



King's Research Portal

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):
Ismail, M. (2014). *Maximin equilibrium*.

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

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

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

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Maximin equilibrium

Mehmet ISMAIL

March, 2014

Abstract

We introduce a new solution concept called maximin equilibrium which extends von Neumann's maximin strategy idea to n -player non-cooperative games by incorporating common knowledge of 'rationality' of the players. Our rationality assumption is, however, stronger than the one of maximin strategy and weaker than the one of Nash equilibrium. Maximin equilibrium, just like maximin strategies, is a method for evaluating the uncertainty that players are facing by playing the game. We show that maximin equilibrium is invariant under strictly increasing transformations of the payoff functions. Notably, every finite game possesses a maximin equilibrium in pure strategies. Considering the games in von Neumann and Morgenstern mixed extension, we show that maximin equilibrium is a generalization of Nash equilibrium. In addition, we demonstrate that maximin equilibria and Nash equilibria coincide in two-player zero-sum games. We propose a refinement of maximin equilibrium called strong maximin equilibrium. Accordingly, we show that for every Nash equilibrium that is not a strong maximin equilibrium there exists a strong maximin equilibrium that Pareto dominates it. In addition, no strong maximin equilibrium is ever Pareto dominated by a Nash equilibrium. Finally, we discuss maximin equilibrium predictions in several games including the traveler's dilemma.

1 Introduction

In their ground-breaking book, von Neumann and Morgenstern (1944, p.555) describe the maximin strategy¹ solution for two-player games as follows.

“There exists precisely one solution. It consists of all those imputations where each player gets individually at least that amount which he can secure for himself, while the two get together precisely the maximum amount which they can secure together. Here the ‘amount which a player can get for himself’ must be understood to be the amount which he can get for himself, irrespective of what his opponent does, even assuming that his opponent is guided by the desire to inflict a loss rather than to achieve again.”

This immediately gives rise to the following question: ‘What happens when a player acts according to the maximin principle but knowing that other players do not necessarily act in order to decrease his payoff?’. We are going to capture this type of behavior by assuming that players are ‘rational’² and letting it be common knowledge among players. In other words, the contribution of the current paper can roughly be considered as incorporating the maximin principle and rationality of the players in one concept calling it maximin equilibrium.

Note that it is recognized and explicitly stated by von Neumann and Morgenstern several times that their approach can be questioned by not capturing the cooperative side of non-zero-sum games. But this did not seem a big problem at that time and it is stated that the applications of the theory should be seen in order to reach a conclusion.³ After more than a half-century of research in this area, maximin strategies are indeed considered to be too defensive in non-strictly competitive games in the literature. Since a maximin strategist plays any game as if it is a zero-sum game, this leads

¹We would like to note that the famous minimax (or maximin) theorem was proved by von Neumann (1928). Therefore, it is generally referred as von Neumann’s theory of games in the literature.

²Throughout the text, we will specify in which context rationality is used to avoid confusion, e.g. rationality in maximin strategies, rationality in Nash equilibrium and so on. The word ‘rationality’ alone will be used when we do not attach any mathematical definition to it (until we formally define).

³For example, see vN-M (1944, p.540).

to an ignorance of her opponent's payoffs and hence the preferences of her opponent. These arguments call for a revision of maximin strategy concept.

Let us consider the following games to support our statement. In the first game⁴, Alfa (he) is the row player and Beta (she) is the column player. There is a unique mixed Nash equilibrium $[(\frac{3}{4}, \frac{1}{4}), (\frac{1}{2}, \frac{1}{2})]$ at which Alfa receives an expected payoff of 3 and Beta receives 4. On the other hand, Alfa has a unique maximin strategy $(\frac{1}{2}, \frac{1}{2})$ which guarantees him to receive an expected payoff of 3. Beta also has a unique maximin strategy $(\frac{1}{4}, \frac{3}{4})$ which guarantees her 4. Although the point we want to make is different, it is of importance to note the historical discussion about this type of games where Nash equilibrium payoffs are equal to the payoffs that can be guaranteed by playing maximin strategies. Harsanyi (1966) postulates that players should use their maximin strategies in those games. Aumann and Maschler (1972) state that they do not know what to recommend but maximin strategies seem preferable.

	a	b
a	5, 1	1, 5
b	1, 13	5, 1

In short, in the games similar as above, the arguments supporting maximin strategies are so strong that it led some game theory giants to prefer them over the unique Nash equilibrium of the game. However, these arguments suddenly disappear and the weakness of maximin strategies can be easily seen if we add a strategy *trick* to the previous game for both players as in the following game. Let the payoffs be as given in Figure 1 with some small $\epsilon > 0$. For every $\epsilon > 0$ there is one completely mixed⁵ and one pure Nash equilibrium $(trick, trick)$ whose payoffs to both players are 0. The third Nash equilibrium $[(\frac{3}{4}, \frac{1}{4}, 0), (\frac{1}{2}, \frac{1}{2}, 0)]$ is the same as the previous game and its payoff vector is $(3, 4)$.

However, notice that the maximin strategies of the previous game have disappeared. The new maximin strategy is *trick* which guarantees 0 while the other strategies guarantee $-\epsilon$ for a player. To elaborate on the reason of this, let us first see why the old Nash equilibrium survived.

⁴It is strategically equivalent to the game in Aumann and Maschler (1972,p.55).

⁵For example, if $\epsilon = 1$ then mixed Nash equilibrium is $[(\frac{3}{20}, \frac{1}{20}, \frac{4}{5}), (\frac{1}{8}, \frac{1}{8}, \frac{3}{4})]$ whose payoff is 0 for both players.

	a	b	<i>trick</i>
a	5, 1	1, 5	$-\epsilon, 0$
b	1, 13	5, 1	$-\epsilon, 0$
<i>trick</i>	0, $-\epsilon$	0, $-\epsilon$	0, 0

Figure 1: A game with an $\epsilon > 0$.

Suppose that players made an agreement (explicitly or implicitly) to play the Nash equilibrium $[(\frac{3}{4}, \frac{1}{4}, 0), (\frac{1}{2}, \frac{1}{2}, 0)]$.⁶ Then, Alfa would make sure that Beta would not unilaterally deviate to the strategy *trick* because Beta is rational (à la Nash). That is, deviating gives 0 to Beta which is strictly less than what she would receive if she did not deviate. Beta would also make sure that Alfa would not make a unilateral deviation to *trick* for the same reason. Therefore, one observes that Nash equilibrium is immune to this sort of additions of strategies whose payoffs are strictly less than the ones of the original game.⁷ It may create additional Nash equilibria though. We extend the maximin principle in such a way that it also becomes immune to these type of ‘tricks’.

To see what actually happened to the maximin strategies of the first game let us look at its profile $[(\frac{1}{2}, \frac{1}{2}, 0), (\frac{1}{4}, \frac{3}{4}, 0)]$ in the second game. Suppose that players agreed (explicitly or implicitly) to play this profile. Alfa would make sure that Beta would not unilaterally deviate to the strategy *trick* if Beta is ‘rational’ because she receives 0 by deviating which is strictly less than what she would receive, namely 4. Beta would also make sure that Alfa would not unilaterally deviate to *trick* if he is ‘rational’. A player might still deviate to *a* or *b* but this is okay for both players since they both guarantee their respective payoffs in this region.⁸ That is, Alfa would guarantee to receive 3 given that Beta is ‘rational’ and Beta guarantees to receive 4 given that Alfa is ‘rational’.

