

CORRELATED EQUILIBRIUM IN QUITTING GAMES

E. SOLAN AND R. V. VOHRA

A quitting game is a sequential game where each player has two actions: to continue or to quit. The game continues as long as all players decide to continue. The moment at least one player decides to quit, the game terminates. The terminal payoff depends on the subset of players who quit at the terminating stage. If the game continues forever, then the payoff for the players is some fixed-payoff vector.

We prove that every quitting game admits a correlated uniform ϵ -equilibrium—a uniform ϵ -equilibrium in an extended game that includes a correlation device that sends one signal to each player before start of play.

1. Introduction. There are many ways to formulate the notion of Nash equilibrium in undiscounted stochastic games. The strongest of these is *uniform ϵ -equilibrium*. A strategy profile is a *uniform ϵ -equilibrium* if for any n sufficiently large, no player could profit more than ϵ over n periods by deviating. A payoff vector is a *uniform equilibrium payoff* if it is the limit (as ϵ goes to 0) of the payoffs that correspond to a sequence of uniform ϵ -equilibrium strategy profiles. Arguments in favor of this formulation of Nash equilibria can be found in Aumann and Maschler (1995).

Existence of uniform equilibrium payoffs in n -player undiscounted stochastic games, while suspected, is still not proven. Existence for important special cases has been established; see, for example, Mertens and Neyman (1981), Vrieze and Thuijsman (1989), Vieille (2000a, b), and Solan (1999).

In all the papers referred to, the uniform ϵ -equilibrium profiles are quite elaborate and admit no easy interpretation. They are history dependent, require the players to use a battery of statistical tests to check for deviations amongst their rivals, and rely on a grim trigger to enforce compliance.

While Nash equilibrium is the most popular solution concept for a game, it is not the only one. For games in strategic form, Aumann (1974) proposes the notion of *correlated equilibria*, which are equilibria in an extended game that includes a correlation device. The device chooses a signal for each player before the start of play, and reveals to each player the signal chosen for him. A correlated equilibrium is a Nash equilibrium of the extended game. For finite games in strategic form, correlated equilibria have a number of appealing properties. They are computationally tractable. Existence is verified by checking a system of linear inequalities rather than a fixed point. The set of correlated equilibria is closed and convex.

For sequential games (ones played in stages), two generalizations of correlated equilibria are possible. One involves a correlation device that sends to each player a signal before the start of *each* round. This signal can depend on the history of past signals as well as past plays. In contrast with the problem of existence of uniform equilibrium payoff, existence of a correlated equilibrium of this kind was proved for every n -player stochastic game by Solan and Vieille (1998a).

Received October 13, 1999; revised February 21, 2001.

MSC 2000 subject classification. Primary: 91A06, 91A15, 91A20.

ORMS subject classification. Primary: Games/group decisions/stochastic.

Key words. Quitting game, absorbing game, correlated equilibrium.

The second generalization is a correlation device that sends a signal to each player just *once* at the start of the first round. It is this form of correlated equilibrium that is the focus of the paper. We show that *every* quitting game admits a correlated equilibrium payoff, where the correlation device sends to each player a single signal *before* the start of play. A *quitting game* is a sequential game in which each player has two actions: to quit or to continue. The game continues as long as all players decide to continue. The moment at least one player decides to quit, the game terminates. The terminal payoff depends on the subset of players who quit at the terminating stage. If the game continues forever, then the payoff for the players is some fixed-payoff vector.

Quitting games are a special case of absorbing games, which are a class of stochastic games. They are the simplest nontrivial stochastic games which are rich enough to confound classical proof techniques. For example, Solan and Vieille (1998b) give an instance of a four-player quitting game in which the classical approach to proving existence of a uniform equilibrium payoff fails. Furthermore, while all strategic form games with a finite number of strategies are known to admit a correlated equilibrium, this result does not extend to games with countably infinite strategy spaces. In particular, every quitting game can be represented in strategic form with countably infinite pure strategies of the type “quit in period t .”

We then generalize the result to probabilistic quitting games. Those are quitting games in which the probability that the game terminates after at least one player quits is strictly positive, but need not be one. Finally, we explain how this result can be extended to absorbing games. Since the proof of this extension is long, the interested reader is referred to Solan and Vohra (1999) for more details.

2. The model and main result. A *quitting game* is a triplet $(N, a, (r^L)_{\emptyset \neq L \subseteq N})$ where

- N is a finite set of players.
- $a \in \mathbf{R}^N$ is a daily payoff.
- For every $\emptyset \neq L \subseteq N$, $r^L \in \mathbf{R}^N$ is a payoff vector if the coalition L quits.

