

Repeated games with endogenous choice of information mechanisms

Citation for published version (APA):

Flesch, J., & Perea y Monsuwé, A. (1999). Repeated games with endogenous choice of information mechanisms. *Mathematics of Operations Research*, 24(4), 785-794.
<https://doi.org/10.1287/moor.24.4.785>

Document status and date:

Published: 01/01/1999

DOI:

[10.1287/moor.24.4.785](https://doi.org/10.1287/moor.24.4.785)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.



Mathematics of Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Repeated Games with Endogenous Choice of Information Mechanisms

János Flesch, Andres Perea Y. Monsuwé,

To cite this article:

János Flesch, Andres Perea Y. Monsuwé, (1999) Repeated Games with Endogenous Choice of Information Mechanisms. Mathematics of Operations Research 24(4):785-794. <https://doi.org/10.1287/moor.24.4.785>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

© 1999 INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

REPEATED GAMES WITH ENDOGENOUS CHOICE OF INFORMATION MECHANISMS

JÁNOS FLESCH AND ANDRES PEREA Y MONSUWÉ

We consider two-player repeated games with nonobservable actions (cf. Lehrer 1989). An information mechanism for a player is a function which assigns a private signal to every action-pair of the one-shot game. In this paper, we extend the model to a situation in which both players can buy an information mechanism before playing the repeated game. Within this model, we provide a characterization of the lower equilibrium payoffs in terms of the one-shot game for the case that both players choose a nontrivial information mechanism with probability one. Moreover, we construct a lower equilibrium in a repeated game in which one of the players strictly randomizes between information mechanisms. It is shown that the corresponding payoffs cannot be induced by a lower equilibrium in which players choose a particular information mechanism with probability one.

1. Introduction. In this paper, we study two player repeated games with nonobservable actions in which both players select an information mechanism before the actual game starts. Here, an information mechanism is a function assigning a private signal to each pair of actions chosen in the one-shot game. The mechanism chosen by a player is not observed by his opponent. We assume that the costs for the information mechanisms are fixed and known to both players. Furthermore, the players always have the ability to buy the trivial information mechanism (revealing no information about the opponent's actions) at price zero. The model can therefore be seen as an extension of the traditional model (with fixed information mechanisms) as studied in Lehrer (1989, 1990, 1991, 1992a, b), Fudenberg and Levine (1989) and Radner (1986), to name just a few. Hillas and Liu (1995) investigate the case in which signals are stochastic instead of deterministic.

We first focus on strategies in which both players choose one specific information mechanism with probability one. Such strategies are called *simple* and lower equilibria (cf. Lehrer 1989) consisting of simple strategies are called *simple lower equilibria*. Our main result provides a characterization of the class of simple lower equilibrium payoffs by means of the one-shot game. As such, it can be viewed as an extension of Lehrer's result to situations where information mechanisms are not fixed but are chosen endogenously. Roughly speaking, the characterization states that the simple lower equilibrium payoffs coincide with the individually rational payoffs generated by some correlation on the one-shot action pairs, such that no player can benefit from switching to another information mechanism and subsequently transforming actions, without being detected.

Another interesting situation would be the case where both players are allowed to randomize over the different information mechanisms that can be chosen. However, this situation is much more complex since both players are uncertain about the mechanism that is used by the opponent. We provide an example in which we construct a lower equilibrium where one of the players strictly randomizes over information mechanisms at the beginning of the game. We prove that the corresponding payoff cannot be sustained by a simple lower equilibrium.

Received April 23, 1997; revised June 20, 1998 and May 10, 1999.

AMS 1991 subject classification. Primary: 90D20, 90D05.

OR/MS subject classification. Primary: Games/group decisions.

Key words. Repeated games, imperfect monitoring, lower equilibria.

2. Model. We consider a two player game in which both players, before playing a repeated game with nonobservable actions, choose an information mechanism. Formally, a *two-player repeated game with nonobservable actions* is a triple $\Gamma = \langle \Gamma_N, l_1, l_2 \rangle$ where

(1) Γ_N is a two-player normal form game with finite action sets A_1, A_2 and payoff-functions v_1, v_2 , and

(2) $l_i: A_1 \times A_2 \rightarrow L_i$ is the so-called *information mechanism* for player i which assigns to every pair of actions a signal in the signals set L_i . The functions l_1, l_2 should be such that $l_1(a_1, a_2) \neq l_1(a'_1, a'_2)$ whenever $a_1 \neq a'_1$ and $l_2(a_1, a_2) \neq l_2(a'_1, a'_2)$ whenever $a_2 \neq a'_2$.

The latter condition on the information mechanisms means that a player always observes his own actions. The information mechanism l_1 is called *trivial* if player 1 only observes his own actions. Formally, this means that

$$l_1(a_1, a_2) = l_1(a_1, a'_2) \quad \text{for all } a_1 \in A_1 \quad \text{and all } a_2, a'_2 \in A_2.$$

We call the information mechanism l_1 *perfect* if it reveals all the opponent's actions, i.e.,

$$l_1(a_1, a_2) \neq l_1(a'_1, a'_2) \quad \text{whenever } a_2 \neq a'_2.$$

In the remainder of this article, a two-player repeated game with nonobservable actions will simply be called a repeated game.

