, \omega_2'^a, \omega_3^a)) = \frac{1 - \varepsilon}{2}, \quad \text{all } a \in A,$$

⁴ This game, which appears in various place in the literature, is sometimes referred to as the “three player matching pennies” game (see [20]). We find this name slightly confusing, given that the “three person matching pennies” game introduced earlier by Jordan [15] is a different, perhaps more natural generalization of matching pennies.

where $\varepsilon > 0$ is small enough, and set $f_3(\omega_3^a) = a$, $f_i(\omega_i^a) = f_i(\omega_i^{a'}) = a$, all $i = 1, 2$ and $a \in A$. The specification of the remaining probabilities is arbitrary. Observe that monitoring is ε -perfect, since the probability that a player receives either ω_i^a or $\omega_i^{a'}$ is at least $1 - \varepsilon$ under action profile a .

Yet players 1 and 2 can secure $(1 - \varepsilon)/2 \rightarrow_{\varepsilon \rightarrow 0} 1/2$, as the discount factor tends to one. Indeed, they can play B if ω_i^a is observed at the previous stage, and C if $\omega_i^{a'}$ is observed at the previous stage, independently of a . Therefore, even under almost-perfect monitoring, the payoff of player 3 in this equilibrium is bounded away from his min max payoff.

Example 1 illustrates that the set of equilibrium payoffs under almost-perfect monitoring may be bounded away from the one under perfect monitoring. In this example, the set of signals is richer under imperfect private monitoring than under perfect monitoring. Therefore, one may argue that the comparison of the min max levels across monitoring structures is not appropriate. In this example, the natural “limiting” monitoring structure, as $\varepsilon \rightarrow 0$, should allow for a private correlation device for players 1 and 2. Indeed, it is the case that the repeated-game min max payoff is a continuous function of the signal distribution for fixed sets of signals and a fixed discount factor.

But restricting the set of signals does not rule out correlation, because it also arises under the *canonical* signal structure in which $\Omega_i = A_{-i}$, as shown by our next example. Nevertheless, our main result implies that, if the monitoring structure is almost-perfect *and* canonical, then both min max payoffs coincide.

Example 2 shows that it is not enough to require that player 3 have perfect information about his opponents’ actions, and/or that the signal structure be canonical.

Example 2 (*Perfect monitoring by player 3, canonical signal structure*). Each player’s set of signals is equal to his opponents’ set of actions: $\Omega_i = A_{-i}$, for all i . Player 3’s information is perfect:

$$q_3^a(a_{-3}) = 1, \quad \forall a \in A.$$

Player 1 perfectly observes player 2’s action, and similarly player 2 perfectly observes player 1’s action. Their signal about player 3’s action is independent of the action profile, but perfectly correlated. In particular:

$$\begin{aligned} q_1^a((a_2, C)) &= q_1^a((a_2, B)) = 1/2, \\ q_2^a((a_1, C)) &= q_2^a((a_1, B)) = 1/2. \end{aligned}$$

Consider the following strategies for players $i = 1, 2$: randomize uniformly in the first period; afterwards, play C if the last signal about player 3 is C , and B otherwise. This guarantees $1/2$.

In this example, player 1 and 2’s signals are uninformative about player 3’s action, but it is easy to construct variations in which their signals are arbitrarily informative, and yet such that the min max payoff is bounded away from the repeated-game min max payoff. One may argue that the problem here is that player 3’s signal set is not nearly rich enough, as it does not include his opponents’ signal about his own action. However, enlarging the signal sets takes us back to our initial example.

The issue is not solved either by requiring that the player’s signals be almost public, a stronger requirement introduced and studied in [19]. Indeed, even under public monitoring, the repeated-game min max payoff may be lower than the min max payoff (see [9, Exercise 5.10]).




Fig. 2. Conditional distributions and garbled conditional distributions.

In both examples, players 1 and 2 are able to secretly coordinate in the context of monitoring structures close to standard structures for the player that is punished. The reader may have guessed by now a condition ruling out any such example: if, conditional on any signal of player 3 that has positive probability for some action profile, player 1 and 2’s signals are independent, then players 1 and 2 cannot secretly correlate. This ensures that the probability distribution over player 3’s opponents’ actions, given his information, corresponds to independent randomizations. The next example shows that this condition is, however, stronger than necessary.

Example 3 (A monitoring structure without conditional independence for which the repeated-game min max payoff equals the stage-game min max payoff). For each player i , $\Omega_i = \{\omega_i, \omega_i'\}$. Probabilities of signal profiles are independent of the action profile. Signals ω_3 and ω_3' have probability 1/2. The distribution of player 1 and 2’s signals, conditional on player 3’s signal, is given in Fig. 2’s left panel. Player 1 and 2’s signals are not independent conditional on any value of player 3’s signal. Yet we claim that, in any game that may be played along with this signal structure, player 3 guarantees her min max payoff.

Why? Observe that player 3 can always decide to “garble” his information, and base his decision on the garbled information, as summarized by two fictitious signals, $\tilde{\omega}_3$ and $\tilde{\omega}_3'$. Upon receiving signal ω_3 , he can use a random device selecting $\tilde{\omega}_3$ with probability 1/5, and selecting $\tilde{\omega}_3'$ otherwise; similarly, upon receiving signal ω_3' , he can use a device selecting $\tilde{\omega}_3'$ with probability 1/5 and $\tilde{\omega}_3$ otherwise.

The right panel of Fig. 2 shows the distribution of player 1 and 2’s signals, conditional on the value of the garbled signal of player 3. Note that player 1 and 2’s signal are independent, conditionally on any value of player 3’s signal.

We now show how, responding to players 1 and 2’s strategies, player 3 can prevent players 1 and 2 from obtaining more than 1/4 in the repeated game. In the first stage of the repeated game, player 3 plays a best-response to the mixed strategies of players 1 and 2. In the second stage, player 3 can garble the signal of the first stage, and play according to this garbled signal only: conditionally on the garbled signal, the signals of players 1 and 2 in the first stage are independent, and so are their actions in the second stage. Therefore, by playing a best-response to the distribution of actions of players 1 and 2 in the second stage given his garbled signal, player 3 ensures that players 1 and 2 do not receive more than 1/4 in the second stage. The construction extends to any repetition of the game, as Corollary 3 establishes more generally.

This example shows the connection between the repeated-game min max payoff and the existence of a garbling of player 3’s signal satisfying conditional independence. To disentangle the role of actions and signals, we first abstract from repeated games and pose our problem as a static Bayesian game with exogenous signals.

3. Static games

As the previous examples show, imperfect monitoring may allow a group of players to secretly coordinate their actions at the expense of another player.

In this section, we examine when such coordination is possible in a static framework in which signals are exogenous and a game is played only once. This allows us to leave aside for now difficulties that arise due to the dynamic dimension of repeated games. For instance, in the repeated game, in each period, player i 's opponents must trade off the generation of such correlation for future use, and the immediate use of the existing correlation at the expense of player i .

The main result of this section characterizes the distributions of signals for which it is possible, for some payoff function of a given player, to drive his payoff below the min max payoff.

We start out by introducing the notions of an information structure q , and of garblings between these information structures in Section 3.1. A central notion is that of an *independence inducing garbling*, which generalizes the idea underlying Example 3. Garblings and independence inducing garblings are defined in informational terms only, without any reference to payoffs. Turning to payoffs and strategic notions in Section 3.2, we define games extended by an information structure, and compare the min max payoff for some player between the game extended by the information structure, and the game without any information structure. We say that an information structure is min max preserving if, in every game, the two values coincide. In Section 3.3 we present the main result of this section, characterizing the min max preserving information structures in terms of independence inducing garbling of signals.

3.1. Information structures and independence inducing garblings

Given any measurable space B , ΔB denotes the set of probability distributions over B , and when B is a subset of a vector space, $\text{co} B$ denotes the convex hull of B . Given a collection of sets $\{B_i\}$, $\prod_i B_i$ denotes the Cartesian product of these sets. When each B_i has a measurable structure, a product distribution Q over $\prod_i B_i$ is one that is obtained by the product of its marginals: $Q(\prod_i B_i) = \prod_i Q(B_i)$ for any collection $(B_i)_i$ of measurable sets in B_i . The set of product distributions over $\prod_i B_i$ is identified with $\prod_i \Delta B_i$.

