

Team-Maxmin Equilibria

Bernhard von Stengel*

Institute for Theoretical Computer Science, ETH Zürich, 8092 Zürich, Switzerland

and

Daphne Koller†

*Computer Science Department, Gates Building 1A, Stanford University,
Stanford, California 94305-9010*

Games and Economic Behavior **21** (1997), 309–321

(page numbers in this manuscript match only approximately)

In a noncooperative game, a team is a set of players that have identical payoffs. We investigate zero-sum games where a team of several players plays against a single adversary. The team is not regarded as a single player because the team members might not be able to coordinate their actions. In such a game, a certain equilibrium can be selected naturally: the *team-maxmin equilibrium*. It assures the team players the best payoff they can hope for, given their inability to coordinate. A team-maxmin equilibrium exists, and in a generic game it is unique. *Journal of Economic Literature* Classification Number: C72.

Keywords. Equilibrium selection, maxmin strategy, noncooperative multi-person game, zero-sum game.

*E-mail: stengel@inf.ethz.ch

†E-mail: koller@cs.stanford.edu

1. Introduction

Many interactions involve *teams* of participants that have coinciding interests but that have to act individually. For example, the individuals in a company work towards a common goal, but often have to make their decisions independently. Coordinating their actions may be impossible or too expensive. Similarly, the performance of a distributed computer system in a given situation depends on individual actions taken by the processors. In many situations, it may be useful to have the processors randomly choose their actions. However, randomization is only efficient if done locally by each processor. In the game of Bridge, the two players on each team are forbidden by the rules to coordinate their actions by secret communication.

These situations can be modeled as noncooperative games. In such a game, we define a *team* as a set of n players with identical payoffs.¹ We are concerned with teams facing a single *adversary*, who is another player in the game. For each outcome of the game, the payoff to the team, received by every team member, is paid n times by the adversary. That is, we extend the classical two-person zero-sum game by considering a multiperson zero-sum game with n team members and the adversary as separate players. (It is not a general multiperson zero-sum game since that requires more than one payoff figure per outcome.) In a corporate setting, the adversary might be the outside world, which in a defensive view acts like a single player. In a distributed system, the adversary may represent a worst-case assumption on the inputs provided to the system. In Bridge, a similar situation applies after the bidding round, where a single player—from the team that won the bid—plays against the opposing team, which consists of two players.

More formally, we have a noncooperative game with $n + 1$ players—the n team members and their adversary. We assume each player has finite number of strategies which (by conditional actions, for example) incorporate all communication allowed by the structure of the game. What is a reasonable way for playing such a game? Since the players on the team have coinciding interests, it is beneficial and reasonable for them to think jointly about their strategies before the game. For example, in the corporate setting, we may assume that the employees are all working from the same corporate operations manual. Hence, we can assume that the players jointly choose a *strategy profile*.

As usual, higher payoffs can typically be achieved if the team members use randomized strategies. However, we do not allow correlated randomization since that involves communication beyond the structure of the game. That is, each team member may only use a separate mixed strategy. This view is particularly appro-

¹There is an economic theory of teams (Marschak and Radner 1976, Kim and Roush 1987) where the definition of a team agrees with ours. However, the results (e.g., Ho and Sun 1974) are of a different character since they usually involve distributions on “states of the world” as information structures.

priate if the game is played repeatedly and the team can get together only once, to decide on the profile of mixed strategies it wants to adopt. One may even think that the team members only have the game description and figure their strategies out by themselves (we will see that this works except in certain degenerate cases).

Since this is an adversarial setting, it is reasonable to assume that the team plays defensively. That is, the team members play such that their worst-case payoff (with respect to the adversary's decisions) is maximized. Thus, we assume that the team adopts a *team-maxmin strategy profile*, that is, a mixed strategy for each team member so that the minimal expected team payoff over all responses of the adversary is maximal. This is particularly reasonable under a worst-case assumption, since the adversary may even get to know the team's profile.

The same defensive approach may be taken by the adversary of the team. In that case, she uses a minmax strategy by randomizing such that the largest expected payoff (which is her cost) for any strategy profile of the team is minimal. If the team has only one member, the game is a two-person zero-sum game where the maxmin payoff is equal to the minmax cost. In that case, the maxmin strategy of the single team player and the minmax strategy of the adversary define an equilibrium. If the team has two or more members, this is generally not the case. Then, the team-maxmin strategy profile usually leads to a payoff that is lower than the minmax cost.