At each stage every player declares, independently of the others, whether he quits or continues. If all players continue the game continues to the next stage. If at least one player quits, the game terminates.

Let t_* be the termination stage and, if $t_* < \infty$, let L_* be the coalition of players who quit at stage t_* . The payoff of each player i is given by a_i if $t_* = \infty$, and $r_i^{L_*}$ if $t_* < \infty$.

A *correlation device* is a pair $((S_i)_{i \in N}, p)$ where S_i is a finite set of *signals* and p is a probability distribution over $S = \prod_{i \in N} S_i$.

Let $C = ((S_i)_{i \in N}, p)$ be a correlation device. The *extended game with correlation device* C is played as follows.

- (1) A signal profile $s = (s_i) \in S$ is chosen according to p .
- (2) Each player i is informed of s_i .
- (3) The original quitting game is played as before.

A (behavioral) *strategy for player i* in the extended game is a function $\sigma_i : S_i \times \mathbf{N} \rightarrow [0, 1]$, where $\sigma_i(s_i, n)$ is the probability that player i quits at stage n if he receives the signal s_i and the game was not terminated before stage n . Denote the strategy space of player i by $\Sigma_i(S_i)$. Note that the strategy space does not depend on the probability distribution p or on the signal spaces of the other players. Note also that if the signal spaces are all singletons, then any strategy in the extended game is also a strategy in the original game without a correlation device.

A *strategy profile* is a vector $\sigma = (\sigma_i)_{i \in N}$ of strategies. Every strategy profile σ and every $n = 1, 2, \dots, \infty$ induce a *payoff vector* $\gamma^n(\sigma) \in \mathbf{R}^N$ in the n -stage game:

$$\gamma_i^n(\sigma) = a_i \mathbf{P}_{p, \sigma}(t_* \geq n) + \sum_{\emptyset \neq L \subseteq N} r_i^L \mathbf{P}_{p, \sigma}(t_* < n, L_* = L).$$

DEFINITION 2.1. Let $\epsilon > 0$. A strategy profile σ is a *correlated (uniform) ϵ -equilibrium* in the extended game if there exists $n_0 \in \mathbf{N}$ such that for every player $i \in N$, every strategy $\tau_i \in \sum_i(S_i)$ of player i , and every $n = n_0, n_0 + 1, \dots, \infty$,

$$(1) \quad \gamma_i^n(\sigma) \geq \gamma_i^n(\sigma_{-i}, \tau_i) - \epsilon.$$

A strategy profile σ is a *(uniform) ϵ -equilibrium* if it is a correlated ϵ -equilibrium and the signal spaces are all singletons.

DEFINITION 2.2. A payoff vector $w \in \mathbf{R}^N$ is a *correlated equilibrium payoff* if for every $\epsilon > 0$ there is a correlation device C and a correlated ϵ -equilibrium σ in the extended game such that $\|\gamma^\infty(\sigma) - w\| < \epsilon$.

Note that by the monotone convergence theorem, $\lim_{n \rightarrow \infty} \gamma^n(\sigma)$ exists for every fixed-strategy profile σ , and is equal to $\gamma^\infty(\sigma)$. Hence a correlated equilibrium payoff is a correlated ϵ -equilibrium payoff in any finite but sufficiently long game.

Our first result is:

THEOREM 2.3. *Every quitting game admits a correlated uniform equilibrium payoff.*

A *probabilistic quitting game* is a quitting game in which the probability that the game terminates if at least one player quits is strictly positive, but not necessarily one.

We generalize the first result to probabilistic quitting games.

THEOREM 2.4. *Every probabilistic quitting game admits a correlated uniform equilibrium payoff.*

In §5 we explain how the existence result can be extended to absorbing games.

3. An example. Consider the following three-player quitting game, where Player 1 chooses a row, Player 2 chooses a column, and Player 3 chooses a matrix.

	Continue		Quit	
	Continue	Quit	Continue	Quit
Continue	2, 2, 0	0, 1, 3*	3, 0, 1*	1, 1, 0*
Quit	1, 3, 0*	1, 0, 1*	0, 1, 1*	0, 0, 0*

Flesch et al. (1997) studied a similar game where the payoff if everyone continues is $(0, 0, 0)$. They proved that the set of equilibrium payoffs of that game coincides with the edges of the triangle $\text{conv}\{(1, 2, 1), (1, 1, 2), (2, 1, 1)\}$. In particular, even though the payoffs in their game are symmetric, the game admits no symmetric equilibrium payoff. The set of equilibrium payoffs of our example coincides with the edges of this triangle as well.