An *information structure* is given by a finite set of signals Ω_i for each player $i = 1, \dots, n$ ($n \geq 1$) along with a probability distribution q over $\Omega := \prod_i \Omega_i$. We denote such an information structure by q .

A *garbling* (for player i) is a family $p = (p_{\omega_i})_{\omega_i \in \Omega_i}$ of probabilities over a measurable space X . The interpretation is that, upon receiving signal ω_i in the information structure q , player i randomly draws a signal in X , according to the probability distribution p_{ω_i} .

The probability q over Ω together with the garbling p induces a probability over $\Omega \times X$, denoted by $p \otimes q$, and given by $(p \otimes q)(\omega, \mathcal{S}_i) = q(\omega) \times p_{\omega_i}(\mathcal{S}_i)$, for every $\omega \in \Omega$ and measurable set \mathcal{S}_i of X .

Definition 1. An *independence inducing garbling* of q for player i is a garbling of q such that, almost surely with respect to the garbled signal x , the distribution of signals of players other than i given x is a product distribution:

$$(p \otimes q)(\cdot | x) \in \prod_{j \neq i} \Delta \Omega_j \quad \text{a.s.}$$

That is, p is independence inducing whenever, conditional on player i 's garbled signal, other players' signals are independently distributed. In Example 3, the device that randomly selects $\tilde{\omega}_3$ or $\tilde{\omega}'_3$ from ω_3 or ω'_3 is an independence inducing garbling.

Note that the existence of an independence inducing garbling for player i is a property of the information structure q only. In particular, it does not involve payoffs.

3.2. Min max preserving information structures

A normal form game G specifies a set of players $j = 1, \dots, n$; for each player j , a finite action set A_j , and a payoff function $g_j: A := \prod_j A_j \rightarrow \mathbb{R}$. We single out some player i , and all payoff functions for players $j \neq i$ are irrelevant for our purposes. We let $S_j := \Delta A_j$ be the set of mixed strategies of player j , and extend g_i to ΔA (and in particular to $S := \prod_j S_j$) using the expectation in the usual way.

The min max payoff of player i in the game G (without signals) is:

$$v_i(G) = \min_{\alpha_{-i} \in \prod_{j \neq i} \Delta A_j} \max_{a_i \in A_i} g_i(a_i, \alpha_{-i}). \tag{1}$$

Given a game G and an information structure q , we denote by $\Gamma(q, G)$ the game obtained from G by adjoining the signal structure q (where the set of signals is given by the domain of q). A (behavioral) strategy $\sigma_j \in \Sigma_j$ for player j in $\Gamma(q, G)$ is a mapping from Ω_j (j 's set of signals in q) to ΔA_j . A profile of strategies $\sigma = (\sigma_j)_j$ induces, for each $\omega \in \Omega$, a profile of mixed strategies $\sigma(\omega) = (\sigma_j(\omega_j))_j \in \prod_j \Delta A_j$. To σ corresponds the payoff for player i in $\Gamma(q, G)$,

$$\gamma_i(\sigma) = \mathbf{E}_q g_i(\sigma(\omega)),$$

where \mathbf{E}_q denotes the expectation under the probability distribution q .

The min max for player i in $\Gamma(q, G)$ is thus

$$V_i(q, G) = \min_{\sigma_{-i} \in \prod_{j \neq i} \Sigma_j} \max_{\sigma_i \in \Sigma_i} \gamma_i(\sigma).$$

For every game G and information structure q , $V_i(q, G) \leq v_i(G)$. To see this, consider a profile $(\alpha^*)_{j \neq i}$ that achieves the minimum in the definition of v_i and the strategies σ_j for $j \neq i$ given by $\sigma_j(\omega_j) = \alpha_j^*$ independently of ω_j . A best-response of player i against these strategies yields exactly $v_i(G)$ to player i .

We are interested in characterizing the information structures q that give scope to strategic correlation between players $j \neq i$, in the sense that they allow for a punishment of player i that is strictly lower than $v_i(G)$. Of course, whether $V_i(q, G) < v_i(G)$ or $V_i(q, G) = v_i(G)$ is not merely a property of q , as it also depends on G . For instance, $V_i(q, G) = v_i(G)$ for every q whenever g_i is constant. Hence, the appropriate notion of an information structure that gives rise to strategic correlation is the one requiring that $V_i(q, G) < v_i(G)$ for *some* game G . If the information structure q does not give rise to such correlation, the min max in G and in $\Gamma(q, G)$ coincide for all games G . This idea is captured by the following definition.

Definition 2. An information structure q is min max preserving (for player i) whenever, for every game G ,

$$V_i(q, G) = v_i(G).$$

Thus, if q is not min max preserving, there exists a game G in which the min max payoff to player i is strictly lower when players have access to signals distributed according to q .

3.3. The equivalence result

We characterize the strategic notion of a min max preserving information structure using the informational notion of an independence inducing garbling.

Theorem 1. *An information structure q is min max preserving for player i if and only if there exists an independence inducing garbling of q for player i .*

Proof. “If” part. Let $p = (p_{\omega_i})_{\omega_i}$ be an independence inducing garbling of q for player i , and let G be a normal form game. Consider strategies $\sigma_{-i} = (\sigma_j)_{j \neq i}$ for players $j \neq i$ in $\Gamma(q, G)$.

For every $x \in X$ such that $(p \otimes q)(\cdot|x)$ belongs to $\prod_j \Delta\Omega_j$, let $\tau(x) \in S_i$ be such that $g_i(\tau(x), (p \otimes q)(\cdot|x)) \geq v_i$. We define a response σ_i for player i to σ_{-i} by $\sigma_i(\omega_i) = \mathbf{E}_{p_{\omega_i}} \tau(x)$. The strategy σ_i consists in first applying the garbling p_{ω_i} to player i 's signal, then choosing the action according to τ , given the resulting garbled signal.

We now verify that σ_i allows player i to defend $v_i(G)$ against σ_{-i} :

$$\begin{aligned} \gamma_i(\sigma_{-i}, \sigma_i) &= \mathbf{E}_\omega g(\sigma(\omega)) \\ &= \mathbf{E}_\omega \mathbf{E}_x g_i(\sigma_{-i}(\omega_{-i}), \tau_i(x)) \\ &= \mathbf{E}_x \mathbf{E}_{(p \otimes q)(\cdot|x)} g_i(\sigma_{-i}(\omega_{-i}), \tau_i(x)) \\ &\geq v_i(G). \end{aligned}$$

“Only if” part. Given $q \in \Delta\Omega$, let ω be a random variable with law q . Player i 's belief on ω_{-i} given signal ω_i is $q(\omega_{-i}|\omega_i)$, which we view as a random variable with values in $\Delta\Omega_{-i}$. We let β_q denote the distribution of this random variable (note that β_q depends only on q , not on the particular choice of random variable ω). This distribution β_q is the distribution of beliefs of player i about the other players' signals induced by q . It characterizes the information about ω_{-i} contained in ω_i .

Assuming that q is min max preserving, we prove the existence of an independence inducing garbling p of q having the further property that the garbled signal is identified with the conditional distribution of signals of players $-i$ given player i 's garbled signal. Let $M := \prod_{j \neq i} \Delta\Omega_j$ be the set of product distributions on $\Omega_{-i} = \prod_{j \neq i} \Omega_j$. We define an M -garbling of q as a garbling p with M as set of garbled signals and such that

$$(p \otimes q)(\omega_{-i}|m) = m \quad \text{a.s. in } m.$$

Obviously, any such M -garbling is independence inducing. Given some M -garbling p of q , let μ_p denote the distribution of the garbled signal, i.e. the marginal of $p \otimes q$ on M . This represents the distribution of beliefs on ω_{-i} of some hypothetical agent informed of m , but not of ω_i .

To prove the existence of an M -garbling, we rely on the following result, which is a characterization of more informative experiments *à la* Blackwell [3].

Lemma 1. *Let q be an information structure and $\mu \in \Delta M$. There exists an M -garbling p of q such that $\mu_p = \mu$ if and only if, for every bounded convex function ψ on $\Delta\Omega_{-i}$,*

$$\mathbf{E}_{\beta_q} \psi \geq \mathbf{E}_\mu \psi.$$

This lemma is an extension of the theorem of [3] to the infinite-dimensional case. See [4,26].