What is the cause of this gap? When the adversary is considering her worst-case scenario, she is assuming that the team chooses its strategy in response to her choice of strategy. That is, the team members "get together" after discovering her strategy, and choose the best possible response to it. In this case, it suffices to consider deterministic responses of the team members, which do not involve correlated randomization. Therefore, the adversary's worst-case assumption does not allow her to take advantage of the fact that the team members cannot coordinate their actions, i.e., the fact that each team member can only respond individually to the adversary's strategy.

In this paper, we show that the adversary can use the team's inability to coordinate in the following way: If the team members use a team-maxmin strategy profile, then the adversary has a mixed strategy so that no team member can increase payoffs by changing his strategy *unilaterally*. In other words, any team-maxmin strategy profile is part of an equilibrium of the game, which we call a *team-maxmin equilibrium*.

In a two-person zero-sum game, a maxmin and a minmax strategy define an equilibrium, and a unique value of the game. Nash (1951) showed that a general noncooperative game has an equilibrium, where the payoff is usually not unique. We provide an intermediate result for our special multi-person zero-sum games. Not all equilibria are team-maxmin equilibria, but a team-maxmin equilibrium has

a unique payoff which is the optimal payoff that separate team players can enforce. Hence, our theorem also provides a way for *equilibrium selection* for these games: a team-maxmin equilibrium can be considered as the natural solution of the game.

In the following section, we give an example showing that for teams with more than one player, the team-maxmin payoff may be lower than the minmax payoff, and which illustrates our theorem and related questions. Then, we state and prove the theorem formally. We summarize it in the concluding section, and discuss the connection with imperfect recall in extensive games.

2. Team-Maxmin Strategy Profiles

We consider a game in normal form with $n + 1$ players. One player is the adversary of a team formed by the remaining n players. Every player has finitely many pure strategies. An n -tuple of pure or mixed strategies, one for each team player, is called a *strategy profile* of the team. Depending on the adversary's strategy and on the team's strategy profile, the adversary has a certain cost that is n times the payoff to each team player.

As an illustrating example, we consider a team with two members, each of whom has two pure strategies. A mixed strategy for the first team player is specified by the probability p , say, that he chooses one of his pure strategies, so that $1 - p$ is the probability for the other pure strategy. The pure strategies themselves correspond to the cases $p = 0$ and $p = 1$. Similarly, a mixed strategy for the second team player is defined by a probability q , where again $q = 0$ and $q = 1$ denote deterministic choices. A strategy profile for the team is given by p, q . The adversary has two strategies, T and B , with the following payoffs to the team resulting from combinations of pure strategies.

p, q	0, 0	0, 1	1, 0	1, 1	
T	3	5	2	1	(1)
B	3	1	5	0	

Suppose first that the team can act as a single player. Then this table represents a two-person zero-sum game, where the column player, that is, the team, has a maxmin strategy where 0, 1 is played with probability $3/7$ and 1, 0 with probability $4/7$. The resulting expected payoff is $23/7$, which is the value of the two-person game. Each column yields at most this expected payoff if the row player plays her minmax strategy where she chooses T with probability $4/7$ and B with probability $3/7$.

However, as we explained in the introduction, the members of the team are not allowed to communicate during the course of the game, except as prescribed by its structure (in this game in normal form, all players move simultaneously). Therefore, the team players must randomize independently. Let p and q be mixed strategies

for team players 1 and 2, respectively. For the strategy profile p, q of the team, let $C_T(p, q)$ and $C_B(p, q)$ denote the expected payoffs to the team, given by

$$\begin{aligned} C_T(p, q) &= 3 + 2q - p - 3pq, \\ C_B(p, q) &= 3 - 2q + 2p - 3pq. \end{aligned}$$

The team's goal is to maximize its worst-case payoff, that is, to maximize the minimum of these two expressions. This minimum is largest if the two expressions are equal, that is, $C_T(p, q) = C_B(p, q)$, or equivalently $p = \frac{4}{3}q$. The resulting payoff $3 + \frac{2}{3}q - 4q^2$ has its maximum at $\bar{q} = \frac{1}{12}$. The *team-maxmin strategy profile* is thus \bar{p}, \bar{q} with $\bar{p} = \frac{1}{9}$. The corresponding team-maxmin payoff (for all possible strategies of the adversary) is $\frac{109}{36}$. This payoff is smaller than the value $\frac{23}{7}$ of the two-person game with the team as a single player.