The symmetric vector $(4/3, 4/3, 4/3) = \frac{1}{3}(1, 3, 0) + \frac{1}{3}(0, 1, 3) + \frac{1}{3}(3, 0, 1)$ is a correlated equilibrium payoff. A naive candidate for a correlation device is the following:

- The correlation device chooses a player. Each player is chosen with probability 1/3.
- The device informs the chosen player that he was chosen, and the other two players that they were not chosen.

We then define a strategy profile, in which all players follow the same strategy:

- If you were chosen, quit at Stage 1 with probability one.
- If you were not chosen, continue forever.

The construction described above is sensitive to two things. The first is the incentives that the chosen player has to never quit. The second is the payoff to an unchosen player from two players quitting at the same stage. If this were large enough, an unchosen player might profit by quitting at the first stage.

The second of these can be accommodated by masking the stage at which the chosen player quits. For example, the chosen player is told to quit in each stage with probability $\epsilon > 0$. Now the unchosen players are ignorant of who the first player is to quit as well as the stage at which he will quit.

Dissuading the chosen player from quitting at a stage other than that prescribed by the device, or continuing indefinitely, is more difficult. In our example Players 1 and 2 are better off if the game never terminates. Thus, if Player 1 is the chosen one, why should he quit? The other two players do not know that he is the chosen one. To deal with this possibility we will ensure that one of the unchosen players can punish Player 1 for his deviation. The idea is to instruct the unchosen players to continue for a certain number of rounds and then quit. To force compliance by Player 1, the payoff to Player 1 by continuing forever should be at most one.

In this example each player i has a *punisher*—a player $j \neq i$ who by quitting yields player i a low payoff. Player 1 is the punisher of Player 3, Player 2 is the punisher of Player 1, and Player 3 is the punisher of 2. A simple modification of the previous equilibrium scheme suggests itself: The device chooses a player uniformly at random to quit at the first stage, and informs his punisher that he should quit at the second stage if the chosen one has not quit at the first stage.

The flaws are obvious. First, the punisher knows who the chosen one is, and might profit by quitting on the first period too. This problem does not arise in this example. Second, the player who is neither the chosen one nor the punisher receives some information too. If Player 3 is neither the chosen one nor the punisher, he can deduce that Player 1 is the chosen one. Therefore, Player 3 would rather quit at the first stage.

To avoid these flaws the device must inform the punisher while masking the identity of the chosen one. One way of doing this is described below.

- The device chooses one player i to be the “designated quitter.” Each player is chosen with probability $1/3$.
- The device chooses a quitting stage Q , uniformly in $\{1, \dots, M\}$, where M is a sufficiently large positive integer, and sends it to the designated quitter.
- The device chooses a punishment stage P , uniformly in $\{M+1, \dots, 2M\}$, and sends it to player $i+1 \bmod 3$, the punisher of i .
- The device sends $P+1$ to the third player.

We now define a pure strategy profile in the extended game, in which all players follow the same strategy:

- Quit at the stage the device recommended.

The chosen player knows that he was chosen, since his quitting stage is at most M , whereas the quitting stages of the other two exceed M . If the chosen player does not quit, he will be punished and get zero. Moreover, the probability that he will correctly guess the quitting stage of his punisher is low. Hence, he has no reason to disobey the recommendation. With high probability the punisher and the third player receive a signal in $\{M+2, \dots, 2M\}$. In this case, the conditional probability that each is a punisher is $1/2$, and therefore the conditional probability that each of their opponents is the designated quitter is $1/2$ as well. In particular, they have no reason to deviate. Thus, this joint probability distribution over $\{1, \dots, 2M+1\}^3$ is a correlated ϵ -equilibrium, provided M is sufficiently large (in our example, provided $M \geq 2/\epsilon$).

4. The proof.

4.1. A condition on the payoffs. For every player i define the *punishment level* of i by

$$(2) \quad p_i = \min_{j \neq i} r_i^j.$$

This is the minimal amount player i may receive if a single player $j \neq i$ quits, while everyone else continues.

We say that i is *punishable* if $p_i \leq r_i^i$, and define the *punisher of i* , j_i , as a player who minimizes the right-hand side of (2). Let N_\star be the set of punishable players.

We denote by $A = \{c_i, q_i\}$ the set of actions of player i . We identify every vector $x \in [0, 1]^N$ with the stationary strategy that indicates each player i to quit at every stage with probability x_i . We identify the pure action c_i (continue) with the mixed action $x_i = 0$ and the pure action q_i (quit) with the mixed action $x_i = 1$. Let $\rho = \max_{L,i} |r_i^L|$. Assume w.l.o.g. that $\rho \geq 1$.

