

A Folk Theorem for Minority Games*

Jérôme Renault
CEREMADE
Université Paris Dauphine
Pl. du Marechal de Lattre de Tassigny
F-75775 Paris cedex 16, France
renault@ceremade.dauphine.fr

Sergio Scarlatti
Dipartimento di Scienze
Università D'Annunzio
Viale Pindaro 42
I-65127 Pescara, Italy
scarlatt@sci.unich.it

Marco Scarsini
Dipartimento di Statistica e Matematica Applicata
Università di Torino
Piazza Arbarello 8
I-10122 Torino, Italy
scarsini@econ.unito.it

January 2003

Abstract

We study a particular case of repeated games with public signals. In the stage game an odd number of players have to choose simultaneously one of two rooms. The players who choose the less crowded room receive a reward of one euro (whence the name “minority game”). Between the stages, only the current majority room is publicly announced. We show that in the infinitely repeated game any feasible payoff can be achieved as a uniform equilibrium payoff, and as an almost sure equilibrium payoff. In particular we construct an inefficient equilibrium where, with probability one, all players choose the same room at almost all stages. This equilibrium is sustained by punishment phases which use, in a unusual way, the pure actions that were played before start of the punishment.

AMS 2000 Subject Classification: 91A20, 91A28.

Keywords: Repeated games, imperfect monitoring, public signals.

*The work of Sergio Scarlatti and Marco Scarsini was partially supported by MIUR-COFIN

1 Introduction

An odd number of players have to choose simultaneously one of two rooms. The players who choose the less crowded room receive a reward of one euro. The others receive nothing. The game is repeated over time. A version of this game was introduced by Arthur (1994) under the name *El Farol's Bar problem* (see also Arthur (1999)). In his paper customers have to decide every weekend whether to go to the bar or stay home. Only customers who make the minority choice are happy. Arthur's paper gave rise to a huge literature on so called *minority games*. The interest in this class of games came especially from theoretical physicists working in statistical mechanics (see e.g. Challet and Zhang (1997), Savit et al. (1999)). They focus on the case of many players and see "these problems as novel examples of frustrated and disordered many-body systems" (Cavagna et al. (1999)). In their models the many agents have limited memory and act according to some evolutionary paradigm without taking into account strategic considerations. The reader is referred to <http://www.unifr.ch/econophysics/minority/> for an extensive list of references.

In our paper we will consider a *repeated* minority game and we will look at it according to the classical rational approach of game theory. Notice that, if after each stage each player observes the players which are in the room she selected, then, by the folk theorem, any feasible payoff is an equilibrium payoff of the repeated game. We study here the following version of a repeated minority game. At each step the players choose an action (one of two rooms). After their choice only a public signal (the majority room) is announced to all players. Therefore they do not observe the actions or the payoffs of the other players. The game is infinitely repeated and the payoffs are not discounted. We use the standard notion of uniform equilibrium, which will turn out to be payoff-equivalent here to that of almost sure equilibrium. We characterize the set of equilibrium payoffs.

Our model is a particular case of repeated games with imperfect observation: The players repeat a known one-shot game and after each stage each player receives a signal depending on the actions played. The reader is referred to Sorin (1992) for a survey of repeated games with complete information. Renault and Tomala (2000) characterized the set of uniform *communication* equilibrium payoffs for any repeated game with imperfect monitoring, but no general characterization exists for (Nash) equilibrium payoffs. Fudenberg and Maskin (1986) proved a folk theorem for a certain class of repeated games with discounting. Lehrer (1989, 1992a,b) dealt with two-person undiscounted repeated games with imperfect observation. More recently Tomala (1998) studied the case of public signals, where all players get the same signal after each stage. In this setup he characterized the set of *pure* uniform equilibrium payoffs. He also provided a characterization of the set of uniform (possibly mixed) equilibrium payoffs in a certain class of games, where all payoffs can be deduced from the public signal (Tomala (1999)).

Since we are interested not only in pure equilibria but in all (possibly mixed) uniform equilibria, the solution to our problem cannot be found in the existing literature. We will prove that a folk theorem holds for our game, i.e. we will show that any feasible payoff is an equilibrium payoff. In particular, we will construct a uniform equilibrium where the payoff of each player is simply zero. This equilibrium can be considered as particularly inefficient, since all feasible payoffs are non negative. It contains a main path and punishment phases. A punishment phase starts when the players suspect that a deviation have occurred. The identity of the possible deviator is not known by the players and it is not possible to punish simultaneously all players suspected of deviation, as done in several recent papers (Tomala (1999); Renault and Tomala (2000)). On the other hand it is possible to punish the deviator, if any, *by replicating some actions previously played in the main path before the punishment phase*. To our knowledge, this kind of punishment is new in the literature. The technical parts of our proofs use statistical techniques due to Lehrer (1990, 1992b), or, more specifically, the variations used by Renault (2000). In our opinion, the construction of our inefficient equilibrium gives insights concerning the difficulty of a general characterization of equilibrium payoffs in repeated games with public signals.

For the sake of simplicity, we first deal with the case of three players. Section 2 contains the model, and the statement of our main result. In Section 3 we define a particular strategy where all players are, at almost all stages with great probability, in the same room. In Section 4 we prove that this strategy is a uniform equilibrium with payoff 0 for each player. In Section 5 we finally extend our result to the case of any odd number of players.

2 The model

If \mathcal{E} is an event, then \mathcal{E}^c is its complementary event. The cardinality of a finite set A will be indicated as $|A|$. If C is a subset of an Euclidean space, $\text{conv}C$ is the convex hull of C .

There are two rooms: L (eft) and R (ight). At each stage, three players have to choose simultaneously one of the two rooms. The player who finds herself in the less crowded room (if any) gains a positive payoff of 1, and the most crowded room is publicly announced before going to the next stage.

2.1 The stage game

The set of players is $N = \{1, 2, 3\}$. For all $i \in N$, we call $A^i = \{L, R\}$ the set of actions for player i , and we put $A = A^1 \times A^2 \times A^3$. For $a = (a^1, a^2, a^3) \in A$ define the payoff function $g^i : A \rightarrow \mathbb{R}$ of player i as

$$g^i(a) = \begin{cases} 0 & \text{if there exists } j \in N \setminus \{i\} \text{ s.t. } a^j = a^i, \\ 1 & \text{otherwise,} \end{cases}$$

It is easy to compute the equilibria of the one-shot game. These are the action profiles such that one player plays L with probability 1 and another player plays R with probability 1, and the action profile where each player plays L and R with equal probability. Consequently, the set of equilibrium payoffs of the one-shot game is just

$$E_1 = \{(1/4, 1/4, 1/4)\} \cup \{(x, 1-x, 0) : x \in [0, 1]\} \\ \cup \{(x, 0, 1-x) : x \in [0, 1]\} \cup \{(0, x, 1-x) : x \in [0, 1]\}.$$

Notice that all payoffs $x = (x^1, x^2, x^3)$ in E_1 satisfy $x^1 + x^2 + x^3 \geq 3/4$. In a Nash equilibrium of the one-shot game, the three players are in the same room with probability at most $1/4$.

Since the stage game will be repeated, we also need notations about what the players observe. We define the set of public signals as $U = \{L, R\}$. The signalling function $\ell : A \rightarrow U$, giving the most crowded room, is formally defined by

$$\ell(R, R, R) = \ell(R, R, L) = \ell(R, L, R) = \ell(L, R, R) = R, \\ \ell(L, L, L) = \ell(L, L, R) = \ell(L, R, L) = \ell(R, L, L) = L.$$

2.2 The repeated game Γ_∞

At each stage $t \geq 1$, each player i (simultaneously with the other players) selects an action $a_t^i \in A^i$. If $a_t = (a_t^1, a_t^2, a_t^3) \in A$ is chosen, the stage payoff of player i is $g^i(a_t)$, and the signal $u_t = \ell(a_t)$ is publicly announced. Then the play proceeds to stage $t+1$. All the players have perfect recall and the whole description of Γ_∞ is common knowledge.