Can the adversary achieve to pay only this lower cost of $\frac{109}{36}$? In order to do so, the adversary needs a strategy that discourages the players from deviating from this team-maxmin strategy profile \bar{p}, \bar{q} . Of course, if the team members can coordinate their actions, the adversary cannot guarantee a lower cost than the minmax cost. However, the adversary can take advantage of the players inability to communicate during the game. Essentially, she needs to find a strategy that, together with \bar{p}, \bar{q} , defines an equilibrium of the original three-person game with the team members as separate players. If such an equilibrium exists, then neither of the team players will have an incentive to deviate unilaterally.

Suppose that the team players use their team-maxmin strategy profile \bar{p}, \bar{q} . Then, the strategies T and B yield the same expected payoff, so the adversary can select both with positive probability. If the mixed strategy of the second team player is fixed at \bar{q} , then the expected payoff to the first team player is a linear function of his strategy p , namely $\frac{19}{6} - \frac{5}{4}p$ for T and $\frac{17}{6} + \frac{7}{4}p$ for B . Therefore, in order to obtain an equilibrium where the best response of the first team player is $p = \bar{p}$ (and not $p = 0$ or $p = 1$), the probabilities for T and B have to be $\frac{7}{12}$ and $\frac{5}{12}$, respectively.

Does this randomized strategy of the adversary define an equilibrium? That is, has the second team player no incentive to deviate either? Somewhat surprisingly, the answer is yes. If \bar{p} is fixed, the expected payoff to the second team player when he plays q is $\frac{26}{9} + \frac{5}{3}q$ for T and $\frac{29}{9} - \frac{7}{3}q$ for B . Indeed, this determines the same probabilities $\frac{7}{12}$ and $\frac{5}{12}$ for T and B as before so that the second team player is indifferent between his two pure strategies and $q = \bar{q}$ is a best response. The main theorem of this paper shows that this is not an accident, but rather an instance of a general phenomenon.

Note that the equilibrium strategy of the adversary is different from her minmax strategy. That minmax strategy guarantees her the lowest cost for all columns in (1). These columns are the pure strategies of the team considered as a single player, but

they also represent special profiles of individual strategies for the team players. That is, there is a gap between the minmax cost and the team-maxmin payoff for all strategy profiles of the team. The equilibrium property is weaker than the minmax property. When using her equilibrium strategy, the adversary exploits the fact that the team players may deviate only unilaterally, and achieves a lower cost.

3. Team-Maxmin Equilibria

We consider now the general case with n team players. We assume that the game is represented in normal form, where each player has a finite set of strategies. (As usual, games in extensive form (Kuhn 1953) can be converted to the normal form.) The mixed strategy of team player k is denoted by the vector x^k of probabilities, with probability x_s^k for his pure strategy s , that is,

$$\sum_s x_s^k = 1 \quad \text{for } k = 1, \dots, n \quad (2)$$

and

$$x_s^k \geq 0 \quad \text{for any pure strategy } s \text{ of any team player } k. \quad (3)$$

For simplicity, we have not named the number of pure strategies of a team member. Any n -tuple $x = (x^1, \dots, x^n)$ of vectors fulfilling (2) and (3) represents a strategy profile of the team. Let m be the number of strategies of the adversary. For her i th strategy, the expected team payoff for the strategy profile x is denoted by $C_i(x)$. A *team-maxmin strategy profile* \bar{x} with team payoff \bar{u} is an optimal solution x, u to the problem: maximize u subject to (2), (3), and

$$u \leq C_i(x) \quad \text{for } i = 1, \dots, m. \quad (4)$$

A team-maxmin strategy profile exists since the minimum of the expected payoffs $C_i(x)$, which is maximized, is a continuous function of x on the compact domain defined by the constraints (2) and (3).

