Leadership with Commitment to Mixed Strategies*

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
CNRS, Maison des Sciences Economiques, 106/112 boulevard de l'Hôpital,
75647 Paris Cedex 13, France,
and
Center for Rationality, The Hebrew University at Jerusalem,
Givat Ram, Jerusalem 91904, Israel
email: zamir@math.huji.ac.il

CDAM Research Report LSE-CDAM-2004-01

February 5, 2004; revised August 30, 2004

Abstract

A basic model of commitment is to convert a game in strategic form to a "leadership game" with the same payoffs, where one player commits to a strategy to which the second player chooses a best reply. This paper studies such leadership games for the mixed extension of a finite game, where the leader commits to a mixed strategy. The set of leader payoffs is an interval (for generic games a singleton), which is at least as good as the set of that player's Nash and correlated equilibrium payoffs in the simultaneous game. This no longer holds for leadership games with three or more players.

Keywords: Bimatrix game, commitment, correlated equilibrium, follower, leader, mixed strategy, Nash equilibrium, Stackelberg solution.

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.

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. In modern parlance, Cournot found a Nash equilibrium in a game where firms choose their quantities *simultaneously*. The *leadership game* of von Stackelberg uses the same payoff functions, but where one firm, the *leader*, moves first, assuming a best reply of the second-moving firm, the *follower*. The "Stackelberg solution" is then a *subgame perfect* equilibrium of this sequential game. The leader-follower issue has been studied in depth in oligopoly theory; see Friedman (1977), Hamilton and Slutsky (1990), Shapiro (1989), or Amir and Grilo (1999) for discussions and references.

This paper studies the leadership game for the *mixed extension* of a finite strategic-form game, one of the most basic models of noncooperative game theory. We provide a complete analysis of two-player games, including nongeneric cases, showing that the possibility to commit never hurts a player. Further results, explained later in this introduction, compare leadership and correlated equilibria, concern games with more than two players, and study the possible follower payoffs.

Our basic setting is to compare the simultaneous version of a two-player game with the corresponding leadership game. In the simultaneous game, both players choose their actions independently and simultaneously, possibly by randomizing with a mixed strategy, which is in general necessary for the existence of a Nash equilibrium. In the leadership game, the leader, player I, say, commits to a mixed strategy x. The follower, player II, is fully informed about x, and chooses her own action, possibly by randomization, with a pure or mixed strategy $y(x)$. The pair of pure actions, and corresponding payoffs, is then chosen independently according to x and $y(x)$ as in the original game. We only consider subgame perfect equilibria of the leadership game where the follower chooses only best replies $y(x)$ against any x, even off the equilibrium path. The set of equilibria that are not subgame perfect seems too large to allow any interesting conclusions. The payoff to the leader in a subgame perfect equilibrium of the leadership game will be called *leader payoff*, his payoff in the simultaneous game *Nash payoff*.

The main result comparing leadership and simultaneous game states that the leader payoff is not worse than the Nash payoff, so commitment power is beneficial. When best replies are *unique*, this is a near-trivial result (and has been observed earlier, e.g., by Simaan and Cruz (1973) or Başar and Olsder (1982, p. 126)): The leader can always commit to his Nash strategy and thus receive at least the payoff in the Nash equilibrium. If there are several Nash equilibria, the leader can choose the equilibrium with the largest payoff to him. The leader may even do better with a different commitment.

Best replies are typically not unique, however, in Nash equilibria with mixed strategies. If the follower's best reply is a properly mixed strategy, she may choose any of her pure best replies, which may be to the disadvantage of the leader. In a zero-sum game, von Neumann's minimax theorem asserts that the leader is not harmed by committing to his mixed strategy. This is a possible motivation for using mixed strategies in a zero-sum game, apart from existence of an equilibrium. Von Neumann and Morgenstern (1947) explicitly define the leadership game of a zero-sum game, first with commitment to pure (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. A commitment to pure strategies only may of course harm the leader, as in “matching pennies”, or any other, even non-zero-sum game with payoffs nearby.

Commitment to mixed strategies are also considered by Rosenthal (1991), who defines “commitment robust” Nash equilibria that remain equilibria in the leadership game. Landsberger and Monderer (1994) also treat commitment to mixed strategies, as discussed in the context of Figure 3 below. Compared to pure strategies, a commitment to mixed strategies is obviously harder to verify. Bagwell (1995) considers games with commitment to pure strategies only, but where the pure strategy is imperfectly observed by the follower. He notes that then the commitment effect vanishes since the leader would always renege on the commitment, given that the follower attributes the differently observed strategy to an erroneous observation. Van Damme and Hurkens (1997) note that the leadership advantage can be re-instated by considering *mixed* equilibria, still in the game where the leader can only commit to a pure strategy. Reny and Robson (2002) consider commitment to mixed strategies as a possible “classical” view of mixed strategies, as done by von Neumann and Morgenstern.

Even when best replies are not unique, the *existence* of a leader payoff that is at least as good as any Nash payoff is again obvious. Namely, the follower may simply respond as in a Nash equilibrium, or, even better, to the leader’s advantage when she is indifferent. In the context of inspection games, Maschler (1966) *postulates* the latter behavior of the follower, calling it “pareto-optimal”. This postulate is unnecessary, as observed by Avenhaus, Okada and Zamir (1991), since on the equilibrium path the follower must choose her best reply that is most favorable to leader, in order to obtain a subgame perfect equilibrium of the leadership game. This is argued in detail for Figure 2 below. Section 2 gives a number examples that illustrate this and other aspects of leadership games.

In generic two-player games, the leader payoff is unique and at least as large as any Nash payoff, as stated in Theorem 3 below. This is due to the fact that the follower’s best replies are *unique almost everywhere*, relative to the set of all mixed strategies of the leader. A favorable reply of the follower can thus be *induced* if necessary.

For nongeneric games, it cannot be true that all leader payoffs are at least as good as all Nash payoffs. The simplest example has one pure strategy for the leader and two best replies by the follower, with different payoffs to the leader. Either best reply defines a Nash equilibrium, commitment does not change the game, and one leader payoff is worse than the other Nash payoff. The obvious fair comparison should involve the *set* of payoffs to the leader. Indeed, we will prove that every leader payoff is at least as large as *some* Nash payoff.

More precisely, our first main result (Theorem 11 below) states: All subgame perfect equilibrium payoffs to the leader (player I) in the leadership game belong to an interval $[L,H]$. The highest possible leader payoff H is at least as high as any Nash equilibrium payoff to player I in the simultaneous game. The lowest possible leader payoff L is at least as high as the lowest Nash equilibrium payoff. In other words, the set of equilibrium payoffs “moves upwards” for the leader (or stays unchanged, as in a zero-sum game). If no pure strategy of the follower is weakly dominated by or payoff equivalent to a different (possibly mixed) strategy, then the leader’s payoff is unique ($L = H$). In particular, this is the case in a generic game. The mathematics for Theorem 11 are developed in Section 3. The three Theorems 3, 10, and 11 are of increasing generality, but are stated separately since they build on each other.

Section 4 shows that the highest (for a generic game, unique) leader payoff H is at least as high as any *correlated equilibrium payoff* of the simultaneous game. This is interesting because commitment can serve as a coordination device, which is a possible motive for considering correlated equilibria. (As shown before, the lowest leader payoff L is at least as large as some Nash equilibrium payoff, and hence as some correlated equilibrium payoff.) We also show that the largest payoff to player I in the “simple extension” of a correlated equilibrium due to Moulin and Vial (1978) may possibly *not* be obtained as a leader payoff. The Moulin–Vial coordination device requires a commitment by *both* players.

In Section 5, we show that commitment to mixed strategies may no longer be advantageous in games with more than two players. The games considered have one leader and k followers, who play an equilibrium among themselves in the subgame induced by the commitment of the leader. In games of a *team* of several players, with identical payoffs, who play a zero-sum game against an adversary (as studied by von Stengel and Koller (1997)), the adversary will never profit when made a leader. On the contrary, the set of subgame perfect equilibrium payoffs in the leadership game will “move downwards” compared to the simultaneous game. The reason is that the commitment helps the team of followers to coordinate their actions, to their advantage. This result holds also for a set of positive measure of generic games nearby.

Section 6 compares our approach to leadership equilibria with the related “Stackelberg” concept in dynamic games (see Başer and Olsder (1982), Mallozzi and Morgan (2002), Morgan and Patrone (2004), and references therein). In that literature, leadership is seen as an optimization problem. A typical “pessimistic” assumption is that if there is more than one best reply of a follower, or more than one induced equilibrium among k followers, the chosen reply is the *worst* one for the leader. The resulting payoff to the leader as a function of his commitment is typically discontinuous. We show in Theorem 13 that the resulting limit payoff is obtained in a subgame perfect equilibrium of the leadership game. The only difference to the view of optimization theory is that the followers do not (and cannot) act according to the described “pessimistic” view in the equilibrium itself. If the pessimistic view is adopted, the leader should choose a nearby, slightly suboptimal, commitment.