By \mathcal{L}_1 and \mathcal{L}_2 we denote the collections of information mechanisms which can be bought by player 1 and 2 respectively. Each information mechanism $l_i \in \mathcal{L}_i$ can be bought at a nonnegative price $c(l_i)$. We assume that \mathcal{L}_1 and \mathcal{L}_2 contain the trivial information mechanism which can be bought at price zero. If the information mechanisms l_1 and l_2 are chosen, the players enter the corresponding repeated game $\Gamma(l_1, l_2)$. Depending on his information mechanism l_i , player i chooses a behavior strategy $\sigma_i(l_i)$ in the repeated game which follows. The new, larger game obtained in this way is called the *extended game*.¹ A *behavior strategy* for player i in the extended game is a pair $\phi_i = (\tau_i, \sigma_i)$ where τ_i is a probability distribution on \mathcal{L}_i and σ_i is a function assigning to every information mechanism l_i in \mathcal{L}_i a corresponding behavior strategy $\sigma_i(l_i)$ for the repeated game. In order to formalize the notion of the behavior strategy $\sigma_i(l_i)$, consider a repeated game which follows after choosing l_i . An n -stage *history* for player i is a sequence $(l_i^1, l_i^2, \dots, l_i^n)$ of player i signals. The set of n -stage histories for player i is denoted by L_i^n . A *behavior strategy* for player i in the repeated game is a sequence $\sigma_i = (\sigma_i^n)_{n=1}^\infty$ where $\sigma_i^1 \in \Delta(A_i)$ and $\sigma_i^n: L_i^{n-1} \rightarrow \Delta(A_i)$ for $n = 2, 3, \dots$.

Let (σ_1, σ_2) be a behavior strategy pair in the repeated game $\Gamma(l_1, l_2)$. For every n and every n -stage history pair (h_1, h_2) , $\mathbb{P}_{\sigma_1, \sigma_2}(h_1, h_2)$ denotes the probability that (h_1, h_2) will occur. The expected average payoff for player i at stage n is given by

$$v_i^n(\sigma_1, \sigma_2) := \frac{1}{n} \sum_{k=1}^n \left[\sum_{(h_1, h_2) \in L_1^{k-1} \times L_2^{k-1}} \mathbb{P}_{\sigma_1, \sigma_2}(h_1, h_2) v_i(\sigma_1^k(h_1), \sigma_2^k(h_2)) \right],$$

where L_i^0 is some arbitrary single-element set. By $v_i(\sigma_1^k(h_1), \sigma_2^k(h_2))$, we denote the

¹ Another possibility would be to allow the players to choose an information mechanism at each stage of the game, before choosing the action. As such, the model would reduce to a special case of a repeated game with unobservable actions as studied by Lehrer (1989, 1992). In this case, a player could check the opponent's actions by choosing a nontrivial mechanism only at a set of stages with density zero and therefore (on average) not paying anything for this. This is not possible in our model: if a player wants to check the opponent by using a nontrivial mechanism then he has to pay the full price in advance. For an example illustrating this difference, see Flesch and Perea y Monsuwé (1999).

expected payoff for player i in the one-shot game, if the mixed actions $\sigma_1^k(h_1)$ and $\sigma_2^k(h_2)$ are played. We define

$$\underline{v}_i(\sigma_1, \sigma_2) = \liminf_{n \rightarrow \infty} v_i^n(\sigma_1, \sigma_2) \quad \text{and} \quad \bar{v}_i(\sigma_1, \sigma_2) = \limsup_{n \rightarrow \infty} v_i^n(\sigma_1, \sigma_2).$$

The payoffs $\underline{v}_i(\sigma_1, \sigma_2)$ and $\bar{v}_i(\sigma_1, \sigma_2)$ are called lower payoff and upper payoff respectively.

Now, suppose that the players play the behavior strategy pair (ϕ_1, ϕ_2) in the extended game. The lower (upper) expected payoff for player i is equal to the difference between the lower (upper) expected payoff in the repeated game and the expected amount to be paid for the information mechanism. We denote the lower and upper expected payoff for player i by $\underline{w}_i(\phi_1, \phi_2)$ and $\bar{w}_i(\phi_1, \phi_2)$ respectively. The behavior strategy pair (ϕ_1, ϕ_2) is called a *lower equilibrium* in the extended game if

- (1) $\underline{w}_i(\phi_1, \phi_2) = \bar{w}_i(\phi_1, \phi_2)$ for $i = 1, 2$, and
- (2) $\underline{w}_i(\phi'_1, \phi_2) \leq \underline{w}_i(\phi_1, \phi_2)$ for all ϕ'_1 and $\underline{w}_2(\phi_1, \phi'_2) \leq \underline{w}_2(\phi_1, \phi_2)$ for all ϕ'_2 .