We are now in position to complete the proof of Theorem 1. Endow $A := \Delta M$, subset of the set of distributions over $D = \Delta \Omega_{-i}$, with the weak* topology. Let B be the set of continuous convex functions on D bounded by 0 and 1, endowed with the uniform topology. Both A and B are compact convex sets.

Assume that q admits no M -garbling for player i . By Lemma 1,

$$\forall \mu \in A, \exists \psi \in B \quad \text{such that} \quad \mathbf{E}_{\beta_q} \psi < \mathbf{E}_\mu \psi. \tag{2}$$

The map $g: A \times B \rightarrow \mathbb{R}$ given by $g(\mu, \psi) = \mathbf{E}_{\beta_q} \psi - \mathbf{E}_\mu \psi$ is bi-linear and continuous, so, by the min max theorem, the two-player zero-sum game in which player I 's action set is A , player II 's action set is B and the payoff to I is given by g has a value v , and (2) implies $v < 0$. There exists an optimal strategy for player II , which is $\psi \in B$ such that

$$\forall \mu \in A, \quad \mathbf{E}_{\beta_q} \psi \leq \mathbf{E}_\mu \psi + v.$$

In particular,

$$\forall m \in M, \quad \mathbf{E}_{\beta_q} \psi \leq \psi(m) + v. \tag{3}$$

Let $\psi' = \psi - \mathbf{E}_{\beta_q} \psi + v/2$. We have $\mathbf{E}_{\beta_q} \psi' = v/2 < 0$, and (3) implies that, for every $m \in M$, $\psi'(m) \geq -v/2 > 0$.

For every $m \in M$, the convexity of ψ' implies the existence of a linear map ϕ_m on D such that $\phi_m(m) > 0$ and $\phi_m \leq \psi'$. Let O_m be an open neighborhood of m such that $\phi_m > 0$ on O_m . Since M is a compact subspace of D and $(O_m)_{m \in M}$ is a covering of M , there exists a finite subcovering $(O_m)_{m \in M_0}$ of M .

We use the family $(\phi_m)_{m \in M_0}$ to construct a game G showing that q is not min max preserving. In G , each player $j \neq i$ has strategy set Ω_j , player i has strategy set M_0 , and i 's payoff function g_i is defined by $g_i(m_0, \omega_{-i}) = \phi_{m_0}(\omega_{-i})$.

Let $m \in M$ be a profile of mixed strategies for players $j \neq i$. For $m_0 \in M_0$ such that $m \in O_{m_0}$, $\mathbf{E}_m g_i(m_0, \omega_{-i}) = \phi_{m_0}(m) > 0$. Hence,

$$v_i(G) = \min_{m \in M} \max_{m_0 \in M_0} g_i(m_0, m) > 0.$$

Now consider G extended by q , and the strategies for players $j \neq i$ that specify as actions in G their signal in q . Given a signal ω_i , a best-response of player i yields an expected payoff of

$$\max_{m_0 \in M_0} \phi_{m_0}(q(\omega_{-i} | \omega_i)).$$

Thus, a best-response strategy for player i yields an expected payoff in $\Gamma(q, G)$ of

$$\begin{aligned} \mathbf{E}_q \max_{m_0 \in M_0} \phi_{m_0}(q(\omega_{-i} | \omega_i)) &\leq \mathbf{E}_q \psi'(q(\omega_{-i} | \omega_i)) \\ &= \mathbf{E}_{\beta_q} \psi' \\ &< 0, \end{aligned}$$

which shows that $V_i(q, G) < 0$.

We have thus shown that if q admits no M -garbling, q is not min max preserving. \square

4. Repeated games with imperfect monitoring

Theorem 1 suggests that the existence of an independence inducing garbling is the appropriate condition for our purposes. However, in the repeated game, the problem is slightly more complicated. After a given history, the private information of a player (what we called his “signal” in the previous section) has two components: the actions that he has taken so far, and the actual signal that he has observed in each period. Therefore, for our purposes, a “signal” in a given period is such a pair. It is clear, however, that, unlike in the previous section, the distribution of this pair is no longer exogenous, both because it contains a player’s own past action, which he chose, and a signal whose distribution his choices of action affected. Since there are as many distributions as action profiles, the condition must be strengthened to the existence of a garbling providing conditional independence, for each possible action profile, whether pure or mixed.⁵ The sufficiency of this condition is established in Theorem 2.

Second, in many applications, conditional independence need not hold exactly (consider, for instance, a monitoring structure that is almost-perfect, but not perfect). Therefore, we wish to allow for small departures from conditional independence, which complicates the analysis, especially since we aim for a bound that is uniform in the discount factor. Since even small departures from conditional independence may allow patient players to accumulate secret correlation over time (as we show in Example 4 below), such uniformity can only be achieved if player i ’s opponents necessarily dissipate this correlation whenever they take advantage of it. Theorem 4, which is the main result of this section, formalizes this logic.

Recall that a stage game G specifies a (finite) action set A_j for each player $i = 1, \dots, n$ and, for each player i , a payoff function $g_i: A := \prod_j A_j \rightarrow \mathbb{R}$. We restrict attention to games G for which $|g_i(a)| \leq 1$ for every player i and action profile a (the specific choice of the upper bound is obviously irrelevant).

Players can use mixed actions $\alpha_i \in \Delta A_i$. Mixed actions are unobservable. No public randomization device is assumed, and there is no communication.

We consider the infinitely repeated game, denoted G^∞ . Periods are indexed by $t = 0, 1, \dots$. In each period, player i observes a private signal ω_i from some finite set Ω_i , whose distribution depends on the action profile being played. Therefore, player i ’s information acquired in period t consists of both his action a_i and his private signal ω_i . Let $s_i = (a_i, \omega_i)$ denote this information, or *signal* for short, and define $S_i := A_i \times \Omega_i$. The monitoring structure determines a distribution over private signals for each action profile. For our purpose, it is more convenient to define it directly as a distribution over $S := S_1 \times \dots \times S_n$. Given action profile $a \in A$, $q^a(s)$ denotes the distribution over signal profiles $s = (s_1, \dots, s_n)$. We extend the domain of this distribution to mixed action profiles $\alpha \in \Delta A$, and write $q^\alpha(s)$. Let q_i^α denote the marginal distribution of q^α over player i ’s signals s_i , and given $s_i \in S_i$ and $\alpha \in \Delta A$ such that $q_i^\alpha(s_i) > 0$, let $q_{-i}^\alpha(\cdot | s_i)$ denote the marginal distribution over his opponents’ signals, conditional on player i ’s signal s_i . From now on, a monitoring structure refers to such a family of distributions q .

Players share a common discount factor $\delta \in (0, 1)$, but as will be clear, its specific value is irrelevant (statements do not require that it be sufficiently large). Repeated game payoffs are discounted, and their domain is extended to mixed strategies in the usual fashion; unless explicitly mentioned otherwise (as will occur), all payoffs are normalized by a factor $1 - \delta$.

Recall that player i ’s min max payoff $v_i(G)$ in G is defined by Eq. (1).

⁵ Conditional independence of signals for each pure action profile does not imply conditional independence for all mixed action profiles.

A private history of length t for player i is an element of $H_i^t := S_i^t$ (let $H_i^0 := \{\emptyset\}$). A (behavioral) private strategy for player i is a family $\sigma_i = (\sigma_i^t)_t$, where $\sigma_i^t: H_i^{t-1} \rightarrow \Delta A_i$. We denote by Σ_i the set of these strategies. Bearing in mind the earlier discussions regarding the omission of any equilibrium consideration here, we define player i 's individually rational payoff v_i^δ in the repeated game as the lowest payoff he can be held down to by any collection $\sigma_{-i} = (\sigma_j)_{j \neq i}$ of independent choices of strategies in the repeated game, provided that player i correctly foresees σ_{-i} and plays a best-reply to it. Formally, the *individually rational* payoff (or min max payoff in the repeated game) is defined as

$$v_i^\delta := \min_{\sigma_{-i} \in \prod_{j \neq i} \Sigma_j} \max_{\sigma_i \in \Sigma_i} \mathbf{E}_\sigma \sum_{t=0}^{\infty} (1-\delta)\delta^t g_i(a_i^t, a_{-i}^t).$$

It is straightforward to see that, for any game G and monitoring structure q , the individually rational payoff does not exceed the min max payoff:

$$v_i(G) \geq v_i^\delta.$$

This is a consequence of the fact that, in the repeated game, players $-i$ can repeatedly play a profile of mixed strategies that achieves the minimum in the definition of $v_i(G)$.