Leadership is advantageous when compared to the simultaneous game, but not necessarily compared to the follower’s situation. Section 7 addresses the payoff to the follower in leadership equilibria. We give an example of a symmetric 3×3 game where the leader payoff is better than the Nash payoff, but where the follower payoff may take any value, depending on a parameter of the game which leaves the best replies of the follower, and the optimal commitment, unchanged. In a separate paper (von Stengel (2003)), it is shown that in a symmetric duopoly game, as considered, for example, by Hamilton and Slutsky (1990), the follower is either worse off than in the simultaneous game, or even better off than the leader.

Section 8 concludes with possible topics for further research.

2 Examples

The well-known *ultimatum game* is a sequential game where player I first offers a split of a unit "pie" into the nonnegative amounts x and $1-x$ for player I and II, respectively, which player II then can accept, whereupon the players receive the payoffs x and $1-x$, or reject, in which case both players receive zero.

I \ II	l	r
T	0, 1	0, 0
B	1, 0	0, 0

Figure 1. Ultimatum game with the leader's demand as probability of B.

The ultimatum game can be cast as the leadership game of the game in Figure 1, with player I as the leader (as throughout the paper), whose own demand x is the probability of playing the bottom strategy B. The left column l means "accept" and r means "reject".

If x can be chosen continuously from the interval $[0,1]$, the unique subgame perfect equilibrium is a textbook example in bargaining (e.g., Binmore (1992, p. 199)): Player II accepts *any* split, even when she receives nothing, $1-x=0$, where she is indifferent between accepting and rejecting, and player I demands the whole unit for himself, $x=1$, on the equilibrium path. The reason is that by subgame perfection, the best reply of player II is to accept (l) whenever $x<1$. Then offering an amount x less than one, which is is player I's payoff, can never be an equilibrium choice since it could be improved to $x+\varepsilon$ with $0<\varepsilon<1-x$. Hence in a subgame perfect equilibrium, player I demands everything ($x=1$), and player II is indifferent. If, in reply to $x=1$, player II rejects with positive probability $y>0$, then player I would receive $1-y$ and could improve his payoff by changing his demand to $1-y/2$, say, which is not an equilibrium choice as argued before. So a subgame perfect equilibrium is given only if player II accepts with certainty, despite being indifferent on the equilibrium path. The same reasoning applies also to the subgame perfect equilibrium of the multiple-round bargaining game due to Rubinstein (1982).

In the simultaneous game shown in Figure 1, player I chooses B in any Nash equilibrium (a positive probability for T would entail the unique best reply l and, in turn, B), and player II can mix l and r with arbitrary probability y, say, for r. The resulting payoffs are $1-y$ for player I, which is any number in $[0,1]$, and 0 for player II. The leader payoff 1 (that is, the payoff to player I in the subgame perfect equilibrium of the leadership game) is at least as good as any Nash payoff to player I (in the simultaneous game).

The effect of leadership is also familiar in the context of *inspection games* (see Maschler (1966), Avenhaus and Canty (1996), Avenhaus, Okada, and Zamir (1991), Wölling (2002), and the survey by Avenhaus, von Stengel, and Zamir (2002)).

I \ II	l	r
T	0, 0	–10, 1
B	–1, 0	–6, –9

Figure 2. Inspection game.

Figure 2 shows a simple example of an inspection game where the inspector, player I, can choose not to inspect (T) or to inspect (B), and the inspectee, player II, can either comply with a legal obligation (l) or cheat (r). The reference strategy pair (T,l) defines the pair of payoffs $(0,0)$ for players I, II, which is the most desirable outcome for the inspector. In all other cases, negative payoffs to the inspector reflect his preference for compliance throughout, rather than catching an inspectee who cheats. An inspection is costly for the inspector, with payoffs $(-1,0)$ for (B,l). The inspectee gains from cheating without inspection, with payoffs $(-10,1)$ for (T,r), but loses when inspected, with payoffs $(-6,-9)$.

The unique Nash equilibrium of the simultaneous game in Figure 2 is in mixed strategies, where player I chooses to inspect (B) with probability $1/10$ and player II cheats (choosing r) with probability $1/5$. The resulting payoff pair is $(-2,0)$.

The leadership game for Figure 2 has a unique best reply by player II when the probability x for inspection is not equal to $1/10$: cheat for $x<1/10$, and comply for $x>1/10$. Since the inspector prefers the inspectee to comply (l) in any case, he will commit to a probability x with $x\geq 1/10$. Since the resulting payoff $-x$ to player I is decreasing in x, a subgame perfect equilibrium requires a commitment to the smallest probability $x=1/10$ for B where player II still responds with l. Then the follower is indifferent, but chooses the reply with the most favorable payoff to the leader, namely compliance. The reason is the same as before: Any positive probability for r would reduce the payoff $-1/10$ to the leader, which he could improve upon by committing to $x=1/10+\varepsilon$ and thus *induce* the follower to comply. As mentioned in the introduction, Maschler (1966) assumes that the inspectee complies when indifferent. Avenhaus, Okada, and Zamir (1991) note that this is the only subgame perfect equilibrium. For a more detailed discussion see Avenhaus, von Stengel, and Zamir (2002, Section 5).

The resulting leader payoff $-1/10$ in the leadership game for Figure 2 is much better for player I than his Nash payoff -2 in the simultaneous game. In the game of Figure 2, the leader commits to the same mixed strategy as in the unique Nash equilibrium of the simultaneous game (this holds for any 2×2 game with a unique completely mixed equilibrium, but is not true for larger games). The follower is indifferent, but chooses the favorable action (here l) for the leader in the leadership game.

Inspection games model a scenario where the inspector is a natural leader. An inspection policy can be made credible, whereas the inspectee cannot reasonably commit to a strategy that involves cheating. We do not try to “endogenize” leadership (as, for example, Hamilton and Slutsky (1990)), but assume that one of the players has commitment power, and study its effect.

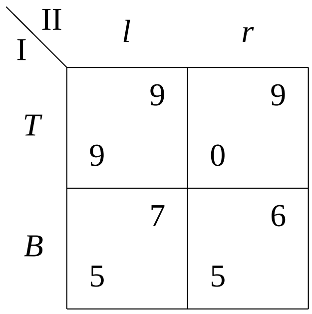

Figure 3. Game with a weakly dominated strategy of player II.

For the game in Figure 3, Landsberger and Monderer (1994) have argued that player I, when offered a choice to commit or not to commit (possibly to a mixed strategy), would choose not to commit, assuming the iterated elimination of weakly dominated strategies as a solution concept. In contrast, we shall argue that the followers’ preference for using only the undominated strategy can be enhanced, rather than weakened, with commitment.

The Nash equilibria of the game in Figure 3 are given by player I choosing T with certainty, and player II mixing between l and r, choosing r with some probability in $[0,4/9]$. The resulting Nash payoffs are any number in $[5,9]$ for player I, and 9 for player II.

In the leadership game corresponding to Figure 3, player I commits to T, and, by the usual reasoning, player II always responds by l, with resulting payoffs 9 for both players. This is arguably better for player I than the simultaneous game. The relationship of Nash to leader payoffs in this game is similar to Figure 1. If the leader has any doubt about the reply of l to his commitment T since the follower is indifferent, the leader can commit to playing B with a small probability ε in order to induce l as a unique best reply.

The game in Figure 4 is an interesting variation of Figure 3. Again, player II has a weakly dominating strategy, in this case r rather than l. The Nash equilibria of the simultaneous game are (B,r) with payoff pair $(5,7)$, or T and a mixture of l and r with

I \ II	l	r
T	9, 9	0, 9
B	5, 6	5, 7

Figure 4. Game with two possible commitments in the leadership game, T and B.

probability at most $4/9$ for r as in Figure 3, with payoffs in $[5,9]$ for player I and 9 for player II.

In the leadership game for Figure 4, the subgame perfect equilibria depend on the reply by the follower against a commitment by the leader to play T with certainty; to any other commitment, the follower responds by playing r. If the follower responds to T, where she is indifferent, with probability at most $4/9$ for r, then the leader gets a payoff in $[5,9]$ and the commitment to T is optimal. However, the leader cannot *induce* the follower to play l, which is preferred by the leader, as in the previous games, since any variation from this commitment will induce the reply r. Instead, the follower may indeed respond to T by choosing r with probability $4/9$ or higher. In that case, the leader maximizes his payoff by committing to B, with resulting payoff pair $(5,7)$. If the follower responds to T by choosing r with probability exactly $4/9$, both T and B are optimal commitments.

In the game of Figure 4, the sets of Nash and leadership payoffs coincide, for both players. The set of leader payoffs is an interval $[5,9]$, where any payoff greater than 5 to the leader depends on the "goodwill" of the follower since her reply to the commitment to T cannot be induced by changing the commitment slightly. The smallest leader payoff 5 cannot be induced by a commitment to T or a slight variation of this commitment, but requires a different commitment to the "remote" strategy B. This smallest leader payoff is found by ignoring pure replies of the follower that are weakly dominated. This is also done in the general proof of Theorem 10 below.