A behavior strategy pair in the extended game is called *simple* if both players but probability one on a specific information mechanism. A lower equilibrium consisting of simple strategies is termed *simple lower equilibrium*.

3. Main result.

3.1. The theorem. In the remainder of this article, we always refer to the extended game when we talk about strategies and equilibria unless we say otherwise. By $\text{SLEP}(l_1, l_2)$ we denote the set of payoffs induced by simple lower equilibria in which the information mechanisms l_1 and l_2 are chosen with certainty. In order to give a characterization of the set $\text{SLEP}(l_1, l_2)$ we need some more definitions. In these definitions, l_1 and l_2 are fixed.

Let l'_1 be an arbitrary information mechanism. In the following definitions we compare pairs $(a_1|l_1)$ consisting of an action a_1 and the information mechanism l_1 with pairs $(a'_1|l'_1)$. A pair $(a'_1|l'_1)$ is called *more informative* than $(a_1|l_1)$, denoted by $(a'_1|l'_1) \triangleright (a_1|l_1)$, if $l_1(a_1, a'_2) \neq l_1(a_1, a'_2)$ implies $l'_1(a'_1, a_2) \neq l'_1(a'_1, a'_2)$ for all a_2, a'_2 . In words, all actions of player 2 that are distinguished by playing action a_1 while using information mechanism l_1 are also distinguished by playing a'_1 while using l'_1 . A pair $(a'_1|l'_1)$ is called *greater* than $(a_1|l_1)$, denoted by $(a'_1|l'_1) \succ^{l_2} (a_1|l_1)$, if $a'_1 \sim^{l_2} a_1$ and $(a'_1|l'_1) \triangleright (a_1|l_1)$. Here, $a'_1 \sim^{l_2} a_1$ means that a'_1 and a_1 are indistinguishable if player 2 uses l_2 . An information mechanism l'_1 is called *greater* than l_1 with respect to l_2 , denoted by $l'_1 \succ^{l_2} l_1$, if for every action a_1 there exists an action a'_1 such that $(a'_1|l'_1) \succ^{l_2} (a_1|l_1)$.

Let \mathcal{B} be the set of probability distributions on $A_1 \times A_2$, which from now on will be called weight matrices. For an information mechanism l'_1 , let $\Psi_1(l'_1)$ be the set of functions $\psi_1: A_1 \rightarrow A_1$ with

$$(\psi_1(a_1)|l'_1) \succ^{l_2} (a_1|l_1)$$

for all $a_1 \in A_1$. In words, if player 1 switches from l_1 to l'_1 , then $\Psi_1(l'_1)$ contains those functions which transform each action into a greater one. Note that $\Psi_1(l'_1)$ is nonempty if and only if $l'_1 \succ^{l_2} l_1$. For a $B \in \mathcal{B}$ and a function $\psi_1 \in \Psi_1(l'_1)$, the weight matrix $\psi_1(B)$ is defined by

$$(\psi_1(B))(a_1, a_2) = \sum_{\bar{a}_1 \in \psi_1^{-1}(a_1)} B(\bar{a}_1, a_2)$$

for all a_1, a_2 . Note that, by definition, $(\psi_1(B))(a_1, a_2) = 0$ if $\psi_1^{-1}(a_1)$ is empty. It can be

seen easily that the function ψ_1 transforming B into $\psi_1(B)$ is a linear (and therefore convex) operator. Furthermore, we define

$$\mathcal{B}_1(l'_1) = \{B \in \mathcal{B} \mid v_1(B) - c(l_1) \geq v_1(\psi_1(B)) - c(l'_1) \text{ for all } \psi_1 \in \Psi_1(l'_1)\}.$$

Intuitively, $\mathcal{B}_1(l'_1)$ contains those weight matrices for which player 1 cannot increase his payoff by deviating to information mechanism l'_1 and transforming his actions into greater ones. By construction, $\mathcal{B}_1(l'_1) = \mathcal{B}$ if $l'_1 \not\succ^{l_2} l_1$ since in this case $\Psi_1(l'_1)$ is empty. In the same way, we define $\mathcal{B}_2(l'_2)$.

THEOREM 3.1. *If l_1 and l_2 are nontrivial,*

$$SLEP(l_1, l_2) = \left[\left[\bigcap_{l'_1 \in \mathcal{L}_1} v(\mathcal{B}_1(l'_1)) \right] \cap \left[\bigcap_{l'_2 \in \mathcal{L}_2} v(\mathcal{B}_2(l'_2)) \right] - (c(l_1), c(l_2)) \right] \cap IR.$$

Here, $v(\mathcal{B}_i(l'_i))$ denotes the set of payoff pairs (v_1, v_2) induced by weight matrices in $\mathcal{B}_i(l'_i)$. The notation $-(c(l_1), c(l_2))$ means that we subtract the cost-pair $(c(l_1), c(l_2))$ from every element of the set. By IR we denote the set of individually rational payoff pairs of the one-shot game. The proof can be found in §3.3. Note that the set $SLEP(l_1, l_2)$ can be empty for some l_1, l_2 .