⁶Note that it is not in general known how players coordinate or agree on playing a specific Nash equilibrium. We would like to see ‘if’ they agree then what happens. As it is stated in Aumann (1990), a player does not consider an agreement as a direct signal that her opponent will follow it. But by making an agreement, she rather understands that the other player is signalling that he wants her to keep it.

⁷If, for example, the payoffs to *trick* were 6 instead of 0, then obviously this would not be the case.

⁸Note that these are the maximin strategies of the 2 by 2 game.

In conclusion, we think that if players prefer playing their maximin strategies in the first game, then they would prefer playing the strategy profile $[(\frac{1}{2}, \frac{1}{2}, 0), (\frac{1}{4}, \frac{3}{4}, 0)]$ in the second game if their ‘rationality’ is commonly known.

In Section 2, we define the rationality assumption we would like to use in this paper. In Section 3, we introduce the concept of maximin equilibrium. Similar to maximin strategies, maximin equilibrium is a method for evaluating the uncertainty that players are facing by playing the game. We show that maximin equilibrium is invariant under strictly increasing transformations of the payoff functions of the players. Moreover, every finite game possesses a maximin equilibrium in pure strategies.

In Section 4, we consider the games in von Neumann-Morgenstern mixed extension and discuss the relationship of maximin equilibrium with the other solution concepts. We show that maximin equilibrium is a coarsening of Nash equilibrium. Moreover, we demonstrate that maximin equilibrium coincides with Nash equilibrium in two-player zero-sum games. We also propose a refinement of maximin equilibrium called strong maximin equilibrium. Accordingly, we show that for every Nash equilibrium that is not a strong maximin equilibrium there exists a strong maximin equilibrium that Pareto dominates it. In addition, no strong maximin equilibrium is ever Pareto dominated by a Nash equilibrium. Furthermore, we show by examples that maximin equilibrium is neither a coarsening nor a special case of correlated equilibrium or rationalizable strategy profiles.

In Section 5, we discuss the maximin equilibrium in extensive form games and in n-player games. All the results provided in Section 3 and in Section 4 hold in n-player games except the one which requires a two-player zero-sum setting. Finally, we discuss maximin equilibrium predictions in several games including the traveler’s dilemma.

2 The framework

In this paper, we use a framework for the analysis of interactive decision making environments as described by von Neumann and Morgenstern (1944, p.11).

One would be mistaken to believe that it [the uncertainty] can be obviated, like the difficulty in the Crusoe case mentioned in

footnote 2 on p. 10, by a mere recourse to the devices of the theory of probability. Every participant can determine the variables which describe his own actions but not those of the others. Nevertheless those ‘alien’ variables cannot, from his point of view, be described by statistical assumptions. This is because the others are guided, just as he himself, by rational principles –whatever that may mean– and no *modus procedendi* can be correct which does not attempt to understand those principles and the interactions of the conflicting interests of all participants.

For simplicity, we assume that there are two players whose finite set of pure actions are X_1 and X_2 respectively. Moreover, players’ preferences over the outcomes are assumed to be a weak order (i.e. transitive and complete) so that we can represent those preferences by the (ordinal) utility functions $u_1, u_2 : X_1 \times X_2 \rightarrow \mathbb{R}$ which depends on both player’s actions. As usual, the notation x in $X = X_1 \times X_2$ represents a strategy profile. In short, a two-player non-cooperative game can be denoted by $(\{1, 2\}, X_1, X_2, u_1, u_2)$. Starting from simple strategic decision making situations, we firstly introduce a deterministic theory of games in this section and in the following one.⁹

For the analysis of a game we need a notion of ‘rationality’ of the players. Moreover, we assume that players maximize utility with respect to the worst case given that the ‘rationality’ of the players is common knowledge. We assume that everything about the game including the previous sentence is common knowledge.¹⁰

In one-player decision making situations the notion of rationality is usually referred as maximizing one’s own utility with respect to her preferences. In multi-player games, however, it is not unambiguous what does it mean to maximize one’s own utility because it simply depends on the other’s actions. Von Neumann proposed an approach to do this: Each player should maximize a minimum utility regardless of the strategy of the other player. Although, in two-player zero-sum games this works quite well, it is considered to be too pessimistic in non-zero-sum games since the preferences

⁹Note that all the definitions we present can be extended in a straight-forward way to the games in von Neumann-Morgenstern mixed extension and to n-player games. See Section 4 and in the Section 5 for more discussion. Note also that interpreting utilities as cardinal which is usual in game theory does not cause any problems.

¹⁰See Lewis (1969) for more details and see Aumann (1976) for a formal definition of common knowledge in a Bayesian setting.

of the players are not necessarily opposing. Due to the fact that rationality of the maximin strategist dictates her to play any game as if it is a zero-sum game, this leads her to ignore the opponent's payoffs and hence the preferences of the opponent. Therefore, one might call maximin strategist being an individually rational player who does not, in general, take the advantage of all the information given to her.

2.1 Definition of rationality

To avoid possible misunderstandings from our rationality definition, one may consider the following scenario in mind which is intended only to facilitate the definition.

Suppose that Alfa and Beta are sitting at a bargaining table and can make non-binding agreements. However, they will submit their strategies in separate rooms simultaneously. By making an agreement on a strategy profile and going back to their rooms, each player faces a decision problem of either keeping the agreement or betraying it by a deviation. Accordingly, we discuss what would players do if they are rational in the sense of maximin strategy, Nash equilibrium and of the notion we define as maximin equilibrium.

Firstly, if the players are rational à la von Neumann, then a player is expected to betray an agreement by deviating to any strategy (including the ones which he gets a strictly less payoff than the agreement payoff). The reason is obvious: Maximin rule imposes no condition on the other player's strategies. Secondly, if we assume that players are rational à la Nash, then a player is expected to betray the agreement unless he is already best replying.

Note that these are two extreme cases that might occur: In the former, there is no condition on the strategy of the other player while there is unique condition in the latter one. We define a rationality concept which is stronger than of maximin strategist but weaker than of Nash strategist. According to our definition, a rational player will be the one who would not betray the agreement by making a non-profitable deviation.¹¹

Let us fix some terminology. As usual, a strategy $x'_i \in X_i$ is said to be a profitable deviation for player i with respect to the profile (x_i, x_j) if $u_i(x'_i, x_j) > u_i(x_i, x_j)$. We call it non-profitable if it is not profitable.

¹¹Note that one can define a bit stronger or a bit weaker notion of rationality but working with the current one suffices for our purposes.

Definition 1. A player is called rational at (the agreement) x if she *does not* make a non-profitable deviation from it.

In other words, when Alfa and Beta leave the bargaining table, Alfa would make sure that Beta would not betray the agreement by deviating non-profitably if Beta is rational and vice versa.

The reader might ask whether we need to assume that pre-play communication is allowed. The answer is that we do not need it; an agreement can be explicit or implicit. The assumptions of rationality and players being maximin decision makers are, basically, the ‘rules’ that players take into account in order to reach an agreement. Since those rules are fixed and common knowledge, players can make the same reasoning in the lack of communication as they do in the presence of a communication.