I \ II	l	r
T	2, 2	0, 1
B	3, 0	1, 1

Figure 5. Commitment to a strictly dominated strategy T by the leader.

Player II will never play a strictly dominated pure strategy, neither in the simultaneous game, nor in a subgame perfect equilibrium when she is the follower in a leadership game, so we can disregard such a strategy. In contrast, player I may commit to a strictly dominated strategy in the leadership game if it induces a reply that is favorable for him, as demonstrated by the game in Figure 5, sometimes called the "quality game" with T and B representing a good or bad service, and l and r the choices of a customer of buying or not buying the service. The simultaneous game has the unique Nash equilibrium (B,r) with payoffs $(1,1)$. In the leadership game, even committing only to a pure strategy would give the subgame perfect equilibrium with commitment to T and reply l, and payoffs $(2,2)$. The optimal commitment to a mixed strategy is to the mixture of T and B with probability $1/2$ each, where the follower is indifferent, but responds by choosing l for the usual reason that otherwise this reply can be induced by committing to play T with probability $1/2+\varepsilon$. The resulting payoff pair is $(5/2,1)$.

I \ II	l	c	r
T	0, 1	0, 1	0, 0
M	0, 0	8, 0	0, 1
B	8, 0	0, 0	1, 1

Figure 6. Game with "equivalent" replies l and r for the follower.

The example in Figure 6 illustrates a point that may arise with general nongeneric games. Here, the pure strategies l and c have identical payoffs for player II, so that player I can never induce player II to play one strategy or the other. The simultaneous game has the unique Nash equilibrium (B,r) with payoffs $(1,1)$, obtained, for example, by iterated elimination of strictly dominated strategies, first T, then l and r, then M.

In the leadership game for Figure 6, player I as leader can induce player II to play l or c by committing to play T with probability at least $1/2$. Against the commitment to the mixed strategy $(1/2,1/2,0)$ for (T,M,B), for example, the follower is indifferent among all her strategies. In particular, she may respond by playing c, giving the leader the good expected payoff 4, so this is a possible payoff to the leader in the leadership game (in fact, the maximum possible). However, she may also respond by playing l or r, with a much smaller payoff which the leader cannot change to a payoff near 4 by a commitment nearby.

In this game, the set of leader payoffs is again an interval, namely $[2,4]$. The lowest possible leader payoff 2 is not even found at an extreme point of a "best reply region" for either pure reply l, c or r (these best reply regions are crucial for the analysis of the leadership game, see Section 3). Rather, it results from a commitment to the mixed strategy $(1/2,1/4,1/4)$, where the follower is again indifferent between all replies. For l and c, the resulting expected payoff to the leader is 2, and this is in fact a subgame perfect equilibrium payoff if player II *randomizes*, with probability $1/2$ for both l and c (which she can always do whenever l and c are best replies, since these strategies are equivalent in the sense of having identical payoffs for her). Player I could also respond to $(1/2,1/4,1/4)$ with r, with payoff $1/4$ to the leader, but this would not be part of a subgame perfect equilibrium since player I could induce l or c as a best reply by committing to $(1/2+2\varepsilon,1/4-\varepsilon,1/4-\varepsilon)$, say.

Equivalent pure strategies like l and c in Figure 6 require the consideration of a *zero-sum game*, in terms of the payoffs of player I, played on the region of mixed strategies of player I where these equivalent strategies of player II are best replies; for details see Theorem 11 and its proof.

3 Leadership games

We consider a bimatrix game with $m\times n$ matrices A and B of payoffs to player I and II, respectively. The set of pure strategies of player I (matrix rows) is denoted by M and the set of pure strategies of player II (columns) by N,

$$M=\{1,\ldots,m\},\qquad N=\{1,\ldots,n\}.$$

The sets of mixed strategies of the two players are called 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,

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

and

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

As elements of X, the pure strategies of player I are the unit vectors, which we denote by e_i for $i\in M$, that is, $e_i\in X$ and

$$e_1=(1,0,\ldots,0),\quad e_2=(0,1,0,\ldots,0),\ \ \ldots,\quad e_m=(0,\ldots,0,1). \tag{1}$$

The *simultaneous game* is played with player I choosing x in X, player II choosing y in Y, and player I and II receiving payoffs xAy and xBy, respectively. A *Nash equilibrium* always means an equilibrium of the simultaneous game. A *Nash payoff* (to any player) is the payoff in any Nash equilibrium.

The *leadership game* is played with strategy set X for player I, called the *leader*, and any function $f\colon X\to Y$ as a strategy of player II, called the *follower*, whereupon the players receive payoffs $xAf(x)$ and $xBf(x)$, respectively. The interpretation of the leadership game is that the leader commits to the strategy x, about which the follower is fully informed, and can choose $f(x)$ in Y separately for each x. A *leadership equilibrium* is any *subgame perfect* equilibrium (x, f) of the leadership game, that is, $f(x')$ is a best reply to any x' in X, even if x' is not the strategy x that the leader commits to. A *leader payoff* is the payoff to the leader in any leadership equilibrium, a *follower payoff* the payoff to the follower in any leadership equilibrium.

For any pure strategy j of player II, both payoffs to player I and to player II 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,

$$A = [A_1 \cdots A_n], \qquad B = [B_1 \cdots B_n].$$

An inequality between two vectors, like, for example, $B_j < By$ 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, the *best reply region* $X(j)$ is the set of those x in X where j is a best reply to x:

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

Any best reply region $X(j)$ is a closed convex polytope, and

$$X = \bigcup_{j \in N} X(j) \tag{2}$$

since any x in X has at least one best reply $j \in N$.

The main power of commitment in a leadership game is that the leader can *induce* the follower to play certain pure strategies, by our assumption that the follower uses only best replies.

Definition 1 *A strategy j in N is called* inducible *if j is the unique best reply to some x in X.*

Let a_j denote the maximal payoff that player I can get when choosing x from $X(j)$:

$$a_j = \max\{ xA_j \mid x \in X(j) \}. \tag{3}$$

If j is inducible, then the leader can get at least payoff a_j in any leadership equilibrium:

Lemma 2 *If j is inducible, then any leader payoff is at least a_j as defined in* (3).

Proof. Consider a leadership equilibrium. Assume that the maximum a_j in (3) is taken at x in $X(j)$, that is, $a_j = xA_j$. Furthermore, let $x' \in X(j)$ be a strategy where j is the unique best reply of player II, so that $x'B_j > x'B_k$ for all k in N, $k \neq j$. Furthermore, $xB_j \geq xB_k$ for all k in N. Then j is also the unique best reply to the convex combination $x(\varepsilon) = (1-\varepsilon)x + \varepsilon x'$, for any $\varepsilon \in (0,1]$, since $X(j)$ is convex and payoffs are linear. By subgame perfection, the follower's reply to $x(\varepsilon)$ is j. The payoff to player I when playing

$x(\varepsilon)$ is then $(1-\varepsilon)xA_j + \varepsilon x'A_j$, which is arbitrarily close to a_j by choosing ε sufficiently small. So if player II's reply on the equilibrium path (that is, to the actual commitment of player I in equilibrium) gave player I a payoff less than a_j, then player I could switch to $x(\varepsilon)$ (with small ε) and thereby get a higher payoff, contradicting the equilibrium property. So the payoff to player I in the leadership equilibrium is at least a_j as claimed. □

If every strategy j of player II is inducible, then the preceding lemma gives an easy way to find a leadership equilibrium: Choose j in N such that a_j is maximal, and let the leader commit to x in $X(j)$ where a_j in (3) is attained, that is, xA_j is maximal. Then the follower responds so that player I receives indeed payoff xA_j. This is stated in the following theorem, which also asserts that the resulting leader payoff is at least as good as any Nash payoff to player I. The preceding proof of Lemma 2 explains this advantage of the leader, the player with "commitment power". Namely, any x such that $xA_j = a_j$ is typically an extreme point of $X(j)$, as x maximizes a linear function on the polytope $X(j)$, and so x typically belongs to several best reply regions $X(k)$. In that case, the follower is indifferent between several best replies, but nevertheless chooses the reply j that is best for the leader, since otherwise the leader would induce this desired reply of the follower by changing x slightly to $x(\varepsilon)$ where j is the unique reply, as argued repeatedly in the examples in Section 2. The following theorem has already been observed by Wölling (2002).

Theorem 3 *Suppose that every j in N is inducible. Then the leader payoff is uniquely given by $L = \max_{j\in N} a_j$, that is,*

$$L = \max_{j\in N} \max_{x\in X(j)} xA_j\,. \tag{4}$$

Furthermore, no Nash payoff to player I exceeds L.

Proof. Player I cannot obtain a higher payoff than L in a leadership equilibrium since player II always chooses best replies. On the other hand, player I will get L by choosing x so that the maximum in (4) is attained, according to Lemma 2.