We wish to identify conditions under which the individually rational payoff and the (stage-game) min max payoff coincide, or are close to one another. The first condition we define is that of conditional independence.

Definition 3. A monitoring structure q satisfies conditional independence for player i if, for every profile of mixed strategies $\alpha \in \prod_j \Delta A_j$, player $-i$'s signals are independent conditional on player i 's signal:

$$\forall s_i \in S_i \quad \text{such that} \quad q_i^\alpha(s_i) > 0, \quad q_{-i}^\alpha(\cdot | s_i) \in \prod_{j \neq i} \Delta S_j.$$

Theorem 2. If the monitoring structure satisfies conditional independence for player i , then, for all $\delta \in [0, 1)$, player i 's individually rational payoff is equal to his min max payoff.

This result is proved in the next section, as an immediate consequence of the first step of the proof of Theorem 4. To state the next important but straightforward extension of Theorem 2, one must, in the spirit of Section 3, introduce the notion of an independence inducing garbling of a monitoring structure.

Definition 4. An *independence inducing garbling* of a monitoring structure q for player i is a garbling p such that p is an independence inducing garbling of the information structure q^α , for every profile of mixed strategies α .

Theorem 2 implies the following:

Corollary 3. If the monitoring structure admits an independence inducing garbling for player i , then, for all $\delta \in [0, 1)$, player i 's individually rational payoff is equal to his min max payoff.

Observe that this corollary generalizes Theorem 2, as if q satisfies conditional independence for player i , it automatically admits an independence inducing garbling for player i .

Another situation in which an independence inducing garbling exists is when if q_{-i}^α is a product distribution, that is, if signals of players $j \neq i$ are independently distributed, because player i can always ignore his signal, and play a best-reply to his prior belief. This may occur for distributions in which q does not satisfy conditional independence for player i .

In fact, Corollary 3 encompasses a set of situations that is much richer than those special cases. However, while it is straightforward to check whether a monitoring structure satisfies conditional independence for any given player, we do not know of a simple algorithm allowing to ascertain whether a monitoring structure admits an independence inducing garbling for this player.

For a given discount factor, since the payoff function in the repeated game is continuous in the monitoring structure, the individually rational payoff is also continuous in the monitoring structure. In particular, if q “almost” satisfies conditional independence for player i , or “almost” admits an independence inducing garbling for player i , then the individually rational payoff for player i is approximately equal to i ’s min max payoff.

Such a result is unsatisfactory because it does not rule out that, for such monitoring structures, there may exist a discount factor, sufficiently close to one, for which the individually rational payoff is bounded away from the min max payoff.

Intuitively, monitoring structures that “almost” admit an independence inducing garbling may still provide “small” amounts of correlation to player i ’s opponents. Over time, these small amounts may accumulate, allowing them to successfully coordinate their play eventually. The next example illustrates this possibility.

Example 4 (*A monitoring structure that “almost” satisfies independence*). The payoff matrix is given by the duenna game. Player 1 and 2’s signal set each has two elements, $\Omega_i = \{\omega_i, \omega'_i\}$. Player 3 receives no signal. The distribution of player 1 and 2’s signals is independent of the action profile, and perfectly correlated. With probability $\epsilon > 0$, the signal profile is (ω_1, ω_2) , and it is equal to (ω'_1, ω'_2) with probability $1 - \epsilon > 0$.

Given $\epsilon > 0$, let H_i^T denote the set of private histories of signals of length T for players $i = 1, 2$, and let $H_i^{>T}$ denote the subset of H_i^T consisting of those histories in which the observed number of signals ω_i exceeds the expectation of this number, $T\epsilon$. Observe that, by the central limit theorem, the probability that a history of length T is in $H_i^{>T}$ tends to $1/2$ as $T \rightarrow \infty$.

Define σ_i^T as the strategy consisting in playing C for the first T periods and in playing C forever after, if the private history is in $H_i^{>T}$, and B if it is not. The payoff to players 1 and 2 from using (σ_1^T, σ_2^T) when player 3 plays a best-response tends to $1/2$ as $\delta \rightarrow 1$, $T \rightarrow \infty$ provided $\delta^T \rightarrow 1$.

This shows that, for any value of $\epsilon > 0$, when $\delta \rightarrow 1$, equilibrium payoffs exist that approach $1/2$ for players 1 and 2. On the other hand, for any fixed δ , players 1 and 2 cannot secure more than $1/4$ as $\epsilon \rightarrow 0$.

As Example 4 shows, the order of limits is important in general. While the set of payoffs is continuous in the monitoring structure for a fixed discount factor, the limit of this set as the discount factor tends to one may be discontinuous in the monitoring structure. Our next result, Theorem 4, shows that such cases are ruled out when player i ’s signals allow to statistically discriminate among action profiles of the other players.

Definition 5 (Identifiability). The monitoring structure q satisfies identifiability for player i if, for all a_{-i} in A_{-i} and α_i in ΔA_i ,

$$q_i^{a_{-i}, \alpha_i} \notin \text{co}\{q_i^{a'_{-i}, \alpha_i} : a'_{-i} \neq a_{-i}, a'_{-i} \in A_{-i}\}.$$

That is, q satisfies identifiability if, for any possibly mixed action of player i , the distribution over his signals that is generated by any pure action profile of his opponents cannot be replicated by some convex combination of other action profiles of theirs. Let d denote the total variation distance between probability measures. Given $\rho > 0$, the monitoring structure q satisfies ρ -identifiability for player i if, for all a_{-i} in A_{-i} and α_i in ΔA_i , and any distribution q'_i in $\text{co}\{q_i^{a'_{-i}, \alpha_i} : a'_{-i} \neq a_{-i}, a'_{-i} \in A_{-i}\}$,

$$d(q_i^{a_{-i}, \alpha_i}, q'_i) > \rho.$$

Thus, the concept of ρ -identifiability captures the distance between the monitoring structure q and the nearest one that fails to satisfy identifiability.

Finally, we need to introduce a measure of the distance between a monitoring structure and the nearest one that satisfies conditional independence for player i . For $\varepsilon > 0$, the monitoring structure q is ε -dependent for player i if, for all action profiles α in $\prod_j \Delta A_j$, there exists a family of product distributions $(q'_{-i}(s^i))_{s_i}$ in $\prod_{j \neq i} \Delta S_j$ such that

$$\mathbf{E}[d(q_{-i}^\alpha(\cdot | s_i), q'_{-i}(s_i))] < \varepsilon.$$

That is, q is ε -dependent for player i if those signals for which the conditional distribution of player i is not close to a product distribution are sufficiently unlikely, given any action profile that corresponds to independent randomizations.

In the sequel, when there is no ambiguity as to which player is considered, we drop the reference to player i when using the expressions “ ρ -identifiability” or “ ε -dependence.”

Theorem 4. For any $\nu > 0$, if q satisfies ρ -identifiability, for some $\rho > 0$, there exists $\varepsilon > 0$ such that, if q is ε -dependent, then, for all $\delta \in [0, 1)$, player i 's individually rational payoff is within ν of his min max payoff.

Theorem 4 strengthens Theorem 2 and provides a continuity result that is uniform in the discount factor. This theorem is important for the literature on the robustness of equilibria for private monitoring structures that are in a neighborhood of perfect, or imperfect public monitoring. Indeed, while almost-perfect monitoring structures need not be ε -dependent for small ε (as expected given Example 1), it is immediate to see that they must be if attention is restricted to canonical signal structures. Therefore, Theorem 4 provides a converse to Theorem 1 in [14].

Theorem 4 extends to distributions for which there exists a garbling that induces an approximate version of independence, provided that the garbled signal satisfies the identifiability condition (that is, the belief of player i , conditional on his garbled signal, should satisfy ρ -identifiability). The generalization is straightforward and omitted.

Note also that the identifiability condition used in Theorem 4 can be weakened. Indeed, we do not need that each action of player i allows for statistical discrimination of his opponents' actions. Rather, it is enough that, for each $\alpha_{-i} \neq \alpha'_{-i}$, there exists one action of player i that discriminates between them. To prove this result, it is enough to consider strategies of player i that play each action with probability at least $\epsilon > 0$ at each stage, then let $\epsilon \rightarrow 0$.