The game Γ_∞ is a game with imperfect monitoring, in that the players do not observe the actions of their opponents, but only a signal (the majority room).

2.3 The equilibria of Γ_∞

A behavioral strategy of player i is an element $\sigma^i = (\sigma_t^i)_{t \geq 1}$, where for all t

$$\sigma_t^i : (A^i \times U)^{t-1} \rightarrow \Delta(A^i).$$

Therefore, for each $t \geq 1$, $\sigma_t^i(a_1^i, u_1, a_2^i, u_2, \dots, a_{t-1}^i, u_{t-1})$ is the lottery played by player i at stage t if she played a_1^i at stage 1, \dots , a_{t-1}^i at stage $t-1$, and the signal was u_1 at stage 1, \dots , u_{t-1} at stage $t-1$.

We call Σ^i the set of behavioral strategies of player i , and $\Sigma = \Sigma^1 \times \Sigma^2 \times \Sigma^3$. A strategy profile $\sigma = (\sigma^1, \sigma^2, \sigma^3) \in \Sigma$ induces a probability measure \mathbb{P}_σ over the set of plays $\Omega = (A \times U)^\infty = \{(a_1, u_1, a_2, u_2, \dots), \forall t \geq 1, a_t \in A, u_t \in U\}$. With an abuse of notation we will indicate as a_t the random variable of the joint action profile in A played at stage t . For all $i \in N$, and for all $T \geq 1$,

$$\gamma_T^i(\sigma) = \mathbb{E}_{\mathbb{P}_\sigma} \left(\frac{1}{T} \sum_{t=1}^T g^i(a_t) \right).$$

Definition 1. The profile σ is a *uniform equilibrium* of Γ_∞ if

- (a) for all $i \in N$, $(\gamma_T^i(\sigma))_T$ converges to some quantity x^i as T goes to infinity (the vector (x^1, x^2, x^3) is called the payoff of σ).
- (b) for all $\epsilon > 0$ there exists T_0 such that for all $T \geq T_0$, σ is an ϵ -Nash equilibrium in the finitely repeated game with T stages, i.e. for all $i \in N$, for all $\tau^i \in \Sigma^i$, $\gamma_T^i(\tau^i, \sigma^{-i}) \leq \gamma_T^i(\sigma) + \epsilon$.

Definition 2. The vector $x \in \mathbb{R}^3$ is an *equilibrium payoff* of Γ_∞ if there exists a uniform equilibrium with payoff x .

We denote by E_∞ the set of equilibrium payoffs of Γ_∞ .

Since all payoffs are nonnegative and $g^1 + g^2 + g^3 \leq 1$, it is clear that E_∞ is a subset of the simplex \mathcal{S} , where

$$\begin{aligned} \mathcal{S} &= \left\{ (x^1, x^2, x^3) \in \mathbb{R}^3 : \text{for all } i \in N, x^i \geq 0, \text{ and } \sum_{i=1}^3 x^i \leq 1 \right\} \\ &= \text{conv}\{(0, 0, 0), (0, 0, 1), (0, 1, 0), (1, 0, 0)\}. \end{aligned}$$

Our main result is the following theorem.

Theorem 3. $E_\infty = \mathcal{S}$.

Since repeating at each stage a Nash equilibrium of the one-shot game is a uniform equilibrium of Γ_∞ , we know that $E_1 \subset E_\infty$. Moreover E_∞ is convex, so the only thing we have to do is to prove the following theorem.

Theorem 4. $(0, 0, 0) \in E_\infty$.

In order to prove the above theorem we need to construct a strategy $\sigma \in \Sigma$ that satisfies the two properties of Definition 1, namely,

- (a) for all $i \in N$, $\lim_{T \rightarrow \infty} \gamma_T^i(\sigma) = 0$,
- (b) for all $\epsilon > 0$, there exists T_0 such that for all $T \geq T_0$, for all $i \in N$,

$$\gamma_T^i(\tau^i, \sigma^{-i}) \leq \epsilon \quad \text{for all } \tau^i \in \Sigma^i.$$

Note that since all payoffs are non negative, (a) is a consequence of (b) here.

3 Construction of the strategy σ

We first give a heuristic description of the uniform equilibrium, $\sigma = (\sigma^1, \sigma^2, \sigma^3)$.

To get a payoff of 0, we need all the players to play with high probability the same action (say L) most of the stages. But if all players play L with probability 1, the deviation of one player, that consists of playing R , will be profitable (in terms of payoffs) and will not be detected (the signal will still be L). Hence some of the players must play R with small but positive probability.

Imagine all players play at each stage R with probability ϵ , where ϵ is small but positive. In order to detect a deviation, we will need a statistical test. If the frequency of stages where R is the most crowded room is higher than it should be, all players will consider that a deviation has occurred and a punishment phase will start. We then need to define an appropriate punishment phase, the difficulty being that the identity of the deviator (if any) is not known by the players. Our main idea is then the following. If player i is deviating, then with great probability at most of the stages where R was the most crowded room, the situation was the following: Player i played R , and exactly one of the other players played R , too. So if the players different from i repeat the actions they have played *at the stages where R was the most crowded*, at most stages one of them will play L and the other will play R . This punishes player i by giving him a payoff of zero.

We now formally construct σ . The set of stages $\{1, 2, \dots\}$ is divided into consecutive blocks of increasing lengths B^1, \dots, B^m, \dots , such that for all $m \geq 1$, $|B^m| = m^{10}$. This is needed because we need the statistical tests to become more and more accurate. The strategy σ consists of a main path and of punishment phases, starting from the main path.

When the play is in the main path, at some block B^m , all players play at each stage t of B^m , independently of what happened before, the mixed action

$$\left(1 - \frac{1}{m}\right) L \oplus \frac{1}{m} R.$$

At the end of such a block, all players can compute the empirical frequency of “ R being the most crowded room” in this block

$$\alpha_m = \frac{1}{|B^m|} |\{t \in B^m, \ell(a_t) = R\}|.$$

Note that if no player deviates at block B^m , by Tchebychev’s inequality α_m should be close to the expectation of “ R being the most crowded room”, which is equivalent to $3/m^2$. The statistical test will be the following:

- If $\alpha_m \leq 1/(m\sqrt{m})$, the test will be considered as passed. The play stays in the main path (and block B^{m+1} is played).

- If $\alpha_m > 1/(m\sqrt{m})$, the test will be considered as failed, and the players will assume that a deviation has occurred. The play will immediately go out of the main path and a punishment phase will start. The punishment phase will last a large number of blocks, but will not be infinite, because there will always be a chance for the punishment to fail. More precisely, the punishment phase will last from the first stage of block B^{m+1} to the last stage of block B^{m^2} . Then, and whatever happens during the punishment phase, the play will go back to the main path at block B^{m^2+1} .

To complete the definition of σ , we have to define what is played in the punishment phases.

Let m be a positive integer, and consider a block B^m where the play is in the main path, such that $\alpha_m > 1/(m\sqrt{m})$, namely, the test fails. Define

$$D = \{t \in B^m : \ell(a_t) = R\}. \quad (1)$$

On the set D we suspect the deviator, if any, to have played R on purpose. We have $|D| = m^{10}\alpha_m$. In order to play the punishment phase at blocks B^{m+1}, \dots, B^{m^2} , each player will have to remember D and the action she played at each stage of D . We order the elements of D so that $D = \{t_1, \dots, t_{|D|}\}$, with $t_1 < t_2 < \dots < t_{|D|}$.