Let (x,y) in $X\times Y$ be a Nash equilibrium of the simultaneous game. Then any pure strategy played with positive probability in y must be a best reply to x, that is, $y_j > 0$ implies $x \in X(j)$. Hence we have for the Nash payoff to player I

$$\sum_{j\in N}(xA_j)y_j = \sum_{j\in N,\, y_j>0}(xA_j)y_j \le \sum_{j\in N} a_j y_j \le \max_{j\in N} a_j = L\,.$$

□

The remainder of this section concerns the case that player II has strategies j that are not inducible. A strategy j is not inducible if it is payoff equivalent to or weakly dominated by a (different) pure or mixed strategy y in Y so that $B_j \le By$. Then whenever j is a best reply to some x in X, so is y and hence all the pure strategies chosen by y with positive probability. (We can assume that y does not choose j, that is, $y_j = 0$, since otherwise we can omit j from the mixture and play the other pure strategies in y with proportionally higher probabilities.) The following lemma shows that this is the only case where j is not inducible.

Lemma 4 *A strategy j is not inducible if and only if $B_j \leq By$ for some y in Y with $y_j = 0$.*

Proof. If j is payoff equivalent to or weakly dominated by the strategy y as stated, then y is a best reply whenever j is, so j is certainly never a unique best reply. To prove the converse, suppose that for no $y \in Y$ with $y_j = 0$ we have $B_j \leq By$. We find an x in X to which j is the unique best reply, using linear programming duality. Consider the following linear program (LP): maximize u with variables u and $y_k \geq 0$ for $k \in N$, $k \neq j$, so that

$$\begin{aligned} -\sum_{k \in N,\ k \neq j} B_k y_k + \mathbf{1}u &\leq -B_j \\ \sum_{k \in N,\ k \neq j} y_k \qquad &= 1, \end{aligned} \tag{5}$$

where $\mathbf{1}$ is a column of m ones. The LP (5) is feasible, with any $y \in Y$ with $y_j = 0$ and sufficiently negative u, and its optimal value u is negative since otherwise $B_j \leq By$. The dual of the LP (5) is to find nonnegative $x = (x_1, \ldots, x_m)$ and unconstrained t so that $-xB_j + t$ is minimal, subject to the equation $x_1 + \cdots + x_m = 1$ (corresponding to the primal variable u), i.e., $x \in X$, and subject to the inequalities $-xB_k + t \geq 0$ for all $k \neq j$. The optimal value of this dual LP is equal to that of (5) and negative. In the corresponding optimal solution with x and t, this means $xB_j > t \geq xB_k$ for all $k \neq j$, that is, j is the unique best reply to x, as desired. □

If j is strictly dominated, then j is never a best reply, and both in the simultaneous game and in the leadership game player II will never choose j, so such a strategy can safely be omitted from the game. (Strictly dominated strategies of player I, on the other hand, may be relevant for the leadership game, as Figure 5 shows.)

Secondly, a strategy j is not inducible if $B_j = B_k$ for some $k \neq j$, since then j and k give the same payoff for any x. If the leader tries to induce the follower to play j, the follower may also play k, which may not be what the leader wants if $A_j \neq A_k$. We will treat this case last.

The remaining possibilities that j is not inducible are therefore that j is is either payoff equivalent to a mixture y of at least two other pure strategies, with $B_j = By$, or that j is weakly dominated by a pure or mixed strategy y, with $B_j \leq By$ but $B_j \neq By$. In the following, we will show that such best reply regions $X(j)$ have lower *dimension* than X itself.

We recall some notions from affine geometry for that purpose. An *affine combination* of points $z^1, \ldots, z^k$ in some Euclidean space is of the form $\sum_{i=1}^k \lambda_i z^i$ where $\lambda_1, \ldots, \lambda_k$ are reals with $\sum_{i=1}^k \lambda_i = 1$. Affine combinations are thus like convex combinations, except that some λ_i may be negative. The *affine hull* of some points is the set of their affine combinations. The points $z^1, \ldots, z^k$ are *affinely independent* if none of these points is an affine combination of the others, or equivalently, if $\sum_{i=1}^k \lambda_i z^i = \mathbf{0}$ and $\sum_{i=1}^k \lambda_i = 0$ imply $\lambda_1 = \cdots = \lambda_k = 0$. A convex set has *dimension d* if and only if it has $d+1$, but no more, affinely independent points.

Since X is the convex hull of the m affinely independent unit vectors $e_1, \ldots, e_m$ in (1), it has dimension $m-1$. Any convex subset Z of X is said to be *full-dimensional* if it also has dimension $m-1$. The following lemma implies that Z is full-dimensional if it is possible to move within Z from some x a small distance in the direction of all m unit vectors e_i, the extreme points of the unit simplex X.

Lemma 5 *Let $x \in X$ and $0 < \varepsilon \le 1$. Then the vectors $(1-\varepsilon)x + \varepsilon e_i$ for $1 \le i \le m$ are affinely independent, and x is a convex combination of these vectors.*

Proof. Consider real numbers $\lambda_1, \ldots, \lambda_m$ with $\sum_{i=1}^m \lambda_i = 0$. Then $\sum_{i=1}^m \lambda_i((1-\varepsilon)x + \varepsilon e_i) = \varepsilon \sum_{i=1}^m \lambda_i e_i$, so if this is the zero vector $\mathbf{0}$, then $\lambda_1 = \ldots = \lambda_m = 0$ since the unit vectors $e_1, \ldots, e_m$ are linearly independent. This proves the claimed affine independence. Secondly, letting $\lambda_i = x_i$ for $1 \le i \le m$ shows $\sum_{i=1}^m \lambda_i((1-\varepsilon)x + \varepsilon e_i) = (\sum_{i=1}^m \lambda_i)(1-\varepsilon)x + \sum_{i=1}^m \lambda_i \varepsilon e_i = (1-\varepsilon)x + \varepsilon x = x$, which is the representation of x as a convex combination of the described vectors. $\square$

The next lemma shows that full-dimensional sets are those that contain an open set, in the topology relative to X. Recall that for $\varepsilon > 0$, the ε-*neighborhood* of a point x in X is the set $\{ y \in X \mid \|y - x\| < \varepsilon \}$ where $\|y - x\|$ is the Euclidean distance of x and y.

Lemma 6 *Let Z be convex, $Z \subseteq X$. Then the following are equivalent:*
(a) *Z is full-dimensional;*
(b) *for some $x \in Z$ and $\varepsilon > 0$, the vectors $(1-\varepsilon)x + \varepsilon e_i$ are in Z for $1 \le i \le m$;*
(c) *Z contains a neighborhood.*

Proof. By Lemma 5, (b) implies (a). To show that (c) implies (b), let x be an interior point of Z. Then for sufficiently small positive ε, each vector $(1-\varepsilon)x + \varepsilon e_i$ has distance $\varepsilon\|x - e_i\|$ from x and hence belongs to a neighborhood of x, and thus to Z.

It remains to show that (a) implies (c). Let, by (a), $z^1, \ldots, z^m$ be affinely independent vectors in Z, and consider the simplex that is the convex hull of these vectors. Then it can be shown that the center of gravity $x = 1/m \cdot (z^1 + \cdots + z^m)$ of that simplex is in the interior of that simplex. Namely, every vector z^k has positive distance from the affine hull of the other $m-1$ vectors, so that x has $1/m$ that distance, call it δ_k. Any y in X that is not in the simplex is an affine combination $y = \sum_{i=1}^m \lambda_i z^i$ with, by affine independence, unique $\lambda_1, \ldots, \lambda_m$, where $\sum_{i=1}^m \lambda_i = 1$ and $\lambda_k < 0$ for at least one k. Then it easy to see that this implies $\|y - x\| > \delta_k$, which shows that the δ-neighborhood of x, with $\delta = \min_k \delta_k$, is included in the simplex and hence in Z. $\square$

The preceding topological condition in (c) shows easily that X is covered by the full-dimensional best reply regions, so these are of particular interest:

Lemma 7 *X is the union of the best-reply sets $X(j)$ that are full-dimensional.*

Proof. Let $F = \{ j \in N \mid X(j) \text{ is full-dimensional} \}$. Consider $Z = X - \bigcup_{j \in F} X(j)$. Suppose that, contrary to the claim, the open set Z is not empty. Then for any j in $N - F$, the set $Z_j = Z - X(j)$ is also open and not empty, since otherwise Z, and hence a neighborhood, would be a subset of the set $X(j)$ which is not full-dimensional, contradicting Lemma 6. Repeating this argument with Z_j instead of Z for the next strategy in $N - F$, and so on, shows that $Z - \bigcup_{j \in N - F} X(j)$, that is, $X - \bigcup_{j \in N} X(j)$, is not empty, contradicting (2). $\square$

The next lemma shows that the intersection of two best reply regions cannot be full-dimensional, except when the regions are identical.