At first sight, it might seem subtle that a player makes a reasoning about a strategy profile which actually includes her own strategy in the above interpretation. However, careful examination reveals that it is not a problem since she is reasoning about a possibility that she makes an agreement. Initially, every strategy profile is a potential agreement.

3 Maximin equilibrium

As it is formulated and explained by von Neumann (1928), playing a game is basically facing an uncertainty which can not be resolved by statistical assumptions. This is actually the crucial difference between games and decision problems. Our aim is to extend von Neumann’s approach on resolving this uncertainty.

Suppose that Alfa and Beta make a non-binding agreement (x_1, x_2) . Alfa faces an uncertainty by keeping the agreement since he does not know whether Beta will keep it. Von Neumann’s method to evaluate this uncertainty is to calculate the minimum payoff of Alfa with respect to all conceivable deviations by Beta.¹² That is, Alfa’s evaluation $v_{x_1x_2}$ (or the utility) of keeping the agreement (x_1, x_2) is $v_{x_1x_2} = \min_{x'_2 \in X_2} u_1(x_1, x'_2)$. Note that for all x'_2 , the evaluation of Alfa for the profile (x_1, x'_2) is the same, i.e. $v_{x_1x_2} = v_{x_1x'_2}$ for all $x'_2 \in X_2$. Therefore, it is possible to attach a unique evaluation $v_{x'_1}$ for every strategy $x'_1 \in X_1$ of Alfa. Next step is to make a

¹²Because, it is assumed that Beta might have a desire to inflict a loss for Alfa. Note that von Neumann used mixed strategies but here we would like to keep it simple.

comparison between those evaluations of the strategies. For that, von Neumann takes the maximum of all such evaluations $v_{x'_1}$ with respect to x'_1 which yields a unique evaluation for the whole game, i.e. the value of the game is $v = \max_{x'_1 \in X_1} v_{x'_1}$. In other words, the unique utility that Alfa can guarantee by facing the uncertainty of playing this game is v . Accordingly, it is recommended that Alfa should choose the strategy $x_1^* \in \arg \max_{x'_1 \in X_1} v_{x'_1}$ which guarantees the value v .

We would like to extend von Neumann's method in such a way that Alfa takes into account the rationality of Beta when making the evaluations and vice versa. Let us construct the approach we take step by step and state its implications.

Firstly, one needs to define the rationality. We have chosen a notion of rationality which allows Beta to keep her agreement or to deviate to a strategy for which she has strict incentives to do so. By this assumption, Alfa can rule out non-profitable deviations of Beta from the agreement (x_1, x_2) which helps decreasing the level of uncertainty he is facing. Now, Alfa's evaluation $v_{x_1 x_2}$ of the uncertainty for keeping the agreement can be defined as the minimum utility he would get under any rational behavior of Beta. That is,

$$v_{x_1 x_2} = \min \left\{ \min_{x'_2 \in B_2(x_1, x_2)} u_1(x_1, x'_2), u_1(x_1, x_2) \right\},$$

where $B_2(x_1, x_2)$ denotes the set of Beta's (strict) better replies with respect to (x_1, x_2) . As a consequence, it is not in general true for a strategy $x'_2 \neq x_2$ that we have the equality $v_{x_1 x_2} = v_{x_1 x'_2}$. Because, the better response set of Beta with respect to (x_1, x_2) is not necessarily the same as the better response set of her with respect to (x_1, x'_2) . Therefore, we can not assign a unique value to every strategy of Alfa anymore. Instead, the evaluation of the uncertainty can be encoded in the strategy profile. As a result, we can no longer refer a strategy in the same spirit of a maximin strategy since a strategy in this setting only makes sense as a part of a strategy profile.¹³ But note that there are two evaluations that are attached to the profile (x_1, x_2) , one from Alfa and one from Beta since she also is doing the similar calculations as him. Accordingly, we define a value function which maps each strategy profile to a vector whose components represent the value of the profile to each player.

¹³This is similar to Nash equilibrium. A strategy in a Nash equilibrium is meaningful only if it is considered with the other strategies in the profile.

The last step is to make comparisons between those evaluations. Our aim is basically to maximize the value function but since it is vector valued there is no obvious way to do this. At first sight, it might seem that we have not done a big improvement so far since we have started with a maximization of payoff vectors of the players but we are still left with a similar problem. However, we can now use the maximization (optimization) methods we know in the literature just as techniques. Two well-known maximization techniques are Pareto optimality and Nash equilibrium.

One may ask why we do not directly apply those maximization methods to the set of payoff vectors of the game. The reason is that applying a maximization method to a game does not necessarily mean evaluating and comparing the uncertainty that a player faces by playing according to what it recommends.

For example, let the strategy profile z be Pareto optimal in the whole domain of a game. Even though the payoff that Alfa receives from the profile z is high, playing her part of the Pareto optimal profile need not be a good (utility maximizing) decision *ex ante*. Because, the uncertainty he faces by playing his part is not evaluated. For example, it could be the case that Beta has a chance to deviate from z which causes a substantial decrease in his payoff. Therefore, it would be better for Alfa if he considers his possible payoffs after such deviations and compare it with the other possible options.

To give another example, let now w be a Nash equilibrium. Again, the fact that w is a Nash equilibrium does not make playing his part of w a good decision for Alfa. Alfa needs to evaluate the uncertainty he is facing by playing his part and compare it with the other options. Actually, it is better to call it a certainty rather than an uncertainty because of the fact that neither Beta nor Alfa has a profitable deviation from it. However, it is still an unevaluated and uncomparing certainty in the game.

Loosely speaking, choosing the Nash equilibrium in the presence of many other strategy profiles is like choosing the sure gamble in the presence of many other uncertain gambles. Even if the outcomes of the uncertain gamble is very high in every state of the world and the outcome of the sure gamble is very low.

Nash's method only compares the outcomes that might occur as a consequence of a player choosing one strategy with the outcomes that might occur as a consequence of an opponent choosing another strategy.¹⁴ Therefore, a

¹⁴If they are in mutual best response, then it is called a Nash equilibrium.

Nash equilibrium completely ignores the outcomes that might occur under any other strategy choices of the players no matter how high their utilities are. If, by any chance, there is only one Nash equilibrium, then this ignorance might lead to a disastrous outcome for both players. One can see this clearly in the traveler's dilemma game which was introduced by Basu (1994).¹⁵ If players play the unique Nash equilibrium, then they ignore a large part of the game which is actually profitable for both of them. At the end of the day, what a utility maximizing player cares is the payoff she receives not the nice property of being sure that her opponent would not deviate had they agreed to play a profile.

Another question one might ask is why we use both Pareto optimality and Nash equilibrium methods for maximization of the value function. The answer is that we simply want to minimize potential disadvantages of using only one of the two. It is well-known that Pareto optimality sometimes yields quite unfair outcomes and that a Nash equilibrium outcome can be quite unsatisfactory. Let us now define the value function formally.

Definition 2. Let $\Gamma = (X_1, X_2, u_1, u_2)$ be a two-player game. A function $v : X \rightarrow \mathbb{R} \times \mathbb{R}$ is called the value function of Γ if for every $i \neq j$ and for all $x = (x_i, x_j) \in X$, i 'th component of $v = (v_i, v_j)$ satisfies

$$v_i(x) = \min\left\{\min_{x'_j \in B_j(x)} u_i(x_i, x'_j), u_i(x)\right\},$$

where the better response correspondence of player j with respect to x is defined as

$$B_j(x) = \{x'_j \in X_j | u_j(x_i, x'_j) > u_j(x)\}.$$

Remark 3. Note that for all x and all i , we have $u_i(x) \geq v_i(x)$. Because, one cannot increase a payoff but can only (weakly) decrease it, by definition of the value function.