Fix $\bar{m} \in \{m+1, \dots, m^2\}$. We now define what σ recommends to play at such block $B^{\bar{m}}$ during a punishment phase.

Let $d \in \mathbb{N}$ be such that

$$d \leq \frac{|B^{\bar{m}}|}{|D|} < d+1.$$

The block $B^{\bar{m}}$ is divided into consecutive sub-blocks $B_1^{\bar{m}}, \dots, B_d^{\bar{m}}, B_{d+1}^{\bar{m}}$ such that for all $d' \in \{1, \dots, d\}$, $|B_{d'}^{\bar{m}}| = |D|$. The role of $B_{d+1}^{\bar{m}}$ will be negligible since $|B_{d+1}^{\bar{m}}| < |D|$. Indeed we have

$$\frac{|B_{d+1}^{\bar{m}}|}{|B^{\bar{m}}|} < \frac{|D|}{\bar{m}^{10}}. \quad (2)$$

With high probability the right hand side of (2) will be small when m is large, even in case of deviation. Consequently we can define σ arbitrarily on such a block $B_{d+1}^{\bar{m}}$.

Let $d' \in \{1, \dots, d\}$. At $B_{d'}^{\bar{m}}$ the strategy σ recommends the players to mimic what happened at stages in D . If $B_{d'}^{\bar{m}} = \{t'_1, \dots, t'_{|D|}\}$ with $t'_1 < t'_2 < \dots < t'_{|D|}$, then σ recommends each player i at each stage $t'_n \in B_{d'}^{\bar{m}}$ (with $n \in \{1, \dots, |D|\}$) to repeat the action she played at stage t_n , i.e. to play $a_{t_n}^i$.

Notice that σ recommends to play exactly the same sequence of actions at each sub-block $B_1^{\bar{m}}, \dots, B_d^{\bar{m}}$.

4 σ is a uniform equilibrium with payoff $(0, 0, 0)$

We first informally discuss the proof.

1. Suppose that all players follow σ . Then at each stage of some block B^m in the main path the probability of R being the most crowded room is equivalent (as m goes to the infinity) to $3/m^2$. Consequently, by Tchebychev's inequality, α_m will be close to $3/m^2$ with high probability. Since $3/m^2 < 1/(m\sqrt{m})$, for m large, the test of block B^m will pass. It will even be possible, by Borel-Cantelli lemma, to show that the set of blocks m such that B^m is not in the main path is almost surely finite. Moreover the (stage) average payoff of some player i at some block B^m in the main path will be close to the probability that she plays R whereas the others play L , hence will be close to $1/m$. This will ensure that the average payoff of each player will go to zero as the number of stages goes to infinity.
2. Suppose that some player (e.g. player 1) deviates from σ . In order for player 1 to have a good payoff at some block B^m in the main path, she should play R a large number of times in this block. We will see that in this case the empirical frequency of “ R being the most crowded room” will be greater than $1/(m\sqrt{m})$ with high probability. Hence a punishment phase will start, and the efficiency of this phase to punish player 1 will only depend on what happened at stages in $D = \{t \in B^m : \ell(a_t) = R\}$.

The set D consists of two kinds of stages: (i) the stages where player 1 played R and exactly one of the other players played R , and (ii) the stages where both players 2 and 3 played R . We will show that, with high probability, the stages of type (ii) are negligible. Consequently, for most of the stages in D , player 2 and player 3 do not play the same action. Hence for most of the punishment stages, player 1's payoff will be zero.

Summing up, player 1 cannot have a good payoff on some block in the main path without being severely punished afterwards with high probability. This will ensure that no deviation is profitable.

To show that σ is a uniform equilibrium, we only need to prove Proposition 6 below. However we will first shortly prove the following Proposition 5 to simplify the exposition of our proof (and because the analogue of Proposition 5 will be needed in Section 5).

Proposition 5. *For all $i \in N$,*

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T g^i(a_t) = 0 \quad \mathbb{P}_\sigma\text{-a.s.}, \quad \text{and} \quad \lim_{T \rightarrow \infty} \gamma_T^i(\sigma) = 0.$$

Proof of Proposition 5. By symmetry, we only consider the case where $i = 1$. Assume that all players play σ . All the probabilities and expectations in the sequel of the proof are computed according to $\mathbb{P} = \mathbb{P}_\sigma$.

For each block m , we define the following events:

$$\mathcal{B}_m = \{\text{the play is in the main path at block } B^m\},$$

$$\mathcal{A}_m = \left\{ \sum_{t \in B^m} \frac{g^1(a_t)}{|B^m|} > \frac{2}{m} \right\}.$$

Fix a block number m where \mathcal{B}_m holds. At each stage of B^m each player plays i.i.d. the mixed action

$$\left(1 - \frac{1}{m}\right) L \oplus \frac{1}{m} R.$$

So at each stage the probability that R is the most crowded room is

$$\eta_m = \frac{1}{m^3} + 3\frac{1}{m^2} \left(1 - \frac{1}{m}\right) \leq \frac{3}{m^2},$$

and the probability that player 1 has a payoff of 1 is

$$\frac{1}{m} \left(1 - \frac{1}{m}\right)^2 + \left(1 - \frac{1}{m}\right) \frac{1}{m^2} = \frac{1}{m} \left(1 - \frac{1}{m}\right).$$

We have, by Tchebychev's inequality

$$\mathbb{P}(\mathcal{A}_m | \mathcal{B}_m) \leq \mathbb{P} \left(\left| \frac{\sum_{t \in B^m} g^1(a_t)}{|B^m|} - \frac{1}{m} \left(1 - \frac{1}{m}\right) \right| > \frac{1}{m} \middle| \mathcal{B}_m \right) \leq \frac{1}{m^8}. \quad (3)$$

Moreover, for m large enough,

$$\begin{aligned} \mathbb{P}(\mathcal{B}_{m+1}^c | \mathcal{B}_m) &= \mathbb{P} \left(\alpha_m > \frac{1}{m\sqrt{m}} \middle| \mathcal{B}_m \right) \\ &= \mathbb{P} \left(\alpha_m - \eta_m > \frac{1}{m\sqrt{m}} - \eta_m \middle| \mathcal{B}_m \right) \\ &\leq \mathbb{P} \left(|\alpha_m - \eta_m| > \frac{1}{2m\sqrt{m}} \middle| \mathcal{B}_m \right) \\ &\leq \frac{4}{m^7}, \end{aligned}$$

Again the last inequality is just Tchebychev. Since $\sum_{m \geq 1} 4/m^7 < \infty$, by Borel-Cantelli lemma we obtain

$$\mathbb{P} \left(\limsup (\mathcal{B}_m \cap \mathcal{B}_{m+1}^c) \right) = 0. \quad (4)$$

Since after a punishment phase the play always comes back to the main path, (4) implies that with probability 1 there exists a block m_1 such that for each $m \geq m_1$, \mathcal{B}_m holds.

By Borel-Cantelli lemma again and (3), we now have $\mathbb{P}(\limsup(\mathcal{A}_m \cap \mathcal{B}_m)) = 0$, hence $\mathbb{P}(\limsup \mathcal{A}_m) = 0$. Hence, with probability 1, there exists a block m_2 such that for all $m \geq m_2$,

$$\frac{\sum_{t \in B^m} g^1(a_t)}{|B^m|} \leq \frac{2}{m}.$$

Since the cardinality of the B^m is polynomial in m , we have

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T g^1(a_t) = 0 \quad \mathbb{P}_\sigma\text{-a.s.}$$

By the bounded convergence theorem we also have that $\lim_{T \rightarrow \infty} \gamma_T^1(\sigma) = 0$. \square

Proposition 6. *For all $\epsilon > 0$ there exists T_0 such that for all $T \geq T_0$, for all $i \in N$*

$$\gamma_T^i(\tau^i, \sigma^{-i}) \leq \epsilon \quad \text{for all } \tau^i \in \Sigma^i.$$

Proof of Proposition 6. Without loss of generality, we consider only deviations by player 1. Fix $\tau^1 \in \Sigma^1$ in all the sequel, and assume that $(\tau^1, \sigma^2, \sigma^3)$ is played. All the probabilities and expectations in the sequel will be with respect to $\mathbb{P} = \mathbb{P}_{\tau^1, \sigma^2, \sigma^3}$.