Lemma 8 *If $X(j)\cap X(k)$ is full-dimensional, then $B_j = B_k$.*

Proof. Consider, by Lemma 6(b), x and $\varepsilon > 0$ so that $(1-\varepsilon)x+\varepsilon e_i$ are in $X(j)\cap X(k)$ for all $i \in M$. Then x, by Lemma 5 a convex combination of these vectors, also belongs to $X(j)\cap X(k)$. Since j and k are best replies to x and to $(1-\varepsilon)x+\varepsilon e_i$ for all $i \in M$, we have $xB_j = xB_k$ and $((1-\varepsilon)x+\varepsilon e_i)B_j = ((1-\varepsilon)x+\varepsilon e_i)B_k$, and therefore $e_iB_j = e_iB_k$, for $i \in M$. This means that the column vectors B_j and B_k agree in all components, as claimed. □

The following lemma shows the close connection between inducible strategies j and full-dimensional best reply sets $X(j)$. The only complication arises when there are other strategies k in N with $B_j = B_k$.

Lemma 9 *Let $j \in N$.*
(a) *If j is inducible, then $X(j)$ is full-dimensional.*
(b) *$X(j)$ is full-dimensional if and only if the only mixed strategies y in Y with $B_j \le By$ are those where $y_k > 0$ implies $B_j = B_k$.*
(c) *If $B_j \neq B_k$ for all $k \in N$, $k \neq j$, and $X(j)$ is full-dimensional, then j is inducible.*

Proof. If j is inducible, then the set $\{x \in X \mid xB_j > xB_k \text{ for all } k \neq j\}$ is not empty and open and a subset of $X(j)$, which is then full-dimensional by Lemma 6. (An easy direct argument with Lemma 5 is also possible.) This shows (a).

To show (b), suppose first that $X(j)$ is full-dimensional, and that there is some y in Y where $B_j \le By$ and, contrary to the claim, $y_k > 0$ and $B_j \neq B_k$ for some k. Then k is a best reply whenever j is, which shows $X(j) \subseteq X(k)$ and thus $X(j) = X(j)\cap X(k)$. This is a full-dimensional set, but then $B_j = B_k$ by Lemma 8, a contradiction. Conversely, suppose that $B_j \le By$ holds only when $y_k > 0$ implies $B_j = B_k$. Consider the game where all strategies k in N, $k \neq j$ with $B_j = B_k$ are omitted. In this game, there is no mixed strategy y with $B_j \le By$ and $y_j = 0$. Then j is inducible by Lemma 4, and hence $X(j)$ is full-dimensional by (a).

In the same manner, (b) and Lemma 4 imply (c). □

Theorem 3 is strengthened as follows. We still exclude the case that there are different pure strategies j and k with $B_j = B_k$.

Theorem 10 *Suppose that $B_j \neq B_k$ for all $j,k \in N$, $j \neq k$. Then the set of all leader payoffs is an interval $[L,H]$, where*

$$L = \max_{j\in N,\ j \text{ inducible}} \ \max_{x\in X(j)} xA_j\,, \qquad H = \max_{j\in N}\ \max_{x\in X(j)} xA_j\,. \tag{6}$$

Some Nash payoff to player I is less than or equal to L, and all Nash payoffs to player I are less than or equal to H.

Proof. By Lemma 9 and Lemma 7,

$$X = \bigcup_{j \in N,\ j \text{ inducible}} X(j). \tag{7}$$

Player II therefore has an inducible best reply to any x in X. Consider now the "inducible-only" game, both in the simultaneous and in the leadership version, where all pure strategies of player II that are not inducible are omitted. Given x in X, any pure best reply j to x in this "inducible-only" game is also a best reply to x in the original game, since this means $xB_k \le xB_j$ for all inducible k, and for any non-inducible best reply l to x in the original game there is also some inducible best reply k to x by (7), which implies $xB_l = xB_k \le xB_j$.

Hence, any Nash or leadership equilibrium of the "inducible-only" game remains such an equilibrium in the original game. By Theorem 3, the leader payoff in the "inducible-only" game is L as in (6). Therefore, L is also a possible leader payoff in the original game, and greater than or equal to some Nash payoff of the original game. The original game cannot have a lower leader payoff than L, even though player II may have additional best replies, by Lemma 2.

The highest possible leader payoff is clearly H as in (6), by letting the follower always choose a best reply in N that maximizes the payoff to the leader, and letting the leader commit to x^*, say, where the maximum H in (6) is taken, with $x^*A_{j^*} = H$, for a best reply j^* to x^* that is best for player I. As argued for Theorem 3, H is greater than or equal to any Nash payoff.

Trivially, $L \le H$. If $L < H$, it remains to show that any P with $L < P < H$ is a possible leader payoff. Then the above reply j^* to x^* is clearly not inducible, so that there is, by (7), an inducible best reply k to x^*, where $x^*A_k \le L$. We let the leader commit to x^*, and let the follower respond to x^* by randomizing between k and j^* with probabilities such the expected payoff to player I, a convex combination of x^*A_k and $x^*A_{j^*}$, is P. To any commitment other than x^*, the follower chooses only inducible replies, so that x^* is indeed the optimal commitment. Then this is a leadership equilibrium with leader payoff P, as claimed. □

Call two pure strategies j and k of player II *equivalent* if $B_k = B_j$, which defines an equivalence relation on N. If two strategies are equivalent, neither of them is inducible, and the previous theorem does not apply, because then (7) is not true. However, Lemma 7 is still true. Indeed, the leader payoffs in the general case (where equivalent strategies are allowed) can be described in terms of the strategies j such that $X(j)$ is full-dimensional, rather than the more special inducible strategies.

The strategies j such that $X(j)$ is full-dimensional are characterized in Lemma 9(b): $B_j \le By$ is only possible for mixed strategies y that assign positive probability to strategies that are equivalent to j, so that then obviously $B_j = By$. Indeed, omitting from the game all strategies that are equivalent to j (except j) would make j inducible, as used in the proof of Lemma 9(b). The possible leader payoffs, however, would not be captured correctly by such an omission, since inducibility is crucial for Lemma 2, and Player II may respond to j by some equivalent strategy instead, as argued for Figure 6.

If there are several equivalent strategies and these define a full-dimensional best reply region, then the leader payoff is determined by the "pessimistic" view that the follower's reply among these equivalent strategies is worst possible for the leader, since the follower is indifferent among all of them. The next theorem, which generalizes the preceding Theorems 3 and 10, states in (8) the lowest possible leader payoff L in terms of this "max-max-min" computation. The first maximization is over all j such that $X(j)$ is full-dimensional, where it in fact suffices to consider one such strategy among all its equivalent strategies (which defines a maximization over the equivalence classes, not stated in (8) to save notation). Then x is maximized over $X(j)$ (which is the same best reply region for all strategies that are equivalent to j). Finally, the payoff xA_k is minimized over all strategies k that are equivalent to j. Notably, L is still at least as large as *some* Nash payoff to player I.

The highest possible leader payoff H in (8) is determined as before in (6). Here, all strategies j in N are considered separately, under the "optimistic" view that player II chooses the reply j that is best for the leader, regardless of whether the strategy has equivalent strategies or whether it defines a best reply region of lower dimension.

Theorem 11 *Let $F = \{\, j \in N \mid X(j)$ is full-dimensional$\}$. The set of all leader payoffs is an interval $[L,H]$, where*

$$L = \max_{j\in F}\ \max_{x\in X(j)}\ \min_{k\in N,\ B_k=B_j} xA_k\,, \qquad H = \max_{j\in N}\ \max_{x\in X(j)} xA_j\,. \tag{8}$$

Some Nash payoff to player I is less than or equal to L, and all Nash payoffs to player I are less than or equal to H.

Proof. For j in N, we use the equivalence class notation $[\,j\,] = \{\,k \in N \mid B_k = B_j\,\}$. Obviously, $[\,j\,]$ is a subset of F whenever $j \in F$.

First, we show that L as defined in (8) is indeed a possible leader payoff. The corresponding leadership equilibrium is constructed as follows. Let $j \in F$, $x \in X(j)$ and $k \in [\,j\,]$ be such that the max-max-min in (8) is achieved with $xA_k = L$. The leader commits to x. The follower responds to x by k, and to any other commitment by choosing a pure best reply that minimizes the payoff to the leader. Then x is the optimal commitment. Any best reply k among the strategies in F is also a best reply when considering all strategies in N, due to Lemma 7, as argued for Theorem 10. Hence, we obtain a leadership equilibrium with leader payoff L.

Furthermore, L is the lowest leader payoff, by an argument similar to Lemma 2. Namely, $X(j)$ is full-dimensional and thus contains an interior point x' by Lemma 6(c). Any convex combination $x(\varepsilon) = (1-\varepsilon)x + \varepsilon x'$ for $\varepsilon \in (0,1]$ is then also in the interior of $X(j)$, so that the follower's best reply to $x(\varepsilon)$ is any $k \in [\,j\,]$. The minimum (over these $k \in [\,j\,]$) of the leader payoffs $x(\varepsilon)A_k$ is a continuous function of ε, and is arbitrarily close to L as ε tends to zero. Hence, the follower must play such that the leader gets at least L on the equilibrium path.