In the above example, the adversary has a randomized strategy that, combined with the given team-maxmin strategy profile, defines an equilibrium. This holds generally, as asserted in the following central definition and theorem of this paper.

Definition. *Consider a noncooperative zero-sum game between a team of n players and an adversary. Then a team-maxmin equilibrium \bar{x}, \bar{y} is an equilibrium (in the $n + 1$ person game) where \bar{y} is a mixed strategy of the adversary and \bar{x} is a team-maxmin strategy profile.*

Theorem. *Any team-maxmin strategy profile is part of a team-maxmin equilibrium.*

Proof. Consider a team-maxmin strategy profile $\bar{x} = (\bar{x}^1, \dots, \bar{x}^n)$ with payoff \bar{u} as in (4). We are looking for a randomized adversary strategy \bar{y} , a vector of probabilities \bar{y}_i for each $i = 1, \dots, m$, so that \bar{x}, \bar{y} is an equilibrium.

For any fixed team member k and adversary's strategy i , we consider the i th expected payoff where only the strategy x^k of that team player may vary. The resulting function of x^k is denoted by

$$D_i^k(x^k, \bar{x}) = C_i(\bar{x}^1, \dots, \bar{x}^{k-1}, x^k, \bar{x}^{k+1}, \dots, \bar{x}^n) \quad \text{for } k = 1, \dots, n.$$

Let e_s denote the probability vector representing the pure strategy s of team player k . Clearly,

$$D_i^k(x^k, \bar{x}) = \sum_s D_i^k(e_s, \bar{x}) x_s^k, \quad (5)$$

that is, the i th expected payoff is a linear function of x^k if the strategies of the other team members are fixed as in \bar{x} .

Consider the following linear program (LP): For all strategy profiles $x = (x^1, \dots, x^n)$ and reals u , maximize nu subject to (2), (3), and

$$nu - \sum_{k=1}^n D_i^k(x^k, \bar{x}) \leq 0 \quad \text{for } i = 1, \dots, m. \quad (6)$$

This LP is feasible with $x, u = \bar{x}, \bar{u}$. We claim that this is already the optimal solution. Suppose otherwise that some x, u fulfill (2), (3), and (6), and $u > \bar{u}$. We show that then $(1 - \varepsilon)\bar{x} + \varepsilon x$ is a strategy profile yielding a larger payoff for all adversary strategies i if ε is positive and sufficiently small, which contradicts the assumption that \bar{x} is a team-maxmin strategy profile. We repeatedly use the multilinearity of $C_i(x^1, \dots, x^n)$ in each argument x^k , which is implied by (5):

$$\begin{aligned} & C_i((1 - \varepsilon)\bar{x} + \varepsilon x) \\ &= C_i((1 - \varepsilon)\bar{x}^1 + \varepsilon x^1, \dots, (1 - \varepsilon)\bar{x}^n + \varepsilon x^n) \\ &= (1 - \varepsilon)C_i(\bar{x}^1, (1 - \varepsilon)\bar{x}^2 + \varepsilon x^2, \dots, (1 - \varepsilon)\bar{x}^n + \varepsilon x^n) + \\ & \quad \varepsilon C_i(x^1, (1 - \varepsilon)\bar{x}^2 + \varepsilon x^2, \dots, (1 - \varepsilon)\bar{x}^n + \varepsilon x^n) \\ &= \dots \\ &= (1 - \varepsilon)^n C_i(\bar{x}) + \varepsilon(1 - \varepsilon)^{n-1} \sum_{k=1}^n D_i^k(x^k, \bar{x}) + \varepsilon^2 E(\varepsilon, x, \bar{x}) \\ &= C_i(\bar{x}) - n\varepsilon C_i(\bar{x}) + \varepsilon \sum_{k=1}^n D_i^k(x^k, \bar{x}) + \varepsilon^2 F(\varepsilon, x, \bar{x}), \end{aligned} \quad (7)$$

where $E(\varepsilon, x, \bar{x})$ and $F(\varepsilon, x, \bar{x})$ are bounded expressions for $\varepsilon \in [0, 1]$. We show that the expression in (7) is greater than \bar{u} for all $i = 1, \dots, m$ if ε is sufficiently small.