3.2. Detectability of deviations to other information mechanisms. In Lehrer's characterization of lower equilibrium payoffs with fixed information mechanisms (Lehrer 1989), a central role is played by checking the opponent's actions. This is used as a tool to prevent the opponent from deviating. Roughly speaking, Lehrer implements this by building in test phases where a player is asked to report the signals he received in the past. In our model, players have the additional possibility to deviate to other information mechanisms. Therefore, new techniques have to be introduced to check whether the opponent has deviated to another information mechanism. This is the aim of this subsection.

Let the information mechanisms l_1 and l_2 be fixed. We say that a deviation from l_1 to l'_1 is *undetectable* if for all behavior strategy pairs (σ_1, σ_2) in $\Gamma(l_1, l_2)$ there exists a behavior strategy σ'_1 in $\Gamma(l'_1, l_2)$ such that (σ'_1, σ_2) induces the same probability distribution on player 2's histories as (σ_1, σ_2) .

By $IH_2(\sigma_1, \sigma_2)$ we denote the set of infinite player 2 histories $h_2 = (l_2^1, l_2^2, l_2^3, \dots)$ having a finite sub-history which is impossible if (σ_1, σ_2) is played. We say that a deviation from l_1 to l'_1 is *detectable* if there exists a behavior strategy pair (σ_1, σ_2) in the repeated game $\Gamma(l_1, l_2)$ such that for all behavior strategies σ'_1 in $\Gamma(l'_1, l_2)$ we have $\mathbb{P}_{\sigma'_1, \sigma_2}(IH_2(\sigma_1, \sigma_2)) = 1$. This means that, after choosing l'_1 , every strategy of player 1 will induce, with probability 1, a finite history for player 2 which could not occur if player 1 had chosen l_1 and played σ_1 .

Suppose that the set of possible signals for player 1 is $L_1 = \{t_1^1, \dots, t_1^{|L_1|}\}$. Since l_1 and l_2 are nontrivial, there exist actions $a_1^0, a_1^1 \in A_1$ and $\bar{a}_2 \in A_2$ such that $l_2(a_1^0, \bar{a}_2) \neq l_2(a_1^1, \bar{a}_2)$. These actions will be used for communication between the players. For an action $a_1 \in A_1$, the *test* $T\langle a_1 \rangle$ is a strategy pair for $1 + |L_1|$ stages defined as follows. At the first stage player 2 plays each of his actions with probability $1/|A_2|$, while at all remaining stages he plays \bar{a}_2 . Player 1 plays action a_1 at the first stage and afterwards answers which signal he received at the first stage. Player 1 codes his answer in the following way. At the $(1 + m)$ th stage he plays action a_1^1 if he received signal t_1^m at the first stage and he plays action a_1^0 otherwise. Tests $T\langle a_2 \rangle$ are defined analogously.

We define the *information mechanism inspection* to be the strategy pair consisting of all possible tests $T\langle a_1 \rangle$ and $T\langle a_2 \rangle$. Since each test is finite and there are only finitely many tests, it follows that this inspection has finite length.

LEMMA 3.2. *Let l_1, l_2 be nontrivial information mechanisms and l'_1 some other information mechanism.*

(a) *If $l'_1 \succ^{l_2} l_1$ then the deviation from l_1 to l'_1 is undetectable.*

(b) *If $l'_1 \not\succ^{l_2} l_1$ then the deviation can be detected by repeating the information mechanism inspection infinitely many times.*

PROOF. (a) Let $l'_1 \succ^{l_2} l_1$ and let (σ_1, σ_2) be a behavior strategy pair in $\Gamma(l_1, l_2)$. We construct a behavior strategy σ'_1 in $\Gamma(l'_1, l_2)$ such that (σ'_1, σ_2) and (σ_1, σ_2) induce the same probability distribution on all finite histories of player 2.

For $n = 1$, we define $\sigma_1^{-1}(a_1) := \sigma_1^1(\psi_1^{-1}(a_1))$ for every $a_1 \in A_1$, where $\sigma_1^1(\psi_1^{-1}(a_1))$ denotes the sum $\sum_{\bar{a}_1 \in \psi_1^{-1}(a_1)} \sigma_1^1(\bar{a}_1)$. Hence, player 1 uses ψ_1 to transform the old mixed action σ_1^1 into the new mixed action σ_1^{-1} . Since $\psi_1(a_1) \sim^{l_2} a_1$ it follows that (σ'_1, σ_2) and (σ_1, σ_2) induce the same probability distributions on the one-stage histories of player 2.