Finally, a large literature has considered a restricted class of strategies, namely public strategies, in the context of games with public monitoring. In such games, the min max payoff in public strategies in the repeated game cannot be lower than the static min max payoff, a result which is not generally true without the restriction on strategies (see again [9, Exercise 5.10]). It is then natural to ask to what extent the ε -dependence assumption can be weakened for such a class of strategies. To be specific, assume that strategies must be a function of the history of private signals ω_i alone, rather than of the history of all signals s_i . Observe that this reduces to public strategies in the case of public monitoring, but is well-defined even under private monitoring. Then Theorem 4 remains valid, if we require that only the restriction of the monitoring structure to private signals be ε -dependent. This is a significantly weaker restriction, which is indeed trivially satisfied if monitoring is public. The proof is a trivial modification of the proof of Theorem 4.

All statements are either trivial or follow from the proof of Theorem 4. The proof of this theorem is rather delicate can be found in Appendix A.

The first part of the proof, presented in Section A.1, reduces the study to a repeated game with an alternate monitoring structure in which the signals to players $-i$ is public among these players (and only among them), and these players are restricted to public strategies (depending only on past realizations of these public signals). When studying the repeated game with the alternate monitoring structure, tools from information theory are brought to bear. This is done in Section A.2. There we show that, under ε -epsilon dependence and ρ -identifiability, it takes time to accumulate sufficient public information for player i 's opponents to successfully correlate their action profile, relative to the time it takes player i to detect which of the plays his opponents have coordinated upon.

5. Concluding comments

In this paper, we provide the necessary and sufficient condition on the information structure for which the lowest equilibrium payoff in any Bayesian game associated with this information structure is no lower than the min max level determined by the payoff matrix only. This provides a sufficient condition under which the stage-game min max payoff is the appropriate lower bound on possible equilibrium payoffs in a repeated game, whether the monitoring is imperfect or not. We also show under which conditions this remains approximately true when the monitoring structure is arbitrarily close to one that satisfies this condition. This provides a condition under which a converse to the folk theorems which can be found in [7,14] hold.

An important question left open is how to actually determine the repeated-game min max payoff when it is below the stage-game min max payoff. Characterizations are only known for some classes of monitoring structures in [16,11,12].

When conditionally on each player's signal, other player's signals are independent, the equilibrium payoff set possesses a natural recursive structure, and methods from dynamic programming can be brought to bear. With three players or more, the paper [10] characterizes the information structures for which conditional on each player's signal, other player's signals are independent. Those are the information structures such that all player's signals are independent conditional on an underlying common-knowledge variable.

The more general question of the characterization of monitoring structures which admit conditional independence garblings for every player, case in which each player's individually rational payoff is given by his min max payoff in the stage game, is left for future research.

Appendix A. Proof of Theorem 4

The proof of Theorem 4 is divided in two parts. In the first part, we replace the private monitoring structure by another one. Player i 's information is unchanged. His opponents' information is now public among them, but it is not simply the information resulting from pooling their individual signals from the original monitoring structure: doing so would enable them to correlate their play in many situations in which they would not be able to do so if their strategy were only based on their own signals. The common information must be "poorer" than that, but we still need to make sure that any probability distribution over plays that could be generated in the original monitoring structure by some strategy profile of player i 's opponents (for some strategy of player i) can still be generated in this alternate monitoring structure. We shall refer to players $-i$ as the team, and it will be understood that their objective is to minimize player i 's payoff.

A.1. Reduction to public strategies

A result that will prove useful in the sequel is the following.

Lemma 2. *If q is a distribution over some product finite set $S := \prod_{k \in K} S_k$, then there exists a product distribution $p \in \prod_j \Delta S_j$ and a "residual" distribution r such that*

$$q = \lambda p + (1 - \lambda)r,$$

for some $\lambda = \lambda(q)$ in $[0, 1]$. Further, for every $v > 0$, there exists $\varepsilon > 0$ such that, if $d(q, q') < \varepsilon$ for some $q' \in \prod_j \Delta S_j$, then we can choose $\lambda > 1 - v$.

Proof. Indeed, if we define λ as the supremum over all such values in the unit interval for which we can write q as a convex combination of distributions p and r , with $p \in \prod_j \Delta S_j$, it follows from the maximum theorem that (i) this maximum is achieved, (ii) it is continuous in q . In fact, since q belongs to a compact metric space, this continuity is uniform, by the Heine–Cantor theorem. The result follows, since $\lambda = 1$ if $q \in \prod_j \Delta S_j$. \square

Given this result, we can view signals in the repeated game as being drawn in three stages. Given the action profile (α_{-i}, a_i) :

- first, the signal s_i is drawn according to the marginal distribution $q_i^{\alpha_{-i}, a_i}$. Given s_i , apply (2) and write

$$q_{-i}^{\alpha_{-i}, a_i}(\cdot | s_i) = \lambda p_{-i}^{\alpha_{-i}, a_i}(\cdot | s_i) + (1 - \lambda)r_{-i}^{\alpha_{-i}, a_i}(\cdot | s_i),$$

where λ depends on $q_{-i}^{\alpha_{-i}, a_i}(\cdot | s_i)$;

- second, a Bernoulli random variable l with $\mathbf{P}\{l = 1\} = 1 - \mathbf{P}\{l = 0\} = \lambda$ is drawn;
- third, if $l = 0$, the signal profile s_{-i} is drawn according to $r_{-i}^{\alpha_{-i}, a_i}(\cdot | s_i)$; if instead $l = 1$, s_{-i} is drawn according to $p_{-i}^{\alpha_{-i}, a_i}(\cdot | s_i)$.

We now use this representation to show that player i 's individually rational payoff is no larger under the original monitoring structure than under an alternate monitoring structure in which player i 's opponents can condition their strategy on the history of values of s_i , l , and, whenever $l = 0$, of s_{-i} . This is non-trivial because player i 's opponents are not allowed to condition their

strategy on their own signals any longer, unless $\lambda = 0$. Yet the conclusion is rather intuitive, for when $\lambda = 1$, the signals of player i 's opponents are independently distributed anyway (conditional on s_i). This result will allow us in the next subsection to view the histories used by player i 's opponents as common. Before stating the result, further notation must be introduced.

Histories. Histories Recall that a *history* of length t in the original game is an element of S^t . The set of plays is $H^\infty = S^\mathbb{N}$ endowed with the product σ -algebra. We define an *extended history* of length t as an element of $(S \times \{0, 1\} \times S')^t$, that is, as a history in the original game augmented by the history of realizations of the Bernoulli variable l . The set of extended plays is $\tilde{H}^\infty = (S \times \{0, 1\})^\mathbb{N}$ endowed with the product σ -algebra.

A private history of length t for player j (in the original game) is an element of $H_j^t = S_j^t$. A *public history* of length t is an element of $H_p^t := S_p^t$, where $S_p := S_i \times \{0\} \times S_{-i} \cup S_i \times \{1\}$; that is, S_p is the set of public signals s_p , where $s_p = (s_i, 0, s_{-i})$ if $l = 0$ and $s_p = (s_i, 1)$ if $l = 1$.

Strategies. A (behavioral) private strategy for player j (in the original game) is a family $\sigma_j = (\sigma_j^t)_t$, where $\sigma_j^t: H_j^{t-1} \rightarrow \Delta A_j$. Let Σ_j denote the set of these strategies. A (behavioral) public strategy for player j is a family $\tau_j = (\tau_j^t)_t$, where $\tau_j^t: H_p^{t-1} \rightarrow \Delta A_j$. Let $\Sigma_{p,j}$ denote the set of these strategies. Finally, a (behavioral) general strategy for player j is a family $\tilde{\sigma}_j = (\tilde{\sigma}_j^t)_t$, where $\tilde{\sigma}_j^t: (S_p \times S_j)^{t-1} \rightarrow \Delta A_j$. Let $\tilde{\Sigma}_j$ denote the set of these strategies.

Note that both $\Sigma_{p,j}$ and Σ_j can naturally be identified as subsets of $\tilde{\Sigma}_j$, but $\Sigma_{p,j}$ and Σ_j cannot be ordered by set inclusion.

A (pure) strategy for player i is a family $\sigma_i = (\sigma_i^t)_t$, where $\sigma_i^t: S_i^{t-1} \rightarrow A_i$.

Any profile of general strategies σ_{-i} for player i 's opponents, together with a strategy σ_i for player i , induces a probability distribution $\mathbf{P}_{\sigma_{-i}, \sigma_i}$ on \tilde{H}^∞ .

