

האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

**LEADERSHIP GAMES WITH CONVEX
STRATEGY SETS**

By

BERNHARD VON STENGEL and SHMUEL ZAMIR

Discussion Paper # 525

November 2009

מרכז לחקר הרציונליות

**CENTER FOR THE STUDY
OF RATIONALITY**

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel

PHONE: [972]-2-6584135 FAX: [972]-2-6513681

E-MAIL: ratio@math.huji.ac.il

URL: <http://www.ratio.huji.ac.il/>

Leadership Games with Convex Strategy Sets*

Bernhard von Stengel

Department of Mathematics, London School of Economics,
Houghton St, London WC2A 2AE, United Kingdom
email: stengel@maths.lse.ac.uk

Shmuel Zamir

Center for the Study of Rationality, The Hebrew University at Jerusalem,
Givat Ram, Jerusalem 91904, Israel
email: zamir@math.huji.ac.il

November 27, 2009

Abstract

A basic model of commitment is to convert a two-player game in strategic form to a “leadership game” with the same payoffs, where one player, the leader, commits to a strategy, to which the second player always chooses a best reply. This paper studies such leadership games for games with convex strategy sets. We apply them to mixed extensions of finite games, which we analyze completely, including non-generic games. The main result is that leadership is advantageous in the sense that, as a set, the leader’s payoffs in equilibrium are at least as high as his Nash and correlated equilibrium payoffs in the simultaneous game. We also consider leadership games with three or more players, where most conclusions no longer hold.

Keywords: Commitment, correlated equilibrium, first-mover advantage, follower, leader, Stackelberg game.

JEL Classification Number: C72

*Support by a STICERD Distinguished Visitorship for Shmuel Zamir to visit the London School of Economics is gratefully acknowledged. We thank Jacqueline Morgan for comments on dynamic games, a referee for suggesting the proof of Theorem 2, and another referee and the associate editor for helpful comments.

1 Introduction

The possible advantage of commitment power is a game-theoretic result known to the general public, ever since its popularization by Schelling (1960). Cournot's (1838) duopoly model of quantity competition was modified by von Stackelberg (1934), who demonstrated that a firm with the power to commit to a quantity of production profits from this leadership position. The leader-follower issue has been studied in depth in oligopoly theory as "Stackelberg leadership"; see Friedman (1977), Hamilton and Slutsky (1990) and the correction to that paper by Amir (1995), Shapiro (1989), or Amir and Grilo (1999) for discussions and references.

We define a leadership game as follows (for details see Section 2). Consider a game of $k + 1$ players in strategic form. Declare one player as *leader* and let his strategy set be X . The remaining k players are called *followers*. Let the set of their partial strategy profiles (with k strategies) be Y , so that $X \times Y$ is the set of full strategy profiles. The leadership game is the extensive game where the leader chooses x in X , the followers are informed about x and choose simultaneously their strategies as $f(x)$ in Y , and all players receive their payoffs as given by the strategy profile $(x, f(x))$. We only consider subgame perfect equilibria of the leadership game where for any x the followers play among themselves a Nash equilibrium $f(x)$ in the game induced by x , even off the equilibrium path. We call $f(x)$ the *response* of the followers to x , which is simply a best reply in the original game if there is only one follower. (The set of equilibria that are not subgame perfect seems too large to allow any interesting conclusions.)

Our aim is to analyze completely leadership games for the mixed extension of a bimatrix game, that is, of a finite two-player game in strategic form. Then there is only one follower ($k = 1$). The leader commits to a mixed strategy x in the bimatrix game. The follower's response $f(x)$ is also a mixed strategy. The pair of pure actions is then chosen independently according to x and $f(x)$ with the corresponding bimatrix game payoffs, and the players maximize expected payoffs as normally.

The payoff to the leader in a subgame perfect equilibrium of the leadership game is called a *leader payoff*. His payoff in a Nash equilibrium of the simultaneous game is called a *Nash payoff*. When considering the simultaneous game, we often have to identify the player who becomes leader in the corresponding leadership game; for simplicity of identification, we call this player also "leader" in the simultaneous game.

For the mixed extension of a bimatrix game, our main result (Corollary 8) states that the set of leader payoffs is an interval $[L, H]$ so that $H \geq E$ for all Nash payoffs E , and $L \geq E$ for at least one Nash payoff E . Furthermore, Theorem 12 states that $H \geq C$ for any correlated equilibrium payoff C to the leader. In this sense, the possibility to commit, by changing a simultaneous game to a leadership game, never harms the leader. However, this no longer holds for two or more followers, where leadership can be disadvantageous (see Remark 5).

One motivation to consider commitment to mixed strategies is the "classical view" of mixed strategies (see also Reny and Robson 2004). This is the view of von Neumann and Morgenstern (1947), who explicitly define the leadership game corresponding to a zero-

sum game, first with commitment to pure strategies (p. 100) and then to mixed strategies (p. 149), as a way of introducing the maxmin and minmax value of the game. They consider the leader to be a priori at an obvious disadvantage. By the minimax theorem, a player is not harmed even if his opponent learns his optimal mixed strategy. Hence, in two-person zero-sum games, commitment to a mixed strategy does not hurt the leader, in line with Corollary 8. The value of a zero-sum game is its unique leadership and Nash payoff.

Important applications of commitment to mixed strategies are inspection games. They model inspections for arms control treaties, tax auditing, or monitoring traffic violations; for a survey see Avenhaus, von Stengel, and Zamir (2002). With costly inspections, such games typically have unique mixed equilibria, and in the corresponding leadership games, the inspector is a natural leader. As observed by Maschler (1966), commitment helps the leader because the follower, who is inspected, acts legally in an equilibrium of the leadership game, but acts illegally with positive probability in the Nash equilibrium of the simultaneous game.

The central observation about leadership games for mixed extensions of bimatrix games is the following. When the leader commits to his mixed strategy in equilibrium, the follower is typically indifferent between several pure best replies. However, the condition of subgame perfection implies that on the equilibrium path, the follower chooses the reply that gives the *best* possible payoff to the leader; otherwise, the leader could improve his payoff by changing his commitment slightly so that the desired reply is unique (which is possible generically).

For inspection games, this reasoning based on subgame perfection was used by Avenhaus, Okada, and Zamir (1991). Maschler (1996) still postulated a benevolent reaction of the follower when she is indifferent (and called this behavior “pareto-optimal”), or else suggested to look in effect at an ε -equilibrium in which the leader sacrifices an arbitrarily small amount to induce the desired reaction of the follower. A similar observation is known for bargaining games, for example in the iterated offers model of Rubinstein (1982). In a subgame perfect equilibrium of this game, the first player makes the second player indifferent between accepting or rejecting the offer, but the second player nevertheless accepts.

Some of our results apply to more general games than mixed extensions of bimatrix games. In particular, we are indebted to a referee who suggested a short proof of Corollary 8 based on Kakutani’s fixed point theorem. Different parts of Corollary 8 hold under assumptions that can be weakened to varying extent. We therefore present these parts separately, as follows.