This is clearly the case if $C_i(\bar{x}) > \bar{u}$. For strategies i with $C_i(\bar{x}) = \bar{u}$, (6) implies that the expression in (7) is at least as large as $\bar{u} + \varepsilon(-n\bar{u} + nu + \varepsilon F(\varepsilon, x, \bar{x}))$, which is greater than \bar{u} for small positive ε because $u > \bar{u}$. This is the indicated contradiction. Thus, \bar{x}, \bar{u} is indeed the optimal solution to the above LP.

Consider the dual of this LP. Corresponding to the n equalities in (2), there are n dual variables z_k without sign constraints. The m dual variables y_i corresponding to the inequalities (6) are nonnegative. With the vectors $z = (z_1, \dots, z_n)$ and $y = (y_1, \dots, y_m)$ of variables, the dual LP says: Minimize $\sum_{k=1}^n z_k$ subject to

$$\sum_{i=1}^m y_i \cdot n = n \quad (8)$$

and, according to (5),

$$z_k - \sum_{i=1}^m y_i D_i^k(e_s, \bar{x}) \geq 0 \quad (9)$$

for any pure strategy s of team player k .

Let \bar{z}, \bar{y} be an optimal solution to the dual program. The randomized strategy we are looking for is \bar{y} , which is a vector of probabilities since it is nonnegative and fulfills (8). If the adversary uses this strategy and the strategies of the other team players are fixed, then any pure strategy s of team player k has expected payoff at most \bar{z}^k since (9) says

$$\bar{z}^k \geq \sum_{i=1}^m \bar{y}_i D_i^k(e_s, \bar{x}). \quad (10)$$

We multiply (10) by \bar{x}_s^k and sum over all pure strategies s of team player k . By (10), interchanging sums, (5), and (4) for \bar{x}, \bar{u} , this yields

$$\bar{z}^k = \sum_s \bar{z}^k \bar{x}_s^k \geq \sum_s \sum_{i=1}^m \bar{y}_i D_i^k(e_s, \bar{x}) \bar{x}_s^k = \sum_{i=1}^m \bar{y}_i C_i(\bar{x}) \geq \sum_{i=1}^m \bar{y}_i \bar{u} = \bar{u}.$$

Since the dual LP has the same value of the objective function as the primal LP, that is, $\sum_{k=1}^n \bar{z}_k = n\bar{u}$, this implies $\bar{z}_k = \bar{u}$ for all k . Therefore, according to (10), any pure strategy s of team player k has expected payoff at most \bar{u} if all other team members play as in \bar{x} and the adversary uses \bar{y} . Team member k gets only that payoff when he plays \bar{x}^k , so this strategy is his best response. Furthermore, against the team-maxmin strategy profile \bar{x} the adversary cannot play better than with \bar{y} . This shows the equilibrium property. \square

Nash (1951) showed that any (finite) game has an equilibrium in randomized strategies. For our games, we have proved this independently, by showing that each game has at least one team-maxmin equilibrium. We note that our proof uses LP duality and not a fixed point argument, which might lead to a simpler proof, as suggested by several colleagues; however, we have not been able to find one.

There may be equilibria which are not team-maxmin equilibria. However, no equilibrium of the game can be better for the team than a team-maxmin equilibrium: Let \bar{u} be the team-maxmin payoff, and let u be the payoff in some other equilibrium. The team-maxmin payoff \bar{u} is defined to be the best payoff that the team can guarantee itself. By playing the other equilibrium strategy profile, the team guarantees itself a payoff of u : since it is an equilibrium, the adversary can only hurt himself by diverging unilaterally, and since the game is zero-sum, that can only help the team. (Note that if we had several adversaries, they might be able to benefit by diverging together.) It follows that $u \leq \bar{u}$.

Corollary. *The team-maxmin equilibria are precisely the equilibria of the game with highest payoff to the team.*

It is important to note that the theorem—and our entire approach—depends heavily on the asymmetry of the situation. It cannot be extended to the case of *two* opposing teams, since, for example, the corollary is clearly asymmetric. It is also important to consider the maxmin strategy profile for the team *as a whole*: The payoff guaranteed by an *individual* maxmin strategy of a team player, which assumes the worst possible behavior also for the other team members, is clearly *less* than or equal to any equilibrium payoff.