For each block m we define the following random variables:

$$\begin{aligned} X_m &= \frac{1}{|B^m|} \sum_{t \in B^m} g^1(a_t), & Z_m &= \frac{1}{|B^m|} \sum_{t \in B^m} \mathbf{1}_{\{a_t^2=R\} \cup \{a_t^3=R\}}, \\ U_m &= \frac{1}{|B^m|} \sum_{t \in B^m} \mathbf{1}_{\{a_t^2=R\} \cap \{a_t^3=R\}}, & x_m &= \frac{1}{|B^m|} \sum_{t \in B^m} \mathbf{1}_{\{a_t^1=R\}}. \end{aligned}$$

We have $X_m \leq U_m + x_m$. We also define the event

$$\begin{aligned} \mathcal{C}_m &= \mathcal{B}_m^c \cup \left(\mathcal{B}_m \cap \mathcal{B}_{m+1} \cap \left\{ X_m \leq \frac{3}{\sqrt{m}} \right\} \right) \\ &\quad \cup \left(\mathcal{B}_m \cap \mathcal{B}_{m+1}^c \cap \left(\bigcap_{m'=m+1}^{m^2} \left\{ X_{m'} \leq \frac{3}{\sqrt{m}} \right\} \right) \right). \end{aligned} \quad (5)$$

Conditionally on $\overline{\mathcal{C}_m}$, one of the following three possibilities is true: Either the play is in a punishment phase, or it is in the main path at B^m , player 1's payoff is low, and it will still be in the main path at B^{m+1} , or an efficient punishment starts at block B^{m+1} .

Lemma 7. *There exists M_1 , independent from τ^1 , such that for all $m \geq M_1$*

$$\mathbb{P}(\mathcal{C}_m) \geq 1 - \frac{2}{m^6}.$$

Proof of Lemma 7. Consider a block B^m , with m large enough, where the play is in the main path. Via Tchebychev's inequality we obtain

$$\begin{aligned}\mathbb{P}\left(Z_m > \frac{3}{m} \middle| \mathcal{B}_m\right) &\leq \frac{1}{m^8}, \\ \mathbb{P}\left(U_m > \frac{2}{m^2} \middle| \mathcal{B}_m\right) &\leq \frac{1}{m^6}.\end{aligned}$$

Hence with high probability player 2 and 3 will not be simultaneously in room R at the same stages.

We now want to estimate the number of stages where R is the most crowded room and exactly two players, including player 1, are in R . Define for all $t \in B^m$

$$Q_t = \mathbf{1}_{\{a_t^1=R\}}, \quad \xi_t = \mathbf{1}_{\{a_t^2=R\} \cap \{a_t^3=L\}} + \mathbf{1}_{\{a_t^2=L\} \cap \{a_t^3=R\}}.$$

The random variables $(\xi_t)_{t \in B^m}$ are i.i.d. (given \mathcal{B}_m) Bernoulli random variables with expectation

$$p_m = 2\frac{1}{m} \left(1 - \frac{1}{m}\right) \geq \frac{1}{m},$$

for $m \geq 2$. The variables $(Q_t)_{t \in B^m}$ may not be independent and may not be independent of $(\xi_t)_{t \in B^m}$, since player 1 is using an arbitrary strategy τ^1 . Nevertheless, for each $t \in B^m$, ξ_t is independent of $(\xi_{t'})_{t' \in B^m, t' < t}$ and $(Q_{t'})_{t' \in B^m, t' < t}$. Hence we can apply a generalization of Tchebychev's inequality due to Lehrer (see Lehrer (1990), Lemma 5.6). For all $\epsilon' > 0$

$$\mathbb{P}\left(\left|\sum_{t \in B^m} \frac{\xi_t Q_t}{m^{10}} - p_m x_m\right| \geq \epsilon' \middle| \mathcal{B}_m\right) \leq \frac{1}{m^{10} \epsilon'^2}.$$

The choice of $\epsilon' = 1/(m\sqrt{m})$ gives

$$\mathbb{P}\left(\left|\sum_{t \in B^m} \frac{\xi_t Q_t}{m^{10}} - p_m x_m\right| \geq \frac{1}{m\sqrt{m}} \middle| \mathcal{B}_m\right) \leq \frac{1}{m^7}.$$

Assume now that there is no punishment phase at block B^{m+1} , i.e. that \mathcal{B}_{m+1} holds. This implies

$$\sum_{t \in B^m} \frac{\xi_t Q_t}{m^{10}} \leq \frac{1}{m\sqrt{m}}, \quad \text{and} \quad p_m x_m \leq \frac{1}{m\sqrt{m}} + \left|p_m x_m - \sum_{t \in B^m} \frac{\xi_t Q_t}{m^{10}}\right|.$$

Assume also that

$$\left|p_m x_m - \sum_{t \in B^m} \frac{\xi_t Q_t}{m^{10}}\right| \leq \frac{1}{m\sqrt{m}} \quad \text{and} \quad U_m \leq \frac{2}{m^2}.$$

Then $p_m x_m \leq 2/(m\sqrt{m})$, so $x_m \leq 2/\sqrt{m}$. Since $X_m \leq U_m + x_m$, we get $X_m \leq 3/\sqrt{m}$ for $m \geq 2$. We have shown that

$$\mathcal{B}_m \cap \mathcal{B}_{m+1} \cap \left(U_m \leq \frac{2}{m^2} \right) \cap \left(\left| p_m x_m - \sum_{t \in B^m} \frac{\xi_t Q_t}{m^{10}} \right| \leq \frac{1}{m\sqrt{m}} \right) \subset \left\{ X_m \leq \frac{3}{\sqrt{m}} \right\}.$$

Consequently

$$\begin{aligned} & \mathbb{P} \left(\mathcal{B}_{m+1} \cap \left(U_m \leq \frac{2}{m^2} \right) \cap \left\{ X_m > \frac{3}{\sqrt{m}} \right\} \middle| \mathcal{B}_m \right) \\ & \leq \mathbb{P} \left(\left| p_m x_m - \sum_{t \in B^m} \frac{\xi_t Q_t}{m^{10}} \right| > \frac{1}{m\sqrt{m}} \middle| \mathcal{B}_m \right) \\ & \leq \frac{1}{m^7}. \end{aligned}$$

We obtain as a first result

$$\begin{aligned} & \mathbb{P} \left(\mathcal{B}_m \cap \left(\left(Z_m > \frac{3}{m} \right) \cup \left(U_m > \frac{2}{m^2} \right) \cup \left(\mathcal{B}_{m+1} \cap \left\{ X_m > \frac{3}{\sqrt{m}} \right\} \right) \right) \right) \\ & \leq \mathbb{P} \left(\left(\left(Z_m > \frac{3}{m} \right) \cup \left(U_m > \frac{2}{m^2} \right) \cup \left(\mathcal{B}_{m+1} \cap \left\{ X_m > \frac{3}{\sqrt{m}} \right\} \right) \right) \middle| \mathcal{B}_m \right) \\ & \leq \frac{1}{m^8} + \frac{1}{m^6} + \frac{1}{m^7} \\ & \leq \frac{2}{m^6} \quad \text{for } m \geq 2. \end{aligned}$$

Therefore if we define the event

$$\mathcal{G}_m = \mathcal{B}_m^c \cup \left(\left(Z_m \leq \frac{3}{m} \right) \cap \left(U_m \leq \frac{2}{m^2} \right) \cap \left(\mathcal{B}_{m+1}^c \cup \left\{ X_m \leq \frac{3}{\sqrt{m}} \right\} \right) \right),$$

we have

$$\mathbb{P}(\mathcal{G}_m) \geq 1 - \frac{2}{m^6}. \quad (6)$$

Assume that \mathcal{G}_m and \mathcal{B}_m hold. Then

- either \mathcal{B}_{m+1} holds, and this implies that $X_m \leq 3/\sqrt{m}$,
- or \mathcal{B}_{m+1}^c holds, and therefore a punishment phase starts at block B^{m+1} .