In Section 2, we give a characterization in Theorem 1 of the lowest leader payoff, using standard assumptions so that Kakutani’s fixed point theorem can be applied, for games with any number of followers. Suppose that they always choose their response to give the worst possible payoff to the leader. In other words, the leader maximizes his payoff under the “pessimistic” view that the followers act to his disadvantage. (This pessimistic view is also used to define a “Stackelberg payoff” to the leader in dynamic games; see Başar and Olsder 1982, equation (41) on p. 136, and p. 141.) The resulting payoff function to

the leader is typically discontinuous and has no maximum (see also Morgan and Patrone (2006) and references). However, the *supremum* of the “pessimistically” computed payoff to the leader is obtained in a subgame perfect equilibrium of the leadership game, as we show in Theorem 1. Subgame perfection implies that on the equilibrium path, the followers’ response is *not* according to the pessimistic assumption but instead yields the supremum payoff to the leader.

In Section 3, we observe in Theorem 2 that the lowest leader payoff is no worse than the lowest Nash payoff. This theorem requires strong assumptions that hold for mixed extensions of bimatrix games (Corollary 3), but not, for example, for mixed extensions of three-player games (Remark 5). Furthermore, the highest leader payoff is obtained when the followers always reply in the best possible way for the leader. It is easy to see that this payoff is at least as high as any Nash payoff. Moreover, if the set of the followers’ responses is connected, then the set of leader payoffs is an interval (Proposition 7).

In Section 4, we consider mixed extensions of bimatrix games. We explicitly characterize the lowest and higher leader payoffs and show how to compute them by linear programming. For generic games, they are equal.

In Section 5, we show that the highest leader payoff H is at least as high as any correlated equilibrium payoff to the leader. This is no longer true for the coarse correlated equilibrium due to Moulin and Vial (1978) that involves commitment by both players, which may give a higher payoff than H .

2 The lowest leader payoff

Our main results concern finite two-player games with commitment to mixed strategies. We regard such games via their mixed extension, where each mixed-strategy simplex becomes a set of new pure strategies. Some of our results hold more generally for any finite number of players with convex and compact strategy sets.

For the general case, we consider a game with finite player set N . Each player $i \in N$ has a convex and compact strategy set S_i , and a continuous payoff function $u_i : \prod_{i \in N} S_i \rightarrow \mathbb{R}$. Let $S_{-i} = \prod_{j \in N - \{i\}} S_j$. For sets X, Y , the set of correspondences (set-valued mappings) $X \rightarrow 2^Y$ is denoted by $X \rightrightarrows Y$. Player i ’s best-reply correspondence $B_i : S_{-i} \rightrightarrows S_i$ is given by

$$B_i(s_{-i}) = \arg \max_{s_i \in S_i} u_i(s_i, s_{-i}), \quad (1)$$

where $\arg \max$ gives the set of all maximizers. Suppose each player’s best-reply correspondence is convex-valued and upper hemi-continuous (uhc), that is, it has a closed graph $\bigcup_{s_{-i} \in S_{-i}} \{s_{-i}\} \times B_i(s_{-i})$. Then by the fixed point theorem of Kakutani (1941), there is a fixed point $(s_i)_{i \in N}$ with $s_i \in B_i(s_{-i})$ for all $i \in N$. This is an equilibrium of the game where each player i plays a best reply s_i to the remaining strategies s_{-i} . These conditions hold for the mixed extension of a finite game where S_i is player i ’s mixed strategy simplex and u_i is his expected payoff function (Nash 1950).

Consider such a game with player set $N = \{1, 2, \dots, k + 1\}$. The corresponding *leadership game* is a two-stage game played as follows. Player 1 is called *leader*, and the k players $2, \dots, k + 1$ are called *followers*. First, the leader chooses and commits to a strategy s_1 in S_1 , which is announced to all followers, who then simultaneously choose their strategies s_2, \dots, s_{k+1} , which are played together with s_1 . The players' payoffs for the strategy profile (s_1, \dots, s_{k+1}) are as in the original game.

For convenience, we write $X = S_1$, and let $Y = S_{-1}$ be the set of partial strategy profiles $y = (s_2, \dots, s_{k+1})$ of the followers. For any strategy profile $(x, y) \in X \times Y$, denote the payoff to the leader by $a(x, y) = u_1(x, y)$. In the leadership game, y may depend on x , so a strategy profile in the leadership game is given by (x, f) where $x \in X$ and $f : X \rightarrow Y$.

We consider only subgame perfect equilibria of leadership games. In such a subgame perfect equilibrium (x^*, f^*) , the followers' response given by $f^*(x)$ in Y is a Nash equilibrium of the game induced by x , for any x in X . Moreover, the leader's commitment x^* is optimal, so $a(x^*, f^*(x^*)) \geq a(x, f^*(x))$ for all $x \in X$. A *leader payoff* is the corresponding payoff $a(x^*, f^*(x^*))$.

For x in X , let the subset $E(x)$ of Y be the set of Nash equilibria of the game induced by x , which is nonempty by Kakutani's theorem. Also, $E(x)$ is the intersection of the followers' best reply correspondences and therefore closed. When there is only one follower ($k = 1$), then $E(x)$ is simply $B_2(x)$.

We define a new correspondence $F : X \rightrightarrows Y$ which expresses a "pessimistic" view of the leader. The correspondence F is the sub-correspondence of E where the followers' response is the worst possible for the leader,

$$F(x) = \arg \min_{y \in E(x)} a(x, y), \quad (2)$$

so $F(x)$ is the set of all y in $E(x)$ that minimize $a(x, y)$. If $k = 1$, then $F(x)$ contains the follower's best replies y to x that minimize the payoff to the leader.

		2				
		<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>
1	<i>T</i>	4	4	0	2	4
	<i>B</i>	8	0	3	1	0

Figure 1 Example of a 2×5 game. In each cell, the payoffs to player 1 and 2 are shown in the bottom left and top right corner, respectively.

We illustrate our considerations with an example, shown in Figure 1, which will be used throughout the paper. We consider the mixed extension of this 2×5 game, with player 2 as the only follower. The set X of mixed strategies of player 1 can be identified with the interval $[0, 1]$ for the probability that player 1 plays strategy B .