The team-maxmin equilibria seem to be the most appropriate solution of the game, since they give the team players the highest payoff they can get with uncoordinated strategies.

In general, a game may have more than one team-maxmin equilibrium. Consider the following game, where the team has two members, each of which have two pure strategies. Their mixed strategies are given by probabilities p and q , respectively. The game is defined by the following payoffs:

p, q	0, 0	0, 1	1, 0	1, 1	
T	0	6	4	0	(11)
B	0	4	6	0	

This game has two team-maxmin strategy profiles: $p, q = 1, 0$ and $p, q = 0, 1$, each with team payoff 4. As shown in our theorem, each of these profiles can be extended to a team-maxmin equilibrium. For the first profile, one obtains an equilibrium if the adversary selects T , for the second if she chooses B . Note that team-maxmin strategy profiles and equilibrium strategies of the adversary cannot be combined arbitrarily like in a two-person zero-sum game.

This situation of multiple team-maxmin equilibria can only occur when the payoff matrix contains special entries. In a *generic* game where the payoffs are not known with perfect accuracy, it is unlikely that any two equilibria yield the

same payoff. For example, one of the payoffs 4 in (11) might be slightly larger, which would make the respective team-maxmin strategy profile unique, without otherwise changing the equilibria of the game in an essential way. The team-maxmin equilibrium can thus be considered as a unique solution of the game except for degenerate cases. It provides the appropriate *equilibrium selection* method for our games.

The optimization problem for determining a team-maxmin strategy profile has nonlinear constraints in (4) since the expected payoff $C_i(x)$ contains products of probabilities. In particular, a team-maxmin strategy profile may involve irrational numbers (see Koller and Megiddo 1992, p. 537). We chose the above example (1) with rationals in the team-maxmin strategy profile for simplicity of exposition. An approach to solving the nonlinear optimization problem in (4) are the Karush-Kuhn-Tucker conditions that characterize locally optimal solutions \bar{x}, \bar{u} to (4) and involve the gradients of the objective function and constraints; for an introduction and references see Kuhn (1991).

The theorem also holds under a weaker assumption. It is fairly easy to see that the proof goes through if \bar{x} and \bar{u} are a *locally optimal* solution to the problem in (4). This means that any strategy profile of the team where the worst-case payoff cannot be improved by a local change (of the entire profile) is part of an equilibrium. For example, this applies to both strategy profiles 1, 0 and 0, 1 in (11) even if one of the payoffs 4 there is slightly changed so that only one of these strategy profiles is globally team-maxmin. On the other hand, if we weaken the assumption in a different way the theorem no longer holds. A strategy profile for the team that is *unilaterally* team-maxmin cannot always be extended to an equilibrium. That is, it does not suffice to assume that no team player can increase the team's worst-case payoff by changing his strategy unilaterally.

A counterexample is (1) for the strategy profile $p, q = 0, 0$ where the payoff for both T and B is 3. When either probability p or q is increased alone, the worst-case payoff gets smaller. However, no randomized strategy of the adversary will yield an equilibrium: T would have to be played with probability at least $2/3$ so that the first team player has no incentive to change his strategy p , and B would have to be played with probability at least $1/2$ so that the second team player does not deviate from $q = 0$, which is not possible. Here, a joint change from the strategy profile 0, 0 yields a larger worst-case payoff, so this strategy profile is not locally team-maxmin.

4. Conclusions

We have considered multiperson zero-sum games of a team against an adversary. These games can model teams in an organization, distributed computing, or parlor games like Bridge. In these situations, it is assumed that the team faces a single (and thus rather powerful) adversary. The team members cannot coordinate their

randomized actions during play, but they can decide in advance which profile of mixed strategies they want to use. Thereby, they use a defensive approach where they get the highest worst-case payoff. If the team uses such a *team-maxmin strategy profile*, then the adversary can exploit this situation by playing so that an equilibrium results. That *team-maxmin equilibrium* is the most reasonable solution to the game. It is therefore the appropriate method for *equilibrium selection*.

The payoff of the team-maxmin equilibrium is unique and can be considered as the *value* of the game. A team-maxmin strategy profile may not be unique. In that case, the adversary usually must know which profile is used in order to play such that a team-maxmin equilibrium results. That is, unlike in a two-person zero-sum game, equilibrium strategies are not interchangeable. However, this is in line with the defensive approach taken by the team, where the adversary may possibly learn about the team's strategy profile. Furthermore, only special games with non-generic payoffs have more than one team-maxmin equilibrium.