Consider D as defined in (1). We have $|D| > m^{8.5}$. As $|D| \leq m^{10} Z_m$, the event $(Z_m \leq 3/m)$, implies $|D| \leq 3m^9$.

Since $(U_m \leq 2/m^2)$, the number of stages in D where player 2 and player 3 play the same action is at most $2m^8$.

Consider a block $B^{\bar{m}}$ with $\bar{m} \in \{m+1, \dots, m^2\}$. Let d be the integer such that $d \leq |B^{\bar{m}}|/|D| < d+1$. At each stage where player 2 plays L and player 3 plays R , or vice versa, player 1's payoff is 0. So the total payoff of player 1 at block $B^{\bar{m}}$ is

$$\bar{m}^{10} X_{\bar{m}} \leq d \cdot 2m^8 + |D| = d|D| \frac{2m^8}{|D|} + |D|.$$

Hence

$$X_{\bar{m}} \leq \frac{d|D|}{|B^{\bar{m}}|} \frac{2}{\sqrt{m}} + \frac{3m^9}{\bar{m}^{10}} \leq \frac{2}{\sqrt{m}} + \frac{3}{m} \leq \frac{3}{\sqrt{m}}, \quad \text{for } m \geq 9.$$

This implies that $\mathcal{G}_m \subset \mathcal{C}_m$. The desired result now follows from (6). \square

Lemma 7 is the key to the proof of Proposition 6. The rest is technical, and very close to the end of the proof in Renault (2000).

Lemma 8. *For all $\epsilon > 0$, there exists M_2 independent of τ^1 , such that for all $m_0 \geq M_2^2$,*

$$\mathbb{E} \left(\frac{\sum_{m=m_0}^{m_0^2} m^{10} X_m}{\sum_{m=m_0}^{m_0^2} m^{10}} \right) \leq 3\epsilon.$$

Proof of Lemma 8. Fix $\epsilon > 0$. Since

$$\sum_{m=1}^{\infty} \frac{1}{m^6} < +\infty,$$

by Lemma 7 one can find M_2 , independent of τ^1 , such that

$$\mathbb{P} \left(\bigcup_{m \geq M_2} \mathcal{C}_m^c \right) \leq \epsilon.$$

One may also assume that for all $m \geq M_2$, we have

$$3/\sqrt{m} \leq \epsilon, \tag{7}$$

and

$$\frac{|B^{m^2}|}{\sum_{m'=m}^{m^2} m'^{10}} \leq \epsilon. \tag{8}$$

Fix now $m_0 \geq M_2^2$, and put

$$Y = \sum_{m=m_0}^{m_0^2} m^{10} X_m.$$

Then

$$\begin{aligned}\mathbb{E}(Y) &= \mathbb{P}\left(\bigcup_{m \geq M_2} \mathcal{C}_m^c\right) \mathbb{E}\left(Y \mid \bigcup_{m \geq M_2} \mathcal{C}_m^c\right) + \mathbb{P}\left(\bigcap_{m \geq M_2} \mathcal{C}_m\right) \mathbb{E}\left(Y \mid \bigcap_{m \geq M_2} \mathcal{C}_m\right) \\ &\leq \epsilon \sum_{m=m_0}^{m_0^2} m^{10} + \mathbb{E}\left(Y \mid \bigcap_{m \geq M_2} \mathcal{C}_m\right).\end{aligned}$$

Assume that for all $m \geq M_2$, \mathcal{C}_m holds. We will show that this implies

$$Y \leq 2\epsilon \sum_{m=m_0}^{m_0^2} m^{10}.$$

By (5) and (7) we have sequences $(X_m)_{m \geq M_2}$ and $(\mathcal{B}_m)_{m \geq M_2}$ such that for all $m \geq M_2$ the following events are true

$$\mathcal{B}_m^c \cup \left(\mathcal{B}_m \cap \mathcal{B}_{m+1} \cap \{X_m \leq \epsilon\} \right) \cup \left(\mathcal{B}_m \cap \mathcal{B}_{m+1}^c \cap \left(\bigcap_{m'=m+1}^{m^2} \{X_{m'} \leq \epsilon\} \right) \right).$$

Since $m_0 \geq M_2^2$, and after a punishment phase the play always comes back to the main path, there necessarily exists some block number m_1 in $\{M_2, \dots, m_0\}$ such that \mathcal{B}_{m_1} holds.

Two cases are possible:

- (I) For all $m \geq m_1$, \mathcal{B}_m holds, and then for all $m \geq m_1$, $X_m \leq \epsilon$ and $Y \leq \epsilon \sum_{m=m_0}^{m_0^2} m^{10}$.
- (II) There exists a first block number $m_2 \geq m_1$ such that $\mathcal{B}_{m_2} \cap \mathcal{B}_{m_2+1}^c$ holds. We have $X_m \leq \epsilon$ whenever $m_1 \leq m < m_2$.

Two sub-cases of (II) are possible

- (i) $m_2 \geq m_0$. For all m such that $m_2 < m \leq m_0^2$, we have $X_m \leq \epsilon$ (the punishment starting from B^{m_2+1} will finish after $B^{m_0^2}$). So, by (8),

$$Y \leq \sum_{\substack{m \in \{m_0, \dots, m_0^2\} \\ m \neq m_2}} m^{10} \epsilon + m_2^{10} \leq 2\epsilon \sum_{m=m_0}^{m_0^2} m^{10}.$$

- (ii) $m_2 < m_0$. For all $m \in \{m_0, \dots, m_0^2\}$, $X_m \leq \epsilon$, and \mathcal{B}_{m_2+1} holds.

We just have to repeat the argument and consider the following sub-sub-cases.

- (a) for all $m \geq m_2^2 + 1$, \mathcal{B}_m holds. Then for all $m \in \{m_0, \dots, m_0^2\}$ we have $X_m \leq \epsilon$ and $Y \leq \epsilon \sum_{m=m_0}^{m_0^2} m^{10}$.
- (b) There exists a first block number $m_3 \geq m_2^2 + 1$ such that $\mathcal{B}_{m_3} \cap \mathcal{B}_{m_3+1}^c$ holds. The only possible block m in $\{m_0, \dots, m_0^2\}$ where we may have $X_m > \epsilon$ is block m_3 . So $Y \leq 2\epsilon \sum_{m=m_0}^{m_0^2} m^{10}$.

In the end we obtain

$$\mathbb{E}(Y) \leq 3\epsilon \sum_{m=m_0}^{m_0^2} m^{10},$$

and Lemma 8 is proved. \square

We can now conclude the proof of Proposition 6.