Proposition 1. For any private strategy profile σ_{-i} , there exists a public strategy profile τ_{-i} such that, for every pure strategy σ_i , $h_i^t \in S_i^t$, $s^{t+1} \in S$,

$$\mathbf{P}_{\tau_{-i}, \sigma_i}(h_i^t) = \mathbf{P}_{\sigma_{-i}, \sigma_i}(h_i^t), \tag{A.1}$$

$$\mathbf{P}_{\tau_{-i}, \sigma_i}(s^{t+1} | h_i^t) = \mathbf{P}_{\sigma_{-i}, \sigma_i}(s^{t+1} | h_i^t) \quad \text{if } \mathbf{P}_{\sigma_{-i}, \sigma_i}(h_i^t) > 0. \tag{A.2}$$

Proof. We first define a *public strategy up to stage t* for player j as a family $\tau_{t,j} = (\tau_{t,j}^{t'})_{t'}$ where

$$\begin{cases} \tau_{t,j}^{t'}: S_p^{t'-1} \rightarrow \Delta A_j, & \text{if } t' \leq t; \\ \tau_{t,j}^{t'}: S_p^{t'-1} \times S_j^{t'-t+1} \rightarrow \Delta A_j, & \text{otherwise.} \end{cases}$$

The proof of the proposition relies on the following lemma. This lemma exhibits a sequence of strategy profiles for player i 's opponents, up to stage t , based on σ_{-i} , that do only depend on the first t public signals, and not on the realizations of the first t private signals (conditional on these public signals). This sequence of strategies is constructed by iterated applications of Kuhn's theorem, as shown in the next lemma.

Lemma 3. For any private strategies σ_{-i} , there exist strategies $(\tau_{t,-i})_t = (\tau_{t,j})_{j \neq i, t}$ where $\tau_{t,j}$ is a public strategy up to stage t for player j and

$$\tau_{0,j} = \sigma_j, \tag{A.3}$$

$$\tau_{t',j}^t = \tau_{t',j}^{t'} \text{ for } t' \geq t, \tag{A.4}$$

$$\mathbf{P}_{\tau_{t+1,-i},\sigma_i}(s^{t+1}, \dots, s^{t+n} | h_p^t) = \mathbf{P}_{\tau_{t,-i},\sigma_i}(s^{t+1}, \dots, s^{t+n} | h_p^t) \tag{A.5}$$

for every $\sigma_i, n, (s^{t+1}, \dots, s^{t+n}) \in S^n$, and $h_p^t \in H_p^t$.

Proof. Define $\tau_{t,-i}$ by induction on t . First let $\tau_{0,-i} = \sigma_{-i}$ so that (A.3) is met. Assume $\tau_{t,-i}$ has been defined. Let $\tau_{t'+1,-i}^t = \tau_{t',-i}^{t'}$ for $t' \leq t$ so that (A.4) is satisfied.

For each history h_p^t and for every $s_j \in S_j$, let $\tilde{\tau}_{t+1,j}[h_p^t, s_j]$ be the private (continuation) strategy defined by $\tilde{\tau}_{t+1,j}[h_p^t, s_j](s_j^1, \dots, s_j^n) = \tau_{t,j}(h_p^t, s_j, s_j^1, \dots, s_j^n)$.

The probability $q^{\tau_{t,-i}(h_p^{t-1}), a_i^t}(s_j | s_p)$ over S_j defines a mixture of private strategies $\tilde{\tau}_{t+1,j}^t[h_p^t, s_j]$, where a_i^t is player i 's action in period t as specified by s_p . By Kuhn's theorem, there exists a private strategy $\tau_{t+1,j}[h_p^t]$ which is equivalent to this mixture. Finally set $\tau_{t+1,j}^{t+n+1}(h_p^t, s_j^1, \dots, s_j^n) = \tau_{t+1,j}^{t+n+1}[h_p^t](s_j^1, \dots, s_j^n)$. Condition (A.5) is met by equivalence of the mixed and the behavioral strategy and because all $(s_j)_j$ are independent conditional on s_p and h_p^t . \square

Back to the proof of Proposition 1, define τ_{-i} by $\tau_j^t(h_p^{t-1}) = \tau_{t,j}^t(h_p^{t-1})$, where $(\tau_{t,-i})_t$ is given by Lemma 3.

From (A.4), for every t' , $\mathbf{P}_{\tau_{t'+1,-i},\sigma_i}$ and $\mathbf{P}_{\tau_{t',-i},\sigma_i}$ induce the same probability on $H_p^{t'}$, and from (A.5), $\mathbf{P}_{\tau_{t+1,-i},\sigma_i}(h_i^t | h_p^t) = \mathbf{P}_{\tau_{t,-i},\sigma_i}(h_i^t | h_p^t)$ for $t' \geq t$. Thus $\mathbf{P}_{\tau_{-i},\sigma_i}(h_i^t) = \mathbf{P}_{\tau_{t,-i},\sigma_i}(h_i^t) = \mathbf{P}_{\tau_{t-1,-i},\sigma_i}(h_i^t) = \dots = \mathbf{P}_{\sigma_{-i},\sigma_i}(h_i^t)$, which gives (A.1).

Also, $\mathbf{P}_{\tau_{-i},\sigma_i}(s^{t+1} | h_i^t) = \mathbf{P}_{\tau_{t+1,-i},\sigma_i}(s^{t+1} | h_i^t) = \frac{\mathbf{P}_{\tau_{t+1,-i},\sigma_i}(s^{t+1}, h_i^t)}{\mathbf{P}_{\tau_{t+1,-i},\sigma_i}(h_i^t)} = \frac{\mathbf{P}_{\sigma_{-i},\sigma_i}(s^{t+1}, h_i^t)}{\mathbf{P}_{\sigma_{-i},\sigma_i}(h_i^t)} = \mathbf{P}_{\sigma_{-i},\sigma_i}(s^{t+1} | h_i^t)$ whenever $\mathbf{P}_{\sigma_{-i},\sigma_i}(h_i^t) > 0$, which gives (A.2). \square

An immediate consequence of Proposition 1 is that the individually rational payoff of player i is necessarily no higher under the alternate monitoring structure in which player i 's opponents use so-called public strategies, as under the original monitoring structure (whether we consider the finitely or infinitely-repeated game, and independently of the discount factor). Note that this already establishes Theorem 2 (and therefore Corollary 3). Indeed, under the assumption of Theorem 2, we have $\mathbf{P}\{l = 1\} = 1$, so that public strategies only depend on s_i , which is known by player i . That is, conditional on his history of private signals, player i can view the choices of continuation strategies of his opponents as independent. Matters are more complicated when the monitoring structure is only ε -dependent, as $\mathbf{P}\{l = 1\} < 1$. Nevertheless, the event $\{l = 0\}$ is unfrequent for ε small enough.

A.2. Measuring secret correlation and its dissipation

In this subsection, we prove Theorem 4, using the public strategies introduced in the previous subsection. The main idea of the proof is that, under the conditions of the theorem, little secret correlation can be generated by team members in the course of the repeated game, and if this correlation is used, there is enough dissipation of this correlation over time. This implies that the individually rational payoff of player i is uniformly close to his min max payoff.

In order to measure the amount of secret correlation generated and dissipated by the team in the course of the repeated game, we rely on the entropy measure of randomness and information. We first derive a bound on the min max payoff in terms of entropy in Section A.2.1. Next, we use this bound to complete the proof of the main result in Section A.2.2.