LEMMA 4.1. *Either the game admits a stationary ϵ -equilibrium for every $\epsilon > 0$ sufficiently small, or there exists a probability distribution ν over N that satisfies (i) $\text{supp}(\nu) \subseteq N_\star$, (ii) $|\text{supp}(\nu)| \geq 2$, and (iii) $\sum_{j \in N} \nu_j r_j^j \geq r_i^i \forall i \in N$.*

PROOF.

Case 1. $a_i \geq r_i^i$ for every player i . The stationary strategy profile x that is defined by $x_i = 0 \forall i \in N$, i.e., no player ever quits, is a 0-equilibrium. Indeed, $\gamma_i^n(x) = a_i$ for each i and every $n \in \mathbf{N} \cup \{\infty\}$, and for every pure deviation y_i of player i , $\gamma_i^n(x_{-i}, y_i)$ is either a_i or $r_i^i \leq a_i$.

Case 2. $N_\star = \emptyset$. If Case 1 does not hold, there is a player i with $a_i < r_i^i$. For every fixed $\epsilon \in (0, 1)$ the following stationary strategy profile y is an ϵ -equilibrium:

$$y_j = \begin{cases} \epsilon/2\rho & j = i, \\ 0 & \text{otherwise.} \end{cases}$$

Indeed, since N_\star is empty, $r_j^j > r_j^i$ for every $j \neq i$. Hence (1) holds w.r.t. y , for every $n \in \mathbf{N} \cup \{\infty\}$ such that $(1 - \epsilon/2\rho)^n < \epsilon/2$.

Define a new quitting game \tilde{G} , which is similar to G , except that the daily payoff is the punishment level $p = (p_i)_{i \in N}$ rather than a .

For every discount factor $\lambda \in (0, 1)$ and every stationary profile x , let

$$\gamma_i^\lambda(x) = \mathbf{E}_x \left[\lambda \sum_{n=1}^\infty (1 - \lambda)^{n-1} (p_i 1_{n \leq t_\star} + r_i^{L_\star} 1_{n > t_\star}) \right]$$

be the expected λ -discounted payoff for player i under x in \tilde{G} . Let $x^\lambda \in [0, 1]^N$ be a stationary λ -discounted equilibrium of \tilde{G} , which exists by Fink (1964), and let $g^\lambda = \gamma^\lambda(x^\lambda) \in \mathbf{R}^N$ be the corresponding λ -discounted equilibrium payoff. By taking a subsequence, we assume that $x^0 = \lim_{\lambda \rightarrow 0} x^\lambda$ and $g^0 = \lim_{\lambda \rightarrow 0} g^\lambda$ exist. Using a similar argument we will assume in the sequel that (finitely many) other limits exist as well.

Case 3. $x_i^0 > 0$ for at least two players. It is well known that in this case x^0 is a terminating stationary 0-equilibrium of \tilde{G} (see, e.g., Vrieze and Thuijsman 1989 or Solan 1999). Since $x_i^0 > 0$ for at least two players, x^0 is a terminating stationary 0-equilibrium of G as well. Indeed, the game terminates with probability one even when a single player deviates, and the termination payoffs are the same in G and \tilde{G} .

Case 4. $x_i^0 > 0$ for exactly one player i . In this case, $g^0 = r^i$.

First we claim that i is punishable. Indeed, since $x_j^0 = 0$ for every $j \neq i$, it follows that as $\lambda \rightarrow 0$, the probability that $|L_\star| \leq 1$ under the stationary strategy profile (x_{-i}^λ, c_i) converges to one. In particular, $\lim_{\lambda \rightarrow 0} \gamma_i^\lambda(x_{-i}^\lambda, c_i)$ is in the convex hull of p_i and $(r_j^j)_{j \neq i}$. Since $p_i \leq r_j^j$ for every $j \neq i$, and by the equilibrium condition,

$$p_i \leq \lim_{\lambda \rightarrow 0} \gamma_i^\lambda(x_{-i}^\lambda, c_i) \leq \lim_{\lambda \rightarrow 0} g_i^\lambda = g_i^0 = r_i^i,$$

and i is punishable. For every fixed $\epsilon \in (0, 1)$ the following stationary strategy y is an ϵ -equilibrium:

$$(3) \quad y_j = \begin{cases} x_i^0 & j = i, \\ x_i^0 \epsilon / 2\rho & j = j_i, \\ 0 & \text{otherwise.} \end{cases}$$

Indeed, (1) is satisfied w.r.t. y for every $n \in \mathbf{N} \cup \{\infty\}$ such that $(1 - x_i^0 \epsilon / 2\rho)^n < \epsilon / 2$.