Fix $\epsilon > 0$. By Lemma 8, there exists a block number M_3 , independent of τ^1 , such that for all $m_0 \geq M_3$

$$\mathbb{E} \left(\frac{\sum_{t \in B^{m_0} \cup \dots \cup B^{m_0^2}} g^1(a_t)}{\sum_{m=m_0}^{m_0^2} m^{10}} \right) \leq \epsilon, \quad \sum_{m=1}^{m_0-1} m^{10} \leq \epsilon \sum_{m=1}^{m_0^2} m^{10}, \quad \sum_{m=m_0^2+1}^{(m_0+1)^2} m^{10} \leq \epsilon \sum_{m=1}^{m_0^2} m^{10}.$$

Define $T_0 = 1 + \max\{B^{M_3}\}$ and let $T \geq T_0$. Define $m(T)$ via $T \in B^{m(T)}$. Then $m(T) \geq M_3^2 + 1$. Define $l \in \mathbb{N}$ such that $l \leq \sqrt{m(T)} - 1 < l + 1$. We have $l \geq M_3$, $l^2 < m(T)$, and $(l+1)^2 \geq m(T)$.

$$\begin{aligned} \mathbb{E} \left(\sum_{t=1}^T g^1(a_t) \right) &= \sum_{t < \min\{B^l\}} g^1(a_t) + \sum_{t \in B^l \cup \dots \cup B^{l^2}} g^1(a_t) + \sum_{t > \max\{B^l\}} g^1(a_t) \\ &\leq \sum_{m=1}^{l-1} m^{10} + \sum_{m=l}^{l^2} m^{10} \epsilon + \sum_{m=l^2+1}^{(l+1)^2} m^{10}, \\ &\leq \epsilon + \epsilon + \epsilon = 3\epsilon. \end{aligned}$$

Proposition 6 is proved since T_0 does not depend on τ^1 . \square

Remark 9. (*Almost sure equilibrium payoffs*) As in Lehrer (1992a), we can define an almost sure equilibrium payoff as a vector $x = (x^1, x^2, x^3)$ in \mathbb{R}^3 such that there exists an (almost sure equilibrium) strategy profile σ satisfying

$$\forall i \in N, \quad \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T g^i(a_t) = x^i \quad \mathbb{P}_\sigma\text{-a.s.}, \quad (9)$$

and

$$\forall i \in N, \forall \tau^i \in \Sigma^i, \quad \limsup_T \left(\frac{1}{T} \sum_{t=1}^T g^i(a_t) \right) \leq x^i \quad \mathbb{P}_{\tau^i, \sigma^i}\text{-a.s.} \quad (10)$$

Take σ to be our inefficient uniform equilibrium just constructed. We only need to prove (10) with $x = (0, 0, 0)$ (because (9) is proved in Proposition 5, or because (9) is a consequence of (10) here).

Fix as before a strategy τ^1 of player 1 and define for each m , \mathcal{C}_m as in (5). By Lemma 7 and Borel-Cantelli lemma, with probability one we can find an integer M_4 , that may depend on τ^1 , such that for all $m \geq M_4$, \mathcal{C}_m holds. Looking at the proof of Lemma 8, this implies that for every $\epsilon > 0$, one can find M_4 such that for all $m_0 \geq M_4^2$,

$$\sum_{m=m_0}^{m_0^2} m^{10} X_m \leq 2\epsilon \sum_{m=m_0}^{m_0^2} m^{10}.$$

So

$$\limsup_T \left(\frac{1}{T} \sum_{t=1}^T g^1(a_t) \right) \leq 2\epsilon \quad P_{\tau^1, \sigma^1}\text{-a.s.}$$

Hence

$$\limsup_T \left(\frac{1}{T} \sum_{t=1}^T g^1(a_t) \right) = \lim_{T \rightarrow \infty} \left(\frac{1}{T} \sum_{t=1}^T g^1(a_t) \right) = 0 \quad P_{\tau^1, \sigma^1}\text{-a.s.}$$

Therefore σ is not only a uniform equilibrium; it is also an almost sure equilibrium, and $(0, 0, 0)$ is an almost sure equilibrium payoff.

It is then easy to see that for this game, the set of almost sure equilibrium payoffs coincides with the set of uniform equilibrium payoffs.

5 An odd number of players

We generalize the model of Section 2 as follows. The set of players is now $N = \{1, \dots, 2n + 1\}$, where n is a fixed positive integer. At each stage, each player gets a payoff of 1 if he is in the minority room, and gets a payoff of 0 otherwise. The signal is again the most crowded room. The previous definitions of equilibrium and equilibrium payoffs extend unambiguously to this general model.

For each subset S of N such that $|S| \leq n$, define e_S as the payoff in \mathbb{R}^N where each player in S gets 1, and each player not in S gets 0. If $S = \emptyset$, then e_S is just the null vector. The set of feasible vectors is now

$$\mathcal{S} = \text{conv}\{e_S, S \subset N \text{ s.t. } |S| \leq n\}.$$

We show that also in this general case the set of uniform equilibrium payoffs and the set of feasible payoffs coincide.

Theorem 10. $E_\infty = \mathcal{S}$.

Proof of Theorem 10. The proof is a generalization of the proof for the three-player case. If S is a subset of N with exactly n elements, e_S is a Nash equilibrium of the one-shot game, hence e_S is also a uniform equilibrium payoff. By convexity, to prove that the set of uniform equilibrium payoffs is \mathcal{S} it is sufficient to show that for any S with $|S| < n$, we can construct a uniform equilibrium with payoff e_S .

Fix a subset S of players such that $|S| < n$. We need to construct a strategy profile $\sigma = (\sigma^i)_{i \in N}$ such that σ is a uniform equilibrium with payoff e_S . The construction of Section 3 generalizes as follows.

If $i \in S$, σ^i is very simple: play R at each stage in $\{1, 2, \dots\}$, independently of what happened before.

Divide the set of stages $\{1, 2, \dots\}$ into consecutive blocks B^1, \dots, B^m, \dots with $|B^m| = m^{10}$ for each m , exactly as in Section 3. The strategy σ consists of a main path and of punishment phases, starting from the main path. When the play is in the main path at some block B^m , each player i in $N \setminus S$ plays i.i.d. at each stage the mixed action

$$(1 - \delta_m)L \oplus \delta_m R, \quad \text{with} \quad \delta_m = m^{\frac{-2}{n+1-|S|}}.$$

Notice that

$$0 < \frac{2}{n+1-|S|} \leq 1,$$

so $\delta_m \geq 1/m$ and $\lim_{m \rightarrow \infty} \delta_m = 0$.

At the end of such a block, all players compute as before the empirical frequency of “ R being the most crowded room” in this block

$$\alpha_m = \frac{1}{|B^m|} |\{t \in B^m, \ell(a_t) = R\}|.$$

Put

$$\theta_m = m^{\frac{-2(n+1/2-|S|)}{n+1-|S|}} = \frac{1}{m^2} m^{\frac{1}{n+1-|S|}} \in \left] \frac{1}{m^2}, \frac{1}{m\sqrt{m}} \right].$$

The statistical test is the following:

- If $\alpha_m \leq \theta_m$, the test is passed. The play stays in the main path (and block B^{m+1} is played).
- If $\alpha_m > \theta_m$, the test fails. Define $D = \{t \in B^m, \ell(a_t) = R\}$. A punishment phase is played from the first stage of block B^{m+1} to the last stage of block B^{m^2} . Then the play goes back to the main phase at block B^{m^2+1} . Punishment are similar to the ones in Section 3. Each block $B^{\bar{m}}$, with $\bar{m} \in \{m+1, \dots, m^2\}$ is divided into sub-blocks $B_1^{\bar{m}}, \dots, B_d^{\bar{m}}, B_{d+1}^{\bar{m}}$, with $|B_1^{\bar{m}}| = \dots = |B_d^{\bar{m}}| = |D|$. At each sub-block $B_{d'}^{\bar{m}}$, with $d' \in \{1, \dots, d\}$, the players play again in the same order the actions they have played at D .