As it is discussed above, the value of a strategy profile to a player represents the minimum payoff that a player could receive under rational behavior of her opponent. To illustrate what a value function of a game looks like, let us consider the game in Figure 2 which is played by Alfa and Beta.

¹⁵The traveler's dilemma is a symmetric game in which players can pick a number from 2 to 100 and the one who picks the lower number receives the Dollar amount equals to her choice plus 2\$ reward and the other receives 2\$ punishment from the lower number. If both choose the same number they get what they choose. For the game matrix and for more discussion see p.15.

		A	B	C	D
$\Gamma =$	A	2, 2	0, 0	1, 1	0, 0
	B	0, 0	90, 80	3, 3	90, 90
	C	1, 100	100, 80	-1, -1	-4, -2
	D	3, 1	75, 0	0, 0	230, 0

		A	B	C	D
$v(\Gamma) =$	A	2, 1	0, 0	1, 1	0, 0
	B	0, 0	90, 80	3, 3	90, 0
	C	1, 1	1, 80	-1, -1	-4, -2
	D	3, 1	3, 0	0, 0	3, 0

Figure 2: A game Γ and its value function $v(\Gamma)$.

Observe that Γ has a unique Nash equilibrium (D,A) whose payoff vector is (3,1). Suppose that pre-game communication is allowed and that Beta is trying to convince Alfa at the bargaining table to make an agreement on playing, for example, the profile (C,B) which Pareto dominates the Nash equilibrium. Alfa would fear that Beta may not keep his agreement and may unilaterally deviate to A leaving her a payoff of 1. Accordingly, the value of the profile (C,B) to Alfa is 1 as shown in the bottom table in Figure 2. Since agreements are not binding, Alfa would not accept this agreement. Now, suppose Alfa offers to make an agreement on (B,B). Beta would not fear a unilateral profitable deviation C of Alfa since his payoff does not decrease in that case. Alfa's payoff does not decrease too in case of a unilateral profitable deviation of Beta to D. In other words, the value of the profile (B,B) is (90, 80) which is equal to its payoff vector in Γ .

Now, we are ready to define the maximin equilibrium.

Definition 4. A strategy profile $x = (x_i, x_j)$ where $i \neq j$ in a two player game Γ is called maximin equilibrium if it satisfies at least one of the following conditions for the value function $v = (v_i, v_j)$:

1. For every player i and all $x' \in X$, $v_i(x') > v_i(x)$ implies $v_j(x') < v_j(x)$.
2. For every i , $x_i \in \arg \max_{x'_i \in X_i} v_i(x'_i, x_j)$.

In addition, x is called a strong maximin equilibrium if it satisfies only the first condition above.

Remark 5. We can interpret the value function v of a game Γ as a game in its own right, that is, $\Gamma_v = (X_1, X_2, v_1, v_2)$ where $v(\Gamma) = (v_1, v_2)$. Then, the set of maximin equilibria in Γ is the union of the set of Pareto optimal strategy profiles and the set of Nash equilibria of the game Γ_v .

Going back to our example, observe that the profile (B,B) is the Pareto dominant profile of the value function of the game Γ shown in Figure 2, so it is a maximin equilibrium. In addition, observe that the Nash equilibrium is also a maximin equilibrium. Moreover, the maximin equilibrium (B,B) has another property which deserves attention. Suppose that players agree on playing it. Alfa has a chance to make a unilateral profitable deviation to C but he can not rule out a potential profitable deviation of Beta to the strategy D. If this happens, Alfa would receive -4 which is less than what he would receive if he did not deviate. But Beta is also in the exactly same situation. As a result, nobody would actually deviate from the agreement (B,B).

An ordinal utility function is unique up to strictly increasing transformations. Therefore, it is crucial for a solution concept (which is defined with respect to ordinal utilities) to be invariant under those operations. The following proposition shows that maximin equilibrium possesses this property.

Proposition 6. *Maximin equilibrium is invariant under strictly increasing transformations of the payoff function of the players.*

Proof. Let $\Gamma = (X_i, X_j, u_i, u_j)$ and $\hat{\Gamma} = (X_i, X_j, \hat{u}_i, \hat{u}_j)$ be two games such that \hat{u}_i and \hat{u}_j are strictly increasing transformations of u_i and u_j respectively. Firstly, we show that the components \hat{v}_i and \hat{v}_j of the value function \hat{v} are strictly increasing transformations of the components v_i and v_j of v , respectively. Notice that,

$$\begin{aligned} B_j(x) &= \{x'_j \in X_j | u_j(x_i, x'_j) > u_j(x)\} \\ &= \hat{B}_j(x) = \{x'_j \in X_j | \hat{u}_j(x_i, x'_j) > \hat{u}_j(x)\}. \end{aligned}$$

It implies that $\arg \min_{x'_j \in B_j(x)} u_i(x_i, x'_j) = \arg \min_{x'_j \in \hat{B}_j(x)} \hat{u}_i(x_i, x'_j)$. It in turn implies that $v_i(x) = \min\{u_i(x_i, \bar{x}_j), u_i(x)\}$ and $\hat{v}_i(x) = \min\{\hat{u}_i(x_i, \bar{x}_j), \hat{u}_i(x)\}$ for some $\bar{x}_j \in \arg \min_{x'_j \in B_j(x)} u_i(x_i, x'_j)$.

Since \hat{u}_i is a strictly increasing transformation of u_i , we have either $v_i(x) = u_i(x_i, \bar{x}_j)$ if and only if $\hat{v}_i(x) = \hat{u}_i(x_i, \bar{x}_j)$ or $v_i(x) = u_i(x)$ if and only if $\hat{v}_i(x) = \hat{u}_i(x)$ for all x_i, x_j and all \bar{x}_j .

	L	R
L	0, 0	2, -2
R	3, -3	1, -1

	L	R
L	11, 11	13, 6
R	16, 5	12, 10

Figure 3: Two ordinally equivalent games.

It follows that showing $v_i(x) \geq v_i(x')$ if and only if $\hat{v}_i(x) \geq \hat{v}_i(x')$ is equivalent to showing $u_i(x) \geq u_i(x')$ if and only if $\hat{u}_i(x) \geq \hat{u}_i(x')$ for all x, x' in X which is correct by our supposition.

Secondly, a profile y is a Pareto optimal profile with respect to v if and only if it is Pareto optimal with respect to \hat{v} because each v_i is a strictly increasing transformation of \hat{v}_i . By the same argument, a profile y is a pure Nash equilibrium in the game Γ^v if and only if it is a pure Nash equilibrium in $\Gamma^{\hat{v}}$. As a result, the set of maximin equilibria of Γ and $\hat{\Gamma}$ are the same. \square

The following proposition shows the existence of strong maximin equilibrium hence the existence of maximin equilibrium with respect to pure strategies. It is especially of importance in games where players can not or are not able to use a randomization device. It might be also the case that a commitment of a player to a randomization device is implausible. In those games, we can make sure that there exists at least one maximin equilibrium.