We assume from now on that $x_i^0 = 0$ for every player i .

It is well known (see, e.g., Vrieze and Thuijsman 1989 or Solan 1999) that

$$(4) \quad g^0 = \alpha p + (1 - \alpha) \sum_{j \in \mathbf{N}} v_j r^j,$$

where

$$\alpha = \lim_{\lambda \rightarrow 0} \frac{\lambda}{(\lambda + (1 - \lambda)(1 - \prod_{j \in \mathbf{N}} (1 - x_j^\lambda)))}, \quad \text{and} \quad v_i = \lim_{\lambda \rightarrow 0} \frac{x_i^\lambda}{\sum_{j \in \mathbf{N}} x_j^\lambda}.$$

Note that $\nu = (\nu_i)_{i \in \mathbf{N}}$ is a probability distribution.

By the equilibrium condition, and since $x_i^0 = 0$ for every $i \in \mathbf{N}$,

$$(5) \quad r_i^i = \lim_{\lambda \rightarrow 0} \gamma_i^\lambda(x_{-i}^\lambda, q_i) \leq \lim_{\lambda \rightarrow 0} g_i^\lambda = g_i^0.$$

Moreover, if $x_i^\lambda > 0$ for every λ sufficiently small we get, as in Case 4,

$$(6) \quad p_i \leq \lim_{\lambda \rightarrow 0} \gamma_i^\lambda(x_{-i}^\lambda, c_i) \leq \lim_{\lambda \rightarrow 0} g_i^\lambda = \lim_{\lambda \rightarrow 0} \gamma_i^\lambda(x_{-i}^\lambda, q_i) = r_i^i,$$

and in particular i is punishable. Note that if $\nu_i > 0$, then $x_i^\lambda > 0$ for every λ sufficiently small; hence i is punishable.

Case 5. $\nu_i = 1$ for some player i . By (6) player i is punishable and one can verify that for every fixed $\epsilon \in (0, 1)$, the stationary strategy y defined in (3) with x_i^0 replaced by $\epsilon / 2\rho$ is an ϵ -equilibrium.

Case 6. $\alpha = 1$. In this case $g^0 = p$. For every $i, j \in \mathbf{N}$ such that $i \neq j$,

$$r_i^j \geq p_i = g_i^0 = \lim_{\lambda \rightarrow 0} g_i^\lambda \geq \lim_{\lambda \rightarrow 0} \gamma_i^\lambda(x_{-i}^\lambda, q_i) = r_i^i.$$

If Case 2 is not satisfied, $N_* \neq \emptyset$. The stationary strategy y that was defined in Case 5 where i is any punishable player is an ϵ -equilibrium.

Case 7. All previous cases do not hold. As mentioned above, player i is punishable once $\nu_i > 0$; hence the first condition holds. If Case 5 is not satisfied, the second condition holds as well. For every nonpunishable player i , $r_i^i < p_i \leq r_i^j$ for every $j \neq i$; hence the third condition holds for nonpunishable players. For every punishable player i , $p_i \leq r_i^j$ for every $j \in \mathbf{N}$. By (5), (4), and since $\alpha < 1$ (if Case 6 is not satisfied),

$$r_i^i \leq g_i^0 = \alpha p_i + (1 - \alpha) \sum_{j \in \mathbf{N}} v_j r_i^j \leq \sum_{j \in \mathbf{N}} v_j r_i^j,$$

and the third condition holds for punishable players as well. \square

4.2. A correlated ϵ -equilibrium.

PROOF. If for every $\epsilon > 0$ sufficiently small the game admits a stationary ϵ -equilibrium, we are done. Otherwise, by Lemma 4.1 there exists a probability distribution ν over punishable players, such that for every player i , $\sum_{j \in N} \nu_j r_i^j \geq r_i^i$ and $\nu_i < 1$.

Fix an $\epsilon > 0$, and let $M > 1/\epsilon$ be some fixed integer. Define a correlation device with signal spaces $S_i = \{1, \dots, 2M + 1\}$ as follows.

- A designated quitter $i \in N$ is chosen according to the probability distribution ν .
- A quitting stage Q is chosen, uniformly from the set $\{1, \dots, M\}$.
- A punishment stage P is chosen, uniformly from the set $\{M + 1, \dots, 2M\}$.
- Player i receives the signal Q , his punisher j_i receives the signal P , and each player $k \neq i, j_i$ receives the signal $P + 1$.