For $n > 1$, we define σ_1^{-n} as follows. If player 1 observes history h_1^{n-1} at stage n , he can compute the conditional probability distribution $\mathbb{P}_{\sigma_1, \sigma_2}(\cdot | h_1^{n-1})$ on the player 1 histories up to stage $n - 1$ which would have occurred if he had played l_1 and σ_1 instead of l'_1 and σ'_1 . The latter is due to the fact that $(\psi_1(a_1) | l'_1)$ is more informative than $(a_1 | l_1)$ for every a_1 . For every h_1^{n-1} we define

$$\sigma_1^{-n}(h_1^{n-1})(a_1) := \sum_{h_1^{n-1}} \mathbb{P}_{\sigma_1, \sigma_2}(h_1^{n-1} | h_1^{n-1}) \cdot \sigma_1^{-n-1}(h_1^{n-1})(\psi_1^{-1}(a_1))$$

for all a_1 . In words, player 1 calculates the probability distribution on histories that he would have observed while using l_1 and σ_1 and for each of those histories transforms the old mixed action into a new one by using ψ_1 . Since ψ_1 transforms actions into indistinguishable ones (from player 2's viewpoint) and player 1 is always able to calculate the probability distribution on histories that would have occurred when using l_1 and σ_1 , it can be shown by induction on n that (σ'_1, σ_2) and (σ_1, σ_2) induce the same probability distribution on the set of n -stage histories of player 2 for all n .

(b) Now, let $l'_1 \not\succ^{l_2} l_1$. Then, by definition, there is an action a_1 for which there is no action a'_1 with $(a'_1 | l'_1) \succ^{l_2} (a_1 | l_1)$. So, by using l'_1 and a'_1 , player 1 either does not play in an indistinguishable way or he cannot distinguish all the actions of player 2 that he could distinguish by playing action a_1 while using l_1 . In the former case, since each player 2 action occurs with strictly positive probability in the information mechanism inspection and the inspection is repeated infinitely many times, player 2 will, with probability 1, receive a signal that he should not get if player 1 played a_1 . In the latter case, player 2 will, with probability 1, receive a wrong answer from player 1 in one of the repetitions of the test $T\langle a_1 \rangle$. \square

3.3. Proof of theorem. Before we prove the characterization of the simple lower equilibrium payoffs as given in Theorem 3.1, we need some technical lemmas. In these lemmas, the nontrivial information mechanism l_1 is fixed and l'_1 is some other information mechanism.

By $V_1(l'_1)$ we define the set of payoff pairs v which can be written as a convex combination

$$v = \sum_{a_1} \alpha^{a_1} v(a_1, p_2^{a_1})$$

with

$$\sum_{a_1} \alpha^{a_1} v_1(a_1, p_2^{a_1}) - c(l_1) \geq \sum_{a_1} \alpha^{a_1} v_1(\psi_1(a_1), p_2^{a_1}) - c(l'_1)$$

for all $\psi_1 \in \Psi_1(l'_1)$. Here, $p_2^{a_1}$ denotes a mixed action for player 2 in the one-shot game, i.e., a probability distribution on A_2 .

LEMMA 3.3. $v(\mathcal{B}_1(l'_1)) = V_1(l'_1)$.

The proof is elementary and can be found in Flesch and Perea y Monsuwé (1999).

It can be checked easily that $v(\mathcal{B}_1(l'_1))$ is a convex and compact set.

Now, suppose that (σ_1, σ_2) is a behavior strategy pair in $\Gamma(l_1, l_2)$ inducing a payoff pair $v \notin v(\mathcal{B}_1(l'_1))$. Then, we can find two disjoint, closed half-spaces K^+ and K^- such that $\text{int}(K^+)$ contains v and K^- contains $v(\mathcal{B}_1(l'_1))$. Here, $\text{int}(K^+)$ denotes the interior of K^+ .

LEMMA 3.4. *There is a $\gamma > 0$ such that $v(B) \in K^+$ implies that $v_1(\psi_1(B)) - c(l'_1) \geq v_1(B) - c(l_1) + \gamma$ for some $\psi_1 \in \Psi_1(l'_1)$.*

PROOF. Let \mathcal{B}^+ be the set of weight matrices B with $v(B) \in K^+$. Let f be the function assigning to every weight matrix B the number

$$\max_{\psi_1 \in \Psi_1(l'_1)} [[v_1(\psi_1(B)) - c(l'_1)] - [v_1(B) - c(l_1)]].$$

By construction, f is continuous and $f(B) > 0$ for all $B \in \mathcal{B}^+$. Since \mathcal{B}^+ is compact, f attains a minimum $\gamma > 0$ on \mathcal{B}^+ . \square

PROOF OF THEOREM 3.1.

(a) “ \subset ” Obviously, $SLEP(l_1, l_2) \subset IR$ since both players can guarantee their individually rational payoff of the one-shot game by choosing the trivial information mechanism and playing the max-min action in every stage of the game.