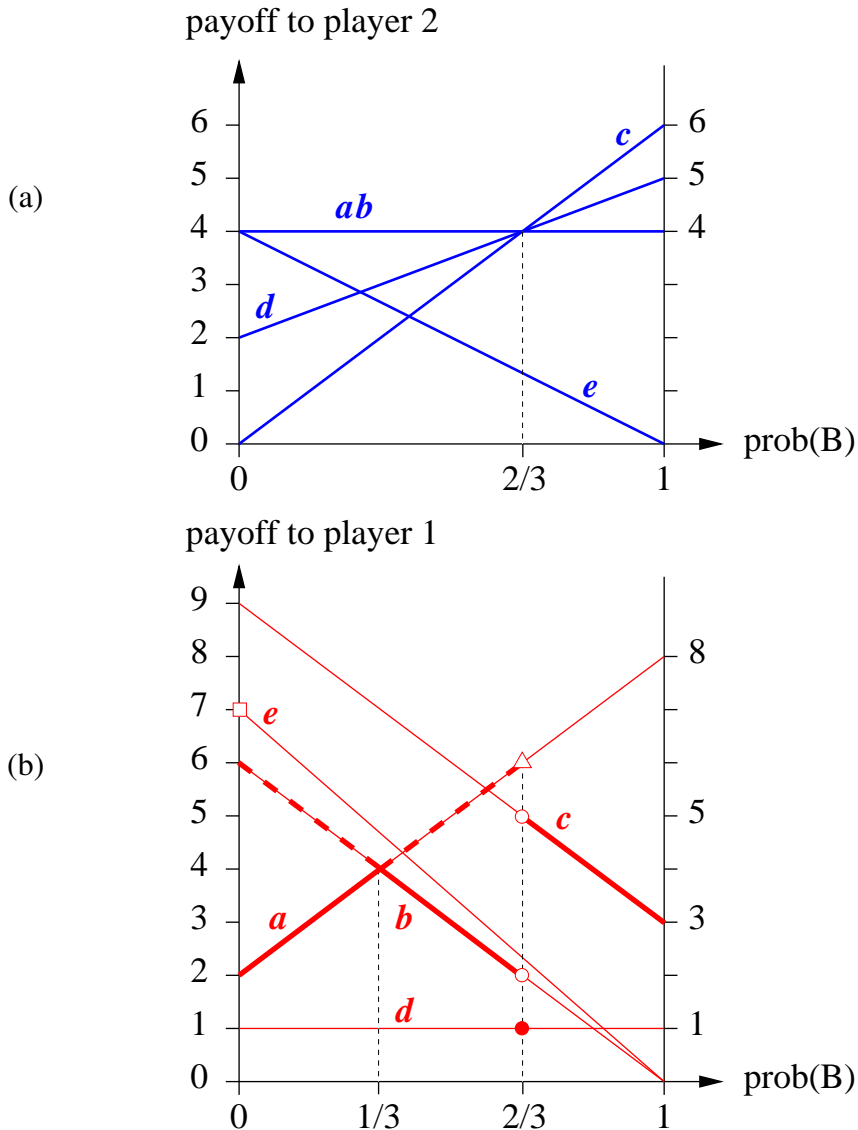


Figure 2 (a) Expected payoffs to player 2, (b) expected payoffs to player 1, as functions of player 1's mixed strategy, in the game of Figure 1. The bold line in (b), including the full dot at $2/3$, indicates the "pessimistic" leader payoff resulting from any reply in $F(x)$ as in (3).

Figure 2 shows the expected payoffs to the two players as a function of the mixed strategy of player 1. The top graph (a) shows the expected payoffs to player 2, from which one can see that the pure best replies of player 2 are

$$\begin{aligned}
 a, b, e & \quad \text{when } \text{prob}(B) = 0, \\
 a, b & \quad \text{when } 0 < \text{prob}(B) < 2/3, \\
 a, b, c, d & \quad \text{when } \text{prob}(B) = 2/3, \\
 c & \quad \text{when } \text{prob}(B) > 2/3.
 \end{aligned}$$

For each of the pure strategies a, b, c, d, e chosen by player 2, the bottom graph (b) in Figure 2 shows the expected payoff to player 1 as a function of his mixed strategy x . The

correspondence $F(x)$ in (2) chooses any best reply of player 2 that minimizes player 1's payoff and is given by

$$\begin{array}{ll}
a & \text{when } \text{prob}(B) < 1/3, \\
\text{any mixture of } a, b & \text{when } \text{prob}(B) = 1/3, \\
b & \text{when } 1/3 < \text{prob}(B) < 2/3, \\
d & \text{when } \text{prob}(B) = 2/3, \\
c & \text{when } \text{prob}(B) > 2/3.
\end{array} \tag{3}$$

The graph of the payoffs to player 1 when player 2 plays according to F is shown in Figure 2(b) with bold lines and the filled-in dot when $\text{prob}(B) = 2/3$ where $F(x) = \{d\}$.

As this example shows, the graph of F is in general not closed so that F is not uhc. The uhc correspondence \bar{F} is defined via the closure of the graph of F , that is, for all $(x, y) \in X \times Y$,

$$y \in \bar{F}(x) \iff (x, y) \in \overline{\bigcup_{x' \in X} \{x'\} \times F(x')}.$$
 \tag{4}

In the example, in order to obtain the uhc correspondence \bar{F} from F according to (4), we have to add the best replies b and c when $\text{prob}(B) = 2/3$. The resulting payoffs to player 1 are 2 for b and 5 for c , shown as the two white dots at the ends of the bold lines for b and c in Figure 2(b).

In the general setup, we define the following payoff to the leader:

$$L = \sup_{x \in X} \min_{y \in E(x)} a(x, y).$$
 \tag{5}

By (2) and (5), the payoff L is the supremum of $a(x, y)$ for $x \in X$ and $y \in F(x)$. In Figure 2(b), it is given by $L = 5$ when $\text{prob}(B)$ approaches $2/3$ from above, where player 2's best reply is c . However, this supremum is not achieved, because for $\text{prob}(B) = 2/3$ the only best reply in $F(x)$ is d with payoff 1 to player 1.

Nevertheless, there is a leader payoff 5 where the leader commits to $\text{prob}(B) = 2/3$ and the follower chooses c , because c is a best reply to the commitment. Moreover, there is no leadership equilibrium with payoff less than 5. Suppose that there is such an equilibrium with leader payoff $5 - \varepsilon$ for some $\varepsilon > 0$. If the leader plays $\text{prob}(B) = 2/3 + \delta$ for $\delta > 0$, the only best reply of the follower is c , with payoff higher than $5 - \varepsilon$ for sufficiently small δ . So the leader can profitably deviate. Hence, this is not an equilibrium.

The following central theorem of this section states that the lowest leader payoff is given by (5). Its proof generalizes the argument made for the preceding example.

Theorem 1 *The payoff L in (5) is the lowest leader payoff.*

Proof. We have to prove that there is a leadership equilibrium with payoff L to the leader, and that there is no lower leader payoff.

First, observe that by (2) which defines F , the payoff $a(x, y)$ is a constant function of y on $F(x)$, so that

$$L = \sup_{x \in X} \max_{y \in F(x)} a(x, y).$$
 \tag{6}

We claim that

$$L = \max_{x \in X} \max_{y \in \bar{F}(x)} a(x, y). \quad (7)$$

Note that the first “max” in (7) does not have to be written as “sup” because \bar{F} is uhc. Also, while $a(x, y)$ is constant in y on $F(x)$, it is no longer constant on $\bar{F}(x)$, so it may happen that $\max_{y \in \bar{F}(x)} a(x, y) > \max_{y \in F(x)} a(x, y)$, as the example for $\text{prob}(B) = 2/3$ shows.

To prove (7), let

$$\bar{L} = \max_{x \in X} \max_{y \in \bar{F}(x)} a(x, y) = a(x^*, y^*)$$

for some $x^* \in X$ and $y^* \in \bar{F}(x^*)$. Clearly, $L \leq \bar{L}$. Consider some sequence (x_n, y_n) that converges to (x^*, y^*) with $y_n \in F(x_n)$ for all n . Then $L \geq a(x_n, y_n)$ by (6) and, because a is continuous, $L \geq a(x^*, y^*)$, which proves (7).