Define a pure strategy profile σ in the extended game, in which all players follow the same strategy:

- If you received the signal s , quit at stage s , and continue in all other stages.

If all players follow σ , then the expected payoff is $\sum_{i \in N} \nu_i r^i$. We will now show that no player can profit too much by deviating to another pure strategy.

Note that the designated quitter i knows he is the designated quitter, since only his signal is at most M . If the designated quitter deviates to a pure strategy that quits for the first time at stage n , instead of quitting at stage Q , his expected payoff is r_i^i if $n < P$, and $r_i^i \leq r_i^j$ if $n > P$. Since the probability that $n = P$ is below ϵ , this means that this deviation can increase his expected payoff by at most $\epsilon\rho$.

Fix a player k who is not the designated quitter. If $M + 1 < P < 2M$ (which happens with probability at least $1 - 2\epsilon$), then conditional on k 's signal, the probability that each player $l \neq k$ is the designated quitter is $\nu_l/(1 - \nu_k)$. In particular, if $M + 1 < P < 2M$, k 's expected payoff conditional on his signal is $\sum_{l \neq k} (\nu_l/(1 - \nu_k)) r_k^l \geq \sum_{l \in N} \nu_l r_k^l$.

We now show that if $M + 1 < P < 2M$, player k cannot profit too much by deviating to a pure strategy that quits for the first time at stage n . Indeed, if $n > Q$ his deviation does not affect his payoff, the probability that $n = Q$ is at most $1/M < \epsilon$, and his payoff if $n < Q$ is $r_k^k \leq \sum_{l \in N} \nu_l r_k^l$. Thus, with probability at least $1 - 2\epsilon$ he cannot profit more than $\epsilon\rho$. Theorem 2.3 is proved. \square

4.3. Probabilistic quitting games.

PROOF OF THEOREM 2.4. In quitting games it suffices for one player to get the signal that he is the designated quitter. In probabilistic quitting games, the game can continue even if the designated quitter quits. Hence, another player needs to be a designated quitter. Since signals are sent only before the start of play, this player needs to know in advance that if the game is not terminated by the first designated quitter, he should do the job.

Lemma 4.1 is still valid in this more general setting. (The definition of ν in this case is slightly different. ν_i is the limit, as λ goes to 0, of the probability that $L_\lambda = \{i\}$ under x^λ .) Assume that for every $\epsilon > 0$ sufficiently small, the game does not admit a stationary ϵ -equilibrium.

Fix $\epsilon > 0$ sufficiently small, and let $M > 1/\epsilon$ be some fixed integer. Let $s(1), s(2), \dots, s(K)$ be a sequence of independent outcomes of the correlation device described in the proof of Theorem 2.3, where K is sufficiently large. Thus, for each k , $s(k) \in \{1, \dots, 2M + 1\}^N$. The signal of player i is the sequence $(s^i(1), s^i(2), \dots, s^i(K))$.

We now construct a strategy profile σ in the extended game in which the players play in rounds. The k th round never ends if no player quit in the first M stages of the round. Otherwise, it ends the moment at least one player quits. After the end of round K , the players play arbitrarily.

We now restrict ourselves to the k th round, for $k = 1, \dots, K$. At the beginning of the round the players forget the history before the round started, and they use the signal $s(k)$.

Player i , who received the signal $s^i(k)$, plays as follows in the k th round:

- If $s^i(k) \leq M$ (player i is the k th designated quitter), he quits at stage $s^i(k)$ of the round.
- If $s^i(k) > M$ (player i is not the k th designated quitter), he continues until stage $s^i(k)$ of the round.
- If some player quit for the first time in this round at stage $s^i(k) - 1$ (this player is the k th punisher), player i continues forever.
- Otherwise (player i is the k th punisher), he quits at stage $s^i(k)$ with probability one, and in each subsequent stage he quits with probability ϵ .

Some points are worth observing.

First, if some player quits before stage $M + 1$ and the game does not terminate, the k th round ends. Thus, if a player who is not the designated quitter quits before stage $M + 1$, the players treat him as if he was the designated quitter: Only the designated quitter knows that that player deviated, but he has no means of transmitting this information to the other players. Similarly, if two players quit at the same stage and the game does not terminate, the play continues as if no deviation occurred.

Second, the punisher quits with probability ϵ at every stage that follows the punishment stage, rather than with probability one, so that the designated quitter would not know when the punisher will actually quit, and use this information to profit by quitting exactly at the same stage.