Similarly to Theorem 10, it easily seen that H is the highest leader payoff, and that the interval $[L,H]$ is the set of possible leader payoffs. Moreover, no Nash payoff to player I exceeds H.

In order to prove that the simultaneous game has a Nash equilibrium with payoff at most L for player I, we modify the given game in two steps. First, the "full-dimensional-only" game is constructed similar to the "inducible-only" game in the proof of Theorem 10, where N is replaced by F, so that player II can only use strategies j where $X(j)$ is full-dimensional. As argued above, any Nash or leadership equilibrium of this game remains an equilibrium in the original game.

In a second step, we consider the "factored" game where the pure strategies of player correspond to the equivalence classes $[j]$ for $j \in F$. For each such $[j]$, player II has a single payoff column B_j, which is by definition the same for any strategy in $[j]$, so we may call it $B_{[j]}$.

The payoff columns to player I in the "factored" game are obtained as certain convex combinations of the original columns A_k for $k \in [j]$. Namely, we consider the "constrained matrix game" where player I chooses x in $X(j)$ and player II mixes among the pure strategies k in $[j]$, with the zero-sum payoff columns A_k to player I. This game has a value (see Charnes (1953)), given by

$$L_j = \max_{x \in X(j)} \min_{k \in [j]} xA_k \,.$$

For completeness, and to clarify the players' strategy sets in this constrained game, we prove this by linear programming duality. Clearly, L_j is the maximal real number u such that $xA_k \geq u$ for all $k \in [j]$ and $x \in X(j)$, where the latter can be written as $x \in X$ and $x(B_j - B_l) \geq 0$ for all $l \in N - [j]$. As a minimization problem, this says: minimize $-u$ subject to

$$\begin{aligned} &xA_k - u &&\geq 0, && k \in [j] \\ &x(B_j - B_l) &&\geq 0, && l \in N - [j] \\ &x\mathbf{1} &&= 1, && \\ &x &&\geq 0 && \end{aligned} \tag{9}$$

where $\mathbf{1}$ is a column of m ones. The dual of this LP uses nonnegative variables z_k for $k \in [j]$ and w_l for $l \in N - [j]$ and an unconstrained t and says: maximize t subject to

$$\sum_{k \in [j]} A_k z_k + \sum_{l \in N - [j]} (B_j - B_l) w_l + \mathbf{1}t \leq \mathbf{0},$$

and, corresponding to the unconstrained variable u in (9),

$$-\sum_{k \in [j]} z_k = -1.$$

Consider an optimal solution to this LP. Then t is equal to the optimum of (9), that is, $t = -u = -L_j$. Furthermore, whenever player II uses the mixed strategy in Y with probabilities z_k for $k \in [j]$, and zero elsewhere, then for any $x \in X(j)$ the expected payoff to player I fulfills (note $w_l \geq 0$)

$$\sum_{k \in [j]} xA_k z_k \;\leq\; \sum_{k \in [j]} xA_k z_k + \sum_{l \in N - [j]} x(B_j - B_l) w_l \;\leq -x\mathbf{1}t = -t = L_j \,.$$

That is, player I indeed cannot get more than L_j for any $x \in X(j)$ when player II plays according the probabilities z_k for $k \in [j]$; call them $z_k^{[j]}$ since they depend on $[j]$. In the "factored" game, the payoff column $A_{[j]}$ to player I for strategy $[j]$ of player II is given by

$$A_{[j]} = \sum_{k \in [j]} A_k z_k^{[j]},$$

so that

$$L_j = \max_{x \in X(j)} \min_{k \in [j]} xA_k = \max_{x \in X(j)} xA_{[j]} . \tag{10}$$

By construction, all payoff columns $B_{[j]}$ for player II in the "factored" game are different. Moreover, all best-reply regions are full-dimensional since we only considered $[j]$ for $j \in F$, and hence all replies are inducible by Lemma 9(c). So Theorem 3 applies, and in some (indeed, any) Nash equilibrium of the "factored" game, the payoff to player I is at most equal to the leader payoff, which by (4), (10) and (8) is equal to

$$\max_{j \in F} \max_{x \in X(j)} xA_{[j]} = \max_{j \in F} L_j = L .$$

Finally, any Nash equilibrium (x, y'), say, of the "factored" game translates to a Nash equilibrium (x, y) of the "full-dimensional-only" game, and hence of the original game, as follows: Player I plays x as before, and player II chooses $k \in [j]$ for $j \in F$ with probability $y_k = y'_{[j]} z_k^{[j]}$. Then player I receives the same expected payoffs as before, so that x is a best reply to y, and since $B_k = B_{[j]}$ for $k \in [j]$, any $k \in F$ so that $y_k > 0$ (and hence $y'_{[j]} > 0$) is a best reply to x, as required. The resulting Nash payoff to player I is at most L, as claimed. □

Generic games do not need the development following Theorem 3. Games with identical payoff columns for player II as in Theorem 11 are obviously not generic. Even games that require the assumptions of Theorem 10 (with non-empty best reply regions that are not full-dimensional, see Lemmas 4 and 9) are not generic. Namely, a strategy that is only weakly but not strictly dominated entails a linear equation among the payoffs, which only holds for a set of measure zero in the space of all games (with independently chosen payoffs). The same holds for games where a strategy is payoff equivalent to a different mixed strategy.

4 Correlated equilibria

Games like the familiar "battle of sexes" illustrate the use of commitment as a coordination device. Coordination can also be achieved by the *correlated* equilibrium due to Aumann (1974) which generalizes the Nash equilibrium. In this section, we first show that the highest leader payoff H as defined in (8) is greater than or equal to the highest correlated equilibrium payoff to player I. Trivially, the lowest leader payoff L in (8) is at least as large as some correlated payoff, since it is at least as large as some Nash payoff.

We consider the *canoncial 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 $j, l \in N$,

$$\sum_{j \in N} z_{ij}\, a_{ij} \ge \sum_{j \in N} z_{ij}\, a_{kj}\,, \qquad \sum_{i \in M} z_{ij}\, b_{ij} \ge \sum_{i \in M} z_{ij}\, b_{il}\,. \tag{11}$$

When a strategy pair (i, j) is drawn with probability z_{ij} according to this distribution by some device or *mediator*, player I is told i and player I is told j. The first constraints in (11) state that player I, 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 II. Analogously, the second inequalities in (11) state that player II, when recommended to play j, has no incentive to switch to l.

Theorem 12 *The largest leader payoff H as defined in* (8) *is greater than or equal to any correlated equilibrium payoff to player I.*

Proof. Consider a correlated equilibrium z with probabilities z_{ij} fulfilling (11) above. Define the marginal probabilities on N by

$$y_j = \sum_{i \in M} z_{ij} \qquad \text{for } j \in N, \tag{12}$$

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 I given that player II 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 $H \ge c_s$, and that c_s is at least the payoff to player I in the correlated equilibrium z, which proves the theorem. Namely, define x in X by $x_i = z_{is}/y_s$ for $i \in M$, let player I commit to x in the leadership game, and let player II respond to x by playing s. According to the second inequalities in (11), s is indeed a best reply to x since the column s of z, which has positive probability y_s, is the distribution on M given by x except for the normalization factor $1/y_s$. For any commitment other than x, choose any best reply of the follower. This may not necessarily define a leadership equilibrium since player I may possibly improve his payoff by a different commitment, so a leadership equilibrium may require a change of the commitment (as, for example, in Figure 5), or may require changing the reply s to x. At any rate, however, the payoff c_s to player I when leader and follower play as described fulfills $c_s \le H$. Furthermore, the correlated equilibrium payoff to player I is an average of the conditional payoffs c_j for $j \in S$ and therefore not larger than their maximum c_s:

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

as claimed. $\square$

Moulin and Vial (1978) define a generalization of correlated equilibria which involves a commitment by both players. We show that it may give a payoff to player I which is higher than any leader payoff.

The *simple extension* of a correlated equilibrium (see Moulin and Vial (1978), p. 203) is also given by a distribution z on strategy pairs (i,j), 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 or herself to playing the recommended strategy, or not to be told the outcome and play some mixed strategy. In the latter case, the player knows only the *marginal probabilities* under z of the choices of the other player (for example, player I would know only y_j in (12)). 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. The respective inequalities are, for all $k \in M$ and $l \in N$,

$$\sum_{i,j} z_{ij}\, a_{ij} \ge \sum_j \Big(\sum_i z_{ij}\Big) a_{kj}\,, \qquad \sum_{i,j} z_{ij}\, b_{ij} \ge \sum_i \Big(\sum_j z_{ij}\Big) b_{il}\,. \tag{13}$$

These inequalities are obviously implied by the incentive constraints (11) of Aumann's correlated equilibrium.

I \ II	p	q	r
P	0, 0	−2, 1	1, −2
Q	1, −2	0, 0	−2, 1
R	−2, 1	1, −2	0, 0

Figure 7. Game with payoff 0 in a "simple extension" of a correlated equilibrium, which is higher than any leader payoff.