Theorem 7. *Every finite game has a strong maximin equilibrium.*

Proof. Since the Pareto dominance relation is transitive a Pareto optimal strategy profile with respect to the value function of a finite game always exists. \square

One might wonder the prediction of pure maximin equilibrium in zero-sum games which possesses no pure Nash equilibrium where mixed strategy extension of the game is not possible. The zero-sum game in Figure 3 is taken from Aumann and Maschler (1972) and the value of the game for player 1 is 1.5 in mixed extension. This game was subject to discussion because of the problem of commitment to the maximin strategies. Before the coin toss, the maximin strategy $(\frac{1}{2}, \frac{1}{2})$ of Alfa guarantees the highest expected payoff,

	100	99	...	3	2
100	100, 100	97, 101	...	1, 5	0, 4
99	101, 97	99, 99	...	1, 5	0, 4
⋮	⋮	⋮	⋮	⋮	⋮
3	5, 1	5, 1	...	3, 3	0, 4
2	4, 0	4, 0	...	4, 0	2, 2

Figure 4: Traveler’s dilemma

however, after the coin toss player 1 still needs to make a decision whether playing according to the outcome of the toss or not. Actually, playing the strategy R guarantees more than playing L after the toss. Confirming the intuitions about the game in Figure 3, the pure maximin equilibrium of this game is (R,R). However, note that if the utility functions of a zero-sum game are ordinal, then the usual intuition of zero-sum games may not hold. For example, the game on the bottom in Figure 3, is ordinally equivalent to the zero-sum game on the top.

Let us now look at the traveler’s dilemma game in Figure 4. The payoff function of a player i if she plays x_i and her opponent plays x_j is defined as $u_i(x_i, x_j) = \min(x_i, x_j) + 2\text{sgn}(x_j - x_i)$ for all x_i, x_j in $\{2, 3, \dots, 100\}$. In this game, the unique strict Nash equilibrium is (2, 2) which is also the only rationalizable strategy profile. To find the maximin equilibria we first need to compute the value of the traveler’s dilemma game. With appropriate calculations, one can observe that profiles (100, 100) and (2, 2) are maximin equilibria. The profile (100, 100) is the strong maximin equilibrium whose value is (97, 97).¹⁶

It is shown by many experiments that the players do not on average choose the Nash equilibrium strategy but choose 100 or something close to 100. Note that changing the reward/punishment level (which is 2 in the example) effects the behavior observed in experiments. Goeree and Holt (2001) found that when reward is high 80% of the subjects choose the Nash equilibrium strategy but when reward is small about 80% of the subjects choose the highest. This finding is a confirmation of Capra et al. (1999), that is, when reward was high the play converged towards the Nash equilibrium over time but when

¹⁶Note that changing the reward/punishment level (which is 2 in the example) does not effect the set of maximin equilibria. However, the strong maximin equilibrium can change depending on the magnitude of the reward.

reward level was small the play converged towards the other extreme. On the other hand, Rubinstein (2007) found (in a web-based experiment without payments) that 55% of 2985 subjects choose the highest amount and only 13% choose the Nash equilibrium where the reward was small.

These results are actually not unexpected. The irony is that if both players choose almost¹⁷ any ‘irrational’ strategy but their Nash equilibrium strategy then they both get strictly more payoff than they would get by playing the Nash equilibrium. Actually, the Nash equilibrium is the only profile which has this property. Therefore, it is difficult to imagine that a player whose aim is to maximize her own payoff would ever play or expect her opponent to play the strategy 2 in this game unless she is a ‘victim’ of game theory.

4 Maximin equilibrium in mixed strategies

4.1 Definition

The mixed extension of a two-player non-cooperative game is denoted by $(\Delta X_1, \Delta X_2, u_1, u_2)$ where ΔX_i is the set of all simple probability distributions (lotteries) over the set X_i . It is assumed that the preferences of the players over the lotteries satisfy weak order, continuity and the independence axioms.¹⁸ As a result, those preferences can be represented by von Neumann-Morgenstern (expected) utility functions $u_1, u_2 : \Delta X_1 \times \Delta X_2 \rightarrow \mathbb{R}$. Let $p \in \Delta X$ denote a strategy profile where $\Delta X = \Delta X_1 \times \Delta X_2$.

We do not need another definition for maximin equilibrium with respect to mixed strategies; one can just interpret the strategies in the Definition 4 as being mixed. But we need to modify the value function as follows to have it well defined. The value function of player i is defined as $v_i(p) = \min\{\inf_{p'_j \in B_j(p)} u_i(p_i, p'_j), u_i(p)\}$ for all $p \in \Delta X$.¹⁹

Harsanyi and Selten (1988, p.70) argue that invariance with respect to positive linear transformations of the payoffs is a fundamental requirement for a solution concept. The reason is obvious: The von Neumann-Morgenstern utility function of a player is unique up to positive linear transformations.

¹⁷If one modifies the payoffs of the game such that $u_i(x_i, 3) = 2.1$ for all i and all $x_i \in \{4, 5, \dots, 100\}$. Then, one can remove ‘almost’ from the sentence.

¹⁸For more information see, for example, Fishburn (1970).

¹⁹The value function is well defined since utility functions are bounded.

The following remark states that maximin equilibrium has this property. We do not provide a proof since it follows essentially the same steps as the proof of Proposition 6.

Remark 8. The maximin equilibria of a game in mixed extension is unique up to positive linear transformations of the payoffs.

The following lemma illustrates a useful property of the value function of a player.

Lemma 9. *The value function of a player is upper semi-continuous.*

Proof. In several steps, we show that the value function v_i of player i in a game $\Gamma = (\Delta X_1, \Delta X_2, u_1, u_2)$ is upper semi-continuous.

Firstly, we show that the better reply correspondence $B_j : \Delta X_i \times \Delta X_j \rightarrow \Delta X_j$ is lower hemi-continuous. For this, it is enough to show the graph of B_j defined as follows is open.

$$Gr(B_j) = \{(p_j, q) \in \Delta X_j \times \Delta X \mid p_j \in B_j(q)\}.$$

$Gr(B_j)$ is open in $\Delta X_j \times \Delta X$ if and only if its complement is closed. Let $[(p_j, q_i, q_j)^k]_{k=1}^\infty$ be a sequence in $[Gr(B_j)]^c = (\Delta X_j \times \Delta X) \setminus Gr(B_j)$ converging to (p_j, q_i, q_j) where $p_j^k \notin B_j(q^k)$ for all k . It follows that we have $u_j(p_j^k, q_i^k) \leq u_j(q^k)$ for all k . Continuity of u_j implies that $u_j(p_j, q_i) \leq u_j(q)$ which means $p_j \notin B_j(q)$. Hence $[Gr(B_j)]^c$ is closed which implies B_j is lower hemi-continuous.

Next, we define $\hat{u}_i : \Delta X_j \times \Delta X_i \times \Delta X_j \rightarrow \mathbb{R}$ by $\hat{u}_i(p_j, q_i, q_j) := u_i(p_j, q_i)$ for all $(p_j, q_i, q_j) \in \Delta X_j \times \Delta X_i \times \Delta X_j$. Since u_i is continuous, \hat{u}_i is also continuous. In addition, we define $\bar{u}_i : Gr(B_j) \rightarrow \mathbb{R}$ as the restriction of \hat{u}_i to $Gr(B_j)$, i.e. $\bar{u}_i = \hat{u}_i|_{Gr(B_j)}$. The continuity of \hat{u}_i implies the continuity of its restriction \bar{u}_i which in turn implies \bar{u}_i is upper semi-continuous.