Next, we show that there is a leadership equilibrium (x^*, f^*) with $L = a(x^*, y^*)$ as leader payoff. The correspondence $E : X \rightarrow Y$ is uhc because its graph is the intersection of the graphs of the followers’ best reply correspondences in the original game (as subsets of $X \times Y$), so the graph of E is closed. Because $F \subseteq E$ (in terms of the graphs of the correspondences), we therefore have

$$\bar{F} \subseteq E. \quad (8)$$

The leadership equilibrium is given by $f^*(x^*) = y^*$ and $f^*(x)$ in $\bar{F}(x)$ chosen arbitrarily for any $x \neq x^*$. This defines a subgame perfect equilibrium because $f^*(x) \in \bar{F}(x) \subseteq E(x)$ for all x , and because $L = a(x^*, y^*)$ by (7).

Finally, we claim there is no leadership equilibrium (\hat{x}, f) with leader payoff less than L . Suppose otherwise, that is, $a(\hat{x}, f(\hat{x})) = L - \varepsilon$ with $\varepsilon > 0$. Consider the above sequence (x_n, y_n) that converges to (x^*, y^*) , and choose n large enough so that $a(x_n, y_n) > a(x^*, y^*) - \varepsilon$. Then if player 1 commits to x_n , he will get a payoff $a(x_n, f(x_n)) \geq a(x_n, y_n) > L - \varepsilon$ because $a(x_n, y_n) \in F(x_n)$, so he can improve his payoff by deviating from his commitment \hat{x} , which is a contradiction. So L is the lowest leader payoff. \square

3 Leader payoff versus Nash payoff

In this section, we compare the leader payoffs with the Nash payoffs to the leader in the simultaneous game. The following result states that the lowest leader payoff L is at least as high as some Nash payoff in the simultaneous game; its proof is inspired by a suggestion of a referee. Recall that $E(x)$ is the set of Nash equilibria of the game among the followers induced by x .

Theorem 2 *Suppose that*

- (a) $E(x)$ is convex for all x in X ,
- (b) $a(x, y)$ is a convex function of y on the convex domain $E(x)$.

Then the lowest leader payoff L in (5) is at least as high as some Nash payoff to the leader in the simultaneous game.

Proof. Player 1's best-reply correspondence $B_1 : Y \rightarrow X$ is uhc and convex-valued. Let $\text{conv } \bar{F}(x)$ be the convex hull of $\bar{F}(x)$ for all $x \in X$, where \bar{F} is the uhc correspondence defined by (4). Consider the correspondence $(x, y) \mapsto B_1(y) \times \text{conv } \bar{F}(x)$. By Kakutani's theorem, it has a fixed point $(\hat{x}, \hat{y}) \in X \times Y$, that is, $\hat{x} \in B_1(\hat{y})$ and $\hat{y} \in \text{conv } \bar{F}(\hat{x})$. Because $E(\hat{x})$ is convex by assumption (a) and $\bar{F}(\hat{x}) \subseteq E(\hat{x})$ by (8), it follows that $\hat{y} \in E(\hat{x})$, so (\hat{x}, \hat{y}) is a Nash equilibrium of the simultaneous game.

Furthermore, $\hat{y} = \sum_{i=1}^k \lambda_i y_i$ for some points y_1, \dots, y_k in $\bar{F}(\hat{x})$ and nonnegative weights $\lambda_1, \dots, \lambda_k$ with $\sum_{i=1}^k \lambda_i = 1$. Then, since $a(x, y)$ is convex in y by assumption (b),

$$a(\hat{x}, \hat{y}) \leq \sum_{i=1}^k \lambda_i a(\hat{x}, y_i) \leq \max_{y \in \bar{F}(\hat{x})} a(\hat{x}, y) \leq \max_{x \in X} \max_{y \in \bar{F}(x)} a(x, y) = L$$

by (7). □

Assumptions (a) and (b) of Theorem 2 are not required for Theorem 1. They are trivially satisfied for the mixed extension of a bimatrix game:

Corollary 3 *For the mixed extension of a bimatrix game, the lowest leader payoff L is at least as high as some Nash payoff to the leader in the simultaneous game.*

Proof. Assumption (a) of Theorem 2 holds because for one follower $E = B_2$ and the best-reply correspondence B_2 is convex-valued. Assumption (b) holds because the expected payoff $a(x, y)$ is linear in y . □

Another case in which (a) and (b) are trivially satisfied is when $E(x)$ is always a singleton. For example, in a two-player Cournot game, the follower's best reply is often unique. The following remark shows that Theorem 2 can fail without condition (b).

Remark 4 *There is a two-player game with continuous payoffs and convex-valued best-reply correspondences so that $E(x)$ is convex for all x , but L is less than any Nash payoff.*

Proof. Suppose that the strategy sets of the two players are $X = Y = [-1, 1]$, that the payoff to player 2 is constant (so any y is a best reply), and that the payoff to player 1 is $a(x, y) = xy - y^2$. Both players' payoffs are linear in their own strategy variable, so the best-reply correspondences are convex-valued. We have $E(x) = Y$ for all x , but $a(x, y)$ is not convex in y .

For any $y \in [-1, 1]$, we have a Nash equilibrium (x, y) if x is a best reply to y , that is, $x = -1$ if $y < 0$, arbitrary x for $y = 0$, and $x = 1$ for $y > 0$. The resulting Nash payoff to player 1 is $|y| - y^2$ which is $|y|(1 - |y|)$ and therefore always nonnegative.

The lowest leader payoff, however, results from player 2 always choosing the worst action for player 1 (which player 1 cannot force to be to his liking because player 2 is

indifferent). This worst reply y , which defines $F(x)$, is $y = 1$ if $x < 0$, it is $y \in \{1, -1\}$ if $x = 0$, and $y = -1$ if $x > 0$, so that the resulting payoff is $x - 1$ if $x \leq 0$ and $-x - 1$ if $x \geq 0$, that is, $-1 - |x|$, which is always negative. The best possible case is $x = 0$ where the leader payoff is -1 , worse than any Nash payoff. \square

In the example in Remark 4, player 2 has a constant payoff, so the leadership game where the follower chooses the worst payoff to the leader is effectively a zero-sum with payoff $xy - y^2$ to the leader. In this game, the max-min value is less than the min-max value. Hence, leadership is disadvantageous compared to playing simultaneously.

A similar situation arises in mixed extensions of games with more than two players. An example are *team games* as investigated by von Stengel and Koller (1997), where k players form a team and receive identical payoffs, which are the negative of the payoffs to player 1. Here, commitment generally hurts player 1 since it allows the opposing team players to coordinate their actions, which is not the case in the simultaneous game.

Remark 5 *There is a mixed extension of a finite three-player game so that L in (5) is less than any Nash payoff.*

	3		
2		p	q
P		-1, 1, 1	0, 0, 0
Q		0, 0, 0	-4, 4, 4
1 :		l	

	3		
2		p	q
P		-4, 4, 4	0, 0, 0
Q		0, 0, 0	-1, 1, 1
1 :		r	

Figure 3 Game between player 1 against the team of player 2 and 3 which has leader payoffs that are worse than any Nash payoff.