By symmetry, it suffices to show that $SLEP(l_1, l_2) \subset v(\mathcal{B}_1(l'_1)) - (c(l_1), c(l_2))$ for every l'_1 . Assume by way of contradiction that $(\phi_1, \phi_2) = ((l_1, \sigma_1), (l_2, \sigma_2))$ is a lower-equilibrium with payoff $v - (c(l_1), c(l_2)) \notin v(\mathcal{B}_1(l'_1)) - (c(l_1), c(l_2))$ for some l'_1 , so $v \notin v(\mathcal{B}_1(l'_1))$. For convenience, we write σ_i instead of $\sigma_i(l_i)$ and \mathbb{P} instead of $\mathbb{P}_{(\sigma_1, \sigma_2)}$. For every $n \in \mathbb{N}$, let $\mathbb{P}^n(a_1, a_2)$ be the probability that the action pair a_1, a_2 is played in stage n and let B^n be the weight-matrix given by $B^n(a_1, a_2) = \mathbb{P}^n(a_1, a_2)$ for all a_1, a_2 . By definition, it holds that

$$v = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N v(B^n) = \lim_{N \rightarrow \infty} v\left(\frac{1}{N} \sum_{n=1}^N B^n\right).$$

Since $v \notin v(\mathcal{B}_1(l'_1))$, we can find sets K^+ and K^- as constructed above. Using the fact that $v \in \text{int}(K^+)$, it follows that $v((1/N) \sum_{n=1}^N B^n) \in K^+$ for large N . We may assume, w.l.o.g., that $v((1/N) \sum_{n=1}^N B^n) \in K^+$ for all N . By Lemma 3.4, there is a $\gamma > 0$ such that for every N we can find a $\psi_1^N \in \Psi_1(l'_1)$ with

$$v_1\left(\psi_1^N\left(\frac{1}{N} \sum_{n=1}^N B^n\right)\right) - c(l'_1) \geq v_1\left(\frac{1}{N} \sum_{n=1}^N B^n\right) - c(l_1) + \gamma,$$

implying that

$$\frac{1}{N} \sum_{n=1}^N v_1(\psi_1^N(B^n)) - c(l'_1) \geq \frac{1}{N} \sum_{n=1}^N v_1(B^n) - c(l_1) + \gamma.$$

For every stage n , let $\bar{\psi}_1^n \in \Psi_1(l'_1)$ be such that $v_1(\bar{\psi}_1^n(B^n))$ is maximal. With the inequality above, it follows that

$$\frac{1}{N} \sum_{n=1}^N v_1(\bar{\psi}_1^n(B^n)) - c(l'_1) \geq \frac{1}{N} \sum_{n=1}^N v_1(\psi_1^n(B^n)) - c(l'_1) \geq \frac{1}{N} \sum_{n=1}^N v_1(B^n) - c(l_1) + \gamma$$

for all N .

Let σ'_1 be the behavior strategy which is obtained from σ_1 by transforming the actions in stage n according to $\bar{\psi}_1^n$. The strategy σ'_1 is constructed in the same way as the strategy σ'_1 in the proof of Lemma 3.2. Since the deviation from σ_1 to σ'_1 cannot be detected, (σ'_1, σ_2) induces the same probability distributions on the player 2 histories as (σ_1, σ_2) . By construction,

$$\begin{aligned} \underline{v}_1(\sigma'_1, \sigma_2) - c(l'_1) &= \liminf_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N v_1(\bar{\psi}_1^n(B^n)) - c(l'_1) \\ &\geq \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N v_1(B^n) - c(l_1) + \gamma = v_1 - c(l_1) + \gamma \end{aligned}$$

which is a contradiction to the fact that (ϕ_1, ϕ_2) is a lower equilibrium.

(b) “ \supset ” Let v be a payoff in the right-hand side. We construct a simple lower equilibrium $((l_1, \sigma_1), (l_2, \sigma_2))$ with payoff v . The implementation of the strategies σ_1 and σ_2 consists of the following six phases which are repeated infinitely many times. The phases 2a, 2b, 3a and 3b are constructed in a similar way as in Lehrer (1989).

Phase 1a. Player 1 plays his information mechanism inspection, as described in §4.2.

Phase 1b. Same for player 2.

Phase 2a. Let l'_1 be an information mechanism in \mathcal{L}_1 . By Lemma 3.3, v can be written as

$$v = \sum_{a_1} \alpha^{a_1} v(a_1, p_2^{a_1})$$

with

$$(3.1) \quad \sum_{a_1} \alpha^{a_1} v_1(a_1, p_2^{a_1}) - c(l_1) \geq \sum_{a_1} \alpha^{a_1} v_1(\psi_1(a_1), p_2^{a_1}) - c(l'_1)$$

for all $\psi_1 \in \Psi_1(l'_1)$.

For every $\epsilon > 0$, let $p_2^{a_1}(\epsilon)$ be the unique mixed action closest to $p_2^{a_1}$ which puts minimal weight ϵ on every action.

If the n th stage is reached, let K be the block containing the next n^n stages. We divide K into sub-blocks K^{a_1} such that the relative length of K^{a_1} in K is close to the corresponding coefficient α^{a_1} . In each stage of block K^{a_1} , player 1 and 2 play a_1 and $p_2^{a_1}(\epsilon)$ respectively, where ϵ has to be chosen small enough. Below, we specify what we mean by “small enough”.

On the other hand, the values for ϵ have to be chosen in such a way that, starting from an arbitrary stage, the probability that a player 2 action is played in the future, is equal to 1.