By the theorem of Berge (1963, p.115)²⁰ lower hemi-continuity of B_j and lower semi-continuity of $-\bar{u}_i : Gr(B_j) \rightarrow \mathbb{R}$ implies that the function $-\bar{v}_i : \Delta X_i \times \Delta X_j \rightarrow \mathbb{R}$ defined by $-\bar{v}_i(q) := \sup_{p_j \in B_j(q)} -\bar{u}_i(p_j, q)$ is lower semi-continuous.²¹ It implies that the function $\bar{v}_i(q) = \inf_{p_j \in B_j(q)} \bar{u}_i(p_j, q)$ is upper semi-continuous.

²⁰We follow a version of the theorem as presented in Aliprantis and Border (1994, p.569).

²¹We use the fact that a function is lower semi-continuous if and only if the negative of the function is upper semi-continuous. Moreover, note that a function is continuous if and only if it is both lower semi-continuous and upper semi-continuous.

The value function of player i defined by $v_i(q) := \min\{\bar{v}_i(q), u_i(q)\}$ is upper semi-continuous because the minimum of two upper semi-continuous functions is also upper semi-continuous. \square

The following theorem shows that strong maximin equilibrium and hence maximin equilibrium exists in mixed strategies.

Theorem 10. *Strong maximin equilibrium exists in mixed strategies.*

Proof. Let us define $v_i^{max} := \arg \max_{q \in \Delta X} v_i(q)$ which is a non-empty compact set because ΔX is compact and v_i is upper semi-continuous by Lemma 9. Since v_i^{max} is compact and v_j is also upper semi-continuous the set $v_{ij}^{max} := \arg \max_{q \in v_i^{max}} v_j(q)$ is non-empty and compact. Clearly, the profiles in v_{ij}^{max} are Pareto optimal with respect to the value function which means v_{ij}^{max} is a non-empty compact subset of the set of strong maximin equilibria in the game Γ . Similarly, one may show that the set v_{ji}^{max} is also a non-empty compact subset of the set of strong maximin equilibria. \square

Now, let us assume that players can use mixed strategies in the game Γ in Figure 2. An interesting phenomenon occurs if we change, ceteris paribus, the payoff of $u_1(C, D)$ from -3 to -4 . Let us call the new game $\hat{\Gamma}$. It has the same pure Nash equilibrium (D,A) as Γ plus two mixed one. The Pareto dominant Nash equilibrium is $[(0, \frac{41}{46}, \frac{5}{46}, 0), (0, \frac{47}{52}, 0, \frac{5}{52})]$ whose expected payoff vector is $(90, 80)$.²² The question arises: What is the ceteris paribus effect of increasing the payoff of $u_1(C, D)$ from -4 to -3 on the nature of this strategic decision making situation that it decreases the payoffs of the players dramatically in Nash equilibria? Note that by passing from $\hat{\Gamma}$ to Γ we just slightly increase Alfa's relative preference of the worst outcome (C,D) with respect to the other outcomes. On the other hand, there is a maximin equilibrium $[B, (0, \frac{28}{31}, 0, \frac{3}{31})]$ in Γ whose value is approximately 80.9 for both players. Moreover, it remains to be a maximin equilibrium with the same value in $\hat{\Gamma}$.²³

²²The other Nash equilibrium is approximately $[(0, 0.01, 0.001, 0.98), (0.20, 0.88, 0, 0.09)]$ whose expected payoff vector is approximately $(88.11, 1.14)$. The Nash equilibria are calculated via Gambit.

²³We have calculated the values via Mathematica 8.0 software. Note that we have given one example of maximin equilibrium whose value is equal for both players, but there can be other maximin equilibria as well.

4.2 The relation of maximin equilibrium with the other concepts

Perhaps, Nash equilibrium is one of the most well-known solution concept in game theory. Nash's path-breaking theorem says that every finite game in mixed extension possesses at least one Nash equilibrium.

Remark 11. The value of a Nash equilibrium is the same as its payoff vector because no player can unilaterally deviate from a Nash equilibrium.

The following proposition shows that maximin equilibrium is a generalization of Nash equilibrium.

Proposition 12. *Every Nash equilibrium is a maximin equilibrium.*

Proof. Let (p_i, p_j) be a Nash equilibrium. By Remark 11, the value of it is the same as its payoff vector. Accordingly, we have $p_i \in \arg \max_{p'_i \in \Delta X_i} v_i(p'_i, p_j)$ for all player i because the values of the other profiles can not increase by Remark 3. Hence, the profile (p_i, p_j) is also a maximin equilibrium. \square

The following two propositions show the Pareto dominance relation between Nash equilibrium and strong maximin equilibrium.

Proposition 13. *For every Nash equilibrium that is not a strong maximin equilibrium there exists a strong maximin equilibrium that Pareto dominates it.*

Proof. If a Nash equilibrium q in a game Γ is not a strong maximin equilibrium, then there exists a strong maximin equilibrium p whose value $v(p)$ Pareto dominates $v(q)$. It implies that p Pareto dominates q in the game Γ since the payoff vector of the Nash equilibrium q is the same as its value by Remark 11. \square

Proposition 14. *No strong maximin equilibrium is ever Pareto dominated by a Nash equilibrium.*

Proof. By contradiction, suppose that a Nash equilibrium q Pareto dominates a strong maximin equilibrium p . It implies that the value of q also Pareto dominates the value of p by Remark 11. But this is a contradiction to our supposition that p is a strong maximin equilibrium. \square

The two propositions above are closely linked but one does not follow from the other. Because, Proposition 13 does not exclude the existence of a Nash equilibrium that is both Pareto dominated by a strong maximin equilibrium and Pareto dominates another strong maximin equilibrium. Proposition 14 shows that this is not the case.

The next proposition illustrates that in zero-sum games a possibly mixed strategy is a Nash equilibrium if and only if it is a maximin equilibrium.

Proposition 15. *In a two-player zero-sum game, a profile is a maximin equilibrium if and only if it is a Nash equilibrium.*

Proof. ‘ \Rightarrow ’ By contraposition, suppose that the strategy profile (p_1, p_2) is not a Nash equilibrium, we show that it can not be a maximin equilibrium. Firstly, we show that the value of (p_1, p_2) is always Pareto dominated. Let (\bar{p}_1, \bar{p}_2) be a Nash equilibrium whose payoff vector is $(w, -w)$. Suppose, without loss of generality, that $u_1(p_1, p_2) < w$ and $u_2(p_1, p_2) > -w$. Next, player 1 can profitably deviate to her maximin strategy \bar{p}_1 from which she will receive at least w . It implies that the value of the profile (p_1, p_2) for player 2 is at most $-w$. By Remark 3, we have $v_1(p_1, p_2) \leq u_1(p_1, p_2)$. It implies that the value of the profile (p_1, p_2) is Pareto dominated by the value of the Nash equilibrium (\bar{p}_1, \bar{p}_2) . Secondly, the value of the profile (\bar{p}_1, p_2) for player 1 is strictly more than the value of (p_1, p_2) , that is $v_1(\bar{p}_1, p_2) > v_1(p_1, p_2)$ because $v_1(\bar{p}_1, p_2) = w$. Since the profile (p_1, p_2) also violates the second condition in the definition of maximin equilibrium, it can not be a maximin equilibrium.