A.2.1. An entropy bound on min max payoffs

Let $\sigma = (\sigma_{-i}, \sigma_i)$ be a strategy profile, where σ_{-i} is a profile of public strategies. Suppose that after stage t , the history for player i is $h_i^t = (s_i^1, \dots, s_i^t)$. Let $h_p^t = (s_p^1, \dots, s_p^t)$ be the public history after stage t . The mixed action profile played by the team at stage $t + 1$ is $\sigma_{-i}^{t+1}(h_p^t) = (\sigma_j(h_p^t))_{j \neq i}$. Player i holds a *belief* about this mixed action, namely he believes that players $-i$ play $\sigma_{-i}^{t+1}(h_p^t)$ with probability $\mathbf{P}_\sigma(h_p^t | h_i^t)$. The distribution of the action profile \mathbf{a}_{-i}^{t+1} of players $-i$ at stage $t + 1$ given the information h_i^t of player i is $\sum_{h_p^t} \mathbf{P}_\sigma(h_p^t | h_i^t) \sigma_{-i}^{t+1}(h_p^t)$, an element of the set ΔA_{-i} of correlated distributions on A_{-i} . Let $X := \prod_{j \neq i} \Delta A_j$ be the set of independent probability distributions on A_{-i} . A *correlation system* is a probability distribution on X and we define $C := \Delta X$ as the set of correlation systems. We identify X with a closed subset of ΔA_{-i} and so C is compact with respect to the weak* topology. Given a correlation system c and $a_i \in A_i$, let (\mathbf{x}, s_p) be a random variable with values in $X \times S_p$ such that the law of \mathbf{x} is c and the law of (s_p, s_i) given $\{\mathbf{x} = x\}$ is $q^{x, a_i}(\cdot)$. We let H denote the entropy function (see, e.g. [5]), and define the entropy variation ΔH as

$$\Delta H(c, a_i) = H(s_p | \mathbf{x}) - H(s_i).$$

The entropy variation is expressed as the difference of two terms. The first term, which we see as an entropy gain, is the conditional uncertainty contained in s_p given \mathbf{x} . The second term, which we interpret as an entropy loss, is the entropy of s_i observed by player i . If \mathbf{x} is finite, from the additivity formula

$$H(\mathbf{x}, s_p) = H(\mathbf{x}) + H(s_p | \mathbf{x}) = H(s_i) + H(\mathbf{x}, s_p | s_i),$$

we have that

$$\Delta H(c, a_i) = H(\mathbf{x}, s_p | s_i) - H(\mathbf{x}).$$

Given a correlation system c , the distribution of the action profile for the team is $x_c \in \Delta A_{-i}$ such that for every $a_{-i} \in A_{-i}$, $x_c(a_{-i}) = \int_X x_{-i}(a_{-i}) dc(x)$. Player i 's best-reply payoff against c is $\pi(c) := \max_{a_i \in A_i} g_i(x_c, a_i)$, and we define $B_i(c) := \operatorname{argmax} g_i(x_c, \cdot) \subseteq A_i$.

Consider the set of feasible vectors $(\Delta H(c, a_i), \pi(c))$ where $a_i \in B_i(c)$ in the (entropy variation, payoff)-plane:

$$V = \{(\Delta H(c, a_i), \pi(c)) \mid c \in C, a_i \in B_i(c)\}.$$

Define w as the lowest payoff given a convex combination of correlation systems under the constraint that the average entropy variation is non-negative:

$$w = \inf \{x_2 \in \mathbb{R} \mid (x_1, x_2) \in \operatorname{co} V, x_1 \geq 0\}.$$

For every correlation system c such that \mathbf{x} is a.s. constant, $\Delta H(c) \geq 0$, and so V intersects the half-plane $\{x_1 \geq 0\}$. It is easy to show that V is compact, so the infimum is achieved.

We are now ready to state the main result of this section.

Proposition 2. For every profile of public strategies σ_{-i} , there exists σ_i such that the induced payoff to player i in every δ -discounted game is no less than w .

Proof. Given σ_{-i} , we inductively define $\sigma_{i,\sigma_{-i}}$ as the strategy of player i that plays stage-best replies to σ_{-i} : At stage 1, $\sigma_{i,\sigma_{-i}} \in \operatorname{argmax}_{a_i} g_i(\sigma_{-i}(\emptyset), a_i)$ where \emptyset is the null history that starts the game. Assume that $\sigma_{i,\sigma_{-i}}$ is defined on histories of length less than t . For every history h_i^t of player i , let $x^{t+1}(h_i^t) \in \Delta A_{-i}$ be the distribution of the action profile of the team at stage $t + 1$ given h_i^t (arbitrary if h_i^t has zero probability) and select $\sigma_{i,\sigma_{-i}}(h_i^t)$ in $\operatorname{argmax}_{a_i} g_i(x^{t+1}(h_i^t), a_i)$. The main step requires the following:

Lemma 4. For every σ_{-i} , $\sigma_{i,\sigma_{-i}}$ defends w in every n -stage game, i.e. for every σ_{-i} , n ,

$$\mathbf{E}_{\sigma_{-i}, \sigma_{i,\sigma_{-i}}} \frac{1}{n} \sum_{t=1}^n g_i(a) \geq w.$$

Proof. The proof follows lines similar to those used by previous papers using entropy methods (see e.g., [21,22,13,11]). Fix a profile of public strategies σ_{-i} for the team and let $\sigma_i = \sigma_{i,\sigma_{-i}}$. Let s_p^t, s_i^t be the random signals to players $-i$ and to player i under $\mathbf{P}_{\sigma_{-i}, \sigma_i}$, $\mathbf{h}_p^t = (s_p^1, \dots, s_p^t)$ and $\mathbf{h}_i^t = (s_i^1, \dots, s_i^t)$ be the public history and the history of player i after stage t . Let $\mathbf{x}^t = \sigma_{-i}^t(\mathbf{h}_i^{t-1})$ and $c^t(\mathbf{h}_i^{t-1})$ be the distribution of \mathbf{x}_m^t conditional on \mathbf{h}_i^{t-1} i.e. $c^t(\mathbf{h}_i^{t-1})$ is the correlation system at stage t after history \mathbf{h}_i^{t-1} . Under $\sigma = (\sigma_{-i}, \sigma_i)$, the expected payoff to player i at stage t given \mathbf{h}_i^t is $\max_{a_i} g_i(\mathbf{E}_{\sigma}[\mathbf{x}^t | \mathbf{h}_i^{t-1}], a_i) = \pi(\mathbf{c}^t)$, from the definition of σ_i . Therefore, the average payoff to player i in the n -stage game is $\mathbf{E}_{\sigma}[\frac{1}{n} \sum_{m=1}^n \pi(\mathbf{c}_m)]$. From the additivity of entropies,

$$\begin{aligned} H(s_p^1, \dots, s_p^t | \mathbf{h}_i^t) &= H(s_i^t | \mathbf{h}_i^{t-1}) + H^t \\ &= H^{t-1} + H(s_p^t | \mathbf{h}_p^{t-1}). \end{aligned}$$

Thus,

$$\begin{aligned} H^t - H^{t-1} &= H(s_p^t | \mathbf{h}_p^{t-1}) - H(s_i^t | \mathbf{h}_i^{t-1}) \\ &= H(s_p^t | \mathbf{x}^t, \mathbf{h}_i^{t-1}) - H(s_i^t | \mathbf{h}_i^t) \\ &= \mathbf{E}_{\sigma} \Delta H(c^t(\mathbf{h}_i^{t-1}), \mathbf{a}_i^t). \end{aligned}$$

Then

$$\sum_{t=1}^n \mathbf{E}_{\sigma} \Delta H(c^t(\mathbf{h}_i^{t-1}), \mathbf{a}_i^t) = H_n \geq 0.$$

Therefore the vector $(\frac{1}{t} \sum_{t=1}^n \mathbf{E}_{\sigma} \Delta H(c^t(\mathbf{h}_i^{t-1}), \mathbf{a}_i^t), \mathbf{E}_{\sigma} \frac{1}{n} \sum_{t=1}^n g_i(a))$ is in $\operatorname{co} V \cap \{x_1 \geq 0\}$. \square

Now we complete the proof of Proposition 2. Recall that the discounted payoff is a convex combination of the average payoffs (see e.g. [17]). It then follows from Lemma 4 that the δ -discounted payoff to player i induced by $\sigma_{-i}, \sigma_{i,\sigma_{-i}}$ is no less than w , hence the result. \square

A.2.2. The case of ε -dependence and ρ -identifiability

Here we prove that under ε -dependence, little secret correlation can be generated per stage by players $-i$. We also show that, under the ρ -identifiability assumption, secret correlation, if utilized, dissipates. Relying on Proposition 2, we then conclude the proof of the main theorem.

It is useful in what follows to relate the number w to the boundary of $\text{co } V$. Define, for each real number h ,

$$u(h) := \inf\{\pi(c) \mid c \in C, \Delta H(c) \geq h\} = \inf\{x_2 \mid (x_1, x_2) \in V, x_1 \geq h\}.$$

Since V is compact, $u(h)$ is well defined. Let $\text{cav } u$ be the least concave function pointwise greater than u . Then

$$w = \text{cav } u(0).$$

Indeed, u is upper-semi-continuous, non-increasing and the hypograph of u is the comprehensive set $V^* = V - \mathbb{R}_+^2$ associated to V . This implies that $\text{cav } u$ is also non-increasing, u.s.c. and its hypograph is $\text{co } V^*$.