Figure 7 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 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 $1/3$, each player getting expected payoff $-1/3$.

For the game in Figure 7, a simple extension of the 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 (13), but is not a correlated equilibrium.

For the leadership game for Figure 7, it suffices to consider only one best reply region, say for the first strategy p of player II. 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 for player I. 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.

In Figure 7, the simple extension of a correlated equilibrium by Moulin and Vial (1978) gives a payoff which is higher than the leader payoff. This concept involves a commitment by *both* players to a correlated device. Moreover, it does not generalize a leadership game. The latter has generically a unique payoff to the leader, whereas the concept by Moulin and Vial has correlated and Nash equilibria of the simultaneous game as special cases.

5 More than two players

In a game with three or more players, it may no longer be advantageous for a player to commit to a mixed strategy if he has the opportunity to do so. If the game has $k+1$ players, any commitment by player I, say, to a mixed strategy induces a game with k players. The natural definition of the leadership game is then to look for a subgame perfect equilibrium where for any commitment of player I the remaining k players, called followers, play an equilibrium of the induced game.

Given any Nash equilibrium of the simultaneous $(k+1)$-player game, a commitment by player I to his equilibrium strategy, with the corresponding replies in that equilibrium played by the other players, should give player I at least the payoff he gets in the simultaneous game. From that perspective, the situation does not seem to differ from the two-player case. Indeed, any Nash payoff to player I is a possible leader payoff in a subgame perfect equilibrium of the leadership game, by the preceding argument. However, there may be additional leader payoffs, all of which are strictly *worse* for the leader, compared to the simultaneous game. That is, the set of payoffs to player I may “move downwards” when introducing commitment, in direct contrast to the two-player case where it “moves upwards” according to Theorem 11.

This situation arises in the *team games* investigated by von Stengel and Koller (1997). These are games of $k+1$ players where player I plays against the remaining k players which form a team because they receive identical payoffs, which are the negative of the payoffs to player I. Here, commitment generally hurts player I since it allows the opposing team to *coordinate* their actions, which is not the case in the simultaneous game. In particular, the team may always reply to a commitment by a profile of *k pure* strategies, which is simply the profile that maximizes their joint payoff. In that case, the leader will commit to a mixed strategy which is his maxmin strategy in the zero-sum game where the other k strategies are chosen by the team acting as a single opponent, where a pure reply suffices in the leadership game. The simultaneous game, in contrast, may require mixed strategies by the team players, who cannot correlate their random choices and are therefore in general worse off than if they acted as a single player. They may choose to play a "team-maxmin" profile of k mixed strategies that maximizes the worst possible payoff to the team. This profile can be completed to a Nash equilibrium of the simultaneous game, as shown by von Stengel and Koller (1997).

I: L

II \ III	p	q
P	$-1,1,1$	$0,0,0$
Q	$0,0,0$	$-4,4,4$

I: R

II \ III	p	q
P	$-4,4,4$	$0,0,0$
Q	$0,0,0$	$-1,1,1$

Figure 8. Game with player I against the team of player II and III which has leader payoffs that are worse than any Nash payoff.

The three-player game in Figure 8 is an example. Player I chooses the left (L) or right (R) panel, and players II and III form the team and have two strategies each. The Nash equilibria in this game are as follows. Suppose that player I chooses R with probability x. Then players II and III each receive expected payoff $1+3x$ for the strategy pair (P,p), and $4-3x$ for (Q,q), and zero otherwise. Any Nash equilibrium of the three-player game induces a Nash equilibrium in this 2×2 coordination game, which is either (P,p), or (Q,q), or the mixture where P and p are each played with probability $(4-3x)/5$, with resulting team payoff $(4-3x)(1+3x)/5$. Since player I wants to minimize that payoff, his best reply to (P,p) is L and to (Q,q) is R, with team payoff 1 in both cases. When player I chooses L (where $x=0$) and players II and III mix by playing P and p each with probability $4/5$, against which L is a best reply, the team does even worse, getting $4/5$; the same applies for the mixed reply against R. For $0<x<1$, the mixed reply of II and III gives an equilibrium only if the resulting payoff cannot be improved by player I by choosing L or R. It is easy to see that if P and p have larger (smaller) probability than Q and q, then L (R) is a best reply. Hence, the only remaining equilibrium is where each player chooses each strategy with probability $1/2$. This is the mentioned "team-maxmin" equilibrium with payoff $5/4$ for the team, and payoff $-5/4$ for player I.

In the leadership game, any commitment to playing R with probability x induces the above game for the team which has three equilibria. Players II and III may coordinate to play their favorable pure equilibrium, namely (Q,q) with team payoff $4-3x$ for $x\in[0,1/2]$ and (P,p) with team payoff $1+3x$ for $x\in(1/2,1]$, say. The optimal commitment is then $x=1/2$. This defines a subgame perfect equilibrium with leader payoff $-5/2$, which is much worse for player I than in any Nash equilibrium. On the other hand, the leader gets his best payoff when the team players play their mixed equilibrium, and the leader commits to either L or R. This is already a Nash payoff, and not improved for the leader by commitment. So commitment worsens the set of payoffs for player I, as claimed.

The game in Figure 8 is of course non-generic. However, the same arguments apply for any other generic game with payoffs nearby.

Can Theorem 12 on correlated equilibria be extended to games with $k+1$ players, for $k \geq 2$? In that case, a natural extension of the leadership game would be to consider correlated equilibria of the game with k players that results from each commitment to a mixed strategy. The resulting "subgame perfect correlated equilibria" are then compared with the correlated equilibria of the original simultaneous game.

In this context, Figure 8 does, at first sight, not seem to give a counterexample since the worst leader payoff $-5/2$ is also a possible correlated equilibrium payoff to player I. (Players II and III correlate by playing (P,p) and (Q,q) each with probability $1/2$, and player I mixes independently between L and R.) However, if player II (or III) is made a leader, she can no longer get payoff $5/2$, since by her commitment to a mixed strategy, player II loses the ability to correlate with player III. Any commitment by player II induces a two-person zero-sum game between players I and III, and the resulting value is maximal for the leader if the players play as in the team-maxmin equilibrium, choosing each pure strategy with probability $1/2$. So Figure 8 shows indeed that, compared to correlated equilibria in the simultaneous game, a player may *strictly lose* by becoming a leader who unilaterally commits to a mixed strategy, if the game has more than two players.

6 Leadership equilibria and Stackelberg problems

In this section, we connect our concept with the closely related notion of "Stackelberg solutions" in the literature on *dynamic games* and optimization theory, as in Başar and Olsder (1982). There the payoffs to the players are usually declared as costs which are minimized, but we keep our view of payoff maximization.

Consider a finite game with $k+1$ players, where the mixed strategy set of the leader is X and the mixed strategy sets of the k followers are $Y^1,\ldots,Y^k$, and let $Y = Y^1 \times \cdots \times Y^k$. Every x in X, representing a commitment by the leader, induces a k-player game where we assume that the k followers play an element of the set $N(x)$ of Nash equilibria of that game, which is a subset of Y. If there is only one follower ($k = 1$), then $N(x)$ is simply the set of best replies to x. For x in X and y in Y, the payoff to the leader is denoted by $a(x,y)$.

Başar and Olsder (1982, p. 136, p. 141) define the *Stackelberg* payoff to the leader as

$$S = \sup_{x \in X} \min_{y \in N(x)} a(x,y)\,. \tag{14}$$

This equation describes the "pessimistic" view that among all the possible equilibria (or best replies if there is only one follower) in $N(x)$, the followers choose that which is worst for the leader. The set of Nash equilibria $N(x)$ is compact, so that it is indeed possible to take the minimum in (14) rather than the infimum. However, $\min_{y\in N(x)} a(x,y)$ is a discontinuous function of x, as for example the inspection game in Figure 2 and its analysis demonstrates. (Başar and Olsder (1982, p. 137) give a similar example.) In the game in Figure 2, the follower is indifferent when the leader chooses B with probability $1/10$. According to (14), the follower should then choose the reply r that is bad for the leader, so the supremum in (14) is not obtained as a maximum.

The discontinuity is usually seen as a problem in the optimization theory literature and addressed by various "regularization" approaches that justify taking a solution that approximates S (see Mallozzi and Morgan (2002) or Morgan and Patrone (2005) and references therein).

The theorem of this section states that the Stackelberg payoff S in (14) in the case of one follower is identical to the lowest leader payoff L in (8), and perhaps more concisely expressed by (14). The main difference is that in general, (14) does *not* describe the follower's behavior in the leadership equilibrium, where the follower usually chooses a reply that is favorable for the leader. Furthermore, S is the lowest possible equilibrium payoff to the leader for any number of followers.

Theorem 13 *Consider the mixed extension of a finite game with $k+1$ players. Then the corresponding leadership game with one leader and k followers has a subgame perfect equilibrium $(\bar{x},\bar{y})$ with $\bar{y}\in N(\bar{x})$ so that $a(\bar{x},\bar{y}) = S$ in* (14). *Any other leader payoff is at least S.*

Proof. The difficulty is that typically $a(\bar{x},\bar{y}) \neq \min_{y\in N(\bar{x})} a(\bar{x},y)$ since $N(\bar{x})$ is usually not a singleton.