Proof. Consider Figure 3. Player 1 chooses the left (l) or right (r) panel, and players 2 and 3 form the team and have two strategies each. The Nash equilibria in this game are the pure equilibria (l, P, p) and (r, Q, q) , both with payoffs $(-1, 1, 1)$, the semi-mixed equilibria $(l, (0.8, 0.2), (0.8, 0.2))$ and $(r, (0.2, 0.8), (0.2, 0.8))$, both with payoffs $(-0.8, 0.8, 0.8)$, and the completely mixed equilibrium $((0.5, 0.5), (0.5, 0.5), (0.5, 0.5))$ with payoffs $(-1.25, 1.25, 1.25)$.

Suppose that the leader commits to the mixed strategy $(1 - x, x)$ in the leadership game. In the correspondence F in (2), players 2 and 3 can coordinate to play their favorable response, namely (Q, q) with payoffs $(3x - 4, 4 - 3x, 4 - 3x)$ if $0 \leq x \leq 0.5$ and (P, p) with payoffs $(-1 - 3x, 1 + 3x, 1 + 3x)$ if $0.5 < x \leq 1$ (for $x = 0.5$ the choice between (Q, q) and (P, p) is arbitrary). The optimal commitment is then $x = 0.5$. This defines a subgame perfect equilibrium with leader payoff $L = -2.5$, which is much worse for player 1 than in any Nash equilibrium of the simultaneous game. \square

In addition to Remark 5, note that the highest leader payoff in the game in Figure 3 is -0.8 , which results when the leader commits to either l or r and the followers respond by playing the mixed equilibrium in the corresponding panel in Figure 3. Because this defines also a Nash equilibrium, the highest leader payoff is not higher than the highest Nash payoff.

The game in Figure 3 is non-generic. However, the same arguments apply for any other generic game with payoffs nearby.

Our next observation concerns the highest leader payoff, for any number of followers and our standard assumptions from Section 2.

Proposition 6 *Let*

$$H = \max_{x \in X} \max_{y \in E(x)} a(x, y). \quad (9)$$

Then H is the highest leader payoff, and $H \geq a(x^, y^*)$ for any Nash equilibrium (x^*, y^*) of the simultaneous game.*

Proof. Clearly, $y^* \in E(x^*)$, which implies

$$H \geq \max_{y \in E(x^*)} a(x^*, y) \geq a(x^*, y^*). \quad \square$$

Next, we note that under certain assumptions the set of leader payoffs is an interval $[L, H]$.

Proposition 7 *Suppose that $E(x)$ is connected for all x in X . Then any payoff in $[L, H]$ with L and H as in (5) and (9) is a possible leader payoff.*

Proof. Let

$$\hat{x} \in \arg \max_{x \in X} \max_{y \in E(x)} a(x, y),$$

so that $H = \max_{y \in E(\hat{x})} a(\hat{x}, y)$. Clearly,

$$\min_{y \in E(\hat{x})} a(\hat{x}, y) \leq \sup_{x \in X} \min_{y \in E(x)} a(x, y) = L,$$

so let $\hat{y} \in E(\hat{x})$ with $a(\hat{x}, \hat{y}) \leq L$. Because $E(\hat{x})$ is connected, and $a(\hat{x}, y)$ is continuous in y , the set $\{a(\hat{x}, y) \mid y \in E(\hat{x})\}$ contains all reals in $[a(\hat{x}, \hat{y}), H]$ and hence in $[L, H]$. Any $a(\hat{x}, \bar{y}) \in [L, H]$ is a leader payoff in the subgame perfect equilibrium (\hat{x}, f) where the followers' response is $f(\hat{x}) = \bar{y}$ and $f(x) \in F(x)$ as in (2) for $x \neq \hat{x}$. \square

The following corollary summarizes our main results for mixed extensions of bimatrix games. An explicit characterization of L and H is given in the next section.

Corollary 8 *For the mixed extension of a bimatrix game, the set of leader payoffs is an interval $[L, H]$ with L and H as in (5) and (9). If l and h are the lowest and highest Nash payoff to the leader in the simultaneous game, then $l \leq L$ and $h \leq H$.*

Proof. The set $E(x)$ is the set of mixed strategies of player 2 that are best replies to x , which is connected, so Proposition 7 applies. \square

4 Leadership in mixed extensions of bimatrix games

In this section, we consider the mixed extension of a bimatrix game. We explicitly characterize the lowest and highest leader payoff L and H in Corollary 8 and show how to compute them by linear programming. For generic bimatrix games, we show that $L = H$.

We consider a bimatrix game with $m \times n$ matrices A and B of payoffs to player 1 and 2, respectively. The players' sets of pure strategies are

$$M = \{1, \dots, m\}, \quad N = \{1, \dots, n\}.$$

Their sets of mixed strategies are denoted by X and Y . For mixed strategies x and y , we want to write expected payoffs as matrix products xAy and xBy , so that x should be a row vector and y a column vector. That is,

$$X = \{ (x_1, \dots, x_m) \mid \forall i \in M \ x_i \geq 0, \sum_{i \in M} x_i = 1 \}$$

and

$$Y = \{ (y_1, \dots, y_n)^\top \mid \forall j \in N \ y_j \geq 0, \sum_{j \in N} y_j = 1 \}$$

As elements of X , the pure strategies of player 1 are the unit vectors, which we denote by e_i for $i \in M$.

For any pure strategy j of player 2, the payoffs to both players depending on $x \in X$ will be of interest. We denote the columns of the matrix A by A_j and those of B by B_j ,

$$A = [A_1 \cdots A_n], \quad B = [B_1 \cdots B_n]. \quad (10)$$

An inequality between two vectors, for example $B_j < B_y$ for some $j \in N$ and $y \in Y$ (which states that the pure strategy j is strictly dominated by the mixed strategy y), is understood to hold in each component.

For j in N , we denote by $X(j)$ the *best reply region* of j . This is set of those x in X to which j is a best reply:

$$X(j) = \{ x \in X \mid \forall k \in N - \{j\} \ xB_j \geq xB_k \}. \quad (11)$$

Let $X^\circ(j)$ denote the interior of $X(j)$ relative to X . Call $X(j)$ *full-dimensional* if $X^\circ(j)$ is not empty. Any best reply region $X(j)$ is a closed convex polytope. If it is full-dimensional, then

$$\overline{X^\circ(j)} = X(j). \quad (12)$$

Because any x in X has at least one best reply j , we also have

$$X = \bigcup_{j \in N} X(j). \quad (13)$$

In the example in Figure 1, we can identify X with the interval $[0, 1]$ for the probability that player 1 plays the bottom row B . Then Figure 2(a) shows that the best reply region of columns a and b is $[0, 2/3]$, of c is $[2/3, 1]$, of d is $\{2/3\}$ and of e is $\{0\}$. The best reply regions of d and e are not full-dimensional.