Notice that if $n = 1$ and $S = \emptyset$, σ is exactly the strategy constructed in Section 3. To conclude, we have to prove that σ is a uniform equilibrium with payoff e_S . In the following computations, “if m is large enough” should be understood as “if m is larger than some constant only depending on n and $|S|$.” We will use the following binomial coefficients:

$$K_1 = \binom{2n+1-|S|}{n+1-|S|}, \quad K_2 = \binom{2n-|S|}{n+1-|S|}, \quad K_3 = \binom{2n-|S|}{n-|S|}$$

A) Assume that all players follow σ . Probabilities are computed according to $\mathbb{P} = \mathbb{P}_\sigma$.

Fix a block number m where $\mathcal{B}_m = \{\text{the play is in the main path at block } B^m\}$ holds. Consider some stage in this block. For R to be the most crowded room at this stage we need at least $n+1-|S|$ players in $N \setminus S$ to play R , hence the probability that R is the most crowded room is

$$\eta_m \leq \binom{2n+1-|S|}{n+1-|S|} \delta_m^{n+1-|S|} = K_1 \frac{1}{m^2}.$$

Note that

$$\theta_m - \eta_m = \frac{1}{m^2} (m^{\frac{1}{n+1-|S|}} - K_1) \geq \frac{1}{m^2},$$

if m is large. Hence by Tchebychev’s inequality and if m is large we get

$$\begin{aligned} \mathbb{P}(\mathcal{B}_{m+1}^c | \mathcal{B}_m) &= \mathbb{P}(\alpha_m > \theta_m | \mathcal{B}_m) \\ &\leq \mathbb{P}\left(|\alpha_m - \eta_m| > \frac{1}{m^2} \middle| \mathcal{B}_m\right) \\ &\leq \frac{1}{m^6}. \end{aligned}$$

By Borel-Cantelli lemma we obtain as in Section 3 that with probability 1 there exists m_1 such that for each $m \geq m_1$, \mathcal{B}_m holds.

Let i be a player in S . At some stage in the main path, the probability that player i ’s payoff is 1 is the probability that L is the most crowded room, hence it is at least $1 - K_1/m^2$. Since $|1 - 1/m - (1 - K_1/m^2)| > 1/(2m)$ for m large, by Tchebychev’s inequality one can prove that

$$\mathbb{P}\left(\frac{1}{|B^m|} \sum_{t \in B^m} g^i(a_t) < 1 - \frac{1}{m} \middle| \mathcal{B}_m\right) \leq \frac{4m^2}{|B^m|} = \frac{4}{m^8}.$$

Again by Borel-Cantelli lemma with probability 1 there will exist a block number m_2 such that for each $m \geq m_2$,

$$\frac{1}{|B^m|} \sum_{t \in B^m} g^i(a_t) \geq 1 - \frac{1}{m}.$$

From this it follows

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T g^i(a_t) = 1 \quad \mathbb{P}_{\sigma}\text{-a.s.}, \quad \text{and} \quad \lim_{T \rightarrow \infty} \gamma_T^i(\sigma) = 1.$$

Let now i be a player in $N \setminus S$. Fix m where \mathcal{B}_m holds. At some stage t in B^m , if player i 's payoff is 1 then either she plays R or R is the most crowded room, hence

$$\mathbb{P}(g^i(a_t) = 1 | \mathcal{B}_m) \leq \delta_m + \frac{K_1}{m^2} \leq 2\delta_m \quad (\text{for } m \text{ large}).$$

Tchebychev's inequality then shows that

$$\mathbb{P} \left(\frac{1}{|B^m|} \sum_{t \in B^m} g^i(a_t) > 3\delta_m \middle| \mathcal{B}_m \right) \leq \frac{1}{\delta_m^2 |B^m|} \leq \frac{1}{m^8}.$$

And as before, $\lim_{T \rightarrow \infty} \gamma_T^i(\sigma) = 0$.

B) It remains to prove that no player can benefit by deviating from σ . Since 1 is the largest possible payoff in the game, we do not have to care about deviations by players in S . We thus only consider a deviation of some player i not in S . By symmetry, we assume that $i = 1 \notin S$ and fix in all the sequel a deviation τ^1 of player 1. We use the probability $\mathbb{P} = \mathbb{P}_{\tau^1, \sigma^{-1}}$. For each m , denote as before the average payoff of player 1 at block B^m as

$$X_m = \frac{1}{|B^m|} \sum_{t \in B^m} g^1(a_t).$$

The definition of \mathcal{C}_m generalizes as follows

$$\begin{aligned} \mathcal{C}_m = & \mathcal{B}_m^c \cup \left(\mathcal{B}_m \cap \mathcal{B}_{m+1} \cap \{X_m \leq 3\sqrt{\delta_m}\} \right) \\ & \cup \left(\mathcal{B}_m \cap \mathcal{B}_{m+1}^c \cap \left(\bigcap_{m'=m+1}^{m^2} \{X_{m'} \leq 3K_2\sqrt{\delta_m}\} \right) \right). \end{aligned} \quad (11)$$

We are going to prove the following analogue of Lemma 7.

Lemma 11. *There exists M_1 , independent from τ^1 , such that for all $m \geq M_1$*

$$\mathbb{P}(\mathcal{C}_m) \geq 1 - \frac{3}{m^6}.$$

Notice that

$$\lim_{m \rightarrow \infty} \sqrt{\delta_m} = \lim_{m \rightarrow \infty} m^{\frac{-1}{n+1-|S|}} = 0.$$

Once Lemma 11 is proved, one can proceed exactly as in the proof of Lemma 8 and as the end of the proof of Proposition 6, and Theorem 10 will be proved. \square

Proof of Lemma 11. For each stage t , we define the random variables ξ^t and U^t with values in $\{0, 1\}$ such that

$\xi^t = 1$ iff there are exactly $n - |S|$ players in $N \setminus (S \cup \{1\})$ that play R at stage t .
 $U^t = 1$ iff there are at least $n + 1 - |S|$ players in $N \setminus (S \cup \{1\})$ that play R at stage t .

If $U^t = 1$, the most crowded room at stage t is R . If $\xi^t = 1$, player 1's payoff is 0, and her action determines the most crowded room at stage t .

For each block number m , we also define

$$\xi_m = \frac{1}{|B^m|} \sum_{t \in B^m} \xi^t, \quad U_m = \frac{1}{|B^m|} \sum_{t \in B^m} U^t, \quad x_m = \frac{1}{|B^m|} \sum_{t \in B^m} \mathbf{1}_{\{a_t^1=R\}}.$$

Again we have $X_m \leq U_m + x_m$.

Fix a block number m where \mathcal{B}_m holds.