Third, the designated quitter i may profit by pretending he is the punisher (e.g., he quits at stage $M + 1$ and continues in all other stages; this is profitable if $a_i > \sum_{j \in N} \nu_j r_j^i$). Therefore, each player l checks whether the player who acts as if he is the punisher indeed quits at the punishment stage (given player l 's information, if l is not the designated quitter the punishment stage is either $s^l(k) - 1$ or $s^l(k)$).

Fourth, if K is sufficiently large, the game terminates with probability at least $1 - \epsilon$ before the end of round K .

Assume M is sufficiently large so that the probability that $s(1), \dots, s(K)$ reveal no information is at least $1 - \epsilon$. We now argue that, conditional on termination before round K and on $s(1), \dots, s(K)$ revealing no information, no player can profit too much by a unilateral deviation.

Under σ , the expected payoff of player i conditional on termination in round k , is r_i^i if i is the k th designated quitter, and $\sum_{j \neq i} (\nu_j / (1 - \nu_i)) r_j^i \geq r_i^i$ otherwise.

It is sufficient to prove that no deviation τ^i of player i that coincides with σ^i until round k is pure in round k , and quits for the first time in that round at stage n , cannot increase the payoff of i too much. n may be greater than $2M + 1$, which means player i does not quit at that round. Let i_k be the k th designated quitter, and Q_k (resp. P_k) the quitting (resp. punishment) stage at the k th round.

Assume first that $i = i_k$ is the designated quitter at round k . If $n \leq M$, i 's payoff remains unchanged. If $M < n < Q_k - 1$ player i either receives r_i^i if the game terminates by him quitting, or he is punished. If $n > Q_k$, then i does not profit, as he is punished in round k , while the probability that $n \in \{Q_k, Q_k - 1\}$ is low.

Assume now that $i \neq i_k$. If $n < P_k$, then i does not profit, since $r_i^i \leq \sum_{j \neq i} (\nu_j / (1 - \nu_i)) r_j^i$. The probability that $n = P_k$ is small, while if $n > P_k$, his expected payoff remains unchanged. \square

5. Absorbing games. Absorbing games are a generalization of quitting games: The set of actions of each player i is some arbitrary finite set A^i , and for each entry a of the multidimensional matrix $A = \times_{i \in N} A^i$ corresponds a daily payoff if this entry is chosen by the players, a probability of termination if that entry is played, and a terminal payoff if the game is terminated through this entry.

We now explain the main ideas behind the construction of the correlation mechanism in absorbing games. The proof itself is long, and the interested reader is referred to Solan and Vohra (1999).

An *absorbing entry* is an entry where the probability of termination is strictly positive. A mixed-action profile is *nonabsorbing* if with probability zero an absorbing entry is played.

In this more general setting, if for every $\epsilon > 0$ sufficiently small there is no uniform ϵ -equilibrium, one can still find a nonabsorbing mixed-action x and a probability distribution ν over absorbing entries such that an analogue of Lemma 4.1 holds.

Some of those absorbing entries require a single player to play some other action than what is prescribed by x , while some of those absorbing entries require that more than one player would play some other action than the one prescribed by x . Forcing the players to use entries of the second type can be done by using standard statistical tests (see, e.g., Vieille 2000b or Solan 1999). Forcing the players to use entries of the first type is more difficult.

Our construction suffers from a basic flaw in the general setting. Given a nonabsorbing mixed-action x , a player may have several punishing actions: one action to punish one player, and another action to punish another player. Since signals are sent only before the start of play, the player should know which action to use if he is the punisher and the designated quitter does not quit. If he has two punishing actions, then the knowledge of the punishing action he should use may reveal some information on the identity of the designated quitter, which may tell the player he is better off by deviating himself.

This difficulty is dealt with in two ways.

If x is actually pure, then by assuming that the payoff function is generic, we may restrict each player i to punish his opponents using only *one* action: the action which is his best absorbing response to x_{-i} . That is, it is the unique action a_i that maximizes i 's terminal payoff under (x_{-i}, a_i) .

If x is not pure, then $\text{supp}(x_{i_0}) \geq 2$ for some player i_0 . In this case, player i_0 can send public information to all other players: To send a bit of information (0 or 1) he plays either x_{i_0} or some fixed small perturbation of x_{i_0} for a long but finite period. If the period is long enough, with high probability the players can read correctly the bit that was sent. By dividing messages into bits, player i_0 can send any (finite) message to the other players. We call such a player a *signaller*.