Theorem 9 Consider the mixed extension of a bimatrix game (A, B) with A_j, B_j as in (10) and $X(j)$ as in (11), for $j \in N$. Let $D = \{j \in N \mid X(j) \text{ is full-dimensional}\}$. Then the interval $[L, H]$ of all leader payoffs in Corollary 8 is given by

$$L = \max_{j \in D} \max_{x \in X(j)} \min_{k \in N: B_k = B_j} xA_k, \quad H = \max_{j \in N} \max_{x \in X(j)} xA_j. \quad (14)$$

We shall first explain Theorem 9 with our example, discuss the easy parts of its proof, and give the full proof afterwards. First consider H in (14). If the leader chooses x in $X(j)$, then the follower can respond with a mixed strategy that assigns positive probability to j , in particular the pure strategy j itself. Then xA_j is the expected payoff to the leader. The highest payoff to the leader is certainly attained with a pure strategy of the follower, and the expression for H in (14) is the highest leader payoff. In the example, Figure 2(b) shows that $H = 7$ where the leader commits to $\text{prob}(B) = 0$ and the follower responds with e .

The characterization of L in (14) is more involved. It states that for the lowest leader payoff L , best reply regions that are not full-dimensional can be ignored. In the example, this concerns both the best reply region of d which has a low payoff to the leader, and the best reply region of e which has a high payoff. The full-dimensional regions of a and b are identical because their two payoff columns are identical, denoted by $B_k = B_j$ in (14), with k, j standing for the columns a, b . Similar to our comments after Remark 4, the leader plays essentially a zero-sum game against the follower on regions $X(j)$ with multiple best replies of the follower, which leads to the inner max-min expression for L in (14). As Figure 2(b) shows, that max-min payoff for the region with best replies a, b is 4 and attained when the leader commits to $\text{prob}(B) = 1/3$, and for the best reply region of c that payoff is 5 and attained for $\text{prob}(B) = 2/3$; this is the lowest leader payoff L . Note that the highest and lowest leader payoff are not obtained for the same commitment of the leader. However, it may happen that the same commitment is used, for example if strategy e was absent in this game, in which case $H = 6$ (shown as a small triangle in Figure 2(b)) where the leader commits to $\text{prob}(B) = 2/3$ and the follower responds with a .

Proof of Theorem 9. In (9), $E(x)$ is the set of best replies to x . Among them, the maximum is already attained for the pure strategies j so that $x \in X(j)$. Hence,

$$H = \max_{x \in X} \max_{y \in E(x)} xAy = \max_{j \in N} \max_{x \in X(j)} xA_j$$

as claimed in (14).

It remains to prove the expression for L . First, we show

$$X = \bigcup_{j \in D} X(j). \quad (15)$$

To see this, let $k \in N - D$, and consider the open set $S = X - \bigcup_{j \in N - \{k\}} X(j)$. Then by (13), S is a subset of $X(k)$ and hence of the set $X^\circ(k)$, which is empty because $k \notin D$, so S is empty. This shows $X = \bigcup_{j \in N - \{k\}} X(j)$ which we now use instead of (13), and continue in this manner for the elements of $N - D$ other than k , to eventually obtain (15).

Secondly, for $j, k \in N$,

$$x \in X^\circ(j) \text{ and } x \in X(k) \implies B_k = B_j. \quad (16)$$

To see this, let $x \in X^\circ(j)$ and $x \in X(k)$. For all $i \in M$ consider the points $z_i = (1 - \varepsilon)x + \varepsilon e_i$ obtained by moving from x in the direction of the unit vectors e_i . The points z_i also belong to $X^\circ(j)$ for sufficiently small $\varepsilon > 0$. By representing x as a convex combination of z_1, \dots, z_m , we prove that not only j but also k is a best reply to z_i for each $i \in M$: Clearly, $z_i B_j \geq z_i B_k$. Suppose that $z_i B_j > z_i B_k$ for some i . Then $x = (x_1, \dots, x_m) = x_1 z_1 + \dots + x_m z_m$ and $x B_j = x_1 z_1 B_j + \dots + x_m z_m B_j > x B_k$ because $x_i > 0$ (since x is in the interior of X), in contradiction to $x \in X(k)$. So $z_i B_j = z_i B_k$ for all $i \in M$, and of course $x B_j = x B_k$. That is, $((1 - \varepsilon)x + \varepsilon e_i) B_j = ((1 - \varepsilon)x + \varepsilon e_i) B_k$, and therefore $e_i B_j = e_i B_k$, for all $i \in M$. So the column vectors B_j and B_k agree in all components, as claimed.

Using (5), and the fact that minimum payoffs are already obtained among pure best replies, and (15),

$$L = \sup_{x \in X} \min_{y \in E(x)} x A y = \sup_{x \in X} \min_{k \in N: x \in X(k)} x A_k = \max_{j \in D} \sup_{x \in X(j)} \min_{k \in N: x \in X(k)} x A_k. \quad (17)$$

Changing $X(j)$ to the smaller set $X^\circ(j)$ and using (16), we get

$$L \geq \max_{j \in D} \sup_{x \in X^\circ(j)} \min_{k \in N: x \in X(k)} x A_k = \max_{j \in D} \sup_{x \in X^\circ(j)} \min_{k \in N: B_k = B_j} x A_k.$$

Given j , the function on the right is the minimum of a fixed finite set of linear functions and therefore continuous. By (12), we obtain

$$\begin{aligned} L &\geq \max_{j \in D} \sup_{x \in X^\circ(j)} \min_{k \in N: B_k = B_j} x A_k = \max_{j \in D} \sup_{x \in X(j)} \min_{k \in N: B_k = B_j} x A_k \\ &\geq \max_{j \in D} \sup_{x \in X(j)} \min_{k \in N: x \in X(k)} x A_k = L \end{aligned}$$

where the last inequality holds because the minimum is taken over a larger set of pure strategies k ; the last equation is just (17). So all inequalities hold as equalities, giving

$$L = \max_{j \in D} \sup_{x \in X(j)} \min_{k \in N: B_k = B_j} x A_k = \max_{j \in D} \max_{x \in X(j)} \min_{k \in N: B_k = B_j} x A_k$$

as claimed in (14). \square

Next, we show how to compute L and H in (14) by linear programming. In contrast to the problem of finding all Nash equilibria of the simultaneous game, the leadership game is therefore easy to solve computationally. In addition, Corollary 8 provides quickly computable bounds on the Nash payoffs.

The computation of H in (14) is straightforward: For each $j \in N$, solve the linear program

$$\max \{x A_j \mid x \in X(j)\} \quad (18)$$

where $X(j)$ is the polyhedral set in (11) defined by the constraints $x \geq 0$, $\sum_{i \in M} x_i = 1$, and $x B_j \geq x B_k$ for $k \neq j$. If these are infeasible, then $X(j)$ is empty and j is never a best reply. The maximum over j with nonempty $X(j)$ of the values obtained in (18) is H .

The following proposition shows how to find L in (14).