For any x, the set of equilibria $N(x)$ is the set of fixed points of a suitable continuous mapping $T_x\colon Y \to Y$, for example the mapping T defined by Nash (1951, p. 288). Furthermore, this mapping is continuous in x as well. Let the continuous function F on $X\times Y$ be defined by $F(x,y) = T_x(y) - y$ (with values in a Euclidean space extending Y). Then $N(x) = \{\, y \mid F(x,y) = 0 \,\}$ and the correspondence

$$\bigcup_{x\in X} \Big(\{x\} \times N(x) \Big) = F^{-1}(0)$$

is also closed and therefore a compact set, as a subset of the compact set $X\times Y$.

Consider a sequence (x_n,y_n) for $n = 1,2,\ldots$ so that for all n,

$$a(x_n,y_n) = \min_{y\in N(x_n)} a(x_n,y)$$

and so that $a(x_n,y_n)$ converges to S in (14). This sequence belongs to $F^{-1}(0)$ and has a convergent subsequence with limit $(\bar{x},\bar{y})$ in $F^{-1}(0)$. Since a is continuous, $a(\bar{x},\bar{y}) = S$ as required.

This shows that the claimed leadership equilibrium exists. If (x^*,y^*) was another leadership equilibrium so that $a(x^*,y^*) \le S - \varepsilon$ for some $\varepsilon > 0$, then for some n in the above sequence we would have $a(x_n,y_n) > a(x^*,y^*)$ and the leader could deviate from x^* to x_n and thereby get a higher payoff, contradicting the equilibrium property. Hence, S is the smallest possible leadership payoff. □

7 Follower payoff

In the leadership game, the leader's payoff is never worse than his Nash payoff in the simultaneous game. The follower may do worse or better, and even profit more from leadership than the leader himself. We show this for the mixed extension of a 3×3 game.

I \ II	l	c	r
T	$-d$ $-d$	$2-d$ 1	$1-d$ 0
M	1 $2-d$	4 4	5 0
B	0 $1-d$	0 5	2 2

Figure 9. Symmetric game with leader payoff $5/2$ and follower payoff $3-d/2$, compared to the unique Nash payoff pair $(2,2)$.

Figure 9 shows a symmetric game, where the payoffs for each player's first strategy depend on a real parameter d. The game has a unique Nash equilibrium (B,r), obtained by iteratively eliminating first the strictly dominated strategies T and l and then M and c. The Nash payoff pair is $(2,2)$. In the leadership game, the follower's strategy l is strictly dominated, and c is a best reply whenever $(2-d)x_1 + 4x_2 \geq (1-d)x_1 + 5x_2 + 2x_3$, that is, $x_1 \geq x_2 + 2x_3$, for the probabilities (x_1,x_2,x_3) for (T,M,B). Thus, the best reply region for c is the convex hull of the extreme points $(1,0,0)$, $(2/3,0,1/3)$, and $(1/2,1/2,0)$ with corresponding payoffs 1, $7/3$, and $5/2$ to player I. Of these, $5/2$ is the maximum, and larger than any payoff when r is a best reply. The leader payoff is therefore $5/2$, corresponding to the commitment to $(1/2,1/2,0)$ with best reply c. The payoff to the follower is $3-d/2$.

For $d=0$, the follower receives 3 in the leadership game, and therefore profits more from the commitment power of the leader than the leader himself. For $d=1$, both leader and follower receive the same payoff $5/2$ in the leadership game. For $d=2$, the follower gets 2 in the leadership game, which is the same as her Nash payoff. Finally, for $d > 2$, the follower gets less in the leadership game than what she would get in the simultaneous game.

Interestingly, many types of duopoly games, like quantity or price competition, have the property that when the follower's payoff is not worse than her Nash payoff, it is already better than the payoff she would get when she was a leader. This is not the case in the game in Figure 9, but relies on strategy sets that are intervals, and certain monotonicity conditions of the players' payoffs, as in Hamilton and Slutsky (1990). Since it would lead too far afield, this result is the topic of a separate paper (von Stengel (2003)).

8 Open questions

As a possible theme for further work, one may consider more general games than mixed extensions of finite games that are known to have Nash equilibria. A general class is given by games fulfilling the concept of "better reply security" by Reny (1999). An example, due to Dufwenberg and Stegeman (2002), is that player I chooses x, player II chooses y, each from $[0,1]$, and they get the payoff pair (x,y), except when $x = 1$ and $y < 1$, where the payoffs are $(0,y)$, and when $x < 1$ and $y = 1$, where the payoffs are $(x,0)$. The unique Nash equilibrium is $(1,1)$ with payoffs $(1,1)$. Here, the strategy pair $(0.99,0.99)$ is much safer and therefore more reasonable, even though it is not a Nash equilibrium, but it is strictly dominated by $(0.999,0.999)$, and so on. The resulting leadership game has also payoffs $(1,1)$, but a subgame perfect equilibrium does not exist because best replies do not exist *off the equilibrium path*: There is no best reply y against a commitment of $x = 0.99$, for example, since the resulting payoff to the follower as a function of y is discontinuous and has no maximum. This does not look like a reasonable objection to analyzing the leadership game. However, we do not analyze this topic further.

Another possible research may relate leadership equilibria to "nonoptimizing" agents, whose behavior represents a form of commitment, in evolutionary games, as studied by Banerjee and Weibull (1995).

References

Amir, R., and I. Grilo (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. and M. Canty (1996), Compliance Quantified. Cambridge Univ. Press, Cambridge.

Avenhaus, R., A. Okada, and S. Zamir (1991), Inspector leadership with incomplete information. In: Game Equilibrium Models IV, ed. R. Selten, Springer, Berlin, 319–361.

Avenhaus, R., B. von Stengel, and S. Zamir (2002), Inspection games. In: Handbook of Game Theory with Economic Applications, Vol. 3, eds. R. J. Aumann and S. Hart, Elsevier, Amsterdam, 1947–1987.

Bagwell, K. (1995), Commitment and observability in games. Games and Economic Behavior 8, 271–280.

Başar, T., and G. J. Olsder (1982), Dynamic Noncooperative Game Theory. Academic Press, London.

Banerjee, A., and J. W. Weibull (1985), Evolutionary selection and rational behavior. In: Learning and Rationality in Economics, eds. A. Kirman and M. Salmon, Blackwell, Oxford, 343–363.

Binmore, K. (1992), Fun and Games: A Text on Game Theory. D. C. Heath, Lexington, Mass.

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.

Dufwenberg, M., and M. Stegemann (2002), Existence and uniqueness of maximal reductions under iterated strict dominance. Econometrica 70, 2007–2023.

Friedman, J. W. (1977), Oligopoly and the Theory of Games. North-Holland, Amsterdam.

Hamilton, J. and S. Slutsky (1990), Endogenous timing in duopoly games: Stackelberg or Cournot equilibria. Games and Economic Behavior 2, 29–46.

Landsberger, M., and D. Monderer (1994), To commit or not to commit? Haifa University and Technion, Haifa, mimeo (unpublished).

Mallozzi, L., and J. Morgan (2002), Mixed strategies for hierarchical zero-sum games. In: Advances in Dynamic Games and Applications, eds. E. Altman and O. Pourtallier, Birkhäuser, Boston, 65–77.

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., and F. Patrone (2005), Stackelberg problems: Subgame perfect equilibria via Tikhonov regularization. In: Annals of the International Society of Dynamic Games, Vol. 7, to appear.

Moulin, H. and J.-P. Vial (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. (1951), Non-cooperative games. Annals of Mathematics 54, 286–295.

Reny, P. J. (1999), On the existence of pure and mixed strategy Nash equilibria in discontinuous games. Econometrica 67, 1029–1056.

Reny, P. J., and A. J. Robson (2002), Reinterpreting mixed strategy equilibria: A unification of the classical and Bayesian views. Mimeo, University of Chicago and University of Western Ontario.

Rosenthal, R. W. (1991), A note on robustness of equilibria with respect to commitment opportunities. Games and Economic Behavior 3, 237–243.

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.

Simaan, M. and J. B. Cruz, Jr. (1973), On the Stackelberg strategy in nonzero-sum games. Journal of Optimization Theory and Applications 11, 533–555.

van Damme, E. and S. Hurkens (1997), Games with imperfectly observable commitment. Games and Economic Behavior 21, 282–308.

von Neumann, J. and O. Morgenstern (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. (2003), Follower Payoffs in Symmetric Duopoly Games. Research Report LSE-CDAM-2003-13, London School of Economics.

von Stengel, B. and D. Koller (1997), Team-maxmin equilibria. Games and Economic Behavior 21, 309–321.

Wölling, A. (2002), Das Führerschaftsprinzip bei Inspektionsspielen. (The Leadership Principle for Inspection Games.) Kovač, Hamburg.