$(\xi^t)_{t \in B^m}$ are i.i.d. (given \mathcal{B}_m) Bernoulli random variables with expectation

$$p_m = \binom{2n - |S|}{n - |S|} \delta_m^{n-|S|} (1 - \delta_m)^n \geq \delta_m^{n-|S|},$$

for m large. Putting $Q^t = \mathbf{1}_{\{a_t^1=R\}}$ for each stage t , Lemma 5.6 of Lehrer (1990) gives

$$P \left(\left| \sum_{t \in B^m} \frac{\xi^t Q^t}{m^{10}} - p_m x_m \right| \geq \theta_m \middle| \mathcal{B}_m \right) \leq \frac{1}{|B^m| \theta_m^2} \leq \frac{1}{m^6}.$$

For some stage t in B^m , the conditional probability (given \mathcal{B}_m) that at least $n + 1 - |S|$ players in $N \setminus (S \cup \{1\})$ play R at stage t is at most

$$\binom{2n - |S|}{n + 1 - |S|} \delta_m^{n+1-|S|} = \frac{K_2}{m^2}.$$

Hence we obtain

$$\mathbb{P} \left(U_m > \frac{2K_2}{m^2} \middle| \mathcal{B}_m \right) \leq \frac{m^4}{K_2^2 |B^m|} \leq \frac{1}{m^6}.$$

Similarly, the conditional probability (given \mathcal{B}_m) that at least $n - |S|$ players in $N \setminus (S \cup \{1\})$ play R at stage t is at most

$$\binom{2n - |S|}{n - |S|} \delta_m^{n-|S|} = K_3 \left(\frac{1}{m} \right)^{\frac{2(n-|S|)}{n+1-|S|}} \leq \frac{K_3}{m}.$$

So we obtain

$$\mathbb{P} \left(\xi_m + U_m > \frac{2K_3}{m} \middle| \mathcal{B}_m \right) \leq \frac{1}{K_3^2 m^8} \leq \frac{1}{m^8}.$$

Assume that

$$U_m \leq \frac{2K_2}{m^2}, \quad \left| \sum_{t \in B^m} \frac{\xi^t Q^t}{m^{10}} - p_m x_m \right| \leq \theta_m,$$

and that there is no punishment after B^m , which implies

$$\sum_{t \in B^m} \frac{\xi^t Q^t}{m^{10}} \leq \theta_m.$$

Then $p_m x_m \leq 2\theta_m$, and

$$x_m \leq \frac{2\theta_m}{\delta_m^{n-|S|}} = 2m^{\frac{-1}{n+1-|S|}} = 2\sqrt{\delta_m}.$$

Since $X_m \leq U_m + x_m$, we obtain that $X_m \leq 3\sqrt{\delta_m}$ for m large. Hence we have

$$\begin{aligned} & \mathbb{P} \left(\mathcal{B}_m \cap \left(\left(U_m > \frac{2K_2}{m^2} \right) \cup \left(\mathcal{B}_{m+1} \cap \{X_m > 3\sqrt{\delta_m}\} \right) \cup \left(\xi_m + U_m > \frac{2K_3}{m} \right) \right) \right) \\ & \leq \mathbb{P} \left(\left(U_m > \frac{2K_2}{m^2} \right) \middle| \mathcal{B}_m \right) + \mathbb{P} \left(\left(U_m \leq \frac{2K_2}{m^2} \right) \cap \mathcal{B}_{m+1} \cap \{X_m > 3\sqrt{\delta_m}\} \middle| \mathcal{B}_m \right) \\ & \quad + \mathbb{P} \left(\left(\xi_m + U_m > \frac{2K_3}{m} \right) \middle| \mathcal{B}_m \right) \\ & \leq \frac{1}{m^6} + \frac{1}{m^6} + \frac{1}{m^8} \leq \frac{3}{m^6}. \end{aligned}$$

So with probability at least $1 - 3/m^6$, the following event holds

$$\mathcal{G}_m = \mathcal{B}_m^c \cup \left(\left(\xi_m + U_m \leq \frac{2K_3}{m} \right) \cap \left(U_m \leq \frac{2K_2}{m^2} \right) \cap \left(\mathcal{B}_{m+1}^c \cup \{X_m \leq 3\sqrt{\delta_m}\} \right) \right),$$

Assume finally that both \mathcal{G}_m and \mathcal{B}_m hold. Then

- either \mathcal{B}_{m+1} holds, and this implies that $X_m \leq 3\sqrt{\delta_m}$,
- or \mathcal{B}_{m+1}^c holds, and therefore a punishment phase starts at block B^{m+1} . We have $U_m \leq 2K_2/m^2$ and $\xi_m + U_m \leq 2K_3/m$. Consider $D = \{t \in B^m, l(a_t) = R\}$. We have $|D| \geq m^{10}\theta_m$, and $|D| \leq (\xi_m + U_m)m^{10}$, so $|D| \leq 2K_3/m^9$. The number of stages in D where player 1 may have a payoff of 1 is at most $U_m m^{10} \leq 2K_2 m^8$.

Consider a punishment block $B^{\bar{m}}$, with $\bar{m} \in \{m+1, \dots, m^2\}$, and let d be the integer such that $d \leq |B^{\bar{m}}|/|D| < d+1$. The total payoff of player 1 at this block is

$$\bar{m}^{10} X_{\bar{m}} \leq 2dK_2 m^8 + |D| = d|D| \frac{2K_2 m^8}{|D|} + |D|.$$

Hence we obtain

$$X_{\bar{m}} \leq \frac{d|D|}{B^{\bar{m}}} \frac{2K_2}{m^2\theta_m} + \frac{2K_3}{m} \leq \frac{2K_2}{m^2\theta_m} + \frac{2K_3}{m}.$$

But

$$\frac{1}{m^2\theta_m} = m^{\frac{-1}{n+1-|S|}} = \sqrt{\delta_m} \geq \frac{1}{\sqrt{m}}.$$

So for m large enough,

$$X_{\bar{m}} \leq 3K_2\sqrt{\delta_m}.$$

Hence we obtain that $\mathcal{G}_m \subset \mathcal{C}_m$. This concludes the proof of Lemma 11. \square

Remark 12. The arguments of Remark 9 can be used here, and one can easily show that σ is also an almost sure equilibrium payoff. The set of almost sure equilibrium payoffs, the set of uniform equilibrium payoffs, and the set of feasible payoffs coincide.

References

- ARTHUR, B. W. (1994) Inductive reasoning and bounded rationality. *Amer. Econom. Rev. Papers Proc.* **84**, 406–411.
- ARTHUR, B. W. (1999) Complexity and the economy. *Science* **284**, 107–109.
- CAVAGNA, A., GARRAHAN, J. P., GIARDINA, I., and SHERRINGTON, D. (1999) A thermal model for adaptive competition in a market. *Phys. Rev. Lett.* **83**, 4429–4432.
- CHALLET, D. and ZHANG, Y.-C. (1997) Emergence of cooperation and organization in an evolutionary game. *Phys. A* **246**, 407–418.
- FUDENBERG, D. and MASKIN, E. (1986) The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* **54**, 533–554.
- LEHRER, E. (1989) Lower equilibrium payoffs in two-player repeated games with nonobservable actions. *Internat. J. Game Theory* **18**, 57–89.
- LEHRER, E. (1990) Nash equilibria of n -player repeated games with semi-standard information. *Internat. J. Game Theory* **19**, 191–217.
- LEHRER, E. (1992a) On the equilibrium payoffs set of two player repeated games with imperfect monitoring. *Internat. J. Game Theory* **20**, 211–226.
- LEHRER, E. (1992b) Two-player repeated games with nonobservable actions and observable payoffs. *Math. Oper. Res.* **17**, 200–224.

- RENAULT, J. (2000) On two-player repeated games with lack of information on one side and state-independent signalling. *Math. Oper. Res.* **25**, 552–572.
- RENAULT, J. and TOMALA, T. (2000) Communication equilibrium payoffs of repeated games with imperfect monitoring. Technical Report 0034, CEREMADE, Université Paris Dauphine.
- SAVIT, R., MANUCA, R., and RIOLO, R. (1999) Adaptive competition, market efficiency, phase transition. *Phys. Rev. Lett.* **82**, 2203–2206.
- SORIN, S. (1992) Repeated games with complete information. In *Handbook of Game Theory with Economic Applications, Vol. I*, 71–107. North-Holland, Amsterdam.
- TOMALA, T. (1998) Pure equilibria of repeated games with public observation. *Internat. J. Game Theory* **27**, 93–109.
- TOMALA, T. (1999) Nash equilibria of repeated games with observable payoff vectors. *Games Econom. Behav.* **28**, 310–324.