Proposition 10 *The lowest leader payoff L in Theorem 9 is computed as follows:*

- (a) *For each $j \in N$, identify the set $N_j = \{k \in N \mid B_k = B_j\}$.*
- (b) *For each $j \in N$, we have $j \in D$ if and only if the following linear program has strictly positive value:*

$$\max\{\varepsilon \mid x \in X, xB_j \geq xB_k + \varepsilon \ (k \in N - N_j)\}. \quad (19)$$

- (c) *For each $j \in D$, the max-min expression for L in (14) is the value u_j of the following linear program:*

$$\max_{x \in X(j)} \min_{k \in N_j} xA_k = u_j = \max\{u \mid xA_k \geq u \ (k \in N_j), x \in X(j)\}, \quad (20)$$

giving $L = \max_{j \in D} u_j$.

Proof. Finding N_j as in (a) is trivial. According to (16), if $X(j)$ has nonempty interior, then any point x in $X^\circ(j)$ has best reply k only if $k \in N_j$, so (19) has a solution (ε^*, x^*) with $\varepsilon^* > 0$. Conversely, $x^*B_j > x^*B_k$ for $k \in N - N_j$ implies that these inequalities hold also for any x in a neighborhood of x^* , that is, for some $x \in X^\circ(j)$. This shows (b).

If $j \in D$ and $x \in X(j)$, then $xA_k \geq u$ for all $k \in N_j$ is equivalent to $\min_{k \in N_j} xA_k \geq u$, and the largest u subject to these inequalities is equal to $\min_{k \in N_j} xA_k$. Furthermore, the maximum such u for all $x \in X(j)$ equals u_j in (20), which shows that L in (14) is $\max_{j \in D} u_j$ as claimed in (c). \square

We have proved that the lowest leader payoff L is at least as high as some Nash payoff (Theorem 2) with the help of Kakutani's fixed point theorem. For the mixed extension of a bimatrix game, one can instead consider for each j in D the "constrained game" (Charnes 1953) where player 1 chooses x in $X(j)$ and player 2 mixes over the set N_j , with zero-sum payoff columns A_k to player 1 for $k \in N_j$. The respective minmax strategies for player 2, for each j obtained as the solution to the dual linear program to (20), can be combined to give a Nash equilibrium with payoff at most L ; for details see von Stengel and Zamir (2004, Theorem 11).

Most of our observations simplify drastically for generic bimatrix games. Generically, payoff matrices do not have identical columns, and best reply regions are either empty or full-dimensional. This is asserted in the following proposition. We use "generic" in the sense that any statement about the game holds also for any game with payoffs sufficiently nearby. For bimatrix games, an explicit alternative definition is *nondegeneracy*, which says that no mixed strategy has more pure best replies than the size of its support (see von Stengel, 2002); nondegeneracy is a generic property.

Proposition 11 *For the mixed extension of a generic bimatrix game, the lowest and highest leader payoff coincide, that is, $L = H$.*

Proof. Consider a generic bimatrix game (A, B) . Then clearly $B_k \neq B_j$ for any $k \neq j$ and in Proposition 10 we have $N_j = \{j\}$ for $j \in N$. By (14), $L = \max_{j \in D} \max_{x \in X(j)} xA_j$.

We claim that any best reply region $X(j)$ is either empty or full-dimensional. To see this, consider the optimal value ε^* of the linear program in (19). If $\varepsilon^* > 0$, then $X(j)$ is full-dimensional (so $j \in D$), and if $\varepsilon^* < 0$, then $X(j)$ is empty because j is never a best reply. The case $\varepsilon^* = 0$ does not hold generically, because the constraints in (19) are independently defined by the payoff columns in B , and the hyperplane defined by the optimal value of ε would change (to positive or negative if the optimum was zero) for any slight variation of a suitable payoff. Because pure strategies where $X(j)$ is empty can be omitted, we have $L = \max_{j \in D} \max_{x \in X(j)} xA_j = \max_{j \in N} \max_{x \in X(j)} xA_j = H$ as claimed. \square

5 Correlated equilibria

In this section, we consider correlated equilibria (Aumann 1974) for bimatrix games. We first show that the highest leader payoff H as defined in (14) is greater than or equal to the highest correlated equilibrium payoff to the leader. This strengthens Corollary 8. Trivially, the lowest leader payoff L in (14) is at least as high as some correlated payoff, because it is at least as high as some Nash payoff.

We consider the *canonical form* of a correlated equilibrium, which is a distribution on strategy pairs. With the notation of the previous section, this is an $m \times n$ matrix z with nonnegative entries z_{ij} for $i \in M$, $j \in N$ that sum to one. They have to fulfill the *incentive constraints* that for all $i, k \in M$ and all $j, l \in N$,

$$\sum_{j \in N} z_{ij} a_{ij} \geq \sum_{j \in N} z_{ij} a_{kj}, \quad \sum_{i \in M} z_{ij} b_{ij} \geq \sum_{i \in M} z_{il} b_{il}. \quad (21)$$

When a strategy pair (i, j) is drawn with probability z_{ij} according to this distribution by some device or mediator, player 1 is told i and player 2 is told j . The first constraints in (21) state that player 1, when recommended to play i , has no incentive to switch from i to k , given (up to normalization) the conditional probabilities z_{ij} on the strategies j of player 2. Analogously, the second inequalities in (21) state that player 2, when recommended to play j , has no incentive to switch to l .

Theorem 12 *In the mixed extension of a bimatrix game, the highest leader payoff H in (14) is greater than or equal to any correlated equilibrium payoff to the leader.*

Proof. Assume that the leader is player 1. Consider a correlated equilibrium z with probabilities z_{ij} fulfilling (21) above. Define the marginal probabilities on N by

$$y_j = \sum_{i \in M} z_{ij} \quad \text{for } j \in N, \quad (22)$$

and let S be the support of this marginal distribution, $S = \{j \in N \mid y_j > 0\}$. For each j in S , let c_j be the conditional expected payoff to player 1 given that player 2 is recommended to (and does) play j ,

$$c_j = \sum_{i \in M} z_{ij} a_{ij} / y_j.$$

Finally, let s in S be a strategy so that $c_s = \max_{j \in S} c_j$.

We claim that $H \geq c_s$, and that c_s is greater than or equal to the payoff to player 1 in the correlated equilibrium z , which proves the theorem. To see this, define x in X by $x_i = z_{is}/y_s$ for $i \in M$. Let player 1 commit to x in the leadership game. Then s is a best reply to x by player 2's incentive constraints (21) for $j = s$, multiplied by the normalization factor $1/y_s$. The corresponding payoff $x A_s$ to player 1 is c_s . This may not necessarily define a leadership equilibrium since player 1 may possibly improve his payoff by a different commitment. At any rate, the payoff c_s to player 1 when leader and follower play as described fulfills $c_s \leq H$. Furthermore, the correlated equilibrium payoff to player 1 is an average of the conditional payoffs c_j for $j \in S$ and therefore not higher than their maximum c_s :

$$\sum_{j \in N, i \in M} z_{ij} a_{ij} = \sum_{j \in S, i \in M} y_j z_{ij} a_{ij} / y_j = \sum_{j \in S} y_j c_j \leq c_s \leq H,$$

as claimed. □

We conclude this paper by considering a generalization of correlated equilibria which involves a commitment by both players. This is the “simple extension” of a correlated equilibrium defined by Moulin and Vial (1978, p. 203), which, following Young (2004), we call *coarse correlated equilibrium*. We show that such a coarse correlated equilibrium may give a payoff to the leader which is higher than any leader payoff in the leadership game.