We repeat this process until every information mechanism I'_1 has been chosen.

Phase 2b. Similar to phase 2a, but now the roles of player 1 and 2 are exchanged.

Phase 3a. Let \mathbf{S} be the set of stages containing all repetitions of phase 2b up to the current stage. If this is the n th repetition of phase 3a, player 1 asks his opponent the question “Which signal did you receive at the n th stage of S ?” Afterwards, player 2 answers the question. The question and answer are implemented in the same way as described in the information mechanism inspection. Since player 2 is supposed to play only pure actions in phase 2b, player 1 knows exactly which signal player 2 should have received at the n th stage of S .

Phase 3b. Same as phase 3a, but now with exchanged roles of player 1 and 2.

At the end of phase 3b, the players return to phase 1a and so on.

If at any moment in the game, a player notices that his opponent has deviated from the prescribed strategy, he will punish him by playing his min-max action forever.

For a set M of stages, $\limsup_{N \rightarrow \infty} (|M \cap \{1, \dots, N\}|/N)$ is called the *upper density* of M , whereas $\liminf_{N \rightarrow \infty} (|M \cap \{1, \dots, N\}|/N)$ denotes the *lower density* of M . By construction, the set of stages belonging to phases 1a, 1b, 3a and 3b has upper density zero. Therefore, these phases have no influence on the final payoff.

The sub-blocks K^{a1} , K^{a2} should be arranged in such a way that the average payoffs converge. The lengths of the sub-blocks K^{a1} , K^{a2} in phase 2a and 2b respectively and the value of ϵ in the same phases can be chosen in such a way that $v(\phi_1, \phi_2) = v$.

Finally, we show that $((l_1, \sigma_1), (l_2, \sigma_2))$ is a lower equilibrium. Suppose that player 1 would deviate to a strategy (l'_1, σ'_1) . If l'_1 is not greater than l_1 , then, with probability one, the deviation will be detected in one of the repetitions of phase 1b and player 1 will be punished. So the only possible way for player 1 to improve his payoff is by deviating to an $l'_1 > l_2$.

Look at an arbitrary block K in phase 2b which corresponds to the information mechanism I'_1 . The length of block K is equal to n^n , where n is the number of preceding stages. Let \mathcal{K} be the collection of blocks K which correspond to I'_1 and let \mathcal{K}^* be the set of final stages of blocks $K \in \mathcal{K}$. Let x_1^n be the player 1 expected payoff at stage n when (σ'_1, σ_2) is played. Since in the long run, the average payoff until some stage in \mathcal{K}^* is completely determined by the payoffs at stages in \mathcal{K} , it follows that

$$\liminf_{N \in \mathcal{K}^*} \frac{1}{N} \sum_{n=1}^N x_1^n = \liminf_{N \in \mathcal{K}^*} \frac{1}{N} \sum_{n \in \mathcal{K} \cap \{1, \dots, N\}} x_1^n.$$

Since

$$\underline{v}_1(\sigma'_1, \sigma_2) = \liminf_N v_1^N(\sigma'_1, \sigma_2) \leq \liminf_{N \in \mathcal{K}^*} \frac{1}{N} \sum_{n=1}^N x_1^n,$$

it suffices to show that

$$\liminf_{N \in \mathcal{K}^*} \frac{1}{N} \sum_{n \in \mathcal{K} \cap \{1, \dots, N\}} x_1^n \leq v_1(\sigma_1, \sigma_2).$$

In other words, it remains to prove that by playing σ'_1 , the average payoff at stages in \mathcal{K} does not exceed $v_1(\sigma_1, \sigma_2)$. If player 1 only deviates on a set of stages with lower density zero,

it cannot increase the lower expected payoff. Therefore, it suffices to consider player 1 deviations on a set of stages with positive lower density.

Suppose that player 1 deviates in \mathcal{H} by transforming actions into distinguishable actions on a set of stages with positive lower density. Then, with probability 1, player 2 will detect this deviation in one of the repetitions of phase 2a.

If player 1 deviates in \mathcal{H} by transforming actions into less informative actions on a set of stages with positive lower density, then, with certainty, this deviation will be detected by player 2 in one of the repetitions of phase 3b.

So the only possible way for player 1 to increase his payoff is by transforming actions into greater actions on a set of positive lower density. In the long run, the expected average payoff in stages of \mathcal{H} converges to

$$v = \sum_{a_1} \alpha^{a_1} v(a_1, p_2^{a_1})$$

with

$$\sum_{a_1} \alpha^{a_1} v_1(a_1, p_2^{a_1}) - c(l_1) \geq \sum_{a_1} \alpha^{a_1} v_1(\psi_1(a_1), p_2^{a_1}) - c(l'_1)$$

for all $\psi_1 \in \Psi_1(l'_1)$. Therefore, player 1 cannot increase his long run average payoff in stages of \mathcal{H} by transforming actions into greater actions, which leads to the conclusion that

$$\liminf_{N \in \mathcal{H}^*} \frac{1}{N} \sum_{n \in \mathcal{H} \cap \{1, \dots, N\}} x_1^n \leq v_1(\sigma_1, \sigma_2).$$