We compare w with the maximum payoff (or the minimum of player i 's payoff) that players $-i$ can obtain in a modified game where $s_p = s_i$, but they could use an entropy of ε_h at each stage. If $s_p = s_i$, $\Delta H(c, a_i) = -I(\mathbf{x}, s_i)$, and the function u' playing the role of u is

$$u'(h) = \min\{\pi(c) \mid (c, a_i) \text{ s.t. } I(\mathbf{x}; s_i) \leq h\}.$$

Lemma 5. For every $\varepsilon_h > 0$, there exists $\varepsilon > 0$ such that if the monitoring is ε -dependent, for every $(c, a_i) : \Delta H(c, a_i) \leq \varepsilon_h - I(\mathbf{x}; s_i)$. In particular, $u(h) \geq u'(h + \varepsilon_h)$ for all h .

Proof. First note that

$$\Delta H(c, a_i) + I(\mathbf{x}; s_i) = H(\mathbf{s}_p | \mathbf{x}, s_i) = \int_x \sum_{s_i} q_i^{x, a_i}(s_i) H(q^{x, a_i}(\mathbf{s}_p | s_i)) dc(x).$$

For every x, s_i ,

$$H(q^{x, a_i}(\mathbf{s}_p | s_i)) = H(\lambda(q_i^{x, a_i}(s_i))) + (1 - \lambda(q_i^{x, a_i}(\cdot | s_i)))H(r_{-i}(\cdot | s_i)).$$

There exists $S'_i \subseteq S_i$ such that $q_i^{x, a_i}(S'_i) \geq 1 - \varepsilon$ and that for $s_i \in S'_i$, $d(q_i^{x, a_i}(\cdot | s_i), \prod_{j \neq i} \Delta S_j) \leq \varepsilon$, and in particular $\lambda(q_i^{x, a_i}(\cdot | s_i)) \geq \eta(\varepsilon)$ by Lemma 2. For $s_i \in S'_i$, $H(\lambda(q_i^{x, a_i}(s_i))) \leq \max\{H(\eta(\varepsilon)), 1\}$ and for all $s_i \in S'_i$, $H(r_{-i}(\cdot | s_i)) \leq \log_2(S_{-i} \times A_{-i})$. Hence

$$\Delta H(c, a_i) + I(\mathbf{x}; s_i) \leq \max\{\eta(\varepsilon), 1\} + \log_2(S_{-i} \times A_{-i}),$$

which proves the first part of the lemma. The second part now follows from the definitions of u and u' . \square

Let us finally define the function u'' that corresponds to u' when $s_i = a_{-i}$, namely when i perfect monitors the actions of his opponents. Formally,

$$u''(h) := \min\{\pi(c), (c, a_i) \text{ s.t. } I(\mathbf{x}; \mathbf{a}_{-i}) \leq h\}.$$

Lemma 6. There exists a continuous function α such that $\alpha(0) = 0$ and $u'(h) \geq u''(\alpha(h))$ for all h .

Proof. Let $\alpha(h) = \max\{I(\mathbf{x}; \mathbf{a}_{-i}) \mid (c, a_i) \text{ s.t. } I(x; s_i) \leq h\}$. It follows from Carathéodory's theorem that we can restrict attention to c with support of size at most 3 in the definition of α , and the sup is actually a max. Assume now c has finite support, $I(\mathbf{x}; s_i) = 0$ implies that s_i and \mathbf{m}_{-i} are independent, and therefore that c is a mass unit and that $I(\mathbf{x}; \mathbf{a}_{-i}) = 0$. Hence $\alpha(0) = 0$. The map α is continuous by the maximum principle. That $u'(h) \geq u''(\alpha(h))$ for all h follows from the definitions of u' , u'' and α . \square

Now we complete the proof of our main result.

Proof of Theorem 4. From Proposition 2 we have $v_i^\delta \geq w$, where by Lemmata 5 and 6, $w \geq \text{cav}(u'' \circ \alpha)(\varepsilon_h)$. The result follows since $\text{cav}(u'' \circ \alpha)$ is continuous and $\varepsilon_h \rightarrow 0$ as $\varepsilon \rightarrow 0$. \square

References

- [1] R.J. Aumann, L.S. Shapley, Long-term competition – A game theoretic analysis, preprint, 1976.
- [2] V. Bhaskar, I. Obara, Belief-based equilibria in the repeated prisoners' dilemma with private monitoring, *J. Econ. Theory* 102 (2002) 40–70.
- [3] D. Blackwell, Comparison of experiments, in: *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press, 1951, pp. 93–102.
- [4] P. Cartier, J.M.G. Fell, P.-A. Meyer, Comparaison des mesures portées par un ensemble convexe compact, *Bull. Soc. Math. France* 92 (1964) 435–445.
- [5] T.M. Cover, J.A. Thomas, *Elements of Information Theory*, Wiley Series in Telecommunications, Wiley, 1991.
- [6] J.C. Ely, J. Välimäki, A robust folk theorem for the prisoner's dilemma, *J. Econ. Theory* 102 (2002) 84–106.
- [7] D. Fudenberg, D.K. Levine, E. Maskin, The folk theorem with imperfect public information, *Econometrica* 62 (5) (1994) 997–1039.
- [8] D. Fudenberg, E. Maskin, The folk theorem in repeated games with discounting or with incomplete information, *Econometrica* 54 (3) (1986) 533–554.
- [9] D. Fudenberg, J. Tirole, *Game Theory*, The MIT Press, Cambridge, MA, 1991.
- [10] O. Gossner, E. Kalai, R.J. Weber, Information independence and common knowledge, *Econometrica* 77 (2009) 1317–1328.
- [11] O. Gossner, T. Tomala, Empirical distributions of beliefs under imperfect observation, *Math. Oper. Res.* 31 (1) (2006) 13–30.
- [12] O. Gossner, T. Tomala, Secret correlation in repeated games with signals, *Math. Oper. Res.* 32 (2007) 413–424.
- [13] O. Gossner, N. Vieille, How to play with a biased coin? *Games Econ. Behav.* 41 (2002) 206–226.
- [14] J. Hörner, W. Olszewski, The folk theorem with private almost-perfect monitoring, *Econometrica* 74 (6) (2006) 1499–1544.
- [15] J.S. Jordan, Three problems in learning mixed strategy equilibria, *Games Econ. Behav.* 5 (1993) 368–386.
- [16] E. Lehrer, Nash equilibria of n -player repeated games with semi-standard information, *Int. J. Game Theory* 19 (1990) 191–217.
- [17] E. Lehrer, S. Sorin, A uniform Tauberian theorem in dynamic programming, *Math. Oper. Res.* 17 (1992) 303–307.
- [18] G. Mailath, S. Morris, Repeated games with almost-public monitoring, *J. Econ. Theory* 102 (2002) 189–229.
- [19] G. Mailath, S. Morris, Coordination failure in repeated games with almost public monitoring, *Theoretical Econ.* 1 (2006) 311–340.
- [20] D. Moreno, J. Wooders, An experimental study of communication and coordination in noncooperative games, *Games Econ. Behav.* 24 (1998) 47–76.
- [21] A. Neyman, D. Okada, Strategic entropy and complexity in repeated games, *Games Econ. Behav.* 29 (1999) 191–223.
- [22] A. Neyman, D. Okada, Repeated games with bounded entropy, *Games Econ. Behav.* 30 (2000) 228–247.
- [23] M. Piccione, The repeated prisoner's dilemma with imperfect private monitoring, *J. Econ. Theory* 102 (2002) 70–84.
- [24] A. Rubinstein, *Equilibrium in supergames*, Center for Research in Mathematical Economics and Game Theory, Research Memorandum 25, 1977.
- [25] T. Sekiguchi, Efficiency in repeated prisoner's dilemma with private monitoring, *J. Econ. Theory* 76 (1997) 345–361.
- [26] V. Strassen, The existence of probability measures with given marginals, *Ann. Math. Statist.* 36 (1965) 423–439.
- [27] B. von Stengel, D. Koller, Team max min equilibria, *Games Econ. Behav.* 21 (1997) 309–321.