A coarse correlated equilibrium is given by a distribution z on strategy profiles, which are chosen according to this commonly known distribution by a mediator. Each player must decide either to be told the outcome of the lottery z and to *commit* himself to playing the recommended strategy, or not to be told the outcome and play some mixed strategy. In equilibrium, the players commit themselves to playing the mediator's recommendation, and do not gain by unilaterally choosing not to be told the recommendation. So a unilaterally deviating player knows only the *marginal probabilities* under z of the choices of the other players. For two players, the respective inequalities are, for all $k \in M$ and $l \in N$,

$$\sum_{i,j} z_{ij} a_{ij} \geq \sum_j \left(\sum_i z_{ij} \right) a_{kj}, \quad \sum_{i,j} z_{ij} b_{ij} \geq \sum_i \left(\sum_j z_{ij} \right) b_{il}. \quad (23)$$

These inequalities are obviously implied by the incentive constraints (21); that is, any correlated equilibrium according to Aumann fulfills (23).

Remark 13 *The payoff to a player in a coarse correlated equilibrium of a two-player game can be higher than any leader payoff in the corresponding mixed extension of the game.*

Proof. Figure 4 shows a variation of the “paper–scissors–rock” game. This game is symmetric between the two players, and does not change under any cyclic permutation of the three strategies. The players' strategies beat each other cyclically, inflicting a loss

		2		
		p	q	r
1	P	0	1	-2
	Q	-2	0	1
	R	1	-2	0

Figure 4 Game with payoff 0 in a coarse correlated equilibrium, which is higher than any leader payoff.

-2 on the loser which exceeds the gain 1 for the winner. The game has a unique mixed Nash equilibrium where each strategy is played with probability $1/3$ and each player gets expected payoff $-1/3$.

For the game in Figure 4, one coarse correlated equilibrium with payoff $(0,0)$ is a lottery that chooses each of (P,p) , (Q,q) and (R,r) with probability $1/3$, and any other pure strategy pair with probability zero. This fulfills (23), but is not a correlated equilibrium.

In the leadership game for Figure 4, it suffices to consider only one best reply region, say for the first strategy p of player 2. The best reply region for p is the convex hull of the points (in X , giving the probabilities for P, Q, R), $(1/3, 1/3, 1/3)$, $(3/4, 0, 1/4)$, $(0, 1/4, 3/4)$, and $(0, 0, 1)$, with respective payoffs $-1/3$, $-1/2$, $-5/4$, and -2 to player 1. The maximum of these leader payoffs is therefore $-1/3$, which is the same for any best reply region because of the symmetry in the three strategies. In this game, leader and Nash payoff coincide. By Theorem 12, the highest correlated equilibrium payoff is also $-1/3$, which is also the lowest correlated equilibrium payoff since it is the maxmin payoff. \square

We have shown that in the game in Figure 4, there is a coarse correlated equilibrium which gives a payoff which is higher than the (unique) leader payoff of the mixed extension of the game. The coarse correlated equilibrium concept involves a commitment by *both* players to a correlated device. However, this concept does not generalize the subgame perfect equilibrium of a leadership game, because it has correlated and Nash equilibria of the simultaneous game as special cases, whereas leadership payoffs are generically unique.

References

Amir, R. (1995), Endogenous timing in two-player games: A counterexample. *Games and Economic Behavior* 9, 234–237.

- Amir, R., Grilo, I. (1999), Stackelberg versus Cournot equilibrium. *Games and Economic Behavior* 26, 1–21.
- Aumann, R. J. (1974), Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1, 67–96.
- Avenhaus, R., Okada, A., Zamir, S. (1991), Inspector leadership with incomplete information. In: *Game Equilibrium Models IV*, ed. R. Selten, Springer, Berlin, 319–361.
- Avenhaus, R., von Stengel, B., Zamir, S. (2002), Inspection games. In: *Handbook of Game Theory with Economic Applications*, Vol. 3, eds. R. J. Aumann and S. Hart, Elsevier, Amsterdam, 1947–1987.
- Başar, T., Olsder, G. J. (1982), *Dynamic Noncooperative Game Theory*. Academic Press, London.
- Charnes, A. (1953), Constrained games and linear programming. *Proc. National Academy of Sciences of the U.S.A.* 39, 639–641.
- Cournot, A. A. (1938), *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. Hachette, Paris.
- Friedman, J. W. (1977), *Oligopoly and the Theory of Games*. North-Holland, Amsterdam.
- Hamilton, J., Slutsky, S. (1990), Endogenous timing in duopoly games: Stackelberg or Cournot equilibria. *Games and Economic Behavior* 2, 29–46.
- Kakutani, S. (1941), A generalization of Brouwer’s fixed point theorem. *Duke Math. J.* 8, 457–459.
- Maschler, M. (1966), A price leadership method for solving the inspector’s non-constant-sum game. *Naval Research Logistics Quarterly* 13, 11–33.
- Morgan, J., Patrone, F. (2006), Stackelberg problems: Subgame perfect equilibria via Tikhonov regularization. In: *Advances in Dynamic Games*, eds. A. Haurie et al., *Annals of the International Society of Dynamic Games*, Vol. 8, Birkhäuser, Boston, 209–221.
- Moulin, H., Vial, J.-P. (1978), Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory* 7, 201–221.
- Nash, J. F. (1950), Equilibrium points in N -person games. *Proc. National Academy of Sciences of the U.S.A.* 36, 48–49.
- Reny, P. J., Robson, A. J. (2004), Reinterpreting mixed strategy equilibria: A unification of the classical and Bayesian views. *Games and Economic Behavior* 48, 355–384.
- Rubinstein, A. (1982), Perfect equilibrium in a bargaining model. *Econometrica* 50, 97–109.
- Schelling, T. C. (1960), *The Strategy of Conflict*. Harvard Univ. Press, Cambridge, Mass.
- Shapiro, C. (1989), Theories of oligopoly behavior. In: *Handbook of Industrial Organization*, Vol. I, eds. R. Schmalensee and R. D. Willig, North-Holland, Amsterdam, 329–414.
- von Neumann, J., Morgenstern, O. (1947), *Theory of Games and Economic Behavior*, 2nd ed. Princeton Univ. Press, Princeton, NJ.
- von Stackelberg, H. (1934), *Marktform und Gleichgewicht*. Springer, Vienna.
- von Stengel, B. (2002), Computing equilibria for two-person games. In: *Handbook of Game Theory*, Vol. 3, eds. R. J. Aumann and S. Hart, North-Holland, Amsterdam, 1723–1759.
- von Stengel, B., Koller, D. (1997), Team-maxmin equilibria. *Games and Economic Behavior* 21, 309–321.

von Stengel, B., Zamir, S. (2004), Leadership with Commitment to Mixed Strategies. Research Report LSE-CDAM-2004-01, London School of Economics.

Young, H. P. (2004), *Strategic Learning and its Limits*. Oxford University Press, Oxford.