Since this holds for every l'_1 , it follows that $\underline{v}_1(\sigma'_1, \sigma_2) \leq v_1(\sigma_1, \sigma_2)$. \square

4. Nonsimple strategies.

EXAMPLE. Consider the repeated game in which the one-shot game is given by

	<i>L</i>	<i>R</i>
<i>U</i>	2, 0	-1, 0
<i>M</i>	3, 0	0, 1
<i>D</i>	3, 0	0, 1

Assume that $\mathcal{L}_1 = \{\text{trivial, perfect}\}$ and $\mathcal{L}_2 = \{\text{trivial, } l_2^1, l_2^2\}$ where

$$l_2^1 = \begin{array}{cc} a & c \\ b & c \\ a & c \end{array} \quad \text{and} \quad l_2^2 = \begin{array}{cc} a & c \\ a & c \\ b & c \end{array}.$$

All mechanisms can be obtained at price zero.

Claim. The payoff (2, 0) is a lower equilibrium payoff if players are allowed to randomize over choosing different information mechanisms but cannot be obtained as a simple lower equilibrium payoff.

Proof of claim. First we construct a lower equilibrium in which player 2 strictly randomizes between l_2^1 and l_2^2 resulting in the payoff (2, 0). Let the strategies for player 1 and 2 be as follows. Player 1 chooses the perfect mechanism with probability one and plays *U* at

every stage, irrespective of the history. Player 2 chooses l_2^1 and l_2^2 both with probability $\frac{1}{2}$, plays L as long as he observes signal a and plays R if he observed signal b or c in one of the previous stages. Clearly, these strategies induce the payoff $(2, 0)$. If player 1 deviates from U to another action at some stage then player 2 will detect this with probability $\frac{1}{2}$. The lower payoff for player 1 by deviating at some stage can therefore not exceed $\frac{1}{2} \cdot 3 + \frac{1}{2} \cdot 0 < 2$. Since player 1 has no incentive to choose the trivial mechanism and player 2 has no incentive for deviation, these strategies constitute a lower equilibrium.

Next we prove that there is no simple lower equilibrium which induces payoff $(2, 0)$. Suppose that player 2 chooses one of his information mechanisms with probability one. Let

$$\mathbb{P}(U) := \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \mathbb{P}^k(U),$$

where $\mathbb{P}^k(U)$ is the probability that player 1 plays U at stage k . If $\mathbb{P}(U) > 0$ then player 1 can always replace the action U by D or M (depending on whether player 2 uses l_2^1 or l_2^2 , respectively) without being detected and increase his payoff. Therefore this cannot be a lower equilibrium. If $\mathbb{P}(U) = 0$ then the strategies cannot induce the payoff $(2, 0)$. Consequently there is no lower equilibrium with payoff $(2, 0)$ in which player 1 and 2 choose one of their mechanisms with certainty. \square

However, if we try to characterize the set of lower equilibrium payoffs in a situation where players may randomize over the information mechanisms that can be bought, we encounter some very difficult problems. This is due to the fact that players are uncertain about the mechanism used by the opponent. Detecting deviations becomes therefore a very complicated matter.

Acknowledgment. The authors thank two anonymous referees for their very helpful comments.

References

- Flesch, J., A. Perea y Monsuwé. 1999. Repeated games with endogenous choice of information mechanisms. Working Paper, University of Maastricht Maastricht, The Netherlands.
- Fudenberg, D., D. Levine. 1991. An approximate Folk Theorem with imperfect private information. *J. Econom. Theory* **54** 26–47.
- Hillas, J., M. Liu. 1995. Repeated games with partial monitoring: The stochastic signaling case. Discussion Paper, State University of New York at Stony Brook, Stony Brook, NY. December 1995.
- Lehrer, E. 1989. Lower equilibrium payoffs in two-player repeated games with nonobservable actions. *Internat. J. Game Theory* **18** 57–89.
- . 1990. Nash equilibria of n -player repeated games with semi-standard information. *Internat. J. Game Theory* **19** 191–217.
- . 1991. Internal correlation in repeated games. *Internat. J. Game Theory* **19** 431–456.
- . 1992a. Correlated equilibria in two-player repeated games with nonobservable actions. *Math. Oper. Res.* **17** 175–199.
- . 1992b. Two-player repeated games with nonobservable actions and observable payoffs. *Math. Oper. Res.* **17** 200–224.
- Radner, R. 1986. Repeated partnership games with imperfect monitoring and no discounting. *Rev. Econom. Stud.* **53** 43–58.
- J. Flesch: Department of Mathematics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands; E-mail: flesch@math.unimaas.nl
- A. Perea y Monsuwé: Departamento de Economía, Universidad Carlos III de Madrid, Calle Madrid 126, 28903 Getafe-Madrid, Spain; E-mail: perea@eco.uc3m.es