While signalling player i_0 plays either x_{i_0} or a small perturbation of that mixed action; hence, if some other player $j \neq i_0$ deviates while player i_0 transmits information, player j 's payoff would be close to his payoff if player i_0 played stationary x_{i_0} . Since such a deviation is detected immediately, one can ensure that it is not profitable.

One way to use this public signalling mechanism is as follows. Before start of play, the device chooses, in addition to the designated quitter i , a permutation π over the players (uniformly from all such permutations) and a verification key V which is chosen uniformly from $\{1, \dots, M\}$ where $M > 1/\epsilon$. It then sends to each player $k \neq i_0$ the permutation π and the verification key V . Player i_0 receives $\pi(i)$, and, if he is *not the designated quitter*, also V .

We then instruct the players to play as follows:

Quitting Phase.

- The designated quitter is supposed to quit prior to stage M .

If at stage M the designated quitter has not quit, his identity should be revealed.

Revelation Phase.

- Player i_0 publicly sends V and $\pi(i)$.

If player i_0 is the designated quitter, he does not know V , and the chance that he can mimic V is low. Hence, with high probability, i_0 's identity as the designated quitter is revealed. If player i_0 is not the designated quitter, he knows V ; hence, everyone knows he is not the

designated quitter, and hence that he plays truthfully. Since all players but i_0 know π , with the additional information of $\pi(i)$ they know the identity of the designated quitter.

Note that i_0 does not know the identity of the designated quitter.

Punishment Phase.

- If the punisher j is not i_0 , player j punishes player i during the following M stages.
- If after M stages no one punished the designated quitter, player i_0 infers that he is the punisher, and he punishes the designated quitter.

If there is only one signaller then the assumption that the game is generic implies that i_0 has a unique absorbing best response to x_{-i_0} , hence, i_0 does not need to know the identity of the designated quitter.

If there are two signallers, then the identity of the designated quitter can be revealed to everyone: The first signaller reveals the identity to everyone but himself, and then the second signaller reveals the identity to the first signaller (after clearing himself from being the designated quitter).

The description provided here is very simplistic: It illuminates the main ideas, but ignores many difficulties that arise in the analysis. In particular, the definition of the punishment level is more involved in the general setting, and more types of signalling are required. A complete analysis can be found in Solan and Vohra (1999).

References

- Aumann, R. J. 1974. Subjectivity and correlation in randomized strategies, *J. Math. Econ.* **1** 67–96.
- , B. M. Maschler, 1995. *Repeated Games with Incomplete Information*, MIT Press, Cambridge, MA.
- Fink, A. M. 1964. Equilibrium in a stochastic n -person game. *J. Sci. Hiroshima Univ.* **28** 89–93.
- Flesch J., F. Thuijsman, K. Vrieze. 1997. Cyclic Markov equilibria in stochastic games. *Internat. J. Game Theory* **26** 303–314.
- Mertens J. F., A. Neyman. 1981. Stochastic games. *Internat. J. Game Theory* **10** 53–66.
- Solan, E. 1999. Three-person absorbing games. *Math. Oper. Res.* **24** 669–698.
- , N. Vieille. 1998a. Correlated equilibrium in stochastic games. D.P. #1226, The Center for Mathematical Studies in Economics and Management Science, Northwestern University, Evanston, IL and Cahiers du CEREMADE #9848.
- , ———. 1998b. Quitting games. D.P. #1227, The Center for Mathematical Studies in Economics and Management Science, Northwestern University, Evanston, IL and Cahiers du CEREMADE #9846.
- , R. Vohra, 1999 Correlated equilibrium, public signalling and absorbing games. D.P. #1272, The Center for Mathematical Studies in Economics and Management Science, Northwestern University, Evanston, IL.
- Vieille, N. 2000a. Equilibrium in 2-person stochastic games I: A reduction. *Israel J. Math.* **119** 55–91.
- . 2000b. Equilibrium in 2-person stochastic games II: The case of recursive games. *Israel J. Math.* **119** 93–126.
- Vrieze, O. J., F. Thuijsman, 1989. On equilibria in repeated games with absorbing states. *Internat. J. Game Theory* **18** 293–310.

E. Solan: Department of Managerial Economics and Decision Sciences, Kellogg Graduate School of Management, Northwestern University, Evanston, Illinois 60208; School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel; e-mail: eilons@post.tau.ac.il

R. Vohra: Department of Managerial Economics and Decision Sciences, Kellogg Graduate School of Management, Northwestern University, Evanston, Illinois 60208; e-mail: r-vohra@nwu.edu