As we argued, the team-maxmin payoff is the appropriate value for the game of a team against an adversary. This value is typically less than the payoff to the team if the individual team players *could* coordinate their actions. Intuitively, the difference between the two payoffs quantifies the potential value of coordinated action of the team members. It therefore seems that this new definition of a value for the game can be used to measure the value of communication within the team. For example, we can quantify the benefits of changing the information flow within a company. If we represent the game in extensive form, increasing the information flow between team members would correspond to refining the information sets in the game. This will result in a refined strategy space for some of the players, and typically also a better team-maxmin equilibrium value. The difference between the old and new values represents the benefit gained from improving the communication structure in the team. If enough communication is added between the team members, they will be able to coordinate all their actions. Then, the team would essentially act as a single player, and the team-maxmin payoff would be identical to the value of the game when viewed as a two-player game between the team and the adversary.

Finally, we argue that the team-maxmin equilibrium is useful for extensive two-person zero-sum games where one of the players has *imperfect recall*. This topic has been of recent interest (Piccione and Rubinstein 1997). In an extensive game, a player makes a decision at an information set. Often, one thinks of an *agent* of the player making that decision. In the *agent normal form* of the game, the agent for each information set is a separate player. That is, the original player is replaced by a team of agents. A strategy profile of the team is then a *behavior strategy* of the player. The player has *perfect recall* if each agent is informed about the previous choices of the other team members. In that case, behavior strategies are as powerful as mixed strategies (Kuhn 1953).

In general, we can always represent a team as a single player with imperfect recall. Then, a game of the type we have been considering corresponds to a two-person zero-sum game where one player (the adversary) has perfect recall and the other (the team) has imperfect recall. The latter may usually achieve less with a behavior strategy than with a mixed strategy. Koller and Megiddo (1992) have argued, however, that a mixed strategy is not in the spirit of imperfect recall since moves that are coordinated across information sets require the player to remember previous moves for the move he is about to make. They argue that, for a zero-sum game, a maxmin behavior strategy seems to be more appropriate. This is precisely a team-maxmin strategy profile of the team of agents. If the adversary has perfect recall (and thus loses no power playing as a team of her own agents), then our theorem shows that the team-maxmin behavior strategy is even a reasonable two-sided solution concept since it is part of an equilibrium of the agent normal form.

Acknowledgments

We thank Susanne Albers, Ehud Kalai, Richard McKelvey, Nimrod Megiddo, Ralph Werchner, and a referee for interesting discussions and useful suggestions.

References

- Ho, Y. C., and Sun, F. K. (1974). "Value of Information in Two-Team Zero-Sum Problems," *J. Optimization Theory Appl.* **14**, 557–571.
- Kim, K. H., and Roush, F. W. (1987). *Team Theory*. Chichester: Ellis Horwood.
- Koller, D., and Megiddo, N. (1992). "The Complexity of Two-Person Zero-Sum Games in Extensive Form," *Games Econ. Behav.* **4**, 528–552.
- Kuhn, H. W. (1953). "Extensive Games and the Problem of Information," in *Contributions to the Theory of Games II* (H. W. Kuhn and A. W. Tucker, Eds.), Annals of Mathematics Studies, Vol. 28, pp. 193–216. Princeton: Princeton Univ. Press.
- Kuhn, H. W. (1991). "Nonlinear Programming: A Historical Note," in *History of Mathematical Programming — A Collection of Personal Reminiscences*, (J. K. Lenstra, A. H. G. Rinnooy Kan, A. Schrijver, Eds.), pp. 82–96. Amsterdam: CWI and North-Holland.
- Marschak, J., and Radner, R. (1976). *Economic Theory of Teams*. New Haven: Yale Univ. Press.
- Nash, J. F. (1951). "Non-cooperative Games," *Annals of Mathematics* **54**, 286–295.
- Piccione, M., and Rubinstein, A. (1997). "On the Interpretation of Decision Problems with Imperfect Recall", *Games Econ. Behav.* **20**, 3–24.