‘ \Leftarrow ’ By Proposition 15, every Nash equilibrium is a maximin equilibrium. □

Rationalizable strategy profiles and correlated equilibrium are two different generalizations of the Nash equilibrium.²⁴ One might wonder the relationship of maximin equilibrium with those concepts. First of all, neither rationalizable strategy profiles nor correlated equilibrium is a coarsening of maximin equilibrium. The reason is that, a strictly dominated strategy can be a part of a maximin equilibrium. For example, the profiles (A, C) and (B, D) are pure maximin equilibria in the game in Figure 5. But notice that the strategy A of player 1 is strictly dominated. Moreover, maximin equilibrium is not a generalization of rationalizable strategy profiles. Consider the game in Figure 6 which is taken from Bernheim (1984) and observe

²⁴See Bernheim (1984) and Pearce (1984) for rationalizability and see Aumann (1974) for correlated equilibrium.

	C	D
A	6, 8	0, 0
B	8, 6	1, 7

Figure 5: Example of a maximin equilibrium which include a strictly dominated strategy.

	A	B	C	D
A	0, 7	2, 5	7, 0	0, 1
B	5, 2	3, 3	5, 2	0, 1
C	7, 0	2, 5	0, 7	0, 1
D	0, 0	0, -2	0, 0	10, -1

Figure 6: A rationalizable strategy profile which is not a maximin equilibrium.

that (A, A) is a rationalizable strategy profile, however, it is not a maximin equilibrium. Correlated equilibrium is also not a special case of maximin equilibrium either because a correlated distribution can not be a maximin equilibrium.

One might wonder whether there is a relationship between the maximin (minimax) decision rule²⁵ in decision theory and the maximin equilibrium. Imagine a one-person game in which the decision maker is to make a choice between several gambles. In that case, maximin equilibrium boils down to expected utility maximization just like maximin strategies and Nash equilibrium. In other words, the decision maker has to choose the gamble with the highest expected utility. However, according to maximin decision rule, a player has to choose the gamble which maximizes the utility with respect to the worst state of the world (whose outcome is the minimum) even though the probability assigned to it is very small.

²⁵See Wald (1950) for maximin decision rule and see Gilboa and Schmeidler (1989) for the axiomatic foundations of it.

		Accept	Reject
	Accept	16, 16	16, 16
	Reject	16, 16	Γ'

		Stay silent	Deny	Confess
$\Gamma' =$	Stay silent	100, 100	110, 105	0, 15
	Deny	105, 110	95, 95	5, 380
	Confess	15, 0	380, 5	10, 10

Figure 7: Modified prisoner's dilemma Γ in extensive form.

5 Maximin equilibrium in extensive form and in n-player games

5.1 Extensive form games

It is possible to apply the logic of maximin equilibrium in extensive form games but we have to be careful with the definition of a profitable deviation. Because, if a player plays a strategy out of an agreement, then the other player could be informed about this before waiting the end of the game. Accordingly, the non-deviator has a chance to deviate too. Thus, defining profitable deviations via its normal form would cause a loss of information. Without getting into the heavy notation of extensive form games, we would like to present a simple example to illustrate the maximin equilibria solutions.

In this example, we modify the famous prisoner's dilemma game as follows. Each player has three options to choose, namely 'stay silent' (S), 'deny' (D) or 'confess' (C). In addition, we add another feature: Before playing the actual game, each player is offered to choose 'Accept' (A) or 'Reject' (R). If one of them accepts, then they will both stay 12 years in prison and the actual game will not be played. If both of them reject, then they will play the game. All this information is common knowledge to both players. This situation can be represented by a two-player symmetric game in extensive form which is shown in Figure 7.²⁶

Firstly, observe that the game Γ' has a unique Nash equilibrium (C,C)

²⁶We assume that players' utility is decreasing in the number of years that will be spent in prison. Notice that if the strategy 'stay silent' is removed from the game Γ' for both players then we would obtain the prisoner's dilemma. Note also that the maximin equilibrium of the prisoner's dilemma is the same as its Nash equilibrium.

with a payoff vector of (10, 10). Accordingly, the unique subgame perfect Nash equilibrium of Γ is (AC, AC) whose payoff vector is (16, 16). For simplicity, let us consider the game in pure strategies. There are four maximin equilibria (C, C), (D, S), (S, D) and (S, S) in the game Γ' . The maximin equilibria of Γ are (AC, AC), (RD, RS), (RS, RD) and (RS, RS) whose values are (16, 16), (5, 110), (110, 5) and (100, 100) respectively.

Suppose that both prisoners are in the same cell and they can freely discuss what to choose before their decision is asked about accepting or rejecting the offer. However, they will make their choices in separate cells, that is, non-binding pre-game communication is allowed. We would expect that players would try to convince each other to play a strategy profile that Pareto dominates (AC, AC) because each of them can simply guarantee receiving 16 no matter what.

A potential agreement in this game seems to be the maximin equilibrium (RS, RS). It simply guarantees a payoff of 100 for a player given that his opponent is rational. First of all, a rational player would not unilaterally deviate from this agreement by playing 'Accept' in the beginning because it is not a profitable deviation. Moreover, if a player wants to play (A), it does not make sense for him to make an agreement since any choice of the opponent does not change the outcome. In the second stage game Γ' , a rational player would not unilaterally deviate to C from the agreement (S, S) since it is a non-profitable deviation. He may rationally deviate to D but this is not a problem (even preferable) for the non-deviator as he gets strictly more than the value of 100. Actually, the fact that they have already rejected receiving 16 is a signal for not playing according to the Nash equilibrium.

5.2 N-player games

In this subsection, we discuss the extensions of the definitions and of the results we have presented so far to n-player games.

Regarding the definition of the value function, one only needs to replace the way v_i is written in Definition 4 to,

$$v_i(p) = \min\left\{\inf_{p'_{-i} \in B_{-i}(p)} u_i(p_i, p'_{-i}), u_i(p)\right\},$$

where $B_{-i}(p)$ is the set of $(n - 1)$ -tuple strategy profiles which can occur under unilateral profitable deviations of all players but i with respect to p ,

that is

$$B_{-i}(p) = \{\hat{p}_{-i} \in \Delta X_{-i} | \hat{p}_k \in B_k(p) \text{ for at least one } k \in \{-i\}\}.$$
²⁷

The definition of the maximin equilibrium does not change. In other words, the set of maximin equilibria in a game Γ is the union of the set of Pareto optimal strategy profiles and the set of Nash equilibria of the game defined by the value function of Γ . Moreover, every result in Section 3 and in Section 4 except the Proposition 15 (which requires a two-player zero-sum setting) is valid in n-player games. That is, maximin equilibrium exists in pure strategies, it is invariant under strictly increasing transformations of the payoff functions of the players and it is a generalization of the Nash equilibrium in mixed strategies. Moreover, for every Nash equilibrium that is not a strong maximin equilibrium there exists a strong maximin equilibrium that Pareto dominates it. Besides, no strong maximin equilibrium is ever Pareto dominated by a Nash equilibrium. The proofs are essentially the same as the ones given in Section 3 and in Section 4.

6 Conclusion

In this paper, we have introduced a new solution concept called maximin equilibrium which extends von Neumann's maximin strategy idea to n-player games by incorporating common knowledge of rationality of the players. The rationality assumption we use is stronger than the one of maximin strategist and weaker than the one of Nash strategist. Maximin equilibrium is a method for evaluating the uncertainty that players are facing by playing the game. In other words, maximin equilibrium is a strategy profile whose value (which is a vector) is maximized with respect to the methods of Pareto optimality and of Nash equilibrium. We showed that maximin equilibrium is invariant under strictly increasing transformations of the payoff functions of the players. Moreover, every finite game possesses a maximin equilibrium in pure strategies.

²⁷Actually, one may define a weaker notion of better response by allowing only one player to unilaterally deviate at a time. In this case, the value function and the maximin equilibrium would still be well defined and all the results would still hold in n-player games. For the time being, we prefer the current definition. In the end, n-player games have more important conceptual problems such as the possibility of forming coalitions which we investigate in our following project.

Considering the games in mixed extension, we showed that maximin equilibrium is a coarsening of Nash equilibrium. Moreover, we demonstrated that the set of Nash equilibria coincides with the set of maximin equilibria in two-player zero-sum games. In addition, we proposed a refinement of maximin equilibrium called strong maximin equilibrium. Accordingly, we showed that for every Nash equilibrium that is not a strong maximin equilibrium there exists a strong maximin equilibrium that Pareto dominates it. Besides, no strong maximin equilibrium is ever Pareto dominated by a Nash equilibrium. We also discussed maximin equilibria predictions in several games including the traveler's dilemma.

The concept introduced in this paper opens up several research directions. In this paper, we have defined the value function in n-player games with respect to unilateral deviations only. Therefore, it is not necessarily immune to coalitional deviations. In our following project, we extend the definition of the value function which takes into account coalitional deviations and we define the maximin equilibrium accordingly. In addition, a detailed analysis of maximin equilibrium in extensive form and in repeated games is the topic of our following project.

References

- C.D. Aliprantis and K.C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. 1994. ISBN 9783540326960.
- R. J. Aumann and M. Maschler. Some thoughts on the minimax principle. *Management Science*, 18(5-Part-2):54–63, 1972. URL <http://ideas.repec.org/a/inm/ormnsc/v18y1972i5-part-2p54-63.html>.
- Robert J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974. ISSN 0304-4068.
- Robert J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):pp. 1236–1239, 1976. ISSN 00905364. URL <http://www.jstor.org/stable/2958591>.
- Robert J. Aumann. Nash equilibria are not self-enforcing. *Economic Decision Making: Games, Econometrics and Optimisation*, pages 201–206, 1990.

- Kaushik Basu. The traveler's dilemma: Paradoxes of rationality in game theory. *The American Economic Review*, 84(2):391–395, 1994. ISSN 00028282. URL <http://www.jstor.org/stable/2117865>.
- C. Berge. *Espaces topologiques: Fonctions multivoques*. Dunod, 1959.
- B. Douglas Bernheim. Rationalizable strategic behavior. *Econometrica*, 52(4):pp. 1007–1028, 1984. ISSN 00129682. URL <http://www.jstor.org/stable/1911196>.
- J. K. Goeree Capra C. M., Holt C. A. and R. Gomez. Anomalous behavior in a traveler's dilemma? *American Economic Review*, 89(3):678–690, 1999. doi: 10.1257/aer.89.3.678. URL <http://www.aeaweb.org/articles.php?doi=10.1257/aer.89.3.678>.
- P.C. Fishburn. *Utility theory for decision making*. Publications in operations research. Wiley, 1970. URL <http://books.google.nl/books?id=lyUoAQAAAJ>.
- Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989. ISSN 0304-4068. doi: [http://dx.doi.org/10.1016/0304-4068\(89\)90018-9](http://dx.doi.org/10.1016/0304-4068(89)90018-9). URL <http://www.sciencedirect.com/science/article/pii/0304406889900189>.
- Jacob K. Goeree and Charles A. Holt. Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, 91(5):1402–1422, 2001. doi: 10.1257/aer.91.5.1402. URL <http://www.aeaweb.org/articles.php?doi=10.1257/aer.91.5.1402>.
- J.C. Harsanyi and R. Selten. *A General Theory of Equilibrium Selection in Games*. Mit Press, 1988. ISBN 9780262582384. URL <http://books.google.cz/books?id=zjwkHAAACAAJ>.
- John C. Harsanyi. A general theory of rational behavior in game situations. *Econometrica*, 34(3):pp. 613–634, 1966. ISSN 00129682. URL <http://www.jstor.org/stable/1909772>.
- D.K. Lewis. *Convention: a philosophical study*. Harvard University Press, 1969. URL <http://books.google.nl/books?id=29LWAAAAMAAJ>.

- McLennan Andrew M. McKelvey, Richard D. and Theodore L. Turocy. Gambit: Software tools for game theory, version 14.0.1. 2013. URL <http://www.gambit-project.org>.
- John F. Nash. Equilibrium points in n-person games. *Proceedings of the national academy of sciences*, 36(1):48–49, 1950.
- John F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2): 286–295, 1951. ISSN 0003486X. URL <http://www.jstor.org/stable/1969529>.
- David G. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):pp. 1029–1050, 1984. ISSN 00129682. URL <http://www.jstor.org/stable/1911197>.
- Robert W Rosenthal. Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic Theory*, 25(1):92–100, 1981. ISSN 0022-0531. doi: [http://dx.doi.org/10.1016/0022-0531\(81\)90018-1](http://dx.doi.org/10.1016/0022-0531(81)90018-1). URL <http://www.sciencedirect.com/science/article/pii/0022053181900181>.
- Ariel Rubinstein. Dilemmas of an economic theorist. *Econometrica*, 74(4): 865–883, 2006. ISSN 1468-0262. doi: 10.1111/j.1468-0262.2006.00689.x. URL <http://dx.doi.org/10.1111/j.1468-0262.2006.00689.x>.
- Ariel Rubinstein. Instinctive and cognitive reasoning: A study of response times. *The Economic Journal*, 117(523):1243–1259, 2007. ISSN 1468-0297. doi: 10.1111/j.1468-0297.2007.02081.x. URL <http://dx.doi.org/10.1111/j.1468-0297.2007.02081.x>.
- L.J. Savage. *The Foundations of Statistics*. Wiley publications in statistics. John Wiley and Sons, 1954. URL <http://books.google.nl/books?id=AGkoAAAAMAAJ>.
- R. Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1):25–55, 1975. ISSN 0020-7276. doi: 10.1007/BF01766400. URL <http://dx.doi.org/10.1007/BF01766400>.

- Reinhard Selten. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträgeit. *Zeitschrift für die gesamte Staatswissenschaft : ZgS.*, 1965.
- J. von Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1928. ISSN 0025-5831. doi: 10.1007/BF01448847. URL <http://dx.doi.org/10.1007/BF01448847>.
- J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1953, third edition, 1944.
- A. Wald. *Statistical decision functions*. Wiley publications in statistics. Wiley, 1950. URL <http://books.google.nl/books?id=nq0gAAAAMAAJ>.
- Inc. Wolfram Research. Mathematica version 8